



SCHLOSS DAGSTUHL  
Leibniz-Zentrum für Informatik

**Jahresbericht**  
***Annual Report***

**2024**





SCHLOSS DAGSTUHL  
Leibniz-Zentrum für Informatik

Jahresbericht  
*Annual Report*  
**2024**

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Germany Telefon: +49 6871 905 0 Fax: +49 6871 905 133 E-Mail: <a href="mailto:service@dagstuhl.de">service@dagstuhl.de</a>
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen
Geschäftsführung	Prof. Dr.-Ing. Holger Hermanns Heike Meißner
Gesellschafter	Gesellschaft für Informatik e. V., Deutschland Universität des Saarlandes, Deutschland Technische Universität Kaiserslautern, Deutschland Karlsruher Institut für Technologie (KIT), Deutschland Technische Universität Darmstadt, Deutschland Universität Stuttgart, Deutschland Universität Trier, Deutschland Johann Wolfgang Goethe-Universität Frankfurt am Main, Deutschland Institut National de Recherche en Informatique et en Automatique (INRIA), Frankreich Centrum voor Wiskunde en Informatica (CWI), Niederlande Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Deutschland
Einbindung	Mitglied der Leibniz-Gemeinschaft
Verantwortlicher Redakteur	Prof. Dr.-Ing. Holger Hermanns
Redaktion	Dr. Marcel Ackermann, Dr. Andreas Dolzmann, Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Heike Clemens, Dr. Michael Didas, Dr. Andreas Dolzmann, Jutka Gasiorowski, Dr.-Ing. Michael Gerke, Dagmar Glaser, Oliver Hoffmann, Dr. Marsha Kleinbauer, Heike Meißner, Petra Meyer, Stephanie Meier, Dr. Florian Reitz, Simone Schilke, Thomas Schillo, Christina Schwarz, Prof. Raimund Seidel, Ph. D., Dr. Michael Wagner und die Autoren der Zusammenfassungen in Kapitel 6
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
ISSN	2199-1995
Online-Version	<a href="https://www.dagstuhl.de/dagpub/2199-1995">https://www.dagstuhl.de/dagpub/2199-1995</a>
Veröffentlichungsdatum	Oktober 2025
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	© Creative Commons BY-ND 4.0

# Vorwort

## Foreword

Oberflächlich betrachtet war 2024 ein ruhiges, ereignisarmes Jahr für Dagstuhl, in dem im Seminarwesen, bei dblp, beim Publikationswesen alles seinen gewohnten, ich traue mich sogar zu sagen, seinen gewohnt erfolgreichen Lauf lief: Gastzahlrekorde, Antragsrekorde, Zugriffrekorde, etc. können im Detail in diesem Bericht gefunden werden.

Unter der Oberfläche aber passierte doch so einiges mehr: Unsere alle sieben Jahre erfolgende Evaluierung kam zu einem erfolgreichen Abschluss mit vielen lobenden Worten. Die Finanzierung ist damit für die nächsten sieben Jahre gesichert. Die Evaluierung machte auch den Weg frei für einen sogenannten Sondertatbestandsantrag, der uns zusätzliche Mittel beschere soll, um eine Abteilung zu betreiben, die sich Open Science in der Informatik widmen soll. Der Antrag wurde gestellt. Darüber entschieden wird im Mai 2025.

Renoviert oder baulich neu gestaltet haben wir im Jahr 2024 nur wenig. Dafür wurde um so mehr Aufwand und Energie in die Konzeption und Planung einer energetischen Neuaufstellung des Standortes in Dagstuhl gesteckt. Wir sind optimistisch, dass wir recht bald über mehr als nur Konzeption und Planung berichten werden können.

Schließlich gab es noch eine außergewöhnliche Aktivität, nämlich die Suche nach einem neuen wissenschaftlichen Direktor, da RS ab 1. Mai 2025 aus Altersgründen ausscheiden wird. Für RS ist es also wohl das letzte Mal, dass er sich mit diesen Prolegomena beschäftigt. Er hofft, Dagstuhl in mindestens so gutem Zustand zu übergeben, wie er es vor 11 Jahren vorgefunden hat, und dass es für seine Nachfolge eine ebenso schöne Aufgabe sein wird wie es für ihn gewesen ist. RS bedankt sich an dieser Stelle bei allen Förderern und Ratgebern, und er bedankt sich herzlich bei allen Mitarbeiterinnen und Mitarbeitern. Sie sind es, die durch ihren Einsatz dem LZI Herz und Seele geben.

Im Namen der Geschäftsleitung

Prof. Raimund Seidel, Ph. D.  
*Wissenschaftlicher Direktor*

Heike Meißner  
*Technisch-administrative Geschäftsführerin*

On the surface, 2024 was a quiet year for Dagstuhl with few noteworthy events. The seminar program, dblp, the publishing activities - they all proceeded in their usual way, I daresay, their accustomed successful way. New record numbers for guests, proposals, accesses, etc. can be found in more detail in this report.

However, beneath the surface there were quite a few things happening: Our septennial evaluation came to a successful completion with quite a lot of praise. Our financing for the next seven years is thus secured. The evaluation also opened the way for a so-called “Sondertatbestandsantrag” (if you love long German composite words), a proposal to expand our base funding in order to allow us to operate a division dedicated to open science in informatics. We will have to wait until May 2025 for a decision on our proposal.

We had very few renovations in 2024. However, we put all the more effort and energy into the conception and planning of an energy saving overhaul of our facilities at Dagstuhl. We are optimistic that very soon we will be able to report on more than just conception and planning.

Finally, there was one unusual activity in 2024, namely the search for a new scientific director, as RS will have to step down from May 2025 for age reasons. For RS this is therefore, in all likelihood, the last time that he has to deal with these prolegomena. He hopes to pass on Dagstuhl in at least as good condition as he inherited it 11 years ago, and he hopes that his successor will find this job as rewarding as RS has experienced it himself. RS thanks all supporters and advisors, and he wants to express his heartfelt thanks to all members of the entire staff and management. They are the ones who through their untiring efforts endow LZI with heart and spirit.

On behalf of the Managing Directors

# Inhaltsverzeichnis

## Contents

	Vorwort	
	<i>Foreword</i> . . . . .	iii
<b>1</b>	<b>Das Zentrum Schloss Dagstuhl</b>	
	<b><i>Schloss Dagstuhl Center</i></b>	<b>1</b>
1.1	Dagstuhls Leitbild	
	<i>Dagstuhl's Mission</i> . . . . .	2
1.2	Neuigkeiten in 2024	
	<i>News from 2024</i> . . . . .	4
<b>2</b>	<b>Seminare und Workshops</b>	
	<b><i>Seminars and Workshops</i></b>	<b>13</b>
2.1	Dagstuhl-Seminare	
	<i>Dagstuhl Seminars</i> . . . . .	14
2.2	Dagstuhl-Perspektiven-Workshops	
	<i>Dagstuhl Perspectives Workshops</i> . . . . .	15
2.3	Einreichung der Anträge und Begutachtungsverfahren	
	<i>Proposal Submission and Review Process</i> . . . . .	15
2.4	Seminar-Programm 2024	
	<i>Seminar-Programm in 2024</i> . . . . .	16
2.5	Angaben zu Teilnehmenden und Organisierenden	
	<i>Participant and Organizer Data</i> . . . . .	17
2.6	Themen und Forschungsgebiete	
	<i>Topics and Research Areas</i> . . . . .	18
2.7	Weitere Veranstaltungstypen	
	<i>Further Event Types</i> . . . . .	22
2.8	Qualitätssicherung	
	<i>Quality Assurance</i> . . . . .	22
2.9	Auslastung des Zentrums	
	<i>Utilization of the Center</i> . . . . .	23
<b>3</b>	<b>Bibliographiedatenbank dblp</b>	
	<b><i>dblp computer science bibliography</i></b>	<b>27</b>
3.1	Offene Bibliographiedaten für die Informatik	
	<i>Open Bibliographic Data in Computer Science</i> . . . . .	28
3.2	Schloss Dagstuhl und dblp	
	<i>Schloss Dagstuhl and dblp</i> . . . . .	29
3.3	Der dblp SPARQL Query Service	
	<i>The dblp SPARQL Query Service</i> . . . . .	30
3.4	dblp-Daten-Dumps	
	<i>dblp dump data releases</i> . . . . .	30
3.5	Statistiken der Datenakquise	
	<i>Data Acquisition Statistics</i> . . . . .	31
3.6	Statistiken der Datenkuration	
	<i>Data Curation Statistics</i> . . . . .	32

3.7	Nutzungsstatistiken <i>Usage Statistics</i> . . . . .	33
<b>4</b>	<b>Dagstuhl Publishing</b>	
	<b><i>Dagstuhl Publishing</i></b> . . . . .	<b>35</b>
4.1	Über Dagstuhl Publishing <i>About Dagstuhl Publishing</i> . . . . .	36
4.2	Dagstuhl Publishing Beirat <i>Dagstuhl Publishing Advisory Board</i> . . . . .	36
4.3	Portfolio <i>Portfolio</i> . . . . .	37
4.4	Infrastruktur <i>Infrastructure</i> . . . . .	41
4.5	Nutzungsstatistiken <i>Usage Statistics</i> . . . . .	42
<b>5</b>	<b>Resonanz</b>	
	<b><i>Feedback</i></b> . . . . .	<b>47</b>
5.1	Resonanz zu Seminaren und Workshops <i>Feedback on Seminars and Workshops</i> . . . . .	48
5.2	Resonanz zur Bibliographiedatenbank dblp <i>Feedback on the dblp computer science bibliography</i> . . . . .	54
<b>6</b>	<b>Die Seminare in 2024</b>	
	<b><i>The 2024 Seminars</i></b> . . . . .	<b>57</b>
6.1	From Proofs to Computation in Geometric Logic and Generalizations . . . . .	60
6.2	Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis . . . . .	61
6.3	Representation, Provenance, and Explanations in Database Theory and Logic . . . . .	63
6.4	Symmetric Cryptography . . . . .	65
6.5	The Emerging Issues in Bioimaging AI Publications and Research . . . . .	66
6.6	Next Generation Protocols for Heterogeneous Systems . . . . .	70
6.7	Reviewer No. 2: Old and New Problems in Peer Review . . . . .	71
6.8	Are Knowledge Graphs Ready for the Real World? Challenges and Perspective . . . . .	74
6.9	Beyond-Planar Graphs: Models, Structures and Geometric Representations . . . . .	76
6.10	Safety Assurance for Autonomous Mobility . . . . .	77
6.11	Triangulations in Geometry and Topology . . . . .	78
6.12	Computational Approaches to Strategy and Tactics in Sports . . . . .	80
6.13	AI for Social Good . . . . .	82
6.14	Reflections on Pandemic Visualization . . . . .	83
6.15	Applied and Combinatorial Topology . . . . .	84
6.16	Robust Query Processing in the Cloud . . . . .	86
6.17	Shapes in Graph Data: Theory and Implementation . . . . .	88
6.18	Logics for Dependence and Independence: Expressivity and Complexity . . . . .	90
6.19	EU Cyber Resilience Act: Socio-Technical and Research Challenges . . . . .	92
6.20	Trustworthiness and Responsibility in AI – Causality, Learning, and Verification . . . . .	94
6.21	Low-Dimensional Embeddings of High-Dimensional Data: Algorithms and Applications . . . . .	97
6.22	Network Calculus . . . . .	101
6.23	Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems . . . . .	103
6.24	Research Software Engineering: Bridging Knowledge Gaps . . . . .	105
6.25	Hardware Support for Cloud Database Systems in the Post-Moore's Law Era . . . . .	108
6.26	Automated Synthesis: Functional, Reactive and Beyond . . . . .	110
6.27	Code Search . . . . .	111
6.28	Computational Metabolomics: Towards Molecules, Models, and their Meaning . . . . .	112
6.29	Resilience and Antifragility of Autonomous Systems . . . . .	114
6.30	Generalization by People and Machines . . . . .	116
6.31	Discrete Algorithms on Modern and Emerging Compute Infrastructure . . . . .	117
6.32	Causal Inference for Spatial Data Analytics . . . . .	118
6.33	Evaluation Perspectives of Recommender Systems: Driving Research and Education . . . . .	120
6.34	Classical-Quantum Synergies in the Theory and Practice of Quantum Error Correction . . . . .	122
6.35	Stochastic Games . . . . .	125

6.36	Designing Computers' Control Over Our Bodies . . . . .	126
6.37	Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis . . . . .	128
6.38	Computational Analysis and Simulation of the Human Voice . . . . .	129
6.39	Teaching Support Systems for Formal Foundations of Computer Science . . . . .	130
6.40	Computational Creativity for Game Development . . . . .	131
6.41	Theory of Randomized Optimization Heuristics . . . . .	133
6.42	A Game of Shadows: Effective Mastery Learning in the Age of Ubiquitous AI . . . . .	135
6.43	Dynamic Traffic Models in Transportation Science . . . . .	137
6.44	Automated Machine Learning For Computational Mechanics . . . . .	139
6.45	Programmable Host Networking . . . . .	141
6.46	Improving Trust between Humans and Software Robots in Robotic Process Automation . . . . .	142
6.47	Art, Visual Illusions, and Data Visualization . . . . .	143
6.48	Learning with Music Signals: Technology Meets Education . . . . .	144
6.49	Resource-Efficient Machine Learning . . . . .	147
6.50	Security and Privacy of Current and Emerging IoT Devices and Systems . . . . .	149
6.51	Proof Representations: From Theory to Applications . . . . .	151
6.52	Leveraging AI for Management Decision-Making . . . . .	152
6.53	Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES) . . . . .	156
6.54	Conversational Agents: A Framework for Evaluation (CAFE) . . . . .	157
6.55	Artificial Intelligence and Formal Methods Join Forces for Reliable Autonomy . . . . .	159
6.56	Next-Generation Secure Distributed Computing . . . . .	161
6.57	Extended Reality Accessibility . . . . .	162
6.58	Explainable AI for Sequential Decision Making . . . . .	163
6.59	Algebraic and Analytic Methods in Computational Complexity . . . . .	165
6.60	Statistical and Probabilistic Methods in Algorithmic Data Analysis . . . . .	167
6.61	Fair Division: Algorithms, Solution Concepts, and Applications . . . . .	168
6.62	Greening Networking: Toward a Net Zero Internet . . . . .	170
6.63	New Tools in Parameterized Complexity: Paths, Cuts, and Decomposition . . . . .	172
6.64	SAT and Interactions . . . . .	174
6.65	Automated Programming and Program Repair . . . . .	176
6.66	Behavioural Metrics and Quantitative Logics . . . . .	178
6.67	Machine Learning Augmented Algorithms for Combinatorial Optimization Problems . . . . .	180
6.68	Machine Learning for Protein-Protein and Protein-Ligand Interactions . . . . .	182
6.69	Reframing Technical Debt . . . . .	183
6.70	Rethinking the Role of Bayesianism in the Age of Modern AI . . . . .	188
6.71	Research Infrastructures and Tools for Collaborative Networked Systems Research . . . . .	191
6.72	Graph Algorithms: Distributed Meets Dynamic . . . . .	193
6.73	Regular Expressions: Matching and Indexing . . . . .	194
6.74	Deep Learning for RNA Regulation and Multidimensional Transcriptomics . . . . .	196
6.75	Human in the Loop Learning through Grounded Interaction in Games . . . . .	198
6.76	Coding Theory and Algorithms for Emerging Technologies in Synthetic Biology . . . . .	201
6.77	Quantum Software Engineering . . . . .	203
<b>7</b>	<b>Öffentlichkeitsarbeit</b>	
	<b><i>Public Relations and Outreach</i></b>	<b>205</b>
7.1	Pressemitteilungen und Medienarbeit <i>Press Releases and Media Work</i> . . . . .	206
7.2	Fortbildung <i>Educational Training</i> . . . . .	206
7.3	„Dagstuhler Gespräche“ <i>“Dagstuhler Gespräche”</i> . . . . .	207
<b>8</b>	<b>Einrichtungen</b>	
	<b><i>Facilities</i></b>	<b>209</b>
8.1	Hauptstandort in Wadern <i>Main Site in Wadern</i> . . . . .	210
8.2	Dagstuhl an der Universität des Saarlandes <i>Dagstuhl at Saarland University</i> . . . . .	213

8.3	Dagstuhl an der Universität Trier <i>Dagstuhl at the University of Trier</i> . . . . .	213
<b>9</b>	<b>Zentrale Dienste</b> <b>Central Services</b>	<b>215</b>
9.1	Bibliothek <i>Research Library</i> . . . . .	216
9.2	IT-Service <i>IT Service</i> . . . . .	217
<b>10</b>	<b>Kunst</b> <b>Art</b>	<b>219</b>
10.1	Dagstuhl als Galerie <i>Dagstuhl as an Art Gallery</i> . . . . .	220
10.2	Dagstuhls permanente Kunstausstellung <i>Dagstuhl's Permanent Art Exhibition</i> . . . . .	221
<b>11</b>	<b>Struktur der Gesellschaft</b> <b>Structure of the Company</b>	<b>223</b>
11.1	Gründung und Gesellschafter <i>Formation and Shareholders</i> . . . . .	224
11.2	Organe der Gesellschaft <i>Organs of the Organization</i> . . . . .	224
11.3	Gremien der Gesellschaft <i>Boards of the Organization</i> . . . . .	226
<b>12</b>	<b>M. Bläser, M. Herbstritt, E. Reindel: Förderverein „Freunde von Dagstuhl“</b> <b>M. Bläser, M. Herbstritt, E. Reindel: Association “Friends of Dagstuhl”</b>	<b>231</b>
<b>13</b>	<b>Statistiken</b> <b>Statistics</b>	<b>235</b>
13.1	Statistiken zu Seminaren und Workshops <i>Statistics on Seminars and Workshops</i> . . . . .	236
13.2	Statistiken zur Bibliographiedatenbank dblp <i>Statistics of the dblp computer science bibliography</i> . . . . .	248
13.3	Statistiken zu Dagstuhl Publishing <i>Statistics of Dagstuhl Publishing</i> . . . . .	252
<b>14</b>	<b>Veranstaltungen 2024</b> <b>Schedule of Events 2024</b>	<b>257</b>
14.1	Dagstuhl-Seminare <i>Dagstuhl Seminars</i> . . . . .	258
14.2	Dagstuhl-Perspektiven-Workshops <i>Dagstuhl Perspectives Workshops</i> . . . . .	263
14.3	GI-Dagstuhl-Seminare <i>GI-Dagstuhl Seminars</i> . . . . .	264
14.4	Lehrveranstaltungen <i>Educational Events</i> . . . . .	264
14.5	Forschungsgruppentreffen <i>Research Group Meetings</i> . . . . .	264



# **1** **Das Zentrum Schloss Dagstuhl** ***Schloss Dagstuhl Center***

## Dagstuhls Leitbild

1.1

## Dagstuhl's Mission

*Schloss Dagstuhl – Leibniz-Zentrum für Informatik* fördert die Informatikforschung auf internationalem Spitzenniveau durch die Bereitstellung von Infrastrukturen zur wissenschaftlichen Kommunikation und für den Austausch zwischen Forschenden. Ziel von Schloss Dagstuhl ist

- die Förderung der Grundlagenforschung und der anwendungsorientierten Forschung auf dem Gebiet der Informatik,
- die wissenschaftliche Fort- und Weiterbildung im Informatikbereich,
- der Wissenstransfer zwischen Forschung und Anwendung der Informatik,
- der Betrieb einer internationalen Begegnungs- und Forschungsstätte für die Informatik.

Die Förderung und Einbindung von Nachwuchswissenschaftlern ist dabei ein wichtiger Teil dieser Aufgabe; ebenso wie der Technologietransfer zwischen Forschung und Industrie.

### ■ Entwicklung des Zentrums

Die Idee zur Gründung eines Tagungszentrum für Informatik wurde Ende der 1980er Jahre geboren, zu einem Zeitpunkt, an dem die Informatikforschung – ursprünglich der Mathematik und den Ingenieurwissenschaften entsprungen – enormen Aufwind erfuhr. Die *Gesellschaft für Informatik* beobachtete damals die zunehmende Nachfrage von Informatikwissenschaftlern am weltbekannten *Mathematischen Forschungsinstitut Oberwolfach* und sah die Notwendigkeit, ein eigens auf die Informatik ausgerichtetes Zentrum einzurichten. Schloss Dagstuhl wurde schließlich 1990 gegründet und entwickelte sich rasch zu einem weltweit renommierten Treffpunkt in der Informatikforschung. Heute beherbergt die Begegnungsstätte (siehe Fig. 1.1) normalerweise jährlich mehr als 3 000 internationale Gäste.

Seit 2005 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 96 Forschungsinstituten, Bibliotheken und Museen.<sup>1</sup> Schloss Dagstuhl wird seit 2006 durch eine Bund-Länder-Förderung finanziert.

Zu dem anfänglich alleinigen Schwerpunkt des Seminarprogramms haben sich in den vergangenen Jahren zwei weitere Geschäftsfelder hinzugesellt: Zum einen der Betrieb der offenen Bibliographiedatenbank dblp, zum anderen die Angebote als Open-Access-Verleger für die Informatikforschenden.

### ■ Seminar- und Workshop-Programm

Schwerpunkt des wissenschaftlichen Programms von Schloss Dagstuhl sind die Dagstuhl-Seminare und die Dagstuhl-Perspektiven-Workshops: Etwa 30 bzw. 45 internatio-

*Schloss Dagstuhl – Leibniz-Zentrum für Informatik* (Leibniz Center for Informatics) pursues its mission of furthering world class research in computer science by facilitating communication and interaction between researchers. The objective of Schloss Dagstuhl is

- to promote basic and application-oriented research in the field of informatics,
- to support advanced, scientific vocational training and to further education in the field of informatics,
- to promote the transfer of knowledge between research into informatics and application of informatics,
- and to operate an international forum and research institute for informatics.

Including and thus promoting young talents is seen as an important part of our efforts, so is promoting the exchange of knowledge and findings between academia and industry.

### ■ History of the Center

The idea behind a seminar center for informatics came about during the late 1980s, when research in computer science was growing rapidly worldwide as an offshoot of mathematics and engineering. At that time the German *Gesellschaft für Informatik* (German Informatics Society) became aware of the growing number of computer scientists at the world-famous *Mathematics Research Institute* in Oberwolfach, Germany, and recognized the need for a meeting venue specific to the informatics community. Schloss Dagstuhl was founded in 1990 and quickly became established as one of the world's premier centers for informatics research. Today, Schloss Dagstuhl (see Fig. 1.1) normally hosts over 3,000 research guests each year from countries across the globe.

Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association, a non-profit research consortium composed of 96 research institutes, libraries and museums throughout Germany.<sup>1</sup> Since 2006 the center has been jointly funded by the German federal and state governments.

Since the very first days of Schloss Dagstuhl, the seminar and workshop meeting program has always been the focus of its programmatic work. In recent years, Schloss Dagstuhl has expanded its operation and also has significant efforts underway in operating the dblp computer science bibliography and in open access publishing for the computer science community.

### ■ Seminar and Workshop Program

The Dagstuhl Seminars and Dagstuhl Perspectives Workshops form the focus of the center's work. Whereas about 30 or 45 established and young researchers gather at

<sup>1</sup> Stand Dezember 2024.  
As of December 2024.



Fig. 1.1  
Aerial photograph of Schloss Dagstuhl.

nale Forschende treffen sich eine halbe bis ganze Woche auf Schloss Dagstuhl, um im Rahmen eines Dagstuhl-Seminars intensiv über ihre aktuelle Forschung zu diskutieren. Darüber hinaus trifft sich in Dagstuhl-Perspektiven Workshops eine kleinere Gruppe von ca. 30 Forschenden der Weltspitze, um über den aktuellen Stand und die zukünftigen Schwerpunkte eines ganzen Forschungsfeldes zu beraten.

Die Seminare und Perspektiven-Workshops werden jeweils von etwa vier ausgewiesenen Forschenden im entsprechenden Gebiet beantragt. Anträge werden durch das wissenschaftliche Direktorium (siehe Kapitel 2.9) begutachtet. Stellenwert bei der Begutachtung haben neben dem eigentlichen Inhalt des Antrags auch die vorgeschlagene Gästeliste sowie die Antragsteller. Nach Annahme finden die entsprechenden Veranstaltungen dann durchschnittlich zwischen 6 und 18 Monaten später statt. Eine Teilnahme ist nur mit einer persönlichen Einladung durch das Zentrum möglich.

Das Seminarzentrum ist im und rund um das 1760 erbaute Schloss Dagstuhl beheimatet und befindet sich in einer ländlichen Gegend im nördlichen Saarland, im Herzen des Dreiländerecks Deutschland, Frankreich und Luxemburg. Es bietet den Gästen eine einzigartige Arbeitsumgebung, die den Austausch mit anderen Gästen in einer wohnlichen Atmosphäre fördert. Gemütliche Sitzecken, ansprechende Essräume, eine herausragenden Informatik-Fachbibliothek, sowie eine Vielzahl von zusätzlichen Arbeits- und Freizeiträumen bieten vielfältige Möglichkeiten, damit sich Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

Nähere Informationen über Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops finden sich in Kapitel 2.

the Dagstuhl Seminars to report on and discuss their current work, smaller groups of about 30 of the international elite of a field gather at the Dagstuhl Perspectives Workshops for the purpose of reflecting on the current status of research and potential development perspectives.

The Dagstuhl Seminars and Perspectives Workshops are characterized by the fact that they are subject to an exacting quality assurance process. A small group of up to four scientists of international standing submit a proposal for a seminar on a specific research topic. The proposal is reviewed by the center's Scientific Directorate (see Section 2.9) with regard to its content, the proposed guest list and those submitting the proposal. The seminars and workshops are held 6 to 18 months later in the seclusion of the center's facilities at Dagstuhl Castle. Participation in a seminar is possible only by way of personal invitation by the center.

Located in a 1760 built manor house in the idyllic countryside of northern Saarland at the heart of the tri-country region formed by Germany, France and Luxembourg, Schloss Dagstuhl offers visitors a unique working environment that encourages guests to interact with each other in tandem with daily life. Lounges, formal and informal dining areas, a world-class research library, and an impressive range of work and leisure rooms offer multiple possibilities for connecting one-on-one outside of the official conference rooms and meeting times.

More information on the Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Chapter 2.

## ■ Bibliographiedatenbank dblp

Bereits seit 2011 betreibt Schloss Dagstuhl in enger Zusammenarbeit mit der Universität Trier die Bibliographiedatenbank dblp. Seit November 2018 ist Schloss Dagstuhl in vollem Umfang alleine für den Betrieb der Datenbank verantwortlich.

Mit mittlerweile mehr als 7,7 Millionen Publikationseinträgen ist dblp die weltweit größte offene Sammlung bibliographischer Daten in der Informatik. Der dblp-Dienst ist darauf ausgerichtet, Forschenden bei ihrer täglichen Arbeit zu unterstützen, etwa bei der Literaturrecherche oder beim Bezug von elektronisch verfügbaren Volltexten. Dabei gilt dblp in der Informatik insbesondere als die Referenzdatenbank für qualitätsgesicherte, normierte Bibliographiedaten. Aber auch Forschungsförderer und Entscheidungsträger unterstützt dblp, etwa durch das Pflegen und öffentlich Verfügbarmachen von personalisierten Publikationsnachweisen. Durch den Betrieb von dblp leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Details über dblp finden sich in Kapitel 3.

## ■ Dagstuhl Publishing

Die Förderung der Kommunikation zwischen den Forschenden in der Informatik gehört zu der zentralen Aufgabe von Schloss Dagstuhl. Wissenschaftliche Veröffentlichungen sind Teil der Forschungskultur, um qualitätsgesicherte Forschungsergebnisse zu diskutieren und zu kommunizieren. Mit seinen Open-Access-Verlagsangeboten unterstützt Schloss Dagstuhl die Forschungsgemeinde dabei, freien Zugang zu den wichtigsten und neuesten Forschungsergebnissen zu erlangen.

Neben Veröffentlichungen, die in engem Bezug zum wissenschaftlichen Programm stehen, verlegt Schloss Dagstuhl auch Konferenzbände und Zeitschriften. Herausragende Reihe ist dabei LIPIcs, in der die Publikationen erstklassiger Konferenzen erscheinen. Alle Angebote der Verlagsabteilung werden durch international besetzte Editorial Boards qualitätsgesichert.

Details zu Dagstuhl Publishing finden sich in Kapitel 4.

## ■ dblp computer science bibliography

Since 2011, Schloss Dagstuhl has been operating the dblp computer science bibliography in close cooperation with the University of Trier. In November 2018, Schloss Dagstuhl alone assumed full responsibility for the operation of the database.

Listing more than 7.7 million articles, dblp is the world's most comprehensive open data collection of computer science research articles. The goal of dblp is to support computer scientists in their daily work, for example when reviewing the literature of a given author or subject area, or when searching for online full-text versions of research articles. The dblp database is often considered to be the reference database for quality-assured and normalized bibliographic metadata in computer science. Additionally, dblp supports funding agencies and decision makers by providing and curating personalized author bibliographies. By operating dblp, Schloss Dagstuhl furthers its mission of promoting the identification, dissemination and implementation of new computer science developments at an internationally recognized level.

More information about the dblp computer science bibliography can be found in Chapter 3.

## ■ Dagstuhl Publishing

Enabling communication between researchers in computer science is part of Dagstuhl's central mission. Scholarly publications belong to the culture of discussing and communicating quality-controlled research results on a global level. Dagstuhl's open-access publishing services hence support the need of the research community to have access to the most important and most recent research results.

In addition to the open documentation of proceedings of its seminar and workshop program, Schloss Dagstuhl also publishes proceedings for computer science conferences and journals. The flagship product of Dagstuhl Publishing is the LIPIcs series, which publishes proceedings of outstanding computer science conferences. The scientific quality of all products is supervised by international editorial boards.

More information on Dagstuhl Publishing can be found in Chapter 4.

## Neuigkeiten in 2024

1.2

## News from 2024

### ■ Förderung von Schloss Dagstuhl

Am 14. Mai 2024 beschloss die GWK, die Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik fortzusetzen. Damit ist die Finanzierung von Schloss Dagstuhl weitere 7 Jahre gesichert. Diese Entscheidung schließt die 2023 begonnene und regulär alle 7 Jahre stattfindende Evaluierung von Schloss Dagstuhl durch die Leibniz-Gemeinschaft ab. Die Arbeit der Abteilungen „Seminarwesen“ und „dblp“ wurden dabei vom Senat der Leibniz Gemeinschaft als „exzellent“ und die der Abteilung „Publishing“ als „sehr gut“ gewürdigt.

### ■ Evaluation of Schloss Dagstuhl

On May 14, 2024, the GWK decided to continue funding Schloss Dagstuhl – Leibniz Center for Informatics. The funding of Schloss Dagstuhl is thus secured for another 7 years. This decision concludes the evaluation of Schloss Dagstuhl by the Leibniz Association, which began in 2023 and regularly takes place every 7 years. The work of the “Seminars” and “dblp” departments were rated by the Leibniz Association Senate as “excellent” and that of the “Publishing” department as “very good”.

## ■ Das Team

Am Ende des Jahres 2024 beschäftigte Schloss Dagstuhl insgesamt 45,69 Vollzeitäquivalente bzw. 64 Angestellte.

Im Jahr 2024 hat das LZI 2 studentische Hilfskräfte im Rahmen der Drittmittelprojekte eingestellt. Im Team für Küche und Hausdienst wurden 7 Mitarbeiterinnen, davon 2 Vollzeitkräfte und 5 teilzeit- bzw. geringfügig Beschäftigte vor allem als Ersatz für Ausscheidende Kräfte eingestellt. Die IT wird seit August durch eine Auszubildende unterstützt.

## ■ Dagstuhls Gremien und Organe

Als Nachfolger für ausscheidende Mitglieder wurden 2024 folgende Personen neu in den Aufsichtsrat berufen: Dr. Christian Berghoff (BMBF), Prof. Dr. Dominik Brodowski (Universität des Saarlandes), Dr. Simone Rehm (Universität Stuttgart) und Prof. Dr. Felix Wolf (Technische Universität Darmstadt) in den Aufsichtsrat aufgenommen; Dominik Brodowski wurde zum stellvertretenden Vorsitzenden gewählt.

Prof. Dr. h. c. Carole Goble, CBE (University of Manchester) wurde als weiteres Mitglied in den Wissenschaftlichen Beirat berufen.

Detaillierte Angaben finden sich in Kapitel 11.

## ■ Vereinbarkeit von Beruf und Familie

Schloss Dagstuhl hatte bereits 2012 das Audit der berufundfamilie Service GmbH erfolgreich durchlaufen und die drei Jahre gültige Zertifizierung „audit berufundfamilie“ erhalten. In dem Auditierungsprozess wird sichergestellt, dass Firmen eine familien- und lebensphasenbewusste Personalpolitik beibehalten oder in einem systematischen Prozess entwickeln. Mit der Bestätigung des wiederum drei Jahre gültigen Zertifikats am 10. Dezember 2024 hat Schloss Dagstuhl erneut die Re-Auditierung erfolgreich abgeschlossen.

## ■ Seminare und Workshops

Mit 77 Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops – hier und im folgenden unter „Seminare“ subsumiert – war das Zentrum wie schon im Vorjahr überdurchschnittlich gut belegt. Mit über 3 200 Gästen, davon über 2 500 bei Seminaren und über 13 500 Übernachtungen war das Schloss Dagstuhl überdurchschnittlich gut nachgefragt und ausgelastet. In den Anzahlen der Seminargästen und Übernachtungen war es das drittbeste Ergebnis in der Geschichte von Dagstuhl.

Mit 172 Anträgen auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops gab es einen neuen Rekord. Auch aufgrund der begrenzten Kapazitäten von Schloss Dagstuhl wurden davon nur 96 Anträge genehmigt. Bei theoretisch 80 möglichen Seminaren führte das zu einer leichten Erhöhung der Vorlaufzeiten. Die hohe Nachfrage sichert weiterhin den Betrieb und ist ein ehrenvoller Beweis des guten Rufs von Schloss Dagstuhl.

## ■ The Team

At the end of 2024, Schloss Dagstuhl had a total of 45.69 full-time equivalent staff members corresponding to 64 employees.

In 2024, LZI employed 2 student assistants as part of externally funded projects. In the kitchen and housekeeping team, 7 employees were hired, including 2 full-time employees and 5 part-time or marginal employees, mainly to replace departing staff. The IT department has been supported by an apprentice since August.

## ■ Dagstuhl Bodies and Boards

The following persons were appointed to the Supervisory Board in 2024 as successors to departing members: Dr. Christian Berghoff (BMBF), Prof. Dr. Dominik Brodowski (Saarland University), Dr. Simone Rehm (University of Stuttgart), and Prof. Dr. Felix Wolf (Technical University of Darmstadt) were appointed to the Supervisory Board; Dominik Brodowski was elected Deputy Chairman.

Prof. Dr. h. c. Carole Goble, CBE (University of Manchester) was appointed as an additional member of the Scientific Advisory Board.

Detailed information can be found in Chapter 11.

## ■ Compatibility of Work and Family

Since 2012, Schloss Dagstuhl has been successfully audited by berufundfamilie Service GmbH and maintains the certificate “audit berufundfamilie” (work and family), which is valid for three years at a time. The auditing process is intended to ensure that the companies’ human resource policy continues to be family- and life-phase-conscious, or that such a policy is developed systematically. The certificate was confirmed and extended for another three years on December 10, 2024 – Schloss Dagstuhl was successful in its re-audit once again.

## ■ Seminars and Workshops

With 77 Dagstuhl Seminars and Dagstuhl Perspectives Workshops - subsumed here and in the following under “Seminars” - the center was, as in the previous year, booked well above average. With more than 3,200 guests, 2,500 of these for seminars, and more than 13,500 overnight stays, Schloss Dagstuhl saw above average demand and utilization. Judging from the number of guests and overnight stays, this was the third best result in Dagstuhl’s history.

There was a record of 172 applications for Dagstuhl Seminars and Dagstuhl Perspectives workshops. In part due to the limits in Schloss Dagstuhl’s capacity, only 96 proposals were accepted. Given the theoretical limit of 80 possible seminars a year, this led to a small increase in lead times. The high demand continues to secure the center’s operation and is an honorable proof of Schloss Dagstuhl’s good reputation.

Prof. Dr.-Ing. Holger Hermanns (Universität des Saarlandes) wurde als zusätzliches Mitglied in das Wissenschaftliche Direktorium berufen.

In den letzten Jahren sind die Personal- und Energiekosten viel stärker gestiegen als die Finanzierung durch Bund und Länder. Daher wurde eine Preissteigerung der Seminargebühren in Höhe von 10 Euro pro Tag (und Nacht) für die Gäste beschlossen.

Mehr Details und Zahlen zum Seminarprogramm finden sich in Kapitel 2.

### ■ Bibliographiedatenbank dblp erzielt neue Rekorde

Die von Schloss Dagstuhl in Trier betriebene Bibliographiedatenbank *dblp computer science bibliography* konnte 2024 erneut einen Rekord bei der Neuaufnahme von Artikeln aufstellen: Annähernd 620 000 neue Artikel aus den verschiedensten Teilgebieten der Informatik wurden indiziert. Ende 2024 listete dblp damit über 7,7 Millionen Publikationen.

Jeden Monat besuchen mindestens 1,4 Millionen individuelle Nutzende aus aller Welt den dblp-Dienst; jeden Tag wurden über 2,4 Million dblp-Webseiten abgerufen. Im Rahmen der Qualitätssicherung wurden vom dblp-Team mehr als 93 000 Fehlerfälle in der Datenbank identifiziert und korrigiert.

### ■ Der neue dblp SPARQL Query Service ist gestartet

Im September 2024 startete die Bibliographiedatenbank dblp ihren neuen *dblp SPARQL Query Service*.<sup>2</sup> Basierend auf einer Kombination der semantischen Informationen des *dblp Knowledge Graph (dblp KG)* mit offen verfügbaren Zitationsdaten ermöglicht der neue Dienst einen einzigartigen Einblick in die komplexen Zusammenhänge der Informatik-Forschungslandschaft.

Die flexible Abfragesprache SPARQL ermöglicht es den Nutzenden, ihre eigenen Anfragen zu erstellen und so neue Sichten zu generieren. Unterstützt wird dies durch intuitive Eingabemaske mit Syntax-Highlighting, Formatierungshilfen und Autovervollständigung, sowie einer öffentlichen SPARQL-API.

Stand Ende 2024 kombiniert der dblp SPARQL Query Service die annähernd 500 Millionen semantische Fakten des dblp KG mit mehr als 800 Millionen offenen Zitationsfakten.

### ■ Dagstuhl Publishing

**Kontinuierliches Wachstum unserer Verlagsaktivitäten.** Auch im zurückliegenden Jahr 2024 waren die Open-Access Aktivitäten von Dagstuhl Publishing weiterhin erfolgreich. Unsere Hauptserie LIPIcs, die sich dem bezahlbaren Open-Access-Publizieren hochwertiger Konferenzbeiträge verschrieben hat, verzeichnete mit 1648 veröffentlichten Artikeln in 38 Bänden ein deutliches Plus

Prof. Dr.-Ing. Holger Hermanns (Saarland University) was named as an additional member of the Scientific Directorate.

In the last years, personnel and energy costs have risen faster than the funding from the federal government and the state governments. Therefore, a price increase of 10 additional Euros per day (and night) for our guests was agreed upon.

See Chapter 2 for more details and statistics regarding the seminar program.

### ■ dblp computer science bibliography is setting new records

The *dblp computer science bibliography* operated by Schloss Dagstuhl in Trier once again set a record for newly included publications: In 2024, almost 620,000 new articles from various sub-fields of computer science were indexed. By the end of 2024, dblp listed more than 7.7 million publications.

Every month, at least 1.4 million distinct users from all over the world were using the dblp web services; more than 2.45 million dblp websites were accessed every day. More than 93,000 error cases have been identified and corrected in the database as part of the continuous quality assurance efforts of the dblp team.

### ■ The new dblp SPARQL Query Service has launched

In September 2024, the dblp computer science bibliography launched its brand new *dblp SPARQL Query Service*.<sup>2</sup> Based on a combination of semantic information from the *dblp Knowledge Graph (dblp KG)* with openly available citation data, the new service provides a unique glimpse into the complex interrelationships of the computer science research landscape.

The flexible query language SPARQL allows users to create their own queries and generate new insights. They are supported by an intuitive query interface, featuring syntax highlighting, formatting aids and auto-completion, as well as a public SPARQL API endpoint.

As of the end of 2024, the dblp SPARQL Query Service combines nearly 500 million semantic statements from the dblp KG with more than 800 million open citation statements.

### ■ Dagstuhl Publishing

**Continuous growth of our publishing activities.** Dagstuhl Publishing's open-access initiatives saw continued success in the past year 2024. Our main series LIPIcs, which is dedicated to affordable open access publishing of high-quality conference proceedings, recorded a significant increase over the previous year with 1648 published articles in 38 volumes (compared to 1455 articles in 34 volumes

<sup>2</sup> <https://sparql.dblp.org>

zum Vorjahr (1455 Artikel in 34 Bänden) und damit einen erneuten Rekord. Ein Beleg für das nach wie vor hohe Ansehen von LIPIcs ist die große Nachfrage von Konferenzen, die sich (erneut) um die Aufnahme in das LIPIcs-Portfolio bewerben. Insbesondere wurden die Publikationsvereinbarungen mit APPROX, RANDOM, ITP, SAND, TYPES, CCC, ICDT, CONCUR und SoCG um weitere fünf Jahre verlängert.

Die wissenschaftliche Zeitschrift “Transactions on Graph Data and Knowledge (TGDK)”, die 2023 ins Leben gerufen wurde, hat ihre Produktion im Jahr 2024 leicht ausgeweitet. Während im Gründungsjahr 12 Artikel veröffentlicht wurden, veröffentlichte TGDK im Jahr 2024 16 Artikel in drei Ausgaben. Die bisher rund 14.000 PDF Downloads sind ein Beleg für das große Interesse der Zeitschrift.

**Verbesserungen der Metadaten.** Im März 2024 wurde die Metadaten-Schnittstelle wieder auf den DROPS-Publikationsserver migriert und erweitert, sodass nun Formate wie Dagstuhl XML, ACM-XML, DOAJ-XML und Schema.org unterstützt werden. Die Schema.org-Metadaten werden auch in den HTML-Code der veröffentlichten Artikel eingebettet, um Suchmaschinen strukturierte Daten für eine bessere Auffindbarkeit zur Verfügung zu stellen.

Unsere Bemühungen, Autoren zur Angabe ihrer ORCID zu ermutigen, führten zu signifikanten Ergebnissen: Bei LIPIcs-Konferenzen, die 2023 und 2024 aufeinanderfolgende Bände veröffentlichten, stieg die ORCID-Abdeckung von 58 % auf 77 %, wodurch eine bessere Autorenenkennung und eine bessere Verknüpfung zwischen den Publikationen gewährleistet wurde.

Anfang 2024 führte Dagstuhl Publishing das Bib2Doi-Tool ein, das Autoren bei der Anreicherung ihrer Bibliografiedateien durch Hinzufügen von DOIs, URLs und fehlenden Metadaten unterstützt. Dieser Service wurde weithin angenommen und hat die Vernetzung veröffentlichter Artikel erheblich verbessert. So stieg beispielsweise der Anteil der Referenzen mit persistenten Identifikatoren in LIPIcs-Bänden, die zwischen Oktober und Dezember 2024 veröffentlicht wurden, im Vergleich zum gleichen Zeitraum im Jahr 2023 von 50 % auf 80 %.

**Veröffentlichung von ergänzenden Materialien und Artefakten.** Dagstuhl Publishing erweiterte seinen Publikations-Workflow um die Möglichkeit, Metadaten für ergänzende Materialien wie Software und Daten einzureichen. Im November 2024 wurden alle im Laufe des Jahres gesammelten ergänzenden Materialien (insgesamt 90 Stück) auf DROPS<sup>3</sup> veröffentlicht. In Zukunft werden diese Materialien zusammen mit den entsprechenden Forschungsartikeln kontinuierlich veröffentlicht.

Zudem sind jetzt mehr als 170 Snapshots der dblp-Datenbank auf DROPS in den Formaten RDF und XML<sup>4</sup> verfügbar, komplett mit DOIs für eine bessere Auffindbarkeit. Zuvor wurden diese Datensätze auf der dblp-Website als Dateien ohne persistente Identifikatoren gehostet. Diese Snapshots, die seit 2015 monatlich archiviert werden, sind nun gemäß den FAIR-Prinzipien öffentlich zugänglich.

Darüber hinaus wurden im Rahmen des Projects FAIR-

in 2023), thus setting a new record. The high demand from conferences applying (again) for inclusion in the LIPIcs portfolio is proof of the consistently strong reputation. Notably, publication agreements with APPROX, RANDOM, ITP, SAND, TYPES, CCC, ICDT, CONCUR, and SoCG have been renewed for another five years.

The scientific journal Transactions on Graph Data and Knowledge (TGDK), launched in 2023, expanded its output slightly in 2024. While 12 articles were published in its founding year, TGDK published 16 articles across three issues in 2024. The 14,000 PDF downloads to date are proof of the magazine’s great popularity.

**Metadata improvements.** In March 2024, the metadata interface has been migrated back to the DROPS publication server and has been extended, now supporting formats such as Dagstuhl XML, ACM-XML, DOAJ-XML, and Schema.org. The Schema.org metadata is also embedded in the HTML of published articles to provide search engines with structured data for improved discoverability.

Our efforts to encourage authors to provide their ORCID have led to significant results: For LIPIcs conferences that published consecutive volumes in 2023 and 2024, ORCID coverage increased from 58 % to 77 %, ensuring improved author identification and better linkage between publications.

In early 2024, Dagstuhl Publishing launched the Bib2Doi tool which assists authors in enriching their bibliography files by adding DOIs, URLs, and missing metadata. This service has been widely adopted and has significantly improved cross-linking of published articles. For instance, references with persistent identifiers in LIPIcs volumes published between October and December 2024 rose from 50 % to 80 % compared to the same period in 2023.

**Supplementary Materials and Artifact Publications.** Dagstuhl Publishing expanded its publication workflow to include metadata submission for supplementary materials such as software and data. Starting in November 2024, all the supplementary materials collected throughout the year (90 in total) were published on DROPS<sup>3</sup>. Moving forward, these materials will be published continuously alongside corresponding research articles.

Furthermore, more than 170 snapshots of the dblp database are now available on DROPS in RDF and XML formats<sup>4</sup>, complete with DOIs for improved discoverability. Previously, these datasets were hosted on the dblp website as files without persistent identifiers. These snapshots, archived monthly since 2015, are now publicly accessible in compliance with FAIR principles.

Additionally, two open-source components for long-term software archiving were released as part of the FAIR-

CORE4EOSC Ende 2024 zwei Open-Source-Komponenten für die langfristige Software-Archivierung veröffentlicht: der SWH Archive Client und der SWH Deposit Client. Diese Tools ermöglichen die Archivierung von Softwareprojekten und deren Metadaten auf der Plattform Software Heritage.

Mehr Informationen zu den Open-Access-Aktivitäten von Schloss Dagstuhl finden sich in Kapitel 4.

## ■ Dagstuhler Gespräche

Nach einer vor allem durch die Pandemie bedingte Pause wurde die erfolgreiche Vortragsreihe „Dagstuhler Gespräche“ als gemeinsame Veranstaltung von Schloss Dagstuhl und der Stadt Wadern wieder aufgenommen. Ziel dieser Gespräche ist es, der interessierten Öffentlichkeit die breite Vielfalt der Informatik und deren praktische Anwendungen im Alltag oder in wirtschaftlichen Prozessen nahezubringen und in einen gemeinsamen Dialog einzusteigen.

Am 26. Januar 2024 hielt Prof. Dr. Friedemann Mattern, emeritierter Professor für Informatik an der ETH Zürich, einen Vortrag zum Thema „Computer, Digitalisierung und die Energie“. Im Vortrag wurden Trends des Energieverbrauchs und anderer Umweltaspekte der immer weiter fortschreitenden Digitalisierung diskutiert und bewertet. Dabei wurde aufgezeigt, dass es vielleicht weniger die offensichtlichen Aspekte sind, die kritisch hinsichtlich der Umweltauswirkungen zu sehen sind, als vielmehr einige problematische indirekte Folgen.

Am 22. November 2024 sprach Prof. Dr. Eduard Arzt, der zwischen 2007 und 2022 einer der beiden wissenschaftlichen Geschäftsführer des Leibniz-Institut für Neue Materialien in Saarbrücken war, zum Thema „Lernen von der Natur – vom Staunen zum Startup“. Der Vortrag beschrieb, wie eine bioinspirierte Technik von der Entwicklung im Labor bis zu einer saarländischen Firmengründung, die eine nachhaltige Technologie auf den Markt bringt, den Weg nahm: Bei der betrachteten Technologie werden feine „haarige“ Mikrostrukturen, ähnlich denen beim Gecko, in der Automatisierungstechnik benutzt, um z. B. kleine Objekte zu greifen und wieder abzusetzen.

Beide Veranstaltungen waren rege besucht und lösten bei den Anwesenden eine rege Anteilnahme an den an den Vortrag anschließenden Gesprächen aus. Die Reihe wird im kommenden Jahr gewiss fortgesetzt werden.

## ■ Lehrerfortbildung

In Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisierte Schloss Dagstuhl 2024 zum 32. Mal eine Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet.

Angeboten wurden fünf Einheiten: „Making und die

CORE4EOSC project in late 2024: the SWH Archive Client and the SWH Deposit Client. These tools enable the preservation of software projects and their metadata on the Software Heritage platform.

More information about the Open Access activities of Schloss Dagstuhl can be found in Chapter 4.

## ■ “Dagstuhler Gespräche”

After a break due primarily to the pandemic, the successful lecture series “Dagstuhler Gespräche” was resumed as a joint event by Schloss Dagstuhl and the town of Wadern. The aim of these lectures is to introduce the general public to the broad diversity of computer science and its practical applications in everyday life or in economic processes and to enter into a joint dialogue.

On January 26, 2024, Prof. Dr. Friedemann Mattern, Professor Emeritus of Computer Science at ETH Zurich, gave a lecture on the topic “Computers, Digitization and Energy”. The lecture discussed and evaluated trends in energy consumption and other environmental aspects of the ever-increasing digitalization. It was shown that it is perhaps not so much the obvious aspects that should be viewed critically in terms of environmental impact, but rather some problematic indirect consequences.

On November 22, 2024, Prof. Dr. Eduard Arzt, who was one of the two scientific directors of the Leibniz Institute for New Materials in Saarbrücken between 2007 and 2022, gave a lecture on the topic “Learning from nature – from amazement to startup”. The lecture described how a bio-inspired technology made its way from development in the laboratory to the founding of a company in Saarland that brings a sustainable technology to market: In the technology under consideration, fine “hairy” microstructures, similar to those found in geckos, are used in automation technology to, for example, grasp and release small objects.

Both events were well attended and sparked lively interest among those present in the discussions that followed the lecture. The series will certainly be continued next year.

## ■ Teacher Training Program

In 2024, Schloss Dagstuhl hosted its teacher training course for the 32nd time. This workshop is specifically designed for computer science and mathematics teachers teaching grades 11 and 12 in Saarland and Rhineland-Palatinate. It is organized in collaboration with the Landesinstitut für Pädagogik und Medien Saarland LPM (Saarland State Institute for Education and Media) and the Pädagogisches Landesinstitut Rheinland-Pfalz PL

<sup>3</sup> <https://drops.dagstuhl.de/entities/collection/supplementary-materials>

<sup>4</sup> <https://drops.dagstuhl.de/entities/collection/dblp>

Informatik“ von Prof. Dr. Nadine Dittert (Universität Koblenz); „Consensus: Fehlertoleranz und verteilte Algorithmen“ von Dr. Christoph Lenzen (CISPA Helmholtz-Zentrum für Informationssicherheit, St. Ingbert); „Open Research Knowledge Graph – Wie KI und Wissensgraphen Forschungsfragen beantworten“ von Dr. Anna-Lena Lorenz (TIB Leibniz-Informationszentrum Technik & Naturwissenschaften, Hannover); „Local-first Software“ von Prof. Dr. Annette Bieniusa (Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Kaiserslautern) und „Beweisen ist auch nur Programmieren“ von Prof. Dr. Sebastian Hack (Universität des Saarlandes, Saarbrücken).

Mehr Informationen zur Öffentlichkeitsarbeit und zu den Weiterbildungsaktivitäten finden sich in Kapitel 7.

### ■ Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch 2024 wurde die Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum<sup>5</sup> (HLF) fortgesetzt. Diese Veranstaltung bringt herausragende Forscher, die mit dem ACM Turing Award, dem Abelpreis, der Fields-Medaille oder mit dem Nevanlinna-Preises ausgezeichnet wurden, mit außergewöhnlich begabten jungen Wissenschaftlern aus aller Welt zusammen. Zwei ausgewählte Teilnehmende des HLF 2024 erhielten in der Woche vor der elften Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Algebraic and Analytic Methods in Computational Complexity“ (24381), je zwei weitere konnten an den Dagstuhl Seminaren „Fair Division: Algorithms, Solution Concepts, and Applications“ (24401) und „Greening Networking: Toward a Net Zero Internet“ (24402) in der Woche nach dem Forum teilnehmen.

Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2025 zugestimmt.

### ■ Spender und Förderer

Schloss Dagstuhl ist allen Personen, Institutionen und Firmen dankbar, die großzügigerweise durch eine Spenden direkt an das Zentrum oder durch eine Spende an den „Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.“ Schloss Dagstuhl unterstützen. Das Zentrum und der Verein können durch die anerkannte Gemeinnützigkeit Spendenquittungen ausstellen.

Details zum Förderverein finden sich im Kapitel 12.

### ■ Baumaßnahmen und Renovierung

In 4 Gästezimmern des 1993 erbauten Trakts, des sogenannten Neubaus, wurden die ersetzten 4 Bäder grundlegend modernisiert. Weitere Bäder werden, jeweils in der Sommerpause der kommenden Jahre, folgen.

Zu 10 Zimmern im selben Gebäude wurden durch

(Rhineland-Palatinate State Institute for Education).

Five units were offered: “Making and Computer Science” by Prof. Dr. Nadine Dittert (University of Koblenz); “Consensus: Fault Tolerance and Distributed Algorithms” by Dr. Christoph Lenzen (CISPA Helmholtz Center for Information Security, St. Ingbert); “Open Research Knowledge Graph – How AI and Knowledge Graphs Answer Research Questions” by Dr. Anna-Lena Lorenz (TIB Leibniz Information Centre for Science & Technology, Hannover); “Local-first Software” by Prof. Dr. Annette Bieniusa (Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Kaiserslautern) and “Proof is just programming too” by Prof. Dr. Sebastian Hack (Saarland University, Saarbrücken).

Further details about public relations and professional training at Schloss Dagstuhl can be found in Chapter 7.

### ■ Joint Outreach with the Heidelberg Laureate Forum

2024 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum<sup>5</sup> (HLF). The HLF brings winners of the ACM Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize together with exceptionally talented young scientists from all over the world. Two participants were selected and participated in the Dagstuhl Seminar “Algebraic and Analytic Methods in Computational Complexity” (24381), taking place during the week before the eleventh edition of the forum. Two more participants each were selected for the Dagstuhl Seminars “Fair Division: Algorithms, Solution Concepts, and Applications” (24401) and “Greening Networking: Toward a Net Zero Internet” (24402) taking place the week after the forum.

Satisfied with the outstanding success of the initiative, all partners have agreed to continue the cooperation in 2025.

### ■ Sponsors and Donors

Schloss Dagstuhl is grateful to every person, institution or company that generously support the center through a direct donation or a donation to the “Friends of Dagstuhl” (“Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.”). Due to their recognized non-profit status, the center and the Friends of Dagstuhl can issue donation receipts.

More details on the friends of Dagstuhl can be found in Chapter 12.

### ■ Construction Work and Renovation

In the four guest rooms of the 1993-built tract, the so-called “Neubau”, the first four bathrooms have been thoroughly modernized. Further bathrooms will follow, in each summer break of the coming years.

New ethernet cabling was installed for 10 rooms

<sup>5</sup> <https://www.heidelberg-laureate-forum.org>

den Keller neue Ethernet-Verkabelung gelegt. In allen 27 Gästezimmern dieses Traktes wurde die neue oder schon vorbereitete Verkabelung bis zum Schreibtisch gezogen, so dass demnächst alle Gästezimmer des Neubaus einen Ethernet-Anschluss am Schreibtisch haben werden.

Weiterhin wurde der Hörsaal „Kaiserslautern“ mit neuen energiesparenden LED-Panels zur Beleuchtung ausgestattet.

Auf dem durchlaufenden Balkon vor den Zimmern des Gästehauses wurden ausziehbare Trennelemente montiert, so dass Gäste sich nach Belieben sowohl auf dem Balkon treffen können, als auch ihre Privatsphäre geschützt wissen.

## ■ Ausstattung

Im Garten von Schloss Dagstuhl wurden an der Mauer zur Oktavie-Allee zwei alte Schultafeln, die mit der Schließung des Arnold-Janssen-Gymnasiums in St. Wendel nicht mehr benötigt wurden, montiert, so dass Gäste auch im Freien ihre Ideen an Tafeln präsentieren können. Schloss Dagstuhl dankt den Steyler Missionaren für die Spende und den Kollegen, die mit ihrem besonderen Engagement dieses neue Angebot erst möglich gemacht haben.

Als kleiner Anreiz für unsere Gäste, alle unsere Räumlichkeiten und Angebote kennen zu lernen, aber auch als kleiner Zeitvertreib, bietet Schloss Dagstuhl vor Ort die sogenannte Geisterjagd an: Über 80 kleine QR-Codes mit einem Geistersymbol sind überall im Schloss verteilt versteckt angebracht. Gäste können diese mit dem Handy scannen und sehen wöchentliche Bestenlisten.

## ■ Bibliothek

Der Bestand der Bibliothek wuchs 2024 durch mehr als 1 300 Neuanschaffungen auf 38 115 Bücher. Digital sind mehr als 7 500 deutschlandweite und internationale Zeitungen und Magazine aus über 120 Ländern in 60 Sprachen zugreifbar.

Zu den meisten Dagstuhl-Seminaren stellt Schloss Dagstuhl alle in der Bibliothek vorhandenen Bücher der anwesenden Gäste zusammen und präsentiert sie in einer separaten Auslage und auf der Bibliothekswebseite. Durchschnittlich werden mehr als 40 Bücher pro Dagstuhl-Seminar präsentiert.

Die umfangreiche Sammlung „Dagstuhls Impact“<sup>6</sup>, die publizierte Ergebnisse aus Dagstuhl-Seminaren dokumentiert, wurde 2024 um zahlreiche weitere Publikationen ergänzt: 163 Artikel, 2 Zeitschriftensonderhefte und 1 Buch konnten hinzugefügt werden.

2024 erhielt die Bibliothek von Verlagshäusern und Einzelpersonen erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 854 Bände als Spende, darunter 848 Monographien des Springer-Verlags im Wert von 70 667,63 €.

Zur Optimierung der Autorenidentifikation werden nun die ORCID-IDs zu neuen Personennamen im Bibliothekskatalog erfasst und es wurden ältere Personendaten umfassend ergänzt.

through the basement of said building. In all 27 guest rooms in this tract, the new or already prepared cabling was run to the desk, so that soon all guest rooms in the “Neubau” will have an ethernet connection at the desk.

Furthermore, the lecture hall “Kaiserslautern” was equipped with new energy-saving LED panels for lighting.

On the continuous balcony in front of the guesthouse rooms, extendable partition elements were installed so that guests can meet on the balcony if they wish, while also knowing that their privacy is protected.

## ■ Facilities

In the garden of Schloss Dagstuhl, two old school blackboards that were no longer needed after the Arnold-Janssen-Gymnasium preparatory highschool in St. Wendel closed were mounted on the wall facing Oktavie-Allee so that guests can present their ideas on blackboards, outdoors as well. Schloss Dagstuhl would like to thank the Steyler Missionaries for their donation and the colleagues whose special commitment has made this new offer possible.

As a small incentive for our guests to get to know all our facilities and offers, but also as a small pastime, Schloss Dagstuhl offers a so-called ghost hunt on site: over 80 small QR codes with a ghost symbol are hidden all over the castle. Guests can scan them with their cell phones and see weekly leaderboards.

## ■ Library

In 2024, the library’s collection grew to 38,115 books through more than 1,300 new acquisitions. Digitally, more than 7,500 German and international newspapers and magazines from over 120 countries in 60 languages are accessible.

For most Dagstuhl Seminars, books available in our library that were written by seminar participants are displayed separately and listed on the library’s website. On average, more than 40 books are presented per Dagstuhl Seminar.

In 2024, the comprehensive collection “Dagstuhl’s Impact”<sup>6</sup>, which documents published results of Dagstuhl Seminars, grew substantially: 163 articles, 2 special journal issues, and 1 book were added.

In 2024, the center’s research library again received a large number of book donations from publishing houses and individuals. The number of donated volumes in 2024 totaled 854, including 848 monographs at a total value of 70,667.63 € donated by Springer-Verlag.

To optimize the identification of authors, the ORCID-IDs are now being added to new personal names in the library catalog, and older personal data has been extensively supplemented.

<sup>6</sup> <https://www.dagstuhl.de/institute/dagstuhls-impact>

## ■ Kunst

Im sogenannten Kreuzgang des Neubaus präsentierte Schloss Dagstuhl 2024 insgesamt vier Kunstausstellungen. Von Januar bis März wurden unter dem Titel „... inzwischen“ Bilder des georgischen Künstlers Shako Berekashvili ausgestellt. Es folgten von April bis Juli unter dem Titel „inside – outside“ Bilder der saarländischen Künstlerin Tina Stein. Die beiden Künstler Annette Marx und Andreas Engel präsentierten von August bis Oktober ausgewählte Werke unter dem Titel „Marx und Engel | Atmen“. Schließlich wurden von Oktober bis Dezember unter dem Titel „Sehnsucht nach dem Norden“ in Zusammenarbeit mit dem Kulturverein Burbach e.V. Bilder des 2022 verstorbenen Künstlers Ulrich Klimmt gezeigt.

Durch großzügige Zuwendungen einzelner Spenden und dem Förderverein von Schloss Dagstuhl von insgesamt 3.440 Euro konnte Schloss Dagstuhl auch 2024 wieder Kunstwerke erwerben. Unter anderem wurden das Bild „Meditations“ von Tina Stein, das „Aquarell Nr. 33“ von Ulrich Klimmt und das Werk „Flügel“ von Shako Berekashvili gekauft. Von den Spenden wird noch ein weiteres Bild von Ulrich Klimmt erworben werden.

Weitere Informationen über das generelle Konzept und die Ausstellungen in 2024 finden sich in Kapitel 10.

## ■ Art

In 2024, Schloss Dagstuhl presented a total of four art exhibitions in the so-called cloister of the “Neubau”. From January to March, pictures by the Georgian artist Shako Berekashvili were exhibited under the title “... inzwischen”. This was followed by pictures by the saarlandian artist Tina Stein from April to July under the title “inside – outside”. The two artists Annette Marx and Andreas Engel presented selected works under the title “Marx und Engel | Atmen” from August to October. Finally, under the title “Sehnsucht nach dem Norden” pictures by the artist Ulrich Klimmt, who died in 2022, were shown in collaboration with the Kulturverein Burbach e.V. from October to December.

Thanks to generous donations from individuals and the Friends of Dagstuhl, totalling 3,440 euros, Schloss Dagstuhl was able to acquire further works of art in 2024. Among others, the picture “Meditations” by Tina Stein, the “Aquarell Nr. 33” by Ulrich Klimmt and the work “Flügel” by Shako Berekashvili were purchased. Another painting by Ulrich Klimmt will be purchased from the donations.

Further information about the art program in general and the exhibitions in 2024 can be found in Chapter 10.



# **2** **Seminare und Workshops** *Seminars and Workshops*

## Dagstuhl-Seminare

### 2.1

## Dagstuhl Seminars

Die Dagstuhl-Seminare haben als wesentliches Instrument der Forschungsförderung Priorität bei der Gestaltung des Jahresprogramms. Hauptziel der Seminare ist die Unterstützung der Kommunikation und des Dialogs zwischen Wissenschaftlern, die an den Forschungsfronten von miteinander verknüpften Forschungsfeldern in der Informatik arbeiten. Die Seminare ermöglichen die Vorstellung neuer Ideen, die Diskussion von aktuellen Problemen sowie die Weichenstellung für zukünftige Entwicklungen. Sie bieten außerdem die Möglichkeit zum Austausch zwischen vielversprechenden Nachwuchswissenschaftlern und internationalen Spitzenforschern in einem speziellen Forschungsgebiet.

Die Teilnahme an den üblicherweise einwöchigen Seminaren ist nur auf persönliche Einladung durch Schloss Dagstuhl möglich. Das Zentrum übernimmt einen Teil der Kosten, sodass die besten Wissenschaftler einschließlich jungen Forschenden und Promovierenden teilnehmen können. Zu den ehemaligen Gästen zählen 29 Preisträger des Turing-Awards, der höchsten Auszeichnung im Bereich der Informatik auf internationaler Ebene.

Charakteristisch für Dagstuhl ist die Etablierung von richtungsweisenden sowie gebietsübergreifenden Seminaren. Manche Themen, die ausgiebig in Dagstuhl diskutiert wurden, entwickelten sich anschließend zu sehr aktiven Forschungsbereichen, die teilweise zu DFG-Schwerpunkten und anderen Förderprogrammen führten. Bei einer Reihe von Forschungsgebieten wurden durch Dagstuhl-Seminare Gruppen zusammengeführt, die zwar an verwandten Problemen und Verfahren forschen, denen aber bisher keine gemeinsame Diskussionsplattform zur Verfügung stand. Dies gilt insbesondere auch für Disziplinen, die nicht zur Informatik gehören. Wichtige Forschungsgebiete, für die in Dagstuhl bereits mehrfach eine intensive Zusammenarbeit mit der Informatik erschlossen und vertieft wurde, sind Sport (seit 2006) und Musik (seit 2009). Die Themen der Dagstuhl-Seminare bieten eine hervorragende und sehr breite Übersicht über die aktuellen Forschungsgebiete der Informatik.

Jedes Dagstuhl Seminar wird gebeten, einen kurze Dokumentation zu erstellen, die eine Zusammenfassung des Seminarverlaufs, eine Kurzübersicht über die gehaltenen Vorträge und eine Zusammenfassung grundsätzlicher Ergebnisse enthält. Diese Berichte, die in der Zeitschrift *Dagstuhl Reports* veröffentlicht werden, gewährleisten eine hohe Sichtbarkeit und eine zeitnahe Kommunikation der Ergebnisse. *Dagstuhl Reports* wird jährlich in einem Band mit 12 Ausgaben veröffentlicht. Jede Ausgabe dokumentiert jeweils die Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops eines Monats. Die *Dagstuhl Reports* sind über die Dagstuhl-Website frei zugänglich.<sup>7</sup>

Kapitel 6 enthält Zusammenfassungen der Dagstuhl-Seminare und Perspektiven-Workshops. Im Kapitel 14 sind alle Veranstaltungen, die 2024 stattfanden, aufgelistet. Auf der Dagstuhl-Website<sup>8</sup> ist das gesamte bereits geplante Programm verfügbar.

Dagstuhl Seminars, the center's key instrument for promoting research, are considered top priority in its annual program. The central goal of the Dagstuhl Seminar program is to stimulate new research by fostering communication and dialogue between scientists working on the frontiers of knowledge in interconnected fields related to informatics. New ideas are showcased, topical problems are discussed, and the course is set for future development in the field. The seminars also provide a unique opportunity for the exchange of research views and findings in a specific cutting-edge field of informatics between promising young scientists and the international elite of the research area.

Participation in these events – which generally last one week – is possible only by way of personal invitation from Schloss Dagstuhl. The center assumes part of the associated costs in order to enable the world's most qualified scientists, including young researchers and doctoral students, to participate. Among Dagstuhl's past guests are 29 winners of the ACM Turing Award, the highest achievable award within the international computer science community.

A distinguishing accomplishment of Dagstuhl is the establishment of pioneering, interdisciplinary seminars that have virtually become institutions themselves. Many of the topics addressed in-depth at Dagstuhl have subsequently developed into highly active research fields, resulting in some cases in DFG priority programs and other grant and funding programs. Dagstuhl Seminars often succeed in bringing together scientists from a range of research areas and disciplines whose work overlaps with respect to issues, methods, and/or techniques, but who have never previously entered into constructive dialogue with one another. This especially applies to disciplines outside of the field of informatics. Key research areas for which in-depth collaboration with informatics specialists was initiated and consolidated at Dagstuhl include sports (since 2006) and music (since 2009). The spectrum of seminar topics provides an excellent and broad overview of the areas currently under discussion in the informatics arena.

Each Dagstuhl Seminar is asked to contribute a record of the seminar proceedings in the form of a Dagstuhl Report. The report gives an overview of the seminar's program, talks, and results in a journal-like manner to allow for a high visibility and timely communication of its outcomes. The periodical *Dagstuhl Reports* is published in one volume with 12 issues per year; each issue documents the Dagstuhl Seminars and Dagstuhl Perspectives Workshops of a given month. *Dagstuhl Reports* are openly accessible and can be downloaded from the Dagstuhl website.<sup>7</sup>

Chapter 6 contains a collection of the summaries of the 2024 Seminars and Perspectives Workshops. Chapter 14 provides a comprehensive list of all events that took place during 2024, and the entire planned seminar program is available on the Dagstuhl website<sup>8</sup>.

<sup>7</sup> <https://www.dagstuhl.de/dagrep/>

## Dagstuhl-Perspektiven-Workshops

2.2

In Ergänzung zu den Dagstuhl-Seminaren werden Dagstuhl-Perspektiven-Workshops veranstaltet, bei denen 25–30 ausgewiesene Forschenden ein bereits fest etabliertes Forschungsgebiet betreffende Tendenzen und neue Perspektiven der weiteren Entwicklung dieses Gebietes diskutieren. Im Gegensatz zu Dagstuhl-Seminaren werden statt aktueller Forschungsergebnisse im Wesentlichen Positionspapiere vorgetragen, welche den aktuellen Stand des Gebietes, offene Probleme, Defizite und vielversprechende Richtungen beschreiben. Der Fokus in den Workshops liegt auf Teilgebieten oder mehreren Gebieten der Informatik. Jeder Workshop hat zum Ziel

- den Stand eines Gebietes zu analysieren,
- Potenziale und Entwicklungsperspektiven bestehender Forschungsfelder zu erschließen,
- Defizite und problematische Entwicklungen aufzudecken,
- Forschungsrichtungen aufzuzeigen und
- Innovationsprozesse anzustoßen.

Der Dagstuhl-Perspektiven-Workshop, der 2024 statt fand, ist in Fig. 2.1 aufgelistet.

Die Ergebnisse der intensiven Diskussionen werden in einem Manifest zusammengefasst, welches die offenen Probleme und die möglichen Forschungsperspektiven für die nächsten 5–10 Jahre aufzeigt. Die Manifeste werden in unserer Fachzeitschrift *Dagstuhl Manifestos*<sup>9</sup> veröffentlicht.

Eine Liste der vergangenen und kommenden Dagstuhl-Perspektiven-Workshops ist auf der Dagstuhl-Website verfügbar.<sup>10</sup>

## Dagstuhl Perspectives Workshops

2

In addition to the traditional Dagstuhl Seminars, the center organizes Dagstuhl Perspectives Workshops. A Perspectives Workshop involves 25–30 internationally renowned senior scientists who wish to discuss strategic trends in a key research area that is already well established and to develop new perspectives for its future evolution. In contrast to Dagstuhl Seminars, Perspectives Workshops do not address current research results but reflect the overall state of a field, identifying strengths and weaknesses, determining promising new developments, and detecting emergent problems and synergies. The workshops tend to focus on subfields or are interdisciplinary in nature, thus covering more than one informatics field. Each workshop aims to

- contribute to an analysis of the present status of a field,
- tap into potentials and development perspectives of existing fields of research,
- detect shortcomings and problematic developments,
- show research directions, and
- trigger innovation processes.

The Dagstuhl Perspectives Workshops held in 2024 are listed in Fig. 2.1.

The results of the in-depth discussions of each workshop are presented in a manifesto detailing open issues and possible research perspectives in that specific field for the coming 5–10 years. The manifestos are published in our periodical *Dagstuhl Manifestos*<sup>9</sup>.

A list of past and upcoming Dagstuhl Perspectives Workshops can be found on our website.<sup>10</sup>

## Einreichung der Anträge und Begutachtungsverfahren

2.3

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch Auswahl der Anträge gewährleistet, die aus Sicht von Schloss Dagstuhl das größte Potential haben, abseits

## Proposal Submission and Review Process

Schloss Dagstuhl maintains the high quality of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop series by identifying those proposals that promise a high potential to engage researchers – often from different

<sup>8</sup> <https://www.dagstuhl.de/seminars/seminar-calendar>

<sup>9</sup> <https://www.dagstuhl.de/dagman>

<sup>10</sup> <https://www.dagstuhl.de/pw-list>

Conversational Agents: A Framework for Evaluation (CAFE)

<http://www.dagstuhl.de/24352>

Reframing Technical Debt

<http://www.dagstuhl.de/24452>

Human in the Loop Learning through Grounded Interaction in Games

<http://www.dagstuhl.de/24492>

Fig. 2.1

Dagstuhl Perspectives Workshops held in 2024.

etablierter Konferenzen neue und wichtige Forschungsprobleme mit Wissenschaftlern aus oft unterschiedlichen Gebieten zu identifizieren und zeitgleich mögliche Methoden und Lösungsansätze zu diskutieren.

Das Zentrum erbittet zweimal im Jahr Themenvorschläge von führenden Wissenschaftlerinnen und Wissenschaftlern aus der ganzen Welt, die ihre Seminaranträge zusammen mit einer vorläufigen Teilnehmerliste einreichen. Die Anträge werden dann vom Wissenschaftlichen Direktorium (siehe Kapitel 2.9) begutachtet und abschließend bei zweitägigen Sitzungen auf Schloss Dagstuhl intensiv diskutiert und über sie entschieden.

Es wird sicher gestellt, dass jedes Dagstuhl-Seminar durch ein starkes Organisatorenteam betreut wird, ein für die Informatik-Community relevantes Thema anspricht, ein kohärentes und gut strukturiertes wissenschaftliches Programm präsentiert und eine Gruppe von geeigneten Teilnehmerinnen und Teilnehmern zusammenbringt, deren kollektive Fachkenntnis einen bedeutenden Durchbruch in dem betreffenden Forschungsfeld ermöglichen kann. Zudem wird auf eine ausgeglichene Repräsentation wissenschaftlicher Gemeinden, geographischer Regionen und besonders auf das Miteinbeziehen junger und weiblicher Wissenschaftler geachtet.

In 2024 wurden mit 172 Anträgen mehr Anträge denn je auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops gestellt, was ein ungebrochen hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Informatikforscherinnen und Informatikforscher zeigte. Davon wurden 96 Anträge, was 56 % entspricht, angenommen. In den vergangenen 7 Jahren variierte die Rate der angenommenen Anträge zwischen 56 % und 69 %. (siehe Fig. 2.2).

Unter den neu genehmigten Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops gab es wie in den vergangenen Jahren wieder verschiedene Konstellationen bzgl. Dauer und Größe (vgl. Fig. 2.3). Der Großteil von 96 neu genehmigten Seminaren wurde für das Seminar-Programm in 2025 eingeplant (hier und im Folgenden wird, sofern nicht anders angegeben, das Wort "Seminar" sowohl für Dagstuhl-Seminare als auch für Dagstuhl-Perspektiven-Workshops verwendet). Jedoch werden 13 der genehmigten Seminare erst in 2026 stattfinden. Damit erhöhte sich die Vorlaufzeit auf 9 bis 21 Monate statt der intendierten 6 bis 18 Monate. Die Vorlaufzeit ist die Zeit zwischen der Sitzung, in der ein Seminar genehmigt wurde, bis zum Anfang des Seminars. Zwischen Beantragung und Sitzung liegen nochmals etwa 8 bis 10 Wochen.

disciplines – in scientific discussions on new and important research problems and their most promising solutions, outside of the existing conferences.

The center solicits topics for new seminars and workshops twice a year from leading researchers worldwide, who submit their proposals together with a list of potential scientists to be invited. The proposals and suggested invitee lists are then reviewed by Dagstuhl's Scientific Directorate (see Section 2.9) and finally discussed and decided during a two-day meeting at Schloss Dagstuhl.

This process ensures that every Dagstuhl Seminar and Dagstuhl Perspectives Workshop is backed by a strong team of organizers, addresses a topic of relevance to the computer science community, presents a coherent and well-structured scientific agenda, and brings together the right group of participants whose collective expertise can lead to a significant breakthrough in the area to be addressed. The balance of research communities and geographical regions, and especially the inclusion of junior and female researchers, are also taken into account during the review process.

In 2024, there were more proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops than ever, with a record of 172 proposals. This demonstrates a sustained high interest in organizing Dagstuhl Seminars and Dagstuhl Perspectives Workshops by computer science researchers. Of the 172 proposals, 96 were accepted, which corresponds to 56 %. In the past 7 years, the rate of accepted proposals varied from 56 % to 69 % (see Fig. 2.2).

Among the newly accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops of 2024, there was – as in the previous years – a wide variation in terms of duration and size (see Fig. 2.3). The majority of the 96 newly approved seminars were included in Dagstuhl's seminar-program for 2025 (here and in the following, the word "seminar" is meant to include both Dagstuhl Seminars and Dagstuhl Perspectives Workshops, if not specified otherwise). However, 13 of the approved seminars will take place only in 2026. Thus, the lead time, increased to 9 to 21 months instead of the intended period of 6 to 18 months. The lead time is the time between the evaluation meeting during which the seminar was approved and the start of the seminar. The time between proposal submission and the evaluation meeting is another 8 to 10 weeks.

## Seminar-Programm 2024

### 2.4

## The Seminar Program in 2024

Grundsätzlich kann Schloss Dagstuhl in jeder Woche zwei Seminare mit insgesamt maximal 70 Teilnehmenden beherbergen – wobei in Ausnahmefällen auch Gäste in naheliegenden Hotels untergebracht werden könnten, sollte eine Überbuchung stattgefunden haben. Durch Belegung von Schloss Dagstuhl durch Veranstaltungen außerhalb des

In principle, Schloss Dagstuhl can host two seminars each week with a total of about 70 participants – although in exceptional cases, guests may also be accommodated in nearby hotels in the case of overbooking. Due to other events at Schloss Dagstuhl outside of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop program,




Fig. 2.2 Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2018–2024.




Fig. 2.3 Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2018–2024. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

Dagstuhl-Seminar und Dagstuhl-Perspektiven-Workshop Programms, wie z. B. die Lehrerfortbildung, GI-Dagstuhl-Seminare und Sommerschulen, sieht Schloss Dagstuhl aber 80 Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops im Jahr als Maximalbelegung an.

In 44 von 49 Wochen, in denen das Tagungszentrum 2024 geöffnet war, fand mindestens ein Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop statt. In 33 Wochen waren es sogar zwei. In 5 Wochen war das Zentrum nur durch andere Veranstaltungen belegt. Es blieb keine Woche unbelegt.

2024 fanden insgesamt 77 Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops statt. In Fig. 2.5 ist die Entwicklung der vergangenen Jahre dargestellt.

like, e.g., the teacher training, GI-Dagstuhl Seminars, and summer schools, Schloss Dagstuhl considers 80 Dagstuhl Seminars and Dagstuhl Perspectives Workshops a year as the maximal capacity.

At least one Dagstuhl Seminar or Dagstuhl Perspectives Workshop was held in 44 of the 49 weeks for which the center was open in 2024. In 33 of those weeks, there were in fact two seminars in parallel. In 5 weeks, there were exclusively other types of events scheduled. There were no weeks without scheduled events.

Altogether, there were 77 Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2024. Fig. 2.5 shows the development in recent years.

## Angaben zu Teilnehmenden und Organisierenden

2.5

## Participant and Organizer Data

Insgesamt nahmen 3241 Wissenschaftler und Wissenschaftlerinnen an Veranstaltungen teil, davon 2556 an Dagstuhl-Seminaren und Perspektiven-Workshops. Eine reguläre remote Teilnahme an Seminaren war 2024 nicht mehr möglich. In einigen wenigen ausgewählten Fällen nahmen aber Teilnehmer an einer Sitzungen des Seminars virtuell teil oder Experten gaben ihr Wissen in einem Videomeeting weiter. Grundsätzlich werden hier und im folgenden alle Teilnehmende gezählt, unabhängig davon, ob sie vor Ort oder online teilgenommen haben. Einen Eindruck, wieviele Teilnehmende vor Ort waren bzw. remote teilgenommen haben, vermittelt Fig. 2.4.

A total of 3241 scientists took part in events, including 2556 in Dagstuhl Seminars and Perspectives Workshops. Regular remote participation in seminars was no longer possible in 2024. In a few selected cases, there were remote participants in a session of a seminar, or an expert that shared his knowledge using the video meeting technology. In general, all participants are counted here, regardless of whether they participated on site or online. Fig. 2.4 illustrates how many participants were on site and how many participated remotely.

Every year, a good mix of scientists with Dagstuhl experience and scientists who have not been to Dagstuhl before

In jedem Jahr nehmen ein guter Mix von Dagstuhl-erfahrenen Wissenschaftlern und Wissenschaftlern, die bisher nicht in Dagstuhl waren, an den Seminaren teil. Erfahrene Forschende, die sich in den vertretenen Wissensgebieten einen exzellenten Ruf erarbeitet haben, nehmen oft nicht zum ersten Mal an einem Dagstuhl Seminar teil. Gerade junge Wissenschaftler, aber auch etablierte Forschende, teils auch aus ganz neuen Gebieten, sind häufig für ihre Seminarteilnahme zum ersten Mal in Dagstuhl. Dabei kann es sogar vorkommen, dass in einem Seminar kein Organisator schon mal in Dagstuhl gewesen ist und evetuell auch (fast) alle Teilnehmenden Schloss Dagstuhl das erste Mal kennen lernen. Gerade diesen Seminaren bietet Schloss Dagstuhl Hilfe bei der Seminarvorbereitung an.

Im Jahr 2024 nahmen mit 48 % (1 187 von 2 452) ein etwas höherer Anteil der Wissenschaftler und Wissenschaftlerinnen das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während mit 17 % der Anteil der Teilnehmenden, die bisher nur einmal in Dagstuhl waren, sich etwas leichter erhöht hat. Mit 8 % ist der Anteil derjenigen, die genau zweimal vorher in Dagstuhl waren, ein wenig gesunken.

Ein beträchtlicher Anteil der Gäste besteht aus jungen Forschenden, die am Anfang ihrer Karriere stehen, und für die der Aufenthalt in Dagstuhl oftmals prägend ist für den weiteren Verlauf ihres Lebenswegs. Etwa 32 % der Gäste der Seminare im Jahr 2024, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als wissenschaftlichen Nachwuchs ein (siehe Fig. 2.6b). Wie in fast allen Vorjahren hatte Schloss Dagstuhl mit etwa einem Drittel eine ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern, die im Laufe der Jahre (mit Ausnahme des Pandemiejahres 2020) relativ konstant geblieben war, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlerinnen und Wissenschaftlern und ihren erfahrenen Kollegen und Kolleginnen zeigt.

Mit 78 % war der Anteil von an Seminaren Teilnehmenden aus dem Ausland 2024 wieder auf einem gewohnt hohen Niveau. Das Diagramm in Fig. 2.6c zeigt die regionale Verteilung der Gäste für 2024 bei Seminaren. Mehr Details können Kapitel 13 entnommen werden.

In 2024 waren etwa 95 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt und rund 36 % aller Organisatoren waren Frauen (siehe Fig. 2.7a). Der Anteil an weiblichen Seminarteilnehmern war mit 26 % erfreulich hoch (siehe Fig. 2.7b).

take part in the seminars. Experienced researchers, who have built an excellent reputation in the represented areas of science, have often been to Dagstuhl before. Especially young scientists, but also established researchers, sometimes from completely new areas, are often in Dagstuhl for the first time. It can even be the case that none of the organizers of a given seminar has been to Dagstuhl before, and possibly even (almost) all of the participants get to know Dagstuhl for the first time. Such seminars are especially offered help by Schloss Dagstuhl with the seminar's preparation.

In 2024, a slightly lower proportion of 48 % of the scientists took part in a Dagstuhl Seminar or Dagstuhl Perspectives Workshop for the first time (1,187 of 2,452 scientists), while the proportion of participants who had only been to Dagstuhl once before increased only slightly to 17 %. With a value of 8 %, the proportion of those who have been to Dagstuhl exactly twice decreased a bit.

A substantial number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of life-long value. Approximately 32 % of the 2024 Dagstuhl Seminar and Dagstuhl Perspectives Workshop survey respondents self-classified as junior (see Fig. 2.6b). As in almost all previous years (with the exception of 2020, the first year of the Covid-19 pandemic), Schloss Dagstuhl had a balanced distribution between junior and senior researchers of about one-third, reflecting the center's determined effort to maintain the “Dagstuhl connection” between brilliant junior scientists and their senior colleagues.

At around 78 %, the proportion of guests with a non-German affiliation in Dagstuhl Seminars was again at its usual high level. The chart in Fig. 2.6c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2024. For a detailed breakdown, please refer to Chapter 13.

In 2024, about 95 % of all organizer teams in our scientific seminar program were mixed with respect to gender and about 36 % of all organizers were women (see Fig. 2.7a). The proportion of female seminar participants was considerably high at 26 % (see Fig. 2.7b).


## Themen und Forschungsgebiete

### 2.6

## Topics and Research Areas

Die thematischen Schwerpunkte der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden von den internationalen Antragstellern identifiziert und dem wissenschaftlichen Direktorium zur Durchführung vorgeschlagen. Hierdurch wird die internationale Forschungscommunity aktiv in die Programmgestaltung eingebunden – zugleich ist gewährleistet, dass aufgrund der Expertise der


The topics of Dagstuhl Seminars and Dagstuhl Perspectives Workshops are identified by researchers from all over the world, who pass on this information to the Schloss Dagstuhl Scientific Directorate in their submitted proposals. The international research community is thus actively involved in shaping Dagstuhl's scientific seminar program, and their expertise ensures that the most important cutting




(a) Distribution of participants of seminars in group A.

(b) Distribution of participants of events in group B.

**Fig. 2.4**  
**Number of participants by attendance status and group.** Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops. Group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).




**Fig. 2.5**  
**Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2018–2024.**  
 Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.




(a) Distribution of the number of previous attendances of participants, according to survey data.

(b) Percentage of junior researchers, according to survey data.



(c) Distribution of the origin of participants by region.

**Fig. 2.6**  
**Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2018–2024.**



(a) Female organizers and organizer teams including women.

(b) Female invitees and participants.

Fig. 2.7

### Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2018–2024.

Antragstellenden in ihren jeweiligen Forschungsgebieten immer brandaktuelle Themenschwerpunkte gesetzt werden.

Im Folgenden sind beispielhaft einige thematische Schwerpunkte und dazugehörige Seminare aufgeführt. Die Aufzählung der Themen und Seminare hat keinen Anspruch auf Vollständigkeit und ist lediglich ein Versuch, einen kurzen Einblick in das facettenreiche Seminar-Programm zu geben. Kapitel 6 bietet mit den Kurzzusammenfassungen der Seminare und Perspektiven-Workshops einen vollständigen Überblick über das wissenschaftliche Seminar-Programm des Jahres 2024.

Wie jedes Jahr deckten unsere Dagstuhl-Seminare große Teile der Informatik ab. Von den theoretischen Grundlagen, etwa *Triangulations in Geometry and Topology* (24072) und *From Proofs to Computation in Geometric Logic and Generalizations* (24021) über mehr angewandte Themen wie *Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis* (24241) und *Regular Expressions: Matching and Indexing* (24472) bis zu praktisch relevanten Themen wie *Computational Analysis and Simulation of the Human Voice* (24242) und *Safety Assurance for Autonomous Mobility* (24071) war wieder alles vertreten. Anwendungen und Grundlagen vom omnipräsenten Thema Machine Learning waren auch Bestandteil einiger Seminare wie zum Beispiel *Automated Machine Learning for Computational Mechanics* (24282) und *Resource-Efficient Machine Learning* (24311). Auch gesellschaftspolitische relevante Fragen werden in Dagstuhl diskutiert, wie z. B. die Seminare *EU Cyber Resilience Act: Socio-Technical and Research Challenges* (24112) und *AI for Social Good* (24082) zeigen. Das Seminar *Reflections on Pandemic Visualization* (24091) schaute eher rückblickend auf ein Thema, während sich das Seminar *Quantum Software Engineering* (24512) mit Technologien der Zukunft beschäftigte. Gemessen an der Interdisziplinarität stach ganz besonders das Seminar *Art, Visual Illusions, and Data Visualization* (24301) heraus, dass Informatikern, Psychologen und Künstler zusammenbrachte und eine beeindruckende Kunstausstellung veranstaltete.

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der im Jahr 2024 veranstalteten Seminare wichtige Themen adressiert hat, die von den beteiligten Wissenschaftlern mit großem Enga-

edge topics are emphasized.

The following overview gives some topical focal points and a few respective seminars from 2024. Neither the list of focal points nor the list of seminars is exhaustive. It merely attempts to offer a brief insight into the multifaceted seminar program of 2024. Chapter 6 provides a complete overview of the 2024 scientific seminar program with the short summaries of the Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

As in every year, the topics of the Dagstuhl Seminars covered large parts of computer science. Everything was represented again, from the theoretical basics, such as *Triangulations in Geometry and Topology* (24072) and *From Proofs to Computation in Geometric Logic and Generalizations* (24021), to more applied topics like *Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis* (24241) and *Regular Expressions: Matching and Indexing* (24472), to practically relevant topics like *Computational Analysis and Simulation of the Human Voice* (24242) and *Safety Assurance for Autonomous Mobility* (24071). Applications and foundations of the omnipresent topic of machine learning were also a part of some seminars like for example *Automated Machine Learning for Computational Mechanics* (24282) and *Resource-Efficient Machine Learning* (24311). Questions of societal relevance are being discussed in Dagstuhl as well, as demonstrated by the seminar *EU Cyber Resilience Act: Socio-Technical and Research Challenges* (24112). The seminar *Reflections on Pandemic Visualization* (24091) looked more back while the seminar *Quantum Software Engineering* (24512) looked at developments for the future. In terms of interdisciplinarity, the seminar *Art, Visual Illusions, and Data Visualization* (24301) stood out, as it brought together computer scientists, psychologists, and artists and organized an impressive art exhibition.

This brief selection of seminars should not draw attention from the fact that each of 2024's seminars addressed important topics which were discussed by the involved researchers with great commitment and hence pushed forward the development in the individual areas.

gument diskutiert wurden und so die weitere Entwicklung in den einzelnen Gebieten wieder ein gutes Stück weitergebracht hat.

## Weitere Veranstaltungstypen

2.7

Neben den Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops finden noch weitere Veranstaltungen im Zentrum statt. Zu diesen Veranstaltungen gehören:

- GI-Dagstuhl-Seminare, die den wissenschaftlichen Nachwuchs zu einem bestimmten Thema zusammenführen. Sie werden in Kooperation mit der GI durchgeführt und von dieser sowie von Dagstuhl gefördert. Anträge auf GI-Dagstuhl Seminare werden vom Vorstand der GIBU (GI Beirat der Universitätsprofessoren) und vom Wissenschaftlichen Direktor von Schloss Dagstuhl begutachtet.
- Weiterbildungsveranstaltungen wie Sommerschulen und Lehrerfortbildungen.
- Forschungsgruppentreffen wie Klausurtagungen von Graduiertenkollegs, GI-Fachgruppen und anderen akademischen Arbeitsgruppen.
- Forschungsaufenthalte von Einzelpersonen, die sich für eine oder mehrere Wochen für intensive Studien nach Dagstuhl in Klausur zurückziehen.

## Qualitätssicherung

2.8

Schloss Dagstuhl befragt die Teilnehmenden der Dagstuhl-Seminare und der Dagstuhl-Perspektiven-Workshops mit Hilfe eines Fragebogens zu ihrer Zufriedenheit mit inhaltlichen und organisatorischen Aspekten ihres Dagstuhlbesuchs. Die Ergebnisse jedes Fragebogens werden im Haus wöchentlich allen Abteilungen zugänglich gemacht, um eine schnelle Reaktion auf Probleme und Wünsche zu erreichen. Gleichzeitig werden anonymisierte Ergebnisse von inhaltlichen Fragen den Teilnehmern eines Seminars per E-Mail mitgeteilt, typischerweise in der Woche nach ihrem Aufenthalt. So erhalten alle Rückmeldungen über den Verlauf des Seminars und Hinweise für die Organisation von zukünftigen Seminaren. In den zur Verfügung gestellten PDF-Dokumente werden die statistischen Ergebnisse mit Hilfe von aussagekräftigen Diagrammen aufbereitet.

Fig. 2.8 zeigt die Zufriedenheit der Antwortenden im Jahr 2024 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 456 Fragebögen, welche die Meinung von etwa 57,% der 2 556 Teilnehmenden repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

Als Teil des Einladungsprozesses der Dagstuhl-Seminare und der Dagstuhl-Perspektiven-Workshops wird die Liste der von dem Organisatorenteam zur Einladung vorgeschlagenen von Schloss Dagstuhl auf eine ausgewogene Zusammensetzung geprüft, bevor Schloss Dagstuhl Einladungen ausspricht. Sind nach der ersten Runde der Einladungen noch Plätze zu vergeben, bittet Schloss Dagstuhl, unter Angabe konkreter Vorgaben zur Verbesserung der Ausgewogenheit der Zusammensetzung sofern diese denn verbessert werden sollte, die Organisatoren um weitere Vorschläge für Einladungen. Mittels einer täglich aktua-

## Further Event Types

In addition to Dagstuhl Seminars and Dagstuhl Perspectives Workshops, Schloss Dagstuhl hosts a number of further events, including:

- GI-Dagstuhl Seminars bring young scholars together to discuss and learn about a specific topic. They are run and sponsored by the German Informatics Society (GI) in association with Schloss Dagstuhl. Proposals for GI-Dagstuhl Seminars are reviewed by the managing board of the GIBU (GI advisory board of computer science professors) and the Scientific Director of Schloss Dagstuhl.
- continuing education courses including summer schools and vocational training for teachers.
- research group meetings including conferences of graduate research training groups, GI specialist groups, and other academic working groups.
- research stays of scientists who wish to use the center as a retreat for several weeks in order to devote themselves to their studies undisturbed.

## Quality Assurance

The center conducts surveys of the participants of the Dagstuhl Seminars and Dagstuhl Perspectives Workshops, the questionnaire containing questions about their satisfaction with the content of the event and the organization of their visit. The results of each questionnaire are made available to all of the center's departments every week, thus enabling a quick response to issues and requests. At the same time, anonymized results of the content questions are made available to the seminar participants via e-mail, typically in the week following their stay at the center. This enables the organizers to receive feedback on how the seminar went and tips for organizing future seminars. In the pdf files with the results, the statistics are visualized using illuminative diagrams.

Fig. 2.8 shows the satisfaction of responding participants in 2024 with regard to selected aspects of their stay. The results were compiled from 1,456 questionnaires, representing the responses of about 57 % of all 2,556 participants. These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

During the invitation process for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, the organizers compile a list of proposed invitees which is reviewed by Schloss Dagstuhl to check it for a balanced composition before Schloss Dagstuhl extends invitations. Should there still be places left after the first round of invitations, Schloss Dagstuhl asks the organizers to propose more invitees, giving guidance on how to balance the composition of the seminar if needed. Via a dedicated webpage that is updated daily, Schloss Dagstuhl gives the organizers direct access to view the status of invitee replies.




Fig. 2.8 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2024. According to survey results.

lisierten Webseite bietet Schloss Dagstuhl allen Organistorenteams einen direkten Einblick in den Status der eingeladenen Gäste bezüglich Zu- oder Absage.

## Auslastung des Zentrums

### 2.9

## Utilization of the Center

Es gab 2024 insgesamt 13 666 Gasttage, wobei 11 760 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Dies entspricht dem drittbesten Ergebnis in der Geschichte Dagstuhls. Im Gegensatz zu den meisten anderen Statistiken zählen hier nur die Tage, an denen die Gäste tatsächlich vor Ort waren.

Es fanden im Berichtsjahr insgesamt 114 Veranstaltungen mit insgesamt 487 Veranstaltungstagen statt; 37 Veranstaltungen mit 124 Tagen entfielen dabei auf Veranstaltungen außerhalb des Dagstuhl-Seminar und Dagstuhl-Perspektiven-Workshop Programm. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2024 ebenso unbelegt wie 2 Wochen im Sommer, welche zu Instandhaltungs- und Verwaltungsarbeiten benötigt wurden, und zwei Wochen zwischen Weihnachten und dem neuen Jahr. In allen anderen Wochen fand mindestens eine Veranstaltung statt.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2024 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite ist ein Kalender<sup>11</sup> verfügbar, in welchem die anstehenden Veranstaltungen

In 2024, there were 13,666 overnight stays in total in 2024, with 11,760 overnight stays for a Dagstuhl Seminar or a Dagstuhl Perspectives Workshop. This is the third best result in Dagstuhl’s history. In contrast to most of the other statistics, only days in which the guests where actually on site are counted here.

The center hosted a total of 114 events over a total of 487 days. Of these, 37 events with 124 days were not part of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop program. See Chapter 13 for further details.

Weekends were kept free in 2024, as well as two weeks in the summer, which were needed for maintenance work to building facilities and administrative work, and one week at Christmas time. At least one event took place in all other weeks.

A comprehensive listing of all events at Schloss Dagstuhl in 2024, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as research meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar<sup>11</sup> of upcoming events and further information and materials on all events: past, present, and future.

eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

### ■ Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 2.9) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Der Wissenschaftliche Direktor gehört dem Wissenschaftlichen Direktorium an. Er empfiehlt dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können das bestehende Wissenschaftliche Direktorium sowie der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt dem Beirat zusammen mit dem Wissenschaftlichen Direktor.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums – mit Ausnahme der des Wissenschaftlichen Direktors – beträgt drei Jahre. Sie beginnt am 1. November des Jahres ihrer Berufung und endet drei Jahre später am 31. Oktober. Wiederberufung ist möglich.

### ■ Scientific Directorate

The Scientific Directorate (see Fig. 2.9) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In its biannual directorate meetings, the Scientific Directorate discusses the proposals and decides which of them to accept or reject.

The Scientific Director is a member of the Scientific Directorate. He recommends to the Supervisory Board the number of Scientific Directorate members. Candidates for the Scientific Directorate may be suggested not only by the shareholders, but also by the Scientific Directorate and the Scientific Advisory Board. The selection of candidates, which are recommended to the Supervisory Board for appointment, is carried out by the Scientific Advisory Board together with the Scientific Director.

The term of office of Scientific Directorate members is three years – with the exception of the Scientific Director. It begins on November 1 of the year of appointment and ends three years later on October 31. Reelections are possible.

---

<sup>11</sup> <https://www.dagstuhl.de/seminars/seminar-calendar/>

Wissenschaftliches Direktorium   Scientific Directorate
Prof. Dr. Elisabeth André Universität Augsburg, Germany
Prof. Dr.-Ing. Franz Baader TU Dresden, Germany
Goetz Graefe, Ph. D. Google Inc., Madison, United States of America
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany
Prof. Dr. Barbara Hammer Universität Bielefeld, Germany
Prof. Dr. Lynda Hardman Centrum Wiskunde & Informatica (CWI), Amsterdam and University of Utrecht, The Netherlands
Prof. Dr.-Ing. Holger Hermanns Universität des Saarlandes, Saarbrücken, Germany   <i>tenure started in October 2024</i>
Dr. Steve Kremer Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France
Prof. Rupak Majumdar, Ph. D. Max Planck Institute for Software Systems, Kaiserslautern, Germany
Prof. Dr.-Ing. Heiko Mantel TU Darmstadt, Germany
Prof. Dr. Lennart Martens Ghent University, Belgium
Prof. Dr. Albrecht Schmidt Ludwig-Maximilians Universität München, Germany
Prof. Dr.-Ing. Wolfgang Schröder-Preikschat Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Heike Wehrheim Carl von Ossietzky Universität Oldenburg, Germany
Prof. Dr. Verena Wolf Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Martina Zitterbart Karlsruher Institut für Technologie, Germany

Fig. 2.9  
Scientific Directorate.



# 3

**Bibliographiedatenbank dblp**

*dblp computer science bibliography*

## Offene Bibliographiedaten für die Informatik

3.1

## Open Bibliographic Data in Computer Science

Moderne Informatik-Forschung benötigt den unmittelbaren und umfassenden Zugriff auf aktuelle Publikationen, um den Bedürfnissen in einer immer komplexer werdenden Forschungslandschaft gerecht zu werden. Auch bei der Einschätzung von Forschungsleistung ist die Verfügbarkeit verlässlicher Publikationsdaten unverzichtbar. Hoch qualitative und vollständige Metadaten sind in der Regel jedoch nur sehr schwer zu erhalten. Freie Suchmaschinen wie etwa Google erlauben einen weiten Einblick in das Internet, besitzen aber keinerlei Qualitätsgarantien oder semantische Organisation. Kommerzielle Datenbanken verkaufen Metadaten als teure Dienstleistung, weisen aber in vielen Fachdisziplinen (wie etwa in der Informatik) nur eine mangelhafte Abdeckung und eine oft ungenügende Datenqualität auf. Insbesondere die einzigartige Publikationskultur der Informatik mit ihrem Schwerpunkt auf Konferenzpublikationen bleibt dabei unberücksichtigt. Universitäten und außeruniversitäre Forschungseinrichtungen bemühen sich oftmals mit immensem personellen und finanziellen Aufwand und unter Belastung der einzelnen Forschenden, eigene Daten zu erheben. Diese Datensätze weisen jedoch zwangsläufig einen lokalen Einschlag auf und vermögen es nicht, ein detailliertes Bild einer Forschungsdisziplin als Ganzes zu zeichnen.

