DAGSTUHL
**REPORTS**

**Volume 15, Issue 5, May 2025**

Adaptive and Scalable Data Structures (Dagstuhl Seminar 25191)
*Michael A. Bender, John Iacono, László Kozma, and Eva Rotenberg* .............. 1

AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation (Dagstuhl Seminar 25192)
*Giuseppe De Giacomo, Marlon Dumas, Fabiana Fournier, Timotheus Kampik, and Lior Limonad* .............. 21

Computational Geometry (Dagstuhl Seminar 25201)
*Maarten Löffler, Eunjin Oh, Jeff M. Phillips, and Alexandra Weinberger* .............. 64

Generative Models for 3D Vision (Dagstuhl Seminar 25202)
*Laura Neschen, Bernhard Egger, Adam Kortylewski, William Smith, and Stefanie Wuhrer* .............. 96

The Constraint Satisfaction Problem: Complexity and Approximability (Dagstuhl Seminar 25211)
*Manuel Bodirsky, Venkatesan Guruswami, Dániel Marx, Stanislav Živný, and Žaneta Semanišinová* .............. 114

Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs (Dagstuhl Seminar 25212)
*Sujoy Bhore, Jie Gao, Hung Le, Csaba D. Tóth, and Lazar Milenković* .............. 134

**ISSN 2192-5283**

*Published online and open access by*
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at
https://www.dagstuhl.de/dagpub/2192-5283

*Publication date*
December, 2025

*Bibliographic information published by the Deutsche Nationalbibliothek*
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at `https://dnb.d-nb.de`.



Digital Object Identifier: 10.4230/DagRep.15.5.i

*Aims and Scope*
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:
- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.



*Contact*
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
https://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 25191

# Adaptive and Scalable Data Structures

**Michael A. Bender*[1], John Iacono*[2], László Kozma*[3], Eva Rotenberg*[4], and Justin Dallant†[5]**

1 **Stony Brook University, US.** bender@cs.stonybrook.edu
2 **ULB – Brussels, BE.** john.iacono@ulb.be
3 **FU Berlin, DE.** lkozma@zedat.fu-berlin.de
4 **Technical University of Denmark – Lyngby, DK.** eva@rotenberg.dk
5 **ULB – Brussels, BE.** justindallant@gmail.com

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 25191 "Adaptive and Scalable Data Structures". Data structures govern the organization and manipulation of data in computing systems across a broad range of applications. The efficiency and scalability of data structures has profound implications, motivating continued research on the entire spectrum from theoretical to practical. As the size and complexity of data sets increases and as the underlying computing infrastructure changes, data structures need to be continually redesigned with scalability in mind. Classical data structures also need reevaluation to better fit the requirements of modern applications. Adaptivity offers a way to design data structures that automatically take advantage of features of the underlying hardware, specific structure and biases in their usage, or side-information, and the limits of data structure adaptivity pose deep research questions. The goal of this seminar was to reflect on these complementary aspects of data structure research and to identify promising research questions. The program provides a snapshot of the current state of research and establishes possible future directions for the field.

**Seminar** May 4–9, 2025 – http://www.dagstuhl.de/25191
**2012 ACM Subject Classification** Theory of computation → Data structures design and analysis; Theory of computation → Design and analysis of algorithms; Theory of computation → Parallel algorithms
**Keywords and phrases** Data structures, Algorithms, Big data, Computational models
**Digital Object Identifier** 10.4230/DagRep.15.5.1

## 1 Executive summary

*Michael A. Bender (Stony Brook University, US)*
*John Iacono (ULB – Brussels, BE)*
*László Kozma (FU Berlin, DE)*
*Eva Rotenberg (Technical University of Denmark – Lyngby, DK)*



### About the seminar

Data structures are fundamental building blocks of computing systems and are deployed in applications of increasing size and sophistication. The design and analysis of *scalable* data structures has been a central goal of computer science from the beginnings of the field. Yet, the study of data structures remains as active as ever, as it reflects changes in the underlying computational infrastructure, as well as the evolving needs of applications.

* Editor / Organizer
† Editorial Assistant / Collector


Adaptive and Scalable Data Structures, *Dagstuhl Reports*, Vol. 15, Issue 5, pp. 1–20
Editors: Michael A. Bender, John Iacono, László Kozma, and Eva Rotenberg
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Adaptivity is a multi-faceted goal that can refer to: data structures seamlessly taking advantage of resources offered by the underlying hardware, or restrictions thereof; the usage of application-specific insights or side-information to improve data structure efficiency; the extent to which the data generating process can *adversarially adapt* to the data structure.

Some of these aspects pose long-standing open questions that continue to inspire research (e.g., the *dynamic optimality conjecture* asks whether binary search trees can optimally *adapt* to their usage pattern). More recent directions of adaptivity to side-information, e.g., from machine learning advice, pose fresh research questions, and can also relate to past work on data structures that are locality-adaptive, robust to errors, or that make use of imprecise or unreliable comparisons. A refined notion related to adversarial adaptivity is that of *history-independence*: the question of whether data structures can avoid leaking information about their past sequence of operations. Besides the clear motivation from the point of view of security and privacy, this concept has recently played a key role in designing state-of-the-art data structures for the prototypical list labeling problem. Participants in the seminar have further explored to what extent history independence can be reconciled with optimal efficiency in fundamental data structures such as priority queues.

This seminar was the 16th in a series of loosely related Dagstuhl Seminars on data structures. It brought together a diverse group of researchers with complementary expertise and varying levels of seniority, to illuminate different aspects of data structure scalability and adaptivity. We also aimed to include researchers who are experts in adjacent, applied areas (data structures in computational geometry, data structures for strings, etc.), to inspire and inform each other and to identify relevant research directions.

The organizers of the seminar put a strong emphasis on discussion and collaboration. Participants were asked to give short (10-15 minutes) or medium (20-25 minutes) talks. Most participants spoke, and the format allowed everyone who wished to speak to do so. We explicitly encouraged forward-thinking talks that include open problems, interesting new directions or challenges, focusing on informal, open-ended discussion, rather than polished presentations of past work. In the program we made a deliberate effort to leave space for loosely structured discussion and collaboration, reserving most of the afternoons for this purpose, and having brief status-report sessions where we all met to discuss what people had been discussing.

## Topics

The presentations and discussions covered a broad range of topics in classical data structures and diverse applications, as well as new directions inspired by various aspects of scalability and adaptivity.

Models of computation were an important theme of the seminar. Afshani (Section 4.10) discussed possible separations concerning I/O operations and total work, Bille (Section 4.1) discussed data structures in the ultra-wide word RAM model, and Meyer (Section 4.2) talked about unstructured parallel I/O on SSDs in the context of graph algorithms.

Parallelism and distributed computing were also the focus of further talks: King (Section 4.4) presented work on a problem in distributed setting with many Byzantine participants, and Chechik (Section 5.8) discussed the maximum independent set problem in the congested clique model.

Several talks focused on questions (old and new) on graph algorithms: Wein (Section 4.8) talked about a new concept of a "DAG cover" extending classical notions of tree covers, Liu (Section 4.9) discussed a batch version of the dynamic connectivity problem, van der Hoog

(Section 5.2) raised a question about greedy chain decomposition on DAGs, Bilò (Section 4.15) revisited the problem of single-source shortest p-disjoint paths, Gudmundsson (Section 4.16) talked about approximate distance oracles for a special class of sparse graphs, and Fineman (Section 4.17) talked about open questions about incremental topological sorting.

Relating to online problems, Agrawal (Section 5.6) discussed results and open questions concerning a scheduling problem with DAG constraints, and Munro (Section 4.11) talked about a (paired) variant of the paging problem.

On compression and compact data structures, Navarro talked about dynamic bitvectors (Section 4.5) and on a graph problem related to grammar-based compression (Section 5.9), and Gawrychowski (Section 5.4) discussed open problems related to Lempel-Ziv compression.

The analysis of data structures and related notions of adaptivity were a key focus of the seminar. Conway (Section 4.3) gave an overview of results and open questions on history-independent priority queues, while Iacono (Section 5.1) discussed the working set property of priority queues in conjunction with efficient decrease-key. Brodal (Section 4.6) highlighted a simple and efficient data structure for storing integers, and Tarjan (Section 5.3) drew attention to the different analyses of path compression, raising the question of an alternative proof; Kuszmaul (Section 4.14) talked about recent results and open questions on quadratic probing hash tables. Kozma (Section 5.10) pointed out a remaining open question in the analysis of classical selection algorithms, and Farach-Colton (Section 5.5) raised a question related to the Karp-Rabin fingerprinting algorithm.

In his talk, Sedgewick (Section 4.13) advocated for an "algorithms science", using cardinality estimation as a case study. Zamir (Section 4.7) presented recent work on using cryptographic techniques to design asymptotically faster algorithms. Bercea (Section 4.18) discussed a possible new model for Bloom filters with predictions. Storandt (Section 5.7) presented data structure questions arising in the context of geometric intersection problems. Goodrich (Section 4.12) presented computational geometry problems in a setting with probabilistic errors, and raised the question of designing data structures that are robust to probabilistic deletions.

Collaboration was structured around informal working groups, with focus on some of the concrete open problems identified during the seminar. These notably included questions on simultaneous work and I/O optimality (see Section 4.10), fully dynamic graph connectivity (see Section 4.9), history independent heaps (see Section 4.3), heaps with the working set property (see Section 5.1) and greedy chain decompositions of DAGs (see Section 5.2). Several promising directions were explored, and results were obtained already during the seminar. It is expected that the collaborations started at the seminar will lead to publications.

## Final Thoughts

The organizers would like to thank the Dagstuhl team for their continuous support and also thank all participants for their contributions to the seminar. The current (16th) seminar was part of a long-running series that has co-evolved with the field, both tracking and inspiring its trends and developments throughout the years.

The seminar fills a unique niche in the computer science research landscape. Although data structural topics are well represented at major venues of theoretical computer science, this series provides essentially the only opportunity for core data structures researchers from around the world to meet and exchange ideas in an informal, collaborative setting, having data structures as the primary focus. (Some other fields, such as computational geometry have a broader tradition of specialized workshops focusing on problem-solving.)

The appreciation by the community for the seminar and the opportunity it offers to personally meet and interact is reflected by the very strong response to invitations and the highly positive subsequent feedback. Respondents particularly praised the relaxed, informal atmosphere and the focus on short talks and loosely structured collaboration around open questions, validating the organizers' efforts in this direction.

Earlier seminars in the series had few female participants. An important focus of the last three seminars was to significantly increase female attendance. In the current seminar, 43% of the invited participants were female, resulting in a 35% female attendance.

Another important goal of the seminar is to encourage the interaction between senior and early career researchers, the latter comprising around a third of the invited participants and eventual attendees. We also made a deliberate effort to include researchers who have not attended the previous seminar(s), these making up a slight majority of both the invitee- and participant lists; we find this to be particularly important for a long-running seminar.