Die *dblp computer science bibliography* leistet auf diesem Gebiet nun bereits seit über 30 Jahren einen substanziellen Beitrag durch die offene Bereitstellung qualitätsgeprüfter und aufbereiteter Publikationsdaten für die gesamte Informatik. Dabei unterstützt dblp die Informatik-Forschung auf gleich mehreren Ebenen, etwa durch:

- Unterstützung der täglichen Forschungsarbeit, etwa bei der Literaturrecherche und dem Bezug von verfügbaren Volltexten
- Unterstützung des wissenschaftlichen Publikationsprozesses durch die Bereitstellung normierter bibliographischer Referenzdaten
- Unterstützung von Forschenden und Institutionen bei der Berichtspflicht durch die Sammlung und Aufbereitung von qualitätsgesicherten Publikationslisten
- Unterstützung von Forschungsfördernden und Entscheidungstragenden durch das öffentliche Verfügbarmachen von nach Daten-Facetten aufgeschlüsselten Publikationsnachweisen
- Ermöglichen einer Betrachtung der Informatik-Disziplin als Ganzes durch die Bereitstellung semantisch reichhaltiger Daten in Forschungs-Wissensgraphen

Darüber hinaus ist der dblp-Datensatz selbst Untersuchungsgegenstand mehrerer tausend Fachartikel.<sup>12</sup> Insgesamt ist dblp daher für die Informatik sowohl als Recherche-Tool, aber auch als Forschungsdatensatz, unverzichtbar geworden.

Modern computer science research requires immediate and comprehensive access to current publications to meet the needs of an ever faster evolving and ever more complex research landscape. Not only in the everyday work of a researcher but also in the assessment of research performance, the availability of reliable bibliographic metadata has become indispensable. However, high-quality and complete metadata is very difficult to obtain. Free search engines like Google allow a broad insight into the Internet but have neither guarantees of quality nor any semantic organization. Commercial databases sell metadata as an expensive service, but in many disciplines (such as in computer science) their coverage is insufficient and the data quality is quite poor. In particular, the unique publication culture of computer science with its emphasis on conference publications remains disregarded. Most universities and non-university research institutions endeavor to collect their own data, yet often consume enormous human and financial resources and impose a burden on the individual researchers. However, these local data sets do inevitably have a local bias and are not suited to draw a detailed picture of a research discipline as a whole.

For over 30 years now, the *dblp computer science bibliography* has substantially contributed to solving this dilemma in the field of computer science by providing open, quality-checked, and curated bibliographic metadata. The dblp web service supports the computer science research community on several levels, for example by:

- supporting researchers in their daily work, e.g., when reviewing the literature or searching for full-text research articles
- supporting the scientific publication process by providing standardized bibliographic reference data
- supporting researchers and institutions in their reporting duties by collecting and editing quality-assured bibliographies
- supporting research funders and decision-makers, e.g., by providing publicly available and explorable bibliographic references
- enabling the study of the computer science discipline as a whole; e.g. by providing semantically meaningful data within a research knowledge graph

In addition, the dblp data set itself is object of study of several thousand research articles.<sup>12</sup> Hence, dblp has become indispensable to the computer science community as both a research tool and a research data set.

<sup>12</sup> Google Scholar liefert zum Suchbegriff „dblp“ über 55 700 Treffer, Semantic Scholar liefert 1 500; im Einzelnen weisen SpringerLink ca. 4 500 Artikel, Elsevier ScienceDirect über 1 750 Artikel, die ACM Digital Library ca. 7 600 Artikel und IEEE Xplore über 4 200 Artikel nach. The search term „dblp“ results in more than 55 700 hits at Google Scholar and 1 500 hits at Semantic Scholar; in particular, SpringerLink lists about 4 500 articles, Elsevier ScienceDirect lists more than 1 750 articles, the ACM Digital Library lists 7 600 articles, and IEEE Xplore lists more than 4 200 articles.

## Schloss Dagstuhl und dblp

3.2

## Schloss Dagstuhl and dblp

3

Bereits seit Ende 2010 engagiert sich Schloss Dagstuhl für die ursprünglich an der Universität Trier entwickelte Bibliographiedatenbank dblp. Zunächst durch ein Projekt im Leibniz-Wettbewerb gefördert, wurde die Datenbank seit Juni 2013 von Schloss Dagstuhl direkt mitfinanziert. Im Zuge der Konsolidierung der Zusammenarbeit wurden dauerhafte Personalstellen im wissenschaftlichen Stab von Schloss Dagstuhl geschaffen, die hauptamtlich mit der Betreuung und Weiterentwicklung von dblp beauftragt sind. Ein eigens gegründeter dblp-Beirat (siehe Fig. 3.1) leistet seit 2011 die wissenschaftliche Aufsicht und unterstützt das dblp-Team mit seiner Expertise.

Pünktlich zum 25-jährigen Jubiläum von dblp erfolgte im November 2018 die endgültige Staffelübergabe des Betriebes der Datenbank von der Universität Trier an Schloss Dagstuhl. Damit einhergehend wurden von Bund und Ländern weitere Mittel für den Betrieb von dblp bereitgestellt und eine neue, eigens eingerichtete Außenstelle von Schloss Dagstuhl auf dem Campus II der Universität Trier angesiedelt. Betrieb und Erforschung der Datenbank erfolgen dabei weiterhin in enger Kooperation mit dem Fach Informatikwissenschaften der Universität Trier.

Das dblp-Team besteht mittlerweile aus 10,9 Vollzeitäquivalenten, welche an der redaktionellen, technischen und wissenschaftlichen Verbesserung der Infrastruktur arbeiten. 2024 wurde mit 619 000 Einträgen eine Rekordzahl an neuen Publikationen indexiert (siehe Abschnitt 3.5). Gleichzeitig wurde ein neuer Rekord bei der Anzahl an behandelten Fehlerkorrekturfällen im Bestand aufgestellt (siehe Abschnitt 3.6). Parallel hierzu wird kontinuierlich an neuen Funktionen und Diensten wie dem neuen *dblp SPARQL Query Service* gearbeitet (siehe Abschnitt 3.3).

The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – has been existing since late 2010. The commitment of Schloss Dagstuhl to dblp, initially funded by a project of the Leibniz Competition, is being funded directly by Schloss Dagstuhl since June 2013. As part of the consolidation of this cooperation, permanent scientific staff positions – assigned full-time to the support and development of dblp – were created. The dblp advisory board (cf. Figure 3.1), established in 2011 at Schloss Dagstuhl, provides scientific supervision and supports dblp with its expertise.

In November 2018, the transfer of the database from the University of Trier to the Leibniz Center for Informatics at Schloss Dagstuhl took place just in time for dblp's 25th anniversary. At the same time, Dagstuhl's funding was increased to support the operation of dblp and a new Schloss Dagstuhl branch office for the dblp team was established on Campus II of the University of Trier. The database will continue to be operated and researched in close cooperation with the University's Department of Computer Sciences.

The dblp team now consists of 10.9 full-time equivalent staff members working on the editorial, technical, and scientific improvement of the infrastructure. In 2024, a record number of 619,000 new publications were indexed in dblp (see Section 3.5). At the same time, the dblp team handled a record number of curation cases in the existing data (see Section 3.6). In parallel to the ongoing curation work, features are continuously extended and improved and new services like the *dblp SPARQL Query Service* are implemented (see Section 3.3).

dblp-Beirat   dblp Advisory Board
Prof. Dr. Hannah Bast University of Freiburg, Germany   <i>Chair</i>
Prof. Dr. Guillaume Cabanac Paul Sabatier University, Toulouse, France
Dr. Martin Fenner Front Matter, Münster, Germany
Prof. Dr. Silvio Peroni University of Bologna, Italy
Lydia Pintscher Wikimedia Deutschland – Association for the Promotion of Free Knowledge e.V., Berlin, Germany
Prof. Dr. Ruzica Piskac Yale University, New Haven, United States of America
Prof. Dr. Rüdiger Reischuk University of Lübeck, Germany
Prof. Dr.-Ing. Ralf Schenkel University of Trier, Germany
Prof. Raimund Seidel, Ph.D. Saarland University, Saarbrücken, Germany

Fig. 3.1  
The dblp Advisory Board in 2024.

## Der dblp SPARQL Query Service

3.3

## The dblp SPARQL Query Service

Seit 2022 arbeitet das dblp-Team aktiv am Aufbau des *dblp Knowledge Graph (dblp KG)*. Als Wissensgraph bietet er eine vollständig semantische Sicht auf alle Metadaten und Beziehungen, die in der dblp computer science bibliography enthalten sind. Der dblp KG findet bereits in der Informatik-Forschungs-Community rege Verwendung, da er Aggregation und Kombination von dblp Daten mit anderen semantischen Datenquellen sehr erleichtert. Doch darüber hinaus ermöglicht er auch dem dblp-Team die Einführung eines neuen Dienstes: Im September 2024 startete der brandneue *dblp SPARQL Query Service*.<sup>13</sup>

Der dblp SPARQL Query Service ermöglicht es, den gesamten semantischen Inhalt von dblp in einer nutzungsfreundlichen Oberfläche zu erkunden. Er ermöglicht es Nutzenden, eigenen Anfragen zu erstellen und so neue Sichten zu generieren, die weit über die begrenzten Möglichkeiten der vorbereiteten dblp-Webseiten oder der textbasierten dblp-Suche hinausgehen. Darüber hinaus ergänzt der Dienst die dblp-Daten mit den offenen Zitationsdaten von OpenCitations.<sup>14</sup> Somit sind dblp-Nutzende nun erstmals in der Lage, eigene Zitationsanalysen mit dblp durchzuführen.

Der Query Service verwendet die leistungsstarken QLever SPARQL-Engine, die von Hannah Bast et al. an der Universität Freiburg entwickelt wird. QLever bietet eine intuitive Eingabemaske mit Syntax-Highlighting, Formatierungshilfe und Autovervollständigung, sowie eine öffentliche SPARQL-API.

Im Gegensatz zu früheren semantischen Suchen für dblp basiert unser Dienst nicht nur auf einer einmaligen Momentaufnahme der dblp-Daten. Stattdessen wird der Dienst täglich mit dem aktuellen Stand synchronisiert. Somit profitiert auch der Query Service von der kontinuierlichen Kuratierung und Disambiguierung des dblp-Teams.

Stand Ende 2024 kombiniert der dblp SPARQL Query Service die annähernd 500 Millionen semantische Triple-Statements des dblp KG mit mehr als 800 Millionen offenen Zitations-Statements.

Since 2022, the dblp team has been actively working on building the *dblp Knowledge Graph (dblp KG)*. As a knowledge graph, it is a fully semantic view on all the metadata and relationships that are contained in the dblp computer science bibliography. The dblp KG has already seen adoption in services and research papers of the computer science community, as it makes sharing dblp's curated data and combining it with other semantic data sources easy. But it also enabled the dblp team to launch a new tool: In September 2024, the team launched its brand new *dblp SPARQL Query Service*.<sup>13</sup>

The dblp SPARQL Query Service allows users to explore the entire semantic content of dblp in a convenient user interface. It enables users to design their own queries and new generate insights well beyond the limited capabilities of the prepared dblp web pages or the text-based dblp search. Furthermore, the query service augments dblp data with open citation data from OpenCitations.<sup>14</sup> Hence, for the first time, users are now able to conduct their own citation analyses using dblp.

The service is powered by the performant QLever SPARQL engine, developed by Hannah Bast et al. at the University of Freiburg. QLever provide an intuitive UI including syntax-highlighting, formatting, and auto-completions, as well as a public SPARQL endpoint API.

In contrast to previous approaches to semantic search on dblp, our query service is not merely based on a one-time snapshot of dblp data. Instead, it is synchronized daily with the current dblp data. That is, just like the regular dblp website, the service benefits from the continuous curation and disambiguation work of the dblp team.

As of end of 2024, the dblp SPARQL Query Service combines and provides the almost 500 million semantic triple statements from the dblp KG together with more than 800 million open citation statements.

## dblp-Daten-Dumps

3.4

## dblp dump data releases

Seit Dezember 2024 werden die monatlichen dblp XML- und RDF-Snapshots<sup>15</sup> über die Dagstuhl-eigene DROPS-Publikationsplattform veröffentlicht. Auch alle früheren Snapshot-Veröffentlichungen seit 2015 wurden auf die neue Plattform migriert.

Der Hauptvorteil der neuen Plattform besteht darin, dass nun alle bestehenden und zukünftigen Datenpakete mit DOI ausgestattet sind und in Übereinstimmung mit den FAIR-Prinzipien veröffentlicht werden. Dies wird die Zitierfähigkeit und Reproduzierbarkeit von Experimenten mit dblp-Daten in Forschungsarbeiten verbessern und die Auffindbarkeit und Interoperabilität der Daten steigern.

In December 2024, dblp started publishing its monthly XML and RDF releases<sup>15</sup> using Dagstuhl's own DROPS publishing platform. All previous snapshot releases since 2015 have also been migrated to the new platform.

The biggest advantage of the new home is that all existing and upcoming data releases are now published with DOI and in compliance with FAIR principles. This should improve the citability and reproducibility for research experiments, as well as improving the discoverability and interoperability of the data.

<sup>13</sup> <https://sparql.dblp.org>

<sup>14</sup> <https://opencitations.net>

<sup>15</sup> <https://drops.dagstuhl.de/entities/collection/dblp>




Fig. 3.2  
**Newly indexed publications in dblp by year and type.** The negative number of new *Incollection* records in 2011 results from relabeling several thousand existing records with the newly introduced *Reference* type. Similarly, in the same year, several thousand *Articles* and *Inproceedings* records have been relabeled as *Informal*.

## Statistiken der Datenakquise

3.5

## Data Acquisition Statistics

Die Bibliographiedatenbank dblp indiziert Publikationen anhand vollständiger Inhaltsverzeichnisse von Konferenzbänden oder Journalausgaben. Mit Hilfe einer eigens entwickelten Software zur Datenextraktion werden Metadaten von Verlagswebseiten ausgelesen und zur weiteren Bearbeitung vorbereitet. Die Metadaten werden anschließend vom dblp-Team redaktionell bearbeitet: Eventuelle Fehler werden korrigiert, mehrdeutige und ungenaue Angaben werden verbessert. Diese Datenpflege wird zwar von Hilfssoftware unterstützt, erfolgt aber vornehmlich händisch durch die jeweiligen Mitarbeitenden.

Dabei erwies sich 2024 erneut als das produktivste Jahr in der Geschichte der Datenbank. So wurden innerhalb von zwölf Monaten mehr als 619 000 neue Publikationen indiziert. Das entspricht mehr als 2 400 neuen Publikationen pro Arbeitstag. Die Neuaufnahmequote übertraf damit die Rekordzahl des Vorjahres um weitere 11,5%.

Die neu aufgenommenen Einträge verteilen sich zu 36,1% auf Konferenzbeiträge, zu 41,2% auf Journalartikel, zu 18,6% auf Preprints und „graue“ Literatur, zu 2,6% auf Monographien und Dissertationen, sowie zu 1,5% auf andere Publikationstypen. Am Ende des Jahres waren mehr als 7,6 Millionen Publikationen aus den verschiedenen Teilgebieten der Informatik indiziert.

Die Aufnahmerate neuer Publikationen in dblp kann Fig. 3.2 entnommen werden.

The dblp computer science bibliography indexes conferences and journals on a per-volume basis. Using dblp's own web harvesting software, bibliographic metadata of journal or proceedings volumes are extracted from the publisher's website. This metadata is diligently checked and corrected by the dblp team. The data-cleaning process is assisted by algorithms but executed almost exclusively by hand.

With more than 619,000 new publications indexed, 2024 proved to be the most productive year in the history of the database. This corresponds to more than 2,400 new records for each working day of the year. The rate of new additions thus exceeded the previous year's record figure by a further 11.5%.

This year's new records consist of 36.1% conference papers, 41.2% journal articles, 18.6% preprints and "grey" literature, 2.6% monographs and PhD theses, and 1.5% further publication types. By the end of 2024, a total of about 7.6 million publications from all subdisciplines of computer science were indexed by dblp.

The rate of new publications added to dblp can be seen in Figure 3.2.




Fig. 3.3

**Curation of existing dblp author bibliographies.** The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp author bibliographies. We distinguish between four curation cases: *Merge* = Two or more synonymous bibliographies have been merged into a single bibliography. *Split* = A single, homonymous bibliography has been split into two or more bibliographies. *Distribute* = A mixed case where records from two or more bibliographies have been redistributed between two or more bibliographies. *Rename* = A case where no actual publications have been reassigned, but the surface form of the author name(s) of a bibliography have been corrected or improved.

## Statistiken der Datenkuration

3.6

## Data Curation Statistics

Ein Hauptziel der intensiven Datenpflege des dblp-Teams ist es sicherzustellen, dass die bereitgestellten Bibliographien so korrekt und vollständig wie möglich sind. Das bedeutet, dass alle Publikationen eines Forschers in nur einer einzigen Bibliographie aufgeführt sein sollen und dass diese Bibliographie auch nur Publikationen des jeweiligen Forschers listen soll. Dies zu gewährleisten kann ziemlich herausfordernd sein, und trotz der Bemühungen des Teams werden regelmäßig Publikationen einer falschen Bibliographie zugeordnet. Aus diesem Grund wird der dblp-Datenbestand kontinuierlich von der dblp-Redaktion überprüft und eventuelle Zuordnungs- oder Daten-Fehler korrigiert.

Während spezielle Algorithmen dem Team helfen, solche Defekte aufzudecken, werden Korrekturen immer auf der Grundlage der Entscheidung eines menschlichen Kuratierenden durchgeführt. Dies ist notwendig, da die verfügbaren Metadaten in der Regel nicht genügend Informationen enthalten, um eine hochpräzise automatisierte Lösung zu erlauben. Oft ist eine manuelle Recherche unter Berücksichtigung externer Ressourcen erforderlich.

Im Laufe des Jahres 2024 wurde ein neuer Rekord von mehr als 93 000 Fehlerfällen bearbeitet. Dies ist eine Steigerung um 30,4% gegenüber dem Vorjahr. Fig. 3.3 zeigt die Anzahl der gelösten Fehlerfälle in den letzten Jahren. Hauptfaktoren für die kontinuierliche Steigerung sind verbesserte interne Werkzeuge zum Erkennen potentieller Fehler und eine verbesserte Verfügbarkeit von ORCID<sup>16</sup>-Daten. ORCIDs sind eindeutige Kennungen, die Forschende mit ihren Publikationen verbinden können. Durch den Vergleich von ORCID-Daten mit dblp-Bibliographien gelingt es, in großer Zahl Zuordnungsfehler aufzudecken.

One main goal of the intensive data curation at dblp is to ensure that our author bibliographies are as correct and complete as possible. This means that all publications of a person should be listed in a single bibliography, and that a bibliography should only list publications from that specific author. It can be quite challenging to ensure this level of quality, and despite the best efforts of the dblp team, we regularly assign publications to the wrong bibliography. Because of this, our editorial team constantly checks our data and corrects such defects.

While specialized algorithms help the team to uncover and identify the nature of defects in our data, corrections are always executed based on the decision made by a human curator. This is necessary since the available metadata usually does not carry enough information to allow for a highly precise automated solution, and often requires a manual investigation taking external resources into account.

In 2024, a new record number of more than 93,000 error cases have been handled and resolved by the dblp team. This is an increase of 30.4% over the previous year. See Fig. 3.3 for a chart of the recent years. Here, the team makes use of efficient internal tools that point out potential errors in the dblp corpus. On major component that enables these tools is the increased availability of ORCID<sup>16</sup> data. ORCID is a unique contributor identifier that authors can attach to their publications. Comparing ORCID data with our bibliographies significantly helps us to identify errors in large numbers.

<sup>16</sup> <https://orcid.org>

	Trier		Dagstuhl		Total		
	2023	2024	2023	2024	2023	2024	%
page views per day	565,541	900,013	1,184,112	1,633,244	1,749,654	2,533,258	+44.8
user sessions (visits) per day	26,469	98,681	57,142	138,991	83,611	237,672	+184.3
page views per user session	21.4	9.1	20.7	11.8	20.9	10.7	-49.1
distinct users (IPs) per month	278,876	1,191,266	692,319	1,579,931	971,196	2,771,198	+185.3
data served per month	1,165.7 GB	1,103.1 GB	1,441.8 GB	1,564.0 GB	2,607.5 GB	2,667.0 GB	+2.3

**Fig. 3.4**  
**Average usage of the public dblp web servers.** Trier = dblp.uni-trier.de, Dagstuhl = dblp.dagstuhl.de or dblp.org. In 2021, the counting method to determine distinct users has been changed in order to avoid double counting of the same IP across different servers. While these figures generally ignore accesses from recognized bots and web crawlers, please note that the numbers given for 2024 are skewed by a currently still unknown amount of automated accesses from scripted agents, as described in the text of Section 3.7.

## Nutzungsstatistiken

3.7

## Usage Statistics

Im Jahr 2024 wurden vom dblp-Team zwei offizielle dblp-Server geführt. Die Informationen auf diesen Servern werden täglich aktualisiert und miteinander synchronisiert:

- Server Trier: dblp.uni-trier.de
- Server Dagstuhl: dblp.dagstuhl.de

Die allgemeine Adresse dblp.org ist dabei ein Alias für den dblp-Server in Dagstuhl. Mittlerweile ist Server Dagstuhl der mit Abstand am meisten genutzte Server bezüglich der Anzahl an Nutzenden dank seiner hohen Platzierung bei der Google Suche.

Obwohl die dblp-Nutzungsstatistiken um alle bekannten Bot- und Crawler-Zugriffe bereinigt sind verzeichneten die dblp-Server im Jahr 2024 einen gewaltigen Anstieg. Die Anzahl täglicher Seitenaufrufe stieg dabei um mehr als 780 000 Zugriffe (+44,8%), die Anzahl der monatlich aktiven Nutzenden sogar um mehr als 1,8 Millionen IPs (+185,3%). Es gibt erheblich Zweifel daran, dass es sich hierbei um ein natürliches Wachstum der normalen dblp-Community handelt. Vielmehr wird vermutet, dass der größte Teil der Steigerung auf neuartige und bisher unerkannte automatisierte Zugriffe zurückgeht. Das dblp-Team ist aktiv damit befasst, die Situation zu analysieren und die Erkennung von automatisierten Zugriffen zu verbessern, um so die Aussagekraft und Vergleichbarkeit der Nutzungsstatistiken zu gewährleisten. Dieser Vorgang ist jedoch noch nicht abgeschlossen.

Fig. 3.4 fasst die durchschnittliche Nutzung aller dblp-Server an Hand der zum Redaktionsschluss vorliegenden Zahlen zusammen. Zukünftige Statistiken werden versuchen, die Zahlen für 2024 zu überarbeiten und bereinigt wiederzugeben.

In 2024, the dblp team maintained two official dblp servers. The data on these servers is updated and synchronized on a daily basis:

- server Trier: dblp.uni-trier.de
- server Dagstuhl: dblp.dagstuhl.de

The main domain name dblp.org is an alias for the dblp server in Dagstuhl. Today, server Dagstuhl is by far the more frequently used server in terms of the number of visitors, thanks to its high ranking on the Google Search.

Although the dblp user statistics have already been adjusted to ignore all known bot and crawler access, the dblp servers recorded a huge increase in 2024. The number of daily page views increased by more than 780,000 (+44.8%), while the number of monthly active users even increased by more than 1.8 million IPs (+185.3%). There are significant doubts that this is a natural growth of the normal dblp community. Rather, it is suspected that the majority of the increase is due to new and currently still undetected automated accesses. The dblp team is actively analyzing the situation and improving the detection of automated accesses in order to ensure the validity and comparability of the usage statistics. However, this process has not been completed yet.

Fig. 3.4 summarizes the average usage of all dblp servers based on the figures available at the editorial deadline. Future statistics will attempt to revise and adjust the figures for 2024.



# **4** **Dagstuhl Publishing** *Dagstuhl Publishing*

## Über Dagstuhl Publishing

4.1

## About Dagstuhl Publishing

Dagstuhl Publishing ist der wissenschaftliche Verlag von Schloss Dagstuhl mit einem inhaltlichen Schwerpunkt auf hochwertigen Forschungsarbeiten in der Informatik. Entsprechend dem Auftrag des LZI, die wissenschaftliche Kommunikation und den freien Austausch zwischen den Forschern zu fördern, ist Dagstuhl Publishing dem Gedanken des Open Access uneingeschränkt verpflichtet: Dies spiegelt sich in den eigenen Angeboten, aber auch im Leitgedanken der Verlagsarbeit wider, den Wandel der Publikationskultur in der Informatik hin zu Open Access und Open Science zu fördern.

Dagstuhl Publishing is the scientific publishing house of Schloss Dagstuhl with a content focus on high-quality research papers in computer science. In line with LZI's mission, namely to promote scientific communication and free exchange between researchers, Dagstuhl Publishing is fully committed to the idea of Open Access: This is reflected in its own offerings, but also in the guiding principle of its publishing activities, which is to promote the transformation of the publication culture in computer science towards open access and open science.

## Dagstuhl Publishing Beirat

4.2

## Dagstuhl Publishing Advisory Board

Um den Verlag und sein Portfolio weiterzuentwickeln und an die Bedürfnisse der Informatikerinnen und Informatiker anzupassen und dabei die Fortschritte im Open-Access- und Open-Science-Umfeld zu berücksichtigen, hat der wissenschaftliche Beirat von Schloss Dagstuhl empfohlen, ein unabhängiges Beratungsgremium für den Verlag zu gründen. Dieses Gremium wurde zu Beginn des Jahres 2023 gegründet.

Aufgabe des Verlagsbeirats ist die Unterstützung und wissenschaftliche Begleitung von Dagstuhl Publishing. Insbesondere

- berät der Beirat in allen Fragen der strategischen Entwicklung des Dagstuhl Verlags,
  - definiert die Richtlinien und Standards, die die Aufnahme neuer Reihen und Zeitschriften in das Portfolio von Dagstuhl Publishing regeln,
  - überwacht die Entwicklung von Dagstuhl Publishing,
  - und unterstützt das Verlagsteam mit seiner Expertise.
- Der Verlagsbeirat besteht aus etablierten Informatikern und Experten für Forschungsinfrastrukturen, siehe Fig. 4.1.

In order to further develop and adapt the publishing house and its portfolio to the needs of computer scientists, taking into account advancement in the Open Access and Open Science environments, Schloss Dagstuhl's Scientific Advisory Board has recommended that an independent consulting body for the publishing house be established. This committee was founded at the beginning of 2023.

The purpose of the Dagstuhl Publishing Advisory Board is to provide support and scientific supervision for Schloss Dagstuhl's publishing department. In particular, the board

- advises in all matters regarding the strategic development of Dagstuhl Publishing,
- defines the guidelines and standards regulating the inclusion of new series and journals into the portfolio of Dagstuhl Publishing,
- monitors the development of Dagstuhl Publishing,
- and supports the publishing team by providing its expertise.

The Publishing Advisory Board consists of established computer scientists and research information infrastructure experts, see Fig. 4.1.

Prof. Dr. Boudewijn Haverkort Tilburg University, the Netherlands   Chair
Prof. Dr. Kurt Mehlhorn MPI für Informatik, Saarbrücken, Germany
Isabella Meinecke Staats- und Universitätsbibliothek Hamburg, Germany
Dr. Anja Oberländer Universität Konstanz, Germany
Prof. Raimund Seidel, Ph.D. Universität des Saarlandes, Saarbrücken, Germany
Dr. Philipp Zumstein Universität Mannheim, Germany

Fig. 4.1

Dagstuhl Publishing Advisory Board.

## Portfolio

4.3

Die Open-Access-Verlagsdienstleistungen von Schloss Dagstuhl werden in der Wissenschaftsgemeinde gut aufgenommen. Im Portfolio des Angebots gibt es zum einen Publikationsserien, die sich auf Veranstaltungen beziehen, die im Schloss Dagstuhl abgehalten wurden (*Dagstuhl Reports* und *Dagstuhl Manifestos*), zum anderen Serien, die Konferenzen und Workshops außerhalb von Schloss Dagstuhl bedienen (*LIPICs* und *OASICs*). Ergänzt wird das Portfolio um die wissenschaftliche Zeitschrift *LITES*, die Serie *DARTS*, in der Forschungsartefakte veröffentlicht werden, und seit 2023 um die wissenschaftliche Zeitschrift *TGDK*.

### ■ Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*<sup>17</sup> dokumentiert, um eine Zitation der Seminare im wissenschaftlichen Kontext zu ermöglichen. Zudem bietet sie auch den Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift erscheint seit 2011 und enthält in monatlichen Ausgaben Berichte zu den Dagstuhl-Seminaren und -Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 2.9) agiert als Herausbergremium für die Reihe.

In 2024 wurde für 84 Dagstuhl-Seminare und -Perspektiven-Workshops ein Bericht in der Reihe *Dagstuhl Reports* veröffentlicht. An dieser Stelle bedanken wir uns ganz herzlich bei den Organisatoren und Kollektoren für die erfolgreiche Zusammenarbeit.

### ■ Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*<sup>18</sup> die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Dagstuhl-Perspektiven-Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 2.9) fungiert hier ebenfalls als Herausbergremium. In 2024 wurde eine Ausgabe mit einem Manifesto veröffentlicht (siehe Fig. 4.2).

<sup>17</sup> <https://www.dagstuhl.de/dagrep>

<sup>18</sup> <https://www.dagstuhl.de/dagman>

## Portfolio

4

The scientific community appreciates the Open Access publishing services offered by Schloss Dagstuhl. The portfolio covers series related to events at Schloss Dagstuhl (*Dagstuhl Reports* and *Dagstuhl Manifestos*) and series for conferences and workshops held outside of Schloss Dagstuhl (*OASICs* and *LIPICs*). The portfolio is supplemented by the scholarly journal *LITES*, by the *DARTS* series which aims at publishing research artifacts, and since 2023 by the scholarly journal *TGDK*.

### ■ Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*<sup>17</sup> which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

The periodical started with the first seminars of January 2011 and publishes, in monthly issues, reports on Dagstuhl Seminars and Perspectives Workshops that took place in a given month. The content is not peer-reviewed. The Scientific Directorate (see Fig. 2.9) acts as editorial board.

In 2024, 84 reports of Dagstuhl Seminars and Dagstuhl Perspectives Workshops have been published. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

### ■ Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*<sup>18</sup> in an Open Access manner. The Scientific Directorate (see Fig. 2.9) acts as the editorial board of the journal. In 2024 one volume with one manifesto was published (see Fig. 4.2).

Current and Future Challenges in Knowledge Representation and Reasoning (Dagstuhl Perspectives Workshop 22282)

Dagstuhl Manifestos, Volume 10, Issue 1, pp. 1–61, <https://doi.org/10.4230/DagMan.10.1.1>

based on Dagstuhl Perspectives Workshop 22282 <https://www.dagstuhl.de/22282>

Fig. 4.2

Manifestos published in the 2024 volume of the journal *Dagstuhl Manifestos*.

### ■ DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*<sup>19</sup> werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2024 wurde die zehnte Ausgabe mit zwei Heften und insgesamt 31 Artefakten veröffentlicht.

Die Veröffentlichung und Bereitstellung von Forschungsdaten und -artefakten ist aktuell ein wichtiges Thema in den wissenschaftlichen Disziplinen und bei den Forschungsfördereinrichtungen. Im Bereich der Informatik wird dieses Thema ebenfalls diskutiert. Im Jahr 2015 gab es zum Beispiel einen Perspektiven-Workshop mit dem Titel „Artifact Evaluation for Publications“<sup>20</sup>, der 2016 durch zwei Seminare ergänzt wurde: „Reproducibility of Data-Oriented Experiments in e-Science“<sup>21</sup> und „Rethinking Experimental Methods in Computing“<sup>22</sup>.

Schloss Dagstuhl unterstützt mit DARTS die Wissenschaftsgemeinde in der Informatik bei dem Wunsch, Forschungsdaten und -artefakte in einer geeigneten Reihe zu veröffentlichen. Hierbei berücksichtigt DARTS insbesondere auch die Publikationskultur in der Informatik mit ihrem Schwerpunkt auf Konferenzbandveröffentlichungen.

### ■ OASlcs: OpenAccess Series in Informatics

Die *OASlcs*-Reihe<sup>23</sup> veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausbergremium (Fig. 4.3), diskutiert sorgfältig alle Anträge, um ausschließlich qualitativ hochwertige sowie professionell durchgeführte Veranstaltungen in die Reihe aufzunehmen und um gegebenenfalls Empfehlungen zur Verbesserung der Veranstaltungsstruktur zu geben.

In 2024 wurden 10 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.7.

### ■ DARTS: Dagstuhl Artifacts Series

The *DARTS* series<sup>19</sup> publishes evaluated research data and artifacts. It is organized as a periodical. In 2024, the tenth volume containing two issues with 31 artifacts in total was published.

The publishing of research data and artifacts is currently in the general focus of the scientific community and funding agencies. In the area of computer science, this topic is also under discussion. For example, in 2015 a Perspectives Workshop on “Artifact Evaluation for Publications”<sup>20</sup> took place which was complemented with two seminars in 2016: “Reproducibility of Data-Oriented Experiments in e-Science”<sup>21</sup> and “Rethinking Experimental Methods in Computing”<sup>22</sup>.

With DARTS, Schloss Dagstuhl is aiming to support the computing research community with a publishing venue dedicated to research data and artifacts. Especially, DARTS takes into account the publication culture in computer science which focuses on conference proceedings publications.

### ■ OASlcs: OpenAccess Series in Informatics

The *OASlcs* series<sup>23</sup> aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 4.3, discusses carefully all submitted proposals to ensure that only significant and professionally organized events are added to the series and that – if applicable – suggestions are given for improving the structure of the event.

In 2024, Dagstuhl published 10 *OASlcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.7.

<sup>19</sup> <https://www.dagstuhl.de/darts>

<sup>20</sup> <https://www.dagstuhl.de/15452>

<sup>21</sup> <https://www.dagstuhl.de/16041>

<sup>22</sup> <https://www.dagstuhl.de/16111>

<sup>23</sup> <https://www.dagstuhl.de/oasics>

Prof. Dr. Daniel Cremers  
TU Munich, Germany

Prof. Dr. Barbara Hammer  
Bielefeld University, Germany

Prof. Dr. Marc Langheinrich  
University of Lugano, Switzerland

Prof. Dr. Dorothea Wagner  
Karlsruhe Institute of Technology, Germany | Chair

Fig. 4.3  
OASlcs Editorial Board.

## ■ LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs-Reihe*<sup>24</sup> veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausbergremium (siehe Fig. 4.4) besteht aus einschlägig bekannten Wissenschaftlern und wird seit Oktober 2023 von Meena Mahajan als Hauptherausgeberin geleitet.

In 2024 wurden 38 Tagungsbände von Konferenzen mit insgesamt 1648 Artikeln veröffentlicht; siehe Fig. 4.8 und 4.9.

Auch im zurückliegenden Jahr 2024 gab es wieder viele Anträge bei LIPIcs, womit die große Nachfrage aus den Vorjahren fortgesetzt wurde. In Fig. 4.5 sind alle Konferenzen aufgelistet, deren Anträge 2024 bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischerweise 5 Jahre) eingegangen wurde. Alle diese Konferenzen haben bereits vorher mit LIPIcs kooperiert.

## ■ LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs series*<sup>24</sup> publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers (see Fig. 4.4) supervises the conferences that are accepted for LIPIcs and is headed since October 2023 by Meena Mahajan.

The series published 38 proceedings of major conferences with 1648 articles in total in 2024; see Fig. 4.8 and 4.9.

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2024 has again seen several applications for LIPIcs, continuing the high interest from the previous years. Fig. 4.5 lists all conferences that have been accepted in 2024 for a cooperation covering several years (typically 5 years). All of these conferences have already cooperated with LIPIcs in the past.

<sup>24</sup> <https://www.dagstuhl.de/lipics>

Prof. Dr. Luca Aceto Reykjavik University, Iceland and GSSI, Italy	Prof. Dr. Meena Mahajan   chair Institute of Mathematical Sciences, Chennai, India
Prof. Dr. Christel Baier TU Dresden, Germany	Prof. Dr. Anca Muscholl University of Bordeaux, France
Prof. Dr. Mikolaj Bojanczyk University of Warsaw, Poland	Prof. Dr. Luke Ong University of Oxford, Great Britain
Prof. Dr. Roberto Di Cosmo Inria and Université de Paris, France	Prof. Dr. Phillip Rogaway University of California, Davis, United States of America
Prof. Dr. Faith Ellen University of Toronto, Canada	Prof. Dr. Eva Rotenberg Technical University of Denmark, Lyngby, Denmark
Prof. Dr. Javier Esparza TU München, Germany	Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Daniel Král' Masaryk University, Brno, Czech Republic	Prof. Dr. Pierre Senellart ENS, Université PSL, Paris, France

Fig. 4.4  
LIPIcs Editorial Board.

APPROX   International Conference on Approximation Algorithms for Combinatorial Optimization Problems accepted for 2024-2028 (Re-Evaluation)
CCC   Computational Complexity Conference accepted for 2025-2029 (Re-Evaluation)
CONCUR   International Conference on Concurrency Theory accepted for 2025-2029 (Re-Evaluation)
ECOOOP   European Conference on Object-Oriented Programming accepted for 2025-2029 (Re-Evaluation)
ICDT   International Conference on Database Theory accepted for 2025-2029 (Re-Evaluation)
ITP   International Conference on Interactive Theorem Proving accepted for 2024-2028 (Re-Evaluation)
RANDOM   International Conference on Randomization and Computation accepted for 2024-2028 (Re-Evaluation)
SAND   Symposium on Algorithmic Foundations of Dynamic Networks accepted for 2025-2029 (Re-Evaluation)
SoCG   Symposium on Computational Geometry accepted for 2025-2029 (Re-Evaluation)
TYPES   International Conference on Types for Proofs and Programs accepted for 2025-2029 (Re-Evaluation)
WABI   International Workshop on Algorithms in Bioinformatics accepted for 2025-2029 (Re-Evaluation)

Fig. 4.5

Conferences accepted in 2024 for publication in LIPIcs.

### ■ TGDK: Transactions on Graph Data and Knowledge

*Transactions on Graph Data and Knowledge (TGDK)*<sup>25</sup> ist eine Open-Access-Zeitschrift, die Original-Forschungsartikel und Übersichtsartikel über graphenbasierte Abstraktionen von Daten und Wissen, aber auch die Techniken, die solche Abstraktionen im Hinblick auf Integration, Abfragen, Schlussfolgerungen und Lernen ermöglichen, veröffentlicht. Die Zeitschrift wurde im Herbst 2023 gegründet und veröffentlichte ihre erste Ausgabe im Dezember 2023.

Die Zeitschrift wird als Diamond Open Access veröffentlicht, d.h. alle Inhalte sind frei zugänglich und es werden keine Gebühren für Leser oder Autoren erhoben. Die Zeitschrift wird von Aidan Hogan, Andreas Hotho, Lalana Kagal und Uli Sattler als Chefredakteure(innen) zusammen mit einem Team von rund 100 assoziierten Redakteurinnen und Redakteuren und Redaktionsmitgliedern herausgegeben.

In 2024 wurden drei Ausgaben von *TGDK* mit insgesamt 16 Artikeln veröffentlicht.

### ■ LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*<sup>26</sup> veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. Die Zeitschrift wurde 2012 gegründet und nahm 2013 ihren Betrieb auf. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern fungiert unter Leitung von

### ■ TGDK: Transactions on Graph Data and Knowledge

*Transactions on Graph Data and Knowledge (TGDK)*<sup>25</sup> is an Open Access journal that publishes original research articles and survey articles on graph-based abstractions for data and knowledge, and the techniques that enable such abstractions with respect to integration, querying, reasoning and learning. The journal was established in Autumn 2023 and published its first issue in December 2023.

The journal is published as Diamond Open Access, i.e. all content is freely accessible and no fees are charged for readers or authors. The journal is edited by Aidan Hogan, Andreas Hotho, Lalana Kagal, and Uli Sattler as editors-in-chief along with a team of around 100 associate editors and editorial board members.

In 2024, three issues of *TGDK* containing 16 articles were published.

### ■ LITES: Leibniz Transactions on Embedded Systems

The *LITES*<sup>26</sup> journal publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was established in 2012 and started operating in early 2013. A broad team of experienced researchers, acting as editorial board headed

<sup>25</sup> <https://www.dagstuhl.de/tgdk>

Björn B. Brandenburg als Herausbergremium und begutachtet alle eingereichten Arbeiten. Die Zeitschrift wird gemeinsam mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*<sup>27</sup> der Fachgesellschaft *European Design and Automation Association (EDAA)*<sup>28</sup> herausgegeben. Die Fachgruppe ist dabei für die Besetzung des Herausbergremiums verantwortlich, während Schloss Dagstuhl die administrativen Aufgaben der Herausbergerschaft übernimmt.

In 2024 wurde ein Artikel in *LITES* veröffentlicht.

by Björn B. Brandenburg, reviews all submitted contributions. The journal is jointly published with the *EMbedded Systems Special Interest Group (EMSIG)*<sup>27</sup> of the *European Design and Automation Association (EDAA)*<sup>28</sup>. The special interest group is responsible for appointing the editorial board, while Schloss Dagstuhl takes over the administrative tasks of the publication.

In 2024, one article was published in *LITES*.

## Infrastruktur

4.4

## Infrastructure

### ■ Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet. Die Bände aus den Reihen *LIPICs* und *OASICs* werden zudem bei Scopus<sup>29</sup> eingereicht, wo sie regelmäßig indiziert werden. Die Reihen *LIPICs* und *OASICs* sowie die Zeitschriften *LITES* und *TGDK* sind zudem im Directory of Open Access Journals (DOAJ) gelistet. Zudem unterstützen die technischen Schnittstellen die Datenakquise durch Google Scholar, was eine gute Sichtbarkeit der Publikationen garantiert.

### ■ Engagement in Arbeitsgruppen

Mitarbeitende von Schloss Dagstuhl waren auch in 2024 in verschiedenen Arbeitsgruppen beteiligt:

- Arbeitsgruppe “Open Access und Publikationsmanagement” der Leibniz-Gemeinschaft
- Arbeitsgruppe “Open Access” des Saarlandes
- Task Force „Research Software Infrastructures” der EOSC (European Open Science Cloud)
- Arbeitsgruppe “Forschungsdaten” der Leibniz-Gemeinschaft

### ■ Publikationsserver: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)<sup>30</sup> werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Nach rund einem Jahr Entwicklungszeit und einer mehrmonatigen Testphase hat Dagstuhl Publishing mit der Inbetriebnahme eines neuen Publikationsservers im Herbst 2023 einen weiteren Meilenstein erreicht. Die neu entwickelte Plattform *DROPS (Dagstuhl Research Online Publication Server)* bietet unter der URL <https://drops.dagstuhl.de> weiterhin freien Zugang zu den über 20 000 Forschungsarbeiten aus dem Bereich der Informatik, die von Dagstuhl Publishing bisher veröffentlicht wurden. Im Zuge der zeitgemäßen Neugestaltung des System wurde insbesondere die Benutzerfreundlichkeit deutlich verbessert. Neben einer schnelleren Navigation

### ■ Indexing

All series of the publication portfolio are listed in *dblp*. The *LIPICs* and *OASICs* volumes are submitted to Scopus<sup>29</sup> where they are regularly indexed. The *LIPICs* and *OASICs* series as well as the journals *LITES* and *TGDK* are also listed in the Directory of Open Access Journals (DOAJ). The technical interface of our publication server enables harvesting according to the Google Scholar guidelines. Google Scholar regularly retrieves metadata and full-texts from our server.

### ■ Involvement in working groups

Staff members of Schloss Dagstuhl were also involved in 2024 in various working groups:

- Working group “Open Access and Publication Management” of the Leibniz association
- Working group “Open Access” of Saarland
- Research Software Infrastructures Task Force of the EOSC (European Open Science Cloud)
- Working group “Research Data” of the Leibniz association

### ■ Publication Server: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)<sup>30</sup>. After around a year of development and several months of testing, Dagstuhl Publishing reached another milestone in autumn 2023 by launching its new publication server. Under the URL <https://drops.dagstuhl.de>, the newly developed platform *DROPS (Dagstuhl Research Online Publication Server)* offers free access to the more than 20,000 research papers in the field of computer science published by Dagstuhl Publishing so far. As part of the modern redesign of the system, the user-friendliness in particular has been significantly improved. In addition to faster navigation and a modernised browsing and search interface, the new platform offers support for mobile

<sup>26</sup> <https://www.dagstuhl.de/lites>

<sup>27</sup> <http://www.emsig.net/>

<sup>28</sup> <https://www.edaa.com/>

<sup>29</sup> <https://www.scopus.com>

und einem modernisiertem Browsing- bzw. Such-Interface bietet die neue Plattform Unterstützung für mobile Geräte, Nutzungsstatistiken auf allen Ebenen sowie einen deutlich erweiterten Metadatenatz für jedes Dokument.

Die folgenden Funktionalitäten bzw. Dienste sind im Jahr 2024 neu hinzugekommen:

- eine erweiterte Metadatenchnittstelle<sup>31</sup>
- die Veröffentlichung von Zusatzmaterialien (wie Software oder Daten) in der Sammlung “Supplementary Materials”<sup>32</sup>
- das nahtlose Browsen in SoftwareHeritage, falls ein Link zu archivierter Software gegeben ist
- die Veröffentlichung der monatlichen dblp Dumps<sup>33</sup> mit persistenten Identifiern.

### ■ Einreichungssystem: DSub

Im Frühjahr 2019 wurde ein von Dagstuhl entwickeltes Einreichungssystem namens DSub in Betrieb genommen. Mit diesem System werden derzeit alle Einreichungen zu den Reihen LIPIcs, OASICs, DARTS, LITES und TGDK entgegengenommen. Unter anderem wurde mit diesem System dem Wunsch einer aktiven Autorenfreigabe der überarbeiteten Dokumente vor der Veröffentlichung entsprochen und die automatische Extraktion der Metadaten aus den LaTeX-Quellen ermöglicht.

Als erstes konkretes Ergebnis laufender Projektbeteiligungen (insbesondere *FAIRCORE4EOSC*<sup>34</sup>) wurde im Frühjahr 2023 eine Benutzeroberfläche in Betrieb genommen, die die Verknüpfung von Konferenzbeiträgen mit Forschungssoftware oder -daten (nach den FAIR-Prinzipien) bei der Einreichung im Dagstuhl Submission System erleichtert. Durch die nahtlose Integration des Dienstes *SoftwareHeritage*<sup>35</sup> können Autoren auf Wunsch vollständig von der Archivierung von Software entlastet werden.

Seit Mitte 2024 ist die Angabe erweiterter Metadaten zu Zusatzmaterialien eines Artikels bereits während der Einreichung möglich, so dass diese Materialien (auf Wunsch der Autoren) schließlich parallel zum Artikel auf DROPS veröffentlicht werden.

devices, usage statistics at all levels and a significantly expanded metadata record for each document.

The following new features and services were added in 2024:

- an extended metadata interface<sup>31</sup>
- publishing supplementary materials (such as software or data) in the collection “Supplementary Materials”<sup>32</sup>
- seamless browsing in Software Heritage if a link to archived software is given
- publishing the monthly dblp dumps<sup>33</sup> with persistent identifiers.

### ■ Submission System: DSub

In spring 2019 a submission system called DSub developed by Dagstuhl was introduced. At present, all submissions for the LIPIcs, OASICs, DARTS, LITES and TGDK series are handled via this system. Among other things, the new system has satisfied the need for active author approval of revised documents prior the publication and enables automatic extraction of metadata from LaTeX sources.

As the first concrete result of ongoing project participations (in particular *FAIRCORE4EOSC*<sup>34</sup>), a user interface was put into operation in spring 2023 that facilitates the linking of conference contributions with research software or data (according to the FAIR principles) during submission in the Dagstuhl Submission System. Seamless integration of the *SoftwareHeritage*<sup>35</sup> service means that authors can be completely relieved of the task of archiving software if they wish.

Since mid-2024, it has been possible to provide extended metadata on additional materials for an article during submission, so that these materials (at the request of the authors) are ultimately published alongside the article on DROPS.

## Nutzungsstatistiken

4.5

## Usage Statistics

Bei der Nutzung des Publikationsserver wird zwischen der Ansicht der Metadaten und der Ansicht/Download des eigentlichen Dokuments unterschieden. Die Nutzungszahlen konnten im Jahr 2024 gegenüber dem Vorjahr noch einmal deutlich gesteigert werden. Die Zahl der Dokumentenaufrufe erhöhte sich gegenüber dem Vorjahr um 25 Prozent auf nunmehr rund 1,43 Millionen Aufrufe.

When using the publication server, a distinction is made between viewing the metadata and viewing/downloading the actual document. The usage figures increased significantly in 2024 compared to the previous year. The number of document views increased by 25 percent to around 1.43 million views.

<sup>30</sup> <https://www.dagstuhl.de/drops>

<sup>31</sup> <https://drops.dagstuhl.de/metadata>

<sup>32</sup> <https://drops.dagstuhl.de/entities/collection/supplementary-materials>


<sup>33</sup> <https://drops.dagstuhl.de/entities/collection/dblp>

<sup>34</sup> <https://faircore4eosc.eu/>

<sup>35</sup> <https://www.softwareheritage.org/>

Fig. 4.6 fasst die jährliche Nutzung des Publikationservers zusammen. Alle Statistiken ignorieren üblicherweise die Zugriffe, die durch bekannte Bot- und Crawler-Software verursacht wurden.

Fig. 4.6 summarises the annual usage of the publication server. All figures should ignore the traffic caused by known bots and crawlers.



Year	PDF Downloads	Abstract
2018	627 988	371 594
2019	737 721	452 321
2020	940 347	656 112
2021	1 242 064	722 787
2022	1 082 872	789 730
2023	1 146 962	1 104 269
2024	1 434 114	1 413 547

(a) Graphical distribution for 2018–2024

(b) Detailed numbers for 2018–2024

**Fig. 4.6**  
Total number of full text accesses (PDF Downloads) and abstract/landing page views for articles published between 2018 to 2024.

Vol. 116   15th Workshop on Parallel Programming and Run-Time Management Techniques for Many-Core Architectures and 13th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-307-2">https://www.dagstuhl.de/dagpub/978-3-95977-307-2</a>
Vol. 117   Fifth Workshop on Next Generation Real-Time Embedded Systems (NG-RES 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-313-3">https://www.dagstuhl.de/dagpub/978-3-95977-313-3</a>
Vol. 118   5th International Workshop on Formal Methods for Blockchains (FMBC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-317-1">https://www.dagstuhl.de/dagpub/978-3-95977-317-1</a>
Vol. 119   The Provenance of Elegance in Computation - Essays Dedicated to Val Tannen <a href="https://www.dagstuhl.de/dagpub/978-3-95977-320-1">https://www.dagstuhl.de/dagpub/978-3-95977-320-1</a>
Vol. 120   13th Symposium on Languages, Applications and Technologies (SLATE 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-321-8">https://www.dagstuhl.de/dagpub/978-3-95977-321-8</a>
Vol. 121   22nd International Workshop on Worst-Case Execution Time Analysis (WCET 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-346-1">https://www.dagstuhl.de/dagpub/978-3-95977-346-1</a>
Vol. 122   5th International Computer Programming Education Conference (ICPEC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-347-8">https://www.dagstuhl.de/dagpub/978-3-95977-347-8</a>
Vol. 123   24th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-350-8">https://www.dagstuhl.de/dagpub/978-3-95977-350-8</a>
Vol. 124   Commit2Data <a href="https://www.dagstuhl.de/dagpub/978-3-95977-351-5">https://www.dagstuhl.de/dagpub/978-3-95977-351-5</a>
Vol. 125   35th International Conference on Principles of Diagnosis and Resilient Systems (DX 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-356-0">https://www.dagstuhl.de/dagpub/978-3-95977-356-0</a>

**Fig. 4.7**  
OASlcs volumes published in 2024.

Vol. 286   27th International Conference on Principles of Distributed Systems (OPODIS 2023) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-308-9">https://www.dagstuhl.de/dagpub/978-3-95977-308-9</a>
Vol. 287   15th Innovations in Theoretical Computer Science Conference (ITCS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-309-6">https://www.dagstuhl.de/dagpub/978-3-95977-309-6</a>
Vol. 288   32nd EACSL Annual Conference on Computer Science Logic (CSL 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-310-2">https://www.dagstuhl.de/dagpub/978-3-95977-310-2</a>
Vol. 289   41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-311-9">https://www.dagstuhl.de/dagpub/978-3-95977-311-9</a>
Vol. 290   27th International Conference on Database Theory (ICDT 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-312-6">https://www.dagstuhl.de/dagpub/978-3-95977-312-6</a>
Vol. 291   12th International Conference on Fun with Algorithms (FUN 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-314-0">https://www.dagstuhl.de/dagpub/978-3-95977-314-0</a>
Vol. 292   3rd Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-315-7">https://www.dagstuhl.de/dagpub/978-3-95977-315-7</a>
Vol. 293   40th International Symposium on Computational Geometry (SoCG 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-316-4">https://www.dagstuhl.de/dagpub/978-3-95977-316-4</a>
Vol. 294   19th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-318-8">https://www.dagstuhl.de/dagpub/978-3-95977-318-8</a>
Vol. 295   5th Symposium on Foundations of Responsible Computing (FORC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-319-5">https://www.dagstuhl.de/dagpub/978-3-95977-319-5</a>
Vol. 296   35th Annual Symposium on Combinatorial Pattern Matching (CPM 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-326-3">https://www.dagstuhl.de/dagpub/978-3-95977-326-3</a>
Vol. 297   51st International Colloquium on Automata, Languages, and Programming (ICALP 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-322-5">https://www.dagstuhl.de/dagpub/978-3-95977-322-5</a>
Vol. 298   36th Euromicro Conference on Real-Time Systems (ECRTS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-324-9">https://www.dagstuhl.de/dagpub/978-3-95977-324-9</a>
Vol. 299   9th International Conference on Formal Structures for Computation and Deduction (FSCD 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-323-2">https://www.dagstuhl.de/dagpub/978-3-95977-323-2</a>
Vol. 300   39th Computational Complexity Conference (CCC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-331-7">https://www.dagstuhl.de/dagpub/978-3-95977-331-7</a>
Vol. 301   22nd International Symposium on Experimental Algorithms (SEA 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-325-6">https://www.dagstuhl.de/dagpub/978-3-95977-325-6</a>
Vol. 302   35th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AofA 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-329-4">https://www.dagstuhl.de/dagpub/978-3-95977-329-4</a>
Vol. 303   29th International Conference on Types for Proofs and Programs (TYPES 2023) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-332-4">https://www.dagstuhl.de/dagpub/978-3-95977-332-4</a>
Vol. 304   5th Conference on Information-Theoretic Cryptography (ITC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-333-1">https://www.dagstuhl.de/dagpub/978-3-95977-333-1</a>
Vol. 305   27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-334-8">https://www.dagstuhl.de/dagpub/978-3-95977-334-8</a>
Vol. 306   49th International Symposium on Mathematical Foundations of Computer Science (MFCS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-335-5">https://www.dagstuhl.de/dagpub/978-3-95977-335-5</a>
Vol. 307   30th International Conference on Principles and Practice of Constraint Programming (CP 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-336-2">https://www.dagstuhl.de/dagpub/978-3-95977-336-2</a>
Vol. 308   32nd Annual European Symposium on Algorithms (ESA 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-338-6">https://www.dagstuhl.de/dagpub/978-3-95977-338-6</a>
Vol. 309   15th International Conference on Interactive Theorem Proving (ITP 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-337-9">https://www.dagstuhl.de/dagpub/978-3-95977-337-9</a>

Fig. 4.8

LIPIcs volumes published in 2024 – Part 1.

Vol. 310   19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-328-7">https://www.dagstuhl.de/dagpub/978-3-95977-328-7</a>
Vol. 311   35th International Conference on Concurrency Theory (CONCUR 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-339-3">https://www.dagstuhl.de/dagpub/978-3-95977-339-3</a>
Vol. 312   24th International Workshop on Algorithms in Bioinformatics (WABI 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-340-9">https://www.dagstuhl.de/dagpub/978-3-95977-340-9</a>
Vol. 313   38th European Conference on Object-Oriented Programming (ECOOP 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-341-6">https://www.dagstuhl.de/dagpub/978-3-95977-341-6</a>
Vol. 314   30th International Conference on DNA Computing and Molecular Programming (DNA 30) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-344-7">https://www.dagstuhl.de/dagpub/978-3-95977-344-7</a>
Vol. 315   16th International Conference on Spatial Information Theory (COSIT 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-330-0">https://www.dagstuhl.de/dagpub/978-3-95977-330-0</a>
Vol. 316   6th Conference on Advances in Financial Technologies (AFT 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-345-4">https://www.dagstuhl.de/dagpub/978-3-95977-345-4</a>
Vol. 317   Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-348-5">https://www.dagstuhl.de/dagpub/978-3-95977-348-5</a>
Vol. 318   31st International Symposium on Temporal Representation and Reasoning (TIME 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-349-2">https://www.dagstuhl.de/dagpub/978-3-95977-349-2</a>
Vol. 319   38th International Symposium on Distributed Computing (DISC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-352-2">https://www.dagstuhl.de/dagpub/978-3-95977-352-2</a>
Vol. 320   32nd International Symposium on Graph Drawing and Network Visualization (GD 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-343-0">https://www.dagstuhl.de/dagpub/978-3-95977-343-0</a>
Vol. 321   19th International Symposium on Parameterized and Exact Computation (IPEC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-353-9">https://www.dagstuhl.de/dagpub/978-3-95977-353-9</a>
Vol. 322   35th International Symposium on Algorithms and Computation (ISAAC 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-354-6">https://www.dagstuhl.de/dagpub/978-3-95977-354-6</a>
Vol. 323   44th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2024) <a href="https://www.dagstuhl.de/dagpub/978-3-95977-355-3">https://www.dagstuhl.de/dagpub/978-3-95977-355-3</a>

Fig. 4.9  
LIPIcs volumes published in 2024 – Part 2.



# **5** Resonanz *Feedback*

## Resonanz zu Seminaren und Workshops

5.1

## Feedback on Seminars and Workshops

### ■ Resonanz von Teilnehmern

Schloss Dagstuhl bekommt viel Lob von seinen Gästen, meistens in mündlicher Form, wenn die Gäste nach einer intensiven Seminarwoche das Schloss verlassen. Manche Gäste nehmen sich jedoch auch die Zeit, uns nachträglich zu schreiben und ihre Eindrücke mit uns zu teilen. Im Folgenden haben wir mit freundlicher Genehmigung der Autoren einen Auszug aus unserer großen Sammlung an Dankeschön-Nachrichten zusammengestellt.

**Luke Lucas (LSE Space – Darmstadt, DE) in an Email sent April 12, 2024**

23461 – Space and Artificial Intelligence | Dagstuhl Seminar | <https://www.dagstuhl.de/23461>

It is difficult for me to describe what an amazing seminar it was, how much I learnt and how invigorated I was by the time I left. And how exciting that ‘thrill’ is. It continues in me, as I seek ways to be involved in the development of AI in space, and have meaningful discussions from the skills gap and training to adoption and implementation. I said to a friend the ‘emotional rush’ of the seminar was almost like falling in love. But not in love with a person, but with science and AI and the exchange of ideas and inspiration.

**Ehsan Yousefzadeh-Asl-Miandoab (IT University of Copenhagen, DK)**

24311 – Resource-Efficient Machine Learning | Dagstuhl Seminar | <https://www.dagstuhl.de/24311>

I had a really nice experience as my first time participation in the Dagstuhl center. The scientific level was good, and we had a good mixture of people from seniority, industry/ academia balance. The food was high quality and wide range of choice.

### ■ Resonanz unserer Organisatoren

Der Erfolg von Schloss Dagstuhl hängt im wesentlichen Maße auch von den Seminarorganisatoren ab, die interessante und neue Themen vorschlagen. Wir sind hoch erfreut, dass die Seminarorganisatoren selber die Angebote und die Umgebung, die wir zur Verfügung stellen, schätzen. Im Folgenden geben wir mit freundlicher Genehmigung der Autoren einige der Kommentare unsere Seminarorganisatoren wieder.

### ■ Feedback from Organizers

The success of Schloss Dagstuhl depends to a large extent on our outstanding seminar organizers, who continually enrich the scientific program with a range of interesting and new topics. We are very glad to be able to provide services and an environment that organizers appreciate. The following comments from organizers are excerpted from the Dagstuhl Report or personal emails to us. We cite them with their kindly permission.

**Organizers of Dagstuhl Seminar 24042**

24042 – The Emerging Issues in Bioimaging AI Publications and Research | Dagstuhl Seminar | <https://www.dagstuhl.de/24042>

our deepest appreciation goes to the exceptional Dagstuhl staff whose support was instrumental in making the seminar a success.

**Organizers of Dagstuhl Seminar 24302**

24302 – Learning with Music Signals: Technology Meets Education | Dagstuhl Seminar | <https://www.dagstuhl.de/24302>

We want to express our gratitude to the Dagstuhl board for giving us the opportunity to organize this seminar, the Dagstuhl office for their exceptional support throughout the organization process, and the entire Dagstuhl staff for their excellent service during the seminar.

**Organizers of Dagstuhl Seminar 24472**24472 – Regular Expressions: Matching and Indexing | Dagstuhl Seminar | <https://www.dagstuhl.de/24472>

We are really thankful to Schloss Dagstuhl for providing an extremely inspiring and professional environment. Scientific talks were interleaved with coffee breaks and cheese tastings which fostered engagements of the participants, and the evening “sessions” in the sauna, wine cellar, and music room offered a relaxed environment to continue chatting about research in a less formal environment.

### ■ Resonanz in Sozialen Netzwerken

Mehr und mehr Gäste nutzen die Möglichkeiten des Webs wie X (vormals Twitter), LinkedIn und Blogs um über ihre Erfahrungen in Dagstuhl zu berichten. Wir geben hier einige ausgewählte Referenzen.

Einige Gäste nutzten den Hashtag „#dagstuhl“ um während ihrer Seminarteilnahme direkt aus dem Seminar zu berichten. So hat etwa die Programming Principles, Logic and Verification group des University College London in mehreren X (vormals Twitter)-Threads<sup>36</sup> vom Dagstuhl Seminar 24341 berichtet. Jakob Engblom bewarb auf LinkedIn seinen englischen Reisebericht zu Dagstuhl<sup>37</sup>. Und Tommy Thompson berichtete auf seinem Blog<sup>38</sup> ausführlich über seinen Dagstuhlbesuch.

### ■ Feedback in Social Media

More and more of our guests are using social media such as X (formerly Twitter), LinkedIn, and blogs to share their experiences of Dagstuhl with others. Below are some selected excerpts.

Some guests used the hashtag “#dagstuhl” to live-tweet their seminar participation. For example, the Programming Principles, Logic and Verification group at University College London reported on Dagstuhl Seminar 24341 in a collection of X (formerly Twitter) threads.<sup>36</sup> Jakob Engblom used LinkedIn to promote his [englisch travel report on Dagstuhl](#)<sup>37</sup>. And Tommy Thompson did a thorough report on his Dagstuhl visit in his [blog](#)<sup>38</sup>.

#### **Dirk Kutscher (The Hong Kong University of Science and Technology – Guangzhou, CN)**

LinkedIn | <https://www.linkedin.com/feed/update/urn:li:activity:7178352635878543360?commentUrn=urn%3A%3Acomment%3A%28activity%3A7178352635878543360%2C7182612985540935681%29&dashCommentUrn=urn%3A%3A%3Acomment%3A%287182612985540935681%2Curn%3A%3A%3Aactivity%3A7178352635878543360%29>

The impact of Dagstuhl cannot be overestimated. Many significant research directions in my field were first conceived by bringing together smart people from across the globe and giving them time to think and discuss. Examples are: Delay-Tolerant Networking, Information-Centric Networking, Compute-First Networking.

#### **Jim Dowling (KTH Stockholm, SE)**

X (formerly Twitter) | [https://twitter.com/jim\\_dowling/status/1766538734218981383](https://twitter.com/jim_dowling/status/1766538734218981383)

Dagstuhl are unique. Our company, @hopsworks, originated in an open problem I learnt about at a dagstuhl.

#### **Niki van Stein Leiden University, NL)**

Google Maps | <https://maps.app.goo.gl/fn7f77muJgfawv5>

Hogwarts for Computer Scientists! The food is great, the service amazing and the new Ghost hunt is brilliant. Very nice location for seminars. The rooms could be a bit more luxurious, especially the beds, but that is all I can complain about.

#### **Marvin Wyrich (Universität des Saarlandes – Saarbrücken, DE)**

LinkedIn | <https://www.linkedin.com/feed/update/urn:li:activity:7215958284741054465?commentUrn=urn%3A%3Acomment%3A%28activity%3A7215958284741054465%2C7215992680802062336%29&dashCommentUrn=urn%3A%3A%3Acomment%3A%287215992680802062336%2Curn%3A%3A%3Aactivity%3A7215958284741054465%29>

<sup>36</sup> [https://x.com/ucl\\_pplv/status/1825517561460466159](https://x.com/ucl_pplv/status/1825517561460466159),  
[https://x.com/ucl\\_pplv/status/1825632552176292039](https://x.com/ucl_pplv/status/1825632552176292039),  
[https://x.com/ucl\\_pplv/status/1826326201084117023](https://x.com/ucl_pplv/status/1826326201084117023),  
[https://x.com/ucl\\_pplv/status/1826610304744345823](https://x.com/ucl_pplv/status/1826610304744345823),  
[https://x.com/ucl\\_pplv/status/1826943669099749720](https://x.com/ucl_pplv/status/1826943669099749720),  
[https://x.com/ucl\\_pplv/status/1826944850517799364](https://x.com/ucl_pplv/status/1826944850517799364)

<sup>37</sup> <https://jakob.engbloms.se/archives/4296>

<sup>38</sup> <https://www.aiandgames.com/i/155937360/attending-dagstuhl-seminar-computational-creativity-for-game-development>

Dagstuhl seminars bring very different people together for a week at a somewhat isolated castle. I know of hardly any other event that manages to encourage so much (productive) exchange in such a short time; primarily on a professional level, but you also simply get to know the researchers much better on a personal level. That is wonderful. [...]

**Ranjit Jhala (University of California – San Diego, US)**

X (formerly Twitter) | <https://x.com/RanjitJhala/status/1811031058739253360>

Not a conference per se and hence probably doesn't apply but @dagstuhl has an \*amazing\* setup!

**Jeremy Siek (Indiana University – Bloomington, US)**

X (formerly Twitter) | <https://x.com/jeremysiek/status/1811063280900690376>

The childcare at dagstuhl was a particular life-saver for Katie and I because it enabled us to attend two different dagstuhl seminars at the same time :)

**Satnam Singh (Groq – Mountain View, US)**

X (formerly Twitter) | <https://x.com/satnam6502/status/1813177755606942110>

A nice touch at Dagstuhl is how they display books written or edited by the participants.

**Roberto Grossi (University of Pisa, IT)**

Google Maps | <https://maps.app.goo.gl/Hxbv2uPriHcbo1aw6>

Great research experience and hospitality, prices are ok to let people around the world to come and work together in Computer Science research

## ■ Resonanz im Fragebogen

Jeder Teilnehmer erhält von uns einen Fragebogen zur Evaluation des vom Teilnehmer besuchten Dagstuhl-Seminars oder Dagstuhl-Perspektiven-Workshops. Durch diese anonyme Befragung erhalten wir ebenfalls eine Menge positiver Kommentare. Im Folgenden zitieren wir hier einige von diesen.

## ■ Survey Feedback

Every participant has the opportunity to fill out a questionnaire about the Dagstuhl Seminar or Dagstuhl Perspectives Workshop they attended for evaluation purposes. Below are some excerpts from the many positive comments we received through this anonymous survey.

24031 – Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis | Dagstuhl Seminar | <https://www.dagstuhl.de/24031>

This was an excellent meeting, and I am fortunate to have this opportunity to come to Dagstuhl. I would definitely attend future meetings at Dagstuhl.

24031 – Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis | Dagstuhl Seminar | <https://www.dagstuhl.de/24031>

love the bees and the honey thanks for caring for food preferences

24032 – Representation, Provenance, and Explanations in Database Theory and Logic | Dagstuhl Seminar | <https://www.dagstuhl.de/24032>

I enjoyed all Dagstuhl workshops that I attended in the past. This workshop stood out as even more friendly and enjoyable.

24052 – Reviewer No. 2: Old and New Problems in Peer Review | Dagstuhl Seminar | <https://www.dagstuhl.de/24052>

Dagstuhl is a treasure; thanks for keeping it going!

24072 – Triangulations in Geometry and Topology | Dagstuhl Seminar | <https://www.dagstuhl.de/24072>

Thanks for everything, Dagstuhl is outstanding. As a researcher I really feel considered when I come to Dagstuhl. This is very important for the productivity and the quality of the work we do there.

24081 – Computational Approaches to Strategy and Tactics in Sports | Dagstuhl Seminar | <https://www.dagstuhl.de/24081>

The seminar offers exactly what academia needs: bring together an international and interdisciplinary group of scientists, providing time for discussion and interaction. It stimulated philosophical discussions, provided time for thought and reflection, opened new routes for collaboration and directions, which is very valuable. Though the location is lovely, it is hard to reach but this has been anticipated and guidance was provided.

24101 – Robust Query Processing in the Cloud | Dagstuhl Seminar | <https://www.dagstuhl.de/24101>

Great place and very unique not only in Germany.

24151 – Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems | Dagstuhl Seminar | <https://www.dagstuhl.de/24151>

I find the idea of these seminars very good and the complete experience is remarkable. I am looking forward to visit Dagstuhl again.

24151 – Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems | Dagstuhl Seminar | <https://www.dagstuhl.de/24151>

Great overall, great organizers, engaging format for exchange – dynamic. Great Dagstuhl Staff. Appreciate the technology upgrades since my first visit more than 15+ years ago.

24162 – Hardware Support for Cloud Database Systems in the Post-Moore's Law Era | Dagstuhl Seminar | <https://www.dagstuhl.de/24162>

Dagstuhl's unique character as a place where people work without distraction for a whole week continues to be very relevant. Please keep doing that!

24172 – Code Search | Dagstuhl Seminar | <https://www.dagstuhl.de/24172>

This was the best industry/academia workshop I have attended and the viewpoints of both greatly strengthened discussions.

24181 – Computational Metabolomics: Towards Molecules, Models, and their Meaning | Dagstuhl Seminar | <https://www.dagstuhl.de/24181>

Dagstuhl is a unique (in my experience) and hugely valuable facility / venue / organisation. I would like to thank the Leibniz Institutes for supporting this excellent opportunity. It is amazing that other research organisations do not provide these types of meetings!

24182 – Resilience and Antifragility of Autonomous Systems | Dagstuhl Seminar | <https://www.dagstuhl.de/24182>

Please continue doing what you already do. Dagstuhl is a rare institution that understands the nature of the scientific process, the need to have a long-term perspective, and the importance of bringing researchers together in the same room for an extended period of time for enabling an exchange of scientific ideas. In this era of social media where many aspects of science have fallen prey to the need for constant marketing in order to grab attention, Dagstuhl stands out as institution that embodies the best aspects of the scientific process—open to public review yet singularly guided by a search for truth.

24201 – Discrete Algorithms on Modern and Emerging Compute Infrastructure | Dagstuhl Seminar | <https://www.dagstuhl.de/24201>

Best aspect: the environment, by which I mean the fact that participants are stuck together for a week so there are many opportunities to chat with people. The opportunities for this are an order of magnitude larger than at a normal conference and they apply to an order of magnitude more people than in a visit to a colleague at a university. I didn't experience any bad aspect so I'll skip choosing the worst.

24241 – Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis | Dagstuhl Seminar | <https://www.dagstuhl.de/24241>

I met many colleagues of mine in Asia who attended the Dagstuhl Seminars who wanted to copy this wonderful environment to their own countries. I hope that they will be successful some time soon in the near future. Thank you very much indeed!

24242 – Computational Analysis and Simulation of the Human Voice | Dagstuhl Seminar | <https://www.dagstuhl.de/24242>

Dagstuhl makes it easy for organizers, I was relaxed all the time (organizer)

24251 – Teaching Support Systems for Formal Foundations of Computer Science | Dagstuhl Seminar | <https://www.dagstuhl.de/24251>

Especially the music room was greatly appreciated.

24272 – A Game of Shadows: Effective Mastery Learning in the Age of Ubiquitous AI | Dagstuhl Seminar | <https://www.dagstuhl.de/24272>

I love the ghost scavenger hunt! It's a phenomenal way to interact with the space, learn about the history of Dagstuhl, and serves as a conversation piece.

24291 – Programmable Host Networking | Dagstuhl Seminar | <https://www.dagstuhl.de/24291>

Calm wonderfulness. Dagstuhl is a very special place

24292 – Improving Trust between Humans and Software Robots in Robotic Process Automation | Dagstuhl Seminar | <https://www.dagstuhl.de/24292>

Thank you so much for providing such an incredibly important and prestigious set-up for advancement of Science.

24292 – Improving Trust between Humans and Software Robots in Robotic Process Automation | Dagstuhl Seminar | <https://www.dagstuhl.de/24292>

The best thing about the seminars is that it provides the perfect setup for connecting with international researchers. For me, it was particularly nice to interact with people from a different research area. The facilities in the castle play a crucial role in being able to do this.

24301 – Art, Visual Illusions, and Data Visualization | Dagstuhl Seminar | <https://www.dagstuhl.de/24301>

This was a life-changing event—fa'bulous!

24351 – Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES) | Dagstuhl Seminar | <https://www.dagstuhl.de/24351>

I was here a few months ago and there were fewer bikes. Thanks for buying new ones!

24371 – Extended Reality Accessibility | Dagstuhl Seminar | <https://www.dagstuhl.de/24371>

Excellent to see Dagstuhl supports parents

24381 – Algebraic and Analytic Methods in Computational Complexity | Dagstuhl Seminar | <https://www.dagstuhl.de/24381>

I am very grateful that gluten-free food is provided. Thank you for that and everything else!

24381 – Algebraic and Analytic Methods in Computational Complexity | Dagstuhl Seminar | <https://www.dagstuhl.de/24381>

In the era of numerous conferences and meetings, Dagstuhl Workshop stands out. I like visiting Dagstuhl, much more than other venues, due to both the quality and quantity of interactions that happen here.

24401 – Fair Division: Algorithms, Solution Concepts, and Applications | Dagstuhl Seminar | <https://www.dagstuhl.de/24401>

The randomized seating for meals is an excellent idea.

24401 – Fair Division: Algorithms, Solution Concepts, and Applications | Dagstuhl Seminar | <https://www.dagstuhl.de/24401>

As usual, Dagstuhl seminars have a level of perfection we don't see anywhere else. The fact that the organizers have already participated to several Dagstuhl seminars (for the one and in general) explains why the quality remains constant and so high.

24411 – New Tools in Parameterized Complexity: Paths, Cuts, and Decomposition | Dagstuhl Seminar | <https://www.dagstuhl.de/24411>

I got to see several state of the art new techniques (often not even published yet) for interesting problems. There was also plenty of time to discuss both old and new projects related to the topic of the seminar. I don't have anything specific in mind that could be improved, besides the weather, which is not in the organizer's hands.

24421 – SAT and Interactions | Dagstuhl Seminar | <https://www.dagstuhl.de/24421>

It is hard to improve Dagstuhl seminars as scientific events. Dagstuhl is a precious resource for the computer science community, let's hope it will continue like this in the future.

24451 – Machine Learning for Protein-Protein and Protein-Ligand Interactions | Dagstuhl Seminar | <https://www.dagstuhl.de/24451>

This was the best meeting I have been to in decades.

24471 – Graph Algorithms: Distributed Meets Dynamic | Dagstuhl Seminar | <https://www.dagstuhl.de/24471>

Thank you so much for the amazing food and the flexibility with dietary restrictions

24472 – Regular Expressions: Matching and Indexing | Dagstuhl Seminar | <https://www.dagstuhl.de/24472>

I love Dagstuhl, it's much better than attending conferences. Much more time to interact, which is the real reason to attend both.

24472 – Regular Expressions: Matching and Indexing | Dagstuhl Seminar | <https://www.dagstuhl.de/24472>

An excellent research atmosphere, with all you need to do research. The right balance between a relaxed environment for research and the opportunities for networking and working. Clearly run by people that understand how this works. Everything is well-planned. I loved the new library at the basement. I cannot identify bad aspects.

24491 – Deep Learning for RNA Regulation and Multidimensional Transcriptomics | Dagstuhl Seminar | <https://www.dagstuhl.de/24491>

I found myself constantly writing down new ideas for projects and collaborations, something that very rarely happens at other meetings. It was extremely nurturing, and I had great discussions, both during the sessions and during leisure time.

24492 – Human in the Loop Learning through Grounded Interaction in Games | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/24492>

Among the best things I would consider the general setup: being in a remote location with a group of like-minded experts to brainstorm on a particular topic. Everything is perfectly arranged, and I can confirm the summary someone gave me at the first seminar I attended: Dagstuhl offers everything you need and nothing that you don't need. Could not be any better. No worst aspect.

24511 – Coding Theory and Algorithms for Emerging Technologies in Synthetic Biology | Dagstuhl Seminar | <https://www.dagstuhl.de/24511>

This was one of the best academic experiences in my life!! Far exceeded my expectations. Thank you to the organisers and to Dagstuhl staff!!!

24512 – Quantum Software Engineering | Dagstuhl Seminar | <https://www.dagstuhl.de/24512>

Great setup, networking opportunities, outstanding facilities, outstanding food

## Resonanz zur Bibliographiedatenbank dblp

5.2

## Feedback on the dblp computer science bibliography

Die Bibliographiedatenbank dblp wird von zahlreichen internationalen Wissenschaftlern hoch geschätzt und erhält viel Lob. Feedback erhalten wir per Mail, durch Gespräche mit Forschern vor Ort in Dagstuhl, oder durch die sozialen Medien.

The dblp computer science bibliography is internationally well known and appreciated. We receive a lot of feedback via mail, through discussions with researchers at Schloss Dagstuhl, and via social media.

### ■ Resonanz in Fachartikeln

### ■ Feedback in scientific articles

**Ariel Rosenfeld, Bar-Ilan University, Israel**

"Is DBLP a Good Computer Science Journals Database?" *Computer* 56(3), March 2023. | <https://doi.org/10.1109/MC.2022.3181977>

Generally speaking, it seems that DBLP provides decent coverage of CS and "CS-adjacent" journals, at least when considering WOS, with the "journals left out" presenting lower impact metrics compared to those selected for indexing by DBLP. At the same time, a significant portion of CS and CS-related literature from Scopus is not indexed by DBLP, and a nonnegligible portion of DBLP's indexed journals seem to have a weak relation (if any) to CS. Factoring in the obvious advantages DBLP brings to the table, such as its open and free access to its data, a simple user interface and its strong CS focus (albeit imperfect), we consider DBLP to be a very solid choice as a first-line CS database. However, DBLP's shortcomings (that is, left out CS contributions or mistakenly indexed non-CS contributions) should be acknowledged and mitigated, when needed, by consulting additional indexing services such as WOS and Scopus.

### ■ Resonanz in Sozialen Netzwerken

### ■ Social Media Feedback

**Sam Power, University of Bristol, UK**

Twitter | [https://twitter.com/sp\\_monte\\_carlo/status/1726663733785837653](https://twitter.com/sp_monte_carlo/status/1726663733785837653)

Have tweeted about it long ago in the past, but DBLP is excellent in certain areas. As I discovered today, on Google Scholar, the bibs for both the VAE paper and Adam paper both point to the arXiv; DBLP nails them both.

**Jonathan Aldrich, Carnegie Mellon University, PA, USA**

Twitter | <https://twitter.com/JAldrichPL/status/172124595328852484>

DBLP is great but it doesn't have nearly the metadata, impact data, or linking facilities that are found in the ACM DL.

**Clément Canonne, University of Sydney, Australia**

Twitter | [https://twitter.com/ccanonne\\_/status/1699004591826997601](https://twitter.com/ccanonne_/status/1699004591826997601)

Google Scholar BibTeX entries are not great, use @dblp\_org whenever possible.

**David Picard, École des Ponts ParisTech, France**

Twitter | [https://twitter.com/david\\_picard/status/1678361583095894019](https://twitter.com/david_picard/status/1678361583095894019)

On est les seuls à utiliser dblp, dixit les bibliothécaires en charge de la bibliométrie. Scopus et autres WOS font référence, avec tout le bien qu'on en pense en info.<sup>39</sup>

<sup>39</sup> engl.: We're the only ones to use dblp, according to the librarians in charge of bibliometrics. Scopus and WOS are the reference, with all the good things we think of them in the news.

**Sebastian Baltes, University of Bayreuth, Germany**Twitter | [https://twitter.com/s\\_baltes/status/1668714268416438273](https://twitter.com/s_baltes/status/1668714268416438273)

Yesterday, I learned that one can use dblp to identify author conflicts based on joint publications. E.g., these are my conflicts for 2019-2023: <https://t.co/Azcs6loAhS> (just expand the "refine by author" list). Thanks to the DBLP team for their help with the query!

**Mark Sanderson, RMIT University, Australia**Twitter | [https://twitter.com/IR\\_oldie/status/1652183494163050496](https://twitter.com/IR_oldie/status/1652183494163050496)

I got a chance to meet @dblp\_MRA who helps run @dblp\_org . Online since 1993, this amazing website enters its 4th decade this year! What an incredible resource DBLP is to the Computer Science community, all free at the point of use.

**Marija Slavkovik, University of Bergen, Norway**Twitter | <https://twitter.com/MSlavkovik/status/1618583680527785984>

What is the path from epistemic trespassing to epistemic citizenship? For computer science, I would say as soon as DBLP finds you ... and if you can teach a 100-level algorithms course.



# 6

**Die Seminare in 2024**

***The 2024 Seminars***

## ■ Algorithms, Complexity, Data Structures

- Beyond-Planar Graphs: Models, Structures and Geometric Representations (24062)
- Triangulations in Geometry and Topology (24072)
- Applied and Combinatorial Topology (24092)
- Low-Dimensional Embeddings of High-Dimensional Data: Algorithms and Applications (24122)
- Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis (24241)
- Theory of Randomized Optimization Heuristics (24271)
- Algebraic and Analytic Methods in Computational Complexity (24381)
- Statistical and Probabilistic Methods in Algorithmic Data Analysis (24391)
- New Tools in Parameterized Complexity: Paths, Cuts, and Decomposition (24411)
- Machine Learning Augmented Algorithms for Combinatorial Optimization Problems (24441)
- Graph Algorithms: Distributed Meets Dynamic (24471)
- Regular Expressions: Matching and Indexing (24472)

## ■ Applications, Interdisciplinary Work

- Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis (24031)
- The Emerging Issues in Bioimaging AI Publications and Research (24042)
- Reviewer No. 2: Old and New Problems in Peer Review (24052)
- Safety Assurance for Autonomous Mobility (24071)
- Computational Approaches to Strategy and Tactics in Sports (24081)
- AI for Social Good (24082)
- Trustworthiness and Responsibility in AI - Causality, Learning, and Verification (24121)
- Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems (24151)
- Computational Metabolomics: Towards Molecules, Models, and their Meaning (24181)
- Evaluation Perspectives of Recommender Systems: Driving Research and Education (24211)
- Classical-Quantum Synergies in the Theory and Practice of Quantum Error Correction (24212)
- Computational Analysis and Simulation of the Human Voice (24242)
- Dynamic Traffic Models in Transportation Science (24281)
- Automated Machine Learning for Computational Mechanics (24282)
- Learning with Music Signals: Technology Meets Education (24302)
- Fair Division: Algorithms, Solution Concepts, and Applications (24401)
- Machine Learning for Protein-Protein and Protein-Ligand Interactions (24451)
- Deep Learning for RNA Regulation and Multidimensional Transcriptomics (24491)
- Coding Theory and Algorithms for Emerging Technologies in Synthetic Biology (24511)
- Quantum Software Engineering (24512)

## ■ Artificial Intelligence, Machine Learning

- Generalization by People and Machines (24192)
- Causal Inference for Spatial Data Analytics (24202)
- Computational Creativity for Game Development (24261)
- Resource-Efficient Machine Learning (24311)
- Leveraging AI for Management Decision-Making (24342)
- Explainable AI for Sequential Decision Making (24372)
- Rethinking the Role of Bayesianism in the Age of Modern AI (24461)
- Human in the Loop Learning through Grounded Interaction in Games (24492)

## ■ Cryptography, Security, Privacy

- Symmetric Cryptography (24041)
- Security and Privacy of Current and Emerging IoT Devices and Systems (24312)
- Next-Generation Secure Distributed Computing (24362)

## ■ Databases, Information Retrieval, Data Mining

- Representation, Provenance, and Explanations in Database Theory and Logic (24032)
- Are Knowledge Graphs Ready for the Real World? Challenges and Perspective (24061)
- Robust Query Processing in the Cloud (24101)
- Shapes in Graph Data: Theory and Implementation (24102)

## ■ Distributed Computation, Networks, Architecture, Systems

- Network Calculus (24141)
- Hardware Support for Cloud Database Systems in the Post-Moore's Law Era (24162)
- Discrete Algorithms on Modern and Emerging Compute Infrastructure (24201)
- Programmable Host Networking (24291)
- Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES) (24351)
- Greening Networking: Toward a Net Zero Internet (24402)
- Research Infrastructures and Tools for Collaborative Networked Systems Research (24462)

## ■ Image Processing, Graphics, Visualization

- Reflections on Pandemic Visualization (24091)
- Art, Visual Illusions, and Data Visualization (24301)

## ■ Human and the Machine

- EU Cyber Resilience Act: Socio-Technical and Research Challenges (24112)
- Resilience and Antifragility of Autonomous Systems (24182)
- Designing Computers' Control Over Our Bodies (24232)
- Teaching Support Systems for Formal Foundations of Computer Science (24251)
- A Game of Shadows: Effective Mastery Learning in the Age of Ubiquitous AI (24272)
- Improving Trust between Humans and Software Robots in Robotic Process Automation (24292)
- Conversational Agents: A Framework for Evaluation (CAFE) (24352)
- Extended Reality Accessibility (24371)

## ■ Software Technology, Programming Languages

- Next Generation Protocols for Heterogeneous Systems (24051)
- Research Software Engineering: Bridging Knowledge Gaps (24161)
- Code Search (24172)
- Automated Programming and Program Repair (24431)
- Reframing Technical Debt (24452)

## ■ Verification, Logic, Formal Methods, Semantics

- From Proofs to Computation in Geometric Logic and Generalizations (24021)
- Logics for Dependence and Independence: Expressivity and Complexity (24111)
- Automated Synthesis: Functional, Reactive and Beyond (24171)
- Stochastic Games (24231)
- Proof Representations: From Theory to Applications (24341)
- Artificial Intelligence and Formal Methods Join Forces for Reliable Autonomy (24361)
- SAT and Interactions (24421)
- Behavioural Metrics and Quantitative Logics (24432)

## 6.1 From Proofs to Computation in Geometric Logic and Generalizations

**Editors:** Ingo Blechschmidt, Hajime Ishihara, and Peter M. Schuster  
**Seminar No.** 24021

Date: January 7–12, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.1

© Creative Commons BY 4.0 license

© Ingo Blechschmidt, Hajime Ishihara, and Peter M. Schuster



**Participants:** Peter Arndt, Steve Awodey, Andrej Bauer, Karim Johannes Becher, Olaf Beyersdorff, Marc Bezem, Ingo Blechschmidt, Ulrik Buchholtz, Gabriele Buriola, Felix Cherubini, Michel Coste, Laura Crosilla, Nicolas Daans, Dominique Duval, Martín H. Escardó, Giulio Fellin, Makoto Fujiwara, Hugo Herbelin, Matthias Hutzler, Hajime Ishihara, Ulrich Kohlenbach, Henri Lombardi, Maria Emilia Maietti, Julien Narboux, Sara Negri, Takako Nemoto, Stefan Neuwirth, Satoru Niki, Paige North, Eugenio Orlandelli, Edi Pavlovic, Iosif Petrakis, Elaine Pimentel, Michael Rathjen, Marie-Françoise Roy, Peter M. Schuster, Monika Seisenberger, Sana Stojanovic-Djurdjevic, Benno van den Berg, Steven J. Vickers

The Dagstuhl Seminar 24021 emerged as a response to the challenges faced by its predecessor, Dagstuhl Seminar 21472, which grappled with pandemic-induced travel restrictions that hindered in-person attendance. The tireless efforts of the Dagstuhl staff notwithstanding, the earlier seminar relied on remote participation, limiting the depth of engagement and interaction among attendees. Many participants advocated for a follow-up seminar on a related topic, which materialized in the form of the present gathering.

Freed from the constraints of travel restrictions, this seminar hosted a dynamic environment characterized by extensive interactions, both structured and informal. Evening sessions provided a platform for casual discussions, enabling participants to delve deeper into topics of interest and forge meaningful connections across different communities within the field. Most notably, the seminar structure allowed ample time for spontaneous working groups.

As compared to the Dagstuhl Seminars the organisers attended in the past, this seminar stands out for an intense interaction between the participants. In the few months after we have already observed push effects to current research in several

directions, e.g. strong negation in constructive mathematics and proof systems, synthetic algebraic geometry especially in the context of homotopy type theory, and topos theory. According to our humble opinion these effects can also be traced back to our stressing of interactive communication formats during the week in Dagstuhl, as there are the lightning talks sessions (one-hour slots of short talks of 5 minutes each) and especially the working groups, unusual both in number and participation.

As a pity, one of the four organisers, Thierry Coquand, could not attend this Dagstuhl Seminar. He still played a decisive role in forming the programme, especially concerning the crucial topics of geometric logic, topos theory and synthetic algebraic geometry with homotopy type theory.

An experimental addition to this year's seminar was the integration of an informal Discord server, serving as a digital hub for participant engagement. This platform not only facilitated pre-seminar planning and coordination but also granted participants a stronger sense of ownership over the seminar proceedings. Through features such as photo-sharing of blackboards and solicitation of talk topics, many participants actively shaped the agenda and direction of discussions.

## 6.2 Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis

6

**Editors: Alessandro Cimatti, Ingo Pill, and Alexander Diedrich**  
**Seminar No. 24031**

Date: January 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.25

© Creative Commons BY 4.0 license

© Alexander Diedrich, Alessandro Cimatti, and Ingo Pill

**Participants:** Kaja Balzereit, Gautam Biswas, Marco Bozzano, Elodie Chanthery, Alessandro Cimatti, Marco Cristoforetti, Philippe Dague, Johan de Kleer, Alexander Diedrich, Kai Dresia, Jonas Ehrhardt, Alexander Feldman, Kenneth D. Forbus, Rene Heesch, Daniel Jung, Lukas Moddemann, Angelo Montanari, Manfred Mücke, Edi Muskardin, Oliver Niggemann, Ingo Pill, Gregory Provan, Belarmino Pulido, Josephine Rehak, Pauline Ribot, Martin Sachenbacher, Anika Schumann, Gerald Steinbauer-Wagner, Markus Stumptner, Anna Sztzyber, Louise Travé-Massuyès, Günther Waxenegger-Wilfing, Katinka Wolter, Franz Wotawa, Marina Zanella, Alois Zoitl



Our goal for this Dagstuhl Seminar was to find approaches that leverage fault diagnosis to build resilient cyber-physical systems through combinations of symbolic, sub-symbolic, and control theoretic approaches.

Cyber-Physical Systems (CPSs), i.e. systems in which mechanical and electrical parts are controlled by computational algorithms, are not only continuously increasing in size and complexity, but they are also required to operate in evolving and uncertain environments, subject to frequent changes and faults. Detecting and correcting faulty behavior is a highly complex task that needs the help of computational algorithms. The constant advances in sensing technology and computational power, as well as the increase in data recording options, enables and also requires us to rely more and more on methods from Artificial Intelligence (AI) for these tasks, i.e. symbolic AI such as planning and reasoning engines, as well as subsymbolic AI like Machine Learning (ML). Sub-symbolic approaches are primarily used to detect symptoms; symbolic reasoning on the other hand provides diagnosis algorithms to identify root causes (from symptoms or observations) or reason about repairs. Furthermore, control engineering methods guide the system back to normal operation (based on the identified root cause). Since these methods come from different fields, they do not always work together in practice.

The research challenge at hand is to combine symbolic a-priori knowledge and learned data, as well as to develop an integrated concept taking both symbolic and sub-symbolic approaches into account. The leading research questions of this seminar are summarised as follows:

- How can a-priori knowledge be combined with data-centric, machine learning-based algorithms?
- Can we integrate a-priori knowledge such as background knowledge about functions, interfaces and operation modes into ML-algorithms to improve model performance?
- Can we use data to learn parts of the symbolic models?

- And can we develop new algorithms which are a synthesis of both worlds, symbolic and subsymbolic?

All of these research questions must be addressed to practical and resilient cyber-physical systems. To tackle these questions, we invited researchers from symbolic AI, sub-symbolic AI, and control engineering to develop a common notion of fault detection and fault handling tasks that takes also the practical needs from industry-scale problems into account. In this regard the seminar also had a secondary function: Traditional symbolic AI diagnosis is located within the Diagnostics community (DX), while sub-symbolic fault diagnosis was traditionally associated with the fault-detection and isolation (FDI) community within the control theory research field. More recently, also the research field of machine learning has created advances with regard to fault diagnosis. Since this seminar brought together researchers from all of these fields, we hope that the seminar created fertile ground for some cross-domain research initiatives.

Besides the individual contributions to the seminar, we used four breakout sessions to brainstorm ideas and next steps following from this seminar:

**1) Breakout Session on Coupling Symbolic and Sub-symbolic Methods for Model Acquisition:** Fusing symbolic methods with sub-symbolic methods in both directions is essential for the creation of resilient systems. The research gap that has been identified is that so far most approaches integrate some symbolic knowledge into the majority of sub-symbolic knowledge, or a small part of sub-symbolic knowledge into a large symbolic knowledge base. But both of these directions have drawbacks and do not automatically lead to models that are well-suited for resilient systems that can be used in practice. One takeaway is the idea to organise a competition that incentivises researchers to develop novel modelling formalisms and diagnosis algorithms that mitigate some of the current drawbacks.

**2) Breakout Session on Causality – How to Generate Knowledge from Data:** The breakout session detailed the importance of high-level causal models in capturing causal relationships within systems. It was discussed where the difficulties in manually crafting these models lie due to their complexity and the even greater challenge of learning causal models directly from data. Crafting causal models manually, one needs a deep understanding of the dependencies. For learning causal models, a large amount of data even for situations which barely occur is needed.

**3) Breakout Session on LLMs for DX – Integrating Large Language Models for Root Cause Diagnosis:** The breakout session featured a comprehensive exploration and discussion on the potential and challenges of using Large Language Models in the topics of the “DX” community. The central aspects that were discussed, revolved around (i) the models themselves and their current and potential future capabilities, (ii) the training data for training and refining LLMs for diagnosis tasks, (iii) potential

application areas, as well as (iv) current, and (v) future trends and topics that should be monitored or covered by the DX community. As a result, the attendees agreed on writing a position paper, which will capture the current potential and drawbacks of LLMs within DX domains.

**4) Breakout Session on Resilient Systems:** For resilient systems we saw that the application and scenario play a significant role when aiming to assess what would be “good” and “bad” behavior for some system. The same goes for the question of whether we would assess the performance of a system in a local or a global context. To this end we identified a set of such relevant scenarios ranking from an energy management scenario at a local home, via the operation of an electric grid, via agents/robots in a collaborative disaster or military scenario, to supply chain management. We also discussed and converged to a definition of resilience that would tailor to all the expressed needs.

## 6.3 Representation, Provenance, and Explanations in Database Theory and Logic

6

**Editors: Pablo Barcelo, Pierre Bourhis, Stefan Mengel, and Sudeepa Roy**  
**Seminar No. 24032**

Date: January 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.49

© Creative Commons BY 4.0 license

© Pablo Barcelo, Pierre Bourhis, Stefan Mengel, and Sudeepa Roy

**Participants:** Antoine Amarilli, Albert Atserias, Pablo Barcelo, Christoph Berkholz, Leopoldo Bertossi, Pierre Bourhis, Florent Capelli, Wolfgang Gatterbauer, Floris Geerts, Amir Gilad, Boris Glavic, Benny Kimelfeld, Ester Livshits, Bertram Ludäscher, Stefan Mengel, Mikaël Monet, Liat Peterfreund, Reinhard Pichler, Cristian Riveros, Sudeepa Roy, Babak Salimi, Pierre Senellart, Christoph Standke, Dan Suciu, Nikolaos Tziavelis, Harry Vinal-Smeeth



### ■ Background and Research area

The Dagstuhl Seminar “*Representation, Provenance, and Explanations in Database Theory and Logic*” (24032) was broadly in *Database Theory*, where the goal is to formalize the theoretical underpinnings of databases and then analyze them with mathematical tools. One of the most fundamental problems in both database theory and systems is efficient query evaluation: given a database and a query, compute the answer to the query on the database. This question has a tight connection to logic, since it has been known for a long time that different fragments of first- or second-order logic can be seen as the core of practical query languages like SQL or Datalog. This seminar focused on three key aspects of query evaluations: *representation*, *provenance*, and *explanations*.

**Representation.** For large datasets, query results can be very large when they are materialized explicitly in the standard form. For efficient query processing and subsequent applications, it is important to *represent* the query answers in a compact fashion. One important form of representations in query evaluation is by *circuits*, which have a long history in complexity theory and AI and can be seen as part of the larger framework of *knowledge compilation* (Darwiche, Marquis, J. Artif. Intell. Res. 2002). Circuits were heavily discussed in several presentations in the seminar. The other aspect of representation that the seminar focused on was the field of *enumeration algorithms* and *direct access*. It first computes a data structure representing the query answers, and then gives an algorithm to extract one answer at a time from the data structure. In this problem, the complexity of the two parts is measured separately: the computation time of the data structure is called the *preprocessing time* and the time of the extraction of each answer is called the *delay*. Typically, the goal of such algorithms is to have a preprocessing time much smaller than the cost of the classical evaluation of the query and very small (ideally constant) delay.

**Provenance.** Data provenance in general refers to how the outputs of a query are generated from the inputs, with a broad goal to enable interpretability, trust, and reproducibility of the queries. A mathematical form of provenance that propagates annotations of inputs to the outputs, called *provenance semirings*, was proposed in a seminal work by Green et al. (PODS 2007). The most specialized case of Boolean semirings captures how an output tuple has been obtained from the inputs with joint usage (joins – translate to conjunctions  $\wedge$ ), and alternative usages (projections or unions – translate to disjunctions  $\vee$ ). Such semirings can be used to understand compactly how outputs are generated from inputs, and have applications in *query evaluation in probabilistic databases* when realization of inputs tuples is uncertain (Dalvi-Suciu, JACM 2012), and in *deletion propagation* or *view update*, to understand how the outputs change if one or more inputs are deleted, without re-computing the query. There are more advanced semirings like tropical semirings that can capture shortest paths in graphs. Compact and efficient knowledge compilations of provenance circuits into *ordered and free binary decision diagrams (OBDDs, FBDDs)*, and more generally as *decomposable deterministic negation normal forms (d-DNNF)* are also important questions in database theory with applications in probabilistic databases (Jha-Suciu, ICDT 2011; Beame et al., ACM Trans. Database Syst. 2017; Monet, PODS 2020).

**Explanations.** While provenance provides one approach to explaining query answers capturing how the query answers are generated, in many applications, other forms of insights as explanations are desired for understanding contributions of inputs, trends and anomalies in the outputs, and deciding next course of actions or recourse. Recently, explanations based on the widely known *Shapley values* from co-operative game theory have been used in database theory to measure the relevance of a certain database fact to a query answer (Deutch et al, SIGMOD 2022; Livshits et al., ICDT 2021), and to measure the relevance of inputs to the outcome of an ML classifier (Arenas et al., AAAI 2021).

Complexity, applications, and algorithms for explanations by Shapley values were heavily discussed in the seminar. Since the naive computation of Shapley values is intractable as it includes a summation over exponentially many subsets, one of the main themes behind this investigation has been the identification of practically relevant classes of database queries for which such explanations can be computed in polynomial time, possibly using knowledge compilation forms. Apart from Shapley values, other forms of explanations, including that of aggregated database query answers (e.g., Roy-Suciu, SIGMOD'14) and connections

of explanations with data privacy, fairness, and causal inference were discussed in the seminar. This way the seminar connected the field of database theory to the field of *responsible data science* that is of paramount importance in real world.

### ■ Acknowledgements

We are grateful to the Scientific Directorate and to the staff of the Schloss Dagstuhl – Leibniz Center for Informatics for their support of this seminar.