All the aspects mentioned were strongly appreciated by the participants in the post-seminar survey. The survey respondents also praised the coverage of a diverse range of research topics, and the continued efforts by the organizers to refine and improve the seminar format.

## 2 Table of Contents

**Executive summary**
*Michael A. Bender, John Iacono, László Kozma, and Eva Rotenberg* . . . . . . . . 1

**Seminar program** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 7

**Overview of Talks** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9

Data Structures on the Ultra-Wide Word RAM
*Philip Bille* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9

Rethinking unstructured (parallel) I/O for SSDs
*Ulrich Carsten Meyer* . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9

History independent priority queues
*Alexander Conway* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9

Distributed collaboration to share the cost of obtaining external data even when the fraction of Byzantine players is large
*Valerie King* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9

Dynamic representation of bitvectors
*Gonzalo Navarro* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 10

A Simple Integer Successor-Delete Data Structure
*Gerth Stølting Brodal* . . . . . . . . . . . . . . . . . . . . . . . . . . . 10

Improving Algorithmic Efficiency using Cryptography
*Or Zamir* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 11

Covering Approximate Shortest Paths with DAGs
*Nicole Wein* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 11

Monte Carlo Batch-Dynamic Connectivity
*Quanquan C. Liu and Valerie King* . . . . . . . . . . . . . . . . . . . . . 11

Simultaneous Work And I/O Optimality
*Peyman Afshani* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 12

Minimizing Cache Misses with Repeated Data: The Cache Pair Problem
*Ian Munro* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 12

Computational geometry with probabilistically noisy primitive errors
*Michael T. Goodrich* . . . . . . . . . . . . . . . . . . . . . . . . . . . . 13

Cardinality Estimation: A Poster Child for Algorithm Science
*Robert Sedgewick* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 13

Quadratic probing hash tables
*William Kuszmaul* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 14

Single-source shortest p-disjoint paths: fast computation and sparse preservers
*Davide Bilò* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 14

A well-separated pair decomposition for low density graphs
*Joachim Gudmundsson* . . . . . . . . . . . . . . . . . . . . . . . . . . . . 15

Incremental topological sort
*Jeremy Fineman* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 15

Data structures and predictions
*Ioana-Oriana Bercea* . . . . . . 15

**Open problems** . . . . . . 16

Working set heaps with decrease-key
*John Iacono* . . . . . . 16

Greedy chain decomposition on DAGs
*Ivor van der Hoog* . . . . . . 16

Top-down analysis of path compression
*Robert Endre Tarjan* . . . . . . 17

Two problems on Lempel-Ziv compression
*Paweł Gawrychowski* . . . . . . 17

Detecting collisions in Karp-Rabin fingerprinting
*Martin Farach-Colton* . . . . . . 17

Scheduling parallel jobs with a DAG of constraints
*Kunal Agrawal* . . . . . . 18

Some Intersection Problems
*Sabine Storandt* . . . . . . 18

Maximum independent set in the congested clique model
*Shiri Chechik* . . . . . . 18

A graph problem with applications to grammar compression
*Gonzalo Navarro* . . . . . . 19

Median of medians in small groups
*László Kozma* . . . . . . 19

**Participants** . . . . . . 20

## 3 Seminar program

**Sunday May 4, 2025**

18:00 *Dinner buffet*

**Monday May 5, 2025**

7:30 *Breakfast*

9:00 *Opening & Introductions*

9:45 *Coffee break*

10:30 *Short presentations (10-15 minutes max.)*
Philip Bille, *Ultrawide word RAM*
Ulrich Meyer, *Utilizing latency on parallel operations on SSDs*
Alex Conway, *History independent priority queues*
Valerie King, *Distributed collaboration to share the cost of obtaining external data even when the fraction of Byzantine players is large*
Gonzalo Navarro, *Dynamic representation of bitvectors*
John Iacono, *Working set heaps with decrease-key*
Gerth Stølting Brodal, *Simple data structures with slightly worse bounds*

12:15 *Lunch*

14:00 Medium presentations (20-25 minutes + questions)
Or Zamir, *Improved runtimes using cryptography*
Nicole Wein, *Covering approximate shortest paths with DAGs*
Quanquan Liu, *Monte Carlo batch-dynamic connectivity*

15:30 *Coffee & Cake*

16:00 Collaboration time

18:00 *Dinner*

**Tuesday May 6, 2025**

7:30 *Breakfast*

9:30 *Short presentations (10-15 minutes max.)*
Peyman Afshani, *Separations between I/O and work*
Ian Munro, *Minimizing cache misses with repeated data*
Ivor van der Hoog, *Greedy chain decomposition on DAGs*
Robert Endre Tarjan, *On the recurrence in Top-down analysis of path compression by Seidel and Sharir*
Pawel Gawrychowski, *Two problems on Lempel-Ziv compression*
Michael Goodrich, *Computational geometry with probabilistically noisy primitive errors*

10:30 *Coffee break*

11:00 Medium presentations (20-25 minutes + questions)
Bob Sedgewick, *Cardinality estimation, a poster child for algorithm science*
William Kuszmaul, *Quadratic probing hash tables*