## 6.4 Symmetric Cryptography

**Editors: Christof Beierle, Bart Mennink, María Naya-Plasencia, and Yu Sasaki**  
**Seminar No. 24041**

Date: January 21–26, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.72

© Creative Commons BY 4.0 license

© Christof Beierle, Bart Mennink, María Naya-Plasencia, and Yu Sasaki

**Participants:** Zahra Ahmadian, Subhadeep Banik, Zhenzhen Bao, Christof Beierle, Yanis Belkheyyar, Ritam Bhaumik, Christina Boura, Anne Canteaut, Patrick Derbez, Christoph Dobraunig, Orr Dunkelman, Avijit Dutta, María Eichseder, Patrick Felke, Henri Gilbert, Lorenzo Grassi, Rachele Heim Boissier, Akiko Inoue, Ryoma Ito, Tetsu Iwata, Ashwin Jha, Antoine Joux, Virginie Lallemand, Nils Gregor Leander, Charlotte Lefevre, Gaëtan Leurent, Willi Meier, Bart Mennink, Kazuhiko Minematsu, Mridul Nandi, María Naya-Plasencia, Patrick Neumann, Léo Perrin, Bart Preneel, Shahram Rasoolzadeh, Christian Rechberger, Yann Rotella, Sondre Rønjom, Dhiman Saha, Yu Sasaki, Ferdinand Sibleyras, Meltem Sonmez Turan, Siwei Sun, Stefano Tessaro, Aishwarya Thiruvengadam, Tyge Tiessen, Yosuke Todo, Aleksei Udovenko, Qingju Wang



IT Security plays an increasingly crucial role in our everyday life and business. Virtually all modern security solutions are based on cryptographic primitives. Symmetric cryptography deals with the case where both the sender and the receiver of a message are using the same key. Due to their good performance, symmetric cryptosystems are the main workhorses of cryptography and are highly relevant not only for academia, but also for industrial activities. For this Dagstuhl Seminar we focused on several topics, which we believe to be of great importance for the research community and, likewise, to have a positive impact on industry and the deployment of secure crypto in the future.

- **Follow Up on Main Results from Last Dagstuhl Seminar.** At the last Dagstuhl Seminar on symmetric cryptography in 2022, the participants were divided into six groups in order to discuss research topics proposed by each participant. The discussions were very productive and there were and will be publications from several groups. We believe that the discussions and results from these 2022 work groups reflect the main interests of the community and are useful topics to continue to discuss at the Dagstuhl Seminar in 2024. Participants at the 2024 Dagstuhl Seminar who also participated in the work groups in 2022 were thus invited to present their finished results.
- **Design and Analysis of Symmetric Crypto for New Applications.** Recently, the design of symmetric-key primitives has started to focus on different types of optimization. Those optimizations could be with respect to performance and with respect to special security requirements. Stated differently, one first considers a target application (such as multi-party computation or non-interactive zero-knowledge proofs), and only then designs symmetric-key primitives for this purpose. This causes a paradigm shift in design criteria. During this seminar, we explored the security of recently introduced ciphers that were designed specifically for such

target applications, and develop novel ciphers with improved security arguments and guarantees.

- **Generic Analysis of Emerging Modes.** Permutation-based cryptography has gained astounding popularity in the last decade, and security proofs are performed in the ideal permutation model. A similar phenomenon is visible in various ideal cipher-based constructions that have appeared recently. In this seminar, we explored how results with different models (such as a standard model and an ideal model) compare from a theoretical perspective, and investigated what cryptanalytical results on certain primitives mean for the targeted construction.

### ■ Seminar Program

The seminar program consisted of short presentations and group meetings. Presentations were about the above topics and other relevant areas of symmetric cryptography, including state-of-the-art cryptanalytic techniques and new designs. The list of abstracts for talks given during the seminar can be found in the full report. Also, participants met in smaller groups and spent a significant portion of the week, each group intensively discussing a specific research topic. There were seven research groups:

- Cryptanalysis of Poseidon;
- Cryptanalysis of TEA-3;
- Exploitation of the wrong key randomization hypothesis non-conformity in key recovery attacks;
- Cryptanalysis of SCARF;
- Differential cryptanalysis and more;
- Key control security;
- Security of sponge combiners.

On the last day of the week the leaders of each group gave brief summaries of achievements. An abstract corresponding to each research group can be found below. Some teams continued working on the topic after the seminar and started new research collaborations.

## 6.5 The Emerging Issues in Bioimaging AI Publications and Research

Editors: Jianxu Chen, Florian Jug, Susanne Rafelski, and Shanghang Zhang

Seminar No. 24042

Date: January 21–24, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.90

© Creative Commons BY 4.0 license

© Jianxu Chen, Florian Jug, Susanne Rafelski, and Shanghang Zhang



**Participants:** Chao Chen, Jianxu Chen, Evangelia Christodoulou, Beth Cimini, Gaole Dai, Meghan Driscoll, Edward Evans III, Matthias Gunzer, Andrew Hufton, Florian Jug, Anna Kreshuk, Thomas Lemberger, Alan Lowe, Shalin Mehta, Axel Mosig, Matheus Palhares Viana, Constantin Pape, Anne Plant, Susanne Rafelski, Ananya Rastogi, Albert Sickmann, Rita Strack, Nicola Strisciuglio, Aubrey Weigel, Assaf Zaritsky, Shanghang Zhang, Han Zhao

### ■ Seminar Structure and Organization

The seminar was divided into three specific directions: ethical considerations in bioimaging AI research and publications, performance reporting on bioimaging AI methods in publications and research, and future research directions of bioimaging AI focusing on validation and robustness. The seminar was structured into two parts: the first half focused on presentations and information sharing related to these three major directions to align experts from different fields, and the second half concentrated on in-depth discussions of these topics.

Given the highly interdisciplinary nature of the seminar, we took two specific steps to facilitate smooth communication and discussion among researchers with diverse backgrounds.

First, about six to eight weeks before the seminar, we sent out a survey to gather potential topics each participant could present within the seminar's overarching theme. We collaborated with several participants to choose or adjust their presentation topics to ensure the effectiveness in this interdisciplinary setting. Based on the survey responses, the presentation and information-sharing

### ■ Presentations, discussions and outcomes

#### ■ Overview of the scientific talks

The seminar began with presentations by editors from Nature Methods and Cell Press, who shared their insights on existing and emerging issues in bioimaging AI publications. Following this, general bioimage analysis validation issues were discussed from both a biological application perspective and an algorithmic metric perspective. These presentations were succeeded by specific application talks demonstrating how AI-based bioimage analysis is utilized and validated in high-throughput biological

portion (the first half) of our seminar began with two keynotes from editors who handle bioimaging AI papers, sharing their insights and the existing efforts by publishers. We then organized all presentations to progress from a focus on biology to bioimaging AI, and finally to AI, ensuring coverage of the full spectrum of necessary knowledge for our in-depth discussions in the second half.

Second, at the beginning of the seminar, we allocated two minutes for each participant for a quick introduction and to briefly rate their experience and expertise on a scale in the range of [B5, B4, B3, B2, B1, 0, A1, A2, A3, A4, A5], with B5 representing pure biology and A5 representing pure AI. Participants could select a single value, multiple values, or a range of values. This was not intended to stereotype participants but to facilitate easier communication. For example, if a participant with experience in the range of B5 to B3 spoke with two others during a coffee break, one with experience from B3 to A1 and the other from A3 to A5, different communication strategies would be necessary for effective discussions. The distribution of self-identified experience is summarized in the histogram below (see Fig. 6.1).

applications [1]. The remainder of day one focused on bioimaging AI validation through explainable AI [2], [3], [4] and existing tools [5], as well as community efforts in deploying FAIR (Findable, Accessible, Interoperable, Reusable) AI tools for bioimage analysis [6].

The second day commenced with several theoretical AI talks introducing key concepts related to model robustness, fairness, and trustworthiness [7]. These were followed by two presentations showcasing state-of-the-art AI algorithms applied in bioimaging [8], [9], and an overview of the application of foundation models in bioimaging [10]. The scientific presentation portion of the seminar concluded with a talk about the pilot work initiated by

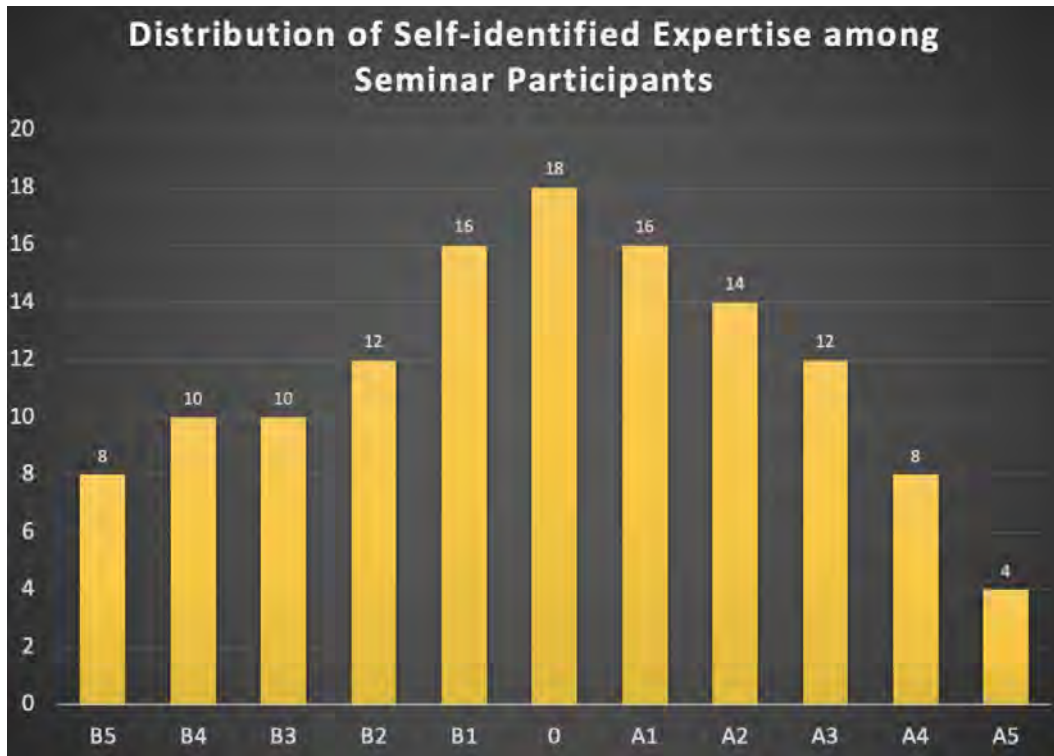


Fig. 6.1 Histogram of the distribution of expertised self-identified by seminar participants.

the EMBO (European Molecular Biology Organization) Press on research integrity and AI integration in publishing and trust. This talk also served as a transition into the in-depth discussions that comprised the second part of the seminar.

■ **Summary of discussions and key outcomes**

After the scientific presentation part of the seminar, the participants naturally reach the agreement on doing the discussion in a four-quadrant manner, as illustrated below in Fig. 6.2.

	“Users” of bioimaging AI	“Makers” of bioimaging AI
In-domain technical considerations	1	2
out-of-domain implications	3	4

Fig. 6.2 The four-quadrant for organizing the in-depth discussion.

Here are some examples of what emerges from discussions in each quadrant.

**I. What are some technical considerations that users of AI should pay attention to?**

When using a specific bioimage analysis model, it is crucial for users to have clear biological questions that align with the technical limitations of the bioimaging AI models. This is known as application-appropriate validation [11]. For example, the trustworthiness or validity of an AI-based microscopy image denoising model may differ significantly between a study that requires merely counting the number of nuclei in an image and one that aims to quantify the morphological properties of the nuclei.

**II. What are some technical considerations that makers of AI should pay attention to?**

When developing a bioimaging AI model, comprehensive evaluations and ablation studies are essential to explicitly demonstrate

the model’s limitations or potential failures. For instance, evaluating a cell segmentation model under different conditions, such as various magnifications, signal-to-noise ratios, cell densities, and possibly different microscope modalities, is highly beneficial. Providing a clear and detailed definition of the conditions under which the model has been evaluated helps users determine whether the model can be directly applied to their images or if it needs retraining or fine-tuning.

**III. What are some important things the users of AI should make sure the makers of AI are aware of or should make clear to the makers of AI?**

One example is the inherent presence of batch effects in biology, such as variations in fluorescence microscopy image quality due to different batches of dyes or slight morphological differences in cells from different colony positions. For effective interdisciplinary collaboration, it would be very helpful if biologists can

clearly describe data acquisition processes and potential batch effects. This enables AI developers to consider these factors in their training sets, validation strategies, and model designs.

#### IV. What do the makers of AI need to make sure the users of AI know?

There is a lot of information that AI method developers need to help biologists think together. For example, in some collaborative projects, AI researchers need to guide their biologist collaborators how to best provide their data. For instance, the data to be analyzed to answer biological questions can be different from special data acquired merely for training the AI models, which could be referred to as a “training assay” [12], i.e., special experimental assays only made for effective model training.

The discussions highlighted in the four quadrants are only examples from the seminar. A follow-up “white paper”-like manuscript based on the full discussions is being planned as a resource for the bioimaging AI community.

A specific topic that emerged from the seminar was the interpretability and explainability of bioimaging AI models. This is evident from the word cloud generated during the discussions, as shown in Fig. 6.3. A follow-up seminar specifically focusing on this topic is being planned for the bioimaging AI community in the coming years.

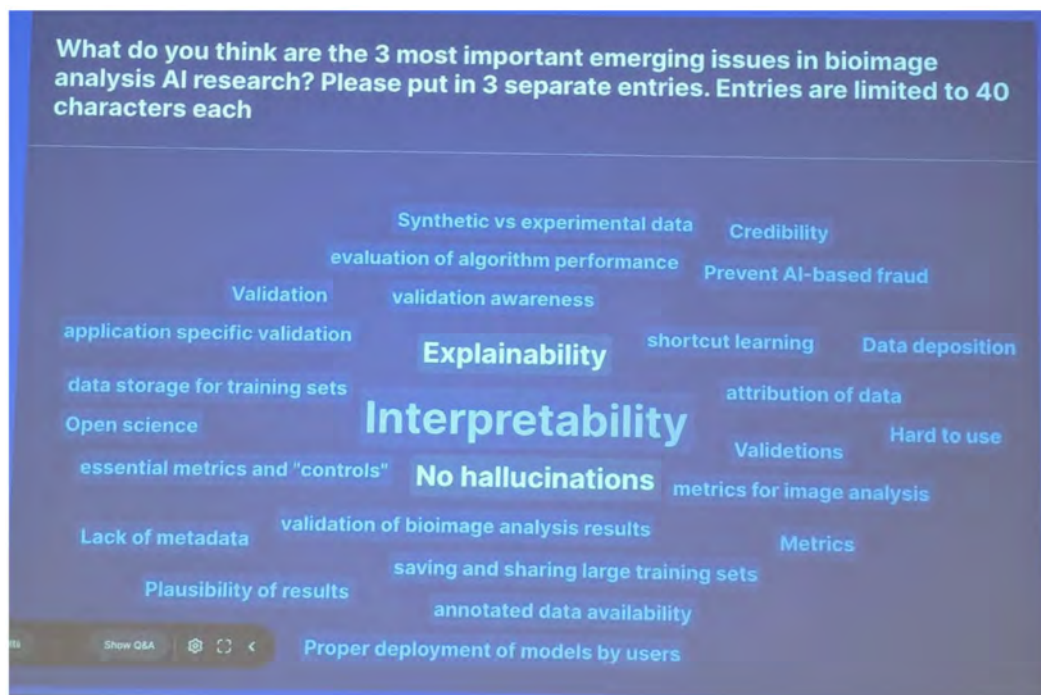


Fig. 6.3 Word cloud generated during the discussion about the emerging issues in bioimaging AI publications and research.

Besides the science program, the seminar provided valuable opportunities for social connections and networking. Due to the pandemic, many researchers had previously only met virtually, making the in-person interactions feel like a reunion. The diversity of research fields among participants, including wet-lab biologists and machine learning theorists with little biology experience, created unique networking opportunities. They would otherwise have rare opportunities to meet in traditional conferences. Biologists expressed that they gained new insights into the theories behind machine learning methods they had used, motivating them to rethink their future research designs. Conversely, machine learning researchers showed strong interest in collaborating with the bioimaging community to address fundamental challenges such as robustness and explainability.

## ■ Conclusions

This Dagstuhl Seminar on “The Emerging Issues in Bioimaging AI Publications and Research” successfully united a diverse group of experts from experimental biology, computational biology, bioimage analysis, computer vision, and AI research. The seminar facilitated in-depth discussions on ethical considera-

tions, performance reporting, and future research directions in bioimaging AI, highlighting the crucial need for interdisciplinary collaboration and communication.

Through structured presentations and interactive discussions, participants underscored the importance of clear communication between AI developers and users, comprehensive model validation, and awareness of biological batch effects. The seminar emphasized the necessity for application-appropriate validation and detailed reporting of AI model conditions to enhance the trustworthiness and applicability of bioimaging AI methods. Furthermore, the seminar provided a valuable platform for social interactions and networking, bridging gaps between researchers from different fields and fostering new collaborations.

In conclusion, the seminar not only advanced discussions on critical issues in bioimaging AI publications but also laid the foundation for ongoing collaboration and innovation in the field. Planned follow-up activities will further contribute to the development and ethical application of AI in bioimaging research. The success of this seminar underscores the importance of continuous communication and cooperation in addressing the emerging challenges in bioimaging AI publications and research.

## Acknowledgement

We are grateful to all seminar participants for their insightful contributions and the engaging discussions they fostered, especially in the interdisciplinary setting with a wide spectrum

of expertise. We also sincerely thank the Dagstuhl Scientific Directorate for the opportunity to organize this event. Finally, our deepest appreciation goes to the exceptional Dagstuhl staff whose support was instrumental in making the seminar a success.

## References

- 1 Z. Cibir et al., “ComplexEye: a multi-lens array microscope for high-throughput embedded immune cell migration analysis,” *Nat. Commun.*, vol. 14, no. 1, p. 8103, Dec. 2023, doi: 10.1038/s41467-023-43765-3.
- 2 Christopher J. Soelistyo and Alan R. Lowe, “Discovering Interpretable Models of Scientific Image Data with Deep Learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024, pp. 6884–6893. [Online].
- 3 Christopher J. Soelistyo, Guillaume Charras, and Alan R. Lowe, “Virtual Perturbations to Assess Explainability of Deep-Learning Based Cell Fate Predictors,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2023, pp. 3971–3980. [Online].
- 4 D. Schuhmacher et al., “A framework for falsifiable explanations of machine learning models with an application in computational pathology,” *Med. Image Anal.*, vol. 82, p. 102594, Nov. 2022, doi: 10.1016/j.media.2022.102594.
- 5 L. M. Moser et al., “Piximi – An Images to Discovery web tool for bioimages and beyond.” Jun. 04, 2024. doi: 10.1101/2024.06.03.597232.
- 6 W. Ouyang et al., “BioImage Model Zoo: A Community-Driven Resource for Accessible Deep Learning in BioImage Analysis,” *Bioinformatics*, preprint, Jun. 2022. doi: 10.1101/2022.06.07.495102.
- 7 D. Guo, C. Wang, B. Wang, and H. Zha, “Learning Fair Representations via Distance Correlation Minimization,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 2, pp. 2139–2152, Feb. 2024, doi: 10.1109/TNNLS.2022.3187165.
- 8 Saumya Gupta, Yikai Zhang, Xiaoling Hu, and Prateek Prasanna, “Topology-aware uncertainty for image segmentation,” presented at the *Advances in Neural Information Processing Systems*, 2024.
- 9 G. Dai et al., “Implicit Neural Image Field for Biological Microscopy Image Compression.” arXiv, 2024. doi: 10.48550/ARXIV.2405.19012.
- 10 A. Archit et al., “Segment Anything for Microscopy,” *Bioinformatics*, preprint, Aug. 2023. doi: 10.1101/2023.08.21.554208.
- 11 J. Chen, M. P. Viana, and S. M. Rafelski, “When seeing is not believing: application-appropriate validation matters for quantitative bioimage analysis,” *Nat. Methods*, vol. 20, no. 7, pp. 968–970, Jul. 2023, doi: 10.1038/s41592-023-01881-4.
- 12 J. Chen et al., “The Allen Cell and Structure Segmenter: a new open source toolkit for segmenting 3D intracellular structures in fluorescence microscopy images,” *Cell Biology*, preprint, Dec. 2018. doi: 10.1101/491035.

## 6.6 Next Generation Protocols for Heterogeneous Systems

**Editors:** Stephanie Balzer, Marco Carbone, Roland Kuhn, and Peter Thiemann

**Seminar No. 24051**

Date: January 28 – February 2, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.108

© Creative Commons BY 4.0 license

© Stephanie Balzer, Marco Carbone, Roland Kuhn, and Peter Thiemann



**Participants:** Sören Auer, Piero Andrea Bonatti, Juan Cano de Benito, Andrea Cimmino, Michael Cochez, John Domingue, Michel Dumontier, Nicoletta Fornara, Irini Fundulaki, Sandra Geisler, Anna Lisa Gentile, Paul Groth, Peter Haase, Andreas Harth, Olaf Hartig, James A. Hendler, Aidan Hogan, Katja Hose, Luis-Daniel Ibáñez, Ryutaro Ichise, Ernesto Jiménez-Ruiz, Timotheus Kampik, Sabrina Kirrane, Manolis Koubarakis, Luis C. Lamb, Julian Padget, Harshvardhan J. Pandit, Heiko Paulheim, Axel Polleres, Philipp D. Rohde, Daniel Schwabe, Oshani Seneviratne, Elena Simperl, Chang Sun, Aisling Third, Ruben Verborgh, Maria-Esther Vidal, Sonja Zillner

This Dagstuhl Seminar followed the earlier Dagstuhl Seminars 17051 “Theory and Applications of Behavioural Types” and 21372 “Behavioural Types: Bridging Theory and Practice”. Whereas Seminar 17051 was focusing on theoretical aspects of behavioural types, and Seminar 21372 focused on bridging the gap with practical application, this seminar was much broader and aimed at extending to other communities such as security and other areas of programming languages.

### ■ Initial preparations

Based on the ideas of our seminar proposal, we established four key general areas: quantitative systems, verification, mechanisation, and security. We assigned each area to a day of the week (from Monday to Thursday) and asked an invitee representative of the area to give an introductory talk. Then, each of these talks was followed by other talks and breakout rooms related to the area. Breakout rooms were established during the seminar based on discussions with the rest of the participants. As a result of this, the first part of the week consisted primarily of talks, while the second part included more time for breakout sessions.

### ■ Activities and outcomes

Throughout the seminar, the participants gathered in focused breakout groups: the findings of the breakout groups are described in more detail in the last part of the report. The participants of several breakout groups have agreed to continue their work and collaboration after the seminar.

In addition to these more structured breakout sessions there were further lively improvised meetings and discussions (especially after dinner) which are not summarised in the report.

Overall, we believe that the seminar activities were a success. At the end of the seminar the participants agreed to remain in contact to continue the discussions, and foster new collaborations. There was strong enthusiasm for organising a follow-up Dagstuhl Seminar in the future, perhaps taking place in about 1–2 years time. One concrete outcome was the submission of a position paper (cf. the working group “Typing Across Heterogeneous Components”) that has been accepted and presented at PLACES 2024 (co-located with ETAPS).

## 6.7 Reviewer No. 2: Old and New Problems in Peer Review

**Editors:** Iryna Gurevych, Anna Rogers, and Nihar Shah  
**Seminar No. 24052**

Date: January 28 – February 2, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.1.130

© Creative Commons BY 4.0 license

© Anna Rogers, Nihar Shah, and Iryna Gurevych

**Participants:** Osama Mohammed Afzal, Koen Dercksen, Nils Dycke, Alexander Goldberg, Iryna Gurevych, Jason Hartline, Tom Hope, Dirk Hovy, Eddie Kohler, Jonathan Kummerfeld, Ilia Kuznetsov, Anne Lauscher, Kevin Leyton-Brown, Sheng Lu, Dorsa Majdi, Mausam, Bahar Mehmani, Margot Mieskes, Aurélie Névéol, Danish Pruthi, Lizhen Qu, Anna Rogers, Roy Schwartz, Nihar Shah, Noah A. Smith, Tamar Solorio, Jingyan Wang, Xiaodan Zhu



### Background

Peer review is the best mechanism for assessing scientific validity of new research that we have so far. But this mechanism has many well-known issues, such as the different incentives of the authors and reviewers, difficulties with preserving reviewer and author anonymity to avoid social biases [1–5], confirmation and other cognitive biases [6–10], that even researchers fall prey to. These intrinsic problems are exacerbated in interdisciplinary fields like Natural Language Processing (NLP), where groups of researchers may vary so much in their methodology, terminology, and research agendas, that sometimes they have trouble even recognizing each other’s contributions as “research” [11].

Our Dagstuhl Seminar covered a range of topics related to organization of peer review in NLP, Machine Learning (ML), and venues more broadly in Artificial Intelligence for intelligent support of peer-reviewing, including the following:

- Improving the paper-reviewer matching by processes/algorithms that take into account both topic matches and reviewer interest in a given research question.
- Peer review vs methodological and demographic diversity in the field.
- Better practices for designing review forms and peer review policies.
- Improving the structural incentives for reviewers.
- Use of NLP and ML for intelligent peer reviewing support: increasing the quality and efficiency of peer review, opportunities and challenges.
- Peer-reviewing and research integrity.

### Goals

We intended for the seminar to serve as a point of reflection on decades of personal experience of the participants in organizing different kinds of peer-reviewed venues in NLP and beyond,

enabling an in-depth discussion of what has been tried, what seems to work and what doesn’t. The objectives of the seminar included collaborative research on the methodological challenges of peer review, NLP and ML for intelligent support of peer-reviewing and actionable proposals, for example for paper-reviewer assignment policies and peer reviewing guidelines and workflows, informed by the experience of participants as chairs, editors, conference organizers, and reviewers.

**Outcomes** The seminar was attended by researchers at different levels of seniority and from a variety of research backgrounds. While a large number of the attendees represented the Natural Language Processing community, about a third represented other communities within the broader sphere of Machine Learning. Most discussions focused on the peer review in the world of ultra-large conferences with thousands of submissions, but we also had a senior representative from fields where journals are most prominent, and hence an opportunity to learn from their experience.

**Knowledge Sharing** The seminar started by contributed talks by a diverse group of participants (see Section 3), which allowed us to share relevant experience and research findings pertinent to the topics of the seminar, across communities. Peer review issues are at most discussed in the business meetings of specific conferences, and there are hardly any opportunities to share this knowledge across communities. Hence, this knowledge-sharing section of the seminar by itself has been unique, and it proved to be useful to establish a common ground and points of reference for subsequent work during the seminar.

**Problem elucidation** After the contributed talks, all the subsequent work was organized into breakout sessions (two running in parallel) on the following topics:

- Integrity issues in peer review (2 sessions)
- Diversity issues in peer review (3 sessions)
- Assisting peer review with NLP (3 sessions)
- Peer review policies (2 sessions)
- Incentives in peer review (3 sessions)
- Paper-reviewer matching (3 sessions)

The work in all these sessions combined brainstorming, establishing common ground and terms, discussing practical solutions for specific problems that were tried in various communities represented by the participants, and ideas for the future. Summaries of work in all the above topics are provided in Section 4 of the full report.

There were also two slots reserved for unstructured breakouts, and every day concluded with an overall summary session in which the leads for various topics summarized the discussions in that day.

**Research program and community formation** The key outcome of the seminar is a white paper with the working title “*What Can NLP do for Peer Review?*”, co-authored by the majority of the participants of the seminar. It formulates the goals and research agenda of assisting peer review with NLP techniques, and we hope that it would play a key role in shaping this research field. This paper is available at [12]. It is accompanied by a repository for tracking research papers in this area, available at <https://github.com/OAfzal/nlp-for-peer-review>.

**Concrete policies** The work in various breakout sessions culminated in the proposal of a new peer review committee for the Association of Computational Linguistics (ACL), that would oversee the systematic research and data-driven peer review

policy development in the NLP community. This proposal has already been formally submitted to the ACL board, and generally approved. The work on formally establishing and announcing the committee will be finished in 2024.

**Research problems and collaborations** This Dagstuhl Seminar also helped surface and crystallize a number of open problems, and alongside, helped establish inter-disciplinary collaborations for working on them, which may not have happened if not for this seminar.

## ■ Next steps

This Dagstuhl Seminar brought together an international, community of NLP and ML researchers from academia and industry to discuss the problems with peer review in large-scale conferences. This is a topic for which various subcommunities have different practices, expectations, and strong opinions, and the seminar brought much discussion throughout all days of the seminar (and also long into the night). This was also a unique opportunity to share the lessons learned the hard way, on issues which are often misconstrued as merely organizational issues. In fact, this is something to be seriously discussed as a research problem, for which much conceptual and empirical work is needed.

We hope that this seminar was the first in a series of events devoted to this topic, and that this inaugural event proves pivotal in the formation of a cohesive research community. The white paper prepared as the main outcome of this seminar aims to galvanize the NLP and ML communities by offering them a wide selection of realistic research problems with peer review as an application area.

## ■ References

- 1 Jürgen Huber, Sabiou Inoua, Rudolf Kerschbamer, Christian König-Kersting, Stefan Palan, and Vernon L. Smith. Nobel and novice: Author prominence affects peer review. *Proceedings of the National Academy of Sciences*, 119(41):e2205779119, October 2022.
- 2 Inna Smirnova, Daniel M. Romero, and Misha Teplitskiy. The bias-reducing effect of voluntary anonymization of authors’ identities: Evidence from peer review, January 2023.
- 3 Andrew Tomkins, Min Zhang, and William D. Heavlin. Reviewer bias in single- versus double-blind peer review. *Proceedings of the National Academy of Sciences*, 114(48):12708–12713, 2017.
- 4 Charvi Rastogi, Ivan Stelmakh, Xinwei Shen, Marina Meila, Federico Echenique, Shuchi Chawla, and Nihar Shah. To ArXiv or not to ArXiv: A study quantifying pros and cons of posting preprints online. *arXiv preprint arXiv:2203.17259*, 2022.
- 5 Emaad Manzoor and Nihar B Shah. Uncovering latent biases in text: Method and application to peer review. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4767–4775, 2021.
- 6 Jian Wang, Reinhilde Veugelers, and Paula Stephan. Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy*, 46(8):1416–1436, October 2017.
- 7 J. A. Garcia, Rosa Rodriguez-Sánchez, and J. Fdez-Valdivia. Confirmatory bias in peer review. *Scientometrics*, 123(1):517–533, April 2020.
- 8 David M. Allen and James W. Howell, editors. *Groupthink in Science: Greed, Pathological Altruism, Ideology, Competition, and Culture*. Springer International Publishing, Cham, 2020.
- 9 Carole J Lee. Commensuration bias in peer review. *Philosophy of Science*, 82(5):1272–1283, 2015.
- 10 Ivan Stelmakh, Nihar B Shah, Aarti Singh, and Hal Daumé III. Prior and prejudice: The novice reviewers’ bias against resubmissions in conference peer review. volume 5, pages 1–17. ACM New York, NY, USA, 2021.
- 11 Anna Rogers and Isabelle Augenstein. What can we do to improve peer review in NLP? In Trevor Cohn, Yulan He, and Yang Liu, editors, *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1256–1262. Association for Computational Linguistics, November 2020.

## ■ References (continued)

12 Ilia Kuznetsov, Osama Mohammed Afzal, Koen Dercksen, Nils Dycke, Alexander Goldberg, Tom Hope, Dirk Hovy, Jonathan K. Kummerfeld, Anne Lauscher, Kevin Leyton-Brown, Sheng Lu, Mausam, Margot Mieskes, Aurélie Névéol, Danish Pruthi, Lizhen Qu, Roy Schwartz,

Noah A. Smith, Tamar Solorio, Jingyan Wang, Xiaodan Zhu, Anna Rogers, Nihar B. Shah, Iryna Gurevych. What Can Natural Language Processing Do for Peer Review?”, CoRR, Vol. abs/2405.06563, 2024.

## 6.8 Are Knowledge Graphs Ready for the Real World? Challenges and Perspective

**Editors: David Chaves-Fraga, Oscar Corcho, Anastasia Dimou, and Maria-Esther Vidal**  
Seminar No. 24061

Date: February 4–9, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.1

© Creative Commons BY 4.0 license

© David Chaves-Fraga, Oscar Corcho, Anastasia Dimou and Maria-Esther Vidal



**Participants:** Marina Aguado, Eva Blomqvist, Carlos Buil-Aranda, Irene Celino, David Chaves-Fraga, Souripriya Das, Christophe Debruyne, Michel Dumontier, Sandra Geisler, Paul Groth, Peter Haase, Olaf Hartig, Katja Hose, Samaneh Jozashoori, Eduard Kamburjan, Craig A. Knoblock, Vanessa López, Edelmira Pasarella, Anisa Rula, Dylan Van Assche

Graphs and knowledge bases have been around for many decades, and research results have had a tremendous impact on areas such as mathematics, artificial intelligence, and databases. However, although the term has been coined by the scientific community, technological developments and astronomical data growth have made knowledge graph (KG) management a fundamental topic in various areas of computer science today. The scientific and industrial communities have responded to the emerging field of knowledge management. As a result, formal frameworks for defining and representing KGs, as well as methods for creating, exploring, and analyzing KGs, have flourished to make KGs a reality. However, despite the tangible results, sustainability is still compromised by the lack of transparent and accountable management of CCs. The real-world application of KGs requires programming paradigms for KG management, transparent data integration and quality assessment techniques, and methods for maintaining access control and privacy. In addition to technological advances, societal adjustments can have a tremendous impact on the management of KGs. The seminar addressed these socio-technical challenges with a mix of invited talks, lightning talks, and small group workshops as follows:

### The Incremental Creation of Knowledge Graphs.

Creating a Knowledge Graph (KG) involves several open research challenges, such as data extraction, data quality, data integration, and data security. It also requires attention to architectural aspects such as scalability and interoperability. A working group was formed to discuss and focus on two main topics: the definition of a general pipeline for KG construction and its relationship to data quality. The main outcome is a standard formalization of the KG construction lifecycle and its associated components. This definition is accompanied by quality measures and provenance tracking of all steps.

### Support of Knowledge Graph Implementation.

Software engineering and programming languages have created approaches and techniques that support complex tasks during software development such as software dependencies, error identification, testing, syntactic validation, software lifecycle, etc. We look into these proposals to determine a set of requirements in software lifecycle management for knowledge graphs. They will improve and facilitate the implementation of knowledge graphs in industrial and complex environments, taking into account the relationships and dependencies between all the artifacts used (ontologies, shapes, mappings, tests, etc.) as well as their evolution and versioning. To achieve this goal, we believe that it is necessary to have a better understanding and general overview of how knowledge graphs are implemented. Therefore, a workshop on this topic has been proposed at ISWC2024<sup>40</sup>. After its celebration, the next step will be to create a community around this topic with researchers and industry stakeholders to standardize and implement the identified challenges/requirements.

### Access Control in Decentralized Knowledge Graphs.

Exploring access control in decentralized Knowledge Graphs has been a relatively underexplored area. Specifically, mechanisms for restricting access to knowledge to safeguard confidential information and personal data, as well as establishing consent models for the processing of personal data, have not received substantial attention within the realm of Knowledge Graph management. Additionally, ensuring compliance with usage policies has been inadequately addressed, particularly in the context of decentralized Knowledge Graphs. During the seminar, a dedicated group convened to deliberate on approaches for managing Knowledge Graphs across a federation of decentralized instances.

<sup>40</sup> <https://w3id.org/soflim4kg>

**A New Generation of Knowledge Engineers.**

Improving the utilization and management of knowledge graphs requires educating a diverse audience about both the social and technical aspects of knowledge work. To address this need, a dedicated working group was established. This group conducted an analysis to identify existing educational resources and gaps in knowledge, exploring how consensus could be

fostered among various stakeholders in the field. Moreover, the group investigated the specific educational requirements tailored to different audiences, including professional students, undergraduates, and postgraduates. By thoroughly examining these aspects, the working group aimed to formulate strategies for enhancing education and understanding in the domain of knowledge graph utilization and management.

## 6.9 Beyond-Planar Graphs: Models, Structures and Geometric Representations

**Editors:** Vida Dujmović, Seok-Hee Hong, Michael Kaufmann, and János Pach  
**Seminar No.** 24062

Date: February, 4–9, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.71

© Creative Commons BY 4.0 license

© Vida Dujmović, Seok-Hee Hong, Michael Kaufmann, and János Pach



**Participants:** Patrizio Angelini, Michael A. Bekos, Therese Biedl, Markus Chimani, Sabine Cornelsen, Giordano Da Lozzo, Giuseppe Di Battista, Emilio Di Giacomo, Walter Didimo, Vida Dujmovic, Stefan Felsner, Henry Förster, Fabrizio Frati, Michael Hoffmann, Seok-Hee Hong, Michael Kaufmann, Stephen G. Kobourov, Jan Kratochvil, Giuseppe Liotta, Anna Lubiw, Fabrizio Montecchiani, Pat Morin, Yoshio Okamoto, János Pach, Maurizio Patrignani, Sergey Pupyrev, Ignaz Rutter, Csaba Tóth, Géza Tóth, Torsten Ueckerdt, Pavel Valtr, Alexander Wolff

Many big data sets in various application domains have complex relationships, which can be modeled as *graphs*, consisting of entities and relationships between them. Consequently, graphs are extensively studied in both mathematics and computer science. In particular, *planar graphs*, which can be drawn without edge crossings in the plane, form a distinguished role in graph theory and graph algorithms. Many structural properties of planar graphs are investigated in terms of excluded minors, low density, and small separators, leading to efficient planar graph algorithms. Consequently, fundamental algorithms for planar graphs have been discovered.

However, most real-world graphs, such as social networks and biological networks, are *nonplanar*. For example, the scale-free networks, which are used to model web graphs, social networks, and biological networks, are globally sparse nonplanar graphs with locally dense clusters and low diameters. To understand such real-world networks, we must solve fundamental mathematical and algorithmic research questions on *beyond-planar graphs*, which generalize the notion of planar graphs regarding topological constraints or forbidden edge crossing patterns.

This Dagstuhl Seminar investigated beyond-planar graphs, in particular, their combinatorial and topological structures (i.e., density, thickness, crossing pattern, chromatic number, queue number, and stack number), computational complexity and algorithmics for recognition, geometric representations (i.e., straight-line drawing, polyline drawing, intersection graphs), and their applications to real-world network visualization.

Compared to the previous two editions of the seminar, we focus more on aspects of combinatorics and geometry. Therefore, we included one new organizer and more participants from the corresponding fields. Thirty-two participants accepted the invitation to participate and arrived on Sunday afternoon.

On Monday morning, the program started with an introduction of all participants, followed by four invited talks to provide fundamental background knowledge on related research fields. We organized an open problems session on Monday afternoon and formed new working groups for research collaboration.

Many new problems related to combinatorics and geometry of beyond-planar graphs have been proposed. Specific open problems include questions regarding the combinatorial structures and topology (e.g.,  $k^+$ -real face graphs, beyond upward planar graphs, sparse universal geometric graphs, local-crossing-critical graphs), the geometric representations (e.g., constrained outer string graphs, rerouting curves on the surface), and applications.

Two progress report sessions were organized on Tuesday and Thursday afternoons to report progress and plans for future publications and follow-up meetings among researchers. From the participants' feedback, the seminar has initiated new research collaboration and led to new research ideas and directions.

Taking this opportunity, we thank Schloss Dagstuhl for providing an environment for fruitful research collaboration.

## 6.10 Safety Assurance for Autonomous Mobility

**Editors: Jyotirmoy Deshmukh, Bettina Könighofer, and Dejan Ničković**  
**Seminar No. 24071**

Date: February 11–16, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.95

© Creative Commons BY 4.0 license

© Dejan Ničković, Jyotirmoy Deshmukh, Bettina Könighofer, and Filip Cano

**Participants:** Houssam Abbas, Ezio Bartocci, Christof Budnik, Filip Cano, Mauricio Castillo-Effen, Chih-Hong Cheng, Patricia Derler, Jyotirmoy Deshmukh, Rayna Dimitrova, Marie Farrell, Michael Fisher, Martin Fränzle, Mahsa Ghasemi, Radu Grosu, Rong Gu, Ichiro Hasuo, Bardh Hoxha, Taylor T. Johnson, Panagiotis Katsaros, Bettina Könighofer, Thierry Lecomte, Lars Lindemann, Assaf Marron, Sayan Mitra, Selma Music, Dejan Nickovic, Necmiye Ozay, Giulia Pedrielli, Doron A. Peled, Stefan Pranger, Andoni Rodríguez, Darko Stern, Alessandro Zanardi, Dirk Ziegenbein



As autonomous mobility systems gain traction worldwide, ensuring their safety, robustness, and dependability has become a paramount concern for their implementation at scale. The Dagstuhl Seminar on “Safety Assurance for Autonomous Mobility” gathered experts from academia, and industry to address the critical challenges and opportunities posed by the rapid growth of autonomous technologies across various mobility domains, including automotive, aerospace, robotics, and railways. This seminar provided a much-needed platform for researchers and practitioners to exchange insights, collaborate on emerging ideas, and build a shared understanding of safety assurance in this rapidly evolving field.

### ■ Seminar Context and Structure

The seminar convened a diverse and multidisciplinary group of participants, each bringing specialized expertise in formal methods, software verification, embedded systems, and transportation safety. We believe that autonomous mobility is a field where progress cannot be just limited to academic research. The ideas and methods developed in theoretical settings build more robust applications, and the challenges faced in industrial settings guide theoretical research towards productive solutions. To reflect this drive, the group of participants was chosen to strike a balance between academia and industry, with many participants having experience in both domains. The seminar was structured around discussions in small working groups, different each day. Each group had a topic or problem to tackle, and the key challenges and state of the art solutions were shared at the end of each day to all participants of the seminar. One of the key objectives of this seminar was to bring together ideas and researchers from academic and industrial backgrounds. To this end, a sessions on Wednesday were focused on the current state of the practice being used in industrial applications, with each industrial partner sharing knowledge about their respective application fields.

### ■ Key Themes and Discussions

- Participants explored various formal methods and verification techniques designed to enhance the reliability of autonomous systems. Discussions highlighted the need to advance state-of-the-art formal verification approaches to accommodate the complexity of modern autonomous systems.
- The seminar also emphasized the importance of building resilient systems capable of functioning reliably in dynamic environments. Discussions tackled challenges related to the various modules (e.g. perception, motion planning, etc.) in autonomous mobile cyber-physical systems, considering how to incorporate robustness into the design phase and beyond.
- The specific challenges unique to each transportation sector were discussed, emphasizing tailored strategies for addressing safety assurance in automotive, aerospace, and railway systems. The cross-sectoral dialogue shed light on shared challenges and provided new perspectives that will inform future efforts.

### ■ Outcomes and Future Directions

The seminar generated a consensus on the urgent need for more research and collaboration across sectors. Participants emphasized the importance of combining expertise from different domains to address the interdisciplinary nature of safety assurance in autonomous mobility.

Moreover, the discussions underscored the potential for ongoing interdisciplinary seminars and follow-up workshops that would ensure continuous engagement among stakeholders. These future events would also provide venues for updating each other on progress, refining safety standards, and accelerating technological advancements in this field.

Overall, the seminar succeeded in creating a collaborative environment that not only identified existing challenges but also laid the groundwork for innovative solutions in safety assurance for autonomous mobility.

## 6.11 Triangulations in Geometry and Topology

**Editors: Maike Buchin, Jean Cardinal, Arnaud de Mesmay, and Jonathan Spreer**  
**Seminar No. 24072**

Date: February 11–16, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.120

© Creative Commons BY 4.0 license

© Maike Buchin, Jean Cardinal, Arnaud de Mesmay, and Jonathan Spreer



**Participants:** Mikkel Abrahamsen, Therese Biedl, Florestan Brunck, Kevin Buchin, Maike Buchin, Benjamin Burton, Jean Cardinal, Hsien-Chih Chang, éric Colin de Verdière, Arnaud de Mesmay, Fabrizio Frati, Alex He, Linda Kleist, Francis Lazarus, Maarten Löffler, Anna Lubiw, Clément Maria, Tim Ophelders, Lionel Pournin, Günter Rote, Saul Schleimer, Lena Schlipf, André Schulz, Eric Sedgwick, Rodrigo I. Silveira, Jonathan Spreer, Stephan Tillmann, Birgit Vogtenhuber, Zili Wang, Erin Moriarty Wolf Chambers, Alexander Wolff

This seminar was a followup to the Dagstuhl Seminars “Applications of Topology to the Analysis of 1-Dimensional Objects” (17072), “Computation in Low-Dimensional Geometry and Topology” (19352), and “Computation and Reconfiguration in Low-Dimensional Topological Spaces” (22062). The common idea behind all of these seminars is to bring together researchers from different communities (such as computational geometry, graph drawing, or geometric topology) with a shared interest in low-dimensional objects (e.g., curves, embedded graphs, knots, or surfaces). The goal of this approach is to foster collaborative work and synergies: The mathematical study of low-dimensional objects has a rich and old history, but research into their algorithmic and combinatorial properties and the underlying computational questions is still young. This makes for a vibrant research situation, giving strong opportunities for interdisciplinary work involving our multiple communities.

The focus of this Dagstuhl Seminar was placed on **triangulations**: partitions of the plane into triangles, or, more generally, of a space into simplices, which are required to meet face-to-face. Triangulations are typically constrained to use a given set of points as vertices and are fundamental tools in many applications such as computer graphics or geographic information systems. Alternatively, a triangulation can be defined on a topological space as a simplicial complex together with a homeomorphism from this simplicial complex to the space. These triangulations play an important role in the study of metrics on surfaces and their moduli space. The multiple facets of triangulations make them an essential and ubiquitous object of study in the combinatorial, algorithmic, geometric and topological properties of low-dimensional spaces, and thus they constitute a fertile ground for collaborations.

The seminar started with a quick introduction of all the participants, and four keynote talks on different aspects of triangulations.

- *Lionel Pournin* gave an overview of flip graphs and their combinatorial and algorithmic properties.
- *Saul Schleimer* surveyed the use of triangulations in 3-manifold topology and the numerous computational challenges that arise from them.
- *Linda Kleist* presented two snapshots on triangulations, first on the computational complexity of linearly embedding simplicial complexes, and then on some hamiltonicity properties of polytopes arising from flipping triangulations.
- *Mikkel Abrahamsen* explained new hardness proofs for algorithmic problems related to packing, covering and partitioning simple polygons with unit squares.

We refer to the abstracts later in this report for more details on these contributions.

These keynote introductory talks were followed with an extended open problem session where we gathered a large collection of open problems. Some of these were circulated in advance of the meeting, many of them were new. The remainder of the week was spent working in small groups actively trying to make progress on the most popular open problems. We made extensive use of the tool “Coauthor,” designed by Erik Demaine (MIT). This allowed for a very efficient recording of the progress made in the different groups. Regular progress reports allowed participants to easily switch between groups during the week, or to start new groups, leading to a very dynamic working environment.

We now quickly survey the different problems that have been worked on during this very productive week:

- **Computational complexity of problems in 3-manifold topology.** This group discussed computational complexity for important algorithmic problems from 3-manifold topology. Examples of problems include showing that a knot is ribbon, and testing 0-efficiency in a 3-manifold triangulation (lead: Eric Sedgwick).

- **Veering triangulations and the flip graph.** This group studied the effect on flip distances of surface triangulations if a veering structure in the associated layered triangulation is (lead: Saul Schleimer).
- **Flip distances.** This group investigated distances in the flip graphs of triangulated convex  $n$ -gons or annuli, both theoretically and algorithmically (lead: Jonathan Spreer).
- **Hardness for simple polygons.** The group considered NP-hard problems on polygons with holes, and how to show that these are also NP-hard on simple polygons. (lead: Lena Schlipf)
- **Catching balls on Curves.** This group considered the problem of characterising curves for which there exist a strategy to catch a ball from any initial configuration. (lead: Maarten Löffler)
- **Computing geodesic paths using edge flips.** This group investigated the FLIPOUT algorithm defined by Crane and Sharp to compute geodesics on intrinsic geometric triangulations and its possible variants. (lead: Hsien-Chih Chang)
- **Rendering a knot without self-intersections.** This group focused on using topological data analysis techniques to produce instructive 3D models of link diagrams (lead: Clément Maria)
- **Shortest cycle separating  $k$  objects from  $n - k$ .** This group investigated the complexity of, given a collection of  $n$  objects, computing the shortest cycle separating  $k$  objects from  $n - k$  objects, and how it behaves with respect to the parameter  $k$ . This was explored in the setting of the plane with obstacles and then planar graphs with obstacles, with an eye towards separating  $k$  handles in graphs embedded on surfaces. (lead: Éric Colin de Verdière)

In summary, this Dagstuhl Seminar provided a very fruitful research environment, allowing participants from very different backgrounds to work together on important open problems. Survey feedback from the participants highlighted how much the emphasis on intensive work in small groups was appreciated. In several of the working groups, significant progress was made, and the results are currently prepared to be submitted for publication. As in the previous meetings, the excellent quality of the Dagstuhl infrastructure and the impeccable support of the Dagstuhl staff provided a seamless experience for all the participants. We are hopeful that this successful experience will lead to follow-up Dagstuhl Seminars on related topics.

## 6.12 Computational Approaches to Strategy and Tactics in Sports

**Editors:** Ulf Brefeld, Jesse Davis, Laura de Jong, and Stephanie Kovalchik

**Seminar No. 24081**

Date: February 18–23, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.164

© Creative Commons BY 4.0 license

© Ulf Brefeld, Jesse Davis, Laura de Jong, and Stephanie Kovalchik



**Participants:** Gabriel Anzer, Felipe Arruda Moura, Pascal Bauer, Joris Bekkers, Luke Bornn, Timothy Chan, Jesse Davis, Laura de Jong, Uwe Dick, Max Goldsmith, Florentina Hettinga, Benjamin Holmes, Mamiko Kato, Matthias Kempe, Hyunsung Kim, Stephanie Kovalchik, Martin Lames, Daniel Link, Jim Little, Patrick Lucey, Jakub Michalczyk, Darren O'Shaughnessy, Sigrid Olthof, David Radke, Pegah Rahimian, Yannick Rudolph, Martin Rumo, Nathan Sandholtz, Raimund Seidel, Joshua Smith, Tim Swartz, Jan Van Haaren, Maaïke Van Roy, Christoph Weber, Hendrik Weber, Albrecht Zimmermann

The rapid growth in spatio-temporal data in sport over the past decade has generated numerous methodological developments from the statistical and machine learning communities. The richness of modern sports data is enabling sports researchers to analyze every action and decision during a competitive event in increasing detail. Two central topics that have emerged from this new phase of methodological research in sport are data-driven approaches to *strategy & tactics*. In a nutshell, *strategy & tactics* allow weaker teams or athletes to win over stronger ones. Therefore, they are one of the most interesting and challenging aspects in sports.

Although both terms describe similar aspects and are even sometimes used interchangeably, they range on different time scales. A *strategy* serves as an overarching umbrella to reach long-term goals. Hence, strategic decisions involve long-term training plans, signing players and coaches, as well as deciding on team formation, pacing, equipment, rotations, or playing philosophy. On a shorter time scale a match/race strategy is the plan made by coaches and athletes before the start of the match or race.

*Tactics*, on the other hand, is rather short-term. Tactics are the execution and adaptations to the planned strategy to have an edge over the opponent during the match or race. Tactics are therefore often broken down into building blocks or patterns that can be easily communicated to athletes. Note that communication is of utmost importance as tactics are invented by the coaching staff while their implementation is the task of the players/athletes. A tactical pattern may thus involve only subgroups of athletes, or subsections of a race and assign concrete tasks for predefined, context-sensitive situations.

The goal of this Dagstuhl Seminar was to bring together a diverse set of researchers from both academia and industry working on these topics. The seminar drew from people with various backgrounds in terms of area of specialization (Artificial Intelligence, Operations Research, Sport Science, Statistics),

role (Academic, Data Provider, Federation, Sports Club) and sport (Australia Rules Football, Baseball, Basketball, Darts, Ice Hockey, Soccer, Speed Skating, Tennis, Wheelchair Rugby).

The seminar was structured around three themes:

**Discovery** The goal of this theme was to discuss different methods that can automatically identify tactical and strategic patterns from spatio-temporal data. Examples were given for problems such as detecting formations, identifying commonly occurring sequences of actions (e.g., passing sequences), discovering player movement trajectories, and deciding where players should aim a tennis serve.

**Evaluation** This theme focused on the challenges and pitfalls associated with trying to evaluate the finding of computational approaches to identifying strategies and tactics. This theme focused on highlighting a number of methodological issues and describing ways to assess the validity of discoveries. There were a number of interesting examples given about how causal analysis could be used to evaluate the efficacy of certain tactics. Finally, the potential and risks for using large language models in sports were also discussed.

**Communication** This theme tackled the problem of how to communicate the findings from tactical studies to an interdisciplinary audience. The emphasis was on how to marry finding from the research literature to things that could be translated into practice. A key point that was made is that it is crucial to think about what types of information will be useful and actionable for practitioners.

The first three days of the seminar focused on one theme, which was introduced with a longer tutorial and then shorter presentations. The final full day of the seminar was open to all topics under the themes and there was a greater focus on presentations from early-career researchers in attendance. The seminar also featured two panels and (small) group discussions about five different topics.

## ■ Results

During the seminar, we identified and agreed upon the following action points aimed at trying to continue integrating the various different communities (Sports Science, Operations Research, Statistics, Artificial Intelligence) working on computational approaches to tactics in sports:

1. We will collect a list of venues where computational approaches to tactics in sports are often published. We will host this on the web: <https://dtai.cs.kuleuven.be/sports/venues/>
2. We will explore setting up a slack or discord channel to facilitate more continuous interaction and the ability to quickly get answers to questions. Joris Bekkers and Jan Van Haaren will take the lead on this point.
3. We have setup a document that contains the biographies, contact details, and topics of interest for all seminar participants that are willing to share their information. That will help people stay in touch.
4. We will strive to setup some basic tutorials that illustrate how to implement standard, concepts that reoccur across sports. For example, many team sports have variants of plus-minus, expected possession value metrics, and expected statistics such as expected goals (soccer, ice hockey) or expected rush yard gained (American Football).
5. We will continue to promote the mailing list for disseminating computational sports-related information (job ads, conference call for papers, etc.) and we will use this list to distribute the report on the seminar to reinforce our thanks to the attendees and excitement about the seminar's outcomes: [ml-ai-4sports@googlegroups.com](mailto:ml-ai-4sports@googlegroups.com)

## 6.13 AI for Social Good

**Editors: Claudia Clopath, Ruben De Winne, Mohammad Emtyiaz Khan, and Jacopo Margutti**  
**Seminar No. 24082**

Date: February 18–23, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.182

© Creative Commons BY 4.0 license

© Ruben De Winne



**Participants:** Asma Atamna, Annabelle Behnke, Siu Lun Chau, Claudia Clopath, Jorn Dallinga, Ruben De Winne, Michael Dhatemwa, Daphne Ezer, Frank Hutter, Roberto Interdonato, Mohammad Emtyiaz Khan, Isabell Klipper, Parvathy Krishnan, Derek Loots, Subhransu Maji, Jacopo Margutti, Marieke Meeske, Krikamol Muandet, N. N., Virginia Partridge, Julia Proskurnia, Lennart Purucker, Jake Robertson, Andrés Roure Cuzzoni, Tom Schaul, Jeremy Springman, Maïna Vergonjanne

AI and ML have made impressive progress in the last few years. Long-standing challenges like Go have fallen and the technology has entered daily use via the vision, speech or translation capabilities in billions of smartphones, and more recently via general uptake of software applications built on large language models. The pace of research progress shows no signs of slowing down, and demand for talent is unprecedented. But as part of a wider AI for Social Good trend, this seminar wanted to contribute to ensuring that the social good does not become an afterthought in the rapid AI and ML evolution, but that society benefits as a whole.

The five-day seminar brought together AI and ML researchers from various universities with representatives from NGOs pursuing various social good goals, such as providing legal aid, providing humanitarian assistance, advocating for gender justice, denouncing growing levels of inequality, and defeating poverty. On these topics, NGOs have rich domain knowledge, just like they have vast networks with (non-)governmental actors in developing countries. Mostly, NGOs have their finger on the pulse of the challenges that the world & especially its most vulnerable inhabitants are facing today, and will be facing tomorrow.

The objective of the seminar was to look at these challenges through an AI and ML lens, to explore if and how these technologies could help NGOs to address these challenges. The motivation was also that collaborations between AI and ML researchers and NGOs could benefit both sides: on the one hand, the new techniques can help with prediction, data analysis, modelling, or decision making. On the other hand, the NGOs' domains contain many non-standard conditions, like missing data, side-effects, or multiple competing objectives, all of which are fascinating research challenges in themselves. And of course, publication impact is substantially enhanced when a method has real-world impact.

The seminar facilitated the exploration of possible collaborations between AI and ML researchers and NGOs through a two-pronged approach. This approach combined high-level talks & discussions on the one hand with a hands-on hackathon on the other hand. High-level talks & discussions focused first on the central concepts and theories in AI and ML and in the NGOs' development work, before diving into specific issues such as generalizability, data pipelines, and explainability. These talks and discussions allowed all participants – in a very short timeframe – to reach a sufficient level of understanding of each other's work. This understanding was the basis to then start investigating jointly through a hackathon how AI and ML could help addressing the real-world challenges presented by the NGOs. At the start of the hackathon, an open marketplace-like setting allowed AI and ML researchers and NGOs to find the best match between technological supply and demand. When teams of researchers and NGOs were established, their initial objective was not to start coding, but to define objectives, assess scope and feasibility.

The intense exchanges during the hackathon allowed NGOs with a lower AI/ML maturity increased to increase understanding of the capabilities of AI/ML and define actions to effectively start working with AI/ML. NGOs that already had a more advanced understanding and use of AI/ML technology prior to the seminar, could take their AI maturity to the next level by trying out new ML approaches, designing and testing tailored ML models, or simply exploring new partnerships. Key to this success of the hackathon – and the seminar at large – was the presence of AI/ML experts whose respective fields of expertise could seamlessly be matched with the various needs of the various NGOs. This excellent group composition also facilitated a productive discussion about topics that cut across the AI for social good field, such as how to properly evaluate AI models that are used for good.

## 6.14 Reflections on Pandemic Visualization

**Editors:** Daniel Archambault, Fintan McGee, Simone Scheithauer, and Tatiana von Landesberger

**Seminar No. 24091**

Date: February 25 – March 1, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.191

© Creative Commons BY 4.0 license

© Daniel Archambault, Fintan McGee, Simone Scheithauer, and Tatiana von Landesberger

**Participants:** Peter Arndt, Steve Awodey, Andrej Bauer, Karim Johannes Becher, Olaf Beyersdorff, Marc Bezem, Ingo Blechschmidt, Ulrik Buchholtz, Gabriele Buriola, Felix Cherubini, Michel Coste, Laura Crosilla, Nicolas Daans, Dominique Duval, Martín H. Escardó, Giulio Fellin, Makoto Fujiwara, Hugo Herbelin, Matthias Hutzler, Hajime Ishihara, Ulrich Kohlenbach, Henri Lombardi, Maria Emilia Maietti, Julien Narboux, Sara Negri, Takako Nemoto, Stefan Neuwirth, Satoru Niki, Paige North, Eugenio Orlandelli, Edi Pavlovic, Iosif Petrakis, Elaine Pimentel, Michael Rathjen, Marie-Françoise Roy, Peter M. Schuster, Monika Seisenberger, Sana Stojanovic-Djurdjevic, Benno van den Berg, Steven J. Vickers



During the recent SARS-CoV-2 pandemic, visualizations were omnipresent, playing a central role in communicating with the public, drawing from multiple data sources, and serving diverse goals. In a short period, public health messages spread globally. In this Dagstuhl Seminar, we brought together 37 experts in visualization, mathematics, modeling, public health, infectious diseases, and psychology from Europe, Asia, Australia, and North America to summarize and discuss their personal insights gained over the three years of the pandemic.

Due to the heterogeneity of expertise and different tasks performed by each individual during the pandemic, and despite intensive collaboration and contact between experts in both fields over the past three years, it was felt necessary to establish some common ground for the work and concepts of both medical and visualization experts. The first two days of the seminar included short talks on public health, infectious diseases, modeling, big data, and an overview of visualizations that participants used or appreciated during the SARS-CoV-2 pandemic. These introductory talks were specially designed to foster discussions and personal exchange, allowing significant time for in-depth conversations. After engaging discussions during and after the

sessions, several key topics emerged across the medical and visualization fields that require deeper reflection. The organizers then clustered these relevant topics into five overarching areas of interest: the use of dashboards during the pandemic, communication to the public, preparedness, data visualization methodology in emergency responses and users tasks and medium.

Over the last two days of the seminar, breakout sessions of six participants were created to work on these key issues. The fruitful discussions in the breakout sessions had as their first output a presentation by each group summarizing their discussions. These discussions will be extended into a second output in the form of several publications that will appear in a Computer Graphics and Applications issue later this year which are now in preparation. Important connections were made between experts in health-care-related disciplines and visualization specialists, addressing a significant need. As a result, one of the organizers was invited to participate in a symposium on automated surveillance of bloodstream infections in Germany, which was promptly accepted. We look forward to hosting more interdisciplinary seminars, which are extremely rewarding and valuable.

## 6.15 Applied and Combinatorial Topology

**Editors:** Paweł Dłotko, Dmitry Feichtner-Kozlov, Anastasios Stefanou, and Yusu Wang  
**Seminar No. 24092**

Date: February 25 – March 1, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.2.206

© Creative Commons BY 4.0 license

© Paweł Dłotko, Dmitry Feichtner-Kozlov, Anastasios Stefanou, and Yusu Wang



**Participants:** Peter Arndt, Steve Awodey, Andrej Bauer, Karim Johannes Becher, Olaf Beyersdorff, Marc Bezem, Ingo Blechschmidt, Ulrik Buchholtz, Gabriele Buriola, Felix Cherubini, Michel Coste, Laura Crosilla, Nicolas Daans, Dominique Duval, Martín H. Escardó, Giulio Fellin, Makoto Fujiwara, Hugo Herbelin, Matthias Hutzler, Hajime Ishihara, Ulrich Kohlenbach, Henri Lombardi, Maria Emilia Maietti, Julien Narboux, Sara Negri, Takako Nemoto, Stefan Neuwirth, Satoru Niki, Paige North, Eugenio Orlandelli, Edi Pavlovic, Iosif Petrakis, Elaine Pimentel, Michael Rathjen, Marie-Françoise Roy, Peter M. Schuster, Monika Seisenberger, Sana Stojanovic-Djurdjovic, Benno van den Berg, Steven J. Vickers

The Dagstuhl Seminar titled “Applied and Combinatorial Topology” brought together researchers in mathematics and computer science to engage in active discussions and exchange of ideas on theoretical, computational, and practical aspects of applied and combinatorial topology. The seminar has led to further the connections between the Discrete Morse Theory, Computational Topology, and Statistics communities and identification of open problems that can be addressed together.

**Context** Applied Topology is a new and rapidly increasing research field within applied mathematics. Its main focus is to utilize topological methods to solve applied problems. The common emphasis of the methods in Applied Topology is on the computational aspect. Application areas include: data analysis, computational biology, network analysis, graph visualization and reconstruction, feature selection, and many more.

**Goals** The Dagstuhl Seminar 24092 (February 25–March 1, 2024) brought together three research communities, namely researchers in discrete Morse theory (DMT), computational topology, and statistics. The aim was to facilitate collaborations that could strengthen the existing interactions of the fields (e.g. Reeb graphs, Mappers and discrete vector fields) and collaborations that may lead to the development of new descriptors of data (e.g. persistence invariants, magnitude functions, etc.), which in turn have the potential to be inputted into statistical methodologies and to provide their efficient implementations.

**Topics** We chose three research topics for which the respective communities will benefit from a knowledge exchange and mutual discussion.

**Discrete Morse theory.** The research field of discrete Morse theory, developed by R. Forman, is a discrete counterpart of a continuous Morse theory. It has recently found many practical

applications both within mathematics (e.g. configuration spaces, and homology computation) as well as outside mathematics, such as computer science (e.g. denoising and mesh compression). The target of discrete Morse theory is to construct a discrete vector field that either simplifies the data at hand, without losing its important features, or to introduce discrete dynamics on the data. The resulting dynamics can be further analyzed to extract certain interesting sets, for instance invariant sets. In both cases, the aim is to simplify the data or to find the important regions of the data. One example of use of Morse theory in data analysis is the Reeb graph of a Morse function. A discrete adaptation of Reeb graphs which uses ideas of partial clustering is known as Mapper and has seen a great deal of success in data analysis. Another example is an adaptation of discrete Morse theory to computations of persistent homology in topological data analysis: the machinery of discrete Morse theory can be used to help reduce the complexity of the evolving topology in the filtrations of datasets. Moreover, it was recently shown that discrete Morse theory can be utilized for simplification and complexity reduction also in the multiparameter persistence setting. Furthermore, discrete Morse theory has been studied in conjunction with persistent homology theory and has found interesting applications, such as reconstruction of grayscale digital images and reconstruction of graphs, see for instance 2D road reconstruction and 3D neuron reconstruction. Another powerful application of Discrete Morse theory is in distributed computing. In both cases, the discrete Morse theory is used to simplify data at hand, and recover their invariants. We believe that using this machinery, one can do even more. We would like to take the paradigm of discrete Morse theory further and directly try to recover certain invariants from the constructed discrete vector fields.

**Computational Topology.** The research field of combinatorial topology originated from the study of topological invariants derived from combinatorial decompositions of spaces

(cf. simplicial approximation theorem), known as simplicial complexes. One of the main examples of such invariants are the Betti numbers. Combinatorial topology was later named algebraic topology due to the switch of focus of the field on its algebraic aspects (as homology groups), which is attributed to Emmy Noether. In the research area of computational topology (also known as topological data analysis), we are interested in studying a single parameter filtration of complexes associated with a data set (viewed as finite metric space) such as the Vietoris-Rips filtration, or multiparameter filtrations of complexes associated to datasets, such as the function-Rips bifiltration and the multicover bifiltration. Those structures can be simplified with the tools of discrete Morse theory. A homology functor is then being applied to those filtrations resulting in a single or multiparameter persistence module. Single-parameter persistence modules are visualized by their persistence diagrams. A well-known invariant of multiparameter persistence modules is the rank invariant which captures important persistence information about multifiltrations of datasets. Recently, there have been some refinements of the rank invariant and also a generalization of the notion of persistence diagram (induced by the rank invariant). Developing algorithms for the efficient computation of multiparameter persistence modules and their rank invariants, is one of the big challenges of computational topology and topological data analysis (TDA).

**Statistics in Topological Data Analysis.** Persistence invariants such as the persistence diagram are equipped with a family of metrics, e.g. the  $\ell^p$ -Wasserstein distances and the bottleneck distance. To make these signatures applicable, one must interface them with standard statistical methods. This has already been done e.g. when developing statistics on persistent diagrams, or other signatures such as persistence landscapes. However, much remains unknown in the case of limits of persistence diagrams when the number of points goes to infinity. A good example of a successful synergy between statistics and combinatorial topology is a process of vectorization of persistence diagrams. This process allowed the community to build multiple applications of persistent homology into many branches of science and engineering. We believe that, if new invariants originated from discrete Morse theory and combinatorial topology are introduced, such as the recently introduced Mapper graph of datasets, a work needs to be done, to incorporate them into existing statistical pipelines, hypothesis testing methods and similar. Moreover, a vectorization method for Mapper graphs needs to be established and their limit behavior (when e.g. the number of points goes to infinity) need to be studied. Also an application in standard statistics will be further explored; It is widely known that one should not rely on summary statistics, but always attempt to visualize the data. However, oftentimes the data are very high dimensional. In this case, Mapper type algorithms may serve as a surrogate of a scatter plot in visualization by providing a graph-based summary of the data. Our aim will be to explore this connection and look for ways of inputting Mappers into standard statistical pipelines, e.g. including concepts of averages and central limit theorems. We will also explore the connections between Mapper and other combinatorial topology concepts via, for instance, discrete Morse theory.

**Participants, Schedule, and Organization** The attendees were strongly encouraged to prepare talks that will include open problems and new research directions. The program for the week consisted of talks of different lengths, open problem sessions, breakout sessions, and summary sessions with the participants. On Monday, we started with an 1-hour session where the participants introduced themselves, and then we had 6 introductory talks, two on Discrete Morse theory, two on computational topology and two on Statistics in TDA. Then, we had an open problem session where participants identified certain open problems and directions for research for the breakout sessions.

Participants chose one or more from the following proposed topics for breakout sessions:

1. Can we compute representatives of generators of persistent homology in less than cubic time? (proposed by Tamal Dey)
2. Optimal Discrete Morse function given a partial matching (proposed by Yusu Wang)
3. Topological information, i.e. “how much topological information remains when going from one to two dimensional filtrations (or from Reeb graphs to Reeb spaces)” (proposed by Bei Wang Phillips)
4. Manifold reconstruction guarantees (proposed by Ulrich Bauer)
5. Algorithmic questions on (multiparameter) persistence (proposed by Fabian Lenzen)
6. Can TDA detect planted cliques? (proposed by Bastian Rieck)
7. Monotonicity of magnitude functions of Euclidean metric spaces (proposed by Sara Kalisnik)
8. General applied topology (proposed by Dmitry Feichtner-Kozlov).

Tuesday to Thursday in the morning we had the lecture talks and we organized breakout sessions on Tuesday and Thursday afternoon. We reserved three rooms for the breakout sessions that ran in parallel, the main seminar room for topics (1)–(5), another room for topics (6)–(7), and a small room for topic (8). On Wednesday afternoon we organized some groups for hiking near Schloss Dagstuhl. Representatives from the working groups summarized the discussions during their breakout sessions and presented it to all participants on Thursday evening and Friday morning.

**Results and Reflection** The seminar successfully facilitated a rich exchange of ideas and expertise among participants. The varied program, including talks, open problem sessions, breakout discussions, and outdoor activities, created an environment conducive to collaborative exploration. Attendees expressed satisfaction with the content and structure of the seminar, indicating a strong interest in future editions. During the breakout sessions, it was encouraging to note that some participants reported preliminary results related to the open problems presented. These early findings sparked lively discussions and provided valuable insights into potential directions for further research. The seminar served as a platform not only for sharing existing knowledge but also for generating new ideas and approaches.

## 6.16 Robust Query Processing in the Cloud

**Editors: Goetz Graefe, Allison Lee, and Caetano Sauer**  
**Seminar No. 24101**

Date: March 3–8, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.1

© Creative Commons BY 4.0 license

© Goetz Graefe, Allison Lee, and Caetano Sauer



**Participants:** Angelos Christos Anadiotis, Manos Athanassoulis, Carsten Binnig, Thomas Bodner, Matthias Böhm, Peter A. Boncz, Nicolas Bruno, Yannis Chronis, Periklis Chrysogelos, John Cieslewicz, Sudipto Das, Thanh Do, Kira Isabel Duwe, Jan Finis, Campbell Fraser, Goetz Graefe, Stefan Halfpap, Alfons Kemper, Kyoungmin Kim, Andrew Lamb, Allison Lee, Viktor Leis, Lucas Lersch, Boaz Leskes, Thomas Neumann, Anisoara Nica, Danica Porobic, Daniel Ritter, Kai-Uwe Sattler, Caetano Sauer, Bernhard Seeger, Knut Stolze, Pinar Tözün, Nga Tran, Immanuel Trummer, Juliane Waack, Marcin Zukowski

The Dagstuhl Seminar on “Robust Query Processing in the Cloud” (24101) assembled researchers from industry and academia for the fifth time to discuss robustness issues in database query performance, this time with a focus on Cloud Computing. The seminar gathered researchers around the world working on indexing, storage, plan generation and plan execution in database query processing, and in cloud-based massively parallel systems with the purpose to address the open research challenges with respect to the robustness of database management systems. Delivering robust query performance is well known to be a difficult problem for database management systems. All experienced DBAs and database users are familiar with sudden disruptions in data centers due to poor performance of queries that have performed perfectly well in the past. The goal of the seminar was to discuss the current state-of-the-art, to identify specific research opportunities in order to improve the state-of-affairs in query processing, and to develop new approaches or even solutions for these opportunities, building upon successes of the past Dagstuhl Seminars [1–6]. The organizers (Goetz Graefe, Allison Lee, and

Caetano Sauer) this time attempted to have a focused subset of topics that the participants discussed and analyzed in more depth. From the proposed topics on algorithm choices, join sequences, storage architectures, database utilities, modern storage hardware, cloud database economics, and benchmarking for robust query processing, the participants formed five work groups: i) robustness benchmarking, ii) economics of query processing in the cloud, iii) storage architectures, iv) out-of-memory query operators, and v) indexing for data warehousing. Upon choosing the topics of interest, the organizers then guided the participants to approach the topic through a set of steps: by first considering related work in the area; then introducing metrics and tests that will be used for testing the validity and robustness of the solution; after metrics, the focus was on proposing specific mechanisms for the proposed approaches; and finally the last step focused on the implementation policies. At the end of the week, each group presented their progress with the hope to continue their work towards a research publication. The reports of work groups are presented in the full report.

### References

- 1 Peter A. Boncz, Yannis Chronis, Jan Finis, Stefan Halfpap, Viktor Leis, Thomas Neumann, Anisoara Nica, Caetano Sauer, Knut Stolze, and Marcin Zukowski. SPA: economical and workload-driven indexing for data analytics in the cloud. In *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*, pages 3740–3746. IEEE, 2023.
- 2 Renata Borovica-Gajic, Stratos Idreos, Anastasia Ailamaki, Marcin Zukowski, and Campbell Fraser. Smooth scan: Statistics-oblivious access paths. In Johannes Gehrke, Wolfgang Lehner, Kyuseok Shim, Sang Kyun Cha, and Guy M. Lohman, editors, *ICDE*, pages 315–326. IEEE Computer Society, 2015.
- 3 Renata Borovica-Gajic, Stratos Idreos, Anastasia Ailamaki, Marcin Zukowski, and Campbell Fraser. Smooth scan: robust access path selection without cardinality estimation. *VLDB J.*, 27(4):521–545, 2018.
- 4 David Justen, Daniel Ritter, Campbell Fraser, Andrew Lamb, Nga Tran, Allison Lee, Thomas Bodner, Mhd Yamen Haddad, Steffen Zeuch, Volker Markl, and Matthias Boehm. POLAR: adaptive and non-invasive

- join order selection via plans of least resistance. *Proc. VLDB Endow.*, 17(6):1350–1363, 2024.
- 5 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In Goetz Graefe and Kenneth Salem, editors, *DBTest*, page 7. ACM, 2011.
- 6 Lukas Vogel, Daniel Ritter, Danica Porobic, Pinar Tözün, Tianzheng Wang, and Alberto Lerner. Data pipes: Declarative control over data movement. In *13th Conference on Innovative Data Systems Research, CIDR 2023, Amsterdam, The Netherlands, January 8-11, 2023*. [www.cidrdb.org](http://www.cidrdb.org), 2023.

## 6.17 Shapes in Graph Data: Theory and Implementation

**Editors:** Shqiponja Ahmetaj, Slawomir Staworko, and Jan Van den Bussche  
**Seminar No. 24102**

Date: March 3–8, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.9

© Creative Commons BY 4.0 license

© Shqiponja Ahmetaj, Slawomir Staworko, and Jan Van den Bussche



**Participants:** Shqiponja Ahmetaj, Iovka Boneva, Angela Bonifati, Anastasia Dimou, Stefania Dumbrava, Nicolas Ferranti, George Fletcher, Benoit Groz, Jan Hidders, Katja Hose, Maxime Jakubowski, George Konstantinidis, José Emilio Labra Gayo, Aurélien Lemay, Leonid Libkin, Wim Martens, Fabio Mogavero, Filip Murlak, Cem Okumus, Nina Pardal, Liat Peterfreund, Axel Polleres, Ognjen Savkovic, Mantas Simkus, Slawomir Staworko, Katherine Thornton, Jan Van den Bussche, Maria-Esther Vidal, Hannes Voigt, Piotr Wiecek

**Research Area and Goals of the Seminar** One of the main reasons for the success of graph databases is that they do not require an elaborate database schema, with accompanying integrity constraints, to be set up in advance. In these classical applications, constraints and schemas are mainly *descriptive*, having as purpose to support the mental map from the real world to the data to be managed in the database. However, the emergence of graph databases is accompanied by a paradigm shift towards new applications where schemas and constraints are used for a *prescriptive* purpose. Here, the goal is to establish a contract between the database and its users, which provides guarantees on the structure and form of data provided. This shift has led to the development of a new class of formalisms based on the notion of *shapes*. Shapes are constraints on nodes in the graph that impose or forbid structural patterns (involving paths, edges, labels, and constant values). Naturally, then, a novel, prescriptive notion of schema emerges, consisting of a set of shapes, together with a targeting mechanism that specifies which nodes should satisfy which shapes. In the world of RDF graphs, two main shape-based formalisms have been proposed: *SHACL* (Shapes Constraint Language), standardized by the W3C, and *ShEx* (Shape Expression schemas). In the world of property graphs (PGs), different systems have their own data definition languages, such as Cypher or GSQL. Moreover, there are recent formal approaches to define schemas for property graphs such as PG-Schema and PG-Keys. The main aim of the Dagstuhl Seminar was to bring together active researchers, both from academia and industry, to report on the most recent results, to discuss the many open problems and research directions that arise from shapes, constraints, and schemas for graph databases, and to initiate new research.

**Organization and Outcomes** The organisers created a schedule based on the entries from a Google document set up before the seminar, inviting participants to add talks, demos, and

research topics. The seminar began with a round of introductions, where participants also asked questions they wanted to be answered during the seminar. The final schedule included 18 contributed talks and 6 short presentations on potential research and discussion topics.

As a major result from the seminar, four working groups were formed on the topics:

1. *What is used in practice for graph data abstractions? What is needed in practice for graph data abstractions?* The group formation was inspired by related questions posed by many participants during the opening introductory round on the first day of the seminar. Several research challenges were discussed and addressing them will call for opening new human-centered research lines in the data management community and beyond.
2. *Repairs and explanations in knowledge graph data management systems in the presence of shape constraints.* The group discussed the problem of assessing and managing data quality in knowledge graphs (KGs). This is a long-standing issue that attracts significant attention both in industry and academia. The new proposals on schemas and shape languages for KGs have introduced new challenges, which involve new methods to verify their validity, to deal with inconsistency, and repair the inconsistent data.
3. *Relating 6NF (Sixth Normal Form) and PG-Schema.* In this working group, two main questions were discussed: (1) Can we show in a systematic manner how schemas for property graphs, as expressed in the proposals of PG-Schema and PG-Keys, can be represented relationally, obtaining highly decomposed (6NF) schemas with key constraints and inclusion constraints such as foreign keys? (2) Can the intent of a graph database application be formalized in a suitable variant of EER (extended Entity-Relationship) diagrams?

4. *Convergence of graph data models and schemas.* The goal of the group was to understand the commonalities and differences between RDF and LPG (labelled property graphs), and their corresponding schema languages, ShEx and SHACL for RDF, and PG-Schema for LPG. The aim is to identify a common core (a small but useful common sublanguage, easily expressible in all three formalisms) and a common superlanguage (a language that captures all three formalisms, yet remains manageable).

The organisers regard the seminar as a very successful scientific event. Members of each working group expressed a clear commitment to staying connected to further investigate these topics. The first two groups specify a vision paper as a specific goal and the result of the group's future efforts and the second two groups aim to produce research papers.

The organisers are grateful to the Scientific Directorate and to the staff for supporting in making this seminar possible.

## 6.18 Logics for Dependence and Independence: Expressivity and Complexity

**Editors: Juha Kontinen, Jonni Virtema, Heribert Vollmer, and Fan Yang**  
**Seminar No. 24111**

Date: March 10–15, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.31

© Creative Commons BY 4.0 license

© Heribert Vollmer, Juha Kontinen, Jonni Virtema, and Fan Yang



**Participants:** Erika, Maria Aloni, Aleksii Ilari Anttila, Christel Baier, Fausto Barbero, Timon Barlag, Dario Della Monica, Arnaud Durand, Fredrik Engström, Hadar Frenkel, Nicolas Fröhlich, Pietro Galliani, Erich Grädel, Jens Gutsfeld, Matilda Häggblom, Miika Hannula, Peter Harremoës, Lauri Hella, Minna Eveliina Hirvonen, Batya Kenig, Søren Brinck Knudstorp, Phokion G. Kolaitis, Juha Kontinen, Cheuk Ting Li, Martin Lück, Yasir Mahmood, Alessio Mansutti, Arne Meier, Till Miltzow, Christoph Ohrem, Ana Oliveira da Costa, Martin Otto, Nina Pardal, Max Sandström, Milan Studený, Marius Tritschler, Jouko Väänänen, Jonni Virtema, Heribert Vollmer, Jef Wijsen, Fan Yang, Martin Zimmermann

Dependence and independence are interdisciplinary notions that are pervasive in many areas of science. They appear in domains such as mathematics, computer science, statistics, quantum physics, and game theory. The systematic development of logical and semantical structures for these notions via the logics of dependence and independence has exposed surprising connections between these areas.

Logics for dependence and independence are new tools for modeling dependencies and interaction in dynamical scenarios. Reflecting this, these logics often have higher expressive power and complexity than classical logics used for these purposes previously. During the past decade, pioneering results on logics for dependence and independence has been disseminated in a spectrum of respected international conferences such as LICS, MFCS, JELIA, LPAR, CSL, AiML, and FSTTCS, and in top journals in the areas of logic and theoretical computer science. Although significant progress has been made in understanding the computational side of these novel logics (see Section 2 in the full report for some examples) still many central questions remain unsolved so far. In addition to addressing the open questions, the seminar also aims at boosting the exchange of ideas and techniques between team-based logics and the application areas.

The complexity and expressivity aspects of logics in propositional, modal and first-order team semantics have been studied extensively during the past decade. Recently, the complexity theoretic focus has turned to the (parameterized) complexity of logically defined counting and enumeration problems as well as algebraic complexity of probabilistic and real-valued logics. Furthermore, the expressivity and complexity of the novel temporal team logics are also not yet well understood.

**Logics for real valued data and probabilistic reasoning** Algorithmically, first-order dependence and independence logic correspond exactly to the complexity class NP and to the

existential fragment of second-order logic (ESO) while inclusion logic corresponds to the complexity class P over ordered finite structures. Recent discoveries on the connections between so-called probabilistic independence logic and logics on real valued data have revealed similar fundamental connections to a computation paradigm that uses real numbers as primitive entities (so-called BSS paradigm). These probabilistic logics have fascinating connections to the area of information theory via the notion of entropy, which can be adopted as a dependency in the probabilistic team semantics framework.

**Applications to Hyperproperties and Formal Verification** An emerging area of applications for team semantics is the area of Hyperproperties. In the field of formal verification an execution of a system is modeled by a trace depicting the evolution of the system over discrete time. Traceproperties, ubiquitous in formal verification, are properties of systems that boil down to verifying that each trace of the system satisfies that property. Hyperproperties on the other hand are properties of systems that cannot be reduced to checking properties of individual execution traces of the system in isolation, but are instead properties of sets of traces. These properties are of vital importance in applications concerning security and information flow. A canonical example here is bounded termination; one cannot check whether there exists a uniform time bound for some action by checking computation traces in isolation. Other examples include security policies such as *non-interference* and *secure information flow*.

**Applications to Formal Linguistics** Team semantics was also proven to be a fruitful tool for formal linguistics, especially for *inquisitive semantics* and the study of *free choice inferences*. Inquisitive semantics is a unified formal framework for analyzing both statements and questions in natural language.

It is known that inquisitive logic essentially adopts team semantics and can thus be viewed as a variant of propositional dependence logic. This connection has already sparked a significant amount of interest and new research at the interface of the two fields. On a different line of research, recently a bilateral modal logic based on team semantics, called BSML, was developed to model free choice inferences in natural language, where an atom NE studied in the context of propositional team logics plays a central role. Very recent works have studied the logical properties of BSML, and promising broader applications of the team semantics method along this line are yet to be explored.

## ■ Organization of the Seminar and Activities

The seminar brought together 42 researchers from mathematical logic, natural language semantics, and theoretical computer science. The participants consisted of both senior and junior researchers, including a number of postdoctoral researchers and advanced graduate students.

Participants were invited to present their work and to communicate state-of-the-art advances. Over the five days of the seminar, 29 talks of various lengths took place. Introductory and tutorial talks of 60 minutes were scheduled prior to the seminar. The remaining slots were filled with shorter talks, mostly scheduled after the seminar commenced. Furthermore the seminar included an open problem session and a concluding perspectives address.

The tutorial talks took place in the beginning of the week in order to establish a common background for the different communities that came together for the seminar. The presenters and topics were:

- Jonni Virtema: Introduction to Team Semantics
- Erika Ábrahám: (Probabilistic) Hyperproperties
- Maria Aloni: Logic and Language: Linguistic Applications of Team Semantics
- Till Miltzow: Existential Theory of the Reals
- Cheuk Ting Li: The Undecidability of Probabilistic Conditional Independence Implication

In addition, the seminar consisted of 24 shorter contributed talks, addressing various topics concerning expressibility, complexity and applications of team-based logics.

A one hour long open problem session was held on Wednesday morning, just before the hike (“Open Problem Walk”). It was moderated by Juha Kontinen. The session was announced

already on Monday morning to give participants the opportunity to register for the session. Besides a couple of shorter contributions on decidability of Team-LTL (by Martin Zimmerman), expressivity of different forms of implications when added to inclusion (predicate) logic (by Jouko Väänänen), and expressivity of propositional independence logic (by Fan Yang), the session consisted of three longer introductions of the following open problems:

- Is PosSLP, the question if a given straight-line program (over the integers with operations of addition, multiplication and subtraction) computes a positive number, solvable in polynomial time? It is conjectured that NP with an oracle to PosSLP equals the complexity class  $\exists\mathbb{R}$ . (Till Miltzow)
- Are all probabilistic conditional independence implications derivable from information inequalities? (Milan Studený)
- Identify tractable fragment for model checking for dependence logic, that is, fragments with an effective syntax that are “natural” and “useful” in the sense that they can express interesting computational problems in a relatively straightforward way, have strictly higher expressive power than first-order logic FO, and have a polynomial-time model checking in data complexity. (Phokion Kolaitis)

The participants were asked to contribute more open problems to a collection in form of an Overleaf project.

The seminar ended with a perspectives address given by Jouko Väänänen just before Friday lunch.

## ■ Concluding Remarks

The seminar achieved its aim of bringing together researchers from various related communities to share state-of-the-art research. Considerable exchange took place between researchers in the application areas of hyperproperties and formal semantics and those working more theoretically on complexity and expressivity questions of team-based logics. The organizers left ample time for interaction outside of this schedule of talks and, as a result, many fruitful discussions between participants took place throughout the afternoons and evenings.

The organizers regard the seminar as a significant success. Bringing together researchers from different areas fostered valuable interactions and led to fruitful discussions. Feedback from the participants was very positive as well.

Finally, the organizers wish to express their gratitude to the Scientific Directorate of the Center for its support of this Dagstuhl Seminar.

## 6.19 EU Cyber Resilience Act: Socio-Technical and Research Challenges

**Editors:** Mila Dalla Preda, Serge Egelman, Anna Maria Mandalari, Volker Stocker, Juan Tapiador, and Narseo Vallina-Rodriguez  
**Seminar No.** 24112

Date: March 10–13, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.52

© Creative Commons BY 4.0 license

© Mila Dalla Preda, Serge Egelman, Anna Maria Mandalari, and Narseo Vallina-Rodriguez



**Participants:** Rainer Böhme, Mila Dalla Preda, Daniel J. Dubois, Carolyn Egelman, Serge Egelman, Hamed Haddadi, Christin Hartung-Kümmerling, François Hublet, Martina Lindorfer, Anna Maria Mandalari, Federica Maria Francesca Paci, Simon Parkin, Sergio Pastrana, Joel Reardon, Anna Schwendicke, Ben Stock, Volker Stocker, Guillermo Suárez-Tangil, Juan Tapiador, Vincent Toubiana, Narseo Vallina-Rodriguez

### ■ Introduction and Motivation

The increasing number of cyberattacks affecting digital products has caused significant security and financial costs to societies. For example, the Mirai attack in 2016 compromised millions of Internet of Things (IoT) devices by exploiting default usernames and passwords, turning them into a botnet army that launched a massive Distributed Denial of Service (DDoS) attack. This attack significantly impacted critical Internet services, causing major outages and disruptions on platforms like Twitter and Netflix [1].

The European Commission has proposed in 2022 the EU Cyber Resilience Act (CRA) to define the legislative framework of essential cybersecurity requirements that product manufacturers must meet when placing any product with digital elements on the internal market, while empowering users to make better security-aware decisions when purchasing and deploying digital products. Following its adoption in 2024, manufacturers will have two years to comply with the new rule, with specific deadlines for different types of products.

The roadmap for CRA adoption follows a multi-phased approach, focusing on high-risk products first and progressively expanding to cover a broader range of digital products over the next few years, aiming to ensure robust cybersecurity standards across the EU. Specifically, during the first year, the focus will be on raising awareness among stakeholders and providing guidance on compliance requirements. The European Commission and national authorities will offer support and resources to help manufacturers understand the new obligations. Then, during the second year, manufacturers and developers will need to ensure that their products meet CRA requirements. This includes implementing necessary security measures, conducting risk assessments, and updating product documentation.

In this scenario, device and software analysis methods – from formal methods to black-box testing – are essential for facilitating compliance at different stages of the product life-cycle, but also

for self-attestation and independent verification and certification. However, the rapid evolution and increasing complexity of new technologies and other socio-technical factors such as developers' awareness and incentives for compliance may add further challenges and barriers to adoption.

On the one hand, it is essential to understand whether regulatory requirements are realistic, unambiguous, and whether they are partially misaligned with technology trends, manufacturers' incentives and goals, and with users' privacy and security awareness. For example, research evidence has shown that many developers do not fully comply with the General Data Protection Regulation (GDPR) and the USA Children Online Privacy Protection Act (COPPA) requirements due to their dependency on obscure third-party components for development support and advertising, economic incentives, poor software engineering habits, or even a lack of awareness about the regulations' existence and scope (and hence their compliance obligations). On the other hand, we need to assess to which extent existing device and software analysis methods are fit for aiding developers and manufacturers in assessing compliance, but also for independent certification by third-parties and regulatory enforcement. Yet, current software and device analysis techniques (e.g., black-box testing) often over-simplify the complexity of digital products and present various scalability and coverage limitations that prevent them from reliably auditing and testing whether observed software properties in digital products comply with regulatory requirements.

This Dagstuhl Seminar united a multidisciplinary group of tech and legal academics, industry actors, and policy experts to share their knowledge and experience to collaboratively explore the complex landscape of research and socio-technical challenges for the adoption and enforcement of the CRA. These challenges arise from developer practices and incentives, user awareness, and the feasibility of existing software analysis methods for certification and enforcement.

## ■ Seminar Structure

The seminar had a dynamic structure during the 3 days, combining dedicated presentations, panels, and multi-disciplinary working groups to encourage active participation and dialogue between different communities and stakeholders. Arriving on Sunday and starting with a welcome dinner at Schloss Dagstuhl. The three-day seminar activities were structured as follows:

■ **Day 1.** The first morning was dedicated to participant introductions, setting common ground on seminar objectives through short elevator pitches by participants, followed by two seminar-like talks and guided discussions. This engaging round of introductions provided a comprehensive overview of the diverse knowledge and skills present in the room, setting the scene for collaborative and constructive discussions. Following these introductions, the seminar continued with an introductory talk by the organizers, a key presentation by Christin Hartung-Kümmerling and Anna Schwendicke from the BSI on the fundamentals, goals, and roadmap of the CRA, and a talk by Vicent Toubina (CNIL) on their experiences with

GDPR implementation and enforcement. Following these, participants engaged in open discussions to identify sub-problems of interest. At the end of the first day, participants formed multidisciplinary discussion groups to summarize seminar outputs and a brainstorm session for identifying three key topics for further discussion: (i) Understanding and Aiding the Developer Ecosystem; (ii) Standardization Efforts; and (iii) Tools for Regulatory Enforcement.

- **Day 2.** The second day continued with the interactive group discussions, finalizing with a final all-hands group to consolidate the outputs of the discussions. The day ended with a social activity involving a guided visit to the Völklingen Ironworks, and a dinner in Saarbrücken.
- **Day 3.** The final day involved several all-hands sessions to identify the main outcomes of the seminar, and research challenges for easing CRA adoption and compliance, ensuring continued progress beyond the seminar.

The full seminar agenda is available on the seminar website.

## ■ References

- 1 Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. Understanding the mirai botnet. In *26th USENIX security symposium (USENIX Security 17)*, pages 1093–1110, 2017.
- 2 Jukka Ruohonen and Kalle Hjerppe. The gdpr enforcement fines at glance. *Information Systems*, 106:101876, 2022.
- 3 Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do cookie banners respect my choice?: Measuring legal compliance of banners from iab europe’s transparency and consent framework. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 791–809. IEEE, 2020.
- 4 Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, Serge Egelman, et al. “won’t somebody think of the children?” examining coppa compliance at scale. In *The 18th Privacy Enhancing Technologies Symposium (PETS 2018)*, 2018.
- 5 Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies*, 2016.
- 6 Noura Alomar and Serge Egelman. Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps. *Proceedings on Privacy Enhancing Technologies*, 2022.
- 7 Michael Backes, Sven Bugiel, and Erik Derr. Reliable third-party library detection in android and its security applications. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 356–367, 2016.
- 8 Andrei Costin, Jonas Zaddach, Aurélien Francillon, and Davide Balzarotti. A {Large-scale} analysis of the security of embedded firmwares. In *23rd USENIX security symposium (USENIX Security 14)*, pages 95–110, 2014.
- 9 Gianluca Anselmi, Anna Maria Mandalari, Sara Lazzaro, and Vincenzo De Angelis. *COPSEC: Compliance-Oriented IoT Security and Privacy Evaluation Framework*. Association for Computing Machinery, New York, NY, USA, 2023.
- 10 Aniketh Girish, Tianrui Hu, Vijay Prakash, Daniel J Dubois, Srdjan Matic, Danny Yuxing Huang, Serge Egelman, Joel Reardon, Juan Tapiador, David Choffnes, et al. In the room where it happens: Characterizing local communication and threats in smart homes. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, pages 437–456, 2023.

## 6.20 Trustworthiness and Responsibility in AI – Causality, Learning, and Verification

**Editors: Vaishak Belle, Hana Chockler, Shannon Vallor, Kush R. Varshney, and Joost Vennekens**

**Seminar No. 24121**

Date: March 17–22, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.75

© Creative Commons BY 4.0 license

© Vaishak Belle, Hana Chockler, Shannon Vallor, Kush R. Varshney, and Joost Vennekens



**Participants:** Nadisha-Marie Aliman, Sander Beckers, Jan M. Broersen, Hana Chockler, Sjur K. Dyrkolbotn, Esra Erdem, Michael Fisher, Leilani H. Gilpin, Gregor Goessler, Till Hofmann, Leon Kester, Ekaterina Komendantskaya, Joshua Loftus, Giuseppe Primiero, Subramanian Ramamoorthy, Judith Simon, Daniel Susser, Kush R. Varshney, Joost Vennekens

### ■ Motivation and research area

How can we trust autonomous computer-based systems? Widely accepted definitions of autonomy take the view of being “independent and having the power to make your own decisions.” While many AI systems fit that description, they are often assembled by integrating many heterogeneous technologies – including machine learning, symbolic reasoning or optimization – and correspondingly the notion of trust is fragmented and bespoke for the individual communities. However, given that automated systems are increasingly being deployed in safety-critical environments whilst interoperating with humans, a system would not only need to be able to reason about its actions, but a human user would need to additionally externally validate the behavior of the system. This seminar tackled the issue of trustworthiness and responsibility in autonomous systems by considering: notions of cause, responsibility and liability, and tools to verify the behavior of the resulting system.

In the last few years, we have observed increasing contributions in terms of manifestos, position papers, and policy recommendations issued by governments and learned societies, touching on interdisciplinary research involving AI ethics. This has primarily focused on “Fairness, Accountability, and Transparency” (FAT) with a majority focus on fairness, as individual and group fairness seems relatively easier to define precisely. On the other hand, DARPA’s XAI agenda has led to a resurgence in diagnostic explanations, but also ignited the question of interpretability and transparency in machine learning models, especially deep learning architectures. Our high-level motivation is that governance and regulatory practices can be viewed not only as rules and regulations imposed from afar but instead as an integrative process of dialogue and discovery to understand why an autonomous system might fail and how to help designers and regulators address these through proactive governance. But before that agenda can be approached, we need to resolve an

important low-level question: how can we understand trust and responsibility of the components that make up an AI system? Autonomous systems will make ‘mistakes’, and accidents will surely happen despite best efforts. How should we reason about responsibility, blame, and causal factors affecting trustworthiness of the system? And if that is considered, what tools can we provide to regulators, verification and validation professionals, and system designers to help them clarify the intent and content of regulations down to a machine interpretable form? Existing regulations are necessarily vague, depending on the nuance of human interpretation for actual implementation. How should they now be made more precise and quantifiable?

The purpose of the seminar was to initiate a debate around these theoretical foundations and practical methodologies with the overall aim of laying the foundations for a “Trustworthiness & Responsibility in AI” framework – a framework for systems development methodology that integrates quantifiable responsibility and verifiable correctness into all stages of the software engineering process. As the challenge, by nature, is multidisciplinary, addressing it must involve experts from different domains, working on creating a coherent, jointly agreed framework. The seminar brought together researchers from Artificial Intelligence (AI), Machine Learning (ML), Robotics (ROB), hardware and software verification (VER), Software Engineering (SE) and the Humanities (HUM), especially Philosophy (PHI), who provided different and complementary perspectives on responsibility and correctness regarding the design of algorithms, interfaces, and development methodologies in AI. From the outset, we wished to especially focus on understanding correctness for AI systems that integrate or utilize data-driven models (i.e., ML models), and to anchor our discussions by appealing to causality (CAU). Causality is widely used in the natural sciences to understand the effect of interventions on observed correlations, allowing scientists to design physical and biological laws. In ML too, increasingly

there is recognition that conventional models focus on statistical associations, which can be misleading in critical applications demanding human-understandable explanations. The concept of causality is central to defining a notion of responsibility, and thus was a key point in our discussions.

### ■ Directions identified and discussed

The seminar involved extensive discussions between AI, ML, ROB, VER, SE, PHI and HUM researchers who have experience in the following research topics:

- Ethical aspects of AI & ML algorithms: explainability and interpretability in AI algorithms, bias & fairness, accountability, moral responsibility. For example, there were discussions on large language models, their black box nature, and capabilities. There was also quite a bit of work on how explanations and causality might be related. Relevant papers that the participants identified included [1, 10].
- The moral and legal concepts of responsibility that underpin trust in autonomous systems, and how these relate to or can be aided by explainability or causal models of responsibility.
- Technical aspects of AI & ML algorithms: explainability and interpretability in AI algorithms, bias & fairness, accountability, quantification of responsibility. There were discussions regarding how visual input and human-in-the-loop models could provide the next frontier of explainability. Relevant papers identified by the participants included [11].
- Complex AI systems: robotics, reinforcement learning, integrated task and motion planning, mixed-initiative systems. There were discussions that suggest that incorporating high-level specifications from humans could considerably enhance the literature. Examples include recent loss function-based approaches and program induction-related directions for reinforcement policies [4, 5].
- Software engineering for AI systems: development methodologies, specification synthesis, formal verification of ML models, including deep learning architectures, software testing, causality. Outside of a range of recent approaches and looking at verifying the robustness properties of newer networks, there was a discussion on enhancing these perspectives by modeling trust. In fact, what exactly trustworthy machine learning might look like and the components it might involve were also discussed. Examples of relevant work include [8, 9, 12].
- Causal analysis of counterexamples and software faults. Causality was a central topic in the discussion, anchoring some of the key perspectives on how trustworthy AI, as well as explanations, could be addressed along with more nuanced notions such as harm. Following Joseph Halpern's talk on how harm could be formalized and related discussions, a number of relevant papers were identified as promising starting points for causal analysis [2, 3].
- Social aspects of AI & society, AI & law, AI & ethics. Examples of related literature include ideas on the types of ethical robots, the ironies of automation, and the notion of

how empathy should apply to explainability among other related topics [6, 7]

### ■ Open questions

Discussions between researchers from these different areas of expertise allowed us to explore topics at the intersection between the main areas, and to ask (and obtain partial answers on) the following questions:

- What sorts of explanations, and more generally, correctness notions are users looking for (or may be helpful for them)? How should these be generated and presented?
- How should we reason about responsibility, blame and causal factors affecting trustworthiness in individual components? How should that be expanded to the overall AI system?
- How do we define and quantify trust? Is trust achieved differently depending on the type of the user? Can trust in AI be achieved only using technology, or do we need societal changes?
- How do users reason about and handle responsibility, blame and cause in their day-to-day activities, and how do we interface those concepts with that of the AI system?
- Do our notions of responsibility and explanations increase user's trust in the technology?
- Who are the users of the technology? We envision different types of users, from policy makers and regulators to developers of the technology, to laypeople – the end-users. Should we differentiate the type of analysis for different categories of users?
- What tools can we provide to regulators, verification and validation professionals and system designers to help them clarify the intent and content of regulations down to a machine interpretable form?
- What tools are available to verify ML components, and do they cover the scope of “correct behavior” as understood by users and regulators?
- What SE practices are relevant for interfacing, integrating and challenging the above notions?
- How can properties of AI systems that are of interest be expressed in languages that lend themselves to formal verification or quantitative analysis?
- What kinds of user interfaces are needed to scaffold users to scrutinise the way AI systems operate?
- What frameworks are needed to reason about blame and responsibility in AI systems?
- How do we integrate research in causal structure learning with low-level ML modules used in robotics?
- How do we unify tools from causal reasoning and verification for assessing the correctness of complex AI systems?
- What challenges arise in automated reasoning and verification when considering the above mixed-initiative systems?
- Given a falsification of a specification, what kind of automated diagnosis, proof-theoretic and causal tools are needed to identify problematic components?
- How broadly will counterfactual reasoning (i.e., “what-if” reasoning) be useful to tackle such challenges?

### ■ References

- 1 Lisanne Bainbridge. Ironies of automation. *Automatica*, 19, 1983.
- 2 Sander Beckers, Hana Chockler, and Joseph Halpern. A causal analysis of harm. *Advances in Neural Information Processing Systems*, 35:2365–2376, 2022.
- 3 Ilan Beer, Shoham Ben-David, Hana Chockler, Avigail Orni, and Richard Treffer. Explaining counterexamples using causality. *Formal Methods in System Design*, 40:20–40, 2012.
- 4 Vaishak Belle and Andreas Bueff. Deep inductive logic programming meets reinforcement learning. In *The 39th*

- International Conference on Logic Programming*. Open Publishing Association, 2023.
- 5 Craig Innes and Subramanian Ramamoorthy. Elaborating on learned demonstrations with temporal logic specifications. *arXiv preprint arXiv:2002.00784*, 2020.
  - 6 William Kidder, Jason D’Cruz, and Kush R Varshney. Empathy and the right to be an exception: What llms can and cannot do. *arXiv preprint arXiv:2401.14523*, 2024.
  - 7 Bran Knowles, Jason D’Cruz, John T. Richards, and Kush R. Varshney. Humble ai. *Commun. ACM*, 66(9):73–79, aug 2023.
  - 8 Ekaterina Komendantskaya and Guy Katz. Towards a certified proof checker for deep neural network verification. *Logic-Based Program Synthesis and Transformation*, page 198.
  - 9 Madsen and Gregor. Measuring human-computer trust. In *11th Australasian Conference on Information Systems*, 2000.
  - 10 James H. Moor. Four types of ethical robot. *Philosophy Now*, 2009.
  - 11 Spinner, Schlegel, Schäfer, and El-Assady. explAiner: A visual analytics framework for interactive and explainable machine learning. *IEEE Trans. on Visualization and Computer Graphics*, 2020.
  - 12 Kush R. Varshney. *Trustworthy Machine Learning*.

## 6.21 Low-Dimensional Embeddings of High-Dimensional Data: Algorithms and Applications

Editors: Dmitry Kobak, Fred Hamprecht, Smita Krishnaswamy, and Gal Mishne  
Seminar No. 24122

Date: March 17–22, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.3.92

© Creative Commons BY 4.0 license

© Dmitry Kobak, Sebastian Damrich, Fred A. Hamprecht, Smita Krishnaswamy, and Gal Mishne



**Participants:** Michael Bleher, Corinna Coupette, Cyril de Bodt, Laleh Haghverdi, Fred Hamprecht, Dmitry Kobak, Smita Krishnaswamy, B.P.F. Lelieveldt, Gal Mishne, Maximilian Noichl, Bastian Rieck, Benjamin M. Schmidt, Guy Wolf

**Remote Participants:** Miguel Á. Carreira-Perpiñán

### 2D embeddings in science

In recent years, high-dimensional “big” data have become commonplace in multiple academic fields. To give some examples, single-cell transcriptomics routinely produces datasets with sample sizes in hundreds of thousands and dimensionality in tens of thousands [1]; single-cell mass spectrometry deals with millions of samples [2]; genomic datasets quantifying single-nucleotide polymorphisms can deal with many millions of features [3]; behavioural physiology produces high-dimensional datasets with tens of thousands of samples [4]. In neuroscience, calcium imaging allows to record time-series activity of thousands of neurons. Many scientific fields that traditionally did not have to deal with high-dimensional data analysis now face similar

challenges; for example, a digital library can yield a dataset with tens of millions of samples and hundreds, if not millions, of features [5].

Such datasets require adequate computational methods for data analysis, including unsupervised data exploration. In fact, exploratory statistical analysis has become an essential tool in many scientific disciplines, allowing researchers to compactly visualise, represent and make sense of their data. It became commonplace to explore low-dimensional embeddings of the data, generated by methods like t-SNE [6] or UMAP [7]. Such visualisation has proven to be a valuable tool for exploring the data, performing quality control, and generating scientific hypotheses (Figure 6.4).




Fig. 6.4  
Example applications in single-cell transcriptomics. Left: cortical neurons [8], sample size  $n = 1.2\text{M}$ . Middle: human brain organoid development [9],  $n = 43\text{K}$ . Right: human blood and bone marrow cells in leukaemia [10],  $n = 70\text{K}$ . Figures from original publications.

Similar algorithms are also applied in artificial intelligence research to visualise massive datasets used to train state-of-the-art artificial intelligence models, such as image-based and text-based generative models. This allows researchers to discover biases and


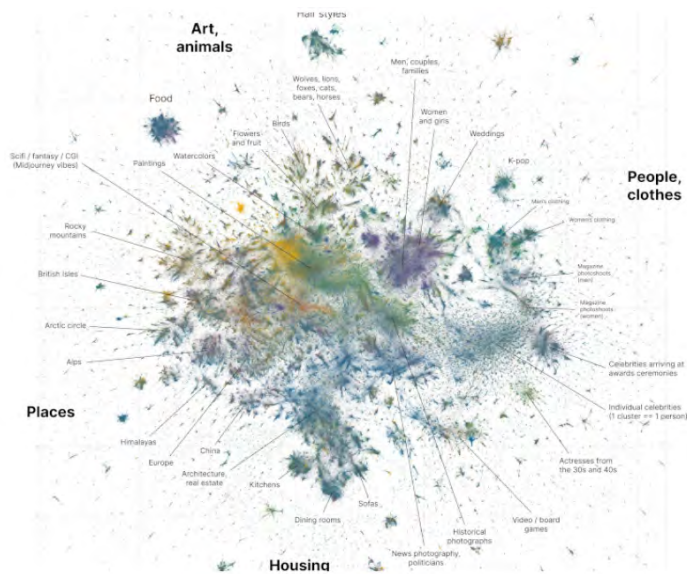


Fig. 6.5

Example applications in artificial intelligence. Left: GPT4All-J training data [11],  $n = 800K$ . Right: image captions from LAION-Aesthetics dataset (figure by David McClure),  $n = 12M$ .



Neighbour embedding methods like t-SNE and UMAP create a low-dimensional map of the data based on the k-nearest-neighbour graph. As a result, they are often unable to reproduce large-scale global structure of the data [12], creating potentially misleading visualizations [13]. Acquisition of increasingly high-dimensional data across scientific fields has sparked widespread interest in employing dimensionality reduction and visualisation methods. However, there is a gap between method developers who propose and implement these algorithms, and domain experts who aim to use them. The purpose of this seminar was to bring together machine learning researchers, theoreticians, and practitioners, to address current gaps in theoretical guarantees and evaluation measures for state-of-the-art approaches, highlight practical challenges in applying these techniques in different domains, brainstorm the solutions, and set up new collaborations to tackle open problems in this vibrant field.

## ■ Seminar topics

The overarching purpose of this Dagstuhl Seminar was to brainstorm open problems and challenges in the field of low-dimensional embeddings, as seen by (i) practitioners; (ii) theoreticians and mathematicians; and (iii) machine learning researchers — leading to new collaborations to tackle these problems. The seminar focused on the following open questions, grouped into four areas.

### ■ Low-dimensional embeddings in actual practice

Single-cell biology, working with large quantities of high-dimensional data and interested in exploratory research, became a field heavily relying on low-dimensional embeddings. But

gaps in the data, to highlight model limitations, and ultimately to develop better models (Figure 6.5). A concise overview of the model's training data can also be helpful for societal oversight and public communication.

embeddings of texts [5], of genomes [14], of graph nodes [15], of chemical structures [16], etc., are also rapidly gaining popularity. Seminar participants discussed and brainstormed which fields in the coming years are likely to generate data amenable for embedding methods, and compared challenges raised by each of these application fields.

Neighbour embeddings have a number of well-known limitations [12]: for example, they can strongly distort the global structure of the data and are unable to represent high-dimensional topological features of the data. These artefacts can lead practitioners to incorrect scientific conclusions or to chasing unfounded hypotheses. We extensively discussed (i) which limitations can be addressed by the new generation of algorithms; (ii) how to diagnose misleading aspects of any given embedding; and (iii) what evaluation metrics are necessary and sufficient for comparing different visualisation techniques.

Moreover, two-dimensional embeddings have been recently criticised as being dangerously misleading [13]. At the same time, they are widely used across many disciplines and can be helpful in actual scientific practice, if used with care [12]. In several talks and multiple discussions, seminar participants talked about specific examples of how and where the embeddings are useful, and which best practices can help to avoid them being misleading.

### ■ Common themes across state-of-the-art algorithms and relevant trade-offs

One common theme in multiple talks and discussions was trade-offs between various embedding algorithms.

First, methods like t-SNE or UMAP are typically used to produce 2D or 3D embeddings, while spectral methods like Laplacian eigenmaps [17] produce low-dimensional embeddings that are often used with more embedding dimensions. This

is less suitable for visualisation but may be better suited for downstream data analysis. Several seminar participants reported successfully applying UMAP to intermediate dimensionality too,

with particular benefits for downstream density-based clustering (using HDBSCAN algorithm).

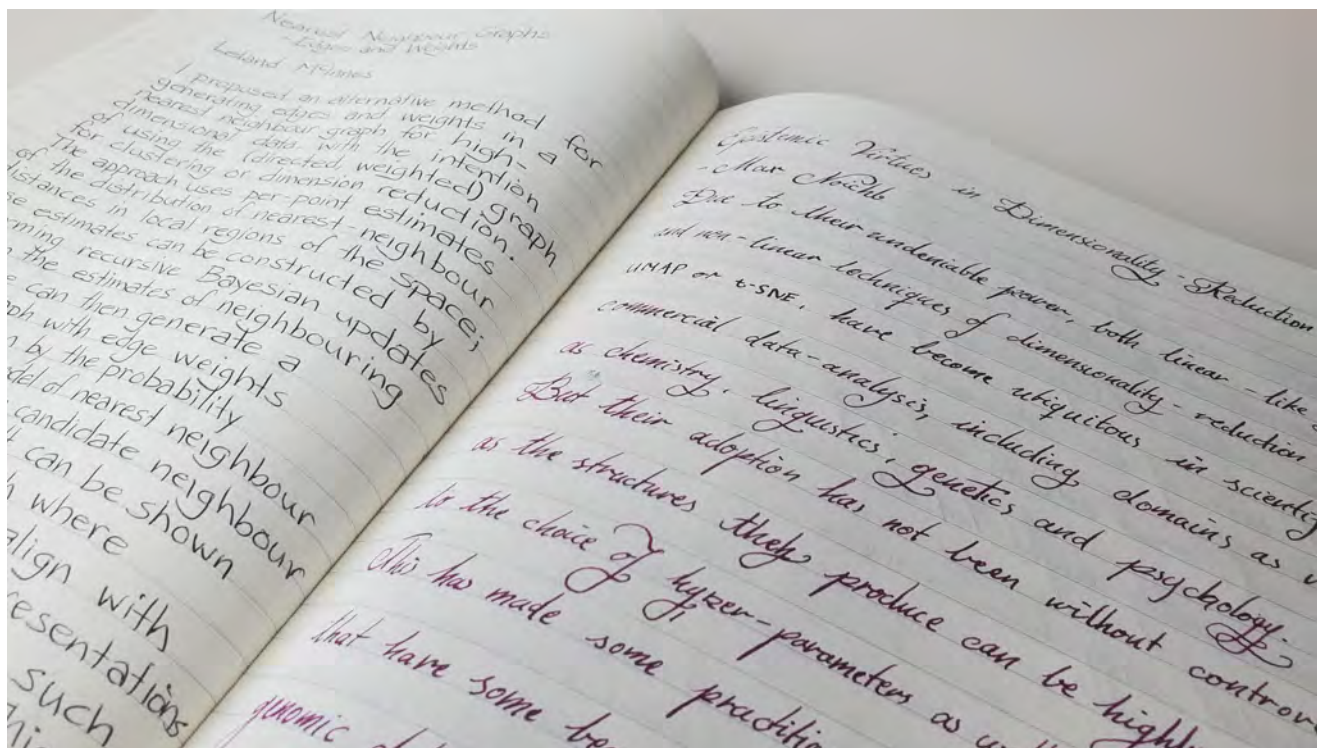


Fig. 6.6  
Handwritten abstracts from our seminar.

Second, all neighbour embedding algorithms operate on the kNN graph of the data but use different loss functions and different attractive/repulsive forces to arrive at the final layout. This yields various trade-offs in the quality of global/local structure preservation [18].

Third, neighbour embedding algorithms are typically run on a kNN graph constructed using pairwise Euclidean distances, but in principle any other metric can be used as well. Specifically, metric design can be useful for incorporating domain knowledge and statistical priors on the data [19,20]. We discussed what kinds of data can profit from using non-Euclidean distance metrics, or from kNN graph post-processing, such as diffusion-based smoothing.

Fourth, more generally, neighbour embeddings are known to be related to the self-supervised learning approach known as contrastive learning [21]. However, despite substantial progress in each of these two fields, they stayed largely disconnected from each other. Seminar participants argued that both contrastive learning and neighbour embedding research can benefit from each other's state-of-the-art approaches, and in particular can be combined to develop new algorithms for visualising textual and/or graph-based data.

Fifth, while neighbour embeddings only aim to preserve nearest neighbours, methods based on MDS aim to preserve all pairwise distances including the large ones. In Isomap [22] and PHATE [23], pairwise distances are obtained as graph distances

on the kNN graph. Isomap uses short path distance, while PHATE uses diffusion-based distance called potential distance. LDLE [24] uses bottom-up manifold learning to align low-distortion local embeddings to a global embedding. We discussed to what extent these approaches can capture both the local and the global structure of the data, and what the advantages and the disadvantages of aiming to preserve global aspects of the data are.

#### ■ Interactive embeddings

Another extensively discussed topic was interactive visualizations of 2D embeddings (in particular see abstracts by Benjamin M. Schmidt and B. P. F. Lelieveldt). While most often low-dimensional embeddings are depicted as static images, they can be powerful tools for *interactive* data explorers. NomicAI has been developing software for in-browser interactive explorers, while the group of B. P. F. Lelieveldt has been working on stand-alone software for interactive exploration of RNA-sequencing data.

#### ■ Perspective paper

During the seminar, participants decided to work together on a perspective paper, provisionally titled like the seminar: "Low-dimensional embeddings of high-dimensional data". During the seminar, we organized several brainstorming sessions on what should the paper cover and how the material should be organized. The writing is currently underway and we hope to be able to release the work some time in summer 2024.

## References

- 1 Kobak, Dmitry and Berens, Philipp *The art of using t-SNE for single-cell transcriptomics*, Nature Communications, 10(1): 1–14, 2019.
- 2 Belkina, Anna C and Ciccolella, Christopher O and Anno, Rina and others *Automated optimized parameters for T-distributed stochastic neighbor embedding improve visualization and analysis of large datasets*, Nature Communications, 10(1): 5415, 2019.
- 3 Diaz-Papkovich, Alex and Anderson-Trocmé, Luke and Ben-Eghan, Chief and Gravel, Simon *UMAP reveals cryptic population structure and phenotype heterogeneity in large genomic cohorts*, PLoS genetics, 15(11): e1008432, 2019.
- 4 Kollmorgen, Sepp and Hahnloser, Richard HR and Mante, Valerio *Nearest neighbours reveal fast and slow components of motor learning*, Nature, 577(7791): 526–530, 2020.
- 5 Schmidt, Benjamin *Stable random projection: Lightweight, general-purpose dimensionality reduction for digitized libraries*, Journal of Cultural Analytics, 3(1), 2018.
- 6 van der Maaten, Laurens and Hinton, Geoffrey *Visualizing data using t-SNE*, Journal of Machine Learning Research, 9(11), 2008.
- 7 McInnes, Leland and Healy, John and Melville, James *UMAP: Uniform manifold approximation and projection for dimension reduction*, arXiv preprint arXiv:1802.03426, 2018.
- 8 Yao, Zizhen and Van Velthoven, Cindy TJ and Nguyen, Thuc Nghi and others *A taxonomy of transcriptomic cell types across the isocortex and hippocampal formation*, Cell, 184(12): 3222–3241, 2021.
- 9 Kanton, Sabina and Boyle, Michael James and He, Zhisong and others *Organoid single-cell genomic atlas uncovers human-specific features of brain development*, Nature, 2019.
- 10 Triana, Sergio and Vonficht, Dominik and Jopp-Saile, Lea and others *Single-cell proteo-genomic reference maps of the hematopoietic system enable the purification and massive profiling of precisely defined cell states*, Nature Immunology, 22(12): 1577–1589, 2021.
- 11 Anand, Yuvanesh and Nussbaum, Zach and Treat, Adam and other *GPT4All: An Ecosystem of Open Source Compressed Language Models*, arXiv preprint arXiv:2311.04931, 2023.
- 12 Wattenberg, Martin and Viégas, Fernanda and Johnson, Ian *How to Use t-SNE Effectively*, Distill, 2016, 10.23915/distill.00002.
- 13 Chari, Tara and Pachter, Lior *The specious art of single-cell genomics*, PLoS Computational Biology, 19(8): e1011288, 2023.
- 14 Diaz-Papkovich, Alex and Anderson-Trocmé, Luke and Gravel, Simon *A review of UMAP in population genetics*, Journal of Human Genetics, 66(1): 85–91, 2021.
- 15 Hu, Yifan *Efficient, high-quality force-directed graph drawing*, Mathematica journal, 10(1): 37–71, 2005.
- 16 Probst, Daniel and Reymond, Jean-Louis *Visualization of very large high-dimensional data sets as minimum spanning trees*, Journal of Cheminformatics, 12(1): 12, 2020.
- 17 Belkin, Mikhail and Niyogi, Partha *Laplacian eigenmaps for dimensionality reduction and data representation*, Neural Computation, 15(6): 1373–1396, 2003.
- 18 Böhm, Jan Niklas and Berens, Philipp and Kobak, Dmitry *Attraction-Repulsion Spectrum in Neighbor Embeddings*, Journal of Machine Learning Research, 23(95), 2022.
- 19 Talmon, Ronen and Coifman, Ronald R *Empirical intrinsic geometry for nonlinear modeling and time series filtering*, Proceedings of the National Academy of Sciences, 110(31): 12535–12540, 2013.
- 20 Mishne, Gal and Talmon, Ronen and Meir, Ron and others *Hierarchical coupled-geometry analysis for neuronal structure and activity pattern discovery*, IEEE Journal of Selected Topics in Signal Processing, 10(7): 1238–1253, 2016.
- 21 Damrich, Sebastian and Böhm, Niklas and Hamprecht, Fred A and Kobak, Dmitry *From t-SNE to UMAP with contrastive learning*, In *The Eleventh International Conference on Learning Representations*, 2023.
- 22 Tenenbaum, Joshua B and Silva, Vin de and Langford, John C *A global geometric framework for nonlinear dimensionality reduction*, Science, 290(5500): 2319–2323, 2000.
- 23 Moon, Kevin R and Van Dijk, David and Wang, Zheng and others *Visualizing structure and transitions in high-dimensional biological data*, Nature Biotechnology, 37(12): 1482–1492, 2019.
- 24 Kohli, Dhruv and Cloninger, Alexander and Mishne, Gal *LDLE: Low distortion local eigenmaps*, Journal of machine learning research, 22(282): 1–64, 2021.

## 6.22 Network Calculus

**Editors: Steffen Bondorf, Anne Bouillard, Markus Fidler, and Jörg Liebeherr**  
**Seminar No. 24141**

Date: April 1–4, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.1

© Creative Commons BY 4.0 license

© Lisa Maile

**Participants:** Michael Beck, Steffen Bondorf, Anne Bouillard, Marc Boyer, Peter Buchholz, Almut Burchard, Georg Carle, Samarjit Chakraborty, Vlad-Cristian Constantin, Hugo Daigormte, Markus Fidler, Anaïs Finzi, Stéphane Gaubert, Damien Guidolin–Pina, Anja Hamscher, Max Helm, Kai-Steffen Jens Hielscher, Yuming Jiang, Kai Lampka, Jean-Yves Le Boudec, Harvinder Lehal, Jörg Liebeherr, Lisa Maile, Mahsa Noroozi, Xi Peng, Stéphan Plassart, Amr Rizk, Giovanni Stea, Mina Tahmasbi Arashloo, Ludovic Thomas, Brenton Walker, Kui Wu, Raffaele Zippo



Our Dagstuhl Seminar brought together leading experts in the field of network calculus to discuss recent research activities, identify future directions, and establish or strengthen cooperation. The seminar fostered collaborations and new ideas by focusing on diversity, incorporating perspectives from research and industry, varying levels and areas of expertise, and participants of different positions and ages. Attendees came from various scientific disciplines, including deterministic and stochastic network calculus, Age-of-Information, industrial communication, and wireless technologies. All combined under the need for Quality of Service (QoS) requirements. Participants contributed scientific presentations and workshops on their current and future research.

Network calculus is a theoretical framework for analyzing networks, mainly, but not limited to communication networks, but applications are also possible in the field of energy systems or even emergency call centers. By modeling data flows and systems as mathematical functions, network calculus allows the calculation of guaranteed (deterministic/stochastic) upper bounds. This analysis can be applied to various elements, including individual hops or entire systems. Additionally, network calculus has been used for certification purposes in avionic networks, such as the Airbus A380 [1], demonstrating its importance and applicability in designing safety-critical systems that require stringent QoS performance guarantees.

The seminar focused on several main areas: integrating algebras and performance metrics, exploring network topology and parallel systems, applying network calculus to emerging technologies, and developing algorithms and tools. New insights in these topics were presented, followed by group work on individual problems. These group sessions explored algorithmic challenges, mathematical and computational complications, and practical applications of network calculus.

In the context of algebras and performance metrics, discussions centered on representing network calculus as a systems

theory under min-plus algebra, as well as integrating min-plus and max-plus algebras. This dual approach could provide a more comprehensive framework for analyzing complex network behaviors. A significant challenge addressed was the issue of rare perturbations, where deterministic network calculus currently lacks the ability to quantify the frequency or rarity of worst-case scenarios. Participants discussed the need for more nuanced metrics to distinguish between frequent low delays and rare high delays.

The seminar also highlighted the powerful modeling capabilities of network calculus in representing various systems, such as links, traffic shapers, and scheduling policies, which can be composed into arbitrary topologies. Advanced results from deterministic network calculus, like pay bursts and multiplexing only once phenomena, have recently been incorporated into stochastic network calculus, necessitating strong cooperation between the deterministic and stochastic communities. Parallel systems, particularly in the context of multi-path transport and fork-join models, were another central point, with discussions addressing the complexities and opportunities inherent in these configurations.

Emerging technologies such as Time-Sensitive Networking (TSN), Deterministic Networking (DetNet), and Ultra-Reliable Low Latency Communications (URLLC) were also a significant focus of the seminar. Network calculus currently receives a lot of research attention due to these emerging topics, which attract high interest from both industry and academia. These technologies are crucial for applications in factory automation, aerospace onboard systems, and automotive in-vehicle networks, where performance guarantees are critical. The seminar emphasized the integration of network calculus with these emerging technologies to ensure robust performance analysis and guarantees, supporting applications that meet stringent reliability and latency requirements.

Recent advancements in network calculus algorithms were

another key theme of the seminar. Participants discussed combining modular and optimization approaches and developing highly parallelizable methods that iteratively improve performance bounds. The use of machine learning for appropriate decomposition in modular approaches was highlighted, allowing for more efficient and scalable solutions to complex network performance challenges. Additionally, the development of tools to rapidly disseminate novel results and facilitate extensive community-based research artifact evaluation and reproducibility verification was emphasized.

Discussions following each presentation helped identify open problems and challenges to be addressed by the community. Participants highlighted the importance of collaboration between the different network calculus communities and experts, the topics which are still open and where our community could focus in the

future, and potentials to make network calculus more accessible to the public.

The seminar also included connecting personal and scientific discussions during a slightly rainy but enjoyable hike, long extended dinners, and quizzes to engage in team building. These social events provided additional opportunities for participants to engage and exchange ideas.

In summary, the Dagstuhl Seminar on network calculus facilitated significant discussions on advancing the field through interdisciplinary collaboration, integrating new technologies, and leveraging new methodologies for optimization. The insights and outcomes from the seminar pave the way for future research and development, ensuring network calculus receives recognition and remains a robust and versatile tool for network analysis and performance guarantees.

## ■ References

- 1 F. Francés, C. Fraboul, and J. Griedu, *Using Network Calculus to optimize the AFDX network*, ERTS 2006: 3rd

European Congress ERTS Embedded real-time software, 2006.

## 6.23 Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems

6

**Editors: Ignacio J. Alvarez, Philip Koopman, Mario Trapp, and Elena Troubitsyna**  
**Seminar No. 24151**

Date: April 7–12, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.23

© Creative Commons BY 4.0 license

© Elena Troubitsyna, Ignacio J. Alvarez, Philip Koopman, and Mario Trapp

**Participants:** Magnus Albert, Ignacio J. Alvarez, Claus Bahlmann, Ensar Becic, Nicolas Becker, Simon Burton, Radu Calinescu, Betty H. C. Cheng, Krzysztof Czarnecki, Niels De Boer, Francesca Favaro, Lydia Gauerhof, Mallory Graydon, Jérémie Guiochet, Hans Hansson, Fuyuki Ishikawa, Aaron Kane, Lennart Kilian, Jörg Koch, Philip Koopman, Lars Kunze, Jonas Nilsson, Ganesh J. Pai, Nick Reed, Jan Reich, Martin Rothfelder, Philippa Ryan, Fredrik Sandblom, Tiziano Santilli, Jan Stellet, Reinhard Stolle, Stefano Tonetta, Mario Trapp, Elena Troubitsyna, Kim Wasson, Alan Wassyn, William H. Widen, Rafael Zalman



### ■ Introduction

The examples of modern autonomous systems include self-driving cars, UAV (drones), underwater vehicles, various industrial and home service robots. In general, autonomous systems are intended to operate without human intervention over prolonged time periods, perceive their operating environment and adapt to internal and external changes.

For example, a self-driving car gathers information from camera and lidar to detect, e.g., pedestrians on the road and plan collision avoidance maneuvers, slowing down or breaking, i.e., avoid hazards. The perception functions process the inputs of various sensors and generate the internal model of the operating environment. By relying on this model, the decision functions plan and execute the actions required to achieve the goals of the mission. In general, they follow the generic “sense-understand-decide-act” behavioral pattern, which is also traditionally adopted in robotics.

Both sensing and decision making usually rely on Artificial Intelligence (AI), in particular Machine Learning (ML). While AI and ML algorithms have already been used in robotics for several decades, their use in safety-critical systems is fairly new and currently not appropriately addressed by safety engineering neither from technological, nor from organizational and legal points of view.

The problem of safe AI has received a significant amount of research and industrial attention over the last few years, but there has been a divergence in the approaches taken by the safety and the ML communities. Moreover, it has become clear that the safety assurance problems cannot be resolved by improving the ML algorithms alone. Hence, the research communities should consolidate their efforts in creating methods and tools enabling a holistic approach to safety of autonomous systems.

This motivated the topic of our Dagstuhl Seminar – exploring the problem of engineering and assuring safety of autonomous

systems from an interdisciplinary perspective. A group of experts from avionics, automotive, machine learning, simulation, verification and validation and safety engineering reviewed the current academic state-of-the-art, industry practices and standardization to determine the latest achievement and challenges in developing and safety assurance for autonomous systems over a broad range of technological, organizational, ethical and legal perspectives.

As a result, the discussions of achievements and challenges in developing and assuring safety of autonomous systems spanned over a broad range of technological, organizational, ethical and legal topics.

### ■ Organisation of the seminar

The seminar brought together researchers and practitioners from different disciplines and application domains. Since, currently, the innovation in autonomous systems is strongly led by industry, a significant number of participants were industrial engineers, who not only shared their best practices but also identified unsolved research problems. In constructive debates, we discussed the results of applying and experimenting with various techniques for engineering safe autonomous systems and identified open research challenges.

To facilitate an open discussion between the participants, and analyze the problem of engineering safe autonomous systems from different points of view, before the seminar, we identified the following general discussion themes:

- Role of formal methods in engineering and assurance of safe autonomous systems
- Regulatory, assurance and standards for safety-critical autonomous systems
- Safety of AI-based system versus normal technical system safety
- Safety and security interactions
- Risk acceptance for autonomous systems

This report presents the summaries of the discussions focused on the specific topics within these themes.

We would like to acknowledge the supporting contributors – the session chairs and scribes that helped to collect the information for this report: Magnus Albert (SICK AG – Waldkirch, DE), Ensar Becic (National Transportation Safety Board, US), Nicolas Becker (Stellantis France – Poissy, FR), Simon Burton (Gerlingen, DE), Radu Calinescu (University of York, GB), Betty H. C. Cheng (Michigan State University – East Lansing, US), Krzysztof Czarnecki (University of Waterloo, CA), Niels De Boer (Nanyang TU – Singapore, SG), Lydia Gauerhof (Bosch Center for AI – Renningen, DE), Jérémie Guiochet (LAAS –

Toulouse, FR), Hans Hansson (Mälardalen University – Västerås, SE), Aaron Kane (Edge Case Research – Pittsburgh, US), Lars Kunze (University of Oxford, GB), Jonas Nilsson (NVIDIA Corp. – Santa Clara, US), Nick Reed (Reed Mobility – Wokingham, GB), Jan Reich (Fraunhofer IESE – Kaiserslautern, DE), Martin Rothfelder (Siemens – München, DE), Philippa Ryan (University of York, GB), Fredrik Sandblom (Zenseact AB – Gothenburg, SE), Stefano Tonetta (Bruno Kessler Foundation – Trento, IT), Kim Wasson (Joby Aviation – Santa Cruz, US), and William H. Widen (University of Miami – Coral Gables, US). In the spirit of Chatham House Rules that prevailed in the meeting, we are not attributing any particular written text to any particular person.

## 6.24 Research Software Engineering: Bridging Knowledge Gaps

Editors: Stephan Druskat, Lars Grunske, Caroline Jay, and Daniel S. Katz

Seminar No. 24161

Date: April 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.42

© Creative Commons BY 4.0 license

© Stephan Druskat, Lars Grunske, Caroline Jay, and Daniel S. Katz

**Participants:** Stuart Allen, David E. Bernholdt, Jeffrey Carver, Mikaela Cashman McDevitt, Neil Chue Hong, Myra B. Cohen, Hannah Cohoon, Ian Cosden, Stephan Druskat, Nasir Eisty, Michael Felderer, Carole Goble, Michael Goedicke, Samuel Grayson, Lars Grunske, Robert Haines, Wilhelm Hasselbring, Toby Hodges, Caroline Jay, Guido Juckeland, Daniel S. Katz, Timo Kehrer, Anna-Lena Lamprecht, Christopher Klaus Lazik, Jan Linxweiler, Frank Löffler, Michael Meinel, Sebastian Müller, Lata Nautiyal, Bernhard Rumpe, Heidi Seibold, Jan Philipp Thiele, Colin Venters, Samantha Wittke, Yo Yehudi, Shurui Zhou



Research Software Engineering (*RSEng*) is the practice of applying knowledge, methods and tools from software engineering in research. Software Engineering Research (*SER*) develops methods to support software engineering work in different domains. The practitioners of research software engineering working in academia – Research Software Engineers (*RSEs*) – are often not trained software engineers. Nevertheless, *RSEs* are the software experts in academic research. They translate research to software, enable new and improved research, and create software as an important output of research [1].

Hypothetically, the *RSEng* community and the *SER* community could benefit from each other. *RSEs* could leverage software engineering research knowledge to adopt state-of-the-art methods and tools, thereby improving *RSEng* practice towards better research software. Vice versa, software engineering research could adopt *RSEng* more comprehensively as a research object, to investigate the methods and tools required for the application of state-of-the-art software engineering in research contexts [2].

There are currently both unknown and known unknowns that make it hard for either community to attain the benefits mentioned above. We call these unknowns *gaps*, and we call methods to discover the known unknowns and to clarify or resolve the (subsequently) known unknowns *bridges*.

To find the gaps between research software engineering and software engineering research, and start building bridges between the two communities with the aim to create mutual benefit through reciprocal collaboration, we organized and held a five-day seminar in April 2024 at Schloss Dagstuhl – Leibniz Center for Informatics, as Dagstuhl Seminar 24161 “Research Software Engineering: Bridging Knowledge Gaps”. Here, we report and document the seminar’s program, outputs, and potential outcomes.

### ■ Seminar participants

In the past, there has been little focused direct communication between the *RSEng* and *SER* communities. Anecdotally, while the German *RSEng* and *SER* conferences co-located in 2023 in Paderborn and shared a break room, there was little, and only informal, exchange in session attendance by participants of either conference, and hence little knowledge exchange.

We specifically organized our Dagstuhl Seminar to serve as a bridge across this type of **communication gap** between the communities, by inviting international experts from both communities as well as individuals who have a track record of working at the intersection between the communities. We also invited experts in adjacent fields such as education and training, and research software funding.

In a pre-seminar survey, we asked participants if they could contribute experience in any of the following areas:

- *Research Software Engineering*: practicing software engineering in a research context
- *Software Engineering*: practicing software engineering
- *Research on Software Engineering*: conducting *SER*
- *Research on Research Software Engineering*: conducting research on *RSEng*
- *Teaching, training, instruction*
- *Research software funding*
- *Community building*

Answers were provided on a 5-point Likert scale: “None” – “A little” – “Some” – “Much” – “A lot”. The results are shown in Figure 6.7. They suggest that we were mostly successful in bringing together participants with at least some relevant expertise, with the possible exception of research software funding, where more than half of participants claimed little or no expertise. This was partly due to invitees with a known background in research software funding being unavailable to attend the seminar.




Fig. 6.7 Survey responses to a question on experience in a given area. A total of 35 participants replied to the survey. Each person answered one of five levels of experience in the respective area.

### Seminar program

The seminar program was prepared to mainly enable direct collaboration of participants from both the SER and the RSEng communities. Sessions were run either in the plenary, or in breakout groups.

After an introductory session where participants introduced themselves and their work, short presentations were given on ideas for collaborative groups to form at the seminar. From the

original set of ideas, participants eventually formed a total of seven breakout groups to work on specific topics. Their work is summarized in the respective working group abstracts in this report. In addition, three additional ideas were brought forward in the course of the seminar in a more informal manner outside of sessions. For these additional ideas, work would either be done in parallel to the sessions, or was planned to be done after the seminar. An overview of all topics is given in Table 6.1.

Table 6.1 Topics of working groups and breakout initiatives worked on at the seminar.

Collaboration	Topic
Working group	Research Software Engineering Training & Education
Working group	Better Architecture, Better Software, Better Research
Working group	Bridging Communities
Working group	Developing a Common Language
Working group	Research Software: Towards Categories and Lifecycles
Working group	Security and usability of research software
Working group	Demystifying research software engineering for research group leaders
Breakout discussion	Software Engineering Equity, Diversity, Inclusion or Accessibility Research in Research Software/RSEng
Breakout project	Short software engineering social videos

The introductory session was rounded off by a primer presentation of research software engineering, and a preliminary discussion of options for the dissemination of seminar outcomes. Based on early feedback from participants, we decided to give up on our original idea of collecting seminar outcomes as chapters of an edited volume – a field manual for research software engineering. Instead, we discussed options for collecting some of

the outputs in a journal special issue, while leaving it generally up to groups to determine the best way for their outputs to be disseminated.

For the remainder of the seminar, work was originally planned to take place in breakout groups, with regular reporting and feedback sessions in the plenary. The third day was planned as an

exception, with two fishbowl sessions scheduled for the morning sessions, and a group excursion in the afternoon.

Based on feedback from participants on the first day of the seminar, we realized that there was a wider **perception gap** between RSEs and software engineering researchers than originally anticipated: software engineering researchers specifically did not feel entirely confident in their understanding of the scope and practice of RSEng, and of what research software projects looked like. Vice versa, RSEs did not feel entirely confident in their understanding of the aims and scope of software engineering research. We aimed to address this gap on the second day of the seminar with an ad-hoc “Ask us anything” plenary session where software engineering researchers asked members of the RSEng community questions about their work and experience, and with a subsequent ad-hoc session where RSEs presented the contexts they work in, setup of RSE groups, and particular research software projects.

In light of the need to address this kind of feedback quickly, we moved to adapt the original schedule for each day the night before, incorporating feedback we have had during the day.

In the same spirit as the “Ask us anything” sessions, we ran two sessions of “mythbusting fishbowls,” where members of the SER, and then the RSEng, community formed a dynamic panel who discussed topics suggested by the respective other community: preconceptions that one community had about the other’s work were fielded via [sli.do](#), voted for by the audience, and the most popular ones then discussed by the panel. Panel members were replaced whenever another community member wanted to contribute to the discussion.

Outputs of the seminar activities are presented in brief in the next section.

## ■ Outputs

A central *outcome* of the seminar is that community members from the SER and RSEng communities started collaborations to identify and bridge gaps between them.

The outputs of the working groups will be invited to be submitted as articles to a special issue “Research Software Engineering: Discovering and Bridging Knowledge Gaps” in *IEEE Computing in Science & Engineering*, to be edited by the seminar organizers and published in 2025.

Beyond collecting outputs in this way, one of the working groups developed and published a website that aims to collect

outputs from the seminar in a way accessible to a wider public and invites contributions from the community: [ser-rse-bridge.github.io](https://ser-rse-bridge.github.io). The website also includes a mapping of terms between SER and RSEng, based on the Guide to the Software Engineering Body of Knowledge [3], which is currently under development (see Section 3.1 of the full report).

Additionally, a series of short videos was produced during the seminar. In these videos, participants introduce central SER and RSEng knowledge concepts in under a minute. These can be used on social media platforms to create interest in these topics with, e.g., students looking to choose their courses.

## ■ Conclusion and future work

In conclusion, Dagstuhl Seminar “Research Software Engineering: Bridging Knowledge Gaps” (24161) was successful in bringing together members of two communities that have a vested interest in research software engineering: research software practitioners and software engineering researchers. Together with software education and community experts, we learned about each others’ work and started conversations and collaborations.

Creating conversations between separate communities and their cultures and codes proved to be challenging at times, e.g., where incentives differed. Where we observed antagonism, or where it was brought to our attention by participants, we tried to defuse it and steer conversations into a constructive direction. We are confident that by and large, this worked well due to flexibility on the side of the participants and a general will to collaborate and progress.

We found that adapting the program to the needs of participants where possible while maintaining the general direction, and arguably intensive workload, helped make the seminar very productive and engaging. Participants have continued to engage after the seminar to identify gaps and potential bridges between the communities.

Future work should focus on the continuation of the efforts started at the seminar, and continued communication and collaboration between the communities. We believe that the seminar marked a starting point for collaboration that can realize future reciprocal benefit for research software engineering and software engineering research in equal measure. Interested parties can refer to the seminar website at <https://ser-rse-bridge.github.io/> for related resources and activities.

## ■ References

- 1 J. Cohen, D. S. Katz, M. Barker, N. Chue Hong, R. Haines, and C. Jay, “The Four Pillars of Research Software Engineering,” *IEEE Software*, vol. 38, no. 1, pp. 97–105, Jan. 2021, doi: 10.1109/MS.2020.2973362.
- 2 M. Felderer, M. Goedicke, L. Grunske, W. Hasselbring, A.-L. Lamprecht, and B. Rumpe, “Toward Research Software Engineering Research,” Zenodo, Jun. 2023. doi: 10.5281/zenodo.8020525.
- 3 P. Bourque and R. E. Fairley, Eds., *Guide to the Software Engineering Body of Knowledge*, Version 3.0. IEEE, 2014.

## 6.25 Hardware Support for Cloud Database Systems in the Post-Moore's Law Era

**Editors:** David F. Bacon, Carsten Binnig, David Patterson, and Margo Seltzer  
**Seminar No. 24162**

Date: April 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.54

© Creative Commons BY 4.0 license

© David F. Bacon, Carsten Binnig, David Patterson, and Margo Seltzer



**Participants:** Anastasia Ailamaki, Gustavo Alonso, David F. Bacon, Lawrence Benson, Carsten Binnig, Alexander Böhm, Helena Caminal, Yannis Chronis, Holger Fröning, Jana Giceva, Mark D. Hill, Ihab Francis Ilyas, Zsolt Istvan, Lana Josipovic, Tim Kraska, Justin Levandoski, Jignesh M. Patel, David A. Patterson, Holger Pirk, Tilmann Rabl, Eric Sedlar, Margo Seltzer, Pinar Tözün, Nandita Vijaykumar, Tianzheng Wang, Lisa Wu Wills, Tobias Ziegler

This Dagstuhl Seminar on the Future of Cloud Database Systems was convened to address the pressing challenges arising from the stagnation in hardware performance gains, historically driven by Moore's and Dennard's laws. As data continues to grow exponentially – propelled by the expansion of autonomous systems, the Internet of Things (IoT), and machine learning – there is an urgent need to rethink the co-design of database systems and hardware. This seminar brought together experts from database systems, hardware architecture, and storage systems to explore innovative approaches to overcoming these scalability bottlenecks and envisioning the future of cloud database systems.

A central theme of the seminar was the growing disconnect between the exponential increase in data and the slowing pace of hardware improvements, leading to what participants referred to as a “scalability wall.” Addressing this challenge requires groundbreaking architectural changes in cloud database systems to support the next generation of applications. One significant area of focus was the potential role of AI-driven hardware and software in reshaping database management systems (DBMS). Participants explored whether AI hardware, such as GPUs and TPUs, could be adapted for database workloads, which traditionally are not compute-bound. Additionally, the concept of leveraging large language models (LLMs) as a new paradigm for databases was discussed, prompting further considerations of the future interplay between AI and DBMS.

To kickstart these discussions, several invited impulse talks were presented, each designed to set the stage for the working groups by exploring possible future scenarios for cloud database systems:

1. **AI Rules:** This talk examined a future where AI hardware and software dominate data centers, fundamentally altering the design and function of DBMS. The discussion centered on how DBMSs might need to evolve in a world where AI is integral to data processing and whether an LLM could serve as a database.

2. **A Disaggregated Future:** This presentations offered a perspective on a future where heterogeneous devices (compute, memory, storage) are connected via ultra-fast networks, creating a fully disaggregated cloud infrastructure. The talk prompted discussions on how DBMS could adapt to and thrive in such an environment.
3. **A Fully Reprogrammable Future:** The talk on this future envisioned a future where all hardware is reprogrammable and customizable at runtime, drastically changing how data processing and storage are handled. The implications for DBMS in such a highly flexible hardware environment were critically examined.
4. **The Pipe Dream:** This session explored the idea of “dreaming up” new DBMS hardware, revisiting the concept of a dedicated database machine. The discussion focused on whether this approach, which has failed in the past, could succeed in the context of modern cloud environments.

Following these impulse talks, the seminar divided into working groups to delve deeper into specific challenges:

1. **Working Group 1:** The Next Order of Magnitude focused on how database technologies can evolve to achieve order-of-magnitude improvements in performance, despite the slowdown in hardware advancements. This group was particularly concerned with managing the exponential growth of unstructured data feeding machine learning models.
2. **Working Group 2:** Memory-Centric DBMS Design advocated for a shift from processor-centric to memory-centric designs, emphasizing the optimization of data access in cloud environments as a solution to the performance bottlenecks caused by traditional architectural models.
3. **Working Group 3:** AI Hardware for Databases investigated how emerging AI hardware, like GPUs and TPUs, could be leveraged for cloud DBMS, even though database workloads

typically do not benefit as much from compute-bound acceleration as other applications do.

4. **Working Group 4:** The last working group explored taking disaggregation to the extreme and considering its impact on systems for cloud DBMSs.

As the seminar progressed, participants emphasized the importance of cross-disciplinary collaboration and knowledge sharing. They worked together to draft a comprehensive paper for publication, summarizing the insights and innovations discussed. The seminar concluded with a focus on the need for continued

innovation in both hardware and software to meet the demands of future cloud database systems.

In summary, the Dagstuhl Seminar provided a crucial platform for reimagining the future of cloud database systems in light of hardware stagnation. By bringing together leading experts from multiple disciplines and sparking deep discussions through targeted impulse talks, the seminar laid the groundwork for the architectural and system-level innovations necessary to overcome the scalability challenges posed by exponential data growth. The insights and collaborative efforts from this seminar will be instrumental in guiding the development of next-generation database systems.

## 6.26 Automated Synthesis: Functional, Reactive and Beyond

**Editors:** S. Akshay, Bernd Finkbeiner, Kuldeep S. Meel, and Ruzica Piskac

**Seminar No.** 24171

Date: April 21–26, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.85

© Creative Commons BY 4.0 license

© S. Akshay, Bernd Finkbeiner, Kuldeep S. Meel, and Ruzica Piskac



**Participants:** S. Akshay, Ashwani Anand, A. R. Balasubramanian, Suguman Bansal, Katrine Bjørner, José Cambronero, Supratik Chakraborty, Deepak D'Souza, Alexis de Colnet, Rayna Dimitrova, Rüdiger Ehlers, Johannes Klaus Fichte, Bernd Finkbeiner, Dror Fried, Mikoláš Janota, Jie-Hong Roland Jiang, Ayrat Khalimov, Alfons Laarman, Benedikt Maderbacher, Pierre Marquis, Kuldeep S. Meel, Tobias Meggendorfer, Jingyi Mei, Guillermo A. Pérez, Ruzica Piskac, Govind Rajanbabu, Subhajit Roy, Mark Santolucito, Ute Schmid, Sanjit A. Seshia, Shetal Shah, Arijit Shaw, Friedrich Slivovsky, Mate Soos, K. S. Thejaswini, Hazem Torfah, Shufang Zhu

In Dagstuhl Seminar 24171, we brought together researchers working in various aspects of automated functional synthesis. This diverse topic encompasses areas ranging from Boolean variants to quantified variants, automated reasoning for general theories, program synthesis, and more. One particular focus was on finding synergies between functional and reactive synthesis communities and investigating the deep connections between these two areas.

On the first day, we started with two introductory tutorials: one on Boolean functional synthesis and another on reactive synthesis, setting the agenda for the entire seminar. This was succeeded by technical presentations on definability and dependency in quantified Boolean formulas. The second day included a tutorial on automated reasoning and synthesis, with an emphasis on theories extending beyond Boolean (e.g., SMT), followed by discussions on quantitative properties. On the third day, we organized a special session with other tool competition organizers to assess the feasibility of a competition or track dedicated to functional synthesis.

The remaining days were filled with diverse technical talks that fell into two categories. The first category included talks that delved deeper into specific aspects of functional synthesis, reactive/LTL synthesis, and specific problems within these fields. The second category introduced new applications or connections, such as quantum applications and functional programming. Discussions during and beyond these talks were further explored in different open and problem sessions. Some of the identified and discussed problems were:

1. How to formalize the Boolean functional synthesis problem at the heart of reactive synthesis? Various problem formulations were discussed, and some benchmarks were created.
2. Can we go beyond Boolean theories and synthesize programs and functions for general SMT? What bottlenecks do we face?
3. How can we find synergy between automated functional synthesis and synthesis using transformers? Specifically,

what is the meeting ground between machine learning and inductive program synthesis techniques, functional synthesis, and automated reasoning?

4. Can the successful lens of knowledge representations and compilations for model counting and Boolean functional synthesis be extended to other settings?
5. Can we synthesize quantum circuits from specifications, thus leading to a theory of automated reasoning for quantum systems?
6. Can reactive synthesis over finite traces utilize techniques developed in automated functional synthesis?

These were among the prominent topics discussed, but the list is by no means exhaustive. Several bottlenecks were identified, such as the need for growth within the community developing these tools before establishing a proper competition. Additionally, there was a recognized necessity for broader and more extensive discussions on benchmarks.

Overall, the seminar fostered a collaborative spirit among theoreticians, tool developers, and experts across different aspects of automated functional synthesis. The seminar was also attended by a large number of early career researchers, postdoctoral fellows, and graduate students who also participated enthusiastically throughout the seminar. The shared optimism generated during this seminar has laid a strong foundation for future advancements. We advocate for the continuation of these valuable discussions and propose organizing further meetings of a similar nature to build on the momentum gained and to explore new frontiers in automated functional synthesis.

In the full report, we provide the abstracts of all the talks, as well as discussion sessions held during the seminar. We thank all the speakers and attendees for their active participation and look forward to attending and organizing more such events in the future.

## 6.27 Code Search

**Editors: Satish Chandra, Michael Pradel, and Kathryn T. Stolee**  
**Seminar No. 24172**


Date: April 21–24, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.108

© Creative Commons BY 4.0 license

© Kathryn T. Stolee, Satish Chandra, and Michael Pradel

**Participants:** Boris Bokowski, José Cambroner, Satish Chandra, Jürgen Cito, Luca Di Grazia, Elena Leah Glassman, Georgios Gousios, Reid Holmes, Ciera Jaspan, Tobias Kiecker, Dongsun Kim, Miryung Kim, Jens Krinke, Julia Lawall, Gabriel Matute, Alexander Neubeck, Michael Pradel, Nikitha Rao, Kathryn T. Stolee, Christoph Treude, Jan Van den Bussche, Rijnard van Tonder, Bogdan Vasilescu, Cristina Videira Lopes, Tobias Welp, Bowen Xu, Svetlana Zemlyanskaya



The 3-day Dagstuhl Seminar on “Code Search” brought together leading experts from academia and industry to discuss and advance the field of code search. This seminar highlighted the critical role of code search in various software engineering activities, from locating where an error was thrown to learning new APIs or programming languages. It also emphasized the importance of search in automated software engineering tasks like automated program repair, code recommendation, and clone detection. The emergence of generative AI tools, which offer alternative methods for finding and reusing code, was also a significant topic of discussion.

Participants explored the implications of code search research on developer productivity, code quality, and software engineering ethics. They examined the diverse tools available for code

search, ranging from internal company tools to open-source platforms like GitHub, and generative AI tools like ChatGPT. The seminar addressed various dimensions of code search, such as appropriate scope for search results, indexing methodologies, and combinations of code search and LLMs, e.g., in the form of retrieval-augmented generation.

In addition to talks and informal discussions, there were several break-out sessions during which participants discussed specific topics in smaller groups and eventually reported back to the other participants. Sections 4.1 of the full report provides an overview of the breakout sessions.

As a result of the seminar, several participants plan to launch various follow-up activities, such as joint publications and transferring promising ideas from academia to industry.

## 6.28 Computational Metabolomics: Towards Molecules, Models, and their Meaning

**Editors: Timothy M. D. Ebbels, Soha Hassoun, Ewy A. Mathé, and Justin J. J. van der Hoof**  
**Seminar No. 24181**

Date: April 28 – May 3, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.124

© Creative Commons BY 4.0 license

© Timothy M. D. Ebbels, Soha Hassoun, Ewy A. Mathé, and Justin J. J. van der Hoof



**Participants:** Wout Bittremieux, Sebastian Böcker, Carl Brunius, Roman Bushuiev, Haley Chatelaine, Ronan Daly, Niek de Jonge, Kai Dührkop, Timothy M. D. Ebbels, Soha Hassoun, Florian Huber, Pär Jonsson, Purva Kulkarni, Jessica Lasky-Su, Alice Limonciel, Liping Liu, Tytus Mak, Ewy A. Mathé, Hosein Mohimani, María Eugenia Monge, Steffen Neumann, Louis-Felix Nothias, Daniel Raftery, Raphael Reher, Stacey N. Reinke, Hannes Röst, Juho Rousu, Robin Schmid, Emma Schymanski, Denise Slenter, Jan Stanstrup, Michael Andrej Stravs, Marynka Ulaszewska-Tarantino, Justin J. J. van der Hoof, Dries Verdegem, Juan Antonio Vizcaino, Ralf Weber, Egon Willighagen, David Wishart, Michael Anton Witting, Nicola Zamboni

Metabolomics is the study of the small molecule composition of biological systems. These small molecules define biological functions, making metabolomics a broadly applied technology in the biomedical, environmental, and biotechnology fields of study. Metabolite measurements are typically produced using mass spectrometry (MS), usually coupled with liquid or gas chromatography (LC or GC), and/or nuclear magnetic resonance (NMR) spectroscopy. New technologies continue to be developed, leading to increased resolution of metabolite species detected, as well as increased sensitivity. The resulting datasets are high throughput and typically yield abundances of thousands of small chemical structures. For these reasons, the field of computational metabolomics continues to grow to address current and imminent issues in data stewardship, processing, analysis and interpretation.

This Dagstuhl Seminar is the 5th seminar in the computational metabolomics series. Previous seminars have addressed key topics in the field. These include leveraging spectral data to annotate and identify routinely measured metabolites, assessing interactions between metabolites and proteins through “metabo-proteomic assays”, as well as implementing multi-omic analyses and interpreting these data through enrichment analyses. This year, we not only dug deeper into the most current and relevant topics but also introduced new topics to the series. Our overall goal was to explore how to improve the utility of metabolomics data and its scientific relevance across many disciplines, alone or in combination with other data types, by leveraging machine learning and deep learning (ML/DL). To accomplish this, our seminar was organized into four categories. The first category was education, which was a newly introduced topic for this edition. Participants recognized the need for resources linking out to available education and training materials and discussed the inherent challenges of teaching multi-disciplinary topics, such as computational metabolomics. Our second category was molecules,

which includes annotations and measurements of metabolites and molecules they associate with, such as proteins, genes, and other metabolites. New areas of emphasis included representation and classification of lipids, polymers and multi-constituent substances, and use of electron-activated disassociation methods for data collection. Our third category was models, which encompasses data quality, uncertainty in annotating metabolites, and relationships between metabolites and other molecules. Sessions in this category were the most prominent and included novel areas of repository-scale analyses, simulations of metabolomic data, and automation of data analysis workflows. Of note, practical sessions on how to best measure scientific impact and how/where to submit data publicly to increase utility of data and ability to develop new computational methods were held. Lastly, our fourth category was meaning, which represents resources, methods and tools that enable visualization and interpretation of large-scale metabolomic data in the context of biological, environmental or other sciences. This included the use and misuse of molecular networking and its current applications in the field. This year, new focus was placed on building interpretable models and leveraging increasingly available information on metabolite annotations, such as pathways and reactions. Of note, recent developments in large language model techniques were discussed throughout our categories, with a special session dedicated to prototyping an LLM with metabolomics-specific content that can be used for training the next generation of metabolomics experts.

As in previous years, and based on positive feedback, the seminar format was again flexible, and topics were finalized on the first day of the seminar. The audience was encouraged to bring forth topics, and attention was paid to rotate who moderated and took notes for the sessions. Moderators, as well as organizers, actively ensured that all participants had a voice in topic sessions. Due to the large number of topics and to keep the groups manageable in size, parallel sessions were held. At the end of

each day, the last group meeting was held to summarize the day's discussion and to finalize the planning for the following day. One new aspect this year was the set-up of a Slack workspace that was used prior to, during, and after the seminar. This workspace facilitated communication and information exchange between participants. Overall, this seminar was highly successful, and participants were highly engaged. Topics and discussions

generated much enthusiasm and concrete next steps for potential collaborations. The field of computational metabolomics is a very active field of study that is somewhat underrepresented in many of the main metabolomics conferences, especially when it comes to tool development. The opportunity to focus on the computational aspects was very well received and will surely continue to grow as the data generation and types are ever increasing.

## 6.29 Resilience and Antifragility of Autonomous Systems

Editors: Simon Burton, Radu Calinescu, and Raffaella Mirandola

Seminar No. 24182

Date: April 28 – May 3, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.4.142

© Creative Commons BY 4.0 license

© Simon Burton, Radu Calinescu, and Raffaella Mirandola



**Participants:** Lee Barford, Amel Bennaceur, Simon Burton, Radu Calinescu, Matteo Camilli, Marc Carwehl, Ana Cavalcanti, Felicita Di Giandomenico, Ada Diaconescu, Kerstin I. Eder, Mario Gleirscher, Vincenzo Grassi, Sebastian Hahner, Andreas Heyl, Antje Loyal, Ravi Mangal, Lina Marsso, Raffaella Mirandola, Gabriel Moreno, Elena Navarro, Shiva Nejati, Diego Perez-Palacin, Ralf H. Reussner, Patrizia Scandurra, Catia Trubiani, Sebastián Uchitel, Gricel Vázquez

The increasing complexity in the environment, tasks, and technology related to autonomous systems results in limitations in the statements that can be made regarding dependability during design time. In particular, these systems may operate within environments for which only incomplete models exist, that may change over time or may be subject to unforeseen interactions and disturbances. As a result, such systems must be engineered to be trustworthy despite residual insufficiencies in their design, and in the presence of unexpected events due to their dynamically evolving operating context.

Related domains concerned with system autonomy in uncertain environments have already taken inspiration from nature to endow artificial systems with self-\* properties (e.g. self-optimisation, -repair, -protection, -configuration, and -adaptation). Such self-\* capabilities enable systems to improve their performance and dependability at runtime while reducing the need for low-level human intervention – properties that are closely related to resilience and antifragility.

This Dagstuhl Seminar aimed to unify the international research on **resilient and antifragile autonomous systems** (RAAS), leading to faster scientific advancements and industrial adoption. To this end, the seminar brought together leading researchers and practitioners with expertise in autonomous system resilience, antifragility, safety, and ethics, from disciplines including computer science, safety science, and ethics, to share and discuss each other’s understanding of, methods for, and open challenges related to RAAS. Initial presentations were used to set the scene by proposing basic definitions, industry perspectives, and engineering views on cyber-resilience. These were followed by group and plenary discussions to explore these concepts in more detail.

A clear set of agreed definitions is essential in order to make progress as a community in this area. *Resilience* can be broadly seen as the ability to absorb disturbances and unexpected

events whilst maintaining essential properties of the system. Using such conditions to harden the system against future events can be viewed as *antifragility*. These definitions highlight that antifragility is a concept referring to systems designed to operate under “open-world” assumptions, where the responsibility of maintaining a given property, despite disturbances (resilience) mostly shifts from design time to runtime, and relies on the presence in the system of some suitable degree of autonomy (self-\* capability). As such, antifragility can be viewed as the ability of a system to *self-improve* its resilience over (run)time. Discussions converged to the idea that in order to define resilience and antifragility, we should build on the work of Control Theory, specifically how systems recover from (potentially previously unknown) disturbances. Thus, we postulated that both resilience and antifragility should be defined over the metrics of settling time, percentage of settling, percentage of overshoot, and percentage of overshoot with respect to the properties of interest in the event of disturbances to the system. Discussions on how to use formal methods to construct systems that guarantee these desired properties generated many challenging questions that are to be followed up in future research.

Initial work in the seminar explored more precise definitions of RAAS that also included the consideration of uncertainty and causality, and where a collection of properties may need to be optimised as a whole. Such trade-offs are particularly evident when considering safety, ethical, and legal aspects of RAAS. In some cases, autonomous systems must remain operational in order to stay safe. A resilient system could remain within its safety bounds when disrupted, whilst maintaining a minimal level of utility. An antifragile system could use repeated disturbances to lower risk over time whilst increasing overall utility. Similar trade-offs and optimisations will be found when considering legal and ethical concerns for RAAS and these could lead to specific technical requirements on the system. For example, for a system

that adapts its function over time, avoiding the loss of agency in human stakeholders needs to be ensured.

Engineering antifragile systems requires specialised consideration in each phase of the traditional software and system development process. This includes requirements, design, implementation, and testing. Artificial Intelligence (AI) – in terms of machine learning, symbolic AI techniques, and combinations thereof – has the potential to provide a basis for both recognising disturbances and deciding the system adaptations needed to mitigate these disturbances. The seminar participants see potential for AI to be used in all phases of the MAPE-K (monitor-analyse-plan-execute supported by knowledge) cycle of self-adaptive systems. Furthermore, a control-theoretic reasoning approach could be used to verify whether a particular adaptation manager pushes the resilience error (i.e., the difference between observed and preferred resilience) below some threshold, or whether the resilience level stabilises at a reference value.

The seminar concluded that much work is still required to advance research in the area of RAAS, and to foster RAAS adoption in industrial applications. This includes:

- Agreeing on terminology and definitions that build upon and extend our traditional understanding of dependable systems;
- Formally defining metrics for resilience and antifragility that can be used to design and verify RAAS;
- Engineering methods and candidate technologies for implementing RAAS;
- Considering the safety, legal, and ethical implications of RAAS, including both their positive potential and their associated risks.

The participants agreed to pursue these important and challenging issues in future collaborations, including joint publications, workshops, and journal special issues.

## 6.30 Generalization by People and Machines

**Editors: Barbara Hammer, Filip Ilievski, Sascha Saralajew, and Frank van Harmelen**

**Seminar No. 24192**

Date: May 5–8, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.5.1

© Creative Commons BY 4.0 license

© Filip Ilievski and Sascha Saralajew



**Participants:** Wael Abd-Almageed, Michael Biehl, Marianna Marcella Bolognesi, Xin Luna Dong, Kenneth D. Forbus, Kiril Gashteovski, Barbara Hammer, Pascal Hitzler, Filip Ilievski, Giuseppe Marra, Pasquale Minervini, Martin Mundt, Axel-Cyrille Ngonga Ngomo, Alessandro Oltramari, Benjamin Paaßen, Gabriella Pasi, Sascha Saralajew, Zeynep G. Saribatur, Ute Schmid, Luciano Serafini, Dafna Shahaf, John Shawe-Taylor, Vered Shwartz, Gabriella Skitalinska, Clemens Stachl, Gido van de Ven, Frank van Harmelen, Thomas Villmann, Piek Vossen, Michael R. Waldmann

The Dagstuhl Seminar consisted of

1. lightning talks, where each participant had 2min for a short introduction and the presentation of a motivating (funny) example of generalization,
2. perspective pitches, where invited researchers from different domains gave a short talk about generalization in their domain (15min talk and 15min discussion),
3. daily discussion breakout sessions, where researchers organized in groups to discuss aspects of generalization and to work on the joint perspectives paper, and
4. plenary sessions, where we discussed the progress and results of the different breakout groups and organizing question around the seminar.

Day 1 featured an introductory session by the organizers and the lightning talks. During day 1, there were two perspective pitches on generalization from the angle of analogy (by *Ken Forbus*) and knowledge representation in symbolic AI (by *Luciano Serafini*). In the afternoon of day 1, the participants discussed generalization in four working groups:

1. types of generalization,
2. methods of generalization,
3. evaluation of generalization, and
4. human-AI teaming.

All teams were comprised of participants with diverse background and interests. The formation of these four groups was informed by a poll on possible topics that was sent to the participants before the seminar, filled by nearly all participants. At the end of day 1, each group provided updates centered around three questions:

1. How is it done today?
2. How well are we doing?
3. What are open challenges and important future directions?

Day 1 ended up with a plenary session during which each of the groups reported on their initial ideas, and received feedback from the other participants.

Day 2 featured four perspective pitches, highlighting the angles of statistical physics (by *Michael Biehl*), cognitive science (by *Ute Schmid*), computational linguistics (by *Vered Shwartz*), and computer vision (by *Wael AbdAlmageed*). In the afternoon, the participants split into the same four working groups as in day 1, with an instruction to organize the list of considerations from day 1. A key goal was to narrow down the scope of each working group and to identify important points to focus on. Day 2 ended with a debrief by the breakout sessions, during which common aspects emerged in different groups.

To facilitate a fruitful end of the seminar, the organizers came up with a set of 4 pillars that each of the groups was supposed to organize their content around, during day 3. These included: theory, context, representation, and foundational models. On day 3, each group provided an attempt to organize their content into these four pillars to the extent possible. Day 3 (half a day) ended with a discussion on the next steps, with a specific goal of writing a joint agenda-setting paper with all participants, targeted at a prestigious venue.

In the meantime, the organizers and Prof. Ute Schmid formed an editorial team that has been leading the process of writing the perspectives paper, and the participants provided two versions of write-up from their group: a long version and a short version. The short versions are limited to 2-3 pages and 20-25 citations, to conform jointly with the restrictions of journals like *Nature Machine Intelligence*. At the time of writing, the editorial team is busy with preparing this submission, with another round of feedback and collaboration scheduled with the participants in August.

## 6.31 Discrete Algorithms on Modern and Emerging Compute Infrastructure

6

**Editors: Kathrin Hanauer, Uwe Naumann, Alex Pothen, and Robert Schreiber**  
**Seminar No. 24201**

Date: May 12–17, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.5.12

© Creative Commons BY 4.0 license

© Kathrin Hanauer, Uwe Naumann, Alex Pothen, and Robert Schreiber

**Participants:** Cevdet Aykanat, David A. Bader, Paolo Bientinesi, Erik Boman, Martin Bücken, Jason Charlesworth, Cedric Chevalier, Alex Crane, Jakob Engblom, John Gilbert, Chris Goodyer, Oded Green, Inge Li Gørtz, Kathrin Hanauer, Bruce Hendrickson, Paul D. Hovland, Nidhi Kaihnsa, Gero Kauerauf, Neil Kichler, Daniel Král', Johannes Langguth, Klaus Leppkes, Quanquan C. Liu, Johannes Lotz, Fredrik Manne, Henning Meyerhenke, Uwe Naumann, Manuel Penschuck, Alex Pothen, Arash Pourdamghani, Henrik Reinstädtler, Eleanor Rieffel, Ilya Safro, Robert Schreiber, Christian Schulz, Angelika Schwarz, George Slota, Sergii Strelchuk, Blair D. Sullivan, Sivan Toledo, Bora Uçar, Nate Veldt, Helen Xu



We are happy to report on a lively and productive scientific discourse on discrete algorithms on modern and emerging compute infrastructure. As always, Dagstuhl presented an ideal setting for this kind of interdisciplinary meeting of experts from diverse backgrounds.

The aim was to identify requirements for

1. discrete algorithms imposed by emerging compute infrastructure;
2. emerging compute infrastructure imposed by discrete algorithms;
3. curricula at universities aiming to educate the next generation of designers of novel discrete algorithms as well as of future compute infrastructure.

We focused on sparse linear algebra and graph algorithms while reaching out to a diverse set of representatives from industry combining expertise in modern accelerators, next-generation silicon, and quantum computing.

Research questions addressed included the following:

1. How should today's discrete algorithms be re-designed in order to meet the requirements of emerging compute infrastructure?
  - a. Can lessons learned while mapping discrete algorithms onto modern compute infrastructure be (partially) generalized for emerging compute infrastructure?

- b. What are implications for (combinations of) deterministic, stochastic, and data-driven methods?
  - c. What impact on the design of discrete algorithms and their implementation can be expected from likely hierarchy / heterogeneity in emerging compute infrastructure?
2. How can emerging compute infrastructure be tailored towards the needs of practically relevant discrete problems and their algorithmic solution?
  - a. How do we support irregularity and dynamics inherent in sparse linear algebra and graph problems by suitable hardware architecture / system software?
  - b. What do suitable programming models / languages look like?
  - c. How could we account for memory-boundedness?

Following individual 5 min introductions, the program consisted of three plenary talks, four plenary focus sessions / panel discussions, and four non-plenary working groups. Refer to the respective abstracts for details.

## 6.32 Causal Inference for Spatial Data Analytics

**Editors: Martin Tomko and Yanan Xin**  
**Seminar No. 24202**

Date: May 12–17, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.5.25

© Creative Commons BY 4.0 license

© Martin Tomko and Yanan Xin



**Participants:** Kevin Credit, Cécile de Bézenac, Simon Dirmeier, Andreas Gerhardus, Totte Harinen, Dominik Janzing, Urmi Ninad, Markus Reichstein, Katerina Schindlerova, Martin Tomko, Jonas Wahl, Jianwu Wang, Levi John Wolf, Yanan Xin, Shu Yang, Andrew Zammit Mangion

Spatial data analytics has undergone a revolution in recent years due to the availability of large, observational spatial datasets and advances in spatially-explicit statistical analysis as well as in machine learning. Despite these improvements, the current spatial data analysis methods primarily center on exploratory, descriptive, and predictive modeling that are grounded in correlational analysis. These approaches fall short of being able to quantify (and sometimes even identify) causal relationships. However, there has been an increasing interest in identifying and quantifying causal relationships in spatial data which are key to designing effective policy interventions in critical applications such as environmental and population science, climate science, epidemiology, urban planning, and traffic management.

Causal inference has been an active field of study in statistics and philosophy for some time. It recently gained traction in the machine learning community as a promising method for enabling more intelligent AI capable of causal reasoning. Yet, the application of existing causal inference methods to the spatial domain is not straightforward, and a theoretical and methodological foundation for spatial causal analysis is in its infancy. Spatial effects, such as spatial dependence and spatial heterogeneity, violate the fundamental assumptions of current causal inference frameworks. Besides, the large sample size, high dimensionality (space, time, attributes), and dynamic properties of spatio-temporal data also pose great challenges in inferring causal effects. Thus, there is a pressing need to accelerate the theoretical development in the field of spatial causal inference and enable a broader adoption of the methodological approaches supported by a well-grounded analytical toolset. Researchers in environmental sciences, spatial econometrics, spatial statistics, theoretical GIScience, and computing/machine learning communities have started making significant, yet thus far disparate efforts contributing to the foundations of spatial causal inference. This lack of interdisciplinary exchange of ideas and a comprehensive

understanding of the potential applications and limitations of spatial causal inference hinders progress across these disciplines.

As machine learning rapidly penetrates various spatial decision-making processes, the time is right to enable cross-discipline conversations around spatial causal inference, and thus maximize the impact of sound methodologies. As AI becomes widely applied to spatial data analysis supporting planning and policy-making, it is imperative to develop approaches that are interpretable, grounded, robust, and responsible. Enabling the conversations between theoretical, computational, and domain experts who are active in causal inference and its application for spatio-temporal systems will accelerate the development of more intelligent and responsible AI for spatial decision-making.

This seminar is convened to initiate conversations across disciplines on these critical questions around spatial causal inference. This five-day seminar covers topics on the definitions and theories of spatial causal inference, methodologies and applications, software and benchmark datasets, and open questions. A detailed program of the seminar is provided in Figure 6.8. A summary of the daily discussions is shown below.

- **Unified Definitions of Spatial Causal Inference.** The discussion focused on the specification of the spatial component in the causal inference process, covering topics on the formalization of spatial causal inference questions, representations (e.g., Spatial DAG), modeling approaches, and practical relevance.
- **Methodological Challenges and Solutions.** Methodological challenges were demonstrated through case studies in environmental science, transportation, advertisement and recommendations, and other social science applications. Based on these case studies, the group explored methods and ideas for modeling spatial confounding, spatial interference, spatial treatments, and evaluation of spatial causal analysis.

- **Open-Source Software and Benchmarks.** The session featured demos of the open-source Python packages `causalml` and `tigramite`. Following the demonstrations, the group dived into the discussions of casual discovery evaluations and establishing benchmarks for spatial causal inference.

- **Open Questions and the Road Ahead.** The group proposed key research questions in the field of spatial causal inference and identified interests for continued collaborations on these topics.

	Day 1	Day 2	Day 3	Day 4	Day 5
Morning	<ul style="list-style-type: none"> <li>• Welcome, programs, short introduction</li> <li>• Definitions and theories</li> </ul>	<ul style="list-style-type: none"> <li>• Presentations of case studies</li> </ul>	<ul style="list-style-type: none"> <li>• Continued discussions on methodologies</li> </ul>	<ul style="list-style-type: none"> <li>• Demos of open-source packages</li> </ul>	<ul style="list-style-type: none"> <li>• Summary of the seminar</li> <li>• Plan for post-seminar activities</li> </ul>
Afternoon	<ul style="list-style-type: none"> <li>• Spatial causal inference vs discovery</li> <li>• Brainstorming: questions/concepts/open problems</li> </ul>	<ul style="list-style-type: none"> <li>• Group discussions on methodologies</li> </ul>	Hiking	<ul style="list-style-type: none"> <li>• Benchmarking and evaluation</li> <li>• Open questions</li> </ul>	Departure

Fig. 6.8 Program of Dagstuhl Seminar: Causal Inference for Spatial Data Analytics (May 12th – 17th, 2024).

As a major outcome of the seminar, key challenges and research questions were identified in the field, as outlined in Section 4.4.5 Open Questions and also detailed in the notes of our daily discussions of the full report. We hope these thoughts and ideas will inspire a broader research interest in spatial causal inference and continue the exchange across disciplines, as well as between academia and industry.

The seminar resulted in the desire to continue these discus-

sions in a series of workshops (the first to take place at ACM SIGSPATIAL 2024) and the need to establish a community ([spatial-causal.org](https://spatial-causal.org)).

In the full report we will first present the position statements prepared by seminar participants on their thoughts related to spatial causal inference. Next, detailed notes of our daily discussions are documented in the report.

## 6.33 Evaluation Perspectives of Recommender Systems: Driving Research and Education

**Editors:** Christine Bauer, Alan Said, and Eva Zangerle  
**Seminar No.** 24211

Date: May 20–24 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.5.58

© Creative Commons BY 4.0 license

© Christine Bauer, Alan Said, and Eva Zangerle



**Participants:** Gediminas Adomavicius, Andrea Barraza-Urbina, Joeran Beel, Toine Bogers, Peter Brusilovsky, Wanling Cai, Tommaso Di Noia, Kim Falk, Bart Goethals, Dietmar Jannach, Joseph Konstan, Maria Maistro, Julia Neidhardt, Denis Parra, Lorenzo Porcaro, Rodrygo Santos, Manel Slokom, Barry Smyth, Helma Torkamaan, Tobias Vente, Lukas Wegmeth, Eva Zangerle

Recommender systems (RS) have become essential tools in everyday life, efficiently helping users discover relevant, useful, and interesting items such as music tracks, movies, or social matches. RS identify the interests and preferences of individual users through explicit input or implicit information inferred from their interactions with the systems and tailor content and recommendations accordingly [13, 16].

Evaluation of RS requires attention at every phase of the system life cycle, including design, development, and continuous improvement during operation. High-quality evaluation is crucial for a system's success in practice. This evaluation can focus on the core performance of the system or encompass the entire context in which it is used [3, 7, 8, 10]. Research typically differentiates between system-centric and user-centric evaluation. System-centric evaluation examines algorithmic aspects, such as the predictive accuracy of recommender algorithms. In contrast, user-centric evaluation assesses the user's perspective, including perceived quality and user experience. Comprehensive evaluation must address both aspects since high predictive accuracy does not necessarily meet user expectations [12].

The topic of evaluation, with all its challenges, is currently very relevant and trending. The PERSPECTIVES workshops (organized at ACM RecSys 2021-2023 [11, 14, 15], co-organized by this seminar's organizers) were highly popular and attracted many participants. This interest is further evidenced by the special issue in ACM Transactions on Recommender Systems [1] on evaluation. Recent calls for more impactful RS research [5, 6, 9, 12] highlight that current evaluation practices are too narrow and may not be practically relevant. [4] advocate for more nuanced evaluation methods that meet industry demands. [9] argue that current practices are insufficient as they often overlook side effects or longitudinal impacts. A recent systematic literature study further reveals that current evaluation methods are limited in experiment design, dataset choice, and evaluation metrics [2].

This seminar on evaluation perspectives of RS brought together researchers and practitioners from diverse backgrounds.

It aimed to discuss current challenges and advance the ongoing discussion on RS evaluation. The seminar began with eight presentations addressing current challenges in evaluation. These talks initiated the general discussion and helped form groups around these topics. As a result, five working groups were established, each focusing on the following areas:

**Working Group 1: Theory of Evaluation** This group focused on the theoretical foundations of RS evaluation. They began by identifying the shortcomings of current evaluation practices and linking these issues to underlying theoretical principles. Key challenges discussed included the selection and configuration of evaluation metrics and the reporting of evaluation results. Section 4.1 of the full report outlines the challenges and theoretical perspectives identified in this group.

**Working Group 2: Fairness Evaluation** This group focused on exploring paradigms and practices for evaluating the fairness of RS. Given the specific nature of fairness metrics and evaluation requirements for different applications, fairness problems, and goals, the group proposed “best meta-practices”, a set of approaches to planning, executing, and communicating rigorous fairness evaluation scenarios. The group's outcome is documented in Section 4.2 of the full report.

**Working Group 3: Best-Practices for Offline Evaluations of Recommender Systems** This working group addressed the topic of offline evaluation, with a specific focus on identifying problems and best practices for this evaluation method. They concentrated on pinpointing the primary challenges related to reproducibility and methodology. Subsequently, they provided guidelines to address these challenges from various perspectives, including those of paper authors, reviewers, editors, and program chairs, as summarized in 4.3 of the full report.

**Working Group 4: Multistakeholder and Multi-method Evaluation** This group examined the challenges and complexities in evaluating multistakeholder scenarios, discussing the key aspects that must be considered in such a nuanced environment. Additionally, they explored the transition from theoretical evaluation frameworks to practical implementation. Section 4.4 of the full report outlines this work.

**Working Group 5: Evaluating the Long-Term Impact of Recommender Systems** This working group concentrated on the long-term perspective and impact of RS and their evaluation. This includes developing suitable long-term measures and conducting social and behavioral research to understand and facilitate aspects such as human behavior, long-term stakeholder goals, and corresponding metrics. Additionally, the group examined practical challenges when evaluating the long-term aspects and impact of RS. This work is presented in Section 4.5 of the full report.

## References

- Christine Bauer, Alan Said, and Eva Zangerle. Introduction to the special issue on perspectives on recommender systems evaluation. *ACM Transactions on Recommender Systems*, 2(1), mar 2024. URL <https://doi.org/10.1145/3648398>.
- Christine Bauer, Eva Zangerle, and Alan Said. Exploring the landscape of recommender systems evaluation: Practices and perspectives. *ACM Transactions on Recommender Systems*, 2(1), mar 2024. URL <https://doi.org/10.1145/3629170>.
- Joeran Beel, Stefan Langer, Marcel Genzmehr, Bela Gipp, Corinna Breiteringer, and Andreas Nürnberger. Research paper recommender system evaluation: a quantitative literature survey. In *Proceedings of the international workshop on reproducibility and replication in recommender systems evaluation*, pages 15–22, 2013.
- Patrick John Chia, Jacopo Tagliabue, Federico Bianchi, Chloe He, and Brian Ko. Beyond ndcg: behavioral testing of recommender systems with relict. In *Companion Proceedings of the Web Conference 2022*, pages 99–104, 2022.
- Dietmar Jannach and Christine Bauer. Escaping the McNamara Fallacy: Towards more impactful recommender systems research. *AI Magazine*, 41(4):79–95, December 2020. ISSN 2371-9621, 0738-4602.
- Paolo Cremonesi and Dietmar Jannach. Progress in recommender systems research: Crisis? what crisis? *AI Magazine*, 42(3):43–54, 2021.
- Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, jan 2004. ISSN 1046-8188. <https://doi.org/10.1145/963770.963772>.
- Dietmar Jannach, Oren Sar Shalom, and Joseph A Konstan. Towards more impactful recommender systems research. In *ImpactRS@ RecSys*, 2019.
- Gourab K Patro, Lorenzo Porcaro, Laura Mitchell, Qiyue Zhang, Meike Zehlike, and Nikhil Garg. Fair ranking: a critical review, challenges, and future directions. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*, pages 1929–1942, 2022.
- Alan Said, Domonkos Tikk, Klara Stumpf, Yue Shi, Martha A Larson, and Paolo Cremonesi. Recommender systems evaluation: A 3d benchmark. In *RUE@ RecSys*, pages 21–23, 2012.
- Alan Said, Eva Zangerle, and Christine Bauer, editors. *Third Workshop: Perspectives on the Evaluation of Recommender Systems (PERSPECTIVES 2023)*, RecSys '23, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400702419. URL <https://doi.org/10.1145/3604915.3608748>.
- Eva Zangerle and Christine Bauer. Evaluating recommender systems: survey and framework. *ACM Computing Surveys*, 55(8):1–38, 2022.
- Bo Xiao and Izak Benbasat. E-commerce product recommendation agents: Use, characteristics, and impact. *MIS quarterly*, pages 137–209, 2007.
- Eva Zangerle, Christine Bauer, and Alan Said, editors. *Perspectives on the Evaluation of Recommender Systems (PERSPECTIVES)*, RecSys '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384582. URL <https://doi.org/10.1145/3460231.3470929>.
- Eva Zangerle, Christine Bauer, and Alan Said, editors. *Second Workshop: Perspectives on the Evaluation of Recommender Systems (PERSPECTIVES 2022)*, RecSys '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392785. URL <https://doi.org/10.1145/3523227.3547408>.
- Francesco Ricci, Lior Rokach, and Bracha Shapira. *Recommender Systems Handbook*. Springer New York, NY, 3rd edition, 2022.

## 6.34 Classical-Quantum Synergies in the Theory and Practice of Quantum Error Correction

**Editors:** Carmen G. Almudéver, Leonid Pryadko, Valentin Savin, and Bane Vasic  
**Seminar No.** 24212

Date: May 20–23, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.5.173

© Creative Commons BY 4.0 license

© Carmen G. Almudéver, Leonid Pryadko, Valentin Savin, and Bane Vasic



**Participants:** Alexei Ashikhmin, Kenneth R. Brown, Michael Epping, Abdul Fatah, Omar Fawzi, Carmen G. Almudéver, Shayan Srinivasa Garani, Francisco García Herrero, Ashutosh Goswami, Mackenzie Hooper Shaw, Liang Jiang, Robert König, Ching-Yi Lai, Mehdi Mhalla, Ioana Moflic, Davide Orsucci, Alexandru Paler, Pavel Panteleev, Leonid Pryadko, Joseph M. Renes, Eleanor Rieffel, Laura Rodríguez Soriano, Joschka Roffe, Pradeep Sarvepalli, Valentin Savin, Emina Soljanin, Matt Steinberg, Barbara Terhal, Bane Vasic, Gilles Zémor

### ■ Background and Motivation: From Classical to Quantum Error Correction and Fault-Tolerance

A fundamental consequence of the mathematical theory of information laid down by Shannon, error correcting codes play a vital role in ensuring the integrity of data in systems exposed to noise or errors. Classical error correcting codes were crucial to the success of modern communications and data storage systems (from the Internet to mobile, satellite, and deep-space communications, and from disk to flash memory storage) and found applications in other areas, such as pattern recognition, group testing, cryptography, or fault-tolerant computing. Likewise, quantum error correcting codes are at the heart of all quantum information processing, from fault-tolerant quantum computing to reconciliation in quantum key distribution, quantum sensing, and reliable optical communications.

Computation in the presence of noise is a long-standing problem, going back to the 1950s and the celebrated works of von Neumann, Elias, Taylor, Kuznetsov, Winograd, Cowan, Dobrushin, Pippenger, and many others. The first attempt to apply general error correction techniques for the design of fault-tolerant computing systems is due to Elias (Computation in the presence of noise, 1962), and one of the first attempts to derive fundamental limits in fault-tolerant computing is due to Winograd and Cowan (Reliable computation in the presence of noise, 1963). These works focused on fault-tolerant classical (Boolean logic based) computation, prior to the advent of ultra-high reliability integrated circuits based on complementary metal-oxide-semiconductor (CMOS) technology, but they still inspire and resonate with current approaches to fault tolerance, *e.g.*, to support the ongoing miniaturization of the emerging data processing and storage devices (technology scaling). In parallel, the last years have seen significant advances in the field of quantum technologies, promising a disruptive impact in information and computing technologies. Basic requirements for quantum computation have been

demonstrated in various technologies, including semiconductor or superconductor materials, photons, trapped ions, etc. Nonetheless, for unleashing the full computational power that quantum computers can bring, a critical task is to protect the quantum computation from the inherent quantum noise. The discovery of quantum error correcting codes in the mid-90s paved the way to noise resilient quantum computation, developed through the works of Calderbank, Shor, Steane, Sloane, Gottesman, Knill, Kitaev, Freedman, Meyer, Preskill, and many others. The integration of quantum error correction (QEC) into the quantum computation led to the development of the fault-tolerant quantum computing framework, aimed at countering the effects of noise on stored quantum information, faulty quantum preparation, faulty quantum gates, and faulty measurements. Such an integration of QEC and fault-tolerance techniques in quantum computing systems is key to the development of a universal large-scale quantum computer, achieving its expected exceptional potential.

While classical and quantum error correction may be regarded as different paradigms, involving different ways of thinking and to a certain extent different research communities, it turns out that they are actually closely related. One may mention here the formalism of quantum stabilizer codes, allowing notably to move from a continuous to a discrete model for quantum error correction, among which of particular interest is the Calderbank-Shor-Steane (CSS) construction of a quantum code from a pair of orthogonal classical binary codes. CSS codes can be alternatively described as chain complexes involving three spaces, where the boundary operators are defined (up to a choice of bases) by the two orthogonal classical codes. This homological point of view is essentially the one adopted by topological constructions, where quantum codes are produced based on cellular decompositions of surfaces (*e.g.*, torus), or higher dimensional manifolds. In parallel, the powerful machinery of abstract homological algebra proved to be very efficient in providing new constructions of

quantum codes, among which of particular interest are codes with constant weight stabilizer generators, referred to as quantum low-density parity-check (qLDPC) codes. The class of qLDPC codes encompasses the above topological constructions, and is the only class of quantum codes known to contain families of codes with both constant non-zero rate and non-zero fault-tolerant error-correction threshold. It is also worth mentioning here the recent constructions of asymptotically good qLDPC codes (with constant rate and relative minimum distance), auguring for practical constructions with increased error correction capacity or reduced qubit overhead. However, unlike their classical counterparts, which are equipped with efficient message-passing decoding algorithms, qLDPC codes are difficult to decode. The decoding of a qLDPC code requires locating not a single most likely error, but the most likely equivalence class of mutually degenerate errors (degeneracy is an inherent characteristic of any qLDPC code), which tends to inhibit the convergence of message-passing algorithms designed for classical codes. Besides, it is also worth mentioning that the time budget available to perform a single error correction round varies with the quantum technology, but a first-order approximation is a period of hundreds of nanoseconds. Hardware implementations meeting such a time constraint will require massive parallel processing, which has to be enabled by both the structure of the quantum code and the decoding algorithm.

To tackle these challenges, this Dagstuhl Seminar aimed at promoting interactions among coding theorists, quantum physicists, mathematicians, and computer and hardware engineers, to discuss achievements, strategies, and remaining gaps in the integration of QEC and fault-tolerance techniques into practical quantum computers, towards a comprehensive and mutual understanding of theory and engineering practice.

#### ■ Topics Covered by the Seminar

**Classical and Quantum LDPC codes.** The quest for low-complexity decoders of classical LDPC codes has resulted to the emergence of soft-decision iterative message passing decoders, *e.g.*, based on belief-propagation (BP) or min-sum (MS) algorithms. In the quantum case, decoding a CSS qLDPC code boils down to decoding the two constituent classical LDPC codes (*e.g.*, assuming separate decoding of  $X$  and  $Z$  errors, which does not preclude taking into account the possible correlations between the two types of error). In homological terms, the goal of the decoder is to find the most likely chain (error) – or more specifically, the most likely class of chains – corresponding to a given boundary (syndrome), where two chains are equivalent if their sum is in the trivial homology class. Maximum-likelihood decoders exist for the toric code (yet, their complexity is too high for practical applications), but they are out of reach for arbitrary topological or qLDPC codes. Developing new approaches to accurate and hardware friendly decoding of quantum codes is a crossroad of theory and practice, and of classical and quantum coding. Presumably, classical-quantum synergies can provide meaningful insights to the theory and practice of qLDPC codes. There are many examples where the theory and practice of qLDPC codes may benefit from classical-quantum synergies, such as devising optimized constructions for short qLDPC codes, improving the decoding performance through modified message-passing or smart post-processing techniques, using knowledge of quantum trapping sets to cope with the code degeneracy, devising machine learning based decoding solutions, conceiving efficient decoding algorithms to exploit soft information on measurement errors, or developing codes and decoding algorithms amenable to single-shot error correction.

Particular challenges discussed during the seminar were broadly related to novel constructions of qLDPC codes and expanding properties of the associated graph, novel decoding algorithms for topological and qLDPC codes, including message-passing based decoding, tensor network decoding, and machine-learning based decoding, applications of quantum error correction in various areas as quantum computing or quantum networks, and the design of entanglement-assisted quantum codes.

**Fault-Tolerant Quantum Computation.** Quantum memory with a topological or, more generally, qLDPC stabilizer code can be implemented with repeated syndrome measurements, where errors are detected by the difference between syndromes measured in consecutive rounds. It is also worth noticing that a QEC with a sufficiently short syndrome measurement cycle is needed throughout the operation of a quantum computer, and measurement circuits have to be designed with fault-tolerance in mind, *e.g.*, to prevent a single error to spread on multiple qubits. More generally, when non-trivial gates are executed on the logical subspace, detection events have to be chosen for each particular circuit. The gate error for the hardware in use, as well as the specific choice of the circuit and of the detection events determines the error model and the structure of the quantum error-correcting code that has to be decoded. Pauli error channels associated with specific gates on specific qubits are most commonly used for decoding. Actual error probabilities may also depend on the parameters chosen for each qubit (*e.g.*, working frequencies chosen for individual qubits in the case of superconducting qubits), as well as variability of the manufacturing. Other important error types include non-Pauli errors (decay, unitary errors, etc.), as well as leakage from the computational subspace. Furthermore, with some hardware, syndrome measurement may contain additional soft information about the measurement outcome. Taking such information into account may dramatically improve the decoding accuracy. While in theoretical analysis such details can often be ignored, in practice, for a quantum computer operating close to the threshold, a relatively small improvement in the decoding accuracy can reduce the required overhead by orders of magnitude, or even be required to attain fault-tolerance.

Particular challenges discussed during the seminar were broadly related to a variety of Pauli error channels, including those derived from Clifford circuits with gate error models customized for specific hardware, related unification of decoding protocols for qubit-based codes, decoding using soft syndrome information, coherent noise and quantum error correction, subsystem and Floquet codes, effective consideration of geometric and connectivity requirements, fault-tolerant quantum computation, and fault-tolerant design of algorithms and protocols.

**From Noisy Intermediate Scale Devices to Large Scale Quantum Computing.** While QEC is the only presently known gateway to reap the benefits of computational quantum algorithms, a robust, scalable, and fully functional QEC technique that allows performing fault-tolerant quantum computations has not been demonstrated experimentally yet. Arguably, QEC is the only technology still lacking to realize a vision of useful large-scale quantum computation. However, there are already a few demonstrations of the potential to protect quantum information on noisy intermediate scale quantum (NISQ) processors based on superconducting qubits, such as: i) the experimental implementation of distance-3 surface code on the Zuchongzhi 2.1 superconducting quantum processor showing that by executing several consecutive error correction cycles, the logical error can be significantly reduced after applying

corrections (Realization of an Error-Correcting Surface Code with Superconducting Qubits); ii) the experimental demonstration that increasing the code distance leads to a better logical qubit performance using an expanded Sycamore device with 72 transmon qubits (Suppressing quantum errors by scaling a surface code logical qubit). NISQ technology may serve as a first step towards demonstrating a certain number of QEC protocols, suitable to the intermediate scale, but which in the long term may also have useful implications for large-scale quantum technologies. Yet, in a large-scale quantum computer, the QEC decoder design faces significant challenges, arising from the need to integrate various system constraints, such as accuracy, bandwidth, latency, power-consumption, or scalability. QEC decoders need to be powerful enough to accurately correct the quantum errors, fast enough to fight against the qubit decoherence, energy efficient to meet stringent power-consumption requirements, and highly scalable to meet the needs of fault-tolerance. Achieving all these constraints is extremely challenging, and might not be possible with existing solutions. Recent research has focused on the design of hardware architectures capable of efficiently accommodating QEC techniques, where considerations such as timing, latency, power, and wiring between the quantum chip and the QEC processor take a prominent place, as they are critical for creating a viable solution.

The main challenges discussed during the seminar ranged from low-qubit overhead fault-tolerant schemes and efficient implementation of small QEC on NISQ processors to scalable modular quantum computing architectures for quantum error

correction and large scale fault tolerance, while also considering software implementation of quality decoders, decoding architectures that lend themselves to high-speed and low energy consumption, and recent progress on the hardware implementation and prototyping of QEC decoders.

#### ■ Organization of the Seminar

The seminar brought together 30 participants, both senior and talented young researchers, from 13 countries (Denmark, Finland, France, Germany, Great Britain, India, Ireland, Netherlands, Russia, Switzerland, Spain, Taiwan, and the United States), with research expertise in relevant areas, *e.g.*, classical and quantum coding theory, hardware architectures and designs of error correcting codes, quantum information processing and software, fault-tolerant quantum computation and fault-tolerant design of algorithms and protocols, quantum technologies, and quantum computer architecture design.

The primary objective of the seminar was to foster an exchange of ideas on challenges faced by quantum error correction, evolving through presentations as well as discussions aimed at realizing the potential of a large community bring diverse viewpoints to the table. In order to facilitate this, the two and a half day program of the seminar comprised a series of 14 invited talks, organized in seven plenary sessions, as well as five time slots for breakout sessions, giving more time for discussions and the organisation of ad-hoc working groups (bringing together a large part of the participants). The full report includes the abstracts of all talks and three working groups.

## 6.35 Stochastic Games

**Editors:** Nathanaël Fijalkow, Jan Kretinsky, and Ann Nowé  
**Seminar No.** 24231

Date: June 2–7, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.1

© Creative Commons BY 4.0 license

© Nathanaël Fijalkow, Jan Kretinsky, and Ann Nowé

**Participants:** Marianne Akian, Xavier Allamigeon, Guy Avni, Christel Baier, Gabriel Bathie, Eline Bovy, Antonio Casares, Florent Delgrange, Dani Dorfman, Kousha Etessami, Nathanaël Fijalkow, Stéphane Gaubert, Kristoffer Arnsfelt Hansen, Sebastian Haslebacher, Sebastian Junges, Marcin Jurdzinski, Jan Kretinsky, Antonin Kucera, Karoliina Lehtinen, Georg Loho, James C. A. Main, Corto Mascle, Théo Matricon, Petr Novotný, Ann Nowé, David Parker, Guillermo A. Pérez, Jakob Piribauer, Roxana Radulescu, Sven Schewe, Guruprerana Shabadi, Aaron Sidford, Mateusz Skomra, Yann Strobecki, Marnix Suilen, K. S. Thejaswini, Patrick Totzke, Pierre Vandenhove, Muthukumar Vidya, Maximilian Weininger, Yinyu Ye



The Dagstuhl Seminar on Stochastic Games brought together leading researchers and practitioners in the field to discuss recent advances, challenges, and future directions. The seminar featured a series of tutorials, invited talks, and contributed talks, which provided a comprehensive overview of the latest developments in Markov decision processes, reinforcement learning, and stochastic game theory. Key results from the seminar include novel insights into branching stochastic games, the development of new algorithms for solving concurrent and population games, and advancements in the theoretical understanding of efficient solutions for Markov decision processes. The seminar fostered lively discussions during open problem sessions and working groups, culminating in a collaborative exploration of open questions and potential research directions.

Overview of the Invited Talks and Tutorials:

- Yinyu Ye: “Progresses and Open Questions on the Markov Decision / Game Process” – Discussed recent advancements and open problems in the field of Markov decision processes and games.

- Dave Parker: “PRISM-games” – Provided a tutorial on PRISM-games, a tool for modeling and analyzing probabilistic systems.
- Kousha Etessami: “Branching MDPs, Branching Stochastic Games, and Generalizations of Newton’s Method” – Explored branching Markov decision processes and games, and introduced generalizations of Newton’s method for these models.
- Aaron Sidford: “Theoretical Advances in Efficiently Solving Markov Decision Processes” – Highlighted recent theoretical progress in solving Markov decision processes more efficiently.
- Sven Schewe: “Automata for Profit and Pleasure” – Discussed the applications of automata theory in both practical and theoretical contexts.

Overall, the seminar was a highly productive event, advancing the collective understanding of stochastic games and fostering future research collaborations.

## 6.36 Designing Computers' Control Over Our Bodies

**Editors:** Nathanaël Fijalkow, Jan Kretinsky, and Ann Nowé

**Seminar No. 24232**

Date: June 2–7, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.1

© Creative Commons BY 4.0 license

© Nadia Bianchi-Berthouze, Mar Gonzalez-Franco, Florian 'Floyd' Mueller, and Misha Sra



**Participants:** Nadia Bianchi-Berthouze, Susanne Boll, Richard James Byrne, Arthur Caetano, Don Samitha Elvitigala, Masahiko Inami, Jarrod Knibbe, Per Ola Kristensson, Xiang Li, Zhuying Li, Joe Marshall, Louise Petersen Matjeka, Florian 'Floyd' Mueller, Minna Nygren, Rakesh Patibanda, Henning Pohl, Sara Price, Harald Reiterer, Aryan Saini, Oliver Schneider, Ambika Shahu, Misha Sra, Jürgen Steimle

**Remote Participants:** Mar Gonzalez-Franco

The Dagstuhl Seminar 24232, titled “*Designing Computers' Control Over Our Bodies*”, was held from June 2 to June 7, 2024, at Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany. This seminar brought together leading experts from diverse fields to explore a new paradigm in human-computer interaction (HCI) where technologies such as electrical muscle stimulation, galvanic vestibular stimulation, and exoskeletons enable computers to exert direct control over the human body [1–5]. The seminar addressed both the opportunities and challenges of this emerging domain, which extends beyond the traditional model of user-directed interaction to consider reciprocal roles where machines also assert control [3, 6, 7].

Over the course of five days, participants engaged in a variety of activities designed to foster interdisciplinary collaboration and develop a comprehensive understanding of this complex interaction paradigm. The seminar’s objectives included the following:

**Identifying Theoretical Gaps.** Participants explored the absence of a structured theoretical framework for articulating and evaluating experiences of being controlled by a machine. Discussions focused on such interactions’ ethical, psychological, and physical dimensions, highlighting the need for a multidisciplinary approach to understand and design for these new dynamics [8–10].

**Design Challenges and Opportunities.** The seminar investigated the design challenges associated with computational control over the human body, considering both the benefits, such as enhanced safety in autonomous systems and increased mobility, and the concerns, including issues of user agency, consent, and trust [11]. Hands-on sessions enabled participants to prototype and critique design concepts that address these complexities.

**Societal Implications.** The discussions extended to the societal impact of these technologies, examining potential applications and unintended consequences. Ethical considerations, particularly around the autonomy of users and the transparency of machine actions, were central to these debates [12]. The seminar aimed to outline guidelines for the responsible development and deployment of such technologies in various contexts, from healthcare to entertainment.

**Envisioning Future Directions.** A key focus of the seminar was envisioning the future landscape of HCI where computational control technologies become more prevalent. Speculative design sessions invited participants to imagine both utopian and dystopian scenarios for the year 2050, facilitating discussions on the potential trajectories of these technologies with- and without-AI, and their integration into daily life.

**Developing a Research Agenda.** The seminar concluded with the development of a research agenda to guide future work in this field. This agenda included key questions to address, potential methodologies for study, and proposals for new theoretical frameworks. Working groups were established to continue collaborative efforts beyond the seminar.

The seminar’s key outcomes included identifying several grand challenges in this emerging field, formulating initial design frameworks, and establishing a network of researchers and practitioners committed to advancing knowledge and practice in computational control technologies. The seminar highlighted the importance of interdisciplinary approaches to understanding and designing for these new interaction paradigms, which blend the digital and physical realms where the computer can control the human body.

The seminar's findings will contribute to the ongoing discourse on the future of HCI, particularly in areas where control dynamics between humans and machines are increasingly blurred. Moving forward, the insights gained from this event will help

shape a more humane technological future, ensuring that innovations in computational control are aligned with ethical and societal values.