12:15 *Lunch*

13:30 Collaboration time

15:30 *Coffee & Cake and status update*

18:00 *Dinner*

**Wednesday May 7, 2025**

| | |
|---|---|
| 7:30 | *Breakfast* |
| 9:00 | *Short presentations* |
| | Martin Farach-Colton, *Detecting collisions in Karp-Rabin fingerprinting* |
| | Kunal Agrawal, *Scheduling parallel jobs with a DAG of constraints* |
| | Sabine Storandt, *Some intersection problems* |
| 10:15 | *Coffee break* |
| 10:45 | *Short presentations* |
| | Shiri Chechik, *Maximum independent set in the congested clique model* |
| | Davide Bilò, *Single-source shortest p-disjoint paths; Fast computation and sparse preservers* |
| | Joachim Gudmundsson, *Approximate distance oracles for lambda-low density graphs* |
| 12:00 | *Group picture* |
| 12:15 | *Lunch* |
| 14:00 | *Hike* (Lake Noswendl) |
| 15:30 | *Coffee & Cake* |
| 18:00 | *Dinner* |

**Thursday May 8, 2024**

| | |
|---|---|
| 07:30 | *Breakfast* |
| 9:00 | *Short presentations* |
| | Jeremy Fineman, *Incremental topological sort* |
| | Gonzalo Navarro, *A graph problem with applications to grammar compression* |
| | Ioana Oriana Bercea, *Oracles for Bloom filters with predictions* |
| | László Kozma, *Median of medians selection* |
| 10:30 | *Coffee break* |
| 11:00 | *Collaboration time* |
| 12:15 | *Lunch* |
| 13:30 | *Collaboration time* |
| 15:30 | *Coffee & Cake and status update* |
| 18:00 | *Dinner* |

**Friday May 9, 2023**

| | |
|---|---|
| 7:30 | *Breakfast and Check-out* |
| 9:00 | *Wrap-up session: outcomes of discussions and future directions* |
| 10:30 | *Coffee Break* |
| 12:15 | *Lunch* |

# 4 Overview of Talks

## 4.1 Data Structures on the Ultra-Wide Word RAM

*Philip Bille (Technical University of Denmark – Lyngby, DK)*



We give an overview of the ultra-wide word RAM and recent data structural results in it.

## 4.2 Rethinking unstructured (parallel) I/O for SSDs

*Ulrich Carsten Meyer (Goethe University – Frankfurt am Main, DE)*



In my talk I highlighted the differences between I/O accesses on traditional hard disks versus flash memory. It appears that there is a potential to use hidden parallelism in many algorithms in order to achieve better performance in the context of unstructured I/O patterns – for example when graph algorithms like SSSP perform random accesses to adjacency lists, which normally only fill a small fraction of the available block size each.

## 4.3 History independent priority queues

*Alexander Conway (Cornell Tech – New York, US)*



History-independence is a property of some data structures by which their memory layout does not reveal any information about the history of specific operations that lead to its current state. Here we consider this concept applied to priority queues, and ask how efficient a history-independent priority queue can be made. We give a short overview of the best known bounds.

## 4.4 Distributed collaboration to share the cost of obtaining external data even when the fraction of Byzantine players is large

*Valerie King (University of Victoria, CA)*


**Joint work of** Valerie King, John Augustine, Soumyottam Chatterjee, Valerie King, Manish Kumar, Shachar Meir, David Peleg
**Main reference** John Augustine, Soumyottam Chatterjee, Valerie King, Manish Kumar, Shachar Meir, David Peleg: "Distributed Download from an External Data Source in Faulty Majority Settings", CoRR, Vol. abs/2412.19649, 2024.
**URL** https://doi.org//10.48550/ARXIV.2412.19649

We consider the Download problem in the Data Retrieval Model, introduced in (DISC'24), where a distributed set of peers, some of which may be Byzantine, seek to learn $n$ bits of data stored at a trustworthy external data source. Each bit of data can be learned by a peer either through a direct (costly) query of the source or through other peers that have already learned it; the goal is to design a collaborative protocol that reduces the maximum number of bits queried by any one peer ("query complexity"). We achieve optimal query complexity in a synchronous fully connected network with resilience to any constant fraction $\beta < 1$ of Byzantine peers, under varying assumptions regarding time and message size.

## 4.5 Dynamic representation of bitvectors

*Gonzalo Navarro (University of Chile – Santiago de Chile, CL)*


**Main reference** Gonzalo Navarro: "Adaptive Dynamic Bitvectors", in Proc. of the String Processing and Information Retrieval – 31st International Symposium, SPIRE 2024, Puerto Vallarta, Mexico, September 23-25, 2024, Proceedings, Lecture Notes in Computer Science, Vol. 14899, pp. 204–217, Springer, 2024.
**URL** https://doi.org//10.1007/978-3-031-72200-4_16

I presented a technique to maintain a bitvector $B[1..n]$ within $n + o(n)$ bits of space, so that operations access, rank, and select, as well as updates to the bitvector (write position, insert bit, delete bit) can all be solved within $O(\log(n/q)/\log\log n)$ amortized time, when there are $q$ queries per update on average. This turns out to be optimal in the cell probe model. This basic tool can be used to obtain similar results on many compact data structures that build on bits. I also showed some experimental results on an implementation of this technique, which show that it is indeed practical. I think the result can be useful for other scenarios, and the main ideas can be inspiring.