## References

- 1 Rakesh Patibanda, Chris Hill, Aryan Saini, Xiang Li, Yuzheng Chen, Andrii Matviienko, Jarrod Knibbe, Elise Van Den Hoven, and Florian 'Floyd' Mueller. Auto-paizo games: Towards understanding the design of games that aim to unify a player's physical body and the virtual world. *Proceedings of the ACM on Human-Computer Interaction*, 7(CHI PLAY):893–918, 2023.
- 2 Pedro Lopes, Sijing You, Lung-Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. Providing haptics to walls & heavy objects in virtual reality by means of electrical muscle stimulation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 1471–1482, 2017.
- 3 Florian Floyd Mueller, Pedro Lopes, Paul Strohmeier, Wendy Ju, Caitlyn Seim, Martin Weigel, Suranga Nanayakkara, Marianna Obrist, Zhuying Li, Joseph Delfa, Jun Nishida, Elizabeth M. Gerber, Dag Svanaes, Jonathan Grudin, Stefan Greuter, Kai Kunze, Thomas Erickson, Steven Greenspan, Masahiko Inami, Joe Marshall, Harald Reiterer, Katrin Wolf, Jochen Meyer, Thecla Schiphorst, Dakuo Wang, and Pattie Maes. Next steps for human-computer integration. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–15, New York, NY, USA, 2020. Association for Computing Machinery.
- 4 Romain Nith, Yun Ho, and Pedro Lopes. Splitbody: Reducing mental workload while multitasking via muscle stimulation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- 5 Rakesh Patibanda, Elise Van Den Hoven, and Florian 'Floyd' Mueller. Towards understanding the design of body-actuated play. In *Extended Abstracts of the 2022 Annual Symposium on Computer-Human Interaction in Play*, CHI PLAY '22, page 388–391, New York, NY, USA, November 2022. Association for Computing Machinery.
- 6 Eun-Soo Jung, Suh-Yeon Dong, and Soo-Young Lee. Neural correlates of variations in human trust in human-like machines during non-reciprocal interactions. *Scientific Reports*, 9, 2019.
- 7 Steve Benford, Richard Ramchurn, Joe Marshall, Max L. Wilson, Matthew Pike, Sarah Martindale, Adrian Hazard, Chris Greenhalgh, Maria Kallionpää, Paul Tennent, and Brendan Walker. Contesting control: journeys through surrender, self-awareness and looseness of control in embodied interaction. *Human-Computer Interaction*, 36(5-6):361–389, 2021.
- 8 Pedro Lopes, Lewis L. Chuang, and Pattie Maes. Physiological i/o. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI EA '21, New York, NY, USA, 2021. Association for Computing Machinery.
- 9 Rakesh Patibanda, Aryan Saini, Nathalie Overdeest, Maria F. Montoya, Xiang Li, Yuzheng Chen, Shreyas Nisal, Josh Andres, Jarrod Knibbe, Elise van den Hoven, and Florian 'Floyd' Mueller. Fused spectatorship: Designing bodily experiences where spectators become players. *Proc. ACM Hum.-Comput. Interact.*, 7(CHI PLAY), oct 2023.
- 10 Rakesh Patibanda, Nathalie Overdeest, Shreyas Nisal, Aryan Saini, Don Samitha Elvitigala, Jarrod Knibbe, Elise Van Den Hoven, and Florian 'Floyd' Mueller. Shared bodily fusion: Leveraging inter-body electrical muscle stimulation for social play. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*, DIS '24, page 2088–2106, New York, NY, USA, July 2024. Association for Computing Machinery.
- 11 Florian 'Floyd' Mueller, Nathan Semertzidis, Josh Andres, Joe Marshall, Steve Benford, Xiang Li, Louise Matjeka, and Yash Mehta. Toward understanding the design of intertwined human-computer integrations. *ACM Trans. Comput.-Hum. Interact.*, 30(5), sep 2023.
- 12 Saul Greenberg, Sebastian Boring, Jo Vermeulen, and Jakub Dostal. Dark patterns in proxemic interactions: a critical perspective. In *Proceedings of the 2014 Conference on Designing Interactive Systems*, DIS '14, page 523–532, New York, NY, USA, 2014. Association for Computing Machinery.

## 6.37 Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis

**Editors:** Tor Dokken, Xiaohong Jia, Géraldine Morin, and Elissa Ross  
**Seminar No.** 24241

Date: June 9–14, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.52

© Creative Commons BY 4.0 license

© Tor Dokken, Xiaohong Jia, Géraldine Morin, and Elissa Ross



**Participants:** Massimo Carraturo, Falai Chen, Tor Dokken, Gershon Elber, Konstantinos Gavriil, Carlotta Giannelli, Ron Goldman, Hans Hagen, Stefanie Hahmann, Kai Hormann, Xiaohong Jia, Bert Jüttler, Panagiotis Kaklis, Shahroz Khan, Myung-Soo Kim, Tae-wan Kim, Stefan Kollmannsberger, Jiri Kosinka, Tom Lyche, Zoë Marschner, Dominik Mokriš, Géraldine Morin, Caitlin Mueller, Georg Muntingh, Suraj R. Musuvathy, Baldwin Nsonga, Daniele Panozzo, Jorg Peters, Jeff Poskin, Helmut Pottmann, David Reeves, Ulrich Reif, Elissa Ross, Péter Salvi, Maria Lucia Sampoli, Espen Sande, Scott Schaefer, Felix Scholz, Gunnar Schulze, Yongjie Jessica Zhang, Jianmin Zheng, Eric Zimmermann

The Dagstuhl Seminar 24241 “Geometric Modeling: Challenges for Additive Manufacturing, Design and Analysis” took place in the week of June 9–14, 2024. This year, the seminar returned to having on-site participants only. In the previous seminar (2021, during the COVID pandemic), a hybrid format was used. In 2021, close to two thirds of participants joined remotely from eleven different time zones. This year, with all participants on site, all could participate in discussions during breaks, meals, and in the evenings. The popular hike on the afternoon of the third day of the seminar was well-attended, combining scientific discussion with a walk through the beautiful surroundings of Dagstuhl. With most participants joining all five days, we benefited from high attendance during all sessions.

One of the challenges when planning a new Dagstuhl Seminar on geometric modeling was to achieve the right balance between renewal and continuation. This was particularly challenging for this iteration of the seminar, as the targeted number of participants was around 40, in comparison to the more than 50 who attended the 2021 seminar (including both on-site and online participation). Consequently, when making the initial list for invitations, some names that have contributed over many seminars were not included. This was necessary to allow new scientists to be invited in the interest of renewal. As always, there were some late cancellations that allowed us to also invite many of those on the reserve list. Industry participation was higher than in previous events. Many of the participants were newcomers who brought new ideas and life into the discussions. For the next seminar, the organizers will once again have to balance renewal and continuation.

As with previous seminars, geometric modeling remained the core topic of the seminar. However, in recent seminars, the focus has shifted from representation of shape for computer aided design to the challenges posed by a wide use of these technologies

in industry and society. In particular, the use of the geometric model is considered within a thorough and complete process in order to design, optimize, and create manufactured 3D content. As the title of the seminar suggests, applications in additive manufacturing, design, and analysis were central. During the seminar, challenges of geometry representation and processing for architecture also arose as a major topic.

The abstracts of the 35 talks presented at the seminar are included in the seminar’s full report as well as the conclusions of the three working groups that addressed challenges in:

- geometric modeling for additive manufacturing,
- geometric modeling for design optimization, and
- computer aided geometric design and isogeometric analysis.

The topics discussed in these working groups ranged from theoretical challenges in spline technology, to the need for improved digital technology in geometric modeling, to manufacturing of novel shape concepts in architecture. An emerging theme was the topic of geometric data, which has become more diverse in shape and nature. At the same time, geometric models are also strongly linked and considered in a context broader than the representation of shape, which includes their physical properties and the capabilities to manufacture these models in a sustainable manner. The reflection, work, and opportunities for geometric design and analysis are now more open, and Dagstuhl meetings and the diversity of participants created an opportunity to further this wide vision of the research field.

As always, both the organizers and participants of the Dagstuhl Seminar appreciated the smooth execution of the event, due to the great support and organization from the research center, the great food and lunch meetings, and the opportunity to discuss scientific challenges in a friendly atmosphere at a beautiful location and great venue.

## 6.38 Computational Analysis and Simulation of the Human Voice

**Editors:** Sten Ternström, Nathalie Henrich Bernardoni, Oriol Guasch, and Peter Birkholz  
**Seminar No. 24242**

Date: June 9–14, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.84

© Creative Commons BY 4.0 license

© Sten Ternström, Nathalie Henrich Bernardoni, Oriol Guasch, and Peter Birkholz

**Participants:** Peter Arndt, Steve Awodey, Andrej Bauer, Karim Johannes Becher, Olaf Beyersdorff, Marc Bezem, Ingo Blechschmidt, Ulrik Buchholtz, Gabriele Buriola, Felix Cherubini, Michel Coste, Laura Crosilla, Nicolas Daans, Dominique Duval, Martín H. Escardó, Giulio Fellin, Makoto Fujiwara, Hugo Herbelin, Matthias Hutzler, Hajime Ishihara, Ulrich Kohlenbach, Henri Lombardi, Maria Emilia Maietti, Julien Narboux, Sara Negri, Takako Nemoto, Stefan Neuwirth, Satoru Niki, Paige North, Eugenio Orlandelli, Edi Pavlovic, Iosif Petrakis, Elaine Pimentel, Michael Rathjen, Marie-Françoise Roy, Peter M. Schuster, Monika Seisenberger, Sana Stojanovic-Djurdjevic, Benno van den Berg, Steven J. Vickers



The human voice is able to produce a very rich set of different sounds, making it the single most important channel for communication human-to-human, and also potentially for human-computer interaction. Spoken communication can be thought of as a stack of layered transport protocols that includes language, speech, voice, and sound. This Dagstuhl Seminar was concerned with the voice and its function as a transducer from neurally encoded speech patterns to sound. This very complex mechanism remains insufficiently explained both in terms of analysing voice sounds, as for example in medical assessment of vocal function, and of simulating them from first principles, as in talking or singing machines. There were four main themes to the seminar:

**Voice Analysis.** Measures derived from voice recordings are clinically attractive, being non-invasive and relatively inexpensive. For clinical voice assessment, however, quantitative objective measures of vocal status have been researched for some seven decades, yet perceptual assessment by listening is still the dominating method. Isolating the properties of a voice (the machine) from those of its owner's speech or singing (the process) is far from trivial. Computational approaches are expected to facilitate a functional decomposition that can advance beyond conventional cut-off values of metrics and indices.

**Voice Visualization.** Trained listeners can deduce some of what is going on in the larynx and the vocal tract, but we cannot easily see it or document it. The multidimensionality of the voice poses interesting challenges to the making of effective visualizations. Most current visualizations are textbook transforms of the acoustic signal, but they are not as clinically or pedagogically relevant as they could be. Can functionally or perceptually informed visualizations improve on this situation?

**Voice Simulation.** Balancing low- and high-order models. A "complete" physics-based computational model of the voice organ would have to account for bidirectional energy exchange between fluids and moving structures at high temporal

and spatial resolutions, in 3D. Computational brute force is still not an option to represent voice production in all its complexity, and a proper balance between high and low order approaches has to be found. We discussed strategies for choosing effective partitionings or hybrids of the simulation tasks that could be suitable for specific sub-problems.

**Data science and voice research.** With today's machine learning and deep neural network methods, end-to-end systems for both text-to-speech and speech recognition have become remarkably successful, but they remain quite ignorant of the basics of vocal function. Yet machine learning and big data science approaches should be very useful for helping us deal with and account for the variability in voices. Rather than seeking for automated discrimination between normal and pathological voice, clinicians wish for objective assessments of the progress of an intervention, while researchers wish for ways to distil succinct models of voice production from multi-modal big-data observations. We have explored how techniques such as domain-specific feature selection and auto-encoding can make progress toward these goals.

**This seminar has resulted** in (1) shared knowledge and data about the science of voice from the perspectives of scientists in fields as diverse as computer science, voice pathology and therapy, clinicians, acoustics and audio engineering, electronics, musicology, speech and hearing sciences, physics and mathematics, (2) identifying areas of common interest where significant progress is being made and needs to be made, such as individual voice variability, physical replicas for modelling and validation, synthesis and computational modelling, motor control, and availability of data and resources, (3) sharing and discussing failures to learn lessons and ideas for future developments, and (4) envisioning the future of progress in human voice analysis and simulation in the medium to long term: what is needed to make a big leap forward in this field? These ideas will be captured by the publication of a collaborative article in a leading voice journal.

## 6.39 Teaching Support Systems for Formal Foundations of Computer Science

**Editors:** Tiffany Barnes, Jan Vahrenhold, and Thomas Zeume  
**Seminar No.** 24251

Date: June 16–21, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.108

© Creative Commons BY 4.0 license

© Thomas Zeume, Tiffany Barnes, and Jan Vahrenhold



**Participants:** Efthimia Aivaloglou, Erik Barendsen, Tiffany Barnes, Rodrigo Duran, Felix Freiburger, Judith Gal-Ezer, Michael Goedicke, Michael Hielscher, Norbert Hundeshagen, Johan Jeuring, Sandra Kiefer, Dennis Komm, Shriram Krishnamurthi, Martin Lange, Alexandra Mendes, Tilman Michaeli, Daphne Miedema, Liat Peterfreund, Sophie Pinchinat, Seth Poulsen, Ramaswamy Ramanujam, Anne Remke, Susan Rodger, Florian Schmalstieg, Marko Schmellenkamp, Francois Schwarzenruber, Thomas Schwentick, Jakob Schwerter, R. Benjamin Shapiro, John Slaney, Vaishnavi Sundararajan, Jan Vahrenhold, Maximilian Weininger, Thomas Zeume, Chenhao Zhang

The primary goal of this Dagstuhl Seminar was to determine how to enable communication between researchers in computing education, builders of systems for teaching formal foundations, as well as instructors of these foundations. While these groups have very similar interests, they also have very different notions, foci, and methods. In particular, participants from the “formal foundations” community talk about the “hardness” of a problem in terms of its computational complexity, participants from the “intelligent tutoring systems” community are concerned with whether or not a system can scale or how to best provide feedback to the learner, and participants from the “computing education research” community study the effectiveness of teaching methods for learning, e.g., the cognitive load, student learning, etc.

Within the first one-and-a-half days of the seminar, tutorials on Formal Foundations of CS, CS Education Research, and Intelligent Tutoring Systems given by experts of the respective domains set the stage for the rest of the seminar. The tutorial on CS Education Research was interspersed with breakout sessions for applying the theoretical content of the tutorial to projects of seminar participants, leading to intense discussions across the

different communities and therefore being very effective also in bridging barriers between the communities. On the afternoon of the first day, tools and tutoring systems in the formal foundation domain were presented in teaser and poster sessions as well.

The rest of the seminar was centered around breakout sessions, whose research and discussion topics were proposed and voted on by participants. There were a few contributed research talks and occasional ad-hoc tutorial-like sessions as they became relevant for the breakout sessions.

Participants noted that the seminar had a very open atmosphere and that the different research communities were eager to learn from each other. This welcoming spirit was also reflected by a music event on one of the evenings where three of the participants gave a concert and a fare-well magician’s show by one of the participants as part of the closing session.

In summary, it was a very fruitful seminar – both with respect to research collaborations and personal interactions. The goal of bringing together the communities and bridging the gaps between them was fully achieved. Several collaborative research projects were initiated during the seminar and are currently being followed-up on.

## 6.40 Computational Creativity for Game Development

**Editors: Duygu Cakmak, Setareh Maghsudi, Diego Perez Liebana, and Pieter Spronck**  
**Seminar No. 24261**

Date: June 23–28, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.130

© Creative Commons BY 4.0 license

© Pieter Spronck, Duygu Cakmak, Setareh Maghsudi, and Diego Perez Liebana

**Participants:** Maren Awiszus, Gabriella A. B. Barros, Paolo Burelli, Duygu Cakmak, Filippo Carnovalini, Alex J. Champandard, M Charity, Michael Cook, João Miguel Cunha, Alena Denisova, Alexander Dockhorn, Anders Drachen, Manuel Eberhardinger, Raluca D. Gaina, James Goodman, Christian Guckelsberger, Greta Hoffmann, Amy K. Hoover, Chengpeng Hu, Leonie Kallabis, Ahmed Khalifa, Antonios Liapis, Simon M. Lucas, Setareh Maghsudi, David Melhart, Gwaredd Mountain, Matthias Müller-Brockhausen, Mirjam Palosaari Eladhari, Diego Perez Liebana, Johanna Pirker, Mike Preuß, Emily Short, Hendrik Skubch, Gillian Smith, Tristan Smith, Pieter Spronck, Anne Sullivan, Tommy Thompson, Tony Veale, Vanessa Volz



Developments in artificial intelligence are currently dominated by deep learning technology, which generates deep neural networks, trained on large data sets, which excel at pattern recognition. Variants of the “classic” deep neural networks have the ability to generate new data with statistical properties similar to the training set. Generative Adversarial Networks (GANs), such as those used by DALL-E and Midjourney, may be used to generate original visual artworks based on a textual description of the desired output. Autoregressive language models, such as those used by ChatGPT, use deep learning to produce text that is often indistinguishable from human-created text. Moreover, artificial intelligence techniques have been used to successfully generate music for many years, and researchers have also experimented with using deep learning to create cooking recipes, personalized fragrances, fashion, and more.

Despite the sometimes astonishing products of such creative artificial intelligence, the results are usually lacking in meaning. While DALL-E and Midjourney produce images that seem impressive, upon further inspection they contain many mistakes which humans would avoid. While ChatGPT can generate human-sounding text in a conversation, it often produces utter nonsense, and cannot write an original coherent story. And, as our own explorations of such techniques during Dagstuhl Seminar 22251 showed, GANs may produce computer game content which looks reasonable at first glance, but is ultimately neither functional nor playable.

While the product of creative artificial intelligence can often be used as a strong basis for humans to build upon, and may as such speed up the creative process, human intelligence and human creativity are almost always a necessary ingredient of the creative process. Moreover, the more relevant the meaning, purpose, and functionality of the product are, the less the creative process benefits from the involvement of artificial intelligence.

Game design and implementation are tasks which require a

high amount of creativity, and which must lead to products which require a high amount of fine-tuned functionality. For example, a game “level” should not only look appealing, it should also be playable (i.e., it must be possible for most players to finish the level) and it should be interesting to play (i.e., the player should feel entertained by playing the level and should experience inherent motivation to finish the level). These are not features which can be acquired simply by “training on big data,” which is what most developments in modern artificial intelligence are based on.

The goal of Dagstuhl Seminar 24261, Computational Creativity for Game Development, was to investigate to what extent modern artificial intelligence techniques can produce meaningful and functional game content, and what changes to or extensions of these techniques can improve this AI-driven creative process.

We like to point out that progress in this area is relevant for a wide range of applications outside the “games” domain. Creativity in artificial intelligence applies to many branches of industry and has a strong impact on society, in which artificially intelligent technology interacts with humans in many shapes and forms. We use games in our research because they are highly-complex but well-defined applications which form safe environments to experiment in. However, solutions found for creative problems in games are often transferable to domains outside games.

This research area lends itself to a wide range of research topics. For the preparation of this seminar, we proposed the following set of sub-topics (many of which were taken up by workgroups):

- **Procedural Content Generation for Games:** Procedural Content Generation (PCG) systems include techniques and methods able to create different types of game elements, such as levels, rules, quests, and characters, among many others. Research in PCG has been prolific in the last decade, but its

presence in the games industry is still far from ubiquitous. A particular interest is set on mixed-initiative systems, which give designers and artist authorial control of the created content and the direction of the algorithm that generates it.

- **Procedural Generation of Games:** An extension of the previous point that deserves its own separate area is the generation of complete games. Same as recent advances in generative systems for music, painting, and long bodies of text, one can research how complete games can be generated from scratch. This includes elements like art, rules, characters, and winning and losing conditions that normally form a game. Automatic generation of new mini-games can open an interesting space of research that merges multiple advances together, but also a useful tool for game designers that will be able to use the generated games as inspiration for new entertainment experiences.
- **Computational Creativity for Narrative Games:** A particular type of game that has become more popular in the last decade is that of narrative games. While there has been some work in using computational creativity methods to generate texts, the adaptation of these techniques to the game development process remains an open area of research.
- **Automatic Generation of Art in Games:** Art is an important part of digital games and takes multiple forms: 3D-models, textures, visual effects, animations, cut-scenes, and so on. Lately, multiple advances on the use of computational creativity have shown the capacity of generating different forms of art, such as images, videos, and even 3D geometries. Examples of systems that generate art are DALL-E, Stable Diffusion, and Midjourney. Research can explore how these and other techniques can be used to generate art for games, including unexplored game art areas, in particular with regards to how this generation can be bound to specific games/genres/restrictions, how it can be integrated into the game development process, and how we can give designers authorial control and modification capabilities over the generated assets.
- **Procedural Generation of Audio for Games:** An important part of automatic generation of content refers to audio. From audio effects (footsteps, heartbeats, weather) to complete sound tracks (background music, melodies, singers), including the generation of different voices for human and non-human game characters, the space for computational creativity to generate this type of art is vast.
- **Computational Creativity for Game Playing:** An unexplored aspect and application of computational creativity is that of generating AI agents that play a game. Traditionally, the objective that leads AI agents in a game to play is to achieve victory, either by reaching a winning state or by maximizing the score they obtain in the game. Some efforts have been made to employ quality-diversity methods to generate different styles of play. Research may explore how we can harness the new developments in computational creativity to generate diverse play styles, including the generation of new strategies or tactics to play games in a different manner.
- **Computational Creativity for Affective Computing:** Affective computing is a discipline that bridges several domains, such as computer science, cognitive science, and psychol-

ogy. It studies the implementation of systems that are able to express, identify, process, and simulate human affects. Research may investigate computational creativity algorithms and methods to provide non-player characters with the possibility of expressing feeling and emotions in a convincing way. This includes, for example, facial expressions and body animations, and it can be applied to human or non-human characters.

- **Automatic Support of Game Development:** Traditionally, computational creativity and the automatic generation methods have focused their efforts on generating the product that creative industries build – be this games (or content for games), art or music, among others. These methods may also be used to aid the process of game development. Examples of the application of this technology include computational creativity for automation of tasks, algorithms for automatic testing of development process (such as code, integration, animations and deployment), production chains, and procedural development processes.
- **Ethical Considerations of Computational Creativity:** The ethical challenges of using computationally creative tools for applications such as game development should not be ignored. The use of automatically-creative tools may have negative effects on the need for artists and designers. Moreover, the automatic creation of games and game content may lead to ethically suspect products. Finally, biases that exist in art and data may be magnified when such products are used to automatically generate new products. These ethical considerations were taken into account in all our explorations of advances in computational creativity.

More than a year-and-a-half passed between writing the proposal and running the seminar. We found it striking how many advancements had been made in the area of Computational Creativity for Games in that period. During the writing of the proposal we were personally convinced that we were proposing an important theme for the seminar. When the seminar took place, we knew that no other theme was this topical.

This seminar was organized around workgroups, which worked in teams and topics proposed by the participants of the seminar in the areas outlined above. These workgroups were accompanied by plenary sessions for group formation, topic debate, and discussions of the deliberation of each group. Workgroups were dynamic, so participants could move between them, and new groups were formed during the week.

A Discord server was set up for coordination and announcements, and it was also used by the different groups for document and link sharing. This also has the benefit of providing a place for discussions after the seminar, easing the communication and further work among the members of each workgroup.

42 participants accepted our invitation to join the seminar; 40 of them attended. The participants were a good mixture of gender, country of origin, junior and senior seniority, and academia or industry. All participants engaged intensively within the seminar, and many expressed how happy they were with what we accomplished, making the seminar a great success.

## 6.41 Theory of Randomized Optimization Heuristics

**Editors:** Anne Auger, Tobia Glasmachers, Martin S. Krejca, and Johannes Lengler  
**Seminar No.** 24271

Date: June 30 – July 5, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.215

© Creative Commons BY 4.0 license

© Anne Auger, Tobias Glasmachers, Martin S. Krejca, and Johannes Lengler

**Participants:** Sumit Adak, Denis Antipov, Dirk Arnold, Asma Atamna, Anne Auger, Dimo Brockhoff, Alexandre Chotard, Benjamin Doerr, Carlos M. Fonseca, Stephan Frank, Oskar Girardin, Armand Gissler, Tobias Glasmachers, Andrea P. Guerreiro, Nikolaus Hansen, Mario Alejandro Hevia Fajardo, Alexander Jungeilges, Kathrin Klamroth, Timo Kötzing, Oswin Krause, Martin S. Krejca, Per Kristian Lehre, Tristan Marty, Daiki Morinaga, Frank Neumann, Andre Opris, Aishwarya Radhakrishnan, Jonathan Rowe, Günter Rudolph, Thomas Sauerwald, Jonathan L. Shapiro, Dirk Sudholt, Andrew M. Sutton, Kento Uchida, Vanessa Volz, Carsten Witt, Christine Zarges, Weijie Zheng



This seminar is part of a biennial seminar series. This year, we envisioned to focus on constraint handling, multivariate estimation-of-distribution algorithms (EDAs), as well as stochastic approximation and Markov chain stability analysis. This vision worked well for constraint handling but not so much for the other two topics, since several key invitees rejected our invitations. Nonetheless, the seminar quickly and organically refocused, and we had plenty of other important topics to discuss instead, as we detail below.

The previous iteration of the seminar had taken place during the peak of the COVID Omicron wave, with lots of restrictions in place. We were glad to see this seminar happening again in the usual format. It was still a bit smaller than usual, due to unfortunate last-minute cancellations, but we managed to partially make up for that by inviting more young researchers. In any case, the group size of about 40 worked very well.

We used morning sessions for talks and group discussions, and afternoon sessions for breakout formats, with the exception of Monday. Several time slots were exempt from the schedule, explicitly leaving free time for individual discussions and/or leisure activities. We kept talks reasonably short, usually at a maximum of 20 minutes, leaving sufficient time for discussion. We had to stop discussions only very rarely so as to respect Dagstuhl's meal schedule. Otherwise, we managed to give sufficient space to each topic by flexibly re-scheduling a few topics on the fly.

We used the first day for group forming and to bring everyone to the same page. We started with a general introduction, a small ice-breaker game, and a series of overview talks. We had talks on optimization heuristics in discrete and continuous search spaces, which represents a classic divide in our community, as well as introductions to multi-objective optimization and constraint handling. Furthermore, we had presentations by Thomas Sauerwald on *Balls into Bins* and by Vanessa Volz on

*Democratising Real-World Problem Tailored Optimisation*, two contributions a bit remote from the core topic of the seminar, so as to build bridges to neighboring fields. Kathrin Klamroth and Oswin Krause continued this series later on with presentations on *A Multiobjective Perspective on Constraint Handling* and on *Evolution in the Quantum world: on tuning quantum dot devices*.

We invited and actually nudged all young researchers to introduce themselves and their research with brief presentations. We had a total of seven such talks, taking place in the morning slots of the following days. This format was generally perceived as useful and fruitful.

Besides junior and outreach talks, there was of course also a lot of activity on core topics of the community. Two selected highlights were the presentation of Armand Gissler's proof of linear convergence of the CMA-ES algorithm on convex quadratic problems by means of Markov chain analysis, and the presentation of Per Kristian Lehre on his SLO hierarchy, a categorization of complexity classes for black-box optimization.

There was a total of 10 breakout sessions, which are all summarized in Section 4 of the full Dagstuhl Report. They covered a broad spectrum of diverse and important topics, and they easily replaced the initial focus topics that we had in mind. This time, we had a fair amount of breakout sessions that garnered the attention of both the discrete and the continuous subcommunity in roughly equal parts. As such, we had sessions on negative drift (which is a tool applicable in either domain), on mixed-integer problems (which requires expertise from both domains), the general structure of discrete problems with the aim to classify them similarly to how continuous problems can be classified, and a session on an abstract framework that can be used by practitioners and analyzed by theoreticians. Moreover, we discussed rather recent hot topics such as multi-objective optimization, co-evolution, and rich surrogate models, assessing which direction to take for the future. This was complemented

by a session on permutation problems, highlighting a domain that only saw little attention so far, as well as a session on how to place the theoretical analysis of randomized search heuristics within the grander spectrum of artificial intelligence. All of these sessions had a larger number of participants and vibrant discussions, showing the interest of the community in these topics.

We established a new format for collecting topics and for scheduling breakout sessions. It was inspired by a workshop that had taken place at the Lorentz Center in Leiden (Netherlands) earlier this year. Instead of collecting topics asynchronously upfront and over lunch breaks, we dedicated a short session to it. To this end, we replaced the usual classroom arrangement with a fully symmetrical setup, asking participants to step up, briefly explain their proposal, write down a title on a sheet of colored paper, and put it up on a wall. Moreover, instead of trying to reach consensus, we left it to session organizers to schedule their

sessions so as to minimize (perceived) overlap. The new process was received as an improvement.

For the first time in this seminar series, we offered a trip to Trier as a social activity on Wednesday afternoon, with a hike taking place in parallel (and in the rain). This worked very smoothly, and it was a great experience, especially for participants from far away, even if they had been to Dagstuhl before.

We are very grateful for the opportunity of organizing a seminar in Dagstuhl. We have to thank for the financial support, for the comfort of rooms directly at the venue, four meals a day, wonderful facilities, and all of that including full service by very friendly, reactive, and always helpful staff. Thank you!

The organizers,

Anne Auger, Tobias Glasmachers, Martin Krejca, Johannes Lengler

## 6.42 A Game of Shadows: Effective Mastery Learning in the Age of Ubiquitous AI

**Editors: Nick Falkner, Juho Leinonen, Miranda C. Parker, Andrew Petersen, and Claudia Szabo**  
Seminar No. 24272

Date: June 30 – July 5, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.6.245

© Creative Commons BY 4.0 license

© Nick Falkner, Juho Leinonen, Miranda C. Parker, Andrew Petersen, and Claudia Szabo

**Participants:** Claudio, Nick Falkner, Armando Fox, Michelle Friend, Petri Ihantola, Johan Jeuring, Hieke Keuning, Natalie Kiesler, Päivi Kinnunen, Tobias Kohn, Juho Leinonen, Andrew James Luxton-Reilly, Stephen MacNeil, Lauri Malmi, Claudia Ott, Miranda C. Parker, Andrew Petersen, Judith Sheard, Jacqueline Smith, Claudia Szabo, Lisa Zhang, Jaromír



The integration of generative AI (GenAI) into education raises significant issues and opportunities, particularly concerning mastery learning and programming education. A primary concern is that students may bypass deep engagement with their learning tasks by relying on AI tools or search engines, which leads to a superficial understanding of the material. This tendency forces instructors to focus more on monitoring for academic dishonesty rather than on effective teaching. To address this, innovative approaches to presenting curricula and materials could foster greater student engagement and reduce the inclination to rely on external aids.

The transformative potential of AI in education is likened in Armando Fox's talk to the early use of movie cameras, which initially focused on replicating existing practices rather than exploring new possibilities. The emphasis is on avoiding mere substitution of traditional methods with AI, and instead, leveraging AI to create entirely new modes of learning. The goal is to integrate AI in a way that complements foundational educational concepts and develops the necessary intellectual frameworks to utilize these new tools effectively.

Mastery learning remains an effective approach in the age of AI, with a deep understanding of the importance of clearly defined learning goals and assessments designed to resist cheating. One proposed method involves students setting their own learning goals, which are then approved by instructors. This approach not only personalizes learning but also utilizes GenAI to assist in developing and refining assignments. GenAI can also serve as a brainstorming partner for educators, providing valuable support in creating innovative class activities and assessments.

AI tools like code co-pilots are nowadays becoming integral for understanding and implementing code. This shift presents challenges for traditional assessment methods, as AI can outperform students in specific coding tasks. It is essential for students to develop strong communication skills about code, as these

will become increasingly crucial in collaborative programming environments where AI tools are prevalent. This suggests that the skill of discussing code should be as rigorously taught as technical coding skills. In addition, providing accurate feedback and identifying essential skills for effective software development remain critical. While LLMs can offer valuable feedback, ensuring that this feedback is accurate and relevant remains a challenge. Additionally, defining the skills necessary for students to develop software with the support of LLMs, such as program specification and refactoring, is crucial for leveraging these tools effectively.

The integration of mastery learning with GenAI presents both significant potential and challenges. While GenAI can enhance personalized learning, assessment creation, and feedback, effective implementation requires careful consideration of the tools used, their alignment with educational goals, and their impact on learning outcomes. Ensuring that these tools are suitable for diverse educational contexts and measuring their effectiveness will be key to successfully adopting mastery learning supported by GenAI, ultimately aiming to improve educational outcomes and better prepare students for future professional challenges.

The seminar was structured into three main sections: lightning and keynote talks, brainstorming, and workshop groups. At the beginning of the seminar, each attendee delivered a lightning talk. Two keynote talks were delivered, as follows:

- Prof Andrew Luxton-Reilly, University of Auckland: "It's the end of the world as we know it, but I feel fine! Teaching and learning with GenAI"
- Dr Claudia Ott, University of Otago: "A Decade of Mastery Learning at Otago – Pitfalls, Challenges & Opportunities"

The lightning talk sessions were followed by a brainstorming session to identify existing challenges and opportunities. This session served a dual purpose, in that it also allowed us to identify

the three main working groups of the seminar. These groups focused on (i) curriculum and pedagogy of mastery learning in the era of GenAI, (ii) university and organisation structures that facilitate the delivery and operationalisation of mastery learning in the era of GenAI, and on (iii) designing courses in a curriculum that is GenAI focused.

### ■ Publications to Date

At the date of the submission of this report, the following papers and posters had been accepted for publication:

- “Models of Mastery Learning for Computing Education”, by Claudia Szabo, Miranda Parker, Michelle Friend, Johan Jeuring, Tobias Kohn, Lauri Malmi, Judithe Sheard was accepted for publication at SIGCSE 2025 as a position paper.
- “Goodbye Hello World – Research Questions for a Future CS1 Curriculum”, Hieke Keuning, Andrew Luxton-Reilly, Claudia Ott, Andrew Petersen and Natalie Kiesler was accepted at Koli Calling 2024 as a poster.

## 6.43 Dynamic Traffic Models in Transportation Science

**Editors:** José Correa, Carolina Osorio, Laura Vargas Koch, and David Watling  
**Seminar No. 24281**

Date: July 7–12, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.1

© Creative Commons BY 4.0 license

© José Correa, Carolina Osorio, Laura Vargas Koch, and David Watling

**Participants:** Umang Bhaskar, Roberto Cominetti, Richard Connors, José R. Correa, Katharina Eickhoff, Andrés Salomón Fielbaum Schnitzler, Gunnar Flötteröd, Martin Gairing, Song Gao, Lukas Graf, Svenja M. Griesbach, Tobias Harks, Martin Hoefer, Takamasa Iryo, Saif Jabari, Max Klimm, Ekkehard Köhler, Hong K. Lo, Gaurav Malik, Jannik Matuschke, Haruko Nakao, Neil Oliver, Tim Oosterwijk, Carolina Osorio, Dario Paccagnan, Britta Peis, Koki Satsukawa, Marco Scarsini, Daniel Schmand, Marc Schröder, Ravi Seshadri, Alexander Skopalik, Nicolas Stier-Moses, Martin Strehler, Laura Vargas Koch, Bernhard von Stengel, Peter Wagner, David Watling, Chiwei Yan, Theresa Ziemke



Transportation road networks and the underlying spatio-temporal dynamics of vehicular traffic are increasingly complex. City infrastructure is increasingly real-time and traffic-responsive, urban mobility services are increasingly flexible, multi-modal and tailored to the heterogeneous needs and preferences of the traveling population, road assets are increasingly shared between the transport of individuals and the transport of freight and goods. Importantly, operators and users of the road network increasingly focus on the sustainability of their choices or service offerings. The key sustainability metrics (e.g., greenhouse gas emissions and energy consumption) are tightly linked to the underlying traffic dynamics.

Dynamic traffic assignment models are, therefore, critical tools for transport planners to be able to predict the spatio-temporal distribution of traffic – and thereby the congestion, environmental and social impacts – of transport policies, for example, in the light of possible changes to the infrastructure (e.g. traffic light control), to the transport services offered (e.g., the provision of public transport fleets and frequencies), or to the prices imposed on travelers (e.g., congestion pricing, public transport fares). The prevailing approaches used in the transportation science literature to predict such distributions can be roughly classified into mathematical approaches based on *dynamic traffic assignment models* (DTA) (using the methodology of flows over time) and *simulation-based approaches* (using large-scale microsimulations). The striking advantage of microscopic simulations over DTA models is that the latter usually ignore practically relevant side constraints such as horizontal queueing of vehicles, the feedback of changing network conditions on user behavior, flexible departure time choice, mode choice, and activity schedule choice. Current simulation tools integrate all these dimensions and many more. The increase in model complexity inherent in simulations, however, is not matched by the existing DTA theory. For most simulation software, there is no “mathematical proof

of concept”. This means it is not known whether equilibria always exist and the heuristics applied in traffic simulation lead to (approximate) equilibria. Moreover, crucial properties for simulations to give meaningful results, such as uniqueness and continuity of equilibria, are only partially understood.

In the Dagstuhl Seminar, we brought together, for the fourth time, leading researchers from three different communities (Simulations, Dynamic Traffic Assignment, and Algorithmic Game Theory) in order to bridge the gap between complex simulation-based models and the existing theory. In the last four seminars, we made substantial progress by answering questions, e.g., on the controllability of traffic flow or the analysis of equilibria with spatial queues. Still, real-world simulation develops faster than mathematical theory, and at the same time, new challenges in modeling future road networks arise, which are characterized by real-time responsiveness, flexible multi-modal urban mobility offerings, automated and autonomous vehicles, coordinated fleets of vehicles, and environmentally sustainable operations. This seminar focused on these new questions with multiple talks on sustainability in traffic simulation in general and more concrete on ride-hailing, ride-sharing, public transportation, and the control of flow via tolls or information design.

For this edition of the seminar, we again welcomed applied researchers from the industry alongside experts in traffic simulations, algorithmic game theory (AGT), and dynamic traffic assignment (DTA). In this issue of the seminar, we strongly profited from the fact that the community has grown together over the last three Dagstuhl Seminars, and we developed a good mutual understanding and a common language. On the one hand, multiple talks and open problems addressed solutions or progress on questions from previous seminars. On the other hand, at this seminar, we focused on sustainability in traffic planning, established new questions, and extended the core community.

Again, the seminar was a great success. The seminar not

only stimulates fruitful collaboration, but by now, we have built a community around this series of seminars with many important research questions and a very supportive and open atmosphere.

We got laudatory feedback from many participants, which is also reflected in the survey conducted by Dagstuhl.

## 6.44 Automated Machine Learning For Computational Mechanics

**Editors:** Elena Raponi, Lars Kotthoff, Hyunsun Alicia Kim, and Marius Lindauer  
**Seminar No. 24282**

Date: July 7–12, 2024 | Dagstuhl Seminar  
 Full report – DOI: 10.4230/DagRep.14.7.17

© Creative Commons BY 4.0 license

© Elena Raponi, Lars Kotthoff, Hyunsun Alicia Kim, and Marius Lindauer

**Participants:** Niels Aage, Fadi Aldakheel, Paolo Ascia, Thomas Bäck, Monica Capretti, Carola Doerr, Elsayed Saber Elsayed Ibrahim Elsayed, Helen Fairclough, Roman Garnett, Frank Hutter, Hyunsun Alicia Kim, Lars Kotthoff, Peter Krause, Melvin Leok, Olaf Mersmann, Charles Mish, Markus Olhofer, Lisa Pretsch, Elena Raponi, Thiago Rios, Gokhan Serhat, Niki van Stein, Marc Zöller



The Dagstuhl Seminar 24282 was organized with the objective of bringing together the automated machine learning (AutoML) community with the computational mechanics (CoMe) community and finding new ways the two communities could help each other. More specifically, this seminar was trying to answer two questions:

1. What are the problems the CoMe community encounters when applying machine learning (ML) methodologies, and how can the AutoML community help overcome them?
2. What characteristics of CoMe benchmarks are currently not well supported by AutoML, and what limitations of AutoML tools are holding back their application in CoMe?
3. How can applications in CoMe inform research directions in AutoML?

To search for an answer to these questions, the seminar was structured with a mixture of talks and group sessions. On one hand, the talks provided an overview of AutoML methodologies that could be valuable to the CoMe community, as well as examples of CoMe applications where ML methods have been applied to address common research questions and challenges. On the other hand, the group sessions offered the opportunity to delve into specific topics, either within individual communities or through cross-disciplinary discussions. The topics selected for the group discussions were intended to identify the barriers preventing the AutoML community from addressing common practical challenges in CoMe test cases, and, conversely, the obstacles hindering the adoption of AutoML methodologies by the CoMe community. More specifically, the topics discussed in groups were:

1. Features and Problem Characterization;
2. Optimization;
3. Integration of Physics;
4. Explainability;
5. Benchmarks and constraints handling.

The seminar was then concluded by a plenary discussion on the outcomes of the given talks and group discussions. During this discussion, we were able to identify a general CoMe workflow that includes and exploits AutoML tools, see Figure 6.9. Moreover, we agreed that there are already many CoMe test cases the AutoML community can use to test their new methodologies, but these are badly disseminated. Therefore, in addition to the newly established collaborations between researchers belonging to the two different fields, one of the major outcomes of the seminar was the recognition that, to bridge the gap between AutoML and CoMe, the first crucial step would be to publish a review paper that compiles and categorizes the currently available datasets and benchmarks.

### ■ Organization of the Seminar

This small Dagstuhl Seminar brought together 23 researchers from both engineering and machine learning/optimization, representing both academia and industry. The group included a mix of senior and junior researchers, creating a diverse and collaborative environment.

Over the course of five days, the mornings featured 14 short presentations, each lasting 15 to 20 minutes. The rest of the seminar was organized in a dynamic and flexible manner, with activities ranging from scientific speed-dating and two-way surveys to the traditional trekking, as well as plenary and parallel discussions on topics chosen by the participants (see the complete seminar schedule in Fig. 6.10). This flexible structure allowed for a more engaging and relaxed schedule, which was appreciated by all attendees. Discussions that began during the day often extended into the evening, where the cozy atmosphere of the Dagstuhl castle played a key role in making everyone feel comfortable during both work-related exchanges and more informal moments.

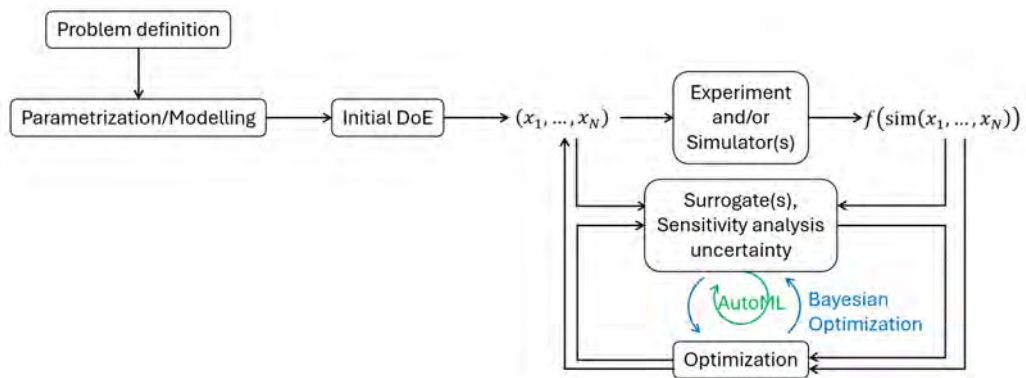


Fig. 6.9 General workflow integrating CoMe problems and AutoML tools.

Dagstuhl Seminar on AutoML for Computational Mechanics					
	Monday June 8	Tuesday June 9	Wednesday June 10	Thursday June 11	Friday June 12
9:00 - 9:10	Introduction	Daily Sync	Daily Sync	Daily Sync	Daily Sync
9:10 - 10:30	Opening Talks: Fadi Aldakheel, Frank Hutter, Thomas Bäck	Talks: Gokhan Serhat, Melvin Leok, Niki van Stein, Carola Doerr	Talks: Roman Garnett, Lisa Pretsch, Paolo Ascia, Alicia Kim	Talks: Charlie Mish, Peter Krause, Niels Aege	Open discussion
10:30 - 11:00	Coffee Break				
11:00 - 12:15	1-slide Intros + Discussion	Two-way survey	Open Discussion	Future projects, collaborations, conferences	Joint report writing + goodbye
12:15 - 13:30	Lunch				
13:30 - 15:30	Speed date + Breakout topics	Open discussion - Breakout topics	Traditional Hike/Trip	Spontaneous group discussions	
15:30 - 16:00	Coffee Break			Coffee Break	
16:00 - 17:30	Breakout Session + Report	Breakout Session + Report		Report	
18:00	Dinner				

Plenary	Parallel
---------	----------

Fig. 6.10 Seminar schedule.

### Outcome

Based on the survey results, the seminar was widely considered a great success by both the organizers and participants. Despite the participants needing some initial time to get used to the languages of the two different research communities, the seminar successfully bridged the gap between the machine learning and computational mechanics communities, fostering valuable cross-disciplinary dialogue. The event provided a platform for sharing cutting-edge research and insights, while also addressing the practical challenges faced when integrating automated machine learning (AutoML) into computational mechanics workflows.

The presentations were thought-provoking and initiated lively discussions throughout the seminar. The collaborative spirit extended beyond the formal sessions, with participants engaging in productive breakout sessions and working groups. These dynamic exchanges not only explored the current state of research but also laid the groundwork for potential future collaborations.

The event demonstrated the clear potential for synergy between these fields, and it is expected that the connections made during the seminar will continue to grow and lead to impactful advances. The organizers extend their sincere thanks to the Scientific Directorate and the Dagstuhl Center administration and staff for their precious support, which was instrumental in the seminar's smooth execution and success.

## 6.45 Programmable Host Networking

**Editors: Gianni Antichi, Katerina Argyraki, Aurojit Panda, and Justine Sherry**  
**Seminar No. 24291**

Date: July 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.35

© Creative Commons BY 4.0 license

© Gianni Antichi, Katerina Argyraki, Aurojit Panda, and Justine Sherry

**Participants:** Peter Arndt, Steve Awodey, Andrej Bauer, Karim Johannes Becher, Olaf Beyersdorff, Marc Bezem, Ingo Blechschmidt, Ulrik Buchholtz, Gabriele Buriola, Felix Cherubini, Michel Coste, Laura Crosilla, Nicolas Daans, Dominique Duval, Martín H. Escardó, Giulio Fellin, Makoto Fujiwara, Hugo Herbelin, Matthias Hutzler, Hajime Ishihara, Ulrich Kohlenbach, Henri Lombardi, Maria Emilia Maietti, Julien Narboux, Sara Negri, Takako Nemoto, Stefan Neuwirth, Satoru Niki, Paige North, Eugenio Orlandelli, Edi Pavlovic, Iosif Petrakis, Elaine Pimentel, Michael Rathjen, Marie-Françoise Roy, Peter M. Schuster, Monika Seisenberger, Sana Stojanovic-Djurdjevic, Benno van den Berg, Steven J. Vickers



Over the past two-decades network link speeds have grown faster than processor clock speeds. Consequently, the software and hardware that applications use to communicate is often the bottleneck. Eliminating this bottleneck requires redesigning the software network stack and hardware NICs that applications use to send or receive messages. Until now changing either was challenging, but recent changes to the software ecosystem have made it possible to prototype, and then deploy new network stack and NIC designs. On the software side, the addition of eBPF to the Linux and FreeBSD kernels have made it easier to modify the network stack *without* needing to change the kernel, while on the hardware side, the availability and adoption of programmable SmartNICs such as Xilinx Alveo and Intel Agilex allows us to now change both NIC design and the NIC-OS interface without needing to upgrade or change server hardware.

However, at present there is no consensus on how NIC-OS-Userspace should be architected, what features they should provide, or how they should be implemented. This is because answering these questions requires combining techniques and ideas from several sub-disciplines including networking, systems, computer architecture, programming languages, and compiler design. This seminar's goal was to bring together

academics and industry practitioners who had expertise in these areas and discuss how to best architect and build platforms that enable accelerator use.

We structured the seminar around four core questions:

- (a) What hardware accelerators are used in practice, and how?
- (b) What interfaces do current OS accelerator frameworks provide, what are their limitations, and how are they used?
- (c) How do hardware accelerators communicate and coordinate with processors?
- (d) What tools do we have available to reason about the performance of both hardware and software accelerators?

The choice of questions was informed by the set of attendees, and by the organizers experiences in this area. Surprisingly we found that there was significant disagreement on the desirability of both hardware and software accelerators, with many of the attendees arguing that they might be a temporary measure while we figure out what future servers look like. This is a different conclusion than what one would arrive at by looking at recent papers in the systems, networking, and architecture communities, and from current industry deployments. Our full report provides a more detailed account of the seminar's discussions.

## 6.46 Improving Trust between Humans and Software Robots in Robotic Process Automation

**Editors:** Adela del Río Ortega, Andrea Marrella, Hajo A. Reijers, and Adriana Wilde  
**Seminar No.** 24292

Date: July 14–19, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.52

© Creative Commons BY 4.0 license

© Adela del Río Ortega, Andrea Marrella, Hajo A. Reijers, Adriana Wilde



**Participants:** Simone Agostinelli, Marco Angelini, Aleksandre Asatiani, Bernhard Axmann, Piercosma Bisconti, Angelo Casciani, Christian Czarnecki, Adela del Río Ortega, Andrea Delgado, José González Enríquez, Glenda Hannibal, Christian Janiesch, Andrés Jiménez Ramírez, Faizan Ahmed Khan, Andrea Marrella, Antonio Martínez Rojas, Artur Modlinski, Ralf Plattfaut, Jana-Rebecca Rehse, Hajo A. Reijers, Manuel Resinas, Michael Rosemann, Flávia Santoro, Stefan Sarkadi, Pnina Soffer, Barbara Weber, Adriana Wilde

This summary provides an overview of the outcomes of our Dagstuhl Seminar “*Improving Trust between Humans and Software Robots in Robotic Process Automation*” (24292). It began with a general introduction to the aim, scope, and context of the Dagstuhl Seminar. The preliminary presentation was followed by a sequence of three invited talks required to clarify the main aspects investigated during the seminar. Specifically, Simone Agostinelli presented an “Introduction to Robotic Process Automation”, and Piercosma Bisconti introduced the basic frameworks to specify trust in autonomous systems in the talk “How to build trust between intrinsic and perceived trustworthiness”. Then, a third talk by Michael Rosemann discussed the existing research efforts to integrate trust into the BPM discipline in his talk: “Managing Trust in Business Processes”.

The background talks were followed by three-minute speeches, during which every participant could give a brief overview of their background, expertise, and personal expectations for the seminar. The seminar participants included experts from the BPM, RPA, HCI and Trust-aware AI communities, with industry representatives and researchers in academia at different levels of seniority.

Then, a plenary brainstorming session began, moderated by the seminar organizers. Every participant in the seminar could suggest a topic they would like to investigate together with the other attendees. The discussion resulted in a set of twenty arguments that were finally distilled together into the five research questions below:

- How trust is formed between humans and RPA technology?
- How to effectively calibrate and manage trust in RPA to align pre-implementation expectations with post- implementation realities?
- How to design trust-aware human-robot interaction in RPA?
- How to identify the right balance between autonomy and trust in RPA systems?
- How to guide organizations in creating a hybrid workforce of SW robots and humans?

From the second day on, time was mostly devoted to working groups, each discussing one of the research questions that were raised in breakout sessions. This activity was accompanied by “lightning talks”, i.e., brief presentations on issues that emerged during the discussions. Periodic intermediate group presentations reported on the advancement of the teams’ work. The final presentations of the results achieved by the working groups and a discussion of the continuation of the freshly established collaborations took place on the final day. The seminar was concluded with the plan for writing a common vision paper to establish the foundations of a new breed of trust-aware systems.

The rest of the report contains all the talks held on the first day of the seminar, the lightning talks and the final working group reports.

## 6.47 Art, Visual Illusions, and Data Visualization

**Editors: Christophe Hurter, Claus-Christian Carbon, Mauro Martino, and Bernice E. Rogowitz**  
**Seminar No. 24301**

Date: July 21–26, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.81

© Creative Commons BY 4.0 license

© Christophe Hurter, Claus-Christian Carbon, Mauro Martino, and Bernice E. Rogowitz

**Participants:** Kim Albrecht, Michael Bach, Michele Banks, Marco Bertamini, Alina Braun, Claus-Christian Carbon, Shanthi Chandrasekar, Oliver Deussen, Peter Eades, Brian D. Fisher, Itay Goetz, Ludwig Hanisch, Christophe Hurter, Sophia Huth, Stephen Kobourov, Jürgen Kornmeier, Bodo Korsig, Michael Kubovy, Karina Kueffner, Cassandra R. Lee, Mauro Martino, Fiona Menzel, Mario Alexis Emilio Michelessa, Arthur I. Miller, Claudia Muth, Alexander Pastukhov, Rebecca Pfiffer, Sylvia Pont, Marius Raab, Brian Rogers, Bernice E. Rogowitz, Arthur Shapiro, Ahna Skop, Dejan Todorovic, Jan Willem Tulp, Boyu Xu, Rebecca Ruige Xu



Art, visual illusions, and data science are seemingly disparate fields, but they are deeply intertwined. Art has always been a means of visual expression and communication of complex narratives. Visual illusions have captivated audiences for centuries with their ability to “trick the eye,” and challenge our perceptions, but also to teach the beholders about cognitive and perceptual functioning in a joyful and compact way (Carbon, 2014). Data science offers powerful tools for analyzing and interpreting diverse and large data sets, using novel visualization metaphors to provide insight and new perspectives. These three disciplines offer different approaches to understanding the relationship between the visual representation and its perception. The goal of this seminar was to provide a forum for exploring synergies between these diverse disciplines. This premise is built on a strong foundation. There is a long history of artists incorporating visual illusions in their art to challenge audience perceptions. Some examples include Robert Delaunay’s explorations of color contrast; Victor Vasarely’s spatial illusions; Bridget Riley’s optical illusions; Salvador Dalí’s spatial sampling illusions; M.C. Escher’s impossible constructions; Giuseppe Arcimboldo’s composite portraiture; Yayoi Kusama’s infinity rooms; Hans Holbein the Younger’s anamorphic art; and Patrick Hughes’s reverse perspective paintings. Contemporary artists now use data science to create works that explore hidden patterns and relationships in complex data sets. Examples include Laurie Frick’s use of self-tracking information gathered from personal devices to create hand-built works and installations; Jenny Odell’s combination of satellite images to create collages

revealing patterns in human-made structures; Refik Anadol’s use of datasets from EEG brainwave patterns and machine learning algorithms to create immersive, data-driven installations exploring memory and cognitive processes; Nathalie Miebach’s intricate, three-dimensional sculptures visualizing climate change data and its impact on weather patterns; Giorgia Lupi and Stefanie Posavec’s hand-drawn visualizations of personal data exchanged through postcards over the course of a year; Aaron Koblin’s visualization of airplane flight paths across the United States, creating mesmerizing patterns that reveal the complexity of air traffic; Onformative’s use of face detection algorithms to search for human-like faces in satellite images from Google Earth, exploring pareidolia and pattern recognition; Ingo Günther’s illuminated globes that visualize various datasets such as population density, energy consumption, and military spending; and Ben Rubin and Mark Hansen’s collection and display of real-time text fragments from internet chat rooms and forums, revealing the ebb and flow of online conversations. Visualization researchers explore spatial, temporal, and color metaphors to represent complex data and relationships (Brooks, 1988), and are increasingly turning to visual arts and perceptual psychology for new ways to communicate complex data in more intuitive and engaging ways. This seminar brought together artists, perceptual and cognitive psychologists, and computer scientists active in the data science fields of visualization and artificial intelligence, who are not only experts in their own fields, but whose work has actively crossed discipline boundaries. Our format encouraged the organic evolution of these connections and synergies.

## 6.48 Learning with Music Signals: Technology Meets Education

**Editors:** Meinard Müller, Cynthia Liem, Brian McFee, and Simon Schwär  
**Seminar No.** 24302

Date: July 21–26, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.115

© Creative Commons BY 4.0 license

© Meinard Müller, Cynthia Liem, and Brian McFee



**Participants:** Vipul Arora, Ching-Yu Chiu, Roger B. Dannenberg, Christian Dittmar, Zhiyao Duan, Mark Gotham, Masataka Goto, Patricia Hu, Jaehun Kim, Katherine M. Kinnaird, Cynthia Liem, Lele Liu, Hanna Lukashevich, Brian McFee, Peter Meier, Alia Morsi, Meinard Müller, Juhan Nam, Alex Ruthmann, Simon Schwär, Sebastian Stober, Bob Sturm, Christopher J. Tralie, Timothy Tsai, Anja Volk, Changhong Wang, Christof Weiß, Jordan Wirfs-Brock

This executive summary provides an overview of our discussions on advancing technology and education in music information retrieval (MIR) and related fields, summarizing the main topics covered in the seminar. We also describe the seminar’s group composition, overall organization, and activities. Finally, we reflect on the most important aspects of the seminar and conclude with future implications and acknowledgments.

### ■ Overview

In the last twenty years, the field of music information retrieval (MIR) has undergone rapid developments in terms of the problems considered, the methodology, and its applications. Using conceptually simple tasks and methods evaluated on small and idealized datasets in its beginnings, MIR now contributes to a wide range of concepts, models, and algorithms that extend our capabilities of accessing, analyzing, understanding, and creating music. Given the complexity and diversity of music, MIR research considers various aspects such as genre, instrumentation, musical form, melodic and harmonic properties, dynamics, tempo, rhythm, and timbre, to name a few. Furthermore, music is inherently multimodal, incorporating speech-like signals (e.g., singing), videos (e.g., of live performances), static images (e.g., scanned music scores), and text (e.g., lyrics and reviews). This wealth of data makes MIR an interdisciplinary and challenging field of research, which closely connects to technical disciplines such as signal processing, machine learning, and information retrieval, as well as mathematics, musicology, psychology, and the digital humanities.

Having the Dagstuhl Seminar 24302 titled “Learning with Music Signals: Technology Meets Education,” our objective was to advance technology and education in MIR and related

disciplines using music as a challenging and instructive multimedia domain. Thinking of data-driven machine learning approaches, we discussed recent deep learning (DL) approaches and their ability to learn from training examples to make accurate predictions for previously unseen data. Furthermore, by learning from the experience of traditional engineering approaches, our aim was to better understand existing and build more interpretable DL-based systems (e.g., through integrating prior knowledge).

Beyond these technically oriented perspectives, an essential focus of our seminar was to approach the concept of learning from other angles, including pedagogical, educational, psychological, and user-centered viewpoints. We argue that music is an essential part of our lives that most people feel connected to. Therefore, music yields an intuitive entry point to support education in technical disciplines. In particular, we explored how music may serve as a vehicle to make learning and teaching signal processing and machine learning an interactive pursuit.

In all perspectives on learning, the question of reproducible research, including open access to data and software, is becoming increasingly important (so future insights can be built upon existing ones). As an overarching topic of our seminar, we discussed questions about open science and good scientific practice. This is a key issue, especially in higher education, where the open exchange of best practices and teaching materials significantly impacts the interdisciplinary and transnational education of the next generation of researchers.

In summary, in our Dagstuhl Seminar we approached and explored the concept of learning from different angles, see also Figure 6.11. Besides considering technological developments, the seminar equally addressed aspects of data and model understanding, transdisciplinary methodology and applications, science communication, and education.

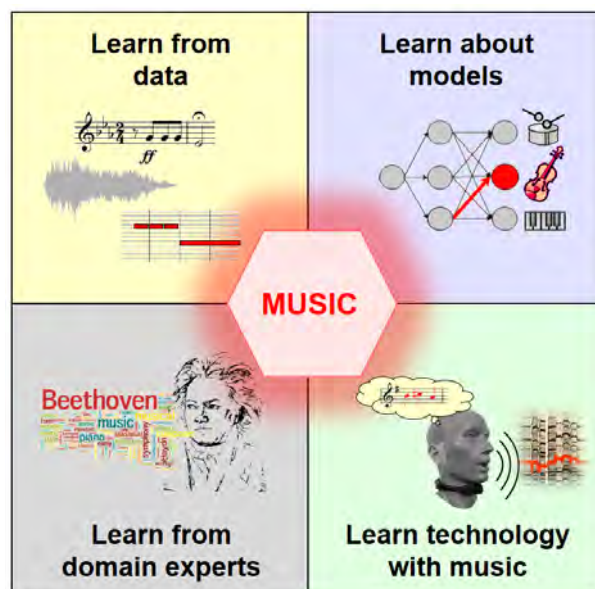


Fig. 6.11

Learning with music signals. The figure illustrates the various perspectives on learning: advancing deep learning techniques, integrating traditional engineering knowledge for interpretability, collaborating with domain experts to understand music corpora, and using music to enhance interactive learning in technical disciplines.

## ■ Participants and Group Composition

In our seminar, we had 28 participants from various locations around the world, including North America (eight participants from the United States), Asia (four participants from China, India, Japan, and South Korea), and Europe (16 participants from France, Germany, the Netherlands, Spain, Sweden, and the United Kingdom). Beyond geographic diversity, many participants had cross-cultural backgrounds and experiences. As naturally happens in international research fields, part of this comes from participants' work experiences in other countries and cultures than those of their country of birth. At the same time, several participants were first-generation for being in research/higher education, had second-generation migration backgrounds (meaning their parents were born in different countries, and often cultures, than themselves) – or, because of the nature of their affiliation, brought extensive experience in teaching with students from such backgrounds.

The seminar was not only international, but also highly interdisciplinary. While most of the researchers specialized in music information retrieval with a technical focus on signal processing and machine learning, we also had participants with backgrounds in musicology, human-computer interaction, science education, mathematics, computer vision, and other fields. This diversity stimulated cross-disciplinary discussions, bringing together experts from both technical and non-technical disciplines and highlighting opportunities for new collaborations. Many participants had strong musical backgrounds, with some even having dual careers in engineering and music, leading to numerous social activities, including playing music together. We also aimed to foster variety in terms of seniority levels, with four Ph.D. students and three postdoctoral participants, as well as gender diversity, with 10 out of 28 participants identifying as female. More than half of the participants (16 out of 28) were attending Dagstuhl for the first time and expressed enthusiasm about the open and retreat-like atmosphere. In conclusion, by bringing together internationally renowned scientists and promising early-career researchers from different fields, our seminar provided support and encouragement for emerging talents on their academic paths.

## ■ Overall Organization and Schedule

Dagstuhl Seminars are known for their flexibility and interactivity, encouraging participants to discuss ideas and raise questions rather than merely presenting research results. In keeping with this tradition, we set the schedule during the seminar, inviting spontaneous contributions focused on future-oriented content. This approach helped us avoid a conference-like atmosphere, where the emphasis is often on past research achievements. Furthermore, instead of sitting in rows, we removed all tables and arranged the seating in a half-circle of chairs, significantly enhancing eye contact and interaction among all participants.

After the organizers provided an overview of the Dagstuhl concept, we began the first day with self-introductions, where each participant shared their background, expectations, and wishes for the seminar. We then proceeded with brief stimulus talks, lasting 15 to 20 minutes, in which selected participants addressed critical questions related to the seminar's overall theme in a non-technical manner. Each talk smoothly transitioned into an open discussion among all participants, with the presenter acting as the moderator. These discussions were well-received and often extended for more than half an hour. The first day concluded with a brainstorming session on central topics reflecting the participants' interests, helping to shape the schedule and format for the following day.

On the subsequent days, we continued with short stimulus talks followed by long and intensive discussion rounds. We also incorporated group discussions, splitting into smaller groups to delve into specific topics in greater depth. The results and conclusions of these parallel group sessions, which lasted between 60 to 90 minutes, were then presented and discussed with the entire group. Additionally, we included panel-like elements featuring moderators, panelists, interviews, surveys, and game-like group activities. On the last day, we concluded the seminar with a session we called "self-outroductions," where each participant presented their personal view on the seminar's results. In summary, thanks to excellent group dynamics and a fair distribution of speaking time, all participants had the opportunity to express their thoughts, effectively avoiding a monotonous conference-like presentation format.

In addition to scientific questions, our seminar also addressed the various challenges that younger colleagues typically face when establishing their research groups and academic curriculum at the beginning of their careers. As previously mentioned, many of our participants had cross-cultural backgrounds, either being born in Asian countries or as second-generation individuals raised in Western cultures. One of the highlights of our Dagstuhl Seminar was a panel discussion on the cross-cultural challenges in academia, especially for individuals with Asian roots living and working in Europe or the US. This deeply personal and enlightening event was facilitated by Dagstuhl's unique environment, which fosters trust and mutual understanding.

While working in technical engineering disciplines, most participants also had a strong background and interest in music. This versatility significantly enriched the seminar's atmosphere, fostering cross-disciplinary interactions and sparking thought-provoking discussions. It also led to intensive joint music-making during breaks and evenings. A particular highlight was the Thursday evening concert organized by Cynthia Liem and Christof Weiß, where various ensembles formed by participants performed a wide variety of music, including classical, Irish folk, and jazz.

## ■ Conclusions and Acknowledgment

At the Dagstuhl Seminar 24302, we used music as a motivating and tangible domain to explore different perspectives on learning. These perspectives stimulated conceptual discussions, laying the groundwork for future projects and academic curricula. We focused on how to teach and pass on new technologies to students, using music as a challenging application domain. With experts in MIR, signal processing, machine learning, software development, science education, and music sciences, our interdisciplinary seminar generated vibrant discussions and highlighted opportunities for new collaborations. Immediate outcomes, such as plans to share research data and software, also emerged from the discussions. We aimed to expose attendees, especially early-career researchers, to new ideas for designing academic curricula in computer science and beyond. Specific areas and topics addressed in this seminar included:

- Contextualized education
- Inclusive education
- Educational software systems
- Interactive software frameworks

- Science communication
- Transdisciplinary methodology and collaborative research
- Computational musicology
- Human-in-the-loop systems for music processing
- Data-driven machine learning for MIR
- Explainable deep learning for MIR
- Integration of musical knowledge
- Hybrid models for MIR
- Differentiable models for MIR
- Data mining, acquisition, measurement, and annotation
- Data/annotation quality
- Data accessibility and copyright issues
- Open science
- Reproducible and sustainable research
- Academic integrity and good scientific practice

Besides the scientific aspect, the social aspect of our seminar was equally important. We hosted an interdisciplinary, international, and interactive group of researchers, consisting of current and future leaders in our field. Many participants were visiting Dagstuhl for the first time and praised the open and inspiring setting. The group dynamics were excellent, with many personal exchanges and shared activities. Some scientists expressed their appreciation for the opportunity to engage in prolonged discussions with researchers from neighboring fields, something often impossible at typical conferences. A standout feature of our seminar was the interaction between younger researchers at the beginning of their academic careers and established researchers and educators. This facilitated a deeply enriching exchange between different generations, promoting mutual trust and understanding. The intensive dialogue between these groups was truly outstanding and highlighted the unique value of our seminar.

In conclusion, our expectations for the seminar were not only met but exceeded, particularly in terms of networking and community building. We want to express our gratitude to the Dagstuhl board for giving us the opportunity to organize this seminar, the Dagstuhl office for their exceptional support throughout the organization process, and the entire Dagstuhl staff for their excellent service during the seminar. In particular, we want to thank Heike Clemens, Andreas Dolzmann, Marsha Kleinbauer, Simone Schilke, and Christina Schwarz for their invaluable assistance in the preparation and organization of the seminar.

## 6.49 Resource-Efficient Machine Learning

**Editors: Oana Balmau, Matthias Boehm, Ana Klimovic, Peter Pietzuch, and Pinar Tözün**  
**Seminar No. 24311**

Date: July 28 – August 2, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.153

© Creative Commons BY 4.0 license

© Oana Balmau, Matthias Boehm, Ana Klimovic, Peter Pietzuch, and Pinar Tözün

**Participants:** Matthias Böhm, Maximilian Böther, Marco Canini, Jerónimo Castrillón-Mazo, Patrick Damme, Khuzaima Daudjee, Pamela Delgado, Thaleia Dimitra Doudali, Jens Hagemeyer, Steven Hand, Dagmar Kainmüller, Fredrik Kjolstad, Ana Klimovic, James Kwok, Manisha Luthra Agnihotri, Peter R. Pietzuch, Tilmann Rabl, Theo Rekatsinas, Ties Robroek, Stefanie Scherzinger, Tom St. John, Foteini Strati, Shinya Takamaeda-Yamazaki, Pinar Tözün, Lluís Vilanova, Eiko Yoneki, Cliff Young, Ehsan Yousefzadeh-Asl-Miandoab



While the capabilities of machine learning models have become more and more impressive in the last decade, one cannot overlook the computational footprint of their end-to-end lifecycle. According to the Stanford AI Index Report [1], the computational complexity of the state-of-the-art language models has increased 7 orders of magnitude since 2017. In turn, this increases the estimated costs to train these models in the cloud by 5 orders of magnitude, and the carbon footprint of training such models are equivalent to 10s of human years. Furthermore, the cost of training is only a fraction of the whole costs. After training, then comes the cost of continuously deploying these models, which depend on the way these models are used for inference and the frequency of retraining to update the models.

The participants of the Dagstuhl Seminar on Resource-Efficient Machine Learning (ML) targeted the computational efficiency challenges for machine learning, especially deep learning, from different angles and by focusing on the different stages. On the first day of the seminar, we split into four groups, each with a specific focus. The groups identified the research questions they want to focus on, delved deeper into the existing work, and identified future steps to continue collaborations beyond the seminar.

The first group, **Resource-Efficient Data Selection** (Section 3 of the full report), targets the efficiency of data selection methods for training deep learning models. Data selection is a preliminary step before any model training, but specifically for fine-tuning tasks, where a pre-trained model must be specialized for a specific task. The effectiveness of a data selection method is typically evaluated by the accuracy it achieves for the given task. The assumption is, as a side effect, if one can achieve a certain accuracy while using less data, this would improve the efficiency of training. This group questions this assumption and asks the following research question: *what are the trade-offs*

*between the computational complexity of a data selection method, its effectiveness in terms of model accuracy, and the end-to-end training efficiency?*

The second group, **The Future of Portable, Extensible, and Composable Machine Learning Systems** (Section 4 of the full report), aims at making the emerging ML systems support a larger diversity of applications more efficiently. At the core of this support lies a departure from the dominant reliance on dense tensor computations. Targeting a larger diversity in applications also requires looking at a larger variety of hardware devices, beyond large accelerators that are highly optimized for dense matrix computations. Targeting such diversity requires co-design and finding the right abstractions across software, compilers, and hardware. The group's vision results in several research challenges with the following overarching research question: *how to design holistic and composable software and hardware frameworks for ML?*

The third group, **Hardware-Software Co-Design for Machine Learning** (Section 5 of the full report), target similar research challenges to the second group, but with a deeper focus on hardware diversity. The group identifies that the conventional way of optimizing machine learning tasks for a certain hardware device is through tight coupling between high-level ML engineering and low-level performance optimizations. Certain high-level optimizations, with a specific hardware in mind, in turn, hinders portability to different hardware devices, resulting in sub-optimal efficiency and missed opportunities for functionality. Therefore, the key research question here is *how does one create a hardware stack for ML that enables better portability across different devices?*

The fourth group, **Workload-Aware Machine Learning Serving** (Section 6 of the full report), focus on ML serving. Serving ML models, especially large language models (LLMs), at

scale is highly costly and requires substantial hardware resources. To achieve more resource- and performance-efficient model ways of serving models, one needs to adaptively determine which specialized model to serve or cache, or how to optimize a model. This adaptivity is highly dependent on the workload needs that

may be dynamic. This group, therefore, aims to answer the following questions: (1) *what is the behavior and needs of the real-world serving workloads* and (2) *how does one build a framework that enables adaptive model serving based on dynamic user and workload needs?*

## ■ References

- 1 The AI Index Report. The AI Index Report - Measuring trends in AI. <https://aiindex.stanford.edu/report/>, 2024.

## 6.50 Security and Privacy of Current and Emerging IoT Devices and Systems

6

**Editors: Bruno Crispo, Alexandra Dmitrienko, Gene Tsudik, and Wenyan Xu**  
**Seminar No. 24312**

Date: July 28 – August 2, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.7.170

© Creative Commons BY 4.0 license

© Bruno Crispo, Alexandra Dmitrienko, Christoph Sendner, Gene Tsudik, and Wenyan Xu

**Participants:** Z. Berkay Celik, Alfred Chen, Bruno Crispo, Ivan De Oliveira Nunes, Xuhua Ding, Alexandra Dmitrienko, Jan-Erik Ekberg, Earlence Fernandes, Kevin Fu, Jorge Guajardo Merchan, David Hock, Murtuza Jadliwala, Yongdae Kim, Farinaz Koushanfar, Veelasha Moonsamy, Surya Nepal, Panagiotis Papadimitratos, Christina Pöpper, Sara Rampazzi, Kasper Rasmussen, Stefanie Roos, Ahmad-Reza Sadeghi, Nader Sehatbakhsh, Christoph Sendner, Gene Tsudik, Markus Wamser, Wenyan Xu



Over the past two decades, there has been a surge in the popularity of Internet-of-Things (IoT) devices and Cyber-Physical Systems (CPS). These devices are now commonplace in private settings, such as homes, offices, and factories, and public spaces like cultural, entertainment, and transportation facilities. They are also extensively used in farming, industrial, and vehicular automation. Furthermore, they are often interconnected and connected to the global Internet. These devices are typically built using low-end microcontroller units (MCUs), which have strict cost, size, and energy constraints and lack security features compared to their higher-end counterparts. As a result, these embedded devices, including sensors, actuators, and hybrids, have become attractive targets for various types of attacks. The focus of these attacks ranges from privacy concerns in the context of sensing, to safety and security issues in the context of actuation, and even zombification, as seen in the infamous Mirai botnet.

The goal of the Dagstuhl Seminar was first to explore the landscape of attacks on current and emerging devices and then to identify and discuss promising research directions for effective countermeasures, both reactive and proactive. The relationship between academic research and industry was also of interest; specifically, to what extent is there a flow of ideas and innovation from the research community and device manufacturers, and what can be done to improve it.

The original proposal for this Dagstuhl Seminar included nine topics. However, once the seminar was approved and attendees were confirmed, it turned out that some of the topics simply did not have sufficient numbers of interested participants. In the end, five topics “survived” to the actual seminar and became – together with one new topic – individual sessions. (The session leader(s) are mentioned in parenthesis.):

1. **Security and Privacy Challenges in IoT-instrumented spaces** (Gene Tsudik) – originally named **Implications of increasing human immersion in instrumented spaces**. The

increasing human immersion in instrumented spaces, such as smart homes, offices, and cities, brings new security and privacy challenges. The proliferation of interconnected devices that collect, store, and transmit personal data creates a larger attack surface for cybercriminals. Additionally, users may not be aware of the privacy implications of these devices and the data they generate. The seminar participants will focus on emerging privacy concerns and discuss the ways to protect users’ personal data while still enabling the functionality of these devices. Another challenge is the lack of standardization, which makes it difficult to create uniform security and privacy mechanisms and ensure the interoperability of different devices.

2. **Realizing Security/Privacy Services Across Hardware and Software Boundaries** (Alexandra Dmitrienko) – originally named **Scope of potential security/privacy services and how they should be realized across the SW/HW boundary**. Currently, microprocessors and many hardware platforms implement many security mechanisms, such as PAC, MTE, BTI, PUF, shadow stacks, and others, in hardware. However, more high-level trusted services (such as remote attestation, authentication, etc.) implemented in software on these platforms only partially, if at all, utilize these available mechanisms. To fully realize the potential of these security features implemented in hardware, a more systematic co-design of hardware and software is required. This approach can result in more efficient implementation of existing trusted services, as well as the design of new ones, thereby enhancing the overall security and privacy of IoT platforms.
3. **The Role of Secure Hardware (Trusted Computing) in IoT Security** (Bruno Crispo) – originally named **The Role of Trusted Computing in providing robust security services for IoT**. Trusted Execution Environments (TEEs) and Roots of Trust (RoTs) are common Trusted Computing tools in

the academic literature related to IoT security. Specific instances have also been implemented in the industry (e.g., ARM TrustZone, Intel SGX, etc.). However, the first generation of TEEs has shown to be vulnerable to security issues. Therefore, major vendors and initiatives like RISC-V are revisiting trusted computing architectures to avoid past pitfalls. Hence, it is important to establish what types of Trusted Computing technologies are needed for different types of devices and under what types of attacks to ensure effective security measures.

4. **Balancing mission-criticality, safety, and security in system design** (Wenyuan Xu) – originally named **Mission-criticality/safety vs. security/privacy**. When it comes to mission-critical systems, striking a balance between safety and security can be challenging. Security measures often come with real-time overhead such as timing delays, interruptions, and increased bandwidth usage. In critical settings where safety is the primary concern, these impacts can have significant consequences. Unfortunately, safety and security requirements are often treated in isolation during the design process, without considering their natural implications and the correlation between them. This separation is also reflected in the evaluation of these systems, with different and separate standards and regulations for assessing safety and security. To address this challenge, it is important to consider these two concerns jointly rather than in isolation. By doing so, unexpected interferences between the two subsystems can be avoided. A comprehensive approach to designing and evaluating mission-critical systems should take into account both safety and security requirements in an integrated manner. This will help ensure that these systems operate reliably and securely, without compromising on safety.
5. **Security Challenges in Unattended (IoT) Environments, e.g., Low-Orbit Satellites** (Wenyuan Xu and Bruno Crispo) – originally named **Space and other challenging (unattended) environments**. Low-orbit satellites are becoming increasingly popular. Deployed on a large scale, they are expected to provide ubiquitous Internet connectivity. However, these satellites operate in challenging, unattended environments that are physically inaccessible to humans. Despite the absence of attacks on low-orbit satellites thus far, it is only a matter of time before they become a target. The central challenge for designers of low-orbit satellite systems is to develop resilient and fault-tolerant security methods that can mitigate attacks from both far away and nearby sources. The remoteness of these satellites makes it difficult to detect and respond to attacks in real time, which increases the importance of designing security measures that can withstand attacks and continue operating even if a compromise occurs.
6. **Addressing the scalability challenge in securing large IoT deployments** (Alexandra Dmitrienko and Gene Tsudik).

## 6.51 Proof Representations: From Theory to Applications

**Editors: Anupam Das, Elaine Pimentel, and Lutz Straßburger**

**Seminar No. 24341**

Date: August 18–23, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.8.1

© Creative Commons BY 4.0 license

© Anupam Das, Elaine Pimentel and Lutz Straßburger

**Participants:** Matteo Acclavio, Bahareh Afshari, Sara Ayhan, Eben Blaisdell, Yll Buzoku, Kaustuv Chaudhuri, Zhibo Chen, Anupam Das, Abishek De, Amy Felty, Alexander Gheorghiu, Marianna Girlando, Rajeev P. Gore, Tao Gu, Andrzej Indrzejczak, Raheleh Jalali, Timo Lang, Anela Lolic, Bruno Lopes, Robin Martinot, Dale Miller, Victor Nascimento, Sara Negri, Carlos Olarte, Edi Pavlovic, Elaine Pimentel, Ian Pratt-Hartmann, Revantha Ramanayake, Alexis Saurin, Peter M. Schuster, Sana Stojanovic-Djurdjevic, Lutz Straßburger, Iris van der Giessen, Fernando Velázquez Quesada, Heinrich Wansing, Richard Zach



Dagstuhl Seminar 24341 was organized in response to growing interest in the representation of formal proofs. Its primary aim was to bring together theorists and practitioners who are exploring various proof representations, with the goal of identifying new applications while simultaneously creating new theoretical directions. A key focus of the seminar was to explore the interface between proof normalization and proof search traditions, by examining proof representations from both perspectives.

The seminar focused on the relation between various new developments in the field, including: the more philosophical direction of proof-theoretic semantics; the upcoming unifying research program of universal proof theory; and the study of non-standard proof formats, such as non-well-founded proofs. Individual talks covered a broad array of topics relevant to these developments, such as the particulars of certain variants on standard proof calculi (calculi including witness operators, calculi for inconsistent logics and logics with numerical quantification, calculi for second-order logic, and “exotic” proof calculi, e.g. labelled calculi for modal and intuitionistic logics), but also various connections of proof theory to computational applications (verification logic, propositional dynamic logic, tableaux using SAT solvers, and interactive theorem provers). To do justice to these complex topics, researchers from proof theory, computational logic, and philosophical logic were invited to collaborate and provide insights. The seminar also identified several research gaps across these areas. One important takeaway was the importance of utilizing different representations of proof systems to address emerging open questions in the field, and to prioritize the need for unification.

The seminar itself was structured to encourage extensive interaction among participants, both formally and informally. Participants were given considerable freedom in preparing their contributions and during the seminar itself, including the option to give talks in a variety of formats. Several in-depth special sessions were organized to focus on the most important developments in the field, and longer individual talks by experts covered specific topics in greater detail. Evening sessions, dubbed “beer talks”, provided a relaxed environment for casual discussions, allowing participants to explore topics of shared interest and build connections across different areas of research.

The discussions from the seminar reflect the key interests of the community and provide valuable topics for ongoing research. At the end of the seminar, participants agreed to stay in contact to continue their discussions and foster new collaborations. Already various initiatives are being taken. For instance, Fernando Raymundo Velázquez Quesada, Carlos Olarte, and Elaine Pimentel began a promising collaboration on Epistemic Propositional Dynamic Logic following discussions at Dagstuhl. They recently held an in-person meeting in Bergen, Norway, and plan to apply for a European grant in the near future. Additionally, collaborations have among others been started by Lutz Straßburger and Revantha Ramanayake on modal logic, and by Lutz Straßburger and Matteo Acclavio on linear logic. We look forward to seeing how the results of these efforts will further shape the field.

## 6.52 Leveraging AI for Management Decision-Making

**Editors:** Stefan Feuerriegel, Foster Provost, and Galit Shmueli

**Seminar No.** 24342

Date: August 18–21, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.8.24

© Creative Commons BY 4.0 license

© Stefan Feuerriegel, Foster Provost, and Galit Shmueli



**Participants:** Kevin Bauer, Margrét Bjarnadóttir, Jessica M. Clark, Theodoros Evgeniou, Carlos Fernández-Loria, Stefan Feuerriegel, Sebastian Gabel, Travis Greene, Jungpil Hahn, Christian Janiesch, Enric Junqué de Fortuny, Nadja Klein, Mathias Kraus, Niklas Kühn, David Martens, Claudia Perlich, Joel Persson, Foster Provost, Galit Shmueli, Sriram Somanchi, Wei Sun, Wouter Verbeke, Michael Vössing, Alona Zharova, Patrick Zschech

### ■ Introduction

Artificial intelligence (AI) has been named a core element of the “fourth industrial revolution” [40]. According to recent estimates by McKinsey & Company, AI has the potential to deliver the added global economic value of \$13–20 trillion annually [5].

AI is increasingly being embraced for decision-making in management, both across a wide array of industries (e.g., health-care, banking, education, manufacturing, retail) and functions (e.g., marketing, accounting, operations, IT). For example, AI can be used for modeling customer behavior [2, 4, 6, 25, 26]. These predictions can also serve as input for better decision-making. Examples include assortment optimization [19, 24], investment decisions [32], scheduling [42], allocation decisions [29, 31, 46], and pricing [1]. AI can predict business failures and thus act as an early warning system for improving service quality [34]. AI can help to locate drivers of low quality and eventually improve product quality [44, 45].

Recent advancements in AI research hold great promise for decision-making in businesses and organizations. Driven by the surge in data availability, computing power, and algorithmic advancement, contemporary AI algorithms are capable of emulating human decision-making and judgment [39]. This places AI in a position to augment and automate a wide range of management decisions within companies and organizations. At the same time, the use of AI for decision making, especially in high-risk applications, poses ethical and legal challenges such as the use of algorithmic risk assessment tools in criminal justice [18]. Likewise, new implications arise from emerging AI acts (e.g., the EU AI Act) with crucial implications for how AI applications must be designed.

A key enabler for data-driven decision-making in business and organizations are new AI technologies [17]. For example, deep learning can empower better decisions in business analytics and operational decision-making [28, 30]. Causal machine

learning (ML) allows for optimal targeting (e.g., of customer coupons) by estimating and subsequently leveraging individualized treatment effects [11, 14, 15, 20, 37]. Probabilistic machine learning fuses methods from a statistical foundation with flexible building blocks from neural networks to yield models that are both flexible and explainable for practitioners in risk management [35, 36, 38]. Further, explainable AI (XAI) has emerged as a principled, user-centered tool not only for explaining black-box prediction models but also for explaining the decisions that are made or recommended by AI systems [3, 16, 27, 33, 45], which can be used to identify root causes of bad quality and thereby inform better decision-making in quality management [44]. Likewise, generative AI [12] and AI fairness [7, 9, 10] offer new research opportunities. Importantly, the aforementioned examples can only be solved effectively through new AI technologies that have been developed in recent years. At the same time, the use of advanced AI on digital platforms, especially the marriage of reinforcement learning with behavior modification techniques, has spurred controversy, creating adverse effects to humans and societies, such as addiction, social discord, and political polarization [22]. Existing and envisioned combinations of prediction and causal behavior modification have implications to platforms and their business customers [23]. These technologies have also created new types of barriers for academic researchers [21, 41].

### ■ Aims of the seminar

The aim of our Dagstuhl Seminar “Leveraging AI for Management Decision-Making” (24342) was to discuss the future of research on AI/ML in businesses and organizations, and how the field should evolve. We especially sought to focus on “rethinking” the field by discussing the current state of AI/ML in businesses and organizations, discussing thought-provoking questions (e.g., is explainable AI really needed in practice? Where can generative AI actually lead to productivity gains? Does the algorithmic

approach to fairness hold much for the future of AI ethics?), and maybe identifying and elevating new important research questions.

Our intended outcome was to reach a joint position (as a group) on what are important and unimportant research directions and what those directions should be going forward. What are the challenges? What are the opportunities? What research questions deserve attention? What questions are getting more attention than they really deserve? Below, we summarize our thoughts where we discuss existing research gaps and make suggestions for the field going forward.

Prior to the seminar, a survey was shared with all participants to identify key topics of interest around which we then designed our discussions. Most participants were primarily interested in topics related to method design and development, as well as the practical applications of AI. Some also expressed interest in evaluating these methods and understanding their impact on organizations. There was a strong focus on exploring the broader implications of AI, particularly in areas like ethics and governance. When asked about specific topics they would like to discuss during the seminar, participants showed the most interest in explainable AI, followed by causal ML, and generative AI. We eventually decided to prioritize the first two – explainable AI and causal ML – thus anticipating that generative AI will naturally arise in all sessions due to its prominence and thus regardless of whether it is a dedicated topic. Many indicated they were also keen to explore how AI can be applied effectively, with discussions centered around overcoming practical challenges in real-world deployments. Other popular topics include the ethics and governance of AI, its economic impact, and the implications for the future of work. However, there was less interest in topics like AI literacy and hybrid work environments. Further, participants also suggested additional topics for discussion. Some highlighted the importance of understanding the behavioral impacts of AI, such as long-term reliance on AI systems and the potential for deskilling human workers.

## ■ Organization of the seminar

We designed our Dagstuhl Seminar as an “un-conference”. By following the format of an un-conference, we eliminated the traditional sequence of research presentations from our agenda. Instead, we aimed to focus on interactivity, collaboration, and co-creation, by making space for discussions of different forms regarding how to shape the field in the future. We held discussions with the full group as well as in smaller break-out groups, where subgroups changed from day to day; we obtained information from individuals via surveys; the schedule also encouraged informal one-on-one or small group discussions while socializing. These various modes of interaction were critical, because our seminar attracted participants from a diverse crowd, from academia and industry, from method research to behavioral research, from marketing to operations. Such diverse researchers typically do not meet or interact, and hence we seized the opportunity to foster novel interactions.

We aimed to learn from each other and create more impact. For example, we actively asked each participant in the get-to-know session to provide a summary statement about their current research and where they would like to go. Throughout the seminar, there were many opportunities to potentially start new collaborations. To spur discussion, we organized short “inspiration exchanges”, which were designed as kick-offs to our breakout sessions. Hence, the idea was primarily to discuss the current state of research and point to gaps and needs to elicit forward-thinking. Here, we selected the topics prior to the seminar based on a survey that was sent to the participants. As a result, we identified two important breakout sessions: (1) AI and causality, and (2) AI and responsibility. We summarize the discussions and findings from both breakout sessions in the full report.

## ■ References

- 1 Svenja Bergmann and Stefan Feuerriegel. Machine learning for predicting used car resale prices using granular vehicle equipment information. *Expert Systems with Applications*, 263:125640, 2024.
- 2 Lennart Baardman, Igor Levin, Georgia Perakis, and Divya Singhvi. Leveraging comparables for new product sales forecasting. *International Journal of Operations & Production Management*, 27(12):2340–2343, 2018.
- 3 Kevin Bauer, Moritz von Zahn, and Oliver Hinz. Expl(ai)ned: The impact of explainable artificial intelligence on cognitive processes. *Information Systems Research*, 2023. forthcoming.
- 4 Ruomeng Cui, Santiago Gallino, Antonio Moreno, and Dennis J Zhang. The operational value of social media information. *Production and Operations Management*, 27(10):1749–1769, 2018.
- 5 Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, and Kate Smaje. *The economic potential of generative AI*. McKinsey & Company, 2023.
- 6 Real Carbonneau, Kevin Laframboise, and Rustam Vahidov. Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, 184(3):1140–1154, 2008.
- 7 Maria De-Arteaga, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Algorithmic fairness in business analytics: Directions for research and practice. *Production and Operations Management*, 31(10):3749–3770, 2022.
- 8 Koen W De Bock, Kristof Coussement, Arno De Caigny, Roman Słowiński, Bart Baesens, Robert N Boute, Tsan-Ming Choi, Dursun Delen, Mathias Kraus, Stefan Lessmann, et al. Explainable ai for operational research: A defining framework, methods, applications, and a research agenda. *European Journal of Operational Research*, 317(2):249–272, 2024.
- 9 Mateusz Dolata, Stefan Feuerriegel, and Gerhard Schwabe. A sociotechnical view of algorithmic fairness. *Information Systems Journal*, 32(4):754–818, 2022.
- 10 Stefan Feuerriegel, Mateusz Dolata, and Gerhard Schwabe. Fair ai: Challenges and opportunities. *Business & information systems engineering*, 62:379–384, 2020.
- 11 Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine learning for predicting treatment outcomes. *Nature Medicine*, 30:958–968, 2024.

- 12 Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. Generative AI. *Business & Information Systems Engineering*, 66(1):111–126, 2024.
- 13 Carlos Fernández-Loría and Foster Provost. Causal classification: Treatment effect estimation vs. outcome prediction. *Journal of Machine Learning Research*, 23(59):1–35, 2022.
- 14 Carlos Fernández-Loría and Foster Provost. Causal decision making and causal effect estimation are not the same... and why it matters. *INFORMS Journal on Data Science*, 1(1):4–16, 2022.
- 15 Carlos Fernández-Loría, Foster Provost, Jesse Anderton, Benjamin Carterette, and Praveen Chandar. A comparison of methods for treatment assignment with an application to playlist generation. *Information Systems Research*, 34(2):786–803, 2023.
- 16 Carlos Fernández-Loría, Foster Provost, and Xintian Han. Explaining data-driven decisions made by AI systems: The counterfactual approach. *Management Information Systems Quarterly*, 46(3):1635–1660, 2022.
- 17 Stefan Feuerriegel, Yash Raj Shrestha, Georg von Krogh, and Ce Zhang. Bringing artificial intelligence to business management. *Nature Machine Intelligence*, 4(7):611–613, 2022.
- 18 Travis Greene, Amit Dhurandhar, and Galit Shmueli. Atomist or holist? A diagnosis and vision for more productive interdisciplinary AI ethics dialogue. *Patterns*, 4(1):100652, 2023.
- 19 Sebastian Gabel, Daniel Guhl, and Daniel Klapper. P2v-map: Mapping market structures for large retail assortments. *Journal of Marketing Research*, 56(4):557–580, 2019.
- 20 Robin M Gubela, Stefan Lessmann, and Björn Stöcker. Multiple treatment modeling for target marketing campaigns: A large-scale benchmark study. *Information Systems Frontiers*, 26(3):875–898, 2024.
- 21 Travis Greene, David Martens, and Galit Shmueli. Barriers to academic data science research in the new realm of algorithmic behaviour modification by digital platforms. *Nature Machine Intelligence*, 4(4):323–330, 2022.
- 22 Travis Greene, Galit Shmueli, and Soumya Ray. Taking the person seriously: Ethically aware is research in the era of reinforcement learning-based personalization. *Journal of the Association for Information Systems*, 24(6):1527–1561, 2023.
- 23 Galit Shmueli, Ali Tafti. How to “improve” prediction using behavior modification, *International Journal of Forecasting*, Vol. 39(2), pp. 541–555, 2023.
- 24 Sebastian Gabel and Artem Timoshenko. Product choice with large assortments: A scalable deep-learning model. *Management Science*, 68(3):1808–1827, 2022.
- 25 Tobias Hatt and Stefan Feuerriegel. Early detection of user exits from clickstream data: A markov modulated marked point process model. In *Proceedings of The Web Conference 2020*, pages 1671–1681, 2020.
- 26 Tobias Hatt and Stefan Feuerriegel. Detecting user exits from online behavior: A duration-dependent latent state model. *arXiv preprint arXiv:2208.03937*, 2022.
- 27 Lukas-Valentin Herm, Kai Heinrich, Jonas Wanner, and Christian Janiesch. Stop ordering machine learning algorithms by their explainability! a user-centered investigation of performance and explainability. *International Journal of Information Management*, 69:102538, 2023.
- 28 Christian Janiesch, Patrick Zschech, and Kai Heinrich. Machine learning and deep learning. *Electronic Markets*, 31(3):685–695, 2021.
- 29 Milan Kuzmanovic, Dennis Frauen, Tobias Hatt, and Stefan Feuerriegel. Causal machine learning for cost-effective allocation of development aid. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
- 30 Mathias Kraus, Stefan Feuerriegel, and Asil Oztekin. Deep learning in business analytics and operations research: Models, applications and managerial implications. *European Journal of Operational Research*, 281(3):628–641, 2020.
- 31 Mathias Kraus, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Data-driven allocation of preventive care with application to diabetes mellitus type II. *Manufacturing & Service Operations Management*, 26(1):137–153, 2023.
- 32 Abdurahman Maarouf, Stefan Feuerriegel, and Nicolas Pröllochs. A fused large language model for predicting startup success. *European Journal of Operational Research*, 2024.
- 33 David Martens and Foster Provost. Explaining data-driven document classifications. *MIS Quarterly*, 38(1):73–100, 2014.
- 34 Christof Naumzik, Stefan Feuerriegel, and Markus Weinmann. I will survive: Predicting business failures from customer ratings. *Marketing Science*, 41(1):188–207, 2022.
- 35 Yilmazcan Ozyurt, Tobias Hatt, Ce Zhang, and Stefan Feuerriegel. A deep Markov model for clickstream analytics in online shopping. In *ACM Web Conference*, 2022.
- 36 Yilmazcan Ozyurt, Mathias Kraus, Tobias Hatt, and Stefan Feuerriegel. AttDMM: An attentive deep Markov model for risk scoring in intensive care units. In *ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021.
- 37 Foster Provost and Tom Fawcett. *Data Science for Business: What you need to know about data mining and data-analytic thinking*. O’Reilly Media, Inc, 2013.
- 38 David Rügamer, Chris Kolb, and Nadja Klein. Semi-structured distributional regression. *The American Statistician*, 78(1):88–99, 2024.
- 39 Yash Raj Shrestha, Shiko M. Ben-Menahem, and Georg von Krogh. Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4):66–83, 2019.
- 40 Klaus Schwab. *The fourth industrial revolution*. World Economic Forum, 2017.
- 41 Galit Shmueli. Research dilemmas with behavioral big data. *Big Data*, 5(2):98–119, 2017.
- 42 Julian Senoner, Bernhard Kratzwald, Milan Kuzmanovic, Torbjørn H Netland, and Stefan Feuerriegel. Addressing distributional shifts in operations management: The case of order fulfillment in customized production. *Production and Operations Management*, 32(10):3022–3042, 2023.

- 43 Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- 44 Julian Senoner, Torbjørn Netland, and Stefan Feuerriegel. Using explainable artificial intelligence to improve process quality: Evidence from semiconductor manufacturing. *Management Science*, 68(8):5704–5723, 2022.
- 45 Julian Senoner, Simon Schallmoser, Bernhard Kratzwald, Stefan Feuerriegel, and Torbjørn Netland. Explainable AI improves task performance in human-AI collaboration. *arXiv:2406.08271*, 2024.
- 46 Daniel Tschernutter and Stefan Feuerriegel. Data-driven dynamic police patrolling: An efficient Monte Carlo tree search. *European Journal of Operational Research*, 321(1):177–191, 2025.

## 6.53 Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES)

**Editors:** Kerstin I. Eder, Timo Höning, Maja Hanne Kirkeby, Daniel Mosse, and Max Plauth  
**Seminar No.** 24351

Date: August 25–30, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.8.36

© Creative Commons BY 4.0 license

© Kerstin I. Eder, Timo Höning, Maja Hanne Kirkeby, Daniel Mosse, and Max Plauth



**Participants:** Gustavo Alonso, Antonio Carlos Schneider Beck Filho, Pierre Bennorth, Liliانا Cucu-Grosjean, Ada Diaconescu, Kerstin I. Eder, Tamar Eilam, Michael Engel, João Paulo Fernandes, Clemens Grelck, Daniel Gruss, Timo Höning, Geerd-Dietger Hoffmann, Romain Jacob, Jonas Juffinger, Fiodar Kazhamiaka, Maja Hanne Kirkeby, Michael Kirkedal Thomsen, Sven Köhler, David Kohnstamm, Julia Lawall, Silverio Martínez-Fernández, Daniel Mosse, Frank Mueller, Jukka K. Nurminen, Robin Ohs, Vinicius Petrucci, Max Plauth, Fabian Rauscher, June Sallou, Joao Saraiva, Andreas Schmidt, Gunnar Schomaker, Wolfgang Schröder-Preikschat, Sibylle Schupp, Jennifer Switzer, Manuel Vögele, Samuel Xavier-de-Souza

The increasing carbon footprint of computing systems represents a critical and immediate challenge for the computing community, particularly as global digitalization accelerates. As computational demands grow – driven by data-intensive applications such as artificial intelligence – addressing the environmental impact of computing has become essential. The Dagstuhl Seminar 24351 on Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES) focused on advancing the field’s understanding of how carbon emissions can be effectively measured, managed, and reduced across all layers of heterogeneous computing systems.

This seminar (PEACHES) brought together leading researchers and practitioners from the fields of computer science, software engineering, and environmental sustainability. The seminar participants explored 5 essential topics through continued

group discussions. The participants explored boundaries, challenges, and possible new methods and techniques for: 1) obtaining **Carbon Transparency**, 2) reaching **Net-Zero in the Age of AI and Machine Learning**, and 3) achieving **Carbon-Aware Computing, Storage, and Communication**. In addition, we discussed and explored **Disruptive Paradigms**, focusing on innovative approaches needed to achieve net-zero carbon goals, rather than relying on incremental improvements. The seminar also featured enlightening **Carbon-aware Computing Hackathons**, introducing and discussing the latest research tools in the area, e.g., to reduce energy consumption via undervolting and tracking carbon emissions in applications. The seminar contributes to advancing the understanding of how software innovations can help mitigate the global environmental impact of computing.

## 6.54 Conversational Agents: A Framework for Evaluation (CAFE)

**Editors: Christine Bauer, Li Chen, Nicola Ferro, and Norbert Fuhr**  
**Seminar No. 24352**

Date: August 25–30, 2024 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.14.8.53

© Creative Commons BY 4.0 license

© Christine Bauer, Li Chen, Nicola Ferro, and Norbert Fuhr

**Participants:** Avishek Anand, Christine Bauer, Timo Breuer, Li Chen, Guglielmo Faggioli, Nicola Ferro, Ophir Frieder, Norbert Fuhr, Hideo Joho, Jussi Karlgren, Johannes Kiesel, Lien Michiels, Andrea Papenmeier, Mark Sanderson, Scott Sanner, Benno Stein, Johanne Trippas, Karin Verspoor, Martijn C. Willemsen



In this Dagstuhl Perspectives Workshop, a general model for the evaluation of CONversational Information ACcess (CONIAC) systems was developed: Conversational Agents Framework for Evaluation (CAFE).

The framework starts from the assumption that a CONIAC system will be able to

- (i) *interact* with users more naturally and seamlessly,
- (ii) *guide* a user through the process of *refining* and *clarifying* their needs,
- (iii) *aid decision-making* by making *personalized recommendations and information* while being able to *explain* them, and
- (iv) *generate, retrieve and summarize* relevant information.

CAFE distinguishes six major elements of an evaluation design:

- **Stakeholder goals.** Stakeholders of a CONIAC system may have diverse goals that might or might not be directly accessible to system designers or evaluators and must often be implicitly inferred in evaluation. CONIAC systems might also have multiple goals ranging from end users having (in-)direct information needs, to platforms deploying CONIAC systems interested in content usage, user engagement, impression generation, and user retention, to name a few.
- **Tasks.** CONIAC involves tasks characterized by an information need (which may be specific or rather vague), human involvement, goal orientation, and mixed initiative between the user and the system. While some tasks and information needs may benefit from introducing a conversationally competent system, others may not, depending on the complexity of the task or need.
- **User aspects.** When developing an evaluation framework for CONIAC systems, it is crucial to consider user-specific aspects, such as preferences, specialized needs, expertise types, and background characteristics, which may make con-

versational systems more beneficial than non-conversational alternatives.

- **Criteria.** The scope of evaluation can range from single-turn interactions to entire conversations and long-term system usage, each requiring different criteria for assessment. Additionally, the temporal dimension, which examines how the system's performance changes over time, is a critical factor that can intersect with both stationary and dynamic properties. Criteria may be system-centric, user-centric, or both. The former regard hardware and software aspects like e. g. efficiency, accuracy, comprehensiveness, and verifiability. For the latter, we can distinguish between conversation-oriented (like e. g. adaptability, coherence, fluency), content-oriented (like e. g. continuance, controllability, perceived accuracy, understandability), and consequences-oriented measures (like e. g. addiction, benevolence, decision quality, confidence, trust).
- **Methodology.** In addition to the standard distinction of user-focused and system-focused methodologies, our evaluation framework categorizes evaluation methodologies also according to the employed time model – a dimension especially relevant for CONIAC. This dimension ranges from stationary methodologies like single-interaction experiments to methodologies like controlled lab studies that allow for continuous measurements such as physiological ones.
- **Measures.** Finally, we allow for measures that typically focus on the system's ability to provide accurate, relevant, and timely information during interactions. Measures include objective measures of effectiveness and subjective notions such as perceived effectiveness or user satisfaction (e. g., self-reported satisfaction). By incorporating both objective as well as subjective (self-reported) measures, evaluators can better understand the system's strengths and areas for improvement.

When designing an evaluation, the first step is to identify the stakeholders and their goals that need to be addressed. Based on the goals, the user tasks to be studied in the evaluation have to be defined, as well as the user aspects to be considered. The central element of an evaluation are the criteria to be focused on,

which can be determined by the stakeholder goals. The chosen criteria restrict the range of possible evaluation methods (e. g. any user-centric criterion requires the involvement of actual users in the evaluation procedure). Finally, an appropriate measure has to be defined for any quantitative criterion.

## 6.55 Artificial Intelligence and Formal Methods Join Forces for Reliable Autonomy

**Editors: Nils Jansen, Mykel Kochenderfer, Jan Kretinsky, Jana Tumova, and Maris Galesloot**  
**Seminar No. 24361**

Date: September 1–6, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.1

© Creative Commons BY 4.0 license

© Nils Jansen, Mykel Kochenderfer, Jan Kretinsky, Jana Tumova, and Maris Galesloot

**Participants:** Eric Atkinson, Joydeep Biswas, Wendelin Böhmer, Jonathan DeCastro, Clemens Dubslaff, Khen Elimelech, Georgios Fainekos, Maris Galesloot, Anna Gautier, Kush Grover, Sofie Haesaert, Nick Hawes, Nils Jansen, Sebastian Junges, Ruya Karagulle, Sydney Katz, Michaela Klauack, Mykel Kochenderfer, Johannes Köhler, Jan Kretinsky, Hanna Kurniawati, Bruno Lacerda, Lars Lindemann, Sara Magliacane, Dimitra Panagou, David Parker, Karinne Ramirez Amaro, Christian Schilling, Philipp Schillinger, Thiago D. Simão, Stephen Smith, Matthijs Spaan, Hazem Torfah, Jana Tumova, Cristian Ioan Vasile, Abhinav Verma, Esen Yel



### ■ Introduction

Artificial intelligence (AI) is a disruptive force. With growing applications in fields like healthcare, transportation, game playing, finance, and robotics, AI systems and methods are entering our everyday lives. Such tight interaction with AI requires serious safety, correctness, and reliability considerations. Recently, the field of safety in AI has triggered a vast amount of research.

The area of formal methods (FM) offers structured and rigorous ways to reason about the correctness of a system. Techniques range from model learning, over testing to formal verification. As an example of the application of verification in AI, solving techniques like SAT or SMT solving help to assess the robustness of neural networks. Model checking is a prominent verification technique that proves the system's correctness with respect to formal specifications. Therefore, it provides ample grounds to address the safety concerns of and provide guarantees in AI systems.

Via a diverse program with ample space for open yet guided discussion, this Dagstuhl Seminar brought together experts from the fields of artificial intelligence, formal methods, and robotics. The predecessor was the Dagstuhl Seminar 18121 *Machine Learning and Model Checking Join Forces*, on the interaction of AI (specifically machine learning) and FM (specifically model checking), mainly from the theoretical perspective. In contrast, this seminar broadened the scope methodologically, including AI and FM generally, but focused more narrowly on the application of reliable robotics in real-world robotics, which then drives the concrete questions towards increased applicability of combined research in AI and FM.

### ■ Formal Methods, AI and Robotics

Formal methods (FM) and AI differ not only in the type of problems they tackle but also in their focus on the desirable properties of the methods. While AI focuses primarily on

*effectiveness and efficiency*, formal methods focus on *reliability and safety*. Obviously, in both cases, the other properties are not negligible for actual real applications. Consequently, over the last decade, achieving the best of both worlds has become even more important. On the one hand, ML experienced a huge boom accompanied by challenges orthogonal to the tradition, such as provable reliability or explainability; on the other hand, FM was transformed by this process, with AI becoming both an object of FM and a useful technique to be used within FM.

In particular, the problem of controlling complex systems is so fundamental that many approaches have evolved independently in several communities, depending on the particular setting and type of the system, e.g., in artificial intelligence, formal methods, control theory, or robotics. Recently, many of the specifics such as uncertainty, safety-criticality, cyber-physical nature, or complex dynamics (combining continuous and discrete aspects) are present all at once, such as in robots. This rich combination of challenges calls for *combining the techniques* of the respective disciplines in fundamentally new ways so that real complex systems can be *efficiently and reliably* controlled and deployed in society.

In addition to the challenges related to integrating AI and formal methods in general, several application-specific challenges arise, reflecting that robots are physical systems: Robot models are highly complex, and often, it is not even possible to derive or use an analytic model (e.g., in contact-rich manipulation). Consequently, the dimensionality and the degree of uncertainty in these systems are very high. These challenges were the inspiration for the topics discussed in this seminar.

### ■ Seminar Structure

**Tutorials** The first half of the week aimed at creating a mutual theoretical understanding and an interdisciplinary view on the topics of the seminar. To this end, the seminar had five plenary

tutorials, each focusing on a different aspect of the subject areas relevant to the seminar. In particular, the following topics were presented:

- Model-Based Reinforcement Learning, by Wendelin Böhmer,
- An Overview of Probabilistic Verification Techniques, by Sebastian Junges,
- Formal Methods & Verification in Robotics & Autonomous Systems Engineering @ Bosch, by Michaela Klauck,
- Robot Learning, by Georgia Chalvatzaki, and,
- Conformal Prediction, by Lars Lindemann.

**Working Groups** In the latter half of the week, there was ample time for focused discussions among participants in breakout sessions. These sessions had various aims, such as further increasing interdisciplinary understanding of the relevant concepts, discussing new research ideas, and identifying open problems. The working groups (breakout rooms) consisted of the following topics, from which a subset is elaborated in more detail in Section 4 of the full report:

- Human Models: Modelling Humans Towards AI Decision-Making,
- Models and Programmatic Reinforcement Learning,

- Partially Observable Markov Decision Processes (POMDPs),
- Robust Markov Decision Processes (Robust MDPs),
- Hierarchies of Specifications,
- Runtime Monitoring (with a particular focus on latent verification),
- Reliability,
- Explainability.

**Talks** In addition to the tutorials and working groups, many participants presented their lines of work, blue-sky ideas, open problems, and big challenges in dedicated presentation slots. The abstracts can be found in the overview section of the talks.

The seminar provided a platform for early-career researchers among the participants to present their research in so-called *spotlight talks* throughout the week. In these talks, the participants presented recent work and their general research direction to the seminar audience. Thanks to the spotlights, participants got an impression of the research directions of the other participants.

In addition to the spotlight talks, the seminar participants disseminated open problems, blue-sky ideas, and big challenges. Many of these talks resulted in topics for further discussion.

## 6.56 Next-Generation Secure Distributed Computing

**Editors:** Aniket Kate, Julian Loss, and Kartik Nayak

**Seminar No.** 24362

Date: September 1–6, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.22

© Creative Commons BY 4.0 license

© Julian Loss

**Participants:** Ittai Abraham, Georgia Avarikioti, Renas Bacho, Adithya Bhat, Sourav Das, Nico Döttling, Sisi Duan, Bryan Ford, Juan A. Garay, Neil Giridharan, Aniket Kate, Aggelos Kiayias, Andrew Lewis-Pye, Chen-Da Liu-Zhang, Giuliano Losa, Julian Loss, Tal Moran, Kartik Nayak, Matthieu Rambaud, Filip Rezabek, Mohammad Sadoghi, Alin Tomescu, Sophia Yakoubov, Sravya Yandamuri, Fan Zhang



This seminar was dedicated toward improving our understanding of techniques and models in the quickly evolving space of distributed systems, in particular as found in the blockchain space. This space intersects in various different areas such as applied cryptography, distributed algorithms, and game theory. As such, our seminar brought together researchers with a background in one or more of these areas to discuss with and learn from each other. Over the course of five days, the seminar participants gave talks about their respective areas followed by a discussion. The talks ranged from presenting recent results to summarizing the progress of a field over several decades.

While the talks took place in the morning sessions, the afternoon sessions were dedicated to breakout discussions on topics related to the talks and chosen by the seminar participants. During these sessions, topics were discussed in greater detail, which led to new insights and research questions. At the end of each seminar day, one of the discussion participants was tasked with summarising the outcome of the discussion to the other seminar participants. This often led to further discussion on some specific points.

Broadly speaking, the topics of these discussions can be summarized as follows.

- Day 1:** The discussions for day one were on topics related to the blockchain space. Three different groups discussed recent advancements in the area of consensus protocols, decentralization, as well as a discussion on rational actors storing secrets in the context of blockchain protocols.
- Day 2:** The second day of discussions focused on topics related to new/non-traditional network and fault models, gossip layers for spreading messages in blockchain systems, as well as further discussion on game theoretic aspects in distributed systems (again with a focus on blockchains).
- Day 3:** There were no afternoon discussions on day 3. Instead, the participants enjoyed an excursion to Saarschleife and a dinner at a local brewery in the evening.
- Day 4:** On day 4, one group discussed the difference in consistency guarantees provided by primitives such as modern blockchain protocols and classic primitives such as byzantine broadcast. Another group discussed the purpose of and possible applications of trusted execution environments for storing secrets in blockchain systems.
- Day 5:** The seminar ended by noon, hence there were no afternoon discussions on day 5.

## 6.57 Extended Reality Accessibility

**Editors: Gerd Bruder, Thies Pfeiffer, Jeanine Stefanucci, and Dylan Fox**  
**Seminar No. 24371**

Date: September 8–13, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.45

© Creative Commons BY 4.0 license

© Gerd Bruder, Thies Pfeiffer, and Jeanine Stefanucci



**Participants:** Bobby Bodenheimer, Gerd Bruder, Sarah Creem-Regehr, Tiare Feuchtner, Faustina Hwang, Kiyoshi Kiyokawa, Katharina Krösl, Torsten Kuhlen, Eike Langbehn, Robert W. Lindeman, Luciana Nedel, Estella Oncins Noguer, Thies Pfeiffer, Steffen Puhl, Daniel Roth, Adalberto Simeone, Anthony Steed, Laura Trutoiu, Radu-Daniel Vatavu, Benjamin Weyers

In the past, when new technologies emerged, making them accessible was always an afterthought. However, retroactive accessibility is never as effective as accessibility that is built-in from the start. Simply put, we must act to make extended reality (XR) accessible now.

The expertise of the organizers and seminar invitees allowed us to focus on specific areas of accessibility including visual, auditory, and motor disabilities or impairments as well as cognitive impairments that may occur with very young or older users. Research questions that we focused on include (but were not limited to) the following:

- How can we make XR more accessible for those who are blind or have uncorrectable low vision?
- How can we display captions or augment XR experiences for deaf individuals?
- How can we make controllers for XR compatible with various motor impairments?
- How can we adjust XR displays to accommodate smaller head sizes of young users?

This Dagstuhl Seminar builds off of the ground-breaking work that has been done in this area, in particular the XR Access Initiative, which was founded to bring together people across industry, academia, advocacy organizations, and government. While these initiatives have been mostly confined to participants from the United States and Canada, this seminar allowed us to expand the reach of these initiatives to a more international audience that will bring new perspectives and solutions. Making XR accessible to everyone will depend on the involvement of researchers from around the world. In the Outcomes section of the full report, we introduce ideas generated at the seminar to make interdisciplinary fields more aware of the field of XR accessibility in order to create more research on the topic but also to increase awareness.

## 6.58 Explainable AI for Sequential Decision Making

**Editors: Hendrik Baier, Mark T. Keane, Sarath Sreedharan, Silvia Tulli, and Abhinav Verma**  
**Seminar No. 24372**

Date: September 8–11, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.67

© Creative Commons BY 4.0 license

© Hendrik Baier, Mark T. Keane, Sarath Sreedharan, Silvia Tulli, and Abhinav Verma

**Participants:** David Abel, Hendrik Baier, Ruth Mary Josephine Byrne, Rebecca Eifler, Claudia Goldman, Bradley Hayes, Tobias Huber, Mark T. Keane, Khimya Khetarpal, Benjamin Krarup, Pat Langley, Simon M. Lucas, Anna Lukina, Samer Nashed, Sriraam Natarajan, Ann Nowé, Ron Petrick, Mark Riedl, Silvia Rossi, Wojciech Samek, Lindsay Sanneman, Julian Siber, Sarath Sreedharan, Mohan Sridharan, Silvia Tulli, Stylianos Loukas Vasileiou, Abhinav Verma



We work with AI and rely on AI for more and more decisions that influence our lives. To serve increasingly urgent goals such as enabling transparency, enhancing collaboration, and increasing trust in AI, the research area of explainable AI (XAI) has rapidly developed in recent years. However, the focus of XAI to date has largely been on explaining the input-output mappings of “black box” models such as neural networks, which have been seen as the central problem for the explainability of AI systems. While these models are certainly important, intelligent behavior often extends over time and needs to be explained and understood as such. The challenge of explaining sequential decision making (SDM), such as that of robots collaborating with humans or software agents engaged in complex ongoing tasks, has only recently gained attention. We may have AIs that can beat us in Go, but can they teach us how to play? We may have search and rescue robots, but can we effectively communicate with them and coordinate missions in the field?

Initial attempts at making the behavior of SDM algorithms more understandable have recently appeared in different fields such as classical AI planning, reinforcement learning, multi-agent systems, or logic-based argumentation – but often focused on ad hoc solutions to specific problems, with an emphasis on area-specific terminology and concepts developed in isolation of other fields. Many of these approaches are also restricted in their scope, for example to explanations of isolated, single actions that do not address the full complexity of SDM; or to summaries of entire agent policies, which are often too high-level to be helpful. To truly trust an AI agent and collaboratively work with it towards human goals, and to increase successful AI adoption and acceptance in many fields from robotics to logistics, and from production planning to smart cities, we need considerable progress in this new field of XAI for SDM (or X-SDM).

Under-researched challenges for X-SDM include for example:

XAI for complex decisions, e.g., on plans or policies instead of single output labels; conversational XAI that continuously interacts with users over time, aiming to understand and support them; contestable and collaborative XAI, which can successfully work with users in areas where neither user nor AI is omniscient nor infallible; and flexible decision-making for XAI, able to adapt to users, respect their autonomy, and go beyond one-size-fits-all explanations. This Dagstuhl Seminar aimed to identify and clarify such challenges that are unique to, or of particular relevance to, explainability in sequential decision-making settings. It used complementary perspectives of researchers from different communities such as reinforcement learning, planning, game AI, robotics, and cognitive science, which historically use different theoretical foundations and computational approaches. The aim of the seminar was to move towards a shared understanding of the field, by first building a taxonomy that unifies our perspectives, and then developing a common roadmap for moving X-SDM forward.

This seminar was organized in two parts: in the first part, breakout groups were formed based on the primary research communities of the participants. These groups were encouraged to identify key terminology and definitions in their area, which were later summarized in plenary sessions, and combined into a first sketch of a shared taxonomy. Based on this understanding of the field of explainable sequential decision making, breakout groups in the second part were then formed by participants interested in a particular aspect of future work, and their results were again summarized and combined in later plenary sessions to develop a first sketch of a roadmap for the field. The seminar was accompanied by spotlight talks throughout, giving insight into the work of individual participants. This report gives an overview of the talks, outlines the discussions of all breakout groups, and summarizes sketches of the taxonomy and the roadmap. The

interest in the seminar was high and the participants were very enthusiastic to contribute and connect across research fields; the most frequently expressed regret was that the seminar did not last long enough to flesh out more details of the work we started. As

the organizers, we therefore consider this seminar a great success and are looking forward to the resulting collaborations and impact to the field.

## 6.59 Algebraic and Analytic Methods in Computational Complexity

**Editors:** Markus Bläser, Shubhangi Saraf, Ronen Shaltiel, and Jacobo Torán  
**Seminar No. 24381**

Date: September 15–20, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.104

© Creative Commons BY 4.0 license

© Markus Bläser, Shubhangi Saraf, Ronen Shaltiel, and Jacobo Torán

**Participants:** Robert Andrews, Vikraman Arvind, Vishwas Bhargava, Markus Bläser, Andrej Bogdanov, Peter Bürgisser, Harry Buhrman, Igor Carboni Oliveira, Prerona Chatterjee, Arkadev Chattopadhyay, Eshan Chattopadhyay, Gil Cohen, Radu Curticapean, Dean Doron, Klim Efremenko, Michael A. Forbes, Anna Gál, Rohit Gurjar, Shuichi Hirahara, William Hoza, Christian Ikenmeyer, Valentine Kabanets, Neeraj Kayal, Pascal Koiran, Antonina Kolokolova, Michal Koucký, Donald Kougang Yombi, Mrinal Kumar, Nutan Limaye, Meena Mahajan, Or Meir, Rafael Mendes de Oliveira, Noga Ron-Zewi, Kilian Rothmund, Rahul Santhanam, Ramprasad Saptharishi, Shubhangi Saraf, Nitin Saxena, Rocco Servedio, Ronen Shaltiel, Amir Shpilka, Jad Silbak, Srikanth Srinivasan, Sathyawageeswar Subramanian University of Cambridge, GB, Sébastien Tavenas, Roei Tell, Thomas Thierauf, Jacobo Torán, Christopher Umans, Jeroen Zuiddam



The seminar brought together more than 50 researchers covering a wide spectrum of complexity theory. The focus on algebraic methods showed the great importance of such techniques for theoretical computer science. We had 27 talks, most of them lasting about 35 minutes, leaving ample room for discussions. We also scheduled several longer talks of a more introductory nature. One of the days, we also had a much appreciated open problem session.

Computational complexity is a fundamental and active area of research that has produced some of the most well known results in theoretical computer science in recent years. Here we discuss a few broad themes that represented the focus areas of our seminar, mentioning some of the given talks in these areas.

### ■ Circuit complexity

Boolean circuits are one of the most fundamental models of computation. Due to their combinatorial nature, they seem more amenable to formal analysis than the uniform models such as Turing machines. The classical lower bound techniques of Razborov and Smolensky are algebraic: they work by first approximating  $AC^0[p]$  circuits (constant-depth circuits with AND, OR, NOT, and counting modulo prime  $p$  gates) by low-degree polynomials, and then proving that certain functions (like Majority) are not well correlated with such polynomials. The Fourier expansion of a Boolean function and its representation as a real multilinear polynomial as well as other analytic tools have been added in the last years to the bag of tools used for the analysis of Boolean circuits. Several seminar talks discussed recent results on circuit complexity.

*Igor Carboni Oliveira* established that efficient indistinguishability obfuscation for multi-output circuits necessarily incurs an additive size overhead of  $\Omega(s/\log s)$  under reasonable complexity assumptions.

In his talk, *William Hoza* studied hardness amplification in the context of two well-known “moderate” average-case hardness results for  $AC^0$  circuits.

### ■ Algebraic complexity

A class of circuits especially suitable for the use of algebraic techniques is that of *arithmetic* or *algebraic circuits*. These are circuit models that compute polynomial functions using gates performing arithmetic operations (additions, subtractions, multiplications, and divisions). Two fundamental complexity measures for arithmetic circuits are *size* and *depth*.

*Robert Andrews* talked about constant-depth arithmetic circuits for linear algebra problems. He explained a new algorithm for the GCD that uses a combination of polynomial interpolation and Newton’s identities for symmetric polynomials.

*Vishwas Bhargava* talked about read-once branching programs, a special class of arithmetic circuits. His focus was the order in which variables can appear in such branching programs.

Exponential lower bounds against sums of ordered set-multilinear branching programs were shown by *Prerona Chatterjee*. Her bound stays superpolynomial for degree  $\omega(\log n)$ , complementing recent results by Bhargava, Dwivedi, and Saxena.

*Chris Umans* talked about the group-theoretic approach to fast matrix multiplication and how to extend the group-theoretic framework to the setting of infinite groups.

In his presentation, *Michael Forbes* explained how to extend the recent breakthrough of Limaye, Srinivasan, and Tavenas, who gave the first super-polynomial lower bounds against low-depth algebraic circuits, to fields of small characteristic. This answers an open question by Limaye, Srinivasan, and Tavenas.

In his talk, *Rafael Mendes de Oliveira* surveyed the connections between Sylvester-Gallai configurations and complexity theory as well as identity testing.

*Peter Bürgisser* studied the computational complexity of robust generalizations of orbit problems, which amount to approximating the distance of orbits in  $C^n$  up to a factor  $\gamma \geq 1$ . These algorithms run in polynomial time if and only if a version of the famous number-theoretic abc-conjecture holds.

In his talk, *Nitin Saxena* showed that the border of bounded-top-fanin depth-3 circuits can be computed by polynomial-size algebraic branching programs, which has applications especially in identity testing and lower bounds.

*Amir Shpilka* proved that for polynomials  $f, g \in C[x]$  such that  $g$  divides  $f$ , the bit complexity of the quotient polynomial  $h = f/g$  is bounded by the sparsity of  $g$  and the sparse representation size of  $f$ .

## ■ Pseudorandomness and derandomization

The theory of pseudorandomness studies explicit constructions and applications of “random-like” objects of combinatorial or algebraic type. The common feature of such objects is that they are easy to construct by random sampling, but it is a very important problem to get efficient *deterministic* constructions.

*Arkadev Chattopadhyay* discussed the power of randomness for computing total Boolean functions in two models of computation, query algorithms and 2-party communication protocols.

In his talk, *Gil Cohen* used analytic techniques to explore a more refined analysis of the zig-zag graph product, that leverages the entire spectrum of the graph.

*Dean Doron* studied almost Chor–Goldreich (CG) random sources, and constructed deterministic condensers with constant entropy gap for them. As a consequence, one can simulate any randomized algorithm with small failure probability using almost CG sources with no multiplicative slowdown.

Seedless condenser were studied by *Eshan Chattopadhyay*, proving several new results for seedless condensers in the context of three related classes of sources.

*Rohit Gurjar* gave a characterization of principal minor equivalence and constructed a deterministic polynomial time algorithm to check if two given matrices are principal minor equivalent, which can be viewed as some sort of identity testing.

## ■ Error correcting codes

Error-correcting codes, particularly those constructed from polynomials, lie at the heart of many of the most significant results in computational complexity, like interactive proofs, PCPs, hardness amplification, or explicit constructions.

In his talk, *Mrinal Kumar* showed that univariate multiplicity codes can be list-decoded up to capacity in nearly linear time.

In her talk, *Noga Ron-Zewi* surveyed the theory behind the recent result that Reed-Solomon codes with random evaluation points are list-decodable up to capacity with optimal list size, and then discussed how it can be extended to the whole class of polynomial ideal codes.

*Jad Sidblak* studied codes against channels that are computationally bounded. He constructed codes with optimal rate that are explicit (assuming circuit lower bounds).

*Srikanth Srinivasan* presented low-degree local correction algorithms for constant-degree polynomials from  $\{0, 1\}^n$  to any field.

*Roei Tell* talked about simulations of interactive proof systems by non-interactive proof systems. He explained how to leverage approaches from recent works in derandomization to address this problem.

## ■ Further applications of algebraic methods

In addition to the focus areas mentioned above, we had a number of talks that considered applications of algebraic methods in other areas. In particular, we had a significant number of contributions that dealt with learning problems.

The talk of *Neeraj Kayal* explored machine learning tasks from the perspective of algebraic complexity, particularly using techniques based on arithmetic circuit lower bounds.

*Pascal Koiran* gave a new constructive uniqueness theorem for tensor decomposition. As a result, he obtained the first efficient algorithm for the so-called overcomplete decomposition of generic tensors of order 3.

*Michal Koucký* presented an efficient randomized algorithm for computing hierarchical Hamming sketches, which has applications in the construction of sketches for the edit distance.

The topic of the talk by *Meena Mahajan* was lower bounds for the polynomial calculus over non-Boolean domains. She showed how to improve a recent work by Sokolov by using an asymmetric gadget, obtaining a stronger lower bound that works over any field.

*Kilian Rothmund* presented an  $NC$  algorithm for testing whether a graph is minimally rigid when the input graph is  $K_{3,3}$ -free.

*Rocco Servedio* showed that the supremum of any centered Gaussian process can be approximated to any arbitrary accuracy by a finite dimensional Gaussian process, where the dimension of the approximator just depends on the target error.

## ■ Conclusion

The talks of the seminar ranged over a broad assortment of subjects with the underlying theme of using algebraic and analytic techniques. It was a very fruitful meeting and has, we hope, initiated new directions in research. Several participants specifically mentioned that they appreciated the particular focus on a common class of techniques (rather than end results) as a unifying theme of the seminar. We look forward to our next seminar.

## 6.60 Statistical and Probabilistic Methods in Algorithmic Data Analysis

Editors: Aristides Gionis, Matteo Riondato, and Eli Upfal

Seminar No. 24391

Date: September 22–27, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.127

© Creative Commons BY 4.0 license

© Aris Gionis, Matteo Riondato, and Eli Upfal

**Participants:** Florian Adriaens, Aris Anagnostopoulos, Luca Becchetti, Fabian Brandt-Tumescheit, Gianmarco De Francisci Morales, Alessandro Epasto, Esther Galbrun, Aristides Gionis, Shahrzad Haddadan, Ravi Kumar, Silvio Lattanzi, Stefano Leonardi, Sebastian Lüderssen, Alessio Mazzetto, Pauli Miettinen, Stefan Neumann, Lutz Oettershagen, Rasmus Pagh, Andrea Pietracaprina, Giulia Preti, Geppino Pucci, Matteo Riondato, Will Rosenbaum, Larry Rudolph, Ilie Sarpe, Mauro Sozio, Mahito Sugiyama, Nikolaj Tatti, Eli Upfal, Fabio Vandin, Sergei Vassilvitskii, Yllka Velaj, Geoffrey Webb, Anthony Wirth



The development of efficient methods for the analysis of data, including graphs, time series, and transactional datasets, is a key challenge in computer science. The availability of larger, richer datasets gathered about all aspects of human life, society, and nature, and which take a variety of forms (e.g., graphs, time series, tabular data) has created the need for such methods, which would find applications in areas such as biology, finance, and computational social sciences. It has also created challenges that must be tackled in order to develop efficient methods, such as the following.

- **Trade-offs between data, computation, and information:** As datasets with *trillions* of points become available, *how much data do we really need* to (approximately) answer thousands or more exploratory analysis questions on them? We need new algorithms that carefully balance the minimum amount of data needed to obtain the desired information at the required accuracy, with the computation needed to extract it.
- **Balancing theory and practice in data science:** How can we satisfy the needs for *theoretical and statistical guarantees* on one side and *practical efficiency* on the other? How do we connect concepts from different disciplines to exploit their power in useful algorithms? Developments from *all areas of computer science and statistics* should be used to obtain data analysis methods that are *efficient from both a computational and a statistical point of view*.
- **Enabling efficient and statistically-sound scientific discoveries:** The practice of science *requires* testing *multiple hypotheses* on the same, often large, data. For example, the Human Protein Reference Database graph has  $\approx 19k$  proteins and  $\approx 37k$  interactions between them. Scientists want to understand the significance of small connected subgraphs in this network, representing pathways in cancer cells. There are  $> 10^{13}$  subgraphs of size 8, each corresponding to an hypothesis. We need *computationally-efficient methods* that

scale to *trillions of hypotheses*, while *controlling for different measures of statistical guarantee on false discoveries*, e.g., the Family-Wise Error Rate (FWER) and the False Discovery Rate (FDR).

- **Enhancing responsible and unbiased data analysis:** There are numerous techniques for analyzing and processing datasets. While these techniques have overall been effective in extracting patterns in datasets, the choice of an appropriate technique and parameters for analyzing a dataset is often more art than science. Different techniques, each implementing a variety of heuristics, may give different results when applied to the same data. Moreover, many of these techniques do not provide an interpretable mechanistic explanation for their performance, and most do not provide a rigorous measure of statistical significance or robustness of the analysis. A principled approach to data analysis, based on well founded mathematical and statistical concepts, enforces objectivity and unbiased results. Such an approach enhances the effectiveness of evidence-based medicine, policy, and social applications of data analysis.

During the Dagstuhl Seminar, the participants engaged in lively presentations of recent results on these topics, ranging from the theory of machine learning, to bias reduction, to pattern mining, to approximation and online algorithms for problems arising in very different settings and with many different applications. An Open Problem session gave the participants a chance to share interesting problems for the community. While the research presentations took the majority of the schedule, there were plenty of chances for the participants to network, socialize, and engage in new and ongoing collaborations, which is, in our opinion, the real advantage of attending a Dagstuhl Seminar in person. Especially positive were the interactions between participants at different stages of their career, from institutions in different countries, and from both academia and industry.

## 6.61 Fair Division: Algorithms, Solution Concepts, and Applications

Editors: Evangelos Markakis, Ruta Mehta, and Yair Zick

Seminar No. 24401

Date: September 29 – October 4, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.145

© Creative Commons BY 4.0 license

© Evangelos Markakis, Ruta Mehta, and Yair Zick



**Participants:** Hannaneh Akrami, Georgios Amanatidis, Haris Aziz, Nawal Benabbou, Umang Bhaskar, Vittorio Bilo, Georgios Birmpas, Paula Böhm, Sylvain Bouveret, Giorgos Christodoulou, Alon Eden, Edith Elkind, Amos Fiat, Sushmita Gupta, Martin Hoefer, Alexandros Hollender, Ayumi Igarashi, Panagiotis Kanellopoulos, Elias Koutsoupias, Jérôme Lang, Evangelos Markakis, Nicolas Maudet, Ruta Mehta, Hervé J. Moulin, Keziah Naggita, Paula Navarette Diaz, Dominik Peters, Alexandros Psomas, Nidhi Rathi, Rebecca Reiffenhäuser, Sanjukta Roy, Ulrike Schmidt-Kraepelin, Alkmini Sgouritsa, Golnoosh Shahkarami, Max Springer, Adaku Uchendu, Rohit Vaish, Giovanna Varricchio, Toby Walsh, Yair Zick

Fair division concerns the allocation of resources to a set of interested entities, according to some fairness criteria. Fair division has remained an active research area over the years, attracting the attention of several scientific disciplines, such as mathematics, computer science, game theory, and political science. Especially within the last two decades, this area has gradually drawn significantly more attention, due to the emergence of novel solution concepts, algorithmic techniques, and promising applications. It has also gained increasing popularity within computer science, since many of the field's key questions are inherently algorithmic. Consequently, there has been a notable growth of the relevant literature by now, spanning numerous fascinating research directions.

The main objective of the seminar was to bring together a leading set of researchers and discuss fundamental economic and computational challenges in fair division. Consequently, the seminar focused on three main categories of research topics, grouped as follows: 1) fundamental questions on the existence and the efficient computation of solution concepts in fair division; 2) the interplay of fair division with other fields such as mechanism design and machine learning; and 3) applications of fair division.

Most of the participants were academics from computer science departments or from research centers. Some were from other disciplines such as economics and mathematics. We also had the pleasure to host 2 HLF (Heidelberg Laureate Forum) participants, that is, 2 young researchers in Computer Science, in line to Dagstuhl's cooperation with the HLF.

The seminar started on Monday with an introductory session, in which the participants shared their main research interests. This session was very well-received by the participants, and it initiated discussions directly from the start. The subsequent program included 3 tutorials of roughly one hour each, along with 22 contributed talks that were solicited from the participants and lasted around 25 minutes each. In many of these talks there were

lively discussions. On most days, the program also provided free time between lunch and the afternoon talks, which was used for research and collaborative meetings.

The program of the tutorials was as follows:

1. On Monday, H. Moulin provided a tutorial on the concept of "Guarantee" in fair division.
2. On Tuesday, A. Psomas gave a tutorial on the theory and practice of fair food allocation.
3. On Wednesday, we had our last tutorial by Y. Zick on the theory and practice of course allocation.

Regarding the contributed talks, a large percentage of the presentations focused on solution concepts in fair division and questions related to their existence, their complexity, and their approximability. There was a substantial interest on the EFX notion, which is currently the most popular and at the same time the most intriguing fairness criterion. For this notion, G. Christodoulou presented results on EFX allocations on graphs, a special class of instances defined using an appropriate graph theoretic representation. A. Sgouritsa also presented new results on pushing the frontier on approximate EFX allocations, showing cases where one can achieve a  $2/3$ -approximation guarantee (an open problem for general instances). E. Markakis also presented families of special cases where one can obtain  $2/3$ -approximation algorithms or even better guarantees. Following a different direction, N. Rathi discussed epistemic notions of EFX allocations and their existence.

Beyond the EFX criterion, there were several talks on other fairness notions. U. Bhaskar presented results on the existence of EF1 allocations and H. Akrami talked about breaking the barrier of  $3/4$  for approximate MMS allocations. Both of these criteria, EF1 and MMS, have been studied extensively within the last years. Furthermore, V. Bilo presented results on achieving envy-freeness via selling items. In a different direction, H. Aziz discussed

the pursuit of *best of both world* fairness, where the goal is to achieve both *ex ante* and *ex post* guarantees. Another model of fair division concerns the case where the items to be allocated are chores instead of goods. There were two presentations covering this topic, namely by R. Vaish, whose talk focused on a model of interval scheduling for chores, and by A. Igarashi, who also presented an actual application for sharing household chores, implemented in Japan. Yet another model was presented by A. Hollender, who discussed open problems for the continuous setting in fair division, referred to as cake cutting. In terms of applying all these notions in relevant areas, S. Roy and N. Benabbou talked about fairness as well as group-fairness in house allocation problems. Finally, E. Elkind talked about incorporating fairness concepts into temporal voting.

Another focus area was the study of incentives for strategic agents and mechanism design aspects. The seminar had 2 sessions on these topics. A. Eden presented results on the recently introduced paradigm of learning-augmented mechanism design for truthful fair division mechanisms via predictions. G. Shahkarami also talked about mechanism design with predictions for facility location problems. G. Amanatidis focused on an

equilibrium analysis of the famous Round-Robin algorithm with strategic agents. In yet another direction, S. Bouveret presented results for strategyproof and fair algorithms based on picking sequences. Finally, T. Walsh talked about the derandomization of mechanisms that retain good normative properties regarding the effects of strategic behavior.

Other directions that were covered by the talks included problems that are motivated by the further development of machine learning and AI. N. Maudet discussed a model of explainable fair division, regarding locally envy-free allocations and R. Mehta talked about fairness and incentives in federated learning.

Finally, J. Lang gave a presentation towards achieving a Sandewallian taxonomy for fair division papers. Essentially this amounts to classifying the works in this field, based on a variety of features, such as the topic or the flavor of the obtained results.

Overall, the seminar was a big success. We believe it will stimulate new and very fruitful collaborations in the near future. We also received very enthusiastic feedback from many participants, which is also reflected in the survey conducted by Dagstuhl.

## 6.62 Greening Networking: Toward a Net Zero Internet

**Editors:** Alexander Clemm, Dirk Kutscher, Michael Welzl, Cedric Westphal, and Noa Zilberman  
**Seminar No. 24402**

Date: September 29 – October 2, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.9.167

© Creative Commons BY 4.0 license

© Alexander Clemm, Dirk Kutscher, Michael Welzl, Cedric Westphal, and Noa Zilberman



**Participants:** Jari Arkko, Carsten Bormann, Wen Cai, Luis Miguel Contreras Murillo, Sawsan El Zahr, Te-Yuan Huang, Tony John, Jong Hoon Kwon, David Oran, Daniel Schien, Farzad Tashtarian, Cedric Westphal, Hammad Zafar

The Dagstuhl Seminar focused on

- identifying and prioritizing the most impactful networking improvements to reduce carbon emissions;
- defining action items for a carbon-aware networking research agenda; and
- fostering and facilitating research collaboration in order to reduce carbon emissions and to positively impact climate change.

As a major result from the seminar, the following problems and topics for future research have been identified:

1. characterizing the Internet footprint on carbon emissions accurately;
2. understanding attributional and consequential accounting of carbon emissions in networked systems; and
3. identifying potential solutions to give network systems more flexibility in better supporting energy grids and connecting to renewable energy sources.

One of the concrete results of this seminar is a list of technologies and research opportunities for which we estimated the potential impact and time horizon.

### ■ Overview of the Seminar: Motivation and Conceptual Framing

What is the path towards a Net Zero Internet? Reducing CO<sub>2</sub> emissions to combat climate change is one of the greatest challenges facing mankind. An overarching goal is to make human activities sustainable, where sustainability is defined as:

*Meeting the needs of the present without compromising the ability of future generations to meet their own needs (UN Brundtland Commission, 1987).*

Networking technologies are a key part of this challenge – as a provider for solutions (such as enabling teleworking and reducing travel), but also as a contributor to CO<sub>2</sub> emissions, for example, through significant power consumption. Thus, it is becoming important to make networks themselves “greener” and devise solutions that result in less carbon intensity while continuing to meet the increasing network demands and service requirements.<sup>41</sup> Complicating the matter, energy consumption is related to CO<sub>2</sub> emissions in ways that vary greatly by network operators (through the use of “clean” energy).

Many of today’s improvements are driven by general advances in computing hardware as well as in transmission technology (antennas, lasers). While this is where the biggest opportunities for energy reduction may lie, it is important to extend questions of greenness to other aspects of carbon reduction, both temporal and spatial, and to other layers in the networking stack – to the data and control plane, to routing and traffic forwarding, to the ways in which networks are organized and deployed. For example, can network protocols be designed in ways that make them more carbon-efficient, e.g. by routing along a path that is primarily powered by renewable energy, so as to reduce CO<sub>2</sub> emissions with proper path selection? Can control planes be designed in a way that allows to rapidly switch off resources that may currently not be needed without compromising other important goals such as network resilience or security? How much overhead is introduced by current architectures and protocols for cloud computing, media streaming, and content distribution, and what is the energy conservation potential of alternative systems? What

<sup>41</sup> Reducing traffic demands could potentially reduce network emissions. However, this is a social issue that was considered outside of the scope of this seminar.

protocol advances could enable greener networking solutions? How can networks be optimized not just for QoS or utilization but for low environmental impact?

This Dagstuhl Seminar brought together researchers of different perspectives with a vested interest on this topic, to exchange ideas and explore the possibilities for synergy in their respective research directions. We identified and prioritized the most impactful networking improvements to reduce carbon emissions, defined action items for a carbon-aware networking research agenda, and fostered/facilitated research collaborations that result in improved techniques to reduce carbon emissions.

This seminar was organized into sessions pertaining to these themes, with presentations from leading researchers drawn from the participants. Researchers in the six main topics – devices, protocols, metrics, measurement, architecture, and management – were invited to participate.

The seminar and the full report follow the same structure: Section 3.3 of the full report discusses the bigger picture of

energy saving and carbon neutrality for applications, systems, and stakeholders. Section 4.4 of the full report discusses potential technical energy-saving approaches for networking functions such as routing, traffic engineering, power saving mechanisms, and technology-specific mechanisms, such as for wireless communications. Section 5.4 of the full report discusses overarching concepts in technology related to cyclical characteristics, such as control loops, i.e., the application of control logic to optimize environmental sustainability of solutions, as well as the actual life-cycle of network equipment. Finally, Section 7 of the full report summarizes the outcomes of the seminar in three tables listing different proposed research directions.

We acknowledge that other topics and perspectives could have been discussed, and that this selection is a result of many parameters: the expertise of the participants, the subjectivity of the organizers, or the amount of time available for discussions. Other ways to frame the conversation could be the basis for another future seminar.

## 6.63 New Tools in Parameterized Complexity: Paths, Cuts, and Decomposition

**Editors:** Fedor V. Fomin, Dániel Marx, Saket Saurabh, and Roohani Sharma  
**Seminar No.** 24411

Date: October 6–11, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.10.1

© Creative Commons BY 4.0 license

© Fedor V. Fomin, Dániel Marx, Saket Saurabh, and Roohani Sharma



**Participants:** Akanksha Agrawal, Aditya Anand, Matthias Bentert, Édouard Bonnet, Joseph Cheriyan, Éric Colin de Verdière, Eduard Eiben, Fedor V. Fomin, Robert Ganian, Petr A. Golovach, Bart Jansen, Eun Jung Kim, Yusuke Kobayashi, Christian Komusiewicz, Tuukka Korhonen, Stefan Kratsch, Madhumita Kundu, Euiwoong Lee, Bingkai Lin, William Lochet, Daniel Lokshtanov, Dániel Marx, Neeldhara Misra, Pranabendu Misra, Jesper Nederlof, Daniel Neuen, Fahad Panolan, Marcin Pilipczuk, Michał Pilipczuk, Thatchaphol Saranurak, Ignasi Sau Valls, Saket Saurabh, Roohani Sharma, Ramanujan Sridharan, Giannos Stamoulis, Vaishali Surianarayanan, Stefan Szeider, Erik Jan van Leeuwen, Magnus Wahlström, Michał Włodarczyk, Jie Xue

### ■ Description of the Seminar: Topics and goals

Parameterized Complexity is an alternative approach of handling computational intractability. The main idea of the approach taken by the Parameterized Complexity is to analyze the complexity in finer detail by considering additional problem parameters beyond the input size and expresses the efficiency of the algorithms in terms of these parameters. In this framework, many NP-hard problems have been shown to be (fixed-parameter) tractable (FPT) when certain structural parameters of the inputs are bounded.

In the last three decades, there has been tremendous progress in understanding which problems are fixed-parameter tractable and which problems are not (under standard complexity assumptions). For all these years the central vision of the Parameterized Complexity has been to provide the algorithmic and complexity-theoretic toolkit for studying multivariate algorithmics in different disciplines and subfields of Computer Science. To achieve this vision several algorithmic and complexity theoretic tools such as *polynomial time preprocessing*, *aka, kernelization*, *color-coding*, *graph-decompositions*, *parameterized integer programming*, *iterative compression*, or *lower bounds methods based on assumptions stronger than  $P \neq NP$*  have been developed. These tools are *universal* as they not only helped in the development of the core of Parameterized Complexity but also led to its success in other subfields of Computer Science such as Approximation Algorithms, Computational Social Choice, Computational Geometry, problems solvable in P (polynomial time) to name a few.

In the last few years several decade old open problems in Parameterized Complexity have been resolved. These have resulted in several new algorithmic tools for the core of Parameterized Complexity. These include tools such as iterative and local improvement methods for graph decomposition, methods arising from extremal combinatorics and graph theory, flow

augmentation, faster algorithms for solving integer programs on  $n$  variables, and new algebraic methods. A natural question is to extend these tools in different directions and explore the limits and applicability of these new tools. Thus,

*the main objective of the Dagstuhl Seminar was to initiate the discussion on extension, limits and applicability of newly developed tools arising from paths, cuts, and decomposition.*

One of the seminar's central goals was to facilitate a fruitful dialogue between researchers working at the core of Parameterized Complexity and those from Mathematical Programming, Computational Linear Algebra, Graph Theory, and Combinatorics, who had contributed to recent advances in parameterized algorithms. The Dagstuhl event enabled participants to explore possibilities for developing new tools and techniques that emerged from this collaboration.

Next, the seminar presented a few concrete examples of newly developed tools and techniques from various domains of Parameterized Complexity, which formed the focal points of discussion.

- **Width Parameters: Treewidth and Twinwidth.**
- **Tools Based on Extremal Combinatorics: Paths and Rainbow Matching.**
- **Cut Based Tool: Flow Augmentation.**
- **Mathematical Programming.**
- **Algebraic Methods.**

### ■ Related Seminars

This Dagstuhl Seminar could be considered as a continuation of the Dagstuhl Seminar series on parameterized algorithms and complexity. The previous seminars in this series are New Horizons in Parameterized Complexity (seminar 19041), Randomization in Parameterized Complexity (seminar 17041),

Optimality and Tight Results in Parameterized Complexity (seminar 14451), Data Reduction and Problem Kernels (seminar 12241), Parameterized complexity and approximation algorithms (seminar 09511), Structure Theory and FPT Algorithmics for Graphs, Digraphs and Hypergraphs (seminar 07281), Exact Algorithms and Fixed-Parameter Tractability (seminar 05301), Fixed Parameter Algorithms (seminar 03311), and Parameterized Complexity (seminar 01311).

### ■ Organization of the Seminar

During the five-day seminar, 48 researchers from theoretical computer science, mathematical optimization, and operations research convened. Attendees ranged from senior scientists to postdoctoral scholars and advanced doctoral candidates, creating a rich environment for both mentorship and innovation.

The seminar featured 22 presentations of varying lengths. Six keynote speakers—Dániel Marx, Michał Pilipczuk, Euiwoong Lee, Tuukka Korhonen, Jie Xue, and Daniel Lokshtanov—delivered 60-minute talks that provided overviews of state-of-the-art methods and showcased recent breakthroughs in their respective

areas. The remaining slots were filled with shorter, 30-minute talks covering various topics.

At the beginning of the week, open problem sessions encouraged participants to share challenges and spark collaborative research. The schedule also included ample free time, which attendees used for productive discussions and joint work, fostering new ideas and potential research partnerships.

### ■ Outcome

Organizers and participants regarded the seminar as a great success. It successfully brought together the relevant research communities, facilitated the sharing of state-of-the-art results, and enabled discussions on the major challenges in the field. The talks were not only excellent but also highly stimulating, prompting active engagement in working groups during afternoons and evenings. Particularly noteworthy was the participation of younger researchers (postdocs and PhD students), who integrated seamlessly and contributed to the seminar's collegial and productive atmosphere.

## 6.64 SAT and Interactions

**Editors: Olaf Beyersdorff, Laura Kovács, Meena Mahajan, and Martina Seidl**  
**Seminar No. 24421**

Date: October 13–18, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.10.22

© Creative Commons BY 4.0 license

© Olaf Beyersdorff, Laura Kovács, Meena Mahajan, and Martina Seidl



**Participants:** Albert Atserias, Olaf Beyersdorff, Ilario Bonacina, Florent Capelli, Leroy Nicholas Chew, Anupam Das, Susanna de Rezende, Katalin Fazekas, Mathias Fleury, Pascal Fontaine, Marlene Gründel, Clemens Hofstadler, Kaspar Kasche, Phokion G. Kolaitis, Wietze Koops, Konstantin Korovin, Laura Kovács, Massimo Lauria, Meena Mahajan, Barnaby Martin, Stefan Mengel, Claudia Nalon, Jakob Nordström, Tomáš Peitl, Florian Pollitt, Michael Rawson, Adrian Rebola-Pardo, Rahul Santhanam, Dominik Alban Scheder, Tanja Schindler, Stephan Schulz, Martina Seidl, Friedrich Slivovsky, Luc Spachmann, Martin Suda, Stefan Szeider, Neil Thapen, Cesare Tinelli, Jacobo Torán, Sophie Tourret, Marc Vinyals, Heribert Vollmer, Andrei Voronkov

The problem of deciding whether a propositional formula is satisfiable (SAT) is one of the most fundamental problems in computer science. Its theoretical significance derives from the Cook-Levin Theorem, identifying SAT as the first NP-complete problem. Since then SAT has become a reference for an enormous variety of complexity statements, among them the celebrated P vs NP problem.

There are many generalisations of SAT to logics such as quantified Boolean formulas (QBF), modal and first-order (FO) logics. Often these logics present harder satisfiability problems (e.g. PSPACE-complete for QBF), but can express many practically relevant problems more succinctly, thus applying to more real-world problems.

Due to its practical implications, intensive research has been performed to solve SAT problems in an automated fashion. The last decades have seen the development of practically efficient algorithms for SAT, QBF, and further logics, and their implementation as solvers, which successfully solve huge industrial instances.

As the fourth in its series, the Dagstuhl Seminar took a broad perspective on the theory of SAT, encompassing propositional logic, QBF, and first-order theorem proving. Its main aim was to bring together researchers from different areas of activity in SAT and first-order logic, including computational complexity, proof complexity, proof theory, theorem proving, and solving, so that they can communicate state-of-the-art advances and embark on a systematic interdisciplinary interaction.

The Dagstuhl Seminar placed particular emphasis on the three following fields: propositional logic (complexity, proof complexity, solving), QBF (proof complexity and solving), and FO theorem proving. A particularly novel feature was the interaction of the communities active in proof complexity and solving of SAT/QBF (the propositional logics) with the first-order theorem proving community. There appeared to be overall

consensus among the participants that this interchange of ideas between SAT solving techniques and first-order theorem proving was very stimulating and might particularly prove useful towards further efficient implementations of first-order proof rules.

To facilitate interactions between participants from the different fields, the seminar included a number of survey talks to introduce neighbouring communities to the main notions, results, and challenges of the represented areas. The following survey talks were given towards the beginning of the seminar:

- Marc Vinyals: SAT Solving and Proof Complexity;
- Friedrich Slivovsky: QBF Solving and Proof Complexity;
- Cesare Tinelli: An introduction to Satisfiability Modulo Theories;
- Stephan Schulz: First-order Theorem Proving.

Each of these surveys was accompanied by one or more sessions with contributed talks dedicated to recent specific results of the field.

The seminar also included an open problem session where participants discussed open research directions and specific problems. The following topics were discussed:

- Stephan Schulz: Can we achieve good engineering and long-term viability of systems?  
Software projects often get abandoned over time, especially if the original authors leave. We are looking for technical and organisational solutions to mitigate this.
- Sophie Tourret: Beyond critical.  
There are established techniques to compare the hardness of propositional formulas. Can we find analogous techniques for first-order logic and SMT? Particularly interesting are cases outside the decidable fragments of these theories.
- Neil Thapen: Where is symmetry breaking in TFNP?  
Symmetry breaking techniques can be understood as the optimisation problem of finding a lexically minimal assignment, which is in TFNP. We know it is in PLS, it might be in

CLS, but can we at least show that it is not in FP? This has implications on the strength of Extended Frege.

- Adrian Rebola-Pardo: How should we design future proof formats?

There is a variety of practical proof formats that need to be suitable for verification and querying. Is it possible to design good universal proof formats? Important considerations include binary encodings, non-clausal representations, specific addition and deletion rules, the needs of incremental SAT solvers, and parallel proof checking.

- Florent Capelli: Properties of a hypergraph measure. The  $\beta$ -hyperorderwidth is a purely graph-theoretic measure on hypergraphs. How does it compare to established hypergraph measures? For example, there is a hypergraph that has  $\beta$ -hyperorderwidth 1, but its incidence graph has treewidth  $n$ . Is this possible the other way around? Can we generalize the definition of  $\beta$ -hyperorderwidth?

The seminar included ample time for discussions and informal interactions, a feature that appeared to be largely welcomed and productively used. On Wednesday afternoon we organised a traditional well-attended hike. On Thursday evening, we had a joyful music night with contributions from Sophie Turret, Ilario Bonacina, Dominik Scheder, Florent Capelli, and Kaspar Kasche. They played music by Marin Marais, Georg Philipp Telemann, Wolfgang Amadeus Mozart, Francis Poulenc, and George Gershwin.

The organisers believe that the seminar fulfilled their original high goals: the talks were well received and triggered many discussions. Many participants reported about the inspiring seminar atmosphere, fruitful interactions, and a generally positive experience. The organisers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and excellent support in the arrangement of a very successful and productive event.

## 6.65 Automated Programming and Program Repair

**Editors: Claire Le Goues, Michael Pradel, Abhik Roychoudhury, and Shin Hwei Tan**  
**Seminar No. 24431**

Date: October 20–25, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.10.39

© Creative Commons BY 4.0 license

© Shin Hwei Tan



**Participants:** Earl T. Barr, Islem Bouzenia, Yuriy Brun, Cristian Cadar, Celso G. Camilo-Junior, Satish Chandra, Chunyang Chen, Tse-Hsun Chen, Zimin Chen, Andreea Costea, Premkumar T. Devanbu, Alexander Frömmgen, Ahmed E. Hassan, Sungmin Kang, Dongsun Kim, Claire Le Goues, Yiling Lou, Fernanda Madeiral, Matías Martínez, Martin Monperrus, Nikhil Parasaram, Michael Pradel, Nikitha Rao, Abhik Roychoudhury, André Silva, Gustavo Soares, Gang (Gary) Tan, Lin Tan, Shin Hwei Tan, Yingfei Xiong, Jinqiu Yang, He Ye, Jooyong Yi, Jie Zhang, Lingming Zhang

Automated tools that generate and improve code promise to fundamentally change software development. For example, there is a recent trend towards automated code generation from large language models, as evidenced by the capabilities of Codex/Copilot, ChatGPT, and GPT-4. These models, and other techniques, such as search-based and semantic analysis-based techniques, have the potential to automate significant parts of today's software development process. In particular, there are promising techniques for automated programming and automated program repair. Automated programming refers to techniques that suggest newly written code, e.g., in the form of code completion tools. The capabilities of such tools have increased from moderately successful single-token predictions just a few years ago to predicting entire functions with relatively high accuracy. Techniques for automated programming include large language models that predict code based on natural language specifications of the intended behavior.

Automated program repair refers to a suite of techniques for automated rectification of errors or vulnerabilities in programs. Automated program repair technologies were originally developed for reducing the debugging effort for manually written code. In other words, automated program repair technologies are meant to boost developer productivity in locating and fixing errors in manually written code. Automated programming and automated program repair strongly overlap in terms of their goals and the techniques used to achieve these goals. Both streams of research aim at generating correct source code while having to cope with limited knowledge of the behavior this source code is meant to have. Since formal specifications of correct program behavior are typically not available, both techniques try to infer specifications from various program artifacts, such as large code corpora, past program versions, natural language documentation, or various executions of the program. To address the challenge of predicting likely correct code, both streams of research combine techniques

from machine learning, search-based approaches, and semantic code analysis.

Despite these similarities, the subcommunities working on automated programming and automated program repair are only partially aware of each other's most recent techniques. This seminar set out to explore the intersection of these two fields in order to foster collaborations between them. For example, to discuss recent work and potential future work in the following directions: (1) Apply program repair to fix code generated by code completion models. The code generated by large language models often leaves significant room for improvement in terms of correctness, thereby raising the question whether automated program repair techniques can be used for last-mile improvement of code that was automatically generated by large language models. (2) Apply the generate-and-validate paradigm from program repair to the code completion problem. For example, such techniques can repeatedly generate code completion candidates and validate them by running test suites. (3) Apply language model-based code generators to the program repair problem. Once the location of a bug has been (heuristically) determined, large language models can predict candidate code snippets for replacing the incorrect code. (4) Use the ability of large language models to infer the intended behavior of code from natural language information embedding in the code. For example, we plan to discuss techniques that specify the intended behavior in the form of assertions or test cases, which can then guide automated program repair. (5) In addition to predicting (fixed) code, generate evidence that the final code is trustworthy. Such evidence may take the form of tests generated along with the code, or other certificates obtained from formal reasoning.

Thus, to discuss topics at the intersection of automated programming and program repair, we had Dagstuhl Seminar 24431. In a five-day seminar with 33 participants from both academia and industries (e.g., Microsoft and Google), we held

a series of talks and three panel discussions. The seminar concluded with more than 20 talks and three panel discussions. Overall, the seminar stimulated quite a few discussions where researchers initiated some future research directions and potential international collaborations.

Before the seminar, all participants had received an invitation to give a talk of a flexible duration (i.e., lightning update that is around 5 minutes, short talk that is around 10 minutes or long talk that is around 25 minutes). More than 20 participants replied positively to the invitation, resulting in a great variety of talks given by many participants. The first day of the seminar (i.e., October 21, 2025) started with an introduction by the organizers and a self-introduction by all the participants. Then, a few short talks and longer talks (more than 25 minutes) were given by participants. The first day ended with a panel discussion on “Benchmarks for LLM Code Generation”. On October 22, 2025, there were several talks followed by a panel discussion on “LLM-beyond just coding-assistance”. On October 23, 2025, several talks took place in the morning, followed by an excursion to Mettlach and Villa Borg after lunch. On October 24, 2025, a few inspiring talks took place, followed by a panel discussion on “Obstacles for deploying program repair techniques”.

Overall, the seminar has received very positive feedback from the participants both personally and formally (via email). Notably, one participant sent an email to one of the organizers saying that “ It was my best Dagstuhl Seminar last October, and I really appreciate your organizing of the seminar once again”, demonstrating that the seminar has been quite successful in leaving a good impression in comparison to other Dagstuhl Seminars that the participants have attended. Meanwhile, a few participants have complimented Dagstuhl on the diversity of the social events held (e.g., excursion, the treetop walk, and sauna), and the babysitting services provided for participants attending the seminar with young children.

In terms of collaborations, there are a few actionable topics for collaborations that have been discussed. An opinion piece of AI Software Engineer titled “AI Software Engineer: Programming with Trust” is now available.<sup>42</sup> Another potential collaboration is

a critical review on benchmarks crafted by AI communities (i.e., SWE-Bench). Meanwhile, AutoCodeRover (presented by one of the organizers in Dagstuhl), which was an NUS spinoff, has on February 19, 2025, been officially acquired by SonarSource, a leader in code quality via its static analysis solutions.

The seminar focused on the following key themes:

- Topics at the intersection of automated programming and automated program repair, analyzing progress in both fields.
- Understanding common mistakes in automatically generated code.
- Discussing the theme of “**Trusted Automated Programming**”, which focuses on:
  - How automatically generated code can be made more trustworthy.
  - How to generate evidence that improvements to auto-generated code maintain trustworthiness.
  - How to decide, based on such evidence, when to incorporate automatically generated code into an existing software project with a stable code-base.
- Important challenges in automated program repair and automated programming in general.
- Using large language models (LLMs) beyond just coding assistance.
- Obstacles in deploying program repair techniques in real-world settings.

The seminar also identified several critical challenges:

1. The problem of curating widely accepted benchmarks for code generation that serve both the Software Engineering and AI communities.
2. The problem of designing evaluation criteria that effectively assess the quality and reliability of auto-generated code.
3. The challenges and opportunities in applying LLM-based techniques beyond traditional automated program repair (APR).
4. The obstacles in training developers to effectively use program repair techniques in real-world software development.

<sup>42</sup> <https://arxiv.org/abs/2502.13767>

## 6.66 Behavioural Metrics and Quantitative Logics

**Editors:** Barbara König, Radu Mardare, Prakash Panangaden, and Jurriaan Rot  
**Seminar No.** 24432

Date: October 20–25, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.10.58

© Creative Commons BY 4.0 license

© Barbara König, Radu Mardare, Prakash Panangaden, and Jurriaan Rot



**Participants:** Giorgio Bacci, Giovanni Bacci, Corina Cirstea, Raphaëlle Crubillé, Josée Desharnais, Justin Hsu, Clemens Kupke, Radu Mardare, Renato Neves, Jurriaan Rot, Wojciech Różowski, Qiyi Tang, Paul Wild

Behavioural equivalences are central in the analysis and verification of concurrent systems. They come with a very well-developed meta-theory including logical characterisations, efficient algorithms, game-theoretic perspectives, algebraic characterisations, and generalisations to a wide variety of state-based systems. Two processes or system states are behaviourally equivalent if they are indistinguishable from the point of view of an external observer. Behavioural equivalences have for instance been used for comparing a system with its specification, for the analysis of cryptographic protocols, for the verification of model transformations, for efficient comparison of non-deterministic automata and in programming language semantics.

In recent years, behavioural metrics and quantitative logics specifying quantitative aspects of systems have received considerable attention. Previous approaches are qualitative and address the question “do they match?” whereas metrics address the question “to what degree? how much?”. Indeed, a metric measures how far apart two systems are in their behaviour, while a quantitative logic evaluates the degree to which a state satisfies a formula. For quantitative systems such as Markov chains, Markov decision processes and probabilistic automata, a behavioural equivalence is often too strict: small deviations in quantitative information, such as probabilities, can cause two states that intuitively behave very much alike to be inequivalent in a formal sense. Indeed, behavioural equivalence misses *robustness* under small perturbations in such systems. Metrics and quantitative logics provide a more nuanced approach than qualitative behavioural equivalence.

As behavioural equivalences, behavioural metrics can be characterized via fixpoint equations. In addition there are alternative characterizations of metrics, either via modal logics or via spoiler-duplicator games. The focus of this Dagstuhl Seminar is in particular on quantitative logics, where evaluating a formula on a state does not yield a boolean value, but a (real) number, intuitively indicating to which degree a state satisfies a formula.

This is connected to a quantitative Hennessy-Milner property which is supposed to hold for the behavioural distance.

There are various applications in model-checking, differential privacy, hybrid systems and learning. Several challenges in this area have been identified: studying suitable metrics and their corresponding logics, generalizing to the setting of coalgebras by parameterizing the branching type of the system under consideration, developing methods for quantitative algebraic reasoning, and finding efficient methods for computing behavioural metrics.

This Dagstuhl Seminar provided a forum for researchers working in all areas of behavioural metrics and quantitative logics, to discuss the state-of-the-art and further developments, and in particular to address applications in various domains, including machine learning.

The topics discussed at the seminar included all aspects of behavioural metrics and quantitative logics. In particular:

- Various approaches to define behavioural metrics, including characterizations via fixpoint equations, logics and games and their relations.
- Quantitative logics and their expressiveness.
- Methods and theories for quantitative equational reasoning.
- Algorithms for (compositionally) computing behavioural metrics or distinguishing formulas and their efficiency.
- Applications of behavioural metrics and quantitative logics, including – but not limited to – model-checking, differential privacy, hybrid systems and learning.

The seminar programme left room to address the above individual topics, but the main aim was to connect these topics, which previously have been studied by different research groups and published in different venues. Establishing such connections proved to be the main way forward, both to solidify the field and to identify new research problems and opportunities for applications. In particular, a key challenge was to connect recent approaches

to quantitative logics, game characterisations, and quantitative equational theories, and to generalise these connections and algorithmic perspectives to enable their application to a wide range of models. We also focused on practical applications and promoted the usefulness of behavioural metrics. The seminar stimulated interactions between more theoretical scientists with researchers working on application-oriented aspects.

We invited four representatives of the different communities to give tutorial talks in order to introduce fundamental concepts and techniques. Specifically, the following four tutorial talks took place on the first day of the seminar:

- Franck van Breugel: Behavioural metrics
- Clemens Kupke: Coalgebras
- Giorgio Bacci: Quantitative equational reasoning
- Pablo Castro: Behavioural metrics for Reinforcement Learning

On Tuesday and Thursday we organized the following working groups in order to discuss more specific topics which were of interest to a substantial part of the participants:

- Generalised Kantorovich-Rubinstein Duality
- Quantitative Algebras/Coalgebras/Modal Logic
- Applications of Behavioural Metrics
- Fixpoints
- Locksteps

The seminar made it possible for many participants to connect and discuss joint collaborations and research projects. The seminar at Dagstuhl was also the opportunity for Matteo Mio, Lutz Schröder, Henning Urbat and Paul Wild to coordinate between the pre-proposal and the full proposal submission for an ANR-DFG project. The full proposal will be submitted in March 2025. This project also involves other applicants from FAU Erlangen-Nürnberg and ENS-Lyon and it aims at furthering our understanding of the theory of quantitative algebras. This topic was the main subject of many talks of this seminar. Further ideas for follow-up events and outreach activities were developed in the wrapup session (for the report on this session see the full Dagstuhl Report).

The organizers would like to thank all the participants and speakers for their inspiring talks and many interesting discussions. Furthermore we would like to acknowledge Florence Clerc who helped to write and prepare this report. A special thanks goes to the Dagstuhl staff who were a great help in organizing this seminar.

## 6.67 Machine Learning Augmented Algorithms for Combinatorial Optimization Problems

**Editors:** Deepak Ajwani, Bistra Dilkina, Tias Guns, and Ulrich Carsten Meyer  
**Seminar No.** 24441

Date: October 27–31, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.10.76

© Creative Commons BY 4.0 license

© Deepak Ajwani, Bistra Dilkina, Tias Guns, and Ulrich Carsten Meyer



**Participants:** Deepak Ajwani, Brandon Amos, Senne Berden, Quentin Cappart, Eanna Curran, Marianne Defresne, Michelangelo Dilligenti, Bistra Dilkina, Adam Elmachtoub, Aaron Ferber, Alexandre Forel, Emma Frejinger, Paul Grigas, Ernestine Großmann, Tias Guns, Mikhail Khodak, Sandra Kiefer, Alex Lindermayr, Michele Lombardi, Nikolai Maas, Irfan Mahmutogullari, Fredrik Manne, Johannes Meintrup, Ulrich Carsten Meyer, Sofia Michel, Nysret Musliu, Mathias Niepert, Siegfried Nijssen, Ryan O Connor, Manuel Penschuck, Adam Polak, María Dolores Romero Morales Copenhagen Business School, DK, Jens Schlöter, Christian Schulz, Darren Strash, Edward Tansley, Sylvie Thiébaux, Ali Vakilian, Jacobus G. M. van der Linden, Maurice Wenig, Neil Yorke-Smith

Combinatorial optimization problems are pervasive across critical domains, driving decades of research across diverse fields, including exact and approximation algorithms, algorithm engineering, operations research, and optimization solvers (such as mixed-integer linear programming and constraint programming solvers). In recent years, there has been a surge in research at the intersection between machine learning, optimization solvers, and algorithm engineering. While significant progress has been made, research efforts in this area remain fragmented across several distinct communities:

- “Learning to scale optimization solvers”: This community focuses on leveraging learning techniques to improve the performance of solvers, such as mixed-integer linear programming, constraint programming, and satisfiability solvers, often using generic, problem-agnostic features. This includes optimizing solver components and exploring “end-to-end” approaches that replace traditional solvers entirely.
- “Algorithm Engineering”: This community utilizes learning techniques to select the most effective components within algorithms, fill in the underspecified parts of the algorithms, and develop novel heuristics, often leveraging problem-specific features and algorithmic insights.
- “Algorithms with predictions”: This community develops algorithms that leverage predictions to improve performance, often treating the prediction system as a black box. These algorithms aim for good performance with accurate predictions while maintaining bounded worst-case performance with inaccurate predictions.
- “Decision-focused learning”: This research area focuses on systems where optimization decisions are made based on machine learning models, necessitating an integrated approach to avoid suboptimal solutions.

This seminar brought together researchers from these diverse

communities, fostering a dialogue on effectively combining algorithm engineering techniques, optimization solvers, and machine learning. The seminar facilitated the development of a shared vocabulary, clarifying similarities and distinctions between concepts across different research areas. Notably, many concepts with different names (e.g., “learning to prune,” “predict then search”, “finding kernel”, “finding a backbone”, “learning to branch”, “learning constraints”) were found to represent similar ideas. Key differences often stem from the source of features, with “learning to prune” typically relying on problem-specific, algorithm-based features, while “learning to branch” often utilizes problem-agnostic features derived from the solver run. This shared understanding fostered significant synergy and cross-fertilization across research areas.

The seminar featured discussions on recent advancements, at the intersection of machine learning, algorithm engineering, and combinatorial optimization. For instance, the tutorial on “ML-Guided Solvers for MIP” focused on recent advances in using contrastive learning for MILP solvers, the talk on “Neural Combinatorial Optimization – One Model to Solve Them All?” described recent advances in foundation models for sequencing combinatorial optimization and the talk on “Graph Learning for Planning” described very impressive results on the usage of learning techniques for planning problems. In addition, many participant talks such as “Progress and Open Questions on DFL and Constrained NNs” and “Machine Learning for Edge Contractions in Multilevel Graph Partitioning” highlighted interesting open problems in the area.

Furthermore, the seminar identified key research directions, including:

- Theoretical frameworks for machine learning augmented combinatorial optimization algorithms
- Learning to fill the underspecified parts of the algorithms
- Reinforcement learning for combinatorial optimization

- Contextual Stochastic Optimization and DFL
- Fast and/or interpretable machine learning techniques for combinatorial problems
- Scalability and Datasets for Decision Focused Learning
- Finding worst-case instances for algorithms using machine learning techniques

These directions are expected to serve as a valuable roadmap for advancing this exciting research area.

## ■ Seminar Overview

The seminar commenced with a vibrant round of introductions, allowing participants to share their research interests and backgrounds. The slides for this round of introductions were collected in advance. To foster interdisciplinary connections, participant order was randomized, encouraging interaction across the diverse research communities represented.

Following this, four insightful tutorials were presented, each from a different research community. These tutorials served to familiarize all participants with the unique vocabulary, key techniques, and recent advancements within each area, providing a common ground for subsequent discussions.

A series of engaging participant talks followed, showcasing cutting-edge research from across the field. The talk schedule was carefully randomized to maximize interaction and foster cross-pollination of ideas among the diverse research communities. This dynamic approach encouraged participants to step outside their immediate research areas and engage with novel perspectives.

Recognizing the importance of collective vision, participants were invited to submit key research directions and open challenges within the field. These valuable contributions were then clustered, enabling the identification of key research themes and areas of significant opportunity.

To delve deeper into these themes, three parallel discussion sessions were held. Each session featured three distinct discussion groups, allowing participants to select the group most aligned with their interests. This dynamic group composition, which changed daily, encouraged participants to engage with a variety of perspectives and fostered a truly inter-community exchange of ideas. It was particularly encouraging to observe that many discussion groups comprised members from all four research communities, demonstrating the successful integration of diverse perspectives throughout the seminar.

## 6.68 Machine Learning for Protein-Protein and Protein-Ligand Interactions

**Editors: Anne-Florence Bitbol, Jennifer Listgarten, and Tomas Pluskal**  
**Seminar No. 24451**

Date: November 3–8, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.11.1

© Creative Commons BY 4.0 license

© Tomas Pluskal, Anne-Florence Bitbol, and Jennifer Listgarten



**Participants:** Anne-Florence Bitbol, Sebastian Böcker, Alexandre Bonvin, Anton Bushuiev, Roman Bushuiev, Alessandra Carbone, Alberto Cazzaniga, Simona Cocco, Francesca Cuturello, Christian Dallago, Arne Elofsson, Sergei Grudinin, Ilia Igashov, Petr Kouba, Jessica Lanini, Andrew Leach, Jennifer Listgarten, Cyril Malbranke, Hiroshi Mamitsuka, Céline Marquet, Simon Mathis, Stanislav Mazurenko, Remi Monasson, Hunter Nisonoff, Armita Nourmohammad, Tomas Pluskal, Burkhard Rost, Juho Rousu, Alexander Schug, Josef Sivic, Martin Steinegger, Aalt-Jan van Dijk, Pablo Varas Pardo, Andrea Volkamer, Martin Weigt, Julius Wenckstern, Bruce Wittmann, Xiaotong Xu, Omri Yakir, Lenka Zdeborova

The Dagstuhl Seminar 24451, titled “Machine Learning for Protein-Protein and Protein-Ligand Interactions”, convened leading experts from computational biology, chemistry, and machine learning (ML) to explore advancements and challenges in understanding biomolecular interactions. This was the first seminar of its kind, and it aimed to address pressing issues such as integrating domain knowledge into ML models, ensuring data availability and quality, and fostering effective interdisciplinary collaboration. The event facilitated discussions on theoretical advancements, practical challenges, and future research directions in leveraging ML for protein science and drug discovery.

A central theme of the seminar was the exploration of ML techniques for predicting protein-protein and protein-ligand interactions with improved accuracy and interpretability. Topics included representation learning, generative modeling, and the role of inductive biases in structuring ML models. The emergence of protein language models (pLMs) was a particular highlight, as they offer sustainability and efficiency advantages over traditional sequence alignment-based methods. These models were recognized for their ability to enhance structural and functional predictions, particularly in cases where evolutionary data is scarce or unreliable.

Another critical discussion focused on benchmarking ML models for biomolecular interactions. Participants debated the limitations of current datasets, particularly in terms of negative data and class imbalances, and emphasized the need for standardized benchmarks. Reliable benchmarking was identified as essential for validating new approaches, particularly in enzyme design and protein interaction network analysis.

The seminar also tackled real-world challenges in applying ML to protein science. One of the most pressing issues was the scarcity of high-quality datasets, particularly for protein-ligand

interactions. Participants proposed new standards for curating balanced datasets and ensuring the inclusion of negative data to improve ML model training. Additionally, discussions highlighted the importance of integrating experimental constraints into ML workflows, which would enhance cost-effectiveness in protein engineering and directed evolution.

The integration of computational and wet-lab research was another key topic. Effective collaboration between these fields remains a major challenge due to differing methodologies and objectives. Strategies were proposed to improve synergy between computational modeling and experimental validation, including the development of interdisciplinary training programs and shared data repositories.

Inspired by previous Dagstuhl meetings, the seminar adopted a discussion-driven format, encouraging open exchanges and collaboration. Several actionable outcomes emerged, including initiatives to enhance data accessibility, develop standardized benchmarking frameworks, and refine ML model architectures for improved performance in biomolecular applications.

The event successfully bridged computational and experimental research, paving the way for future innovations in protein science and drug discovery. Moving forward, participants emphasized the need for continued interdisciplinary collaboration, the development of more reliable datasets, and the refinement of ML techniques to better capture the complexity of protein interactions. The outcomes of this seminar are expected to significantly influence the future of machine learning applications in biology and chemistry, setting the stage for groundbreaking advancements in the field.

## 6.69 Reframing Technical Debt

**Editors: Paris Avgeriou, Ipek Ozkaya, Heiko Kozirolek, Zadia Codabux, and Neil Ernst**  
**Seminar No. 24452**

Date: November 3–8, 2024 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.14.11.16

© Creative Commons BY 4.0 license

© Paris Avgeriou, Ipek Ozkaya, Heiko Kozirolek, Zadia Codabux, and Neil Ernst

**Participants:** Apostolos Ampatzoglou, Paris Avgeriou, Lodewijk Bergmans, Markus Borg, Alexandros Chatzigeorgiou, Marcus Ciolkowski, Zadia Codabux, Stefano Dalla Palma, Florian Deißeböck, Philippe-Emmanuel Douziech, Neil Ernst, Daniel Feitosa, Michael Felderer, Collin Green, Ciera Jaspan, Ron Koontz, Heiko Kozirolek, Christof Momm, Brigid O’Hearn, Ipek Ozkaya, Klaus Schmid, Carolyn Seaman, Tushar Sharma, Guilherme Horta Travassos, Roberto Verdecchia, Marion Wiese



Technical Debt is considered the “silent killer” of software projects. In a recent survey, software developers claimed they spend 13.5 hours weekly on Technical Debt, one-third of their working time<sup>43</sup>. The industry now has widespread consensus that managing Technical Debt should be treated as a core software engineering practice. Companies are increasingly incorporating Technical Debt management practices into their development processes [2], [6]. Many software quality tools now incorporate features to help software engineering teams visualize and triage quality issues within their code bases to assist Technical Debt management.

This advancement in industrial practice is paired with very vibrant research. The research community has produced a substantial body of knowledge on the topic of Technical Debt, especially in investigating the problem and its manifestation, understanding its urgency and impact, and proposing solutions to address it. This research is often performed in collaboration with industry, indicating its practical relevance but also the aligned interests of researchers and practitioners. The maturity of research in Technical Debt is evidenced by the increasing number of publications and systematic literature reviews since the first research paper was published in 2010 [1]. A recent tertiary study reviewing secondary studies in managing technical debt reported 532 unique research studies [3].

Despite the significant progress made in industry practices, commercial tools, and the research on the body of knowledge on Technical Debt Management, there are five critical gaps.

**Lack of tools.** While there is widespread use of static analysis tools to detect code quality issues, not all of them indicate Technical Debt, and there is significant confusion around their use. For example, several commonly used industry tools report

different results for simple measures such as size, complexity, file cycles, and package cycles [4]. In addition, the software engineering community is unclear on what a Technical Debt management tool is.

**Difficulty in understanding the impact of AI on Technical Debt.** There is tremendous momentum around developing software with AI-based tools whose implications on new forms of Technical Debt are unclear. Consequently, while the industry is driving the simplification of code generation with AI-augmented tools to develop vast new quantities of code and develop it fast, it also risks incurring large volumes of unanticipated Technical Debt. Furthermore, the reality of industrial software engineering is dealing with a significant amount of legacy software. In these heritage systems, accumulated Technical Debt exists, but teams often lack resources to address the underlying issues. A large portion of that Technical Debt was introduced many years prior; it remains hidden and undocumented and involves design or architecture issues that cannot be mined by analyzing source code.

**Lack of theory on Technical Debt economics.** The bottom line of Technical Debt management is the Return On Investment (ROI) for the business and value, as difficult as it is to quantify the financial bottom line reliably. Existing research, while aimed to address complicated concepts such as the principal and interest of Technical Debt, fails to recognize that a new software economics approach is needed to concretely capture the cost of ownership and value aspects of Technical Debt.

**Not focusing on architecture.** Industry challenges consistently emphasize the difficulties around managing architecture-level Technical Debt where the tradeoffs are more implicit

<sup>43</sup> <https://stripe.com/files/reports/the-developer-coefficient.pdf>

and complex. However, we observe the proverbial “lamp post” research, as researchers tend to focus on the easy part (code) rather than the more valuable but also more challenging part (architecture).

**Narrow scope and data.** There is not sufficient emphasis on the social side of Technical Debt Management. This is reflected in the limited research datasets: they only contain code artifacts and are not validated with human subjects. In addition, data sets that allow understanding of the socio-technical aspects of Technical Debt are lacking.

This Dagstuhl Perspectives Workshop brought together researchers, tool vendors, and software practitioners to analyze these gaps and reframe the field of technical debt, focusing on the following key challenge areas and guiding research questions:

- *Technical Debt as value-creation*: How can the value of Technical Debt be managed in an informed and conscious manner to meet business goals while still avoiding prohibitively high-interest payments in the future?
- *Elevating the role of architecture*: How can software architecture considerations be integrated into Technical Debt Management to elevate it beyond code-level issues?
- *Next-generation tooling*: How can new and existing tools better focus on Technical Debt Management, including taking advantage of AI-based capabilities where appropriate?
- *New perspectives on data collection*: How can forms of data beyond code analysis, such as architecture, documentation artifacts, and business artifacts, inform Technical Debt Management?
- *Socio-technical aspects*: Given that Technical Debt is often a matter of how individuals and teams operate, how can social aspects be integrated into Technical Debt Management?

#### ■ Workshop Format

Before the workshop, a key reading list was distributed to the attendees to familiarize them with the most recent advances in Technical Debt Management. The workshop itself featured four elements:

- *Plenary Sessions* which included sharing organizational details, lightning talks on the five key challenges by all participants, reporting from the break-out groups, brainstorming on the vision, and planning ahead.
- *Break-out Sessions* which included brainstorming, discussions, and write-up sessions of the state of the art, as well as the values, beliefs, and principles of the manifesto, and finally, the roadmap towards realizing the manifesto. Each session had five break-out groups that corresponded to the aforementioned five key challenge areas.
- *Open Spaces* which allowed participants to propose and subsequently discuss additional related topics in a group setting.
- *Interactive and Inclusive Sessions* following some of the Liberating Structures<sup>44</sup>. These included activities such as “Impromptu Networking” to mine important challenges, a “User Experience Fishbowl” to share successes and concerns in both industry and research, a “Chaos Cocktail Party” to prioritize intermediate results, and a game to strengthen team bonding.

The participants in this workshop were carefully selected to obtain a balanced representation of the three main stakeholder groups in software engineering: (1) researchers whose work focuses on Technical Debt, (2) practitioners regularly dealing

with Technical Debt, and (3) tool vendors who develop tools to address different aspects of Technical Debt Management. A list of participants is provided at the end of the full report.

During the workshop, the lightning talks allowed the participants to get to know each other’s backgrounds, work, and goals. The break-out sessions made up the majority of the workshop agenda and were the most intense sessions, producing hundreds of post-it notes, which were grouped into themes and subsequently transcribed to inform the resulting manifesto document. Each day, the break-out groups were slightly re-shuffled, with one or two new members joining each group to offer a fresh point of view. The outcomes of each break-out group were peer-reviewed by another group to offer feedback from a different perspective.

Various Open Spaces were proposed by workshop participants that covered a number of topics that complemented the five workshop topics: one open space dealt with “green” Technical Debt focusing on sustainability, another with different facets of Technical Debt (e.g., people, processes, business), and a third discussed how to create purpose-built Technical Debt benchmarks in a way they can be commonly accepted. The interactive sessions using the Liberating Structures were effective in energizing the participants after lunch and allowing a more engaging format to lead the discussions. As the workshop progressed, it became clear that the participants understood each other’s perspectives better and became more efficient in their brainstorming and in-depth discussions.

#### ■ Identifying Technical Debt Challenges

Following the lightning talks, the workshop agenda included an *impromptu networking* session. This activity aimed to provide participants an opportunity to interact with each other through an engaging activity and to elicit the challenges that the group collectively perceived as important.

Participants were asked to answer the following question on small cards: *What is the top priority challenge the Technical Debt community should make progress on to improve its management by software development teams?* The Technical Debt community here refers to researchers, tool vendors, and practitioners who are actively working on Technical Debt and its management. The participants then passed around their cards to read as many answers as possible. There were three impromptu pauses during this reading session, where the participants paired up and rated the challenge statements. This activity at the end of three rounds resulted in all cards having a rating by three different pairs. The top ten challenges the participants elicited are the following (in decreasing rating):

- Integrating tooling into existing development ecosystems in a non- or less-disruptive fashion
- Providing practical guidance to software development teams about when to accept Technical Debt
- Expressing the importance of Technical Debt to users, leaders, funding groups
- Helping C-level executives understand the strategic value of proper Technical Debt Management
- Developing techniques to prioritize technical debt so that the issues with the highest ROI can be addressed first to support efficiency, motivate developers/management to actually address Technical Debt, and avoid waste.
- Convincing business stakeholders to invest in Technical Debt repayment (not all Technical Debt has to be repaid) or not, including explaining the benefits in terms of money, but also risk management

<sup>44</sup> <https://www.liberatingstructures.com/>

- Extracting the context that underpins a team’s Technical Debt, such as process, Technical Debt culture, and domain
- Easily determining the impact of Technical Debt reduction measures (e.g., regarding productivity gains, cost reduction, compliance mistakes) to support decision-making for product owners
- Developing tools that capture a commonly agreed Technical Debt viewpoint aligned with the company culture
- Transitioning from low-level code issues to architecture-level technical debt (architecture is defined as “expensive to change”)

#### ■ Industry/Academia Perspectives Session

A user experience “fishbowl” session on the first day highlighted several key insights about measuring and managing Technical Debt, capturing the different perspectives in industry and academia. While metrics were deemed often incomplete, subjective, or difficult to act upon, participants noted that the context around Technical Debt, for example, expressed in source code comments (also known as *self-admitted debt* [5]) could be utilized more in the future. Measuring Technical Debt and creating feedback loops can provide valuable data at varying levels, from developers to C-level executives. However, blindly optimizing specific metrics can lead to unintended consequences, such as wasted effort or undue pressure on teams. Empirical research today tends to focus heavily on source code analysis, missing the broader development and system context, which makes capturing a holistic view of Technical Debt challenging.

Looking forward, the fish bowl session underscored the importance of fostering collaboration between industry and academia as well as including interdisciplinary perspectives. Industry participants highlighted the relevance of team roles (e.g., having a dedicated role in overseeing Technical Debt Management), team rules (e.g., having design guidelines), and management dynamics (e.g., educating non-technical management on Technical Debt), while academic perspectives leaned toward analyzing easily accessible artifacts, such as source code. Both groups agreed on the need to address the human aspects of Technical Debt and the variability of metrics, which some view as a problem and others as an opportunity.

To improve Technical Debt Management, industry participants recommended appointing a “Technical Debt champion” within software development teams/organizations to educate and advocate for best practices, especially towards management. They also suggested adapting research methods from other disciplines, such as social sciences, creating more purpose-built Technical Debt Management tools that seamlessly integrate into developer and decision-maker workflows, and demonstrating how Technical Debt impacts innovation to secure buy-in from upper management. These approaches aim to create more actionable strategies for addressing Technical Debt while maintaining alignment with organizational goals. The fishbowl session was thus instrumental in eliciting some of the participant’s core values, beliefs, and principles regarding Technical Debt, which is at the heart of the manifesto.

#### ■ Open Space on the Multi-Dimensionality of Technical Debt

Technical Debt emerges from decisions or other phenomena (e.g., deprecation of a library) throughout the software development cycle, jeopardizing software maintenance and evolution and affecting most development workflows. However, managing Technical Debt takes significant effort and time. One of the main

reasons is that different context variables and factors (peopleware, processes, products, business, etc.) influence how Technical Debt is managed or even perceived. It certainly goes far beyond source code, and there is no single metric that can accurately capture the concept. This Open Space discussed ways to measure Technical Debt, taking those multiple perspectives into account, as well as evidence-based metrics required to support decision-making.

The participants agreed that collecting both objective and subjective data is necessary, as well as a way to aggregate these two types. This aggregation should take stakeholders’ concerns into account (what stakeholders want to know to make decisions). Quantitative measures can be contextualized by qualitative ones. The combination of measures must represent the context and match the organization; context is key.

Representation of Technical Debt also requires multiple metrics; a single number is hard to understand and interpret and may hide or obscure the complexities underlying Technical Debt. In addition, graphical representations are useful but must be followed by an explanation (recommendation) to provide value to stakeholders (CEO, Project Owner, developers, and so on).

Beyond metrics, Technical Debt items must include other contextual information. How exactly does it hinder development? Why was it incurred in the first place? This kind of context should drive data collection. Sometimes, specific Technical Debt items need different metrics and contextual information as they affect development differently.

There is still much tacit knowledge involved in managing technical debt, especially in identifying, measuring, prioritizing, and repaying it. Therefore, experienced software developers have a relevant role in mentoring their more junior colleagues in understanding the software structure and behavior, as well as past design decisions. Traceability between software artifacts (e.g., between source code and documentation) can facilitate making this tacit knowledge explicit. However, in general, software professionals must be prepared to spend time understanding the context behind technical debt items.

Based on the aforementioned issues, the Open Space participants identified several research directions:

- How to identify meaningful objective metrics that should be automatically collected?
- How to measure and manage developer “turnover” (e.g., the departure of developers with knowledge of legacy code)?
- How to capture current developers’ perspectives about Technical Debt?
- How to record the design decisions that incur Technical Debt?
- How to aggregate different unit metrics?
- What is the interplay between Indicators of Technical Debt and Metrics of Technical Debt?
- How to incorporate a human-in-the-loop within the Technical Debt management process?
- How to adequately capture the balance between Technical Debt repayment and the development of new features, especially regarding quality, risks, and costs?

#### ■ Manifesto

The core outcome of this workshop is a report published in the Dagstuhl Manifestos series<sup>45</sup>, titled *Reframing Technical Debt*. This manifesto report summarizes the current landscape of Technical Debt methods and tools and explains where current industrial practice is struggling, where current research has made progress, and where it falls short. This motivates the actual manifesto, which condenses the outcomes of the workshop into a series of **values, beliefs, and principles**.

<sup>45</sup> <https://drops.dagstuhl.de/entities/journal/DagMan>

The manifesto reports subsequently explains the stated values, beliefs, and principles in detail. Finally, it outlines a roadmap with concrete milestones and a call to action addressed at researchers, software practitioners, tool vendors, Technical Debt advocates, policymakers and other government agencies, leading technology companies, funding agencies, industry leaders, and industry associations. The manifesto is signed by the workshop participants.

### ■ Concept Map

During the lightning talks and plenary discussions, the organizers captured key concepts discussed and how they related to each other as a concept map. This concept map (Figure 6.12), though not exhaustive, offers a perspective on the multidisciplinary nature of Technical Debt, its management, and its essential role in software engineering.

The rest of this report includes abstracts from workshop attendees summarizing their relevant work, which informed the discussions during the workshop.

## ■ References

- 1 Nanette Brown, Yuanfang Cai, Yuepu Guo, Rick Kazman, Miryung Kim, Philippe Kruchten, Erin Lim, Alan MacCormack, Robert Nord, Ipek Ozkaya, Raghvinder Sangwan, Carolyn Seaman, Kevin Sullivan, and Nico Zazworka. Managing technical debt in software-reliant systems. In *Proceedings of the FSE/SDP Workshop on Future of Software Engineering Research, FoSER '10*, page 47–52, New York, NY, USA, 2010. Association for Computing Machinery.
- 2 Ciera Jaspan and Collin Green. Defining, measuring, and managing technical debt. *IEEE Softw.*, 40(3):15–19, May 2023.
- 3 Helvio Jeronimo Junior and Guilherme Horta Travassos. Consolidating a common perspective on technical debt and its management through a tertiary study. *Information and Software Technology*, 149:106964, 2022.
- 4 Jason Lefever, Yuanfang Cai, Humberto Cervantes, Rick Kazman, and Hongzhou Fang. On the lack of consensus among technical debt detection tools. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, pages 121–130, 2021.
- 5 Aniket Potdar and Emad Shihab. An exploratory study on self-admitted technical debt. In *2014 IEEE International Conference on Software Maintenance and Evolution*, pages 91–100. IEEE, 2014.
- 6 Wolfgang Trumler and Frances Paulisch. How “specification by example” and test-driven development help to avoid technical debt. In *2016 IEEE 8th International Workshop on Managing Technical Debt (MTD)*, pages 1–8, 2016.

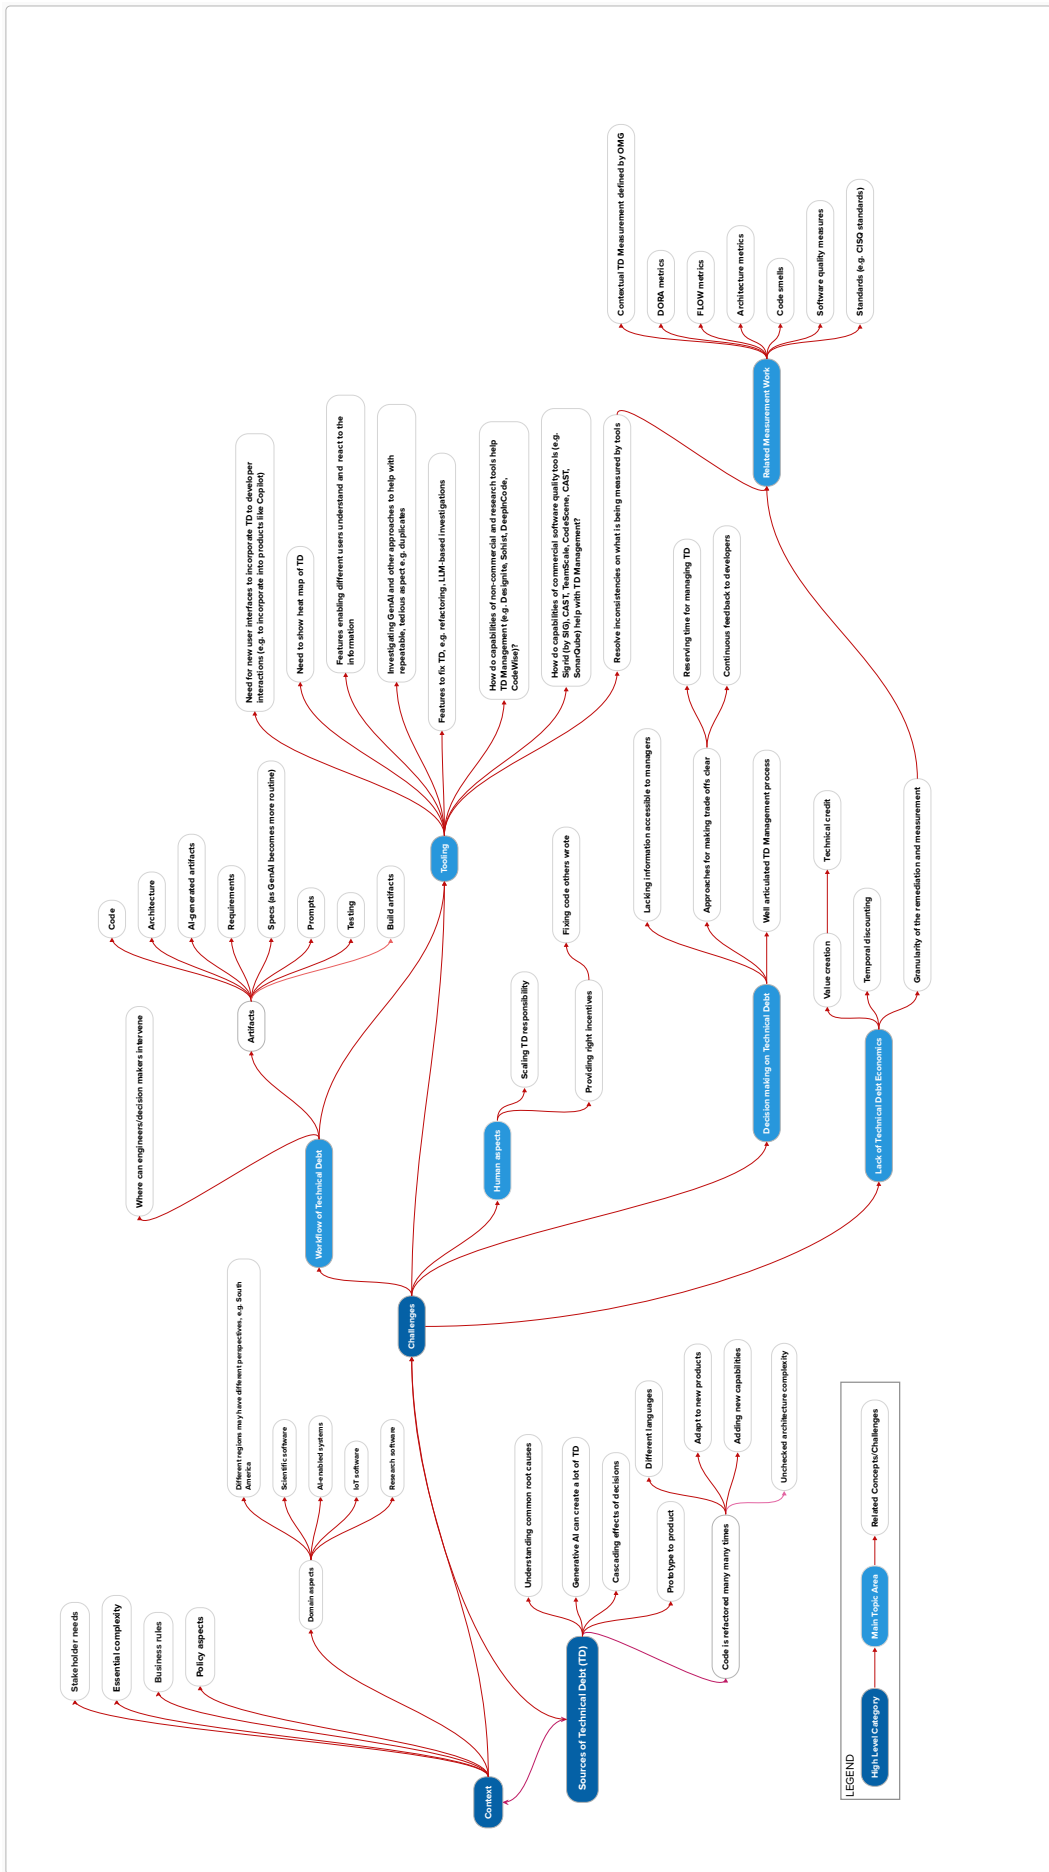


Fig. 6.12  
Technical Debt Concept Map.

## 6.70 Rethinking the Role of Bayesianism in the Age of Modern AI

**Editors:** Vincent Fortuin, Mohammad Emtiyaz Khan, Mark van der Wilk, and Zoubin Ghahramani

**Seminar No. 24461**

Date: November 10–15, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.11.40

© Creative Commons BY 4.0 license

© Vincent Fortuin



**Participants:** Laurence Aitchison, Alexander A. Alemi, Pierre Alquier, Julyan Arbel, Thang Bui, Kamélia Daudel, Gintare Karolina Dziugaite, Carl Henrik Ek, Maurizio Filippone, Katharine Fisher, Vincent Fortuin, Pablo García Arce, Erin Grant, Philipp Hennig, Alexander Immer, Desi Ivanova, Theofanis Karaletsos, Mohammad Emtiyaz Khan, Jeremias Knoblauch, Yingzhen Li, Thomas Möllenhoff, Kevin Murphy, Eric Nalisnick, Roi Naveiro Flores, Theodore Papamarkou, Guiomar Pescador Barrios, Tom Rainforth, Daniel Roy, Tim Rudner, Maja Rudolph, David Rügamer, Jan-Willem van de Meent, Tycho van der Ouderaa, Mark van der Wilk, Mariia Vladimirova, Florian Wenzel, Sinead Williamson, Andrew G. Wilson

The Dagstuhl Seminar “Rethinking the Role of Bayesianism in the Age of Modern AI” (24461) was convened to explore the contemporary role of Bayesian methods in artificial intelligence, particularly in light of the remarkable advancements in large-scale deep learning. While Bayesian Deep Learning (BDL) holds the promise of addressing key limitations of traditional deep learning, such as uncertainty estimation, encoding prior knowledge, and preventing catastrophic failures, it frequently falls short of its potential in practical applications. This discrepancy arises from several fundamental challenges. These challenges include the difficulty of computing accurate posterior approximations, the scarcity of flexible prior distributions, and the lack of suitable benchmarks for evaluating Bayesian models. Furthermore, misconceptions regarding the scope of Bayesian methods often lead researchers to harbor unrealistic expectations and overlook simpler, non-Bayesian alternatives like bootstrap methods, post-hoc uncertainty scaling, and conformal prediction. Such over-expectations, followed by under-delivery, may cause researchers to lose faith in Bayesian approaches. The central question addressed by the seminar was: In this era of AI where scaling seems to solve many problems, what is the unique role of Bayesian methods? The goal was to redefine the promises and challenges of Bayesian approaches, identify areas where they can outperform non-Bayesian methods, and highlight key application domains where their strengths can be best leveraged. By bringing together researchers from diverse backgrounds, the seminar aimed to chart a path for future research to innovate, enhance, and strengthen the real-world impact of BDL. The seminar recognized that while non-Bayesian methods seem to be solving problems that Bayesians once hoped to solve with Bayesian methods, it was important to re-examine the value and potential of the Bayesian approach.

### ■ Structure of the Seminar

The seminar was designed to foster an interactive and collaborative environment, incorporating three distinct types of events: workgroup sessions, guided discussions, and final plenary discussions.

**Workgroup Sessions.** These sessions revolved around overarching questions pertaining to the main advantages of Bayesian methods, the challenges hindering their adoption, and the practical areas where they can make the most difference. The workgroups always featured between one and three short input talks from different participants, which then informed the subsequent discussions. The workgroups were structured around three main questions:

- What are the main benefits of Bayes that are hard to achieve otherwise?
- What are the most pressing challenges in its adoption?
- What are the most impactful ways for Bayes to make a difference in practice?

**Guided Discussions.** The guided discussions were designed to examine contentious issues and encourage debate, focusing on three key motions:

- “*Bayes’ theorem is broken for making predictions with large models*”
- “*We can build subjective Bayesian priors for NNs that we actually believe in*”
- “*Bayes is useless if we cannot scale to LLMs*”

**Final plenary discussions.** The final discussions focused on the bigger picture and the next steps for researchers in the field. They centered around three main themes:

- What can we do to encourage researchers to join the BDL community and how can we support and uplift each other within the community?
- How can we measure progress in the field and find promising application areas that would convince practitioners to use Bayesian methods?
- What are some grand long-term challenges for which we could hope Bayesian methods to make a difference and potentially outperform standard deep learning?

## ■ Insights from the Working Groups

**Talks.** The seminar featured a series of presentations covering a wide range of topics related to Bayesian methods. The participants contributed these talks based on a pre-seminar poll regarding the group's interests, which informed the working group's discussions. They included discussions on:

- The distinction between *aleatoric and epistemic uncertainty*. This included a detailed look at how these terms are often used inconsistently, leading to issues in the literature. The discussion also covered how to estimate these uncertainties in practice and how to best decompose total uncertainty.
- The difference between predictive and parameter uncertainty. The discussion here considered how to search the space of predictions and how to judge explanations without relying on predictions.
- Developing benchmarks for Bayesian methods. This included a discussion on whether current uncertainty measures are useful for model comparison and selection, and whether new benchmarks are needed.
- The roles of prediction and explanation in science. The discussion focused on how machine learning has changed the landscape of prediction and explanation, and the role of Bayesian approaches in these areas.
- Bayesian foundation models. This discussion considered how probabilistic thinking can help us understand foundation models and whether deep learning technologies can help advance probabilistic methods.
- Bayesian Neural Network (BNN) architectures. This included a look at model selection using the marginal likelihood, and whether uncertainty helps to avoid overfitting.
- Pseudo-posteriors. This session explored methods like likelihood tempering and robust loss functions to address model misspecification.
- Bayesian methods for sequential learning. This included discussions of new algorithms for deep learning and how to apply them in dynamic settings.
- The geometry of BNN posteriors. The discussion focused on the challenges for Bayesian inference in deep learning, such as the intractability of posterior distributions and the existence of multiple minima.
- Partial stochasticity in BNNs. This talk explored scalable variational approximations based on subnetworks and whether a fully Bayesian treatment of NNs is necessary.
- Teaching Bayesian ML. This session covered the decisions academics make when teaching Bayesian ML, what to include and what to omit, and the value of diversity in teaching approaches.
- The relationship between Bayesian theory and practice. This presentation explored non-Bayesian justifications for Bayesian updating, the challenges in modeling complex data, and the value of trying out models to see which ones work best.

Together these workgroup sessions yielded important insights into the potential and challenges associated with Bayesian methods along the three main themes of the seminar:

**Benefits of Bayes.** Participants highlighted several core benefits, including the ability to quantify uncertainty, update models, perform model selection, and obtain improved point estimates. The quantification of uncertainty was noted as a key advantage, although it was admitted that it can sometimes be achieved by other means. In contrast, model updating was seen as a critical unique benefit, allowing models to adapt to new data without complete retraining.

**Challenges in Adoption.** Significant challenges were identified, particularly in the areas of scalability, and prior and model misspecification. These challenges pose barriers to the wider adoption of Bayesian methods. Scalability was a significant concern, as many Bayesian methods are computationally expensive. Prior misspecification was also identified as a major issue, as it can bias the results negatively and hamstring many of the benefits of the Bayesian approach. Finally, model misspecification also presents problems, as many models do not perfectly fit real-world data.