## 4.6 A Simple Integer Successor-Delete Data Structure

*Gerth Stølting Brodal (Aarhus University, DK)*


**Main reference** Gerth Stølting Brodal: "A Simple Integer Successor-Delete Data Structure", in Proc. of the 23rd International Symposium on Experimental Algorithms, SEA 2025, July 22-24, 2025, Venice, Italy, LIPIcs, Vol. 338, pp. 8:1–8:16, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.SEA.2025.8

A very simple decremental data structure for maintaining a set of integers was presented, that supports initializing the set followed by $d$ deletions and $s$ successor queries in arbitrary order in total time $O(n + d + s(1 + \log_{\max(2,s/n)} \min(s, n)))$. The data structure consists of a single array of integers. The data structure is essentially a special case of the classic union-find data structure with path compression but with unweighted linking (i.e., without linking by rank or size).

## 4.7 Improving Algorithmic Efficiency using Cryptography

*Or Zamir (Tel Aviv University, IL)*


**Joint work of** Vinod Vaikuntanathan, Or Zamir
**Main reference** Vinod Vaikuntanathan, Or Zamir: "Improving Algorithmic Efficiency using Cryptography", CoRR, Vol. abs/2502.13065, 2025.
**URL** https://doi.org//10.48550/ARXIV.2502.13065

Cryptographic primitives have been used for various non-cryptographic objectives, such as eliminating or reducing randomness and interaction. We show how to use cryptography to improve the time complexity of solving computational problems. Specifically, we show that under standard cryptographic assumptions, we can design algorithms that are asymptotically faster than existing ones while maintaining correctness in a black-box manner.

## 4.8 Covering Approximate Shortest Paths with DAGs

*Nicole Wein (University of Michigan – Ann Arbor, US)*


**Joint work of** Sepehr Assadi, Gary Hoppenworth, Nicole Wein
**Main reference** Sepehr Assadi, Gary Hoppenworth, Nicole Wein: "Covering Approximate Shortest Paths with DAGs", in Proc. of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025, pp. 2269–2280, ACM, 2025.
**URL** https://doi.org//10.1145/3717823.3718298

I will talk about our newly defined notion of a "DAG cover". It is a directed analog of a tree cover, which is closely related to a probabilistic tree embedding. A DAG cover of a general directed graph $G$ is a small collection of DAGs so that for all pairs of vertices $s$,$t$, some DAG in the collection provides low distortion for $\mathrm{dist}(s, t)$. I will discuss upper and lower bounds for DAG covers in various parameter regimes, and pose some open problems.

## 4.9 Monte Carlo Batch-Dynamic Connectivity

*Quanquan C. Liu (Yale University – New Haven, US) and Valerie King (University of Victoria, CA)*


**Joint work of** Lalia Cann, Valerie King, Quanquan C. Liu

Connectivity is an important problem with countless real-world applications. In particular, the graphs we perform connectivity queries on are often *dynamic* where edges are inserted or deleted with high frequency. The dynamic connectivity problem then seeks to answer connectivity queries while an online sequence of edge insertions and deletions occur in the graph. In this paper, we give the first parallel batch-dynamic implementation with provably large efficiency and parallelism guarantees. Furthermore, our guarantees show that the runtimes are small, even in the worst case. To do this, we simplify, adapt, and parallelize the sequential Monte Carlo algorithms of Kapron, King, and Mountjoy [SODA 2013], Gibb, Kapron, King, and Mountjoy [2015], and Wang for the batch-dynamic setting. These algorithms provably require very small, $\mathrm{poly}(\log n)$ time per update and query *even in the* *worst-case.* Our simplification uses only Euler Tour (ET) trees, rather than more complicated path compression data structures, making it simultaneously simpler to implement and faster than more complicated structures.

We present an open problem regarding whether we can eliminate the level data structure altogether in both the sequential dynamic algorithm and our algorithm.

## 4.10 Simultaneous Work And I/O Optimality

*Peyman Afshani (Aarhus University, DK)*



We study a number of algorithmic problems in the I/O model of computation, aiming to obtain an algorithm which achieves the optimal number of I/O operations in the I/O model as well as achieving the optimal amount of CPU work in the classical RAM model. In some cases, this is already known to be impossible (Afshani, Brodal, Sitchinava, manuscript, 2025). However, if that is the case, the goal would be to obtain an optimal work-I/O trade-off. The problems that we study are the following: deferred data structures and red-blue sorting.

For the first problem, the input is an unsorted list of values which needs to be stored for online queries. However, the total number of queries is not known and the goal is to minimize the total amount of preprocessing time done. It is a classical result of Karp and Motwani (SIAM JoC'88) that if the total number of queries is $k$ (which is unknown in advance), then this can be done in $O(n \log k)$ time for a number of different queries such as predecessor search queries. A similar performance can be achieved in the I/O model however, achieving both seems to be impossible and thus this leads to work-I/O trade-off for this problem.

For the second problem, the input is a set of n values, each associated with one of the two colors, red or blue. The goal is to sort the items using the minimum number of comparisons, however, for two elements $u$ and $w$ if there is no other element of the opposite color that lies between them, then the determining the order of u and w is not necessary. When there is only one value $v$ of one color, then the problem reduces to simply the splitting the values with respect to pivot $v$ which can be done in linear time, as well as linear I/Os and thus simultaneous work and I/O optimality is possible in this case. It has been observed that in the worst-case, this can be done in general. However, obtaining an adaptive bound on the number of comparisons is a more interesting problem for which obtaining such a simultaneous optimality does not seem possible.

## 4.11 Minimizing Cache Misses with Repeated Data: The Cache Pair Problem

*Ian Munro (University of Waterloo, CA)*


**Joint work of** Ian Munro, Amin Khodaee

We consider the generalization of the problem of minimizing cache misses in the situation in which data items are stored on multiple cache lines across memory (i.e. there are multiple copies of the same objects). The issue is, then what line to evict when a new one must be uploaded. In the usual (single copy) situation, it is well known that (i) the minimum number of cache misses for the off line problem (i.e. all requests are known in advance) can easily be determined and (ii) in the on line case, one can come within a factor of two of this bound given twice the number of cache lines. We show that if each datum is stored on two cache lines (so we have a "choice" of which line to bring in,) the off line version of the problem is NP-hard, and indeed, assuming the unique games conjecture, coming within twice the optimal using twice the given number of line is also NP-hard. We give an easy solution to the on line version that is within four times as many cache misses as the optimal given quadruple the space.

## 4.12 Computational geometry with probabilistically noisy primitive errors

*Michael T. Goodrich (University of California – Irvine, US)*


**Joint work of** David Eppstein, Michael T. Goodrich, Vinesh Sridhar
**Main reference** David Eppstein, Michael T. Goodrich, Vinesh Sridhar: "Computational Geometry with Probabilistically Noisy Primitive Operations", in Proc. of the 19th International Symposium on Algorithms and Data Structures, WADS 2025, August 11-15, 2025, York University, Toronto, Canada, LIPIcs, Vol. 349, pp. 24:1–24:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.WADS.2025.24

Much prior work has been done on designing computational geometry algorithms that handle input degeneracies, data imprecision, and arithmetic round-off errors. We take a new approach, inspired by the noisy sorting literature, and study computational geometry algorithms subject to noisy Boolean primitive operations in which, e.g., the comparison "is point q above line L?" returns the wrong answer with some fixed probability. We propose a novel technique called path-guided pushdown random walks that generalizes the results of noisy sorting. We apply this technique to solve point-location, plane-sweep, convex hulls in 2D and 3D, dynamic 2D convex hulls, and Delaunay triangulations for noisy primitives in optimal time with high probability.

## 4.13 Cardinality Estimation: A Poster Child for Algorithm Science

*Robert Sedgewick (Princeton University, US)*


**Joint work of** Svante Janson, Jérémie O. Lumbroso, Robert Sedgewick
**Main reference** Svante Janson, Jérémie O. Lumbroso, Robert Sedgewick: "Bit-Array-Based Alternatives to HyperLogLog", in Proc. of the 35th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms, AofA 2024, June 17-21, 2024, University of Bath, UK, LIPIcs, Vol. 302, pp. 5:1–5:19, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.
**URL** https://doi.org//10.4230/LIPICS.AOFA.2024.5

This talk emphasizes the role of algorithm science in the decades-long development of state-of-the-art cardinality estimation algorithms, from the seminal papers of Flajolet and coauthors to the recent bit-array based methods of Lumbroso, Janson and Sedgewick.

## 4.14 Quadratic probing hash tables

*William Kuszmaul (Carnegie Mellon University – Pittsburgh, US)*


**Joint work of** William Kuszmaul, Zoe Xi
**Main reference** William Kuszmaul, Zoe Xi: "Towards an Analysis of Quadratic Probing", in Proc. of the 51st International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024, Tallinn, Estonia, LIPIcs, Vol. 297, pp. 103:1–103:19, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.
**URL** https://doi.org//10.4230/LIPICS.ICALP.2024.103

Since 1968, one of the simplest open questions in the theory of hash tables has been to prove anything nontrivial about the correctness of quadratic probing. We make the first tangible progress towards this goal, showing that there exists a positive-constant load factor at which quadratic probing is a constant-expected-time hash table. Our analysis applies more generally to any fixed-offset open-addressing hash table, and extends to higher load factors in the case where the hash table examines blocks of some size $B = \omega(1)$.

## 4.15 Single-source shortest p-disjoint paths: fast computation and sparse preservers

*Davide Bilò (University of L'Aquila, IT)*


**Joint work of** Davide Bilò, Gianlorenzo D'Angelo, Luciano Gualà, Stefano Leucci, Guido Proiett, Mirko Rossi
**Main reference** Davide Bilò, Gianlorenzo D'Angelo, Luciano Gualà, Stefano Leucci, Guido Proietti, Mirko Rossi: "Single-Source Shortest p-Disjoint Paths: Fast Computation and Sparse Preservers", in Proc. of the 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference), LIPIcs, Vol. 219, pp. 12:1–12:21, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
**URL** https://doi.org//10.4230/LIPICS.STACS.2022.12

Let $G$ be a directed graph with $n$ vertices, $m$ edges, non-negative edge costs, and a fixed source vertex $s$. Given a target vertex $t$, the problem of computing a set $S_t$ of $p$ pairwise edge-disjoint (simple) paths from $s$ to $t$ in $G$ having minimum total cost (if such paths exist) can be solved in $O(p(m + n \log n))$ time using the Successive Shortest Path algorithm. We are interested in constructing a compact data structure that, when queried with any target vertex $t$, can quickly return $S_t$.

For $p = 1$, we can precompute a shortest-path tree of $G$ rooted at $s$ and build a data structure of size $O(n)$ that answers queries in time linear in the size of the output. The preprocessing time is $O(m + nlogn)$ when using Dijkstra's algorithm. For $p = 2$, Suurballe and Tarjan [Networks 1984] designed a data structure with the same asymptotic trade-offs among preprocessing time, query time, and size. For general values of $p$, Bilò et al. [STACS 2022] designed a $O(pn(pn + n \log n))$-time algorithm that computes a subgraph $H$ of $G$ of size at most $p(n - 1)$ containing $S_t$ for every target vertex $t$. The size of $H$ is optimal when $G$ is $p$-edge-outconnected from $s$. This result implies the existence of a data structure with preprocessing time $O(pn(pn + n \log n))$, size $O(pn)$, and query time $O(p(pn + n \log n))$. Improving the preprocessing and query times remains an interesting open problem.