**Impactful Applications.** Sequential learning was emphasized as an area where Bayesian methods have the potential to make a substantial impact. The ability of Bayesian methods to update beliefs over time and adapt to new data makes them well-suited to sequential learning tasks.

## ■ Insights from the Guided Discussions

The guided discussions brought to light differing opinions and perspectives on critical issues within the Bayesian community.

**Bayes' Theorem and Large Models.** A central debate revolved around the applicability of Bayes' theorem to large models. The "pro" side contended that while mathematically sound, the epistemological assumptions of Bayes' theorem do not translate well to complex neural networks (NNs). They argued that NNs lack clearly defined priors. Furthermore, they noted that simpler, more direct methods like point estimates or conformal predictions are often more cost-effective and practical. The "con" side, however, argued that any limitations are due to implementation issues and not the theorem itself. They also noted the value of Bayesian methods when fine-tuning models with small datasets, emphasizing that Bayes provides a flexible framework. The core of the debate was whether the practical constraints of large models should limit the application of Bayesian methods, or whether the flexibility of Bayesian approaches could be adapted to these large models. This discussion highlighted the need for a nuanced understanding of the strengths and limitations of Bayesian methods in different contexts.

**Subjective Bayesian Priors.** The discussion on subjective priors for NNs explored the significance of priors, particularly for out-of-distribution data, and the difficulties in defining them effectively. Some participants emphasized that priors should be based on domain expertise, while others questioned the mathematical basis for using subjective priors on neural networks. The discussion highlighted the challenge of balancing subjective knowledge with the need for mathematical rigor. It was also noted that priors on function spaces might be easier to specify than priors on model parameters, and that designing priors to bias solutions toward the data was an area worth exploring.

**Scaling to LLMs.** A significant point of contention was whether Bayesian methods are still relevant if they cannot scale to LLMs. The “pro” side argued that the need for scalability to LLMs is paramount for Bayes to stay relevant in the field. The “con” side countered that Bayes should not be limited to large models; it also plays a crucial role in small-data problems and scientific experiments. It was suggested that LLMs themselves could be used as priors and diffusion models as inference algorithms, highlighting the possibility of using modern AI tools within a Bayesian framework. This debate emphasized the need to re-evaluate the role of Bayesian methods in the context of rapidly advancing AI technologies, and whether the Bayesian approach can be adapted to new tools.

## ■ Insights from the Final Discussions

The concluding discussions synthesized the key findings from the seminar and outlined future directions for the community.

**Community Building.** There was a strong consensus on the need to foster inclusivity within the Bayesian community, encompassing all levels of seniority, as well as industry and academia. It was stressed that a positive outlook on the Bayesian toolkit in reviews was also crucial. The community should view Bayesian methods as a set of useful tools, rather than as a rigid ideology. The importance of mentorship and support for junior researchers was also noted, as well as the value of bringing in people who may be implicitly Bayesian without realizing it.

**Benchmarks and Applications.** Participants emphasized the importance of moving beyond traditional vision-based benchmarks to include decision-making and sequential learning tasks. The community needs to focus on identifying applications that highlight the unique advantages of Bayesian methods and create tools that can be used in impactful applications. The use of scoring rules for decision-making was also suggested, as it allows for a clear understanding of the value of improvements and highlights utility in downstream decisions as the key metric for the success of predictive systems. The discussion also highlighted the need to consider applications relevant to the current state of AI and other sciences, rather than relying on past applications.

**Grand Challenges.** Discussions on grand challenges included developing a Bayesian equivalent of AlphaFold, addressing the ARC challenge, and incorporating LLMs as priors. Data efficiency was highlighted as a key strength of Bayesian methods,

with the potential to significantly reduce the amount of data required for training. The group also raised important questions about the nature of reasoning and compositionality in models, as well as the challenge of building robust and trustworthy AI systems. The need for causal inference was also noted as critical for many real-world applications. The discussion also covered the possibility of using LLMs to learn structured models and programs.

## ■ Next Steps

The seminar concluded with the identification of several concrete steps to advance the field of Bayesian deep learning.

**Benchmarks.** The community should develop benchmarks that are challenging for deep learning but can be addressed using Bayesian methods, with a focus on sequential and active learning. Data efficiency should also be a focus when creating benchmarks. Furthermore, existing benchmarks should be evaluated for adaptation, especially those that move beyond vision-based tasks.

**Research.** Future research should move beyond traditional likelihood metrics and instead prioritize posterior predictive checks to ensure that models are making good predictions. Researchers should also seek to communicate the importance of decision outcomes and ensure that the metrics align with practical goals. There was also a call to focus on the principles behind Bayesian methods rather than just scaling, and to allow for alternative Bayesian inference frameworks (such as the martingale posterior).

**Organization.** There is a need to establish a benchmark track at the yearly Symposium on Advances in Approximate Bayesian Inference (AABI), as well as to continue fostering connections within the Bayesian deep learning community through communicative tools (e.g., slack, Notion). The community should also explore the possibility of creating a Bayesian summer school and a virtual seminar, and to seek integration with with the International Society for Bayesian Analysis (ISBA), possible through the foundation of a Bayesian deep learning chapter. There is also a desire to create a yearly Bayesian AI event to foster community. Furthermore, there was a call to share teaching resources to help standardize and improve international higher education in Bayesian machine learning.

## 6.71 Research Infrastructures and Tools for Collaborative Networked Systems Research

6

**Editors: Georg Carle, Serge Fdida, Kate Keahey, and Henning Schulzrinne**  
**Seminar No. 24462**

Date: November 10–13, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.11.60

© Creative Commons BY 4.0 license

© Georg Carle, Serge Fdida, Kate Keahey, and Henning Schulzrinne

**Participants:** Tom Barbette, Terry Benzel, Georg Carle, Hakima Chaouchi, Walid Dabbous, Yuri Demchenko, Serge Fdida, Sebastian Gallenmüller, Jorge Gasos, Michael Goedicke, Cheikh Ahmadou Bamba Gueye, Tobias Hoßfeld, Kate Keahey, Wolfgang Kellerer, Raymond Knopp, Jim Kurose, Deep Medhi, Jelena Mirkovic, Andrew W. Moore, Paul Michael Ruth, Damien Saucez, Björn Scheuermann, Henning Schulzrinne, Jörg Widmer, Walter Willinger, Adam Wolisz, Ellen Zegura, Martina Zitterbart



Research infrastructures should evolve towards advanced scientific instruments that offer a vital insight to the underlying information in improvising the understanding of science methodologies and practices as they reliably and precisely help the scientist to measure the subject of their investigations. The Dagstuhl Seminar participants strongly agreed that large-scale research infrastructures are essential for providing scientists with access to specialized, advanced resources enabling cutting-edge experiments. As a result, the following key conclusions were drawn:

### A. Strategic Investment & Community Engagement:

Research infrastructures represent a vital and long-term investment that demands active participation from research communities, sustained human capital development, and financial sustainability.

### B. Open Access & Data Sharing:

While open access to shared physical infrastructure is essential, access to open research data is equally critical. Digital sharing of scientific results accelerates innovation, enhances reproducibility, and strengthens FAIR (Findable, Accessible, Interoperable, and Reusable) data sharing through metasearch services.

### C. Amplified Impact & Network Effects:

Research infrastructures inherently complement and amplify each other, creating a synergistic network effect. This interconnectedness fosters a more rigorous scientific approach and methodology, driving greater collaboration and knowledge advancement.

The results of the seminar include the following key recommendations:

## Key Recommendations

- 1. Define clear scientific objectives:** Research infrastructures must explicitly articulate their scientific goals and establish a well-defined set of research questions to address.
- 2. Foster a strong scientific community:** The success of research infrastructures depends on strong community engagement. Support measures are essential to strengthen and sustain the scientific community. The effort of their support teams should be better recognized.
- 3. Implement EasyFAIR principles:** Adopting an EasyFAIR framework – offering comprehensive and automated support for researchers – is crucial to ensuring and leveraging the FAIR principles (Findability, Accessibility, Interoperability, and Reusability) of digital assets. Additionally, open research data and reproducibility should be mandated by funding agencies and scientific societies. Scientists making an effort to share their research data should be rewarded.
- 4. Enhance reproducibility:** Reproducibility is a critical priority. Concrete methodologies must be established to ensure comparability of experimental results across different research infrastructures.
- 5. Multi-year investment strategy:** Research infrastructures should be properly articulated, designed and supported according to a longer-term roadmap and a sustained investment strategy.
- 6. Establish common abstractions:** Standardized models should be widely adopted for describing experiments and associated frameworks, including information models, data models, and ontologies.
- 7. Improve findability and accessibility:** The discovery and accessibility of research infrastructures and testbeds should be enhanced through comprehensive catalogs detailing available hardware and functionalities. It is also essential to assess

how new and planned infrastructures contribute to scientific diversity.

8. **Define standardized evaluation criteria:** A clear and well-defined set of evaluation criteria is necessary to assess the relevance and impact of research infrastructures. The outcomes of the “Testbed Evaluation” World Cafe (cf. Section 7 of the full report) provide a valuable basis for these criteria. Standardized assessment frameworks should be established for different categories of testbeds.
9. **Optimize user experience:** Usability for researchers must be a priority. An innovative metric, *Time to First Experiment (TTFE)*, can be used to measure the efficiency of infrastructures in enabling rapid experimentation and adoption. Education and training should be a strong component of research infrastructures.
10. **Ensure interoperability and openness:** Strong support for interoperability between testbeds is crucial as well as using open components as often as possible, as it supports the ability to easily port experiments across different infrastructures.
11. **Promote flexibility and adaptability:** Provisioning of ready-to-use experimental platforms: Instead of providing merely the fundamental resources of an experiment, the testbed should provide a experimental templates that researchers can use to perform their own research. This experimental template or “*blueprint*” can be used by researchers to answer specific research questions. These blueprints are not static but *malleable*, i.e., researchers are encouraged to adapt and extend them to fit their needs. The concept of *malleability* includes facilitating the modification of software artifacts and support of composability.
12. **Support sustainable development goals (SDGs):** Large-scale research infrastructures contribute directly to the SDGs by optimizing the efficiency of hardware resource usage and improving workflows from experiment design to result dissemination. Additionally, insights gained from research findings can enhance global technical infrastructure, further supporting SDG objectives.

## 6.72 Graph Algorithms: Distributed Meets Dynamic

**Editors: Keren Censor-Hillel, Yasamin Nazari, Eva Rotenberg, and Thatchaphol Saranurak**  
**Seminar No. 24471**

Date: November 17–22, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.11.92

© Creative Commons BY 4.0 license

© Thatchaphol Saranurak, Keren Censor-Hillel, Yasamin Nazari, and Eva Rotenberg

**Participants:** Sepehr Assadi, Alkida Balliu, Aaron Bernstein, Sayan Bhattacharya, Joakim Blikstad, Sebastian Brandt, Karl Bringmann, Keren Censor-Hillel, Yi-Jun Chang, Shiri Chechik, Keerti Choudhary, Martín Costa, Tijn de Vos, Michal Dory, Aditi Dudeja, Faith Ellen, George Giakkoupis, Seth Gilbert, Christoph Grunau, Magnús M. Halldórsson, Adam Karczmarz, Fabian Daniel Kuhn, Jakub Lacki, Quanquan C. Liu, Yannic Maus, Danupon Nanongkai, Yasamin Nazari, Alexandre Nolin, Krzysztof Nowicki, Dennis Olivetti, Ami Paz, Richard Peng, Vijaya Ramachandran, Eva Rotenberg, Thatchaphol Saranurak, Shay Solomon, Jukka Suomela, Jan van den Brand, Virginia Vassilevska Williams, David Wajc, Tianyi Zhang, Anna Zych-Pawlewicz



This seminar brought together researchers from two research communities on distributed and dynamic graph algorithms. In dynamic settings, algorithms receive updates to the graph, and should update the solution efficiently. In many distributed settings, computation takes place in a decentralized manner over a network or multiple machines. The input is distributed among multiple computing devices that need to communicate to find a solution to a given problem. The first connection between these models that was explored in the seminar is the technical toolbox. In recent years, there has been a growing number of results in both areas that use similar algorithmic tools and technical ideas. Examples of these common technical tools include expander-based techniques, truncating Vizing chains for improved edge colouring, algebraic techniques based on matrix multiplication, spanners and hopsets, clustering coresets, etc.

While some of these techniques have found applications in one or both of these models, there are often technical challenges in applying them in different settings and knowledge of models of specific tools is needed. This is why the seminar aimed to increase collaborations and discussions between the two communities. The second type of connection is from a modeling point of view,

motivated by practice. In many practical big data settings the input is both distributed and dynamic at the same time. First, it is stored in a decentralized way over many machines, and second, the input changes over time. Few talks focused on the models that combine multiple computational settings.

The format of the seminar was as follows: the first day included an introduction session in addition to an ice-breaker “research speed-dating” event. In this ice-breaker event, participants were paired together in several short sessions with an attempt to pair participants who are not already familiar with each other’s work. We had 4 plenary overview/survey talks, and the rest of the presentations were 20 minute talks and 10 minutes of questions and discussions. We also had two engaging open problem sessions in addition to various allocated times for collaboration.

We aimed to invite a diverse group of participants, both based on demographic factors and geographic affiliation, as well as seniority level. We had participants from Europe, North America, and Asia, and with experiences ranging from PhD students and postdocs to academic and industry research with varied level of seniority.

## 6.73 Regular Expressions: Matching and Indexing

**Editors:** Inge Li Gørtz, Sebastian Maneth, Gonzalo Navarro, and Nicola Prezza  
**Seminar No.** 24472

Date: November 17–22, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.11.108

© Creative Commons BY 4.0 license

© Inge Li Gørtz, Sebastian Maneth, Gonzalo Navarro, and Nicola Prezza



**Participants:** Antoine Amarilli, Hideo Bannai, Ruben Becker, Giulia Bernardini, Philip Bille, Manuel Cáceres, Davide Cenzato, James Davis, Dominik D. Freydenberger, Pawel Gawrychowski, Adrián Gómez Brandón, Inge Li Gørtz, Roberto Grossi, Moshe Lewenstein, Konstantinos Mamouras, Sebastian Maneth, Wim Martens, Yasuhiko Minamide, Gonzalo Navarro, Nicola Prezza, Cristian Riveros, Markus L. Schmid, Teresa Steiner, Michelle Sweering, Simon Rumle Tarnow, Brink van der Merwe

The Dagstuhl Seminar “Regular Expressions: Matching and Indexing” (24472) took place from November 17th to 22nd, 2024. The goal of this seminar was to bring together researchers from various research directions dealing with algorithmic aspects of regular expressions and finite automata. Regular expressions and finite automata lie at the foundations of Computer Science and have been used since the sixties in basic problems like compiler design. The key algorithmic challenge is regular expression matching, that is, efficiently identifying words of a regular language within a sequence or within the paths of a labeled graph. Over the years, there have been numerous algorithmic advances around the topic, while at the same time their applications have spread over too many different areas like information retrieval, databases, bioinformatics, security, and others, which not only make use of standard results but also pose new and challenging variants of the regular expression matching problem. The use of regular expressions has made its way even into current standards like SQL:2016 and SPARQL.

The seminar was meant to bring together expertise from communities involved with regular expressions that, despite working on closely-related topics, do not frequently meet: graph databases, compressed data structures, and streaming algorithms. The following specific points were addressed.

**String Indexing for Complex Patterns.** The classic and well-studied string indexing problem is to preprocess a string such that subsequent pattern-matching queries (locate all occurrences of the query string within the indexed strings) can be supported efficiently. The goal is to achieve a fast query time in terms of the length of the query string while keeping the memory needed for the index small. Several string indexes that support “complex” pattern matching queries have recently appeared. Examples include indexes that support pattern matching with wildcards, gaps, elastic degenerate, and approximate string matching. Such complex patterns may

be viewed as special restricted cases of regular expressions. Thus, studying these in the context of regular expressions may lead to a more unified understanding of these and new research directions.

**Advanced Regular Expression Operators.** Nearly all existing methods for regular expression matching consider regular expressions constructed using the concatenation, union, and Kleene star operators. In this setting, the complexity of the problem is well understood, and upper bounds with near-matching conditional lower bounds are known. However, more “advanced” operators that allow increased expressiveness are widely used in practice. These include backreferences (matching a previously matched subexpression), intervals (matching a subexpression a given number of times), character classes (matching a character to a set of characters), variable length gaps (matching any string of a length within a specified interval). In contrast to the classic setting, regular expression matching with these operators is much less well-understood. While some of these are known to be much harder from a complexity-theoretic viewpoint (i.e., regular expression matching with backreferences is, in general, NP-hard), recent developments have led to new methods handling many practically essential variants of the problem. Examples include deterministic regular expressions, pattern matching with variables, pattern matching with variable length gaps, and elastic degenerate strings. This enables a new study area that is theoretically challenging and practically relevant.

**Indexing Automata and Regular Expressions.** The counterpart of the “string indexing for complex patterns” problem is to index/pre-process a regular expression to speed up tasks such as deciding membership of a string in its accepted language. An equivalent formulation of the problem, which is lately receiving much attention in the research community,

is to index finite-state automata for *path queries*: given a string at query time, decide whether the string occurs in some walk on the automaton (equivalently, it is a substring of some string in the automaton's accepted language). While recent research has shown quadratic (size of the regular expression/automaton multiplied by the length of the query string) conditional lower bounds for the problem, particular classes of regular expressions and automata do admit efficient solutions: deterministic regular expressions, elastic degenerate strings, and Wheeler automata are examples of this kind. Even more interestingly, other lines of research showed that the problem can be parameterized, thus classifying *all* regular expressions and automata according to their *propensity* to support membership/path queries: examples of such parameters include the number of strings in the regular expression (equivalently, the number of union symbols and Kleene stars), the co-lexicographic width of an automaton, and the density (amount of nondeterminism) of a regular expression. These directions opened an exciting new research area of both theoretical and practical interest with applications including bioinformatics (with the problem of indexing large pan-genomic graphs for path queries) and graph databases.

**Graph Databases and Regular Path Queries.** Another flavor of regular expression matching arises in the context of graph databases, where the data consists of a labeled directed graph and queries involve different forms of subgraph pattern matching. The SPARQL standard, as well as several other alternative graph query languages, support in particular “regular path queries (RPQs)”, which are essentially regular expressions that are to be matched against the sequence of labels of graph paths. The basic strategy is to traverse the product graph between the database graph and the automaton of the regular expression. This leads to a quadratic time complexity, but other strategies based on graph traversals, pre-indexing the graph, or multiplication of sparse matrices, perform better in many practical cases. The most basic query aims at returning the endpoints of the paths found, but others ask for all the paths, the shortest paths, and so on. There has also been work on lower bounds on this problem, particularly relating it to the well-known AGM bounds that hold for matching a given subgraph shape. Strategies developed for regular expression matching on strings can be leveraged to solve RPQs, but this requires a cross-fertilization between the

stringology and the database community in order to transfer the knowledge: most solutions implemented in actual graph database systems are just heuristics that have been shown to be clearly inferior to solutions grounded on more solid algorithmic foundations.

**Regular Expression Matching on Streams.** Another relevant research area, motivated by big data applications, is that of regular expression matching on data streams: pre-process a regular expression  $R$  so that, later, we can identify all suffixes  $S[1, i]$  of an incoming stream  $S$  such that  $S[j, i]$  matches  $R$  for some  $j \leq i$ , using space sub-linear in  $|S|$  and  $|R|$ . This problem generalizes a seminal work by Porat and Porat (2009) on exact pattern matching and on the  $k$ -mismatches problem on streams. A subsequent prolific line of works improved the original bounds and extended these results to edit distance, to the dictionary problem, and to wildcard matching. Recent research has shown that the problem admits solutions for arbitrary regular expressions using space polynomial in the logarithm of the stream's length and in the number of union symbols and Kleene stars in the regular expression. This shows that parameterized approaches (as the ones for the problem of indexing automata and regular expressions) are indeed useful also in the streaming setting, and therefore strongly motivates an exchange of ideas between the two sub-fields.

The seminar fully satisfied our expectations. The 26 participants from 14 countries (Chile, Denmark, Finland, France, Germany, Great Britain, Israel, Italy, Japan, Netherlands, Poland, South Africa, Spain, and US) gave invited survey talks covering the state of the art in scientific fields related with the topics of the seminar. The talks presented works related with algorithms for regular expression matching (covering classic algorithms, extensions of regular expressions, indexing, streaming, and compressed data structures), regular expressions in graph databases, and practical tools used at the industrial scale.

We are really thankful to Schloss Dagstuhl for providing an extremely inspiring and professional environment. Scientific talks were interleaved with coffee breaks and cheese tastings which fostered engagements of the participants, and the evening “sessions” in the sauna, wine cellar, and music room offered a relaxed environment to continue chatting about research in a less formal environment.

## 6.74 Deep Learning for RNA Regulation and Multidimensional Transcriptomics

**Editors: Annalisa Marsico, Uwe Ohler, Igor Ulitsky, and Kathi Zarnack**  
Seminar No. 24491

Date: December 1–6, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.12.1

© Creative Commons BY 4.0 license

© Kathi Zarnack, Annalisa Marsico, Uwe Ohler, and Igor Ulitsky



**Participants:** Rolf Backofen, Pavel Baranov, Mathieu Blanchette, Charlotte Capitanchik, Christoph Dieterich, Florian Erhard, Eduardo Eyras, Julien Gagneur, Jonathan Göke, Marko Jovanovic, Jan Philipp Junker, Julian König, Frederick Korb, Claudia Kutter, Gioele La Manno, Markus Landthaler, Liana Lareau, Martin Lewinski, Mo Lotfollahi, Yael Mandel-Gutfreund, Miguel Ángel Manzanares Serrano, Annalisa Marsico, Jin-Wu Nam, Evgenia Ntini, Uwe Ohler, Yaron Orenstein, Michal Rabani, Katie Rothamel, Alexander Sasse, Hagen Tilgner, Jernej Ule, Igor Ulitsky, Li Yang, Kathi Zarnack, Jianyang Zeng

The Dagstuhl Seminar 24491 “Deep Learning for RNA Regulation and Multidimensional Transcriptomics” brought together an interdisciplinary group of computer scientists, computational biologists, and experimentalists to discuss current challenges and emerging opportunities at the intersection of RNA biology and artificial intelligence. Across a broad range of sessions, participants showcased recent advances in deep learning models, RNA sequencing methods, and systems biology approaches, revealing novel insights into the epitranscriptome, RNA structure and function, RNA–protein interactions, regulatory mechanisms, and disease biology.

A central theme of the seminar was the exploration of transcriptome complexity. This extends well beyond coding sequences, as non-coding RNAs, such as long non-coding RNAs (lncRNAs), microRNAs and enhancer RNAs, play essential roles in modulating gene expression. Toward the molecular code of lncRNA function, several speakers presented high-throughput approaches to study the spatiotemporal behavior of lncRNAs and dissect their functions in cis and trans. In addition, several presentations emphasized how translating ribosome footprints illuminate new regulatory events in non-coding regions, while new deep learning models decipher the interplay between motif composition and translation outcomes.

Another major focus was the promise of AI-driven analyses of vast repositories of RNA-sequencing data. Participants showcased deep learning strategies to identify regulators of splicing, alternative polyadenylation, RNA stability, and translation. From the single-cell perspective, novel methods—particularly those combining metabolic labeling or spatial transcriptomics with deep learning—offered unprecedented resolution into gene regulatory cascades in development and disease. These high-throughput approaches highlighted how combinatorial readouts of RNA modifications, transcript isoforms, and protein–RNA binding events can be turned into powerful predictive frameworks through deep neural networks.

At the level of RNA–protein interactions, discussions centered on data from experiments like individual-nucleotide resolution UV crosslinking and immunoprecipitation (iCLIP) and multiplexed profiling platforms that revealed how RBPs coordinate transcript processing and decay. Deep learning architectures such as RBPNet and panRBPNet showcased how to predict RBP binding sites directly from primary sequence, integrating background noise models and improved interpretability. Complementarily, methods that quantify the uncertainty of predictions or predict changes in splicing and translational efficiency broaden the scope and robustness of computational pipelines.

Disease contexts, including metabolic liver disorders, neuroblastoma, and other malignancies, highlighted the biomedical importance of precisely quantifying transcriptomic complexity and RNA-regulatory events. By capturing single-cell and subcellular heterogeneity, these advanced molecular methods can inform novel therapeutic strategies. Participants presented case studies on the roles of specific lncRNAs in cellular pathways, microproteins in cardiomyocytes, and the impact of RNA modifications on stability and translation.

Complementary to the scientific presentations, the Dagstuhl Seminar included several panel discussions and dropout sessions that carved out twelve critical research tasks for RNA-focused machine learning research. These tasks encompass understanding RNA molecules by identifying functional lncRNAs and their mechanisms of action, mapping chemical modifications and their effects, predicting RNA secondary and tertiary structures, and representing RNA conservation and functional regions. Furthermore, the challenge of understanding interactions and complexes was addressed, including context-aware prediction of regulatory targets, modeling RNA-protein complex formation including condensates, and constructing combinatorial maps of cellular compartments and their RNA-regulatory networks. In addition, the discussions highlighted several domain-specific challenges for machine learning in RNA biology as biological data is inherently

complex, often messy, and continually evolving, making data processing and curation particularly demanding. Benchmarking emerged as a crucial aspect of model evaluation as participants stressed the need for unbiased datasets, developing benchmarks for non-model organisms, and organizing community challenges to foster competition and innovation. Future directions emphasized were improving data curation and infrastructure, collaborative model development, transfer learning for cross-species applications, and creating practical, efficient models accessible to the academic community.

In sum, the Dagstuhl Seminar provided a rich forum for

bridging computational and experimental frontiers in RNA biology. Emerging deep learning algorithms are revealing complex layers of RNA regulation, while large-scale, high-resolution data-spanning ribosome profiling, single-cell labeling, nanopore direct RNA sequencing, and other transcriptomic techniques—continuously expand our understanding. Altogether, the meeting underscored the need for continued collaboration between computational method developers and bench biologists to realize the transformative potential of deep learning in decoding RNA-based regulation and its application to human health.

## 6.75 Human in the Loop Learning through Grounded Interaction in Games

**Editors: Raffaella Bernardi, Julia Hockenmaier, Udo Kruschwitz, and Massimo Poesio**  
**Seminar No. 24492**

Date: December 1–6, 2024 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.14.12.28

© Creative Commons BY 4.0 license

© Massimo Poesio, Raffaella Bernardi, Julia Hockenmaier, and Udo Kruschwitz



**Participants:** Malihe Alikhani, Elisabeth André, Raffaella Bernardi, Marc-Alexandre Côté, Simon Dobnik, Haishuo Fang, Jonathan Ginzburg, Ryuichiro Higashinaka, Julia Hockenmaier, Nikolai Illykh, Prashant Jayannavar, Alexander Koller, Udo Kruschwitz, Sharid Loáiciga, Catharine Oertel, Diego Perez Liebana, Massimo Poesio, Matthew Purver, David Schlangen, Carina Silberer, Edwin Simpson, Alessandro Suglia, Alane Suhr, Sina Zarrieß, Andrew Zhu

**Background and Motivation** Over the past few years, there has been a decisive move in Artificial Intelligence (AI) towards human-centered intelligence and towards AI models that can learn through interaction. An important reason for this shift has been the appearance of the latest generation of Large Language Models such as InstructGPT, ChatGPT, BARD, or Lamda-2 [12, 13, 17] capable of a step-increase in performance. A good part of the success of these models is due to the adoption of training regimes involving a combination of supervised learning and learning from interaction with humans, such as Reinforcement Learning Through Human Feedback [2, 13]. And particularly the most recent among such models, such as GPT-4, are not simply language models, but are trained with multimodal data and are capable of producing output in different modalities as well. However, these models still suffer from a number of widely discussed issues, such as hallucinations.

In parallel, there has also been substantial progress on grounded interaction – developing models aware of the situation in which they operate (a physical world in the case of robots, a virtual world in the case of artificial agents) and able to, e.g., understand / produce references to this situation [1, 4, 7, 8, 16] perhaps through negotiation [3]. However, the communication between the interactive learning and grounded interaction communities is still limited [10].

One domain considered particularly promising to study learning through grounded interaction with human agents is virtual world games: games in which conversational agents impersonating characters can learn to perform tasks, or improve their communicative ability, by interacting with human players in platforms such as Minecraft or Light [6, 9, 11, 15, 19, 24]. Games have been shown to be a promising platform for collecting data from thousands of players [20, 23]; virtual worlds approach the complexity of the real world; and virtual agents operating in such virtual worlds need to be able to develop a variety of interactional skills to be perceived as “real” [14].

This Dagstuhl Perspectives Workshop aimed, first of all, to bring together the communities working on the related areas of learning through interaction, (conversational) agents in games, dialogue and interaction, and collecting judgments from crowds through games, to make each community aware of the most recent developments in the other areas. We also intended to discuss current challenges, and whether advances in one area (e.g., grounded interaction) can benefit other areas (e.g., interactive learning).

**Directions Identified and Discussed** The workshop involved extensive discussions between researchers working in all the fields that contribute to the research area. After in-depth presentations of:

- The State of the Art (SOTA) in The Grounded and Communication Task Performance Abilities of (Embodied/Multimodal, Conversational) AI Agents (by group led by Hockenmaier)
- The Games and Multimodal Platforms Useful for Conversational AI Agents (group lead: Bernardi) and
- Current Approaches to Human (and Artificial Agent)-in-the-Loop Learning for AI Agents (group lead: Suglia),

and presentation of some of the most recent relevant research by the participants, we identified a few research directions particularly worth discussing in depth, and formed working groups around them. These included:

- **Complex Interaction.**

A common assumption in many computational models is that dialogue consists of a linear sequence of turns in which two agents alternately exchange information. Each turn is assumed to depend only on the last turn of the other participant. However, human conversation requires more complex forms of interaction spanning multiple turns to solve real-world tasks.

Complexity occurs for several reasons: Dialogue is done by multiple people. They start from different information states, have different perspectives, and cannot see what is in each other's minds. They have a social relationship that they have to manage. Their interaction happens in real time, across multiple modalities, in the presence not only of various kinds of noise but of fundamental asymmetries in what the participants can perceive, know and understand.

To successfully overcome these asymmetries and solve tasks through such interactions, the interaction scheme needs to offer a number of functions (see below). Among humans, these are exemplified by a variety of phenomena that depend not just on sequential information exchange but on more complex structures, with richer models of the local and global interaction context. It is not clear to what degree current LLM-based models of dialogue can cope with them, and how much this limits their ability to collaborate efficiently with humans.

The working group on Complex Interaction reviewed some of these complex interaction phenomena, gave pointers to the literature, and discussed ways in which future interactive systems might handle them.

#### ■ Game Design for Grounded Interaction.

Existing games and platforms used to evaluate and develop conversational agents are extremely diverse in their setting, goals, and complexities, and they are being developed in different subfields of AI, NLP, and computational linguistics [5]. Furthermore, within these subfields, games are designed for different purposes and, to some extent, classified using different taxonomies.

The explosion and diversity of games raise new research questions for these communities that we suggest should be explored in future research:

- Q1: How can games and game benchmarks be designed more systematically, such that they lead to a deeper understanding of games and the skills that games are testing? How do we generalize skills and agents' abilities across games?
- Q2: What role does the complexity of the game have? And how do we measure it?
- Q3: How do we evaluate agents within and across games? In particular, how do we evaluate whether the skills trained / tested with a game transfer to real world applications?

#### ■ Perspectives for Language Learning from Human Interaction.

Most of the recent AI breakthroughs have been in non-interactive settings: classical NLP tasks, math reasoning, etc. This was mainly due to the large availability of evaluation datasets in those domains. This is now changing, we see several new types of more realistic benchmarks that necessitate interacting within a given environment (WebArena [25], Webshop [22], OSWorld [21], AppWorld [18]). Learning paradigms that deal with interactivity need to be used.

Games are a convenient tool that enables us to construct scenarios that constructively approximate real scenarios. For example, the complexity of the games (what is observed vs what is learned, the search space for ML) can be iteratively increased or decreased and hence different learning methods can be studied in a more systematic and comparative fashion. Secondly, games, not being samples from real-world interactions but being close approximations of such interaction are a good way to engage human interactors to provide behavioural information and create consistent environments where data-collection (and if needed also data annotation) can be systematically performed.

The working group on Perspectives for Language Learning from Human Interaction produced a classification scheme of current ML approaches for learning from interaction, identifying a number of open questions, including:

- Q4: How can agents learn to have/recognise intentions?
- Q5: What are the tasks/games that can facilitate the acquisition of these skills?

#### ■ Perceptual Grounding for Embodied Conversational Agents.

This group built their discussion around the research hypothesis that interactivity plays a major role in human intelligence; interactivity has multiple aspects that we spelled out: 1) interacting with an environment (manipulate objects, act on them), 2) interacting with others through language, and 3) interacting with others while acting in an environment. Through such multimodal embodied experiences humans develop their cognitive intelligence (in other words, understanding the state of affairs in the world) and social intelligence (understanding the mechanisms of interactions).

## ■ References

- 1 Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. Babyai: A platform to study the sample efficiency of grounded language learning. In *7th International Conference on Learning Representations, ICLR*, 2019.
- 2 Paul F. Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, page 4299–4307, 2017.
- 3 Herbert H. Clark and Susan E. Brennan. Grounding in communication. In L. B. Resnick, J. Levine, and S. D. Behrend, editors, *Perspectives on Socially Shared Cognition*. APA, 1990.
- 4 Nicholas FitzGerald, Yoav Artzi, and Luke Zettlemoyer. Learning distributions over logical forms for referring expression generation. In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard, editors, *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925, Seattle, Washington, USA, October 2013. Association for Computational Linguistics.
- 5 Roberto Gallotta, Graham Todd, Marvin Zammit, Sam Earle, Antonios Liapis, Julian Togelius, and Georgios N. Yannakakis. Large language models and games: A survey and roadmap. ArXiv, 2024.
- 6 Matthew Johnson, Katja Hofmann, Tim Hutton, and David Bignell. The malmo platform for artificial intelligence experimentation. In *Proc. of IJCAI*, pages 4246–4247, 2016.
- 7 Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. ReferItGame: Referring to objects in photographs of natural scenes. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of*

- the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 787–798, Doha, Qatar, October 2014. Association for Computational Linguistics.
- 8 Casey Kennington and David Schlangen. A simple generative model of incremental reference resolution for situated dialogue. *Computer Speech and Language*, 41:43–67, 2017.
  - 9 Julia Kiseleva, Alexey Skrynnik, Artem Zholus, Shrestha Mohanty, Negar Arabzadeh, Marc-Alexandre Côté, Mohammad Aliannejadi, Milagro Teruel, Ziming Li, Mikhail Burtsev, Maartje ter Hoeve, Zoya Volovikova, Aleksandr Panov, Yuxuan Sun, Kavya Srinet, Arthur Szlam, Ahmed Awadallah, Seungeun Rho, Taehwan Kwon, Daniel Wontae Nam, Felipe Bivort Haiek, Edwin Zhang, Linar Abdrazakov, Guo Qingyam, Jason Zhang, and Zhibin Guo. Interactive grounded language understanding in a collaborative environment: Retrospective on Iglu 2022 competition. In *Proceedings of the NeurIPS 2022 Competitions Track*, PMLR, volume 220, pages 204–216, 2022.
  - 10 Jayant Krishnamurthy and Thomas Kollar. Jointly learning to parse and perceive: Connecting natural language to the physical world. *Transactions of the Association for Computational Linguistics*, 1:193–206, 2013.
  - 11 Anjali Narayan-Chen, Prashant Jayannavar, and Julia Hockenmaier. Collaborative dialogue in Minecraft. In *Proc. of the 57th Annual Meeting of the ACL*, pages 5405–5415, 2019.
  - 12 OpenAI. Chatgpt: A large-scale open-domain chatbot. Blog post, 2022.
  - 13 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
  - 14 David Schlangen. Dialogue games for benchmarking language understanding: Motivation, taxonomy, strategy. arXiv:2304.07007 [cs.CL], 2023.
  - 15 Arthur Szlam, Jonathan Gray, Kavya Srinet, Yacine Jernite, Armand Joulin, Gabriel Synnaeve, Douwe Kiela, Haonan Yu, Zhuoyuan Chen, Siddharth Goyal, Demi Guo, Danielle Rothermel, C. Lawrence Zitnick, and Jason Weston. Why build an assistant in Minecraft? arXiv: 1907.09273, 2019.
  - 16 Alberto Testoni and Raffaella Bernardi. “I’ve seen things you people wouldn’t believe”: Hallucinating entities in GuessWhat?! In Jad Kabbara, Haitao Lin, Amandalynne Paullada, and Jannis Vamvas, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 101–111, Online, August 2021. Association for Computational Linguistics.
  - 17 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
  - 18 Harsh Trivedi, Tushar Khot, Mareike Hartmann, Ruskin Manku, Vinty Dong, Edward Li, Shashank Gupta, Ashish Sabharwal, and Niranjan Balasubramanian. Appworld: A controllable world of apps and people for benchmarking interactive coding agents, 2024.
  - 19 Jack Urbanek, Angela Fan, Siddharth Karamcheti, Saachi Jain, Emily Dinan, Tim Rocktäschel, Douwe Kiela, Arthur Szlam, Samuel Humeau, and Jason Weston. Learning to speak and act in a fantasy text adventure game. arXiv preprint arXiv:1903.03094, March 2019.
  - 20 Luis von Ahn. Games with a purpose. *Computer*, 39(6):92–94, 2006.
  - 21 Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments, 2024.
  - 22 Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents, 2023.
  - 23 Juntao Yu, Silviu Paun, Maris Camilleri, Paloma Carretero Garcia, Jon Chamberlain, Udo Kruschwitz, and Massimo Poesio. Aggregating crowdsourced and automatic judgments to scale up a corpus of anaphoric reference for fiction and wikipedia texts. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, page 767–781, Dubrovnik, Croatia, 2023. Association for Computational Linguistics (ACL).
  - 24 Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and Prithviraj Ammanabrolu. I cast detect thoughts: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 11136–11155, Toronto, CAN, 2023. Association for Computational Linguistics.
  - 25 Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents, 2024.

## 6.76 Coding Theory and Algorithms for Emerging Technologies in Synthetic Biology

**Editors: R. B., Olgica Milenkovic, Zohar Yakhini, and Yonatan Yehezkeally**  
**Seminar No. 24511**

Date: December 15–20, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.12.46

© Creative Commons BY 4.0 license

© R. B., Olgica Milenkovic, Zohar Yakhini, and Yonatan Yehezkeally

**Participants:** Roei Amit, Iryna Andriyanova, R. B., Anisha Banerjee, Daniella Bar-Lev, Jessica Bariffi, Salim El Rouayheb, Ohad Elishco, Nick Goldman, Alexandre Graell i Amat, Francesca Granito, Robert Grass, Jasper Groen, Anina Gruica, Serge Kas Hanna, Cai Kui, Olgica Milenkovic, Lior Nissim, Tzachi Pilpel, Nimesh Pinnamaneni, Inbal Preuss, Roni Rak, João Ribeiro, Eirik Rosnes, Omer Sabary, Benno Schwikowski, Ilan Shomorony, Roman Sokolovskii, Mark Somoza, Kasra Tabatabaei, Jennifer Tang, Emanuele Viterbo, Van Khu Vu, Frederik Walter, Zhiying Wang, Eitan Yaakobi, Zohar Yakhini, Yonatan Yehezkeally



Designing DNA-based storage systems intrinsically requires joint efforts between biologists, chemists, engineers, and computer scientists, as prominent differences from classical storage media exist at all stages. For example, cost-effective synthesis introduces insertion and deletion errors on top of the well-understood substitution errors occurring in classical media, and much less is known for correcting these. Error-correction techniques could also be affected by the targeted application due to the intrinsic properties of the stored data and the effects of the different types of errors (e.g., this phenomenon is observed when storing images in DNA-based storage systems). Further, strands in the storage container are not ordered in the memory, thus, during sequencing, it is not possible to distinguish which strand is being read, making error correction even more challenging. Lastly, sequencing techniques allow observing many copies of erroneous sub-sequences of the stored strands, which can be leveraged to reconstruct the stored strands more efficiently.

The first large-scale experiments that demonstrated the potential of *in vitro* DNA storage were reported by Church et al., who recovered 643KB of data [1], and Goldman et al., who accomplished the same task for a 739 KB message [2]. However, both of these groups did not recover the entire message successfully due to the lack of using the appropriate coding solutions to correct errors. Most published studies report that either substitutions or deletions are the most prominent error types in DNA-based storage systems, depending upon the specific technology for synthesis and sequencing. Thus, coding-theoretic aspects of DNA-based storage systems have received significant attention recently. However, these theoretical works have not yet led to viable storage technologies.

The progress in enabling information storage in DNA has been driven by the progress in using synthetic DNA in more general applications. In [3], the authors demonstrated how high throughput synthesis can be used to understand, optimize, and

fine-tune the functionality of biological systems. This work involved careful design of the reagents – the composition of a large library of candidates – as well as rigorous statistical data analysis to support the result. This Dagstuhl Seminar, therefore, covered general design and analysis frameworks for high throughput experiments. In particular, one specific (and prominent) type of high throughput experiments that is strongly related to synthetic DNA is CRISPR screening. These experiments involve the silencing or the activation of a large number of elements in genomes to allow for optimizing and tuning certain outcomes, ranging from growth rates in plants and bovine cultures to insinuating immune responses in cancer patients.

Informed by this observation, this seminar ultimately aimed at forging closer connections between information theorists, computer scientists, data scientists, biologists, and chemists to: (i) drive joint progress in coding-theoretic techniques specifically tailored to the emerging synthesis sequencing technologies; (ii) have a better understanding of, and initiate innovation in, the application of computer science techniques for high throughput experimental synthetic biology; and (iii) shape an application-driven design of low-error cost-effective DNA-based storage systems. The seminar schedule was flexibly designed, allowing participants to present their research and expertise while interactively accommodating audience input. The plenaries exposed participants to assorted underlying fields, namely genetic code, CRISPR and gene editing, bio-informatics, informatics and machine learning for medical applications, the utility of coding theory for DNA-based information systems, and market data-storage applications. Meanwhile, working groups enabled participants to leverage interdisciplinary backgrounds and share their knowledge and expertise to envision holistic solutions to contemporary challenges. To advertise the research of junior participants, the schedule included a handful of short talks to showcase their results. During the discussions, a couple of

participants noticed the implications of their research on the discussed topics; therefore, “pop-up” talks were scheduled for those participants to share their thoughts. Throughout, much fun

was had, and connections were forged in a myriad of ice-breaking activities.

## ■ References

- 1 G. M. Church, Y. Gao, and S. Kosuri, “Next-generation digital information storage in DNA,” *Science*, no. 6102, pp. 1628–1628, Sep. 2012.
- 2 N. Goldman, P. Bertone, S. Chen, C. Dessimoz, E. M. LeProust, B. Sipos, and E. Birney, “Towards practical, high-capacity, low-maintenance information storage in synthesized DNA,” *Nature*, no. 7435, pp. 77–80, Jan. 2013.
- 3 E. Sharon, Y. Kalma, A. Sharp, T. Raveh-Sadka, M. Levo, D. Zeevi, L. Keren, Z. Yakhini, A. Weinberger, and E. Segal, “Inferring gene regulatory logic from high-throughput measurements of thousands of systematically designed promoters,” *Nature biotechnology*, vol. 30, no. 6, pp. 521–530, 2012.

## 6.77 Quantum Software Engineering

**Editors: Shaukat Ali, Johanna Barzen, Andrea Delgado, Hausi A. Müller, and Juan Manuel Murillo**

**Seminar No. 24512**

Date: December 15–20, 2024 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.14.12.63

© Creative Commons BY 4.0 license

© Shaukat Ali, Johanna Barzen, Andrea Delgado, Hausi A. Müller, and Juan Manuel Murillo

**Participants:** Shaukat Ali, Paolo Arcaini, Antonio Brogi, José Campos, Schahram Dustdar, Michael Falkental, Sebastian Feld, Michael Felderer, José Manuel García Alonso, Ignacio García Rodríguez de Guzmán, Kostas Magoutis, Wolfgang Mauerer, Andriy Miranskyy, Anila Mjeda, Hausi A. Müller, Juan Manuel Murillo, Yehuda Naveh, Ricardo Pérez-Castillo, Antonio Ruiz Cortés, Shinobu Saito, Ina Schaefer, Laura Schulz, Ulrike Stege, Robert Wille, Manuel Wimmer, Lei Zhang



Under the umbrella of quantum information theory, quantum algorithms have been proposed that solve polynomial-time problems for which no classical algorithm solves in that order of complexity. The entry into the NISQ era with the development of the first quantum computers and simulators makes testing such algorithms on real computers possible. This generates new expectations and promotes growing interest in Quantum Computing. Researchers are proposing solutions in the quantum domain to optimize existing algorithmic approaches regarding the number of qubits or circuit depth, propose solutions to manage decoherence and quantum errors, or propose new algorithms for unexplored problems. All these efforts, coupled with the fact that numerous media reports on quantum computing achievements appear, are inviting researchers and practitioners from industry, academia and government to investigate the utility of quantum computing to explore the possibilities it can offer. Some countries have launched specific programs encouraging industry and academia to examine together the applicability in different sectors, the current limitations of quantum technologies and their expected evolution over time. The objective is to prepare for the future uptake of quantum computing in industry in many application domains.

All of the above is evidence that quantum software is already a reality. It is commonly accepted that future software will be hybrid, integrating both classical and quantum components because each quantum algorithm requires classical preprocessing or postprocessing. With this, interest is also beginning to appear from researchers in the field of software engineering. The question arises whether the software engineering body of knowledge acquired over the last decades is applicable in this new computing era. The first forums focusing on the field of Quantum Software Engineering (QSE) have already appeared. Examples are Q-SE at ICSE, QSW at IEEE Services or Q-SET at IEEE QCE. Many interesting works have been presented in these forums on approaching certain software engineering practices now in the

context of quantum computing. Topics such as building hybrid software architectures, orchestrating quantum systems, building models that enable automatically generating quantum software, and creating quantum software with better quality attributes are being addressed in these forums.

Although all these works cover interesting specific aspects that different research teams have addressed, there is a need for a deeper reflection on the nature and scope of QSE. This reflection should address questions such as the differences between classical and quantum software engineering, about which features of quantum software can be addressed with the known classical techniques, which features cannot be addressed with classical techniques and why, which classical techniques should be extended or modified to address quantum software development and which new software engineering techniques should be developed to address quantum software from an engineering perspective.

As the area of QSE can be very large, the Dagstuhl Seminar on “Quantum Software Engineering” focused on three main work topics described below.

**Quantum Software Design, Modelling, and Architecturing** Each software development methodology embodies its own design artefacts. The question arises as to which design artefacts of the methodologies suit quantum parts in hybrid systems. We may not have such artifacts; we need new ones or known artifacts with new features. In the same way, it could be asked if the conventional classical methodologies are adequate for the development of hybrid systems or if it is necessary to define or adapt our known methodologies. In particular, Model-driven engineering (MDE) has been one of the areas of research that has attracted significant interest over the years and has affected many changes in software engineering. It is now necessary to explore how its principles are adapted to capture the quantum nature of hybrid systems. Some of the contributions that could emerge are

new design techniques to address hybrid software development, differentiation and integration of classical and quantum parts during the design, new design objects suitable to the characteristics of quantum computing, specific domain languages for quantum systems, platform-specific models for hybrid systems, reverse engineering of legacy systems for building new hybrid systems or model transformation for hybrid systems. Finally, how to compose hybrid systems from simpler classical and quantum parts is one of the aspects that we should most rapidly address to facilitate the construction of large-scale computational systems involving quantum computation. There are interesting aspects to be addressed here, such as the orchestration of quantum programs, techniques to deal adequately with the classical preprocessing and postprocessing needed by quantum algorithms, and techniques for the interaction between quantum and classical code that do not avoid the benefits provided by service-oriented computing.

**Adaptive Hybrid Quantum Systems** Quantum systems combine classical computers with quantum processors to realize hybrid solutions to practical computing problems, including optimization, machine learning and nature simulation. Classical computers are involved in every layer of the quantum stack, from high-level algorithm design to low-level qubit control. The success of quantum computing relies on the seamless integration of classical and quantum components. Optimizing performance requires considering the interactions between different layers of the quantum stack. Thus, the quantum stack demonstrates the inherent hybrid nature of quantum systems.

Quantum system adaptivity is equally essential to the success of quantum computing. Adaptive quantum systems dynamically optimize the use of quantum and classical resources during computation at different levels of the quantum stack. Adaptive systems rely on various technologies to achieve self-\* properties, such as self-healing, self-optimizing, self-configuring, or self-adjusting behavior. Key technologies to adapt different levels of the quantum stack include quantum runtimes, parameterized circuits, mid-circuit measurements, dynamic circuits, and circuit weaving to achieve real-time quantum error correction, dynamic calibration of qubit properties, noise mitigation, parameter adjustments in variational algorithms, qubit entanglement routing, or adaptive circuit decomposition. These highly dynamic quantum systems will significantly benefit from software engineering technologies such as feedback loops, models at runtime, assurance at runtime, autonomic systems, digital twins, software-defined infrastructure, and compute & storage clouds.

The rapid proliferation of distributed quantum computing (DQC) amplifies the need for adaptive quantum system technologies to help optimize entanglement distribution across quantum chips, manage network dynamics & latency, and problem decomposition across the network. Adaptive technologies are the key to unlocking the full potential of DQC and enabling the development of large-scale, fault-tolerant, and resilient quantum computers.

**Quantum Software Quality Assurance** The construction of software from an engineering perspective always demands that the product obtained has adequate quality attributes. Building

a product that is not profitable due to its final characteristics or the cost necessary for its construction is not a valid solution. Quantum software cannot suffer from the lack of these quality attributes. So, it is necessary to guarantee attributes such as reusability or maintainability. How to promote these attributes during the process of building new quantum software is still a task that needs to be addressed. Related to the task of guaranteeing that software meets the required quality attributes is the discipline of software testing. Software testing groups together a set of practices aimed at examining software artefacts and their operation to guarantee that the software has adequate quality attributes and behaves adequately with respect to what it was conceived for. In classic software, many types of testing can be practised, such as white-box, black-box, grey-box, dynamic, or static. It is necessary to determine whether all these testing types are applicable during quantum software development. Also, when dealing with quantum software, it is necessary to study the need for different testing, at least at two levels. On the one hand, testing the classical software that generates the quantum circuit requires further investigation, where classical testing techniques can be applied. On the other hand, testing quantum circuits requires studying new testing techniques. Related to the above, formal methods for ensuring correctness and checking quantum software are also engaging in this topic.

**Seminar agenda and development** The seminar agenda started on the first day with a session dedicated to personal presentations, where attendees described their focus area and the related problems that interest them. Next, the rest of the day was devoted to the topic “When Software Engineering meets Quantum Mechanics.” Five interesting talks were given on this topic.

The talks continued on the second day to address “Quantum Software Engineering and its Challenges.” Eleven talks were given on this topic.

From the third day onwards, the seminar focused on discipline-specific group discussions. To build groups of interest for all attendees, the theme of the groups and dynamics were discussed on the morning of the third day.

During each day, coffee breaks were scheduled with sufficient time for interaction and free discussion among the participants. In addition, at the end of the first, second and fourth days, we orchestrated sessions for plenary discussion on topics of relevance to the participants. During the seminar, a board was available for interested attendees or groups of attendees to write down the topics they would like to address during those highly interactive plenary sessions.

During the off-schedule time, the participants enjoyed much social interaction, including long singing sessions accompanied by the excellent guitar player Prof. Hausi A. Müller. This singing was not unrelated to the discipline of quantum computing. Far from it, a song written by Prof. Müller was rehearsed and performed multiple times by the attendees in the renowned Dagstuhl music room and on stage at the closing session on Friday morning. The lyrics of this song are also an outcome of the seminar and are therefore included as a panel discussion at the end of the full report.

# **7** **Öffentlichkeitsarbeit** *Public Relations and Outreach*

## Pressemitteilungen und Medienarbeit

7.1

## Press Releases and Media Work

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Vermittlung des Konzepts von Schloss Dagstuhl ist dabei ebenfalls ein Thema. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl<sup>46</sup> abrufbar.

Schloss Dagstuhl hat sich zur allgemeinen Anlaufstelle für Journalisten etabliert, die über bestimmte Informatikthemen, aber auch über Schloss Dagstuhl berichten möchten.

Schloss Dagstuhl verbreitet Neuigkeiten rund um sein Programm über soziale Netzwerkdienste wie LinkedIn, wo interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben werden. Inzwischen hat Schloss Dagstuhl bei LinkedIn über 800 Follower, wobei im Berichtsjahr im Durchschnitt fast jeden Tag einer dazu kam, und viele unserer Gäste teilen dort ihre Dagstuhlerfahrungen. Bei LinkedIn wird auch schon seit vielen Jahren eine eigene Gruppe Friends of Schloss Dagstuhl mit über 600 Mitgliedern unterhalten, mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Seitdem Dagstuhl X nicht mehr aktiv nutzt, sind einige neue und daher noch kleine Social-Media-Präsenzen, etwa für dblp auf Mastodon und Bluesky, hinzugekommen.

Regular press releases showcase and disseminate information about current computer science topics in a comprehensible manner and clarify the concept behind Schloss Dagstuhl. Press releases and media reports that come to the center's attention are available on the Schloss Dagstuhl website<sup>46</sup>.

Schloss Dagstuhl has become a port of call for journalists seeking to report on specific computer science topics and/or on Schloss Dagstuhl itself.

Schloss Dagstuhl spreads news about its program via social networking services such as LinkedIn, where interesting news about Schloss Dagstuhl is announced. Schloss Dagstuhl now has over 800 followers on LinkedIn, with an average of almost one new follower a day during the report year, and many of our guests share their Dagstuhl-experiences there. For many years, Schloss Dagstuhl has also maintained its own Friends of Schloss Dagstuhl group on LinkedIn with over 600 members, with the aim of supporting the networking of participants in Dagstuhl seminars. Since Dagstuhl no longer actively uses X, some new and therefore still small social media presences have been added, for example for dblp on Mastodon and Bluesky.

## Fortbildung

7.2

## Educational Training

### Lehrerfortbildung

Seit nunmehr 34 Jahren engagiert sich Schloss Dagstuhl im schulischen Bereich durch die Organisation einer jährlichen Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem Bildungscampus Saarland und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Diese beiden Institute unterstützen die 32. Fortbildung 2024 auch finanziell, indem sie die Kosten der Referenten tragen.

Die Lehrerfortbildung dauert drei Tage; an jedem Tag werden in jeweils 3-stündigen Vorträgen zwei Informatikthemen vorgestellt. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch können seit 2011 bis zu fünf Lehrer aus anderen Bundesländern teilnehmen. Mehr Informationen zur Lehrerfortbildung 2024 gibt es auf der Webseite der

### Teacher training

Since 34 years, Schloss Dagstuhl hosts an annual teacher training workshop specifically designed for teachers of upper secondary students working in the Saarland or the Rhineland Palatinate.

The workshop is organized together with the Bildungscampus Saarland and the Pädagogisches Landesinstitut Rheinland-Pfalz (PL). These two institutes support the 32<sup>nd</sup> training workshop financially as well in 2024, by assuming the costs of speakers.

The workshop lasts three days; each day two computer science topics are presented in a three hour presentation each. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, since 2011 up to five teachers of other federal states can participate. Details on the workshop in 2024 are available at the event webpage<sup>47</sup>.

<sup>46</sup> <https://www.dagstuhl.de/institute/press>

Veranstaltung<sup>47</sup>.

7

## „Dagstuhler Gespräche“

7.3

## “Dagstuhler Gespräche”

Um die Türen des Schlosses etwas weiter für die Allgemeinheit und die Region zu öffnen, hat Schloss Dagstuhl zusammen mit der Stadt Wadern die Veranstaltungsreihe *Dagstuhler Gespräche* ins Leben gerufen. Der interessierten Öffentlichkeit werden hier Themen aus dem breiten Spektrum der Informatik sowie ihre praktische Anwendung im Alltag oder in wirtschaftlichen Prozessen anschaulich in Form eines Impulsvortrages näher gebracht, um danach in einen gemeinsamen Dialog einzusteigen. An den Dagstuhler Gesprächen nehmen Entscheider und Gestalter aus Wirtschaft, Politik und der Informatik teil, aber auch Interessierte aus der Bevölkerung sind herzlich eingeladen.

Am 26. Januar 2024 hielt Prof. Dr. Friedemann Mattern, emeritierter Professor für Informatik an der ETH Zürich, einen Vortrag zum Thema „Computer, Digitalisierung und die Energie“. Im Vortrag wurden Trends des Energieverbrauchs und anderer Umweltaspekte der immer weiter fortschreitenden Digitalisierung diskutiert und bewertet. Dabei wurde aufgezeigt, dass es vielleicht weniger die offensichtlichen Aspekte sind, die kritisch hinsichtlich der Umweltauswirkungen zu sehen sind, als vielmehr einige problematische indirekte Folgen.

Am 22. November 2024 sprach Prof. Dr. Eduard Arzt, der zwischen 2007 und 2022 einer der beiden wissenschaftlichen Geschäftsführer des Leibniz-Institut für Neue Materialien in Saarbrücken war, zum Thema „Lernen von der Natur – vom Staunen zum Startup“. Der Vortrag beschrieb, wie eine bioinspirierte Technik von der Entwicklung im Labor bis zu einer saarländischen Firmengründung, die eine nachhaltige Technologie auf den Markt bringt, den Weg nahm: Bei der betrachteten Technologie werden feine „haarige“ Mikrostrukturen, ähnlich denen beim Gecko, in der Automatisierungstechnik benutzt, um z. B. kleine Objekte zu greifen und wieder abzusetzen.

In order to open its doors a bit further for the general public and the local region, Schloss Dagstuhl, together with the town of Wadern, has initiated the *Dagstuhler Gespräche* (“Dagstuhl conversations”) event series. The interested public will be introduced to a broad spectrum of topics from computer science, as well as to practical applications of those topics in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

On January 26, 2024, Prof. Dr. Friedemann Mattern, Professor Emeritus of Computer Science at ETH Zurich, gave a lecture on the topic “Computers, Digitization and Energy”. The lecture discussed and evaluated trends in energy consumption and other environmental aspects of the ever-increasing digitalization. It was shown that it is perhaps not so much the obvious aspects that should be viewed critically in terms of environmental impact, but rather some problematic indirect consequences.

On November 22, 2024, Prof. Dr. Eduard Arzt, who was one of the two scientific directors of the Leibniz Institute for New Materials in Saarbrücken between 2007 and 2022, gave a lecture on the topic “Learning from nature – from amazement to startup”. The lecture described how a bio-inspired technology made its way from development in the laboratory to the founding of a company in Saarland that brings a sustainable technology to market: In the technology under consideration, fine “hairy” microstructures, similar to those found in geckos, are used in automation technology to, for example, grasp and release small objects.

<sup>47</sup> <https://www.dagstuhl.de/24503>



# 8

## Einrichtungen *Facilities*

Das Zentrum verfügt über drei Standorte; der Hauptstandort ist Schloss Dagstuhl in Wadern, an dem auch die Veranstaltungen stattfinden. Die Vorbereitung der Seminare findet hauptsächlich am Standort Saarbrücken statt. Am Standort Trier arbeiten die meisten Mitarbeiter der Bibliographiedatenbank dblp. Dagstuhl Publishing ist in Saarbrücken und Wadern angesiedelt.

Durch die neuen Räumlichkeiten für Schloss Dagstuhls Mitarbeiter ist das Konferenzzentrum in Wadern inzwischen der zentrale Treffpunkt aller Projekt- und wissenschaftlichen Mitarbeiter, die ansonsten getrennt über die verschiedenen Standorte verteilt beziehungsweise im Homeoffice arbeiten. Hier können sie sich austauschen und gemeinsam vor Ort an den aktuellen Projekten und Aufgaben arbeiten.

The institute operates over three locations: the main site is Schloss Dagstuhl in Wadern, where the events also take place. The staff in charge of seminar preparation are located at the Saarbrücken site. Most of the employees operating the dblp computer science bibliography are based at the Trier site. Dagstuhl Publishing is located in Saarbrücken and Wadern.

Thanks to the new premises for Schloss Dagstuhl's employees, the conference center in Wadern is now the central meeting place for all project and scientific staff, who are otherwise working in separate locations or in home office. Here, they can exchange ideas and work together on current projects and tasks on site.

## Hauptstandort in Wadern

8.1

## Main Site in Wadern

Der Hauptstandort in Wadern umfasst das historische Schloss (gebaut um 1760) mit einem Anbau aus den 1970ern, einem 1993 fertiggestellten Erweiterungsbau, in dem sich Forschungsbibliothek, Hörsäle, Gästezimmer, Büros und Infrastruktur befinden, und ein 2012 fertiggestelltes Gästehaus mit Gästezimmern, einem Konferenzraum und Räumlichkeiten der Gebäudeverwaltung. Alle Einrichtungen in Wadern werden, abgesehen von je zwei Wochen im Sommer und Winter, die für größere Instandhaltungsarbeiten genutzt werden, ganzjährig für Veranstaltungen genutzt.

Die Kapazitäten von Dienstleistungen und Räumlichkeiten zur Veranstaltung von Seminaren sind genau aufeinander abgestimmt: Das Zentrum hat 70 Gästezimmer, davon sind in 16 Zimmern zwei Betten oder eine zusätzliche Schlafcouch, sodass insgesamt 88 Personen übernachten können. Die Doppelzimmer werden überwiegend von Paaren oder Elternteilen mit Kindern genutzt. Bei Normalbetrieb finden parallel zwei Veranstaltungen mit etwa 28 und etwa 42 Teilnehmern statt, wobei jedem Seminar ein Hörsaal für 38 beziehungsweise 50 Personen zur Verfügung steht. In den seltenen Fällen, in denen die vorhandenen Zimmer nicht ausreichen, können Seminargäste in Doppelzimmern oder einem nahegelegenen Hotel untergebracht werden. In vielen Wochen des Jahres sind alle verfügbaren Zimmer auch tatsächlich von Gästen belegt.

The main site in Wadern comprises the historic manor house (built around 1760) with an extension from the 1970s; a facility completed in 1993, which is housing a research library, lecture halls, guest rooms, offices, and infrastructure; and a guest house completed in 2012 with guest rooms, a conference room, and garages for facility management. All facilities at the site in Wadern are operated all year round, with the exception of two weeks in summer and two weeks in winter – which are used for larger maintenance tasks.

The capacities of services and facilities for hosting seminars at the main site are well coordinated: the site has 70 guest rooms, including 16 double rooms, such that the total capacity for people staying overnight is 88. The double rooms are mainly used by couples or parents with children. During normal operation, two seminars with around 28 and 42 participants are hosted in parallel, each using a lecture hall with 38 and 50 seats, respectively. In the rare case that there are not enough rooms available, it is possible to book seminar guests into double rooms or a nearby hotel. In many weeks of the year, all available room are occupied by guests.

### ■ Tagungsräume

Schloss Dagstuhl bietet drei Hörsäle unterschiedlicher Größe. Während der kleinste mit Tischen in U-Form für 24 Personen ausgelegt ist, kann der größte ohne Tische bis zu 60 Personen fassen. Jeder Hörsaal ist mit einem Beamer ausgestattet. Dank „Airservers“ ist eine drahtlose Videübertragung von Rechnern der Teilnehmer auf den Beamer mittels Airplay, Google Cast oder Miracast möglich. Daneben kann auch ein vorhandener MS-Windows-Laptop zur Präsentation genutzt werden.

Alle Hörsäle sind auch für Hybridveranstaltungen aus-

### ■ Conference Facilities

Schloss Dagstuhl offers three lecture halls of different sizes. While the smallest one is designed for 24 people with tables in a U-shape, the largest one can hold up to 60 people without tables. Every lecture hall is equipped with a projector. An “Airservers” enables wireless video transmission from participants' computers to the projector via Airplay, Google Cast, or Miracast. In addition, an existing MS Windows laptop can also be used for presentations.

All lecture halls are equipped for hybrid events, which can include a combination of on-site participants and those

gestattet, bei denen ein Teil der Teilnehmer vor Ort sind während andere remote per Videokonferenztechnik zugeschaltet sind. Dabei ist es egal, ob die Remoteteilnehmer mittels eines persönlichen Rechners teilnehmen, oder einen mit entsprechender Technik ausgestatteten Konferenzraum nutzen. Die Technik dafür umfasst eine Audioanlage mit Lautsprechern und verschiedenen Mikrofonen. Neben einem Kopfbügelmikrofon und Handmikrofon sind mehrere Mikrofone an den Decken montiert, sodass alle Teilnehmer frei diskutieren können. Für die Videoübertragung stehen eine fernsteuerbare Kamera in Richtung des Vortragenden und der Tafeln sowie eine fest ausgerichtete Kamera auf das Publikum zur Verfügung. Weiterhin ist ein zusätzlicher Beamer beziehungsweise ein 85 Zoll Display zur Anzeige der Remote-Teilnehmer vorhanden.

Neben den Hörsälen gibt es im Zentrum sechs Seminarräume. Davon sind zwei mit modernen Beamern ausgestattet, während in einem ein großes Plasmadisplay montiert ist. Weitere Beamer auf Rollwagen stehen zusätzlich zur flexiblen Benutzung in allen Räumen zur Verfügung.

Alle Hörsäle und andere Tagungsräume sind mit Tafeln und/oder Whiteboards ausgestattet. Weitere mobile Flipcharts stehen flexibel für alle Veranstaltungen bereit. Die drei Hörsäle sind jeweils mit mehreren Tafeln ausgestattet, die beiden kleineren davon zusätzlich mit Whiteboards, während in den anderen Tagungsräumen jeweils große Whiteboards an den Wänden montiert sind. Daneben gibt es über das ganze Zentrum verteilt weitere Räume, in denen Gäste sich in entspannter Atmosphäre treffen und diskutieren können.

Insbesondere am Abend zieht es viele Gäste in den Weinkeller und die Cafeteria, zwei der gemütlichsten Räume im Haus und hervorragend geeignet für die Fortsetzung einer produktiven Diskussion in angenehmer Atmosphäre. Ein Whiteboard im Weinkeller und zwei Tafeln im Außenbereich ermöglichen den Gästen auch in informeller Atmosphäre ihre Gedanken visuell zu präsentieren.

### ■ Dagstuhl's Küche

Die Mahlzeiten sind ein wichtiger Bestandteil des wissenschaftlichen Programms von Schloss Dagstuhl. Die Sitzordnung wird absichtlich stets zufällig gemischt, um eingefahrene Gruppen aufzuteilen und Gäste zu ermuntern, während ihres Aufenthalts möglichst viele verschiedene Kolleginnen und Kollegen kennenzulernen. Große Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten.

Dagstuhl's Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden jeden Tag von unseren Mitarbeitern der Küche frisch zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dies steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber der Mehrheit der internationalen Gäste des Zentrums durchaus entgegen.

Sowohl die Zutaten als auch die Gerichte wechseln saisonal. Über das Jahr hinweg wird eine ausgewogene Mischung an regionalen und internationalen Spezialitäten

participating remotely via video conferencing technology. The technology offered includes an audio system with speakers and various microphones. In addition to a headset microphone and a hand-held microphone, several microphones are mounted on the ceiling to enable fluid discussions between all participants. For the video transmission, we offer a remote-controlled camera facing the speaker and the boards, as well as a fixed camera capturing the on-site audience. Furthermore, an additional projector or an 85-inch display is available for showing the remote participants.

In addition to the lecture halls, the center has six meeting rooms. Two are equipped with up-to-date projectors and one has a large plasma display on the wall. Mobile projectors are available for use in all other rooms.

All lecture halls and meeting rooms are equipped with blackboards and/or whiteboards. Additional mobile flipcharts are available for all events. Whereas the three main lecture halls are equipped with several blackboards and the two smaller ones of them with whiteboards as well, big whiteboards are provided in the other meeting rooms. The center also offers a variety of other spaces where guests can sit and work together in a relaxed atmosphere.

Particularly in the evening, guests gravitate towards the wine cellar and the "old" cafeteria, two of the coziest places on the premises and ideal for continuing a productive discussion in a pleasant atmosphere. The wine cellar contains a whiteboard and there are two blackboards in the outdoor area near the dining hall, so guests can spontaneously pitch their ideas visually in an informal atmosphere.

### ■ Dagstuhl's Kitchen

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are mixed deliberately in order to break up cliques and encourage guests to talk to as many different people as possible during the course of their stay. Large tables in the dining hall promote collaborative interaction during meals.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy, and tasty meals. Everything is freshly prepared each day by the kitchen's staff. The focus is on lighter fare during the day in order to aid scientists' concentration, and on a warm meal in the evening, breaking with the German tradition of a cold evening meal while matching the internationality of the center's guests.

Both ingredients and dishes vary with the seasons. In general, the kitchen tries to keep meals lighter in the summertime and heavier in the winter, offering a blend of regional and international dishes year-round that include some new recipes and many tried-and-true Dagstuhl favorites. The kitchen works in accordance with the HACCP concept (Hazard Analysis and Critical Control

aus neuen sowie bewährten und beliebten Rezepten angeboten. Im Allgemeinen sind die angebotenen Gerichte im Sommer etwas leichter und im Winter ein wenig nahrhafter. Die Küche arbeitet nach dem HACCP-Konzept (Hazard Analysis and Critical Control Points) und hält sich an die Kennzeichnungspflicht von Allergenen, zu der alle lebensmittelverarbeitenden Betriebe verpflichtet sind. Des Weiteren achten wir auf deklarationsfreie Zusatz- und Konservierungsstoffe.

Alle Gäste, die aus medizinischen oder ethischen Gründen Einschränkungen bei der Speiseauswahl haben, können sich vor dem Seminar bei Schloss Dagstuhl melden. Unsere Küchenmitarbeiter erarbeiten gerne individuelle Lösungen für jeden Gast, soweit es irgend möglich ist. Gäste, die koscheres Essen benötigen, haben die Möglichkeit, mitgebrachte abgepackte Speisen selbst zu erhitzen.

Um unseren Gästen trotz eines begrenzten Budgets eine ausgewogene Qualität anbieten zu können, bietet unsere Küche ein Frühstücksbuffet, sowie im allgemeinen ein Mittagsbuffet und ein Menü am Abend. Abhängig von den personellen Kapazitäten und der Auslastung des Zentrums wird aber manchmal Mittags ein Menü serviert oder abends stehen an einer Theke ein reguläres und ein vegetarisches Menü zur Auswahl. Unser Restaurant mit den großen Fenstern zum Garten des Hauptgebäudes bietet etwa 80 Personen Platz. Hier herrscht eine entspannte und fast familiäre Atmosphäre, was nicht zuletzt auf unsere freundlichen und engagierten Mitarbeiter zurückzuführen ist.

Kleine und große Pausen unterbrechen auf angenehme Weise die tägliche Routine und die anstrengenden Diskussionen. In der kleinen Kaffeepause am Vormittag steht in den Pausenräumen der beiden großen Vortragssäle ein Kaffeevollautomat zur Zubereitung von Kaffee, Kakao und Tee zur Verfügung. In der großen Kaffeepause am Nachmittag wird den Gästen im Speiseraum neben heißen Getränken auch frisch gebackener Kuchen angeboten. Darüber hinaus gibt es im Gästehaus, der „alten“ Cafeteria und dem Weinkeller jeweils einen weiteren Kaffeevollautomaten. Im Kiosk vor der Cafeteria können Gäste Snacks erwerben. Abends gibt es in der Cafeteria und im sogenannten Weinkeller einen Gruß aus der Küche, bestehend aus Brot und einer Käseauswahl.

### ■ Kinderbetreuung

Schloss Dagstuhl bietet Teilnehmern, die mit Kindern anreisen müssen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser von Schloss Dagstuhl subventionierte Service kann gegen ein geringes Entgelt im Voraus gebucht werden. Alternativ ist es Eltern auch möglich, eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Schloss Dagstuhl kommt für die Unterkunft und Verpflegung der Kinder auf. Wenn statt Inanspruchnahme der Kinderbetreuung von Schloss Dagstuhl eine Betreuungsperson mitreist, hat diese ebenfalls freien Aufenthalt. Schloss Dagstuhl hält einen Aufenthaltsraum für Kinder vor. Dieser ist unter anderem mit einer Sitzgruppe, einem Spielhaus, einer Spielhöhle und Kinderzimmerteilern, die etwa mit einer Kugelbahn oder einem Spiegel ausgestat-

Points) and adheres to the mandatory labeling of allergens, which is required of all food processing establishments. Food additives and conservatives for which labeling is non-mandatory are also carefully monitored.

All guests with special dietary requirements due to ethical or health reasons can announce their needs prior to the events. Our kitchen staff will then work out individual solutions if at all possible. Guests who need kosher meals can bring their own ready-to-eat meals and heat them up for themselves.

In order to offer our guests a balanced quality despite a limited budget, our kitchen offers a breakfast buffet, as well as a lunch buffet and a menu in the evening. Depending on the personnel capacities and the overall booking situation, however, sometimes a set menu is served at lunchtime or in the evening there is a regular and a vegetarian menu to choose from at a bar. The large dining-hall, seating up to 80 people, opens onto the castle garden and patio, and offers a relaxed, familiar atmosphere.

Small and large breaks are a welcome interruption to the daily routine and intense discussions. During the small coffee break in the morning, a fully-automatic coffee machine is available in the break rooms of the two large lecture halls; offering coffee, hot chocolate, and tea. During the longer coffee break in the afternoon, hot drinks together with freshly baked cake are served in the dining hall. There are also fully-automatic coffee machines in the guest house, in the “old” cafeteria, and in the wine cellar. Guests can buy small snacks at the kiosk in front of the “old” cafeteria. Bread and cheese is served in the “old” cafeteria and the wine cellar every night.

### ■ Childcare

Schloss Dagstuhl offers participants who need to travel with children childcare options. A certified nanny service, which is subsidized by Schloss Dagstuhl, can be booked in advance for a small fee. Alternatively, parents can bring an accompanying person to look after their child or children and when the nanny service is not used, this person stays free of charge. In both cases, Schloss Dagstuhl covers the cost of accommodation and meals of the children. Schloss Dagstuhl has a recreation room for children. It includes a seating area, a playhouse, a play cave, and children’s room dividers equipped with a marble run, a mirror, and so on. This room gives the children a pleasant place to play indoors as well.

Dagstuhl’s offered childcare services for parents was

tet sind, versehen. Er erlaubt einen angenehmen Aufenthalt der Kinder auch im Innenraum.

Das Angebot der Kinderbetreuung wurde auch 2024 rege genutzt. Insgesamt 33 Teilnehmern von 28 Veranstaltungen mit 39 Kindern konnte eine Seminarerteilnahme durch die Kinderbetreuung ermöglicht werden. 16 Kinder davon wurden durch Schloss Dagstuhl's Kinderbetreuerin betreut.

### ■ Freizeit und Ambiente

Die Freizeitanlagen auf Schloss Dagstuhl wurden so gestaltet, dass sie auf unterschiedliche Art und Weise sowohl tagsüber als auch abends die Kommunikation zwischen den Seminarernehmern fördern. Die Mischung aus Arbeit und Freizeit in entspannter, familiärer Atmosphäre ist ein wichtiger Bestandteil des Dagstuhl-Konzepts. Gäste leben und arbeiten zusammen in einem Komplex aus drei Gebäuden, im Zentrum das historische Schloss, wo sie rund um die Uhr freien Zugang zu den zahlreichen Freizeiträumen und -anlagen haben. Musikalische Gäste können ihre Fertigkeiten im barocken Musiksaal zu Gehör bringen, wo ein Flügel und diverse andere Instrumente wie z. B. zwei Konzertgitarren zur Verfügung stehen. Unser Zentrum verfügt außerdem über eine Sauna, einen Billardtisch, Tischfußball, Mountainbikes, eine Dartscheibe, einen Freizeitraum mit Fitnessgeräten und Tischtennis sowie einen Außenbereich mit Volleyballnetz.

## Dagstuhl an der Universität des Saarlandes

8.2

Das Büro in Saarbrücken befindet sich auf dem Campus der Universität des Saarlandes im Gebäude E11. Die Räumlichkeiten werden hauptsächlich vom Sachbearbeitungsteam und von dem Teil des wissenschaftlichen Stabs genutzt, der für die Seminarbetreuung zuständig ist. Zusammen mit dem Wissenschaftlichen Direktor Raimund Seidel, der einen Lehrstuhl an der Universität des Saarlandes in Saarbrücken hat, sind so alle, die bei der Gestaltung des Wissenschaftlichen Programms eine wesentliche Rolle spielen, an einem Standort.

In den Räumen der Geschäftsstelle arbeiten aber auch ein paar Mitarbeiter des Verlagsservices Dagstuhl Publishing und der Bibliographiedatenbank dblp.

## Dagstuhl an der Universität Trier

8.3

Die meisten für die Bibliographiedatenbank dblp zuständigen Mitarbeiter haben ihren Standort an der Universität Trier. Die Ende 2010 zunächst auf Basis zweier Projekte gestartete Zusammenarbeit zwischen Schloss Dagstuhl und der Universität Trier wurde im November 2018 in eine offizielle und permanente Außenstelle von Schloss Dagstuhl auf dem Campus der Universität Trier überführt.

again well used in 2024. A total of 33 participants from 28 events with 39 children were able to attend a seminar thanks to the childcare. Of these, 16 children were cared for by Schloss Dagstuhl's nanny.

### ■ Leisure Facilities

Leisure facilities at Schloss Dagstuhl are designed to encourage and support communication among seminar participants in different settings throughout the day and evening. This work/life continuum within a relaxed, informal setting is an important part of the Dagstuhl concept. Guests live and work together in a complex of three buildings, the historical manor house ("Schloss") in the middle, and enjoy full access to the center's many unique rooms and facilities around the clock. Musically talented guests are welcome to exercise their skills in the baroque music room on the upper floor of the historical main building, which features a grand piano and various other instruments, e.g., two concert guitars. Schloss Dagstuhl also has a full sauna, a pool table, table football facilities, mountain bikes, a dartboard, and a recreation room with gym equipment and table tennis as well as outdoor sports grounds featuring a volleyball net.

## Dagstuhl at Saarland University

The Dagstuhl Office in Saarbrücken is located on the campus of Saarland University in building E11. The site mainly houses some administrative staff and the part of the scientific staff responsible for seminar support. Together with the Scientific Director, Raimund Seidel, who holds a chair at Saarland University, all those who play an essential role in the design of the scientific program are in one location.

However, some employees of Dagstuhl's publishing service and bibliography database also work at this location.

## Dagstuhl at the University of Trier

The scientific and editorial staff working on the *dblp computer science bibliography* is mainly located on the the University of Trier. Initially based on a project-based cooperation between Schloss Dagstuhl and the University of Trier which was first established in 2010, in November 2018, an official and permanent Schloss Dagstuhl branch office has been established on the campus of the University

Dabei profitiert das dblp-Team von der engen Zusammenarbeit mit der Abteilung Informatikwissenschaften und als externer Partner im Digital Research and Bibliographic Meta Data Lab des Center for Informatics Research and Technology (CIRT).

of Trier. In Trier, the dblp team benefits from the close cooperation with the University's department of computer science, and as an external partner in the Center for Informatics Research and Technology (CIRT) lab for Digital Research and Bibliographic Meta Data.

# **9** **Zentrale Dienste** *Central Services*

Schloss Dagstuhl verfügt über zwei zentrale Dienste: die IT-Abteilung und eine Forschungsbibliothek. Beide Einrichtungen befinden sich am Hauptstandort in Wadern.

Schloss Dagstuhl has two central services: the IT service and a research library, which are both located at the main site in Wadern.

## Bibliothek

9.1

## Research Library

Zur wissenschaftlichen Literatur- und Informationsversorgung der Seminarteilnehmer unterhält Schloss Dagstuhl eine hervorragende Forschungsbibliothek für Informatik.

Die Bibliothek ist für Wissenschaftler vor Ort rund um die Uhr und für externe Wissenschaftler nach Absprache zugänglich. Zur digitalen Informationsinfrastruktur gehören ein Online-Bibliothekskatalog, ein modernes Discovery-System zur Artikelrecherche sowie zahlreiche Angebote für den Online-Zugriff auf wissenschaftliche Publikationen. Ein Bildschirmarbeitsplatz sowie drei Tablets stehen in der Bibliothek zur Literaturrecherche zur Verfügung.

Für jedes Seminar wird eine individuelle Buchausstellung zusammengestellt, bestehend aus Büchern, die von Seminarteilnehmern verfasst oder herausgegeben wurden. Die anwesenden Autoren werden gleichzeitig gebeten, ihre Bücher zu signieren. Zur Optimierung der Autorenidentifikation werden die ORCID-IDs der Personennamen im Bibliothekskatalog erfasst.

Außerdem wird der Name eines jeden Seminarteilnehmers in der Online-Teilnehmerliste mit seinen oder ihren in der dblp-Literaturdatenbank erfassten Veröffentlichungen verlinkt. Diese Maßnahmen ermöglichen den Seminarteilnehmern einfachen und schnellen Zugriff auf seminarrelevante Literatur.

Die Bibliothek verfügt über einen umfangreichen Buchbestand, der Zugriff auf aktuelle Forschungspublikationen wie Konferenzbände und wissenschaftliche Zeitschriften erfolgt ausschließlich digital.

- Der Buchbestand orientiert sich am wissenschaftlichen Seminarprogramm. Bei Neuanschaffungen liegt der Fokus auf Büchern, die einen Bezug zu Dagstuhl-Seminaren oder Perspektiven-Workshops haben oder von Seminarorganisatoren oder -teilnehmern verfasst wurden. Außerdem erhält die Bibliothek zahlreiche Bücher als Spenden von Verlagen und Autoren. Aktuell verfügt die Bibliothek über etwa 38 000 Informatikbücher. Die Metadaten werden standardisiert erfasst und mit Hyperlinks angereichert, die durch persistente Adressierung (DOIs) verlässlich verlinkt sind.
- Beiträge in Konferenzbänden verkörpern den wichtigsten Teil der Literatur in der Informatik. Die Bibliothek hat die kompletten ACM- und IEEE-Proceedings elektronisch abonniert. Ältere Bände stehen teilweise auch in Druckform zur Verfügung. Die Verlagsgruppe Springer Nature spendet der Bibliothek alle Bände der Reihe Lecture Notes in Computer Science (LNCS) in Druckform. Die Bibliothek verfügt somit über Druckexemplare aller veröffentlichten Bände ab dem Band 1.
- Wissenschaftliche Fachzeitschriften sind eine wesent-

Schloss Dagstuhl maintains an excellent research library for computer science to provide seminar participants with scientific literature and information.

The library is accessible to on-site researchers around the clock and to external researchers by appointment. The digital information infrastructure includes an online library catalog, a modern discovery system for article research, as well as numerous options for online access to scientific publications. A computer and three tablets are available in the library for researching literature.

For each seminar, an individual book exhibition is compiled, consisting of books written or edited by seminar participants. The authors who are present at the seminar are asked to sign their own books. In order to optimize the author identification, the ORCID-IDs of the authors' names are recorded in the library catalog.

In addition, the name of each seminar participant will be linked in the online list of participants with their publications recorded in the dblp literature database. These measures provide seminar participants with easy and quick access to the literature relevant to the seminar.

The library maintains an extensive collection of books. Access to current research publications such as conference proceedings and scientific journals is exclusively digital.

- The book collection is oriented towards the scientific seminar program. New acquisitions focus on books which are related to Dagstuhl Seminars and Perspectives Workshops or which were written by seminar organizers or participants. In addition, the library receives numerous books as donations from publishers and authors. Currently, the library has more than 38,000 books on computer science. The metadata are recorded in a standardized way and enriched with hyperlinks, which are reliably linked by permanent addressing (DOIs).
- Contributions in conference proceedings represent the most important part of the literature in computer science. The library has subscribed to the complete ACM and IEEE proceedings electronically. Earlier volumes are also partly available in printed form. The Springer-Nature publishing group donates all printed volumes of the series Lecture Notes in Computer Science (LNCS) to the library. The library thus has print copies of all published volumes from volume 1 onwards.
- Scientific journals are essential for excellent research. Journals often publish extended versions of results that were previously published in conference proceedings. The library provides access to several thousand digital scientific journals. Most of them are included in journal packages licensed in cooperation with nationwide consortia, such as DFG-funded national and alliance

- liche Voraussetzung für exzellente Forschung. Häufig werden in Zeitschriften erweiterte Fassungen von Ergebnissen veröffentlicht, die zuvor in Konferenzbänden publiziert wurden. Die Bibliothek bietet Zugriff auf mehrere Tausend digitale Fachzeitschriften. Die meisten sind in Zeitschriftenpaketen enthalten, die in Kooperation mit deutschlandweiten Konsortien lizenziert sind, beispielsweise DFG-geförderte National- und Allianzlizenzen, Projekt DEAL, sowie von der Leibniz-Gemeinschaft geförderte Konsortiallizenzen.
- Die Bibliothek ermöglicht den benutzerfreundlichen Online-Zugriff auf über 7 500 deutschlandweite und internationale Zeitungen und Magazine aus über 120 Ländern in 60 Sprachen.

## ■ Zusammenarbeit

Schloss Dagstuhl's Forschungsbibliothek ist mit zahlreichen überregionalen Bibliotheksdatenbanken vernetzt. Der komplette Zeitschriftenbestand ist in der Zeitschriftendatenbank (ZDB) nachgewiesen. Zusätzlich ist der Bestand an elektronischen Zeitschriften in der kooperativen bundesweiten Elektronischen Zeitschriftenbibliothek (EZB) erfasst. Darüber hinaus wird der komplette Monographienbestand im K10plus, der gemeinsamen Katalogisierungsdatenbank von GBV und SWB mit über 180 Millionen Nachweisen, nachgewiesen.

Diese Datenbanken bilden die Grundlage für den deutschlandweiten und internationalen Leihverkehr der Bibliotheken. Somit steht der Zeitschriftenbestand auch standortübergreifend und überregional für Fernleihzwecke zur Verfügung.

Außerdem besteht eine enge Zusammenarbeit zwischen Schloss Dagstuhl und der Saarländischen Universitäts- und Landesbibliothek (SULB), der Campusbibliothek für Informatik und Angewandte Mathematik an der Universität des Saarlandes sowie der Bibliothek des Leibniz-Instituts für Neue Materialien (INM), die sich alle in Saarbrücken befinden.

Schloss Dagstuhl's Fachbibliothek ist institutionelles Mitglied des Deutschen Bibliotheksverbandes (DBV). Die Bibliothekarin Petra Meyer ist persönliches Mitglied im Berufsverband Information Bibliothek e.V. (BIB).

## ■ Spenden an die Bibliothek

Die Bibliothek von Schloss Dagstuhl profitiert von zahlreichen Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2024 Buchspenden vom Springer Verlag und auch von Seminarteilnehmern. Autorenexemplare werden ebenso dankbar entgegengenommen. Insgesamt erhielt das Zentrum im Berichtszeitraum 854 Bände, darunter 848 Monographien des Springer Verlags.

## IT-Service

9.2

Die IT-Abteilung bietet umfassenden Support für alle internen Vorgänge an den drei Standorten. Darüber hinaus betreut sie die IT-Infrastruktur und -Dienste und bietet

licenses, as well as consortium licenses funded by the Leibniz Association, and Projekt DEAL.

- The library enables user-friendly online access to over 7,500 Germany-wide and international newspapers and magazines from over 120 countries in 60 languages.

## ■ Collaboration

Schloss Dagstuhl's research library is connected to numerous national library databases. The complete journal inventory is recorded in the Zeitschriftendatenbank (ZDB). In addition, the inventory of electronic journals is recorded in the cooperative nationwide Electronic Journals Library (Elektronische Zeitschriftenbibliothek, EZB). Furthermore, the complete Monograph collection is recorded in K10plus, the joint cataloging database of GBV and SWB with over 180 million records.

These databases form the foundation for the libraries' nationwide and international lending system. Thus the journal collections are also available for inter-library loan purposes across locations and regions.

There is also a close cooperation between Schloss Dagstuhl and the Saarland University and State Library (SULB), the Campus Library for Computer Science and Applied Mathematics at Saarland University, and the library of the Leibniz Institute for New Materials (INM), all of which are located in Saarbrücken.

Schloss Dagstuhl's specialized library is an institutional member of the German Library Association (Deutscher Bibliotheksverband, DBV). The librarian, Petra Meyer, is a personal member of the Professional Association Information and Libraries (Berufsverband Information Bibliothek e.V., BIB).

## ■ Library Donations

The Dagstuhl Informatics Research Library benefits from numerous book donations. In 2024, the Informatics Research Library received book donations from Springer Verlag and also from seminar participants. The center is also grateful for donations of author's copies. The center received a total of 854 volumes during the reporting period, including 848 monographs from Springer Verlag.

## IT Service

The IT service provides comprehensive support for all internal operations at all three sites. Moreover, it provides IT infrastructure, services, and support for all guests of

Unterstützung für alle Gäste bei Dagstuhl-Veranstaltungen.

Der IT-Service umfasst u.a.:

- Internetzugang über Ethernet und WLAN in allen Räumen. Für den WLAN-Zugang bietet Schloss Dagstuhl persönliche Accounts an und ist auch an der *eduroam*-Initiative beteiligt<sup>48</sup> (eine praktische Alternative für Gäste, die bereits einen *eduroam*-Account haben). Innerhalb sämtlicher Einrichtungen stellt Schloss Dagstuhl ein weitläufiges Netzwerk von Zugangspunkten zum Drahtlosnetzwerk zur Verfügung, das aktiv überwacht und regelmäßig erweitert wird. Die Verbindung zum (externen) Internet wird durch zwei redundante 375 Mbit/s-Leitungen sichergestellt, betrieben durch den DFN e.V. (Deutsches Forschungsnetz).
- Fahrbare ebenso wie fest montierte Präsentationsmöglichkeiten in den Tagungsräumen, teilweise mit der Option der drahtlosen Videoübertragung von den Computern der Gäste.
- Ein Terminal vor der Rezeption, das den Check-in und das Schreiben von Ersatzzugangskarten zu den Gästezimmern erlaubt, auch wenn die Rezeption gerade unbesetzt ist.
- Ein Terminal vor der Cafeteria, das den schnellen Ausdruck von Tickets erlaubt. Dieses Terminal kann im Bedarfsfall auch für das Schreiben von Zimmerzugangskarten benutzt werden.
- Zugang zu einem Multifunktionsgerät, das Kopierer, Scanner und Drucker in einem vereint.
- Technischen Support für Seminarteilnehmer und Mitarbeiter von Schloss Dagstuhl.

Der IT-Service verwaltet Server für alle Abteilungen, u.a., die drei zentralen Webserver:

- <https://www.dagstuhl.de>, der allgemeine Informationen über das Zentrum mit seinen drei Abteilungen sowie Informationen für Teilnehmer des Seminarprogramms bereitstellt,
- <https://drops.dagstuhl.de>, Dagstuhls Publikationsplattform DROPS, und
- <https://dblp.dagstuhl.de> (auch <https://dblp.org>), die dblp computer science bibliography

Darüber hinaus stellt der IT-Service für interne Arbeitsabläufe Tools für eine kollaborative Arbeitsumgebung bereit und wartet diese.

Dagstuhl events.

This service includes – among others – the following:

- Internet access via Ethernet and Wi-Fi throughout all rooms. For Wi-Fi access, Schloss Dagstuhl offers personal accounts and also takes part in the *eduroam* service<sup>48</sup> (which is a comfortable option for guests with existing *eduroam* accounts). Within its facilities, Schloss Dagstuhl provides a vast network of professional-grade wireless network access points that is actively monitored and extended regularly. External internet access for Schloss Dagstuhl is provided through two redundant 375 Mbit/s connections that are managed by DFN e.V. (National Science Network).
- Mobile and stationary presentation facilities in meeting rooms, partially including the option to have wireless presentations from the guest computers.
- One terminal in front of the reception, allowing for self check-in and the writing of replacement room cards, even if there is no one at the reception desk.
- A terminal in front of the “old” cafeteria that can be used to print tickets. This terminal can also be used to write room cards.
- Access to a multifunction device that combines a copier, scanner, and printer in one.
- Technical support for both seminar participants and Dagstuhl staff.

The IT service manages servers for Schloss Dagstuhl’s divisions, including the three main web-severs:

- <https://www.dagstuhl.de>, providing general information of the center with its three divisions as well as information for participants of our seminar program,
- <https://drops.dagstuhl.de>, Schloss Dagstuhl’s publishing platform DROPS, and
- <https://dblp.dagstuhl.de> (also <https://dblp.org>), the dblp computer science bibliography.

Furthermore, for internal work procedures, the IT service provides and maintains tools for a collaborative work environment.

<sup>48</sup> *eduroam* (education roaming) ist ein weltweiter roaming Zugangsdienst der für die internationale Forschungs- und Bildungsgemeinde entwickelt wurde, siehe <https://www.eduroam.org>.  
*eduroam* (education roaming) is a world-wide roaming access service developed for the international research and education community, see <https://www.eduroam.org>.

# 10 Kunst *Art*

## Dagstuhl als Galerie

10.1

## Dagstuhl as an Art Gallery

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des Flurs sowie die hervorragende Ausleuchtung mit starken Kontrasten zwischen Tag und Nacht bieten den Künstlern sehr gute Möglichkeiten, ihre Werke darzustellen. Die Kunstwerke an den Wänden des schmalen Gangs durchbrechen die Nüchternheit des Neubaus in anregender und angenehmer Weise. Die wechselnden Ausstellungen bieten einen erfrischenden und dynamischen Kontrast zu der ständigen Kunstsammlung von Schloss Dagstuhl.

Prof. Reinhard Wilhelm, ehemaliger wissenschaftlicher Direktor des Zentrums, fungiert nach seinem Eintritt in den Ruhestand im April 2014 weiterhin als Betreuer der Ausstellungsaktivitäten von Schloss Dagstuhl. Das Zentrum veranstaltet jährlich etwa drei bis vier Kunstausstellungen für jeweils zwei bis drei Monate.

Während lange Jahre einzelne Künstler und Sammler ihre Werke im Zentrum ausstellten, wurden zwischen 2016 und 2018 im Rahmen eines Kooperationsvertrages zwischen der Saarland-Sporttoto GmbH (kurz Saartoto), der Hochschule der Bildenden Künste Saar (kurz HBKsaar) und Schloss Dagstuhl drei Ausstellungen organisiert. Als bedeutender Förderer von Künstlern besitzt Saartoto einen großen Bestand an Kunstwerken..

Im August 2022 wurde für die Dauer von 2 Jahren, die Zusammenarbeit zwischen Schloss Dagstuhl und dem Institut für aktuelle Kunst im Saarland, vertreten durch den Institutsleiter Dr. Andreas Bayer, beschlossen, ruhte jedoch ab 2023 aus personellen Gründen. Die Zusammenarbeit beinhaltet die kuratorische Ausgestaltung von mindestens zwei Ausstellungen/Jahr in Schloss Dagstuhl. Dankenswerterweise unterstützt der *Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.* die Zusammenarbeit.

In 2024 fanden die vier folgenden von Reinhard Wilhelm vermittelten Ausstellungen (siehe Fig. 10.1) statt. Die jeweils aktuellen Ausstellungen sind nach Anmeldung auch für die interessierte Öffentlichkeit zugänglich. Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite<sup>49</sup> von Dagstuhl.

Art exhibitions are regularly organized in the so-called cloister of the new building. The spacious surroundings, excellent lighting, and dramatic day-to-night contrast offer artists a unique exhibition space. Arranged along the corridor walls, the artworks offset the otherwise ascetic nature of the new building. These temporary exhibits offer a fresh and dynamic counterpoint to the center's permanent collection, which can be found scattered throughout the three buildings.

Prof. Reinhard Wilhelm continue to supervise the Schloss Dagstuhl art exhibitions following his retirement as the center's Scientific Director in April 2014. The center holds approximately three to four art exhibits per year, with each exhibit generally running for two to three months.

For many years Dagstuhl's exhibitions were organized by artists and individual collectors. In the period from 2016 to 2018 three exhibitions had been organized in the course of a cooperation between between Saarland-Sporttoto GmbH (Saartoto for short), Hochschule für Bildende Künste Saar (university of art and design; HBKsaar for short), and Schloss Dagstuhl. Being a major art sponsor, Saartoto is in possession of a substantial art collection.

In August 2022, the cooperation between Schloss Dagstuhl and the Institut für aktuelle Kunst im Saarland (institute for contemporary art), represented by the Institute's director Dr. Andreas Bayer, was agreed for a period of two years. The cooperation includes the curatorial organization of at least two exhibitions/year at Schloss Dagstuhl, but was paused since 2023 due to staffing reasons. The association *Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.* ("Friends of Dagstuhl" for short) gratefully supports the cooperation.

The following four exhibitions conveyed by Reinhard Wilhelm (cf. Fig. 10.1) took place in 2024. Current exhibitions are also open to the interested public by appointment. Information and current news can be found on Dagstuhl's art art webpage<sup>49</sup>.

<sup>49</sup> <https://www.dagstuhl.de/de/guests/art>

»... inzwischen«	Works from Shako Berekashvili   January 15 to March 3, 2024
»inside – outside«	Works from Tina Stein   April 23 to July 26, 2024
»Marx und Engel   Atmen«	Works Annette Marx and Andreas Engel   August 26 to October 4, 2024
»Sehnsucht nach dem Norden«	Works from Ulrich Klimmt   October 21 to December 19, 2024

Fig. 10.1  
Art exhibition in 2024.

### ■ »... inzwischen«

Werke des georgischen Künstlers Shako Berekashvili waren in der ersten Ausstellung des Jahres zu sehen. Dank Spenden konnte das Bild „Flügel“ für Dagstuhl's permanente Kunstaussstellung erworben werden.

### ■ »inside – outside«

In der Ausstellung wurden Bilder der saarländischen Künstlerin Tina Stein gezeigt. Dank Spenden konnte die permanente Kunstaussstellung um das Bild „Meditations“ ergänzt werden.

### ■ »Marx und Engel | Atmen«

In der Ausstellung wurden Bilder der Künstlerin Annette Marx und des Künstlers Andreas Engel gezeigt.

Marx und Engel arbeiten sehr unterschiedlich, bedienen sich allerdings ähnlicher Mittel. Während Engel die klassische Collage, also das Zerschneiden von Fotos und Zeitungsbildern zu neuen Bildern entwickelt, überwiegt bei Marx die Farbe, und doch collagiert sie. Einzelne Stoff fetzen sind auf fast jedem Gemälde von ihr angeordnet und lenken den Blick. So gegensätzlich Werkgruppen erscheinen, so vereinen sie sich in ihrer düsteren Wiedergabe des Weltgeschehens, aber jedes Bild atmet eine elementare, prinzipielle Zuversicht.

### ■ »Sehnsucht nach dem Norden«

In Kooperation mit dem Kulturverein Burbach e. V. (Saarbrücken) wurden Bilder des 2022 verstorbenen Künstlers Ulrich Klimmt ausgestellt. Dank großzügiger Spenden konnte das Werk „Aquarell Nr. 33“ von Schloss Dagstuhl angekauft werden. Der Erwerb eines weiteren Bildes ist angedacht.

## Dagstuhl's permanente Kunstaussstellung

10.2

Die von Gästen immer wieder positiv hervorgehobene Kunstsammlung geht auf den Gründungsdirektor Professor Wilhelm zurück. Seine Idee war es, den 1995 neueröffneten Speisesaal und den etwa ein Jahr älteren Neubau durch Kunstwerke zu beleben. Dazu startete er die oben beschriebenen Kunstaussstellungen. Unter Mitwirkung der Künstler wird aus ausgewählten Ausstellungen ein Werk ausgewählt, für das dann Spender gesucht werden. In den letzten 25 Jahren kamen so ungefähr 180 Kunstwerke zusammen. Auch durch diese Initiative angeregt und verstärkt erhielt Dagstuhl in den vergangenen Jahren weitere Spenden von Künstlern und Mäzenen. Schloss Dagstuhl dankt den 5 Spendern und dem Förderverein „Freunde von Schloss Dagstuhl“, die den Kauf der Werke erst ermöglicht haben. Die meisten Arbeiten kommen in den Räumen des Zentrums in Wadern, einige auch in der Geschäftsstelle in Saarbrücken, sehr gut zur Geltung.

### ■ »... inzwischen«

Paintings by the Georgian artist Shako Berekashvili were displayed in the first exhibition of the year. Thanks to donations, the painting “Flügel” was acquired for Dagstuhl's permanent art exhibition.

### ■ »inside – outside«

The exhibition showed paintings by the Saarland artist Tina Stein. Thanks to donations the painting “Meditations” was added to Dagstuhl's permanent art collection.

### ■ »Marx und Engel | Atmen«

The exhibition showed pictures by the artists Annette Marx and Andreas Engel.

Marx and Engel work very differently, but use similar means. While Engel uses the classic collage, i.e. cutting up photos and newspaper images to create new pictures, Marx's work is dominated by colour, and yet she also uses collage. Individual shreds of fabric can be found in almost all of her paintings, drawing the viewer's eye. As contradictory as the groups of works appear, they are united in their gloomy depiction of world events, but each painting breathes an elementary, fundamental confidence.

### ■ »Sehnsucht nach dem Norden«

In cooperation with the Kulturverein Burbach e. V. (Saarbrücken), paintings by the artist Ulrich Klimmt, who died in 2022, were exhibited. Thanks to generous donations, the work “Aquarell Nr. 33” could be purchased by Schloss Dagstuhl. The purchase of a further painting is being considered.