## 4.16 A well-separated pair decomposition for low density graphs

*Joachim Gudmundsson (The University of Sydney, AU)*


**Joint work of** Joachim Gudmundsson, Sampson Wong

Low density graphs are considered to be a realistic graph class for modelling road networks. It has advantages over other popular graph classes for road networks, such as planar graphs, bounded highway dimension graphs, and spanners. We believe that low density graphs have the potential to be a useful graph class for road networks, but until now, its usefulness is limited by a lack of available tools. In this talk, we show two new fundamental tools for low density graphs, that is, a well-separated pair decomposition and an approximate distance oracle. We believe that by expanding the algorithmic toolbox for low density graphs, we can help provide a useful and realistic graph class for road networks, which in turn, may help explain the many efficient and practical heuristics available for road networks.

## 4.17 Incremental topological sort

*Jeremy Fineman (Georgetown University – Washington, DC, US)*



This talk gives a short overview on incremental topological sort. Here m edges are added to an n-vertex graph incrementally, and the goal is to maintain a topological sort. The current state of the art includes the following total update times, ignoring log factors:

- Dense graphs (large m): $\tilde{O}(n^2)$ from Bender, Fineman, Gilbert, and Tarjan. This work also includes an extension to maintain a strongly connected components (SCCs).
- Sparse graphs (small m): $\tilde{O}(m^{4/3})$ from Battacharya and Kulkarni. Bernstein, Dudeja, and Pettie later obtained the same bound for SCCs.

Finally, Chen, Kyng, Liu, Meierhans, and Probst Gutenberg obtained an algorithm that has total update time of $m^{1+o(1)}$ for the related problem of cycle detection, but their algorithm does not maintain a topological sort. Moreover, their algorithm is not combinatorial.

We have reason to believe that it should be possible to obtain a bound of $\tilde{O}(n\sqrt{m})$ for incremental topological sort with a combinatorial algorithm. Ignoring log factors, this would be at least as good as the best dense algorithm for all m and n, and it would beat the best sparse algorithm for graphs of sufficient density.

## 4.18 Data structures and predictions

*Ioana-Oriana Bercea (KTH Royal Institute of Technology – Stockholm, SE)*



The Daisy Bloom filter is an approximate memembership data structure (filter) that has access to knowledge of the distribution of the input and query distribution. While we understand how to set optimal parameters in this setting, this has yet to translate into practical gains. We highlight some questions in this direction, with the exciting prospect of bridging the gap between classical results on instance optimality and emerging directions such as algorithms with predictions.

# 5 Open problems

## 5.1 Working set heaps with decrease-key

*John Iacono (ULB - Brussels, BE)*



Priority queues are one of the oldest abstract data types in computer science, and those that support the decrease-key operation, such a Fibonacci Heaps, have been instrumental in a number of efficient graph algorithms such as single source shortest paths (SSSP). Recent progress in the beyond-worst-case analysis of SSSP has shown the utility of having a priority queue that supports constant time-decrease key as well as the working set property, one variant of which requires that the cost of an extract-min be proportional to the logarithm of the number of operations performed on the heap while that data item was present in it. No such pointer-model data structure exists with both the working-set property and constant-time decrease-key. Also briefly discussed were some variants of what the working set could mean for priority queues.

## 5.2 Greedy chain decomposition on DAGs

*Ivor van der Hoog (Technical University of Denmark – Lyngby, DK)*



Given a directed graph $G$ with $n$ vertices and $m$ edges, we consider the problem of computing a *longest directed path decomposition*: a recursive partitioning of the graph into sets of vertices, where each set consists of the vertices along some longest directed path in the remaining graph. Observe that, since the longest path need not be unique, the decomposition is not unique either.

We present a simple two-stage algorithm that computes such a decomposition in time $O(nm)$, improving over the previous best-known bound of $O(n^{2.5})$ when the number of edges $m$ is subquadratic. The algorithm proceeds as follows:

- First, perform a DFS to find a longest directed path in $O(m)$ time. If its length is at least $n$, remove its vertices and recurse.
- Otherwise, the graph has height at most $n$. We compute a height decomposition and construct a flow-DAG by connecting a super-source to all sources in $G$, and a super-sink to all vertices of maximum height. A blocking flow in this DAG can be computed in $\tilde{O}(m)$ time, and its removal eliminates all longest paths. Since the height decreases by at least one per iteration and is initially at most $\sqrt{n}$, this phase takes $\tilde{O}(m\sqrt{n})$ time overall.

## 5.3 Top-down analysis of path compression

*Robert Endre Tarjan (Princeton University, US)*


**Main reference** Raimund Seidel, Micha Sharir: "Top-Down Analysis of Path Compression", SIAM J. Comput., Vol. 34(3), pp. 515–525, 2005.
**URL** https://doi.org//10.1137/S0097539703439088

Seidel and Sharir gave a proof of the inverse-Ackermann-function upper bound for path compression based on a beautiful top-down recurrence. The analysis does not use Ackermann's function explicitly, but uses a related function they call "J." Problem: Give a simple, direct proof of the inverse-Ackermann function bound using their top-down recurrence and the classical definition of Ackermann's function.