## Dagstuhl's Permanent Art Exhibition

The art collection, continually praised by guests, was initiated by Founding Director Professor Wilhelm. It was his idea to use works of art in order to enliven the new building as well as the dining room opened in 1994 and 1995, respectively. To this end, Professor Wilhelm launched the exhibitions described above. Assisted by the artists, an artwork from selected exhibitions is chosen and donors are drummed up. Thus, approximately 180 works of art have been acquired over the last 25 years. Additionally, this initiative has increasingly encouraged artists and patrons to make donations. Schloss Dagstuhl would like to thank the five donors and the “Friends of Schloss Dagstuhl” association, without whom the purchase of the works would not have been possible. Most works of art are shown in the center's rooms in Wadern, some are placed in the office in Saarbrücken.



# 11

## **Struktur der Gesellschaft** *Structure of the Company*

## Gründung und Gesellschafter

11.1

## Formation and Shareholders

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit elf Gesellschaftern (siehe Fig. 11.1) organisiert. Dies sind einerseits die vier Gesellschafter, die Schloss Dagstuhl gegründet haben, nämlich die Gesellschaft für Informatik e. V. (GI), die Universität des Saarlandes, die Technische Universität Kaiserslautern und das Karlsruher Institut für Technologie (KIT). Als vier weitere Gesellschafter wurden 1994 die Technische Universität Darmstadt, die Johann Wolfgang Goethe-Universität Frankfurt am Main, die Universität Stuttgart und die Universität Trier aufgenommen. Drei international renommierte Forschungsinstitute, das Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), das Centrum Wiskunde & Informatica (CWI, Niederlande) und die Max-Planck-Gesellschaft (MPG, Deutschland) wurden 2005/2006 als weitere Gesellschafter aufgenommen.

Aufgrund eines Beschlusses der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (heute Gemeinsame Wissenschaftskonferenz) wurde das Zentrum mit Wirkung zum 1. Januar 2006 als Serviceeinrichtung für die Forschung in die gemeinsame Forschungsförderung von Bund und Ländern aufgenommen. Es ist seit 2005 Mitglied der Leibniz-Gemeinschaft. Entsprechend wurde 2008 der Name des Zentrums von vormals „Internationales Begegnungs- und Forschungszentrum für Informatik“ in „Schloss Dagstuhl – Leibniz-Zentrum für Informatik“ geändert.

Schloss Dagstuhl wurde im Juli 2009 erstmals durch die Leibniz-Gemeinschaft evaluiert. Die Stellungnahme der Evaluierungs-Kommission vom März 2010 war sehr positiv: Schloss Dagstuhl widme sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen. Schloss Dagstuhl wurde wie auch schon 2016 im Jahr 2023 erneut mit hervorragendem Ergebnis evaluiert. In der Stellungnahme des Senats der Leibniz-Gemeinschaft aus dem Jahre 2024 wurde das Veranstaltungsprogramm und die Beteiligung an der Literaturdatenbank dblp als „exzellent“ bewertet, während der Bereich Open Access (Publishing) als „sehr gut“ bewertet wurde.

Schloss Dagstuhl is operated as a non-profit organization by eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e. V. (GI),<sup>50</sup> the Universität des Saarlandes, the Technische Universität Kaiserslautern, and the Karlsruher Institut für Technologie (KIT). In 1994, the organization was extended to include four new associates: the Technische Universität Darmstadt, the Johann Wolfgang Goethe-Universität Frankfurt am Main, the Universität Stuttgart, and the Universität Trier. Finally, in 2005 and 2006, three internationally renowned research institutes joined the association: the Institut National de Recherche en Informatique et en Automatique (INRIA, France), the Centrum Wiskunde & Informatica (CWI, Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung<sup>51</sup> (today Joint Science Conference) the center has been classified as a research service institution for joint funding by the German federal and state governments since January 2006. Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association and changed its name accordingly from “Internationales Begegnungs- und Forschungszentrum für Informatik”<sup>52</sup> to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”<sup>53</sup> in 2008.

In July 2009, Schloss Dagstuhl was evaluated for the first time by the Leibniz Association. The March 2010 findings of the evaluation commission were very positive, and established that the center has shown outstanding commitment to its designated task of supporting the international computer science research community by providing a seminar center for academic events. In 2023, Schloss Dagstuhl was evaluated again, with excellent results, as it was in 2026. In the Leibniz Association Senate report from 2024, the seminar program and the cooperation with the computer science bibliography dblp were both rated as “excellent”, whereas the Open Access Publishing was rated “very good.”

## Organe der Gesellschaft

11.2

## Dagstuhl Bodies

Die drei Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, die stellvertretend für die Gesellschaft als juristische Person handeln, sind die folgenden:

- Gesellschafterversammlung
- Aufsichtsrat
- Geschäftsführung

Details zu den Organen sind den folgenden Abschnitten zu

The three bodies of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, which act for the company as a legal entity, are the following:

- Shareholders' Meeting
- Supervisory Board
- Management

Detailed information is given in the sections below.

<sup>50</sup> engl.: German Informatics Society

<sup>51</sup> engl.: Federal/State Government Commission for Educational Planning and Research Promotion

<sup>52</sup> engl.: International Conference and Research Center for Computer Science

<sup>53</sup> engl.: Schloss Dagstuhl – Leibniz Center for Informatics

entnehmen.

## ■ Die Gesellschafterversammlung

Die Gesellschafter beschließen über alle Änderungen an der Gesellschaft, insbesondere über die Aufnahme weiterer Gesellschafter, über die Änderung des Gesellschaftsvertrags und über ihre Auflösung. Die Gesellschafter bestätigen unter anderem auch die von Gesellschaftern neu entsandten Mitglieder in den Aufsichtsrat sowie die Berufung und Abberufung der Geschäftsführer. Derzeit haben anteilig nach der Höhe der Geschäftsanteile alle Gesellschafter die gleiche Anzahl von Stimmen, außer der Gesellschaft für Informatik, die die dreifache Anzahl besitzt. Beschlüsse werden entweder in der mindestens einmal jährlichen stattfindenden Gesellschafterversammlung gefasst oder durch schriftliche Stimmabgabe.

## ■ Der Aufsichtsrat

Der Aufsichtsrat ist verantwortlich dafür, dass die Geschäftsführung die Ziele der Gesellschaft rechtmäßig, zweckmäßig und wirtschaftlich sinnvoll erfüllt. Er wirkt in allen wesentlichen Angelegenheiten der Gesellschaft betreffend Forschung und Finanzplanung mit.

Die 12 Mitglieder des Aufsichtsrats (siehe Fig. 11.2) setzen sich aus vier Repräsentanten der Gesellschaft für Informatik, je einem Vertreter der drei Gründungsuniversitäten, zwei Vertretern der später hinzugekommenen vier Universitäten und je einem Vertreter des Bundes und der beiden Bundesländer Saarland und Rheinland-Pfalz, in denen Schloss Dagstuhl formal seinen Sitz hat, zusammen. Die reguläre Amtszeit der Aufsichtsratsmitglieder beträgt mindestens vier volle, abgeschlossene Geschäftsjahre und endet mit der Entlastung für das vierte Geschäftsjahr. Die Vertreter der Universitäten in Darmstadt und Stuttgart wechseln im Allgemeinen Amtszeit für Amtszeit mit denen der Universitäten in Frankfurt und Trier ab.

Der Aufsichtsrat entscheidet über die Berufung und Abberufung der Geschäftsführer sowie der Mitglieder des Wissenschaftlichen Direktoriums, des Wissenschaftlichen Beirates und des Kuratoriums. Alle Beschlüsse, die die Finanzen oder das Vermögen der Firma betreffen, benötigen seine Zustimmung. Beschlüsse von forschungspolitischer Bedeutung und Beschlüsse mit erheblichen finanziellen Auswirkungen können nicht gegen die Stimmen der Vertreter des Bundes und der beiden Sitzländer gefasst werden. Der Aufsichtsrat entscheidet zudem über die Erteilung einer Prokura.

## ■ Die Geschäftsführung

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH hat zwei Geschäftsführer (siehe Fig. 11.3), die gemeinsam die Gesellschaft vertreten. Die Geschäftsführung besteht aus dem *Wissenschaftlichen Direktor* und dem *Technisch-administrativen Geschäftsführer*.

Der Wissenschaftliche Direktor ist verantwortlich für die wissenschaftlich-fachliche Zielsetzung und die Programmgestaltung, und ist zudem Mitglied und Vorsitzender des Wissenschaftlichen Direktoriums. Seit Mai 2014 ist Prof. Raimund Seidel, Ph.D., der wissenschaftliche Direk-

## ■ Shareholders' Meeting

All changes to the company, in particular the inclusion of new associates, the revision of the Shareholders' agreement, and the dissolution of the company, are decided by the shareholders. Shareholders also confirm new members forwarded by them to the Supervisory Board and the appointment or recall of the managing directors. In accordance with their shares, all shareholders currently have the same number of votes except the Gesellschaft für Informatik, which has three times the number of votes of the other shareholders in proportion to its larger number of shares. Decisions are made in shareholders' meetings which take place at least once a year, or via a written vote.

## ■ Supervisory Board

The Supervisory Board is responsible for ensuring that the management complies with the center's objectives in a legally and economically meaningful manner. The board is involved in all essential matters with regard to research and financial planning.

The 12-member board (see Fig. 11.2) is composed of four representatives of the Gesellschaft für Informatik, one representative from each of the three founding universities, two representatives of the four universities that subsequently joined, and one representative from each of the German federal government and the two host state governments of Saarland and Rhineland-Palatinate. The Supervisory Board members typically hold office for at least four full fiscal years. The term of office ends with the approval for the fourth fiscal year. In general, representatives of the universities in Darmstadt and Stuttgart and of the universities in Frankfurt and Trier rotate after each term of office.

The Supervisory Board formally appoints and recalls the managing directors and members of the Scientific Directorate, Scientific Advisory Board, and Industrial Curatory Board. Furthermore, all decisions regarding financial issues and company assets must be approved by the Supervisory Board. Consent cannot be given against the votes of the represented (federal) state governments if the matter affects political issues in the area of science or has considerable financial weight. The Supervisory Board also holds decision power with respect to the granting of power of attorney.

## ■ Management

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH has two managing directors (see Fig. 11.3) who jointly represent the company. These are the *Scientific Director* and the *Technical Administrative Director*.

The Scientific Director is in charge of drafting the company's scientific goals and program planning, and is also a member and the chairperson of the Scientific Directorate. Since May 2014, Prof. Raimund Seidel, Ph.D., is the Scientific Director of Schloss Dagstuhl.

tor von Schloss Dagstuhl.

Der Wissenschaftliche Direktor wird dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Dieser Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Seit Juli 2014 hat Heike Meißner diese Position inne.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, Heike Meißner has been holding this position.

## Gremien der Gesellschaft

11.3

## Dagstuhl Boards

Die Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH werden durch drei Gremien unterstützt. Es sind die folgenden:

- Wissenschaftliches Direktorium
- Wissenschaftlicher Beirat
- Kuratorium

Details zu den Gremien werden in den folgenden Abschnitten ausgeführt.

The bodies of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH are supported by the following bodies:

- Scientific Directorate
- Scientific Advisory Board
- Industrial Curatory Board

Detailed information about these boards can be found in the sections below.

### ■ Der Wissenschaftliche Beirat

Die Aufgaben des Wissenschaftlichen Beirats (siehe Fig. 11.4) werden nicht nur durch den Gesellschaftsvertrag festgelegt, sondern auch durch die Empfehlungen der Leibniz-Gemeinschaft. Im Sinne dieser wirkt der Wissenschaftliche Beirat auf zwei Wegen bei der Qualitätssicherung mit. Zum einen berät er die Leitung in Fragen der Forschungs- und Entwicklungsplanung, nimmt Stellung zu den Programmbudgets und gibt Empfehlungen zum Ressourceneinsatz. Er unterstützt weiterhin den Aufsichtsrat bei wichtigen Entscheidungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung (SAE) der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat sollte aus sechs bis zwölf international angesehenen, im Berufsleben stehenden Wissenschaftlern aus dem In- und Ausland bestehen. Die Amtszeit der Mitglieder beträgt vier Jahre und beginnt am 1. Januar des auf ihrer Berufung folgenden Jahres und endet vier Jahre später am 31. Dezember. Eine einmalige Wiederberufung ist möglich. Der Beirat wählt aus seiner Mitte einen Vorsitzenden. Der Wissenschaftliche Beirat tagt einmal im Jahr. Mitglieder des Beirats werden vom Aufsichtsrat auf Vorschlag des Beirats ernannt.

### ■ Das Kuratorium

Das Kuratorium (siehe Fig. 11.5) erfüllt eine Transmissionsfunktion zwischen Schloss Dagstuhl und den Forschungsabteilungen und Entwicklungslaboren der Indus-

### ■ Scientific Advisory Board

The tasks of the Scientific Advisory Board (see Fig. 11.4) are not only defined by the Shareholders' Agreement, but also by the recommendations of the Leibniz Association. The latter stipulates two different ways in which the Scientific Advisory Board is involved in quality assurance. On the one hand, the board offers advice to the management with regard to research as well as development planning and issues comments on the program budget draft, making recommendations on the use of resources. It also assists the Supervisory Board in the making of important decisions with regard to the future development of the institute as well as the acquisition of management staff. On the other hand, it carries out an audit of the entire institute between two evaluations by the Senatsausschuss Evaluierung (SAE, Senate Committee Evaluation) of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board, and the SAE.

The Scientific Advisory Board should consist of six to twelve internationally reputable, well established scientists and academics from Germany and abroad. The term of office for members is four years and can be prolonged once. It begins on January 1 of the year after the appointment and ends four years later on December 31. The Scientific Advisory Board members elect a chairperson from their midst. The board convenes once a year. Members are appointed by the Supervisory Board in accordance with the suggestions of the Scientific Advisory Board.

### ■ Industrial Curatory Board

The Industrial Curatory Board (see Fig. 11.5) performs a transmissional function between the center and the industrial R&D departments and laboratories. Its role

trie. Es hat die Aufgabe, die Akzeptanz des Zentrums in Verwaltung, Industrie und Wirtschaft abzusichern und als Förderungsorganisation die wirtschaftliche Basis des Zentrums zu verbreitern. Mitglieder des Kuratoriums werden vom Aufsichtsrat ernannt.

Nach seiner Geschäftsordnung hat das Kuratorium mindestens fünf Mitglieder, deren Amtszeit vier Jahre beträgt. Sie beginnt am 1. Januar des auf ihrer Berufung folgenden Jahres und endet vier Jahre später am 31. Dezember. Eine einmalige Wiederberufung ist möglich. Die Mitglieder des Kuratoriums unterstützen das Zentrum dabei, aktuelle Themen zu identifizieren und dazu geeignete zugkräftige Organisatoren aus der Industrie zu gewinnen. Sie werden ebenso gebeten, geeignete Personen aus der Industrie als Teilnehmer von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops zu benennen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

### ■ Das Wissenschaftliche Direktorium

Im im Gesellschaftsvertrag ist noch das Wissenschaftliche Direktorium als „sonstige Einrichtung“ aufgeführt, da dieses ursprünglich das einzige Beratungsgremium von Schloss Dagstuhl war. Mit der Aufnahme in die Leibniz-Gemeinschaft wurden dann als abteilungsübergreifende Beratungsgremien der Wissenschaftliche Beirat (siehe Abschnitt 11.3) installiert. Details zum Wissenschaftlichen Direktorium, als abteilungsspezifisches Gremium des Seminarwesens, finden sich in Kapitel 2.9.

is to secure acceptance of Schloss Dagstuhl within the business, industry and administrative communities, and as a promotional organization to broaden the economic basis of the center. Board members are appointed by the Supervisory Board.

According to its rules of procedure, the Industrial Curatory Board consists of at least five members whose term of office is four years. It begins on January 1 of the year after the appointment and ends four years later on December 31. A one-off reappointment for a second term is possible. The board members help the center to identify current R&D topics for seminars and locate attractive organizers in industry. The Industrial Curatory Board is regularly called upon to propose suitable participants for Dagstuhl Seminars and Dagstuhl Perspectives Workshops known to it from its activities. It convenes once a year together with the Scientific Advisory Board.

### ■ Scientific Directorate

In the company agreement, the Scientific Directorate is listed as an “other entity” because it was originally the only advisory body at Schloss Dagstuhl. When the centre was accepted into the Leibniz Association, the Scientific Advisory Board (see Section 11.3) was then installed as a cross-departmental advisory body. Details of the Scientific Directorate, as a department-specific body for the seminar department, can be found in chapter 2.9.

Gesellschafter   Associates
Centrum Wiskunde & Informatica (CWI), The Netherlands
Gesellschaft für Informatik e. V., Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Karlsruher Institut für Technologie (KIT), Germany
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Berlin, Germany
Technische Universität Darmstadt, Germany
Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Germany
Universität des Saarlandes, Germany
Universität Stuttgart, Germany
Universität Trier, Germany

Fig. 11.1

Associates.

Aufsichtsrat   Supervisory Board
Dr. Christian Berghoff Bundesministerium für Bildung und Forschung, Bonn, Germany   Representative of the German federal government   <i>tenure started in May 2024</i>
Prof. Dr. Dominik Brodowski Universität des Saarlandes, Saarbrücken, Germany   Representative of Universität des Saarlandes   <i>tenure started in May 2024</i>
Dr. Marc Brüser Ministerium für Wissenschaft und Gesundheit, Mainz, Germany   Representative of Rhineland-Palatinate
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany   Representative of Karlsruher Institut für Technologie
Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen Technische Universität Berlin, Germany   Representative of Gesellschaft für Informatik e. V.   Chairman of the Supervisory Board
Prof. Dr. Volker Lindenstruth Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany   Representative of Johann Wolfgang Goethe-Universität Frankfurt am Main   <i>tenure ended in May 2024</i>
Christian Mees Ministerium für Wirtschaft, Innovation, Digitales und Energie, Saarbrücken, Germany   Representative of Saarland
Prof. Dr. Arnd Poetzsch-Heffter Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, Germany   Representative of Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau
Christine Regitz SAP SE   Representative of Gesellschaft für Informatik e. V.
Dr. Simone Rehm Universität Stuttgart, Germany   Representative of Universität Stuttgart.   <i>tenure started in May 2024</i>
Prof. Dr. Ralph Schenkel Universität Trier, Germany   Representative of Universität Trier   <i>tenure ended in May 2024</i>
Prof. Dr. Manfred J. Schmitt Universität des Saarlandes, Saarbrücken, Germany   Representative of Universität des Saarlandes   <i>tenure ended in May 2024</i>
Prof. em. Dr.-Ing. Dr.-Ing. h. c. Roland Vollmar Karlsruher Institut für Technologie, Germany   Representative of Gesellschaft für Informatik e. V.   <i>tenure ended in October 2024</i>
Cornelia Winter Gesellschaft für Informatik e. V., Bonn, Germany   Representative of Gesellschaft für Informatik e. V
Prof. Dr. Felix Wolf Universität Darmstadt, Germany   Representative of Universität Darmstadt   <i>tenure started in May 2024</i>

Fig. 11.2

Supervisory Board members.

**Geschäftsführung | Management**

Heike Meißner (Technisch-administrative Geschäftsführerin | Technical Administrative Director)  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern, Germany

Prof. Raimund Seidel, Ph. D. (Wissenschaftlicher Direktor | Scientific Director)  
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern and Universität des Saarlandes, Saarbrücken, Germany

Fig. 11.3

**Management.****Wissenschaftlicher Beirat | Scientific Advisory Board**

Prof. Dr. Anja Feldmann  
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Prof. Dr. h.c. Carole Goble, CBE  
University of Manchester, Great Britain | *tenure will start in January 2025*

Prof. Dr. Ir. Dr. h. c. (AAU) Joost-Pieter Katoen (PDEng)  
RWTH Aachen, Germany

Prof. Dr. Laurent Romary  
Centre Marc Bloch, Berlin, Germany and INRIA, Paris, France

Prof. Alistair Sinclair, Ph. D.  
University of California, Berkeley, United States of America

Prof. Dr. Lothar Thiele  
ETH Zürich, Switzerland | Chair of the Scientific Advisory Board

Fig. 11.4

**Scientific Advisory Board.****Kuratorium | Industrial Curatory Board**

Dr. Tim Harris  
Microsoft Research, Cambridge, United Kingdom

Dr. Alexander Keller  
NVIDIA, Berlin, Germany

Dr. Jaroslaw Kutylowski  
DeepL, Köln, Germany

Dr.-Ing. Christof Leng  
Google, Darmstadt, Germany

Dr.-Ing. Michael Perscheid  
SAP Innovation Center, Potsdam, Germany

Fig. 11.5

**Industrial Curatory Board.**



**12**

**Förderverein „Freunde von  
Dagstuhl“**

***Association “Friends of Dagstuhl”***

## ■ Förderverein „Freunde von Dagstuhl“

Markus Bläser (Universität des Saarlandes, Germany)  
 Marc Herbstritt (Herbolzheim, Germany)  
 Erich Reindel (Universität des Saarlandes, Germany)

Seit Mitte 2014 gibt es den „Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.“. Der sehr technische und holprig klingende Name spiegelt dabei exakt den Vereinszweck wider: die Förderung von Wissenschaft und Forschung im Leibniz-Zentrum für Informatik in Schloss Dagstuhl. Für die Webpräsenz wurde allerdings ein wesentlich geschmeidigerer Name gewählt: „Friends of Dagstuhl“ (<https://www.friends-of-dagstuhl.de>).

Der Verein ist darauf ausgerichtet, finanzielle Mittel zur erfolgreichen Umsetzung des Vereinszwecks zu beschaffen und bereitzustellen. Außerdem verwaltet der Verein treuhänderisch das Stiftungsvermögen der nicht rechtsfähigen „Stiftung Informatikzentrum Schloss Dagstuhl“ und vertritt diese im Rechts- und Geschäftsverkehr. Das Stiftungsvermögen ist langfristig angelegt. Die strategische Aufsicht über die Stiftung obliegt einem Stiftungsrat (siehe Fig. 12.1). Der Verein selbst wird von einem Vorstand (siehe Fig. 12.2 und Fig. 12.3) geleitet.

Aus Mitteln einer großzügigen Zuwendung aus dem Nachlass eines verstorbenen Freundes von Schloss Dagstuhl finanzierte der Verein in den vergangenen Jahren einen Shuttle-Service an den An- und Abreisetagen. Dadurch wurde die Anbindung von Schloss Dagstuhl an den Bahnhof Türkismühle deutlich erleichtert. Außerdem engagierte sich der Verein auch 2024 beim Erwerb und bei der Inventarisierung von Kunstobjekten im Schloss und in seinem Umfeld. Die Unterstützung des Vereins belief sich für diese beiden Maßnahmen auf insg. rd. 10 000 €. Aktuell befasst sich der Verein mit den Möglichkeiten zur Stiftung eines Dagstuhl-Awards, durch den in Zukunft herausragende Leistungen in der Informatik oder auch herausragende Dagstuhl-Seminare ausgezeichnet werden könnten.

Nachdem sich Ende 2023 der langjährige Vorsitzende Professor Holger Hermanns und die langjährige stellvertretende Vorsitzende Angelika Müller von Brochowski aus der Vorstandsarbeit zurückgezogen haben, wurden im Jahr 2024 alle notariellen und banktechnischen Schritte durchgeführt, so dass der neue Vorstand inzwischen voll handlungsfähig ist.

Weitere Informationen zum Verein, aber auch Mitgliedschaftsanträge finden Sie unter <http://www.friends-of-dagstuhl.de>.

## ■ Association “Friends of Dagstuhl”

The registered association for support of Schloss Dagstuhl – Leibniz Center for Informatics (“Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.”) has existed since mid-2014. While very technical and rather clumsy, the name nevertheless reflects the precise purpose of the association: the promotion of science and research at the Leibniz Center for Informatics at Schloss Dagstuhl. The name chosen for the website (<http://www.friends-of-dagstuhl.de>), “Friends of Dagstuhl”, rolls off the tongue much more easily.

The association is geared towards acquiring and providing funds for the successful execution of its purposes as well as holding the funds entrusted to it. Therefore, the Dagstuhl Foundation (“Stiftung Informatikzentrum Schloss Dagstuhl”) was integrated into the association as a dependent foundation. Since late 2014, the Friends of Dagstuhl represent the Dagstuhl Foundation in legal and business transactions and manage the foundation assets under the strategic supervision of a foundation council (see Fig. 12.1). The association is chaired by a board (see Fig. 12.2 and Fig. 12.3).

In recent years, the association financed a shuttle service on arrival and departure days, using funds from a generous donation from the estate of a deceased friend of Schloss Dagstuhl. This made the journey between Schloss Dagstuhl and Türkismühle railway station much easier. The association is also involved in the acquisition and inventory of art objects in Schloss Dagstuhl and on its premises. The association’s support for these two measures amounted to a total of approximately € 10,000. The association is currently looking into the possibility of establishing a Dagstuhl Award, which could be used to recognize outstanding achievements in computer science or outstanding Dagstuhl Seminars in the future.

After the long-standing Chair Professor Holger Hermanns and the long-standing Deputy Chair Angelika Müller von Brochowski stepped down from the Executive Board at the end of 2023, all notarial and banking formalities were completed in 2024, meaning that the new Executive Board is now fully operational.

Further information about the association as well as the membership application form can be found at <http://www.friends-of-dagstuhl.de>.

**Stiftungsrat | Foundation council**

Prof. Dr. Holger Hermanns  
Universität des Saarlandes, Saarbrücken, Germany

Prof. Dr. Dr. h.c. mult. Kurt Mehlhorn  
Max Planck Institute for Informatics (MPII), Saarbrücken, Germany

Prof. Dr. Dorothea Wagner  
Karlsruher Institut für Technologie (KIT), Germany

Fig. 12.1

The council of the foundation “Informatik-Zentrum Schloss Dagstuhl”

**Vorstand des Vereins | Board of the association**

Prof. Dr. Markus Bläser (Chairperson)  
Universität des Saarlandes, Saarbrücken, Germany

Dr. Marc Herbstritt (Secretary and first deputy chairperson)  
Saarbrücken, Germany

Erich Reindel (Treasurer and second deputy chairperson)  
Universität des Saarlandes, Saarbrücken, Germany

Fig. 12.2

The board of the association “Friends of Dagstuhl”




Fig. 12.3

The board of the association “Friends of Dagstuhl”, f.l.t.r.: Erich Reindel, Prof. Dr. Markus Bläser, and Dr. Marc Herbstritt.



# 13 Statistiken

## *Statistics*

## Statistiken zu Seminaren und Workshops

13.1

## Statistics on Seminars and Workshops

In diesem Abschnitt werden statistische Daten zum wissenschaftlichen Programm und der Zusammenstellung der Teilnehmer aufgeführt. Durch die Covid-19 Pandemie bedingt, weichen die veranstaltungs- und teilnehmerbezogenen statistischen Werte der Jahre 2020 und 2021 erheblich von denen der anderen Jahre ab. Details sind in den Abschnitten weiter unten aufgeführt.

Die Diagramme und Tabellen sind dabei wie nachfolgend beschrieben gegliedert.

**Antrags-bezogene Daten:** Die Anzahl eingereicherter Anträge von Dagstuhl Seminaren und Dagstuhl Perspektiven Workshops sowie deren Akzeptanzraten sind in Fig. 13.1 dargestellt. Die sehr hohe Anzahl von Anträgen 2020 kann unter anderem auf die pandemiebedingte auf Juni 2020 verschobene zweite Antragsfrist zurückgeführt werden. In dieser Runde wurde mit 92 Anträge so viele wie nie zuvor gestellt. Die sehr niedrige Anzahl 2021 ist ausschließlich durch eine beschränkte erste Antragsrunde, in der nur wegen der Pandemie abgesagte Seminare neu beantragt werden durften, begründet. In dieser Runde wurden schließlich alle Anträge, die ja schon einmal positiv begutachtet worden waren, genehmigt, was den überdurchschnittlich hohen Anteil der genehmigten Seminare erklärt. Fig. 13.2 zeigt, wie die akzeptierten Seminare und Workshops sich bezüglich Größe und Länge aufgliedern.

**Veranstaltungs-bezogene Daten:** Fig. 13.3 zeigt die Anzahl der verschiedenen Veranstaltungstypen. Der dramatische Einbruch 2020 und die immer noch geringe Anzahl von Veranstaltungen in 2021 sind ebenso durch die Covid-19 Pandemie bedingt. In 2020 hatte Schloss Dagstuhl wegen der Pandemie ab Mitte März für etwa 5 Monate geschlossen. Aber auch nach Wiedereröffnung sind 2020 fast alle geplanten Veranstaltungen abgesagt worden. In den Jahren 2021 und 2022 bot Schloss Dagstuhl zum Ausgleich auch hybride Seminare, in denen Teilnehmer sowohl vor Ort waren als auch online über Audio-/Videoübertragung zugeschaltet waren, an. Als Spezialfall waren in reinen online Seminaren gar keine Teilnehmer vor Ort. Alle diese Typen werden in der Tabelle hier mitgezählt. Fig. 13.4 gibt aber einen Überblick, wie viele Veranstaltungen on-site, hybrid oder online stattgefunden haben. Daten zu der Anzahl der durchgeführten Seminare gegliedert nach der ursprünglich genehmigten Größe und Dauer sind in Fig. 13.5 angegeben. Manche Organisatorenteams von – insbesondere rein online stattfindenden – Seminaren hatten sich entschlossen, die Dauer teilweise erheblich zu verkürzen.

Zum Ausgleich der Einschränkungen von online oder hybrid stattfindenden Seminaren wurden bei diesen auch mehr Teilnehmer als ursprünglich genehmigt zugelassen. Fig. 13.6 zeigt Anzahl und Anteil der eingeladenen Seminarteilnehmer, welche die Einladung annehmen bzw. ablehnen. Der in 2021 auffallende,

This section provides statistical data about the scientific program and the composition of program participants. Due to the Covid-19 pandemic, the statistical values related to the program and participants of 2020 and 2021 deviate considerably from those of other years. Details are provided in the sections below.

Charts and tables are structured as described below.

**Proposal-related data:** Fig. 13.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops as well as the respective acceptance rates. The exceptionally large number of proposals in 2020 is due to the fact that the second submission deadline was postponed until June 2020, among other reasons. That proposal round saw 92 submissions, more than ever before. The very low number in 2021 is exclusively due to the fact that the first submission round was limited in that only resubmissions of proposals for seminars that had been cancelled due to the pandemic were allowed. Eventually, all proposals were accepted – they had obviously been accepted before – which explains the above-average share of accepted seminars. Size and duration of accepted seminars and workshops are displayed in Fig. 13.2.

**Event-related data:** Fig. 13.3 illustrates the number of different event types. The sharp drop in 2020 and the continuously low number of events in 2021 were also caused by the Covid-19 pandemic. From mid-March 2020, Schloss Dagstuhl was closed for approximately 5 months as a consequence of the pandemic. Even after re-opening, almost all scheduled events in 2020 were cancelled. In 2021 and 2022, Schloss Dagstuhl offered hybrid seminars where some participants were on site while others joined the seminar online via audio/video transmission. In special cases, there were also online-only seminars where there are no participants on site. The table includes all of these event types. However, Fig. 13.4 provides an overview of how many events took place on site, in a hybrid setting or online, respectively.

Fig. 13.5 illustrates data regarding the number of seminars that took place, categorised by originally approved size and duration. Some teams of seminar organizers – especially in the context of online-only seminars – decided to shorten the duration, in some cases drastically.

In order to compensate for the limitations of hybrid and online seminars, larger numbers of participants than originally approved were allowed to join. Fig. 13.6 shows the number and proportion of invited seminar participants who accepted or declined the invitation. It is evident that the share of invitees who accepted their invitation was significantly higher in 2021. This may in part be due to the fact that some seminars offered

deutlich höhere Anteil an Eingeladenen, die zugesagt haben, mag zum einen an der Möglichkeit der Online-Teilnahme eines Teils der angebotenen Seminare liegen, aber genauso auch daran, dass zu manchen Seminaren noch sehr kurzfristig vor dem Seminar Personen eingeladen wurden, deren Zusage vor der Einladung abgesprochen wurde. Auch dies wurde, zum Teil, durch die Möglichkeit der Online-Teilnahme ermöglicht.

Die Verteilung der Annahmerate pro Seminar, also der Quotient aus Teilnehmer und allen Eingeladenen, ist in Fig. 13.7 dargestellt. Durch die pandemiebedingten Besonderheiten des Seminarbetriebs in und den deutlich kleineren Stichproben aus 2020 und 2021 sind die Daten dieser Jahre nicht im Vergleich der anderen Jahre zu interpretieren.

Fig. 13.8 zeigt dagegen, wie viel Prozent der zugesagten Größe (gemessen an der Personenanzahl) tatsächlich von einem Seminar belegt wurde. Auch hier sind die Daten der letzten beiden Jahre nicht mit den vorherigen Daten zu vergleichen. Werte über 100 % bedeuten, dass mehr Personen als ursprünglich genehmigt am Seminar teilgenommen haben. Während dies bei online zugeschalteten Teilnehmer problemlos möglich ist, war die Gesamtanzahl der Teilnehmer im Haus und damit auch für die Veranstaltungen strikter limitiert.

**Teilnehmer-bezogene Daten:** Grundsätzlich werden hier alle Teilnehmer gezählt, unabhängig davon, ob sie vor Ort oder online teilgenommen haben. Vor allem durch ausfallende Veranstaltungen aber auch durch schlechtere Beteiligung insbesondere an Veranstaltungen die ausschließlich vor Ort stattfanden, sind die Zahlen aus 2020 und 2021 deutlich geringer als die der Vorjahre. Die Teilnehmerzahlen – abhängig vom Veranstaltungstyp – gibt Fig. 13.9 an. Einen Eindruck, wieviele Teilnehmer vor Ort waren bzw. remote teilgenommen haben, vermittelt Fig. 13.10.

Fig. 13.11 zeigt die Verteilung der Herkunftsländer unserer Gäste.

**Umfrage-bezogene Daten:** Hier stellen wir ausgewählte Daten unserer fortlaufenden Befragung von Teilnehmern an Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops dar. Ein Überblick über die Ergebnisse der regelmäßigen Gästebefragungen kann Fig. 13.12 entnommen werden. Die Anzahl von früheren Seminarbesuchen kann man Fig. 13.13 entnehmen. Fig. 13.14 gibt Auskunft über die Altersstruktur der Teilnehmer. Während Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops sich primär an Forscher aus Universitäten und Forschungseinrichtungen richten, sind auch Anwender und Forscher aus der Industrie stets willkommen. Die Verteilung ihres Anteils ist in Fig. 13.15 gezeigt.

**Auslastungs-bezogene Daten:** Die Auslastung des Zentrums wird schließlich in Fig. 13.16 an Hand der Übernachtungen und ihrer Verteilung über die einzelnen Wochen getrennt nach Veranstaltungstypen aufgezeigt. Im Gegensatz zu den anderen Statistiken beziehen sich diese Daten ausschließlich auf Gäste vor Ort.

online participation, but may also have been caused by issuing invitations to some seminars at very short notice when the invitees had informally agreed to attend in advance. This was also partly made possible through offering online participation.

The acceptance rate distribution per seminar, i.e., the quotient of the number of participants and the total number of invitees, is illustrated in Fig. 13.7. Due to the particularities regarding the seminars and the significantly smaller samples in 2020 and 2021 caused by the pandemic, the data from those years is not to be interpreted in relation to previous years.

In contrast, Fig. 13.8 visualizes the percentage of the reserved space (in terms of number of people) that was actually used by seminar participants. Again, the data from the last two years cannot be compared to previous data. Values above 100 percent show that more people than originally approved attended the seminar. While this was easy to facilitate with regard to online participants, the total number of participants on site and therefore the maximum number of on-site participants allowed for the respective events was limited more strictly.

**Participant-related data:** Basically, all participants are counted here, regardless of whether they participated on site or online. The numbers from 2020 and 2021 are significantly lower than the ones in previous years due especially to event cancellations, but also due to lower participation rates, particularly with regard to on-site only events.

Fig. 13.9 shows the number of participants broken down by event type. Fig. 13.10 illustrates how many participants were on site and how many participated remotely.

Fig. 13.11 shows the distribution of our guests' country affiliations.


**Survey-related data:** In this section we present selected data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results gained from routine participants surveys for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Fig. 13.12. Fig. 13.13 displays how often participants have attended seminars in the past. Fig. 13.14 provides data on the seniority of participants. While Dagstuhl Seminars and Dagstuhl Perspectives Workshops mainly target academic researchers, researchers and developers from industry are also always welcome. The distribution of their share is shown in Fig. 13.15.

**Utilization-related data:** Finally, Fig. 13.16 illustrates utilization on the basis of overnight stays hosted at Schloss Dagstuhl – broken down by event type – as well as their distribution by week. In contrast to the other statistics, this data exclusively refers to guests on site.

**Gender-related data:** Fig. 13.17 illustrates gender distribution with respect to organizer teams of Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In contrast, Fig. 13.18 shows the proportion of women in seminar proposals with respect to both the teams

**Geschlechter-bezogene Daten:** Fig. 13.17 enthält Daten zur Geschlechter-Verteilung in der Seminarleitung. Dagegen zeigt Fig. 13.18 die Quote von Frauen bei der Beantragung von Seminaren sowohl bezüglich der Teams als auch bezüglich der gesamten Antragsteller. Die Abbildungen Fig. 13.19 und Fig. 13.20 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler. Die Verteilung der Rate der weiblichen Teilnehmer in den einzelnen Seminaren wird in Fig. 13.21 aufgezeigt. Diese Statistiken werden von unsere Geldgeber im Rahmen der Frauenförderung gefordert. Die Daten basieren vorwiegend auf der Selbstauskunft von Gästen, dass sie weiblich sind.

and the proposers overall. Fig. 13.19 and Fig. 13.20 mainly illustrate the share of female participants and invitees, respectively. The distribution of the share of women among seminar participants is displayed in Fig. 13.21. Please note that our funding agencies require such statistics as part of their efforts to promote gender equality. The statistics are predominantly based on guests' self-reporting that they are female.




(a) Chart for 2018–2024

Year	Proposals		Accepted		Rejected	
	#	%	#	%	#	%
2018	136		81	59.6	55	40.4
2019	103		58	56.3	45	43.7
2020	153		87	56.9	66	43.1
2021	74		57	77.0	17	23.0
2022	93		64	68.8	29	31.2
2023	111		67	60.4	44	39.6
2024	172		96	55.8	76	44.2

(b) Detailed numbers for 2018–2024

Fig. 13.1

Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.




(a) Chart for 2018–2024

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2018	10	22	0	49	81
2019	5	17	1	35	58
2020	15	27	0	45	87
2021	2	15	1	39	57
2022	9	19	0	36	64
2023	7	23	1	36	67
2024	14	29	0	53	96

(b) Detailed numbers for 2018–2024

Fig. 13.2

Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2018–2024. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.




(a) Chart for 2018–2024


Year	DS	PW	GI	EDU	RGM	Total
2018	63	2	2	2	40	109
2019	71	2	2	7	29	111
2020	13	0	0	0	24	37
2021	47	0	0	1	28	76
2022	75	2	0	2	26	105
2023	76	1	2	3	25	107
2024	74	3	0	4	33	114

(b) Detailed numbers for 2018–2024

**Fig. 13.3**  
**Number of all events held at Dagstuhl, by type.** DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.



(a) Graphical distribution of seminars in group A in 2022–2024 by year and form




(b) Graphical distribution of events in group B in 2022–2024 by year and event form

Year	Group A			Group B		
	On-Site	Hybrid	Online	On-Site	Hybrid	Online
2022	59	16	2	23	5	0
2023	76	1	0	30	0	0
2024	77	0	0	35	2	0

(c) Detailed numbers for 2022–2024 by event form

**Fig. 13.4**  
**Number of all events held at Dagstuhl, by event form and group.** Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops. Group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).




(a) Chart for 2018–2024

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2018	9	14	1	41	65
2019	9	23	0	41	73
2020	1	2	0	10	13
2021	9	13	1	24	47
2022	6	21	1	49	77
2023	9	25	0	43	77
2024	9	25	1	42	77

(b) Detailed numbers for 2018–2024

**Fig. 13.5**  
**Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2018–2024.** Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.




(a) Chart for 2018–2024

Year	Invitees		Accepted		Declined	
	#	%	#	%	#	%
2018	4692	49.4	2320	49.4	2372	50.6
2019	5143	48.6	2498	48.6	2645	51.4
2020	964	45.9	442	45.9	522	54.1
2021	3022	62.7	1894	62.7	1128	37.3
2022	4950	46.0	2278	46.0	2672	54.0
2023	5494	44.3	2433	44.3	3061	55.7
2024	5286	48.4	2556	48.4	2730	51.6

(b) Detailed numbers for 2018–2024

Fig. 13.6

Total number of invitees, accepted and declined invitations for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.




(a) Chart for 2018–2024

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2018	21.8	77.0	51.2	12.0
2019	24.8	77.5	49.8	11.4
2020	14.3	66.7	45.8	17.4
2021	23.0	85.9	63.2	12.5
2022	22.2	90.9	47.9	14.1
2023	23.9	69.6	45.8	9.9
2024	21.9	81.7	49.8	11.6

(b) Detailed numbers for 2018–2024

Fig. 13.7

Acceptance rate distribution per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2018–2024. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.




(a) Chart for 2018–2024

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2018	60.0	113.3	90.3	10.2
2019	57.8	120.0	89.1	10.7
2020	30.0	131.1	79.7	27.9
2021	44.4	153.3	108.0	24.6
2022	44.4	133.3	76.1	18.7
2023	55.6	106.7	82.8	8.2
2024	53.3	113.3	86.7	9.5

(b) Detailed numbers for 2018–2024

Fig. 13.8

Occupancy rate distribution per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2018–2024. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.




(a) Chart for 2018–2024


Year	DS		PW		GI		EDU		RGM		Total #
	#	%	#	%	#	%	#	%	#	%	
2018	2268	70.8	52	1.6	50	1.6	99	3.1	733	22.9	3202
2019	2450	74.3	48	1.5	50	1.5	282	8.5	469	14.2	3299
2020	442	65.1	0	0.0	0	0.0	0	0.0	237	34.9	679
2021	1894	84.6	0	0.0	0	0.0	31	1.4	314	14.0	2239
2022	2229	79.3	49	1.7	0	0.0	97	3.5	436	15.5	2811
2023	2406	79.6	27	0.9	40	1.3	119	3.9	430	14.2	3022
2024	2483	76.6	73	2.3	0	0.0	139	4.3	546	16.8	3241

(b) Detailed numbers for 2018–2024

**Fig. 13.9**  
**Number of participants.** DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.



(a) Graphical distribution of participants of seminars in group A in 2018–2024 by year and attendance status



(b) Graphical distribution of participants of events in group B in 2018–2024 by year and attendance status

Year	Group A			Group B		
	On-Site	Online	Total	On-Site	Online	Total
2018	2320	0	2320	882	0	882
2019	2498	0	2498	801	0	801
2020	419	23	442	237	0	237
2021	529	1365	1894	233	112	345
2022	1895	383	2278	491	42	533
2023	2395	38	2433	588	1	589
2024	2549	7	2556	682	3	685


(c) Detailed numbers for 2018–2024 by attendance status

**Fig. 13.10**  
**Number of participants by attendance status and group.** Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops. Group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).


Country	A	B	Total
Germany	552	486	1038
United States	474	41	515
United Kingdom	230	26	256
Netherlands	142	26	168
France	151	11	162
Italy	105	6	111
Austria	86	13	99
Canada	90	8	98
Switzerland	68	16	84
Sweden	53	12	65
Denmark	58	3	61
Israel	59	2	61
Belgium	54	2	56
Spain	51	2	53
Australia	44	2	46
Japan	41	1	42
India	38	0	38
Norway	24	7	31
Finland	28	1	29
Czech Republic	23	2	25
China	22	0	22
Ireland	16	1	17
Poland	16	1	17
Singapore	16	0	16
Republic of Korea	14	0	14

Country	A	B	Total
Brazil	12	1	13
Portugal	11	2	13
Chile	10	0	10
Greece	9	0	9
Luxembourg	9	0	9
Hong Kong	7	0	7
Hungary	4	2	6
Serbia	4	2	6
New Zealand	5	0	5
Slovenia	3	2	5
United Arab Emirates	5	0	5
Malta	4	0	4
Argentina	3	0	3
Bulgaria	0	3	3
Iceland	1	2	3
Taiwan	3	0	3
Iran	2	0	2
Kenya	0	2	2
Turkey	2	0	2
Romania	1	0	1
Russian Federation	1	0	1
Rwanda	1	0	1
Saudi Arabia	1	0	1
Senegal	1	0	1
South Africa	1	0	1
Uruguay	1	0	1
Total	2556	685	3241

(a) Details for 2024 by country



(b) Graphical distribution of seminars in group A in 2018–2024 by year and region




(c) Graphical distribution of events in group B in 2018–2024 by year and region

Region	Group A		Group B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	1154	45.1	142	20.7	1296	40
Germany	552	21.6	486	70.9	1038	32
North America	564	22.1	49	7.2	613	18.9
Asia	208	8.1	3	0.4	211	6.5
Australia	49	1.9	2	0.3	51	1.6
South America	26	1	1	0.1	27	0.8
Africa	3	0.1	2	0.3	5	0.2
Total	2556	100	685	100	3241	100

(d) Details for 2024 by region

Fig. 13.11

**Number of Dagstuhl guests by country of origin.** Group A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants. Group B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).




(a) Graphical distribution for 2024

	2018	2019	2020	2021	2022	2023	2024	2024 – Detailed Numbers					
	Ø	Ø	Ø	Ø	Ø	Ø	Ø	1	2	3	4	5	total
would come again	4.9	4.9	4.9	4.9	4.9	4.9	4.9	1	1	7	123	1299	1431
found inspiration	4.6	4.6	4.6	4.5	4.6	4.6	4.6	6	4	32	436	969	1447
found collaboration	4.2	4.3	4.3	4.1	4.4	4.4	4.4	5	10	122	574	709	1420
found insight from neighboring fields	4.2	4.3	4.2	4.2	4.4	4.3	4.3	7	31	130	567	688	1423
found new research direction	4.2	4.2	4.2	4.2	4.3	4.3	4.3	4	25	130	672	595	1426
group composition	4.5	4.6	4.5	4.5	4.6	4.6	4.6	1	7	32	446	958	1444
integration of junior researchers	4.3	4.3	4.3	4.3	4.4	4.3	4.4	6	36	106	493	789	1430
new professional contacts	3.8	3.8	3.6	3.5	3.8	3.9	4.0	26	103	282	484	504	1399
exchange between academia and industry	4.4	4.4	4.3	4.3	4.3	4.4	4.4	5	20	88	299	570	982
advance information from Dagstuhl	4.4	4.5	4.5	4.3	4.4	4.4	4.5	0	2	97	489	821	1409
advance information from organizers	4.3	4.3	4.3	4.2	4.2	4.2	4.4	1	11	169	495	716	1392
number and length of talks	4.3	4.3	4.3	4.3	4.4	4.4	4.4	3	39	85	516	778	1421
opportunity for one on one talks	4.6	4.6	4.6	4.1	4.7	4.7	4.7	2	22	31	303	1066	1424
flexibility of schedule	4.3	4.5	4.4	4.2	4.4	4.5	4.5	4	41	94	432	859	1430
open and honest discussion	4.7	4.7	4.7	4.6	4.8	4.8	4.8	2	5	17	262	1149	1435
outing	4.3	4.2	4.2	4.4	4.4	4.4	4.4	9	22	119	383	662	1195
venue	4.7	4.7	4.7	4.8	4.8	4.8	4.8	1	3	29	280	1124	1437
conference facilities	4.7	4.7	4.7	4.6	4.7	4.7	4.7	1	4	47	318	1066	1436
IT facilities	4.4	4.4	4.3	4.3	4.2	4.3	4.3	5	24	155	457	591	1232
staff support	4.8	4.8	4.7	4.7	4.8	4.8	4.8	1	0	26	206	1139	1372
meals	4.1	4.2	4.1	4.3	4.2	4.3	4.3	6	25	167	591	632	1421
rooms	4.4	4.4	4.3	4.4	4.4	4.4	4.4	3	28	154	510	725	1420
leisure facilities	4.6	4.6	4.6	4.5	4.6	4.6	4.6	0	6	43	386	881	1316
library services	4.6	4.6	4.5	4.5	4.5	4.5	4.6	0	0	37	153	325	515

(b) Averages for 2018–2024 and detailed numbers for 2024: 1 = very low, 2 = low, 3 = medium, 4 = high, 5 = very high

Fig. 13.12 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.




(a) Graphical distribution for 2018–2024

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	
2018	557	41	219	16	148	11	425	32	1349
2019	615	41	230	15	144	10	503	34	1492
2020	61	29	31	15	25	12	90	43	207
2021	283	43	96	15	63	10	211	32	653
2022	470	42	156	14	107	9	398	35	1131
2023	509	40	177	14	140	11	450	35	1276
2024	620	44	222	16	138	10	442	31	1422

(b) Detailed numbers for 2018–2024

Fig. 13.13

Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants and their previous instances of attendance at Dagstuhl Seminars or Dagstuhl Perspectives Workshops, according to our guest survey.




(a) Chart for 2018–2024

Year	Junior		Senior		Total
	#	%	#	%	
2018	401	35.7	722	64.3	1123
2019	385	31.9	823	68.1	1208
2020	53	28.5	133	71.5	186
2021	195	34.5	370	65.5	565
2022	322	32.6	665	67.4	987
2023	311	28.6	775	71.4	1086
2024	389	31.7	839	68.3	1228

(b) Detailed numbers for 2018–2024

Fig. 13.14

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.




(a) Chart for 2018–2024

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2018	0.0	41.7	11.1	10.4
2019	0.0	41.7	11.4	10.7
2020	0.0	100.0	14.6	25.9
2021	0.0	31.3	8.5	10.1
2022	0.0	46.7	10.1	11.1
2023	0.0	71.4	12.6	14.0
2024	0.0	60.7	10.8	13.4

(b) Detailed numbers for 2018–2024

Fig. 13.15

Distribution of the ratio of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2018–2024, according to our guest survey. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation. Occupation in business includes “industrial research”, “industrial development”, and “self employed”.



(a) Chart for 2018–2024


(b) Graphical distribution for 2022–2024 by week

Year	DS	PW	GI	EDU	RGM	Total
2018	10270	182	250	231	1740	12673
2019	11127	225	239	1004	1144	13739
2020	1984	0	0	0	614	2598
2021	2397	0	0	150	576	3123
2022	8435	252	0	218	1104	10009
2023	10842	121	195	338	1130	12626
2024	11400	360	0	441	1465	13666

(c) Detailed numbers for 2018–2024

Fig. 13.16

Number of overnight stays at Schloss Dagstuhl. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.




(a) Chart for 2018–2024

(b) Detailed numbers for 2018–2024

Fig. 13.17

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with organizer teams including women.




(a) Chart for 2018–2024

Year	Proposer Teams			Proposers		
	Total	incl. Women		Total	Women	
	#	#	%	#	#	%
2018	136	102	75.0	522	140	26.8
2019	103	84	81.6	411	127	30.9
2020	153	123	80.4	593	178	30.0
2021	74	57	77.0	296	93	31.4
2022	93	78	83.9	367	124	33.8
2023	111	104	93.7	438	158	36.1
2024	172	156	90.7	677	244	36.0

(b) Detailed numbers for 2018–2024

Fig. 13.18

Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with proposer teams including women.




(a) Chart for 2018–2024

Year	Invitees			Participants		
	Total	Women		Total	Women	
	#	#	%	#	#	%
2018	4692	1086	23.1	2320	453	19.5
2019	5143	1366	26.6	2498	603	24.1
2020	964	218	22.6	442	99	22.4
2021	3022	794	26.3	1894	456	24.1
2022	4950	1593	32.2	2278	635	27.9
2023	5494	1648	30.0	2433	631	25.9
2024	5286	1576	29.8	2556	676	26.4

(b) Detailed numbers for 2018–2024

Fig. 13.19

Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.




(a) Chart for 2018–2024

Year	Female Invitees	Accepted		Declined	
	#	#	%	#	%
2018	1086	453	41.7	633	58.3
2019	1366	603	44.1	763	55.9
2020	218	99	45.4	119	54.6
2021	794	456	57.4	338	42.6
2022	1593	635	39.9	958	60.1
2023	1648	631	38.3	1017	61.7
2024	1576	676	42.9	900	57.1

(b) Detailed numbers for 2018–2024

Fig. 13.20

Female invitees to Dagstuhl Seminars and Dagstuhl Perspectives Workshops.



Year	Min (%)	Max (%)	Avg (%)	Std (%)
2018	5.7	47.8	19.8	9.2
2019	4.7	60.6	24.6	10.5
2020	2.6	38.5	22.7	9.2
2021	8.2	57.9	24.5	11.4
2022	8.0	54.5	28.5	10.7
2023	9.3	51.9	26.3	10.3
2024	11.5	48.1	26.9	7.5

(a) Chart for 2018–2024

(b) Detailed numbers for 2018–2024

Fig. 13.21

**Distribution of the share of women among participants per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2018–2024.** Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.


## Statistiken zur Bibliographiedatenbank dblp

13.2

## Statistics of the dblp computer science bibliography

Dieser Abschnitt enthält statistische Daten zur Bibliographiedatenbank dblp. Fig. 13.22 listet die durchschnittlichen Nutzungszahlen der letzten Jahre auf. Ein Überblick über die Anzahl der Neuaufnahmen in den dblp Datenbestand kann Fig. 13.23 entnommen werden. Fig. 13.24–13.26 geben Auskunft über die kontinuierliche Datenkuration und -anreicherung des Bestandes.

This section provides statistical data about the dblp computer science bibliography. Fig. 13.22 shows the average usage statistics of the dblp servers in the past years. An overview of the rate of new publications added to the dblp database can be found in Fig. 13.23. Information about the continuous data curation and enrichment of existing records can be found in Fig. 13.24–13.26.




(a) Chart for 2018–2024

	Trier		Dagstuhl		Total		
	2023	2024	2023	2024	2023	2024	%
page views per day	565,541	900,013	1,184,112	1,633,244	1,749,654	2,533,258	+44.8
user sessions (visits) per day	26,469	98,681	57,142	138,991	83,611	237,672	+184.3
page views per user session	21.4	9.1	20.7	11.8	20.9	10.7	-49.1
distinct users (IPs) per month	278,876	1,191,266	692,319	1,579,931	971,196	2,771,198	+185.3
data served per month	1,165.7 GB	1,103.1 GB	1,441.8 GB	1,564.0 GB	2,607.5 GB	2,667.0 GB	+2.3

(b) Detailed numbers for the past two years

Fig. 13.22

**Average usage of the two public dblp servers.** Trier = dblp.uni-trier.de, Dagstuhl = dblp.dagstuhl.de or dblp.org. All figures exclude traffic caused by recognized bots and web crawlers. In 2021, the counting method to determine distinct users has been changed in order to avoid double counting of the same IP across different servers. While these figures generally ignore accesses from recognized bots and web crawlers, please note that the numbers given for 2024 are skewed by a currently still unknown amount of automated accesses from scripted agents, as described in the text of Section 3.7.



(a) Chart for 1996–2024

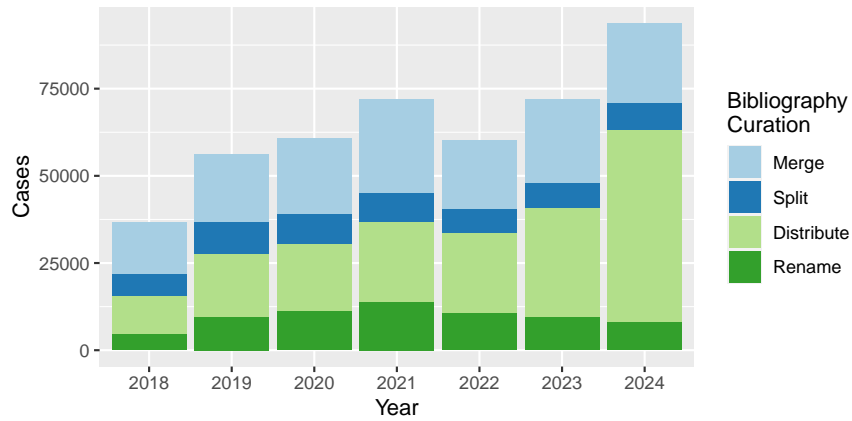
Year	Book		Article		Inproceedings		Incollection		Editor		Reference	
	#	%	#	%	#	%	#	%	#	%	#	%
2018	5,841	1.4	148,732	35.5	215,099	51.3	1,607	0.4	3,746	0.9	61	0.0
2019	7,577	1.6	170,310	37.1	211,713	46.1	5,749	1.3	4,103	0.9	3,847	0.8
2020	4,616	0.9	199,991	39.4	217,418	42.9	5,915	1.2	3,583	0.7	324	0.1
2021	3,847	0.7	214,795	41.3	216,029	41.5	1,373	0.3	3,487	0.7	0	0.0
2022	20,194	3.8	211,924	39.9	209,017	39.3	1,662	0.3	3,727	0.7	0	0.0
2023	21,948	4.0	209,877	37.9	213,980	38.6	1,264	0.2	3,875	0.7	45	0.0
2024	16,109	2.6	252,282	40.8	223,591	36.1	538	0.1	3,521	0.6	0	0.0

Year	Data		Informal		Withdrawn		Total
	#	%	#	%	#	%	#
2018	465	0.1	42,261	10.1	1,427	0.3	419,239
2019	888	0.2	54,649	11.9	458	0.1	459,294
2020	828	0.2	74,043	14.6	333	0.1	507,051
2021	652	0.1	80,115	15.4	299	0.1	520,597
2022	661	0.1	82,786	15.6	1,721	0.3	531,692
2023	1,110	0.2	98,967	17.9	2,973	0.5	554,039
2024	5,389	0.9	114,997	18.6	2,237	0.4	618,664

(b) Detailed numbers for 2018–2024

Fig. 13.23

**Newly indexed publications in dblp.** In chart (a), the negative number of new *Incollection* records in 2011 results from relabeling several thousand existing records with the newly introduced *Reference* type. Similarly, in the same year, several thousand *Articles* and *Inproceedings* records have been relabeled as *Informal*.




(a) Chart for 2018–2024

Year	Merge		Split		Distribute		Rename		Total #
	#	%	#	%	#	%	#	%	
2018	14,906	40.6	6,282	17.1	11,014	30.0	4,524	12.3	36,726
2019	19,595	34.9	9,192	16.4	17,795	31.7	9,562	17.0	56,144
2020	21,636	35.7	8,606	14.2	19,326	31.9	11,083	18.3	60,651
2021	27,204	37.7	8,254	11.4	22,877	31.7	13,785	19.1	72,120
2022	19,739	32.8	6,813	11.3	23,065	38.4	10,508	17.5	60,125
2023	24,089	33.5	7,008	9.7	31,443	43.7	9,340	13.0	71,880
2024	22,832	24.4	7,880	8.4	55,016	58.7	8,013	8.5	93,741

(b) Detailed numbers for 2018–2024

Fig. 13.24

**Curation of existing dblp author bibliographies.** The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp author bibliographies. We distinguish between four curation cases: *Merge* = Two or more synonymous bibliographies have been merged into a single bibliography. *Split* = A single, homonymous bibliography has been split into two or more bibliographies. *Distribute* = A mixed case where records from two or more bibliographies have been redistributed between two or more bibliographies. *Rename* = A case where no actual publications have been reassigned, but the surface form of the author name(s) of a bibliography have been corrected or improved.




(a) Chart for 2018–2024

Year	Homepages
2018	40,018
2019	44,519
2020	48,216
2021	51,756
2022	55,572
2023	59,192
2024	62,756

(b) Detailed numbers for 2018–2024

Fig. 13.25

**Linked and verified academic homepages in dblp author bibliographies.** A single author bibliography may be linked to multiple academic homepages. These figures exclude linked external IDs which are given in Figure 13.26.



(a) Chart for 2018–2024

Year	ORCID		Google Scholar		Wikidata		Wikipedia (en)		ACM DL		IEEE Xplore		zbMATH	
	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2018	30,939	36.4	10,182	12.0	15,565	18.3	4,366	5.1	2,144	2.5	0	0.0	2,145	2.5
2019	61,976	41.1	13,726	9.1	30,022	19.9	5,547	3.7	2,448	1.6	104	0.1	3,363	2.2
2020	81,675	41.8	16,338	8.4	45,216	23.1	6,275	3.2	3,166	1.6	557	0.3	3,607	1.8
2021	99,459	41.2	20,060	8.3	51,744	21.4	6,668	2.8	9,642	4.0	3,013	1.2	7,486	3.1
2022	119,683	40.3	24,443	8.2	59,514	20.0	7,312	2.5	10,306	3.5	10,269	3.5	9,524	3.2
2023	146,690	41.7	28,108	8.0	66,337	18.9	8,050	2.3	11,035	3.1	12,979	3.7	10,874	3.1
2024	175,457	44.0	31,532	7.9	71,325	17.9	8,523	2.1	11,578	2.9	15,685	3.9	12,214	3.1

Year	Math Genealogy		LOC		GND		VIAF		ISNI		Twitter		Other		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2018	65	0.1	2,862	3.4	2,385	2.8	938	1.1	1,231	1.4	1,567	1.8	10,670	12.5	85,059
2019	3,071	2.0	5,680	3.8	3,782	2.5	1,060	0.7	1,383	0.9	2,183	1.4	16,349	10.8	150,694
2020	3,590	1.8	6,289	3.2	4,241	2.2	1,061	0.5	1,385	0.7	2,613	1.3	19,358	9.9	195,371
2021	5,048	2.1	6,877	2.8	4,777	2.0	1,068	0.4	1,390	0.6	3,183	1.3	21,233	8.8	241,648
2022	7,060	2.4	8,234	2.8	5,938	2.0	1,069	0.4	1,399	0.5	3,786	1.3	28,652	9.6	297,189
2023	8,892	2.5	9,513	2.7	12,611	3.6	1,071	0.3	1,407	0.4	4,043	1.1	30,229	8.6	351,839
2024	10,011	2.5	10,260	2.6	13,388	3.4	1,071	0.3	1,412	0.4	4,138	1.0	31,873	8.0	398,467

(b) Detailed numbers for 2018–2024

Fig. 13.26

Linked and verified external person IDs in dblp author bibliographies. A single bibliography may be linked to multiple external IDs.

## Statistiken zu Dagstuhl Publishing

13.3

## Statistics of Dagstuhl Publishing

Dieser Abschnitt enthält statistische Daten zum Publikationswesen von Schloss Dagstuhl.

Ein Überblick über die Entwicklung der seminarbezogenen Veröffentlichungen kann den ersten zwei Diagrammen und Tabellen entnommen werden. Fig. 13.27 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen und Fig. 13.28 die der Publikationen in der Reihe Dagstuhl Manifestos.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 13.29 fasst die Daten in der Reihe OASICs und Fig. 13.30 die der Reihe LIPICs zusammen.

Die Kennzahlen der Zeitschriften LITES und TGDK können Fig. 13.31 bzw. Fig. 13.32 entnommen werden.

Fig. 13.34 zeigt die Anzahl der Online-Zugriffe auf die PDFs und Abstracts/Landing Pages der veröffentlichten Artikel.

Die verschiedenen Publikationsserien wurden in unterschiedlichen Jahren zwischen 2009 und 2023 gegründet. Wir stellen in den Statistiken dennoch stets den Zeitraum 2018–2024 dar.

In this section the statistical data of Dagstuhl Publishing is presented.


The first two figures present the development of the seminar-focused series: Fig. 13.27 summarizes the data of the periodical Dagstuhl Reports and Fig. 13.28 the data of the Dagstuhl Manifestos series.

The statistical data of the service-focused series are presented afterwards. Fig. 13.29 presents numbers related to OASICs and Fig. 13.30 numbers related to LIPICs.

We summarize the publications of the journal LITES in Fig. 13.31 and of the journal TGDK in Fig. 13.32.

The Fig. 13.34 shows the number of online accesses of the PDFs and abstracts/landing pages of the published all articles.

Please note that the publication series were established in different years in the period between 2009 and 2023. However, we always consider the period 2018–2024.




(a) Graphical distribution for 2018–2024

Year	Articles	Pages
2018	78	1938
2019	75	1959
2020	49	1322
2021	28	770
2022	53	1276
2023	80	2126
2024	84	2268

(b) Detailed numbers for 2018–2024

Fig. 13.27  
Statistics about Dagstuhl Reports published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Articles	Pages
2018	5	141
2019	0	0
2020	1	36
2021	2	63
2022	0	0
2023	0	0
2024	1	63

(b) Detailed numbers for 2018–2024

**Fig. 13.28**  
Statistics about Dagstuhl Manifestos published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Volumes	Articles	Pages
2018	7	107	1312
2019	9	120	1576
2020	11	133	1754
2021	11	161	2284
2022	10	95	1306
2023	9	120	1524
2024	10	134	2006

(b) Detailed numbers for 2018–2024

**Fig. 13.29**  
Statistics about OASlcs volumes published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Volumes	Articles	Pages
2018	32	1387	21876
2019	29	1208	20032
2020	32	1352	24562
2021	32	1333	23960
2022	36	1444	25998
2023	35	1508	26982
2024	38	1648	30214

(b) Detailed numbers for 2018–2024

**Fig. 13.30**  
Statistics about LIPIcs volumes published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Articles	Pages
2018	4	134
2019	5	188
2020	0	0
2021	2	64
2022	10	271
2023	0	0
2024	0	0

(b) Detailed numbers for 2018–2024

**Fig. 13.31**  
Statistics about LITES articles published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Articles	Pages
2018	0	0
2019	0	0
2020	0	0
2021	0	0
2022	0	0
2023	13	398
2024	16	452

(b) Detailed numbers for 2018–2024

**Fig. 13.32**  
Statistics about TGDK articles published between 2018 to 2024.




(a) Graphical distribution for 2018–2024

Year	Articles	Pages
2018	19	90
2019	22	82
2020	24	92
2021	18	100
2022	33	141
2023	30	113
2024	31	137

(b) Detailed numbers for 2018–2024

**Fig. 13.33**  
Statistics about DARTS artifacts published between 2018 to 2024.



(a) Graphical distribution for 2018–2024

Year	PDF Downloads	Abstract
2018	627,988	371,594
2019	737,721	452,321
2020	940,347	656,112
2021	1,242,064	722,787
2022	1,082,872	789,730
2023	1,146,962	1,104,269
2024	1,434,114	1,413,547

(b) Detailed numbers for 2018–2024

Fig. 13.34

Total number of full text accesses (PDF Downloads) and abstract/landing page views for articles published between 2018 to 2024.



# **14** **Veranstaltungen 2024** *Schedule of Events 2024*

**Dagstuhl-Seminare**

14.1

**Dagstuhl Seminars****24021 – From Proofs to Computation in Geometric Logic and Generalizations**

Ingo Blechschmidt (Universität Augsburg, DE), Thierry Coquand (University of Gothenburg, SE), Hajime Ishihara (Toho University – Chiba, JP), Peter M. Schuster (University of Verona, IT)  
January 7–12, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24021>

**24031 – Fusing Causality, Reasoning, and Learning for Fault Management and Diagnosis**

Alessandro Cimatti (Bruno Kessler Foundation – Trento, IT), Alexander Diedrich (Helmut-Schmidt-Universität – Hamburg, DE), Meir Kalech (Ben Gurion University – Modiin, IL), Maria Krantz (Helmut-Schmidt-Universität – Hamburg, DE), Ingo Pill (Silicon Austria Labs – Graz, AT)  
January 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24031>

**24032 – Representation, Provenance, and Explanations in Database Theory and Logic**

Pablo Barcelo (PUC – Santiago de Chile, CL), Pierre Bourhis (CNRS – CRISTAL, Lille, FR), Stefan Mengel (CNRS, CRIL – Lens, FR), Sudeepa Roy (Duke University – Durham, US)  
January 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24032>

**24041 – Symmetric Cryptography**

Christof Beierle (Ruhr-Universität Bochum, DE), Bart Mennink (Radboud University Nijmegen, NL), Maria Naya-Plasencia (INRIA – Paris, FR), Yu Sasaki (NTT – Tokyo, JP)  
January 21–26, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24041>

**24042 – The Emerging Issues in Bioimaging AI Publications and Research**

Jianxu Chen (ISAS – Dortmund, DE), Florian Jug (Human Technopole – Milano, IT), Susanne Rafelski (Allen Institute for Cell Science – Seattle, US), Shanghang Zhang (Peking University, CN)  
January 21–24, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24042>

**24051 – Next Generation Protocols for Heterogeneous Systems**

Stephanie Balzer (Carnegie Mellon University – Pittsburgh, US), Marco Carbone (IT University of Copenhagen, DK), Roland Kuhn (Actyx AG – Kassel, DE), Peter Thiemann (Universität Freiburg, DE)  
January 28 to February 2, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24051>

**24052 – Reviewer No. 2: Old and New Problems in Peer Review**

Iryna Gurevych (TU Darmstadt, DE), Anna Rogers (IT University of Copenhagen, DK), Nihar Shah (Carnegie Mellon University – Pittsburgh, US)  
January 28 to February 2, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24052>

**24061 – Are Knowledge Graphs Ready for the Real World? Challenges and Perspective**

David Chaves-Fraga (University of Santiago de Compostela, ES), Oscar Corcho (Technical University of Madrid, ES), Anastasia Dimou (KU Leuven, BE), Maria-Esther Vidal (TIB – Hannover, DE)  
February 4–9, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24061>

**24062 – Beyond-Planar Graphs: Models, Structures and Geometric Representations**

Vida Dujmovic (University of Ottawa, CA), Seok-Hee Hong (The University of Sydney, AU), Michael Kaufmann (Universität Tübingen, DE), János Pach (Alfréd Rényi Institute – Budapest, HU & EPFL – Lausanne, CH)  
February 4–9, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24062>

**24071 – Safety Assurance for Autonomous Mobility**

Jyotirmoy Deshmukh (USC – Los Angeles, US), Bettina Könighofer (TU Graz, AT), Dejan Nickovic (AIT – Austrian Institute of Technology – Wien, AT), Ufuk Topcu (University of Texas – Austin, US)  
February 11–16, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24071>

**24072 – Triangulations in Geometry and Topology**

Maike Buchin (Ruhr-Universität Bochum, DE), Jean Cardinal (ULB – Brussels, BE), Arnaud de Mesmay (CNRS, Gustave Eiffel University – Marne-la-Vallée, FR), Jonathan Spreer (University of Sydney, AU)  
February 11–16, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24072>

**24081 – Computational Approaches to Strategy and Tactics in Sports**

Ulf Brefeld (Universität Lüneburg, DE), Jesse Davis (KU Leuven, BE), Laura de Jong (Deakin University – Melbourne, AU), Stephanie Kovalchik (Zelus Analytics – Austin, US)  
February 18–23, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24081>

**24082 – AI for Social Good**

Claudia Clopath (Imperial College London, GB), Ruben De Winne (Oxfam Novib – The Hague, NL), Mohammad Emtiyaz Khan (RIKEN – Tokyo, JP), Jacopo Margutti (510 / Netherlands Red Cross – The Hague, NL)  
February 18–23, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24082>

**24091 – Reflections on Pandemic Visualization**

Daniel Archambault (Newcastle University, GB), Fintan McGee (Luxembourg Inst. of Science & Technology, LU), Simone Scheithauer (Universitätsmedizin Göttingen, DE), Tatiana von Landesberger (Universität Köln, DE)

February 25 to March 1, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24091>

**24092 – Applied and Combinatorial Topology**

Pawel Dlotko (Polish Academy of Science, PL), Dmitry Feichtner-Kozlov (Universität Bremen, DE), Anastasios Stefanou (Universität Bremen, DE), Yusu Wang (University of California, San Diego – La Jolla, US)

February 25 to March 1, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24092>

**24101 – Robust Query Processing in the Cloud**

Anastasia Ailamaki (EPFL – Lausanne, CH), Goetz Graefe (Google – Madison, US), Allison Lee (Snowflake – San Mateo, US), Caetano Sauer (Salesforce – München, DE)

March 3–8, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24101>

**24102 – Shapes in Graph Data: Theory and Implementation**

Shqiponja Ahmetaj (TU Wien, AT), Slawomir Staworko (relationalAI – Berkeley, US), Jan Van den Bussche (Hasselt University, BE)

March 3–8, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24102>

**24111 – Logics for Dependence and Independence: Expressivity and Complexity**

Juha Kontinen (University of Helsinki, FI), Jonni Virtema (University of Sheffield, GB), Heribert Vollmer (Leibniz Universität Hannover, DE), Fan Yang (Utrecht University, NL)

March 10–15, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24111>

**24112 – EU Cyber Resilience Act: Socio-Technical and Research Challenges**

Mila Dalla Preda (University of Verona, IT), Serge Egelman (ICSI – Berkeley, US), Anna Maria Mandalari (University College London, GB), Narseo Vallina-Rodriguez (IMDEA Networks Institute – Madrid, ES)

March 10–13, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24112>

**24121 – Trustworthiness and Responsibility in AI – Causality, Learning, and Verification**

Vaishak Belle (University of Edinburgh, GB), Hana Chockler (King's College London, GB), Sriraam Natarajan (University of Texas at Dallas – Richardson, US), Shannon Vallor (University of Edinburgh, GB), Kush R. Varshney (IBM Research – Yorktown Heights, US), Joost Vennekens (KU Leuven, BE)

March 17–22, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24121>

**24122 – Low-Dimensional Embeddings of High-Dimensional Data: Algorithms and Applications**

Fred Hamprecht (Universität Heidelberg, DE), Dmitry Kobak (Universität Tübingen, DE), Smita Krishnaswamy (Yale University – New Haven, US), Gal Mishne (University of California, San Diego – La Jolla, US)

March 17–22, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24122>

**24141 – Network Calculus**

Steffen Bondorf (Ruhr-Universität Bochum, DE), Anne Bouillard (Huawei Technologies – Boulogne-Billancourt, FR), Markus Fidler (Leibniz Universität Hannover, DE), Jörg Liebeherr (University of Toronto, CA)

April 1–4, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24141>

**24151 – Methods and Tools for the Engineering and Assurance of Safe Autonomous Systems**

Ignacio J. Alvarez (Intel – Hillsboro, US), Philip Koopman (Carnegie Mellon University – Pittsburgh, US), Mario Trapp (TU München, DE), Elena Troubitsyna (KTH Royal Institute of Technology – Stockholm, SE)

April 7–12, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24151>

**24161 – Research Software Engineering: Bridging Knowledge Gaps**

Stephan Druskat (German Aerospace Center (DLR), Berlin, DE), Lars Grunske (HU Berlin, DE), Caroline Jay (University of Manchester, GB), Daniel S. Katz (University of Illinois Urbana-Champaign, US)

April 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24161>

**24162 – Hardware Support for Cloud Database Systems in the Post-Moore's Law Era**

David F. Bacon (Google – New York, US), Carsten Binnig (TU Darmstadt, DE), David A. Patterson (University of California – Berkeley, US), Margo Seltzer (University of British Columbia – Vancouver, CA)

April 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24162>

**24171 – Automated Synthesis: Functional, Reactive and Beyond**

S. Akshay (Indian Institute of Technology Bombay – Mumbai, IN), Bernd Finkbeiner (CISPA – Saarbrücken, DE), Kuldeep S. Meel (University of Toronto, CA), Ruzica Piskac (Yale University – New Haven, US)

April 21–26, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24171>

**24172 – Code Search**

Satish Chandra (Google – Mountain View, US), Michael Pradel (Universität Stuttgart, DE), Kathryn T. Stolee (North Carolina State University – Raleigh, US)

April 21–24, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24172>

**24181 – Computational Metabolomics: Towards Molecules, Models, and their Meaning**

Timothy M. D. Ebbels (Imperial College London, GB), Soha Hassoun (Tufts University – Medford, US), Ewy A. Mathé (National Institutes of Health – Bethesda, US), Justin J. J. van der Hoof (Wageningen University & Research, NL)

April 28 to May 3, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24181>

**24182 – Resilience and Antifragility of Autonomous Systems**

Simon Burton (Gerlingen, DE), Radu Calinescu (University of York, GB), Raffaella Mirandola (KIT – Karlsruher Institut für Technologie, DE)

April 28 to May 3, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24182>

**24192 – Generalization by People and Machines**

Barbara Hammer (Universität Bielefeld, DE), Filip Ilievski (VU Amsterdam, NL), Sascha Saralajew (NEC Laboratories Europe – Heidelberg, DE), Frank van Harmelen (VU Amsterdam, NL)

May 5–8, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24192>

**24201 – Discrete Algorithms on Modern and Emerging Compute Infrastructure**

Kathrin Hanauer (Universität Wien, AT), Uwe Naumann (RWTH Aachen, DE), Alex Pothen (Purdue University – West Lafayette, US), Robert Schreiber (Cerebras Systems – Palo Alto, US)

May 12–17, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24201>

**24202 – Causal Inference for Spatial Data Analytics**

Fernando Perez Cruz (ETH Zürich, CH), Jakob Runge (DLR – Jena, DE & TU Berlin, DE), Martin Tomko (University of Melbourne – Carlton, AU), Yanan Xin (ETH Zürich, CH)

May 12–17, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24202>

**24211 – Evaluation Perspectives of Recommender Systems: Driving Research and Education**

Christine Bauer (Paris Lodron Universität Salzburg, AT), Alan Said (University of Gothenburg, SE), Eva Zangerle (Universität Innsbruck, AT)

May 20–24, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24211>

**24212 – Classical-Quantum Synergies in the Theory and Practice of Quantum Error Correction**

Carmen G. Almudéver (Technical University of Valencia, ES), Leonid Pryadko (University of California at Riverside, US), Valentin Savin (CEA – Grenoble, FR), Bane Vasic (University of Arizona – Tucson, US)

May 20–23, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24212>

**24231 – Stochastic Games**

Nathanaël Fijalkow (CNRS – Talence, FR), Jan Kretinsky (Masaryk University – Brno, CZ), Ann Nowé (Free University of Brussels, BE), Uri Zwick (Tel Aviv University, IL)

June 2–7, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24231>

**24232 – Designing Computers' Control Over Our Bodies**

Nadia Bianchi-Berthouze (University College London, GB), Mar Gonzalez-Franco (Google – Seattle, US), Florian 'Floyd' Mueller (Monash University – Clayton, AU), Misha Sra (University of California – Santa Barbara, US)

June 2–7, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24232>

**24241 – Geometric modeling: Challenges for Additive Manufacturing, Design and Analysis**

Tor Dokken (SINTEF – Oslo, NO), Xiaohong Jia (Chinese Academy of Sciences, CN), Géraldine Morin (IRIT – University of Toulouse, FR), Elissa Ross (Metafold 3D – Toronto, CA)

June 9–14, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24241>

**24242 – Computational Analysis and Simulation of the Human Voice**

Peter Birkholz (TU Dresden, DE), Oriol Guasch Fortuny (Ramon Llull University – Barcelona, ES), Nathalie Henrich Bernardoni (University Grenoble Alpes, FR), Sten Ternström (KTH Royal Institute of Technology – Stockholm, SE)

June 9–14, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24242>

**24251 – Teaching Support Systems for Formal Foundations of Computer Science**

Tiffany Barnes (North Carolina State University – Raleigh, US), Jan Vahrenhold (Universität Münster, DE), Thomas Zeume (Ruhr-Universität Bochum, DE)

June 16–21, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24251>

**24261 – Computational Creativity for Game Development**

Duygu Cakmak (Creative Assembly – Horsham, GB), Setareh Maghsudi (Ruhr-Universität Bochum, DE), Diego Perez Liebana (Queen Mary University of London, GB), Pieter Spronck (Tilburg University, NL)

June 23–28, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24261>

**24271 – Theory of Randomized Optimization Heuristics**

Anne Auger (INRIA Saclay – Palaiseau, FR), Tobias Glasmachers (Ruhr-Universität Bochum, DE), Martin S. Krejca (Ecole Polytechnique – Palaiseau, FR), Johannes Lengler (ETH Zürich, CH)

June 30 to July 5, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24271>

**24272 – A Game of Shadows: Effective Mastery Learning in the Age of Ubiquitous AI**

Nick Falkner (University of Adelaide, AU), Juho Leinonen (Aalto University, FI), Miranda C. Parker (San Diego State University, US), Andrew Petersen (University of Toronto Mississauga, CA), Claudia Szabo (University of Adelaide, AU)

June 30 to July 5, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24272>

**24281 – Dynamic Traffic Models in Transportation Science**

José R. Correa (University of Chile – Santiago de Chile, CL), Carolina Osorio (HEC – Montréal, CA & Google Research – Mountain View, US), Laura Vargas Koch (Universität Bonn, DE), David Watling (University of Leeds, GB)

July 7–12, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24281>

**24282 – Automated Machine Learning for Computational Mechanics**

Hyunsun Alicia Kim (UC – San Diego, US), Lars Kotthoff (University of Wyoming – Laramie, US), Marius Lindauer (Leibniz Universität Hannover, DE), Elena Raponi (Leiden University, NL)

July 7–12, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24282>

**24291 – Programmable Host Networking**

Gianni Antichi (Polytechnic University of Milan, IT), Katerina Argyraki (EPFL – Lausanne, CH), Aurojit Panda (New York University, US), Justine Sherry (Carnegie Mellon University – Pittsburgh, US)

July 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24291>

**24292 – Improving Trust between Humans and Software Robots in Robotic Process Automation**

Adela del Río Ortega (University of Sevilla, ES), Andrea Marrella (Sapienza University of Rome, IT), Hajo A. Reijers (Utrecht University, NL), Adriana Wilde (University of Southampton, GB)

July 14–19, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24292>

**24301 – Art, Visual Illusions, and Data Visualization**

Claus-Christian Carbon (Universität Bamberg, DE), Christophe Hurter (ENAC – Toulouse, FR), Mauro Martino (MIT-IBM Watson AI Lab – Cambridge, US), Bernice E. Rogowitz (Visual Perspectives – New York, US)

July 21–26, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24301>

**24302 – Learning with Music Signals: Technology Meets Education**

Cynthia Liem (TU Delft, NL), Brian McFee (New York University, US), Meinard Müller (Universität Erlangen-Nürnberg, DE)

July 21–26, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24302>

**24311 – Resource-Efficient Machine Learning**

Oana Balmau (McGill University – Montréal, CA), Matthias Böhm (TU Berlin, DE), Ana Klimovic (ETH Zürich, CH), Peter R. Pietzuch (Imperial College London, GB), Pinar Tözün (IT University of Copenhagen, DK)

July 28 to August 2, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24311>

**24312 – Security and Privacy of Current and Emerging IoT Devices and Systems**

Bruno Crispo (University of Trento, IT), Alexandra Dmitrienko (Universität Würzburg, DE), Gene Tsudik (University of California – Irvine, US), Wenyuan Xu (Zhejiang University – Hangzhou, CN)

July 28 to August 2, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24312>

**24341 – Proof Representations: From Theory to Applications**

Carlos Areces (National University – Córdoba, AR), Anupam Das (University of Birmingham, GB), Elaine Pimentel (University College London, GB), Lutz Straßburger (INRIA Saclay – Île-de-France, FR)

August 18–23, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24341>

**24342 – Leveraging AI for Management Decision-Making**

Stefan Feuerriegel (Ludwig-Maximilians-Universität München, DE), Foster Provost (New York University, US), Galit Shmueli (National Tsing Hua University – Hsinchu, TW)

August 18–21, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24342>

**24351 – Power, Energy, and Carbon-Aware Computing on Heterogeneous Systems (PEACHES)**

Kerstin I. Eder (University of Bristol, GB), Timo Höning (Ruhr-Universität Bochum, DE), Maja Hanne Kirkeby (Roskilde University, DK), Daniel Mosse (University of Pittsburgh, US), Max Plauth (UltiHash – Berlin, DE)

August 25–30, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24351>

**24361 – Artificial Intelligence and Formal Methods Join Forces for Reliable Autonomy**

Nils Jansen (Ruhr-Universität Bochum, DE), Mykel Kochenderfer (Stanford University, US), Jan Kretinsky (Masaryk University – Brno, CZ), Jana Tumova (KTH Royal Institute of Technology – Stockholm, SE)

September 1–6, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24361>

**24362 – Next-Generation Secure Distributed Computing**

Christian Cachin (Universität Bern, CH), Aniket Kate (Purdue University – West Lafayette, US & Supra Research – West Lafayette, US), Julian Loss (CISPA – Saarbrücken, DE), Kartik Nayak (Duke University – Durham, US)

September 1–6, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24362>

**24371 – Extended Reality Accessibility**

Gerd Bruder (University of Central Florida – Orlando, US), Thies Pfeiffer (Hochschule Emden/Leer, DE), Jeanine Stefanucci (University of Utah, US)

September 8–13, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24371>

**24372 – Explainable AI for Sequential Decision Making**

Hendrik Baier (TU Eindhoven, NL), Mark T. Keane (University College Dublin, IE), Sarath Sreedharan (Colorado State University – Fort Collins, US), Silvia Tulli (Sorbonne University – Paris, FR), Abhinav Verma (Pennsylvania State University – University Park, US)

September 8–11, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24372>

**24381 – Algebraic and Analytic Methods in Computational Complexity**

Markus Bläser (Universität des Saarlandes – Saarbrücken, DE), Shubhangi Saraf (University of Toronto, CA), Ronen Shaltiel (University of Haifa, IL), Jacobo Torán (Universität Ulm, DE)

September 15–20, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24381>

**24391 – Statistical and Probabilistic Methods in Algorithmic Data Analysis**

Aristides Gionis (KTH Royal Institute of Technology – Stockholm, SE), Matteo Riondato (Amherst College, US), Eli Upfal (Brown University – Providence, US)

September 22–27, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24391>

**24401 – Fair Division: Algorithms, Solution Concepts, and Applications**

Evangelos Markakis (Athens University of Economics and Business, GR), Ruta Mehta (University of Illinois – Urbana-Champaign, US), Yair Zick (University of Massachusetts Amherst, US)

September 29 to October 4, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24401>

**24402 – Greening Networking: Toward a Net Zero Internet**

Alexander Clemm (Los Gatos, US), Dirk Kutscher (HKUST – Guangzhou, CN), Michael Welzl (University of Oslo, NO), Cedric Westphal (Futurewei – Santa Clara, US), Noa Zilberman (University of Oxford, GB)

September 29 to October 2, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24402>

**24411 – New Tools in Parameterized Complexity: Paths, Cuts, and Decomposition**

Fedor V. Fomin (University of Bergen, NO), Dániel Marx (CISPA – Saarbrücken, DE), Saket Saurabh (The Institute of Mathematical Sciences – Chennai, IN), Roohani Sharma (MPI für Informatik – Saarbrücken, DE)

October 6–11, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24411>

**24421 – SAT and Interactions**

Olaf Beyersdorff (Friedrich-Schiller-Universität Jena, DE), Laura Kovács (TU Wien, AT), Meena Mahajan (The Institute of Mathematical Sciences – Chennai, IN), Martina Seidl (Johannes Kepler Universität Linz, AT)

October 13–18, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24421>

**24431 – Automated Programming and Program Repair**

Claire Le Goues (Carnegie Mellon University – Pittsburgh, US), Michael Pradel (Universität Stuttgart, DE), Abhik Roychoudhury (National University of Singapore, SG), Shin Hwei Tan (Concordia University – Montreal, CA)

October 20–25, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24431>

**24432 – Behavioural Metrics and Quantitative Logics**

Barbara König (Universität Duisburg-Essen, DE), Radu Mardare (University of Strathclyde – Glasgow, GB), Prakash Panangaden (McGill University – Montréal, CA), Jurriaan Rot (Radboud University Nijmegen, NL)

October 20–25, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24432>

**24441 – Machine Learning Augmented Algorithms for Combinatorial Optimization Problems**

Deepak Ajwani (University College Dublin, IE), Bistra Dilkina (USC – Los Angeles, US), Tias Guns (KU Leuven, BE), Ulrich Carsten Meyer (Goethe University – Frankfurt am Main, DE)

October 27–31, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24441>

**24451 – Machine Learning for Protein-Protein and Protein-Ligand Interactions**

Anne-Florence Bitbol (EPFL – Lausanne, CH), Jennifer Listgarten (University of California – Berkeley, US), Tomas Pluskal (IOCB – Prague, CZ)

November 3–8, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24451>

**24461 – Rethinking the Role of Bayesianism in the Age of Modern AI**

Vincent Fortuin (Helmholtz AI – Neuherberg, DE), Zoubin Ghahramani (Google – Mountain View, US), Mohammad Emtiyaz Khan (RIKEN – Tokyo, JP), Mark van der Wilk (University of Oxford, GB)

November 10–15, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24461>

**24462 – Research Infrastructures and Tools for Collaborative Networked Systems Research**

Georg Carle (TU München – Garching, DE), Serge Fdida (Sorbonne University – Paris, FR), Kate Keahey (Argonne National Laboratory, US), Henning Schulzrinne (Columbia University – New York, US)

November 10–13, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24462>

**24471 – Graph Algorithms: Distributed Meets Dynamic**

Keren Censor-Hillel (Technion – Haifa, IL), Yasamin Nazari (VU Amsterdam, NL), Eva Rotenberg (Technical University of Denmark – Lyngby, DK), Thatchaphol Saranurak (University of Michigan – Ann Arbor, US)

November 17–22, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24471>

**24472 – Regular Expressions: Matching and Indexing**

Inge Li Gørtz (Technical University of Denmark – Lyngby, DK), Sebastian Maneth (Universität Bremen, DE), Gonzalo Navarro (University of Chile – Santiago de Chile, CL), Nicola Prezza (University of Venice, IT)

November 17–22, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24472>

**24491 – Deep Learning for RNA Regulation and Multidimensional Transcriptomics**

Annalisa Marsico (Helmholtz Zentrum München, DE), Uwe Ohler (Max-Delbrück-Centrum – Berlin, DE), Igor Ulitsky (Weizmann Institute – Rehovot, IL), Gene Yeo (University of California – San Diego, US), Kathi Zarnack (Universität Würzburg, DE)

December 1–6, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24491>

**24511 – Coding Theory and Algorithms for Emerging Technologies in Synthetic Biology**

R. B. (TU München, DE), Olgica Milenkovic (University of Illinois – Urbana Champaign, US), Zohar Yakhini (Reichman University – Herzliya, IL), Yonatan Yehezkeally (TU München, DE)

December 15–20, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24511>

**24512 – Quantum Software Engineering**

Shaukat Ali (Simula Research Laboratory – Oslo, NO), Johanna Barzen (Universität Stuttgart, DE), Andrea Delgado (Oak Ridge National Laboratory, US), Hausi A. Müller (University of Victoria, CA), Juan Manuel Murillo (University of Extremadura – Cáceres, ES)

December 15–20, 2024 | Dagstuhl Seminar | <https://www.dagstuhl.de/24512>

## Dagstuhl-Perspektiven-Workshops

14.2

## Dagstuhl Perspectives Workshops

### 24352 – Conversational Agents: A Framework for Evaluation (CAFE)

Christine Bauer (Paris Lodron Universität Salzburg, AT), Li Chen (Hong Kong Baptist University, HK), Nicola Ferro (University of Padova, IT), Norbert Fuhr (Universität Duisburg-Essen, DE)

August 25–30, 2024 | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/24352>

### 24452 – Reframing Technical Debt

Paris Avgeriou (University of Groningen, NL), Zadia Codabux (University of Saskatchewan – Saskatoon, CA), Heiko Koziol (ABB – Mannheim, DE), Ipek Ozkaya (Carnegie Mellon University – Pittsburgh, US)

November 3–8, 2024 | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/24452>

### 24492 – Human in the Loop Learning through Grounded Interaction in Games

Raffaella Bernardi (University of Trento, IT), Julia Hockenmaier (University of Illinois – Urbana-Champaign, US), Udo Kruschwitz (Universität Regensburg, DE), Massimo Poesio (Queen Mary University of London, GB & Utrecht University, NL)

December 1–6, 2024 | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/24492>

## GI-Dagstuhl-Seminare

14.3

## GI-Dagstuhl Seminars

There were no such meetings in 2024.

## Lehrveranstaltungen

14.4

## Educational Events

### 24254 – Summer School “Data Management Techniques”

Goetz Graefe (Google – Madison, US)

June 16–20, 2024 | Educational Event | <https://www.dagstuhl.de/24254>

### 24383 – IGAFIT Workshop for Algorithms Postdocs in Europe (AlgPiE by IGAFIT 2024)

Artur Czumaj (University of Warwick – Coventry, GB), Aditi Dudeja (Paris Lodron Universität Salzburg, AT), Nick Fischer (INSAIT – Sofia, BG), Laura Vargas Koch (Universität Bonn, DE), Isabella Ziccardi (Bocconi University – Milan, IT)

September 15–20, 2024 | Educational Event | <https://www.dagstuhl.de/24383>

### 24393 – Institute 2024 – Artificial and Human Intelligence

Mehul Bhatt (Örebro University, SE), Paul Hemeren (University of Skövde, SE), Vasiliki Kondyli (Jagiellonian University – Kraków, PL), Árni Kristjánsson (University of Iceland – Reykjavik, IS), Jakob Suchan (Constructor University – Bremen, DE)

September 22–27, 2024 | Educational Event | <https://www.dagstuhl.de/24393>

### 24503 – Lehrerfortbildung in Informatik

Sebastian Connette (Bildungscampus Saarland, DE), Michael Gerke (Schloss Dagstuhl – Saarbrücken, DE), Hannes Heusel (Eduard-Spranger Gymnasium Landau, DE & Pädagogisches Landesinstitut Rheinland-Pfalz – Speyer, DE)

December 11–13, 2024 | Educational Event | <https://www.dagstuhl.de/24503>

## Forschungsgruppentreffen

14.5

## Research Group Meetings

### 24013 – Unknown Data: Mining and Consolidating Research Dataset Metadata on the Web

Benedikt Maria Beckermann (Schloss Dagstuhl – Trier, DE)

January 3–5, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24013>

### 24014 – Ontologie, Linguistik, Terminologie, Logik

Johannes Busse (HAW Landshut, DE), Anatol Reibold (Universität Leipzig, DE & Cardiso GmbH – Frankfurt/Main, DE)

January 2–5, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24014>

**24016 – Semantics of Software Systems**

Andreas Zeller (CISPA – Saarbrücken, DE)

January 2–5, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24016>**24023 – Workshop on Historical Cryptology**

Bernhard Esslinger (Universität Siegen, DE), Beáta Megyesi (Stockholm University, SE)

January 8–12, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24023>**24024 – Open Machine Learning 2024 Winter Workshop**

Pieter Gijsbers (TU Eindhoven, NL), Joaquin Vanschoren (TU Eindhoven, NL)

January 7–12, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24024>**24133 – GIBU 2024: GI-Beirat der Universitätsprofessor\*innen**

Lars Grunske (HU Berlin, DE)

March 24–26, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24133>**24134 – Towards A Unified Interface For Modern Probabilistic Model Checking Tools**

Sebastian Junges (Radboud University Nijmegen, NL), David Parker (University of Oxford, GB)

March 24–27, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24134>**24143 – Generative AI for Knowledge Engineering**

Achim Rettinger (Universität Trier, DE), Steffen Thoma (FZI – Karlsruhe, DE)

April 2–3, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24143>**24144 – Paradigms of ML: Links between human-, model- and data-centric learning**

Diego Botache (Universität Kassel, DE), Kristina Dingel (Universität Kassel, DE), David Meier (Helmholtz-Zentrum – Berlin, DE)

April 1–5, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24144>**24145 – Klausurtagung Telematik Karlsruhe (KIT)**

Michael König (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)

April 3–5, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24145>**24153 – AI for Realistic Science: Advances and Challenges in Large-Scale Experimental Facilities**

Ian T. Foster (Argonne National Laboratory – Lemont, US), Tony Hey (Science and Technology Facilities Council – Didcot, GB), James A Sethian (University of California – Berkeley, US), Arjun Shankar (Oak Ridge National Laboratory, US), Jeyan Thiyagalingam (Rutherford Appleton Lab. – Didcot, GB)

April 7–12, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24153>**24173 – bwInfoSec Retreat**

Marc Herbstritt (Universität Freiburg, DE), Vincent Heuveline (Universität Heidelberg, DE)

April 24–26, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24173>**24193 – Kolloquium zum GI-Dissertationspreis 2023**

Rüdiger Reischuk (Universität zu Lübeck, DE)

May 5–8, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24193>**24213 – 2nd Overall Project Meeting DFG FOR 5359**

Marius Kloft (RPTU Kaiserslautern-Landau, DE), Steffen Reithermann (RPTU Kaiserslautern-Landau, DE)

May 23–24, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24213>**24223 – Joint Meeting of the German Research Training Groups / Gemeinsamer Workshop der Graduiertenkollegs**

Helen Bolke-Hermanns (RWTH Aachen, DE), Jan-Christoph Kassing (RWTH Aachen, DE), Michael Schaub (RWTH Aachen, DE), Michael Scholkemper (RWTH Aachen, DE)

May 26–29, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24223>**24259 – Demands of Programming-Language Agnostic Test Case Reduction**

Michael Philippsen (Universität Erlangen-Nürnberg, DE)

June 16–28, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24259>**24263 – Klausurtagung 2024 Forschungsgruppe Internet Architecture (INET)**

Iris Wagner (MPI für Informatik – Saarbrücken, DE)

June 23–26, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24263>

**24264 – Applied Machine Intelligence 2024**

Johannes Busse (HAW Landshut, DE), Thomas Hoppe (Fraunhofer FOKUS – Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE)

June 26–28, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24264>

**24343 – Simulation of Quantum Matter with a Focus on Long-Range Interactions**

Andreas Buchheit (Universität des Saarlandes – Saarbrücken, DE), Benedikt Fauseweh (TU Dortmund, DE)

August 21–23, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24343>

**24399 – Highly reliable decision-making under uncertainty**

Bismark Singh (University of Southampton, GB)

September 22–28, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24399>

**24403 – MCO Book Project: Computational Methods for Multiobjective Optimization**

Kathrin Klamroth (Universität Wuppertal, DE), Boris Naujoks (TH Köln, DE)

October 2–4, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24403>

**24404 – Interdisciplinary Perspectives on Information**

Jens Allwood (University of Gothenburg – Mölndal, SE), Paavo Pylkkänen (University of Helsinki, FI), Michael Richter (Universität Leipzig, DE)

October 2–4, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24404>

**24413 – Aligning Technology Architectures with Cross-Domain Metadata Models**

Michelle Edwards (University of Guelph, CA), Arofan Gregory (DDI Alliance / CODATA – Jaffrey, New Hampshire, US), Simon Hodson (CODATA – Paris, FR), Steven McEachern (UK Data Service – Colchester, GB), Hilde Orten (Sikt – Bergen, NO), Joachim Wackerow (Ludwigshafen, DE)

October 6–11, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24413>

**24423 – Evaluating and Refining Cross-Domain Metadata Exchange Frameworks**

Michelle Edwards (University of Guelph, CA), Arofan Gregory (DDI Alliance / CODATA – Jaffrey, New Hampshire, US), Simon Hodson (CODATA – Paris, FR), Steven McEachern (UK Data Service – Colchester, GB), Hilde Orten (Sikt – Bergen, NO), Joachim Wackerow (Ludwigshafen, DE)

October 13–18, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24423>

**24443 – Software Engineering Forschungsmethoden Training**

Sven Apel (Universität des Saarlandes – Saarbrücken, DE), Sebastian Baltes (Universität Bayreuth, DE), Ben Hermann (TU Dortmund, DE)

October 27–31, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24443>

**24463 – Explainable Decision-Making**

Wolfgang Maaß (UdS – Saarbrücken, DE & DFKI – Saarbrücken, DE)

November 13–15, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24463>

**24464 – GPU-Accelerated Temporal Graph Learning with pathpyG**

Ingo Scholtes (Universität Würzburg, DE)

November 13–15, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24464>

**24483 – A long-term strategy for NFDI for Data Science and Artificial Intelligence**

Christine Hennig (Fraunhofer FOKUS – Berlin, DE), Sonja Schimmler (Fraunhofer FOKUS – Berlin, DE)

November 28–29, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24483>

**24484 – Turtle Hack: Programmieren statt programmiert werden**

Jacqueline Staub (Universität Trier, DE)

November 24–29, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24484>

**24485 – Workshop of the ELLIS Natural Language Processing Program**

Iryna Gurevych (TU Darmstadt, DE), André F. T. Martins (IST – Lisbon, PT), Edoardo Maria Ponti (University of Edinburgh, GB), Ivan Titov (University of Edinburgh, GB)

November 24–27, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24485>

**24486 – Comparing Continuous Optimizers (COCO): 2nd Code and Documentation Sprint**

Dimo Brockhoff (Inria Saclay Centre and IP Paris, FR), Nikolaus Hansen (INRIA Saclay – Palaiseau, FR), Olaf Mersmann (HS Bund f. Öffentl. Verwaltung – Brühl, DE), Tea Tusar (Jozef Stefan Institute – Ljubljana, SI)

November 24–29, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24486>

**24487 – Research Retreat Cologne Information Retrieval Group**

Timo Breuer (TH Köln, DE), Philipp Schaer (TH Köln, DE)

November 24–27, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24487>**24504 – Redaktionssitzung inf-schule.de**

Bernd Fröhlich (Nikolaus-von-Kues-Gymnasium – Bernkastel-Kues, DE), Hannes Heusel (Eduard-Spranger Gymnasium Landau, DE &amp; Pädagogisches Landesinstitut Rheinland-Pfalz – Speyer, DE)

December 10–11, 2024 | Research Group Meeting | <https://www.dagstuhl.de/24504>







© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH  
Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2024 | ISSN 2199-1995  
<https://www.dagstuhl.de>