## 5.4 Two problems on Lempel-Ziv compression

*Paweł Gawrychowski (University of Wroclaw, PL)*



We ask two questions about the widely used Lempel-Ziv compression algorithm, and give an overview of some known results relating to these questions:

- Can pattern matching on a Lempel-Ziv compressed text be done in $O(n + m)$ time, for a string of compressed length $n$ and a pattern of length $m$?
- Can we design a data-structure using $O(n)$ space which allows for $O(\log N)$-time indexing of characters in a string of length $N$ with a Lempel-Ziv compressed representation of length $n$?

## 5.5 Detecting collisions in Karp-Rabin fingerprinting

*Martin Farach-Colton (NYU – New York, US)*



The Karp-Rabin Fingerprint method is a clean and simple way to achieve fast algorithms for many string matching problems. The idea is to treat the characters of a string as the coefficients of a polynomial and to take the value of this polynomial evaluated at $\Sigma$, the size of the alphabet. In order to keep the number of bits down (and thus to be able to perform constant-time manipulations), the value is computed modulo a prime number whose value is some polynomial of $n$, the size of the longest string under consideration. For any substring, this modular evaluation of the induced polynomial is the Karp-Rabin fingerprint of that substring.

Although taking the mod does speed up the fingerprinting algorithm, it induces collisions where two substrings of the same length might evaluate to the same fingerprint even though they do not match.

The open question is: how quickly can one take a string and a prime and check of any pair of substrings (of the same length) of the string have colliding fingerprints? In other words, are there two distinct substrings for which the given prime induces a collision? A quadratic-time algorithm is straightforward using suffix trees. Can we achieve a linear time algorithm or at least an $O(n \log n)$-time algorithm?

It's ok if we reject some polynomials that don't, in fact, induce a collision, as long as we do it with small probability.

## 5.6 Scheduling parallel jobs with a DAG of constraints

*Kunal Agrawal (Washington University – St. Louis, US)*



We give an overview of known results and open questions regarding the scheduling of parallel tasks, each of which is represented as a directed acyclic graph of precedence constraints (thus allowing for both parallelism between tasks as well as intra-task parallelism).

## 5.7 Some Intersection Problems

*Sabine Storandt (Universität Konstanz, DE)*



The talk introduces the following problems in the realm of intersection reporting data structures:

- Construct an efficient segment-circle intersection data structure under the assumption that the n input segments form a connected graph. The general data structure requires $O(n^{1.5})$ preprocessing time and $O(n^{0.5} + k)$ query time where k is the output size.
- Given a set of geometric objects, preprocess it such that for a given query range pairs of input objects intersection inside the query range can be reported efficiently. While for many types of objects and ranges, 3-SUM hardness can be shown, half-plane queries for segments and lines escape that lower bound. However, faster query times might be possible in that case and more general scenarios as the input objects being polylines is wide open.

## 5.8 Maximum independent set in the congested clique model

*Shiri Chechik (Tel Aviv University, IL)*



We ask how many rounds of communication are necessary to compute a maximum independent set of a graph $G$ in a particular model of distributed computation, known as the congested clique model. In this model there are n players identified with the vertices of $G$, each player initially knows the vertices adjacent to itself, and can send $b$ bits of information to every other player in each round of communication.

We give a short overview of the known results.

## 5.9 A graph problem with applications to grammar compression

*Gonzalo Navarro (University of Chile – Santiago de Chile, CL)*



One can represent a text $T[1..u]$ with a context-free grammar of size n that produces (only) $T$. This representation requires $O(n \log n)$ bits. However, every structure providing random access to T using the grammar requires $\Theta(n \log u)$ further bits of space, to record the expansion length of the nonterminals. The question is whether we can remove this term and still access the text efficiently. One choice is to sample $o(n)$ nodes so that the others' expansion lengths can be obtained in polylog time from the sampled ones. This boils down to a problem sparse on acyclic graphs.

## 5.10 Median of medians in small groups

*László Kozma (FU Berlin, DE)*



I briefly mentioned an open problem raised by Chen and Dumitrescu in connection to the "median of medians" selection algorithm of Blum, Floyd, Pratt, Rivest, and Tarjan from 1973. This linear time selection algorithm is based on arranging the input elements in groups of five, leading to the natural question of whether groups of three are also sufficient. The classical analysis only yields a superlinear bound in this case. Chen and Dumitrescu point out that contrary to common belief, this does not necessarily imply that the running time with groups of three is indeed superlinear, and constructing a "difficult instance" appears to be difficult; thus, a gap remains in our understanding of this fundamental algorithm.

## Participants

- Peyman Afshani
Aarhus University, DK
- Kunal Agrawal
Washington University – St. Louis, US
- Hideo Bannai
Institute of Science Tokyo, JP
- Michael A. Bender
Stony Brook University, US
- Ioana Oriana Bercea
KTH Royal Institute of Technology – Stockholm, SE
- Philip Bille
Technical University of Denmark – Lyngby, DK
- Davide Bilò
University of L'Aquila, IT
- Gerth Stølting Brodal
Aarhus University, DK
- Shiri Chechik
Tel Aviv University, IL
- Alexander Conway
Cornell Tech – New York, US
- Justin Dallant
UL – Brussels, BE
- Aditi Dudeja
Paris Lodron Universität Salzburg, AT
- Faith Ellen
University of Toronto, CA
- Martin Farach-Colton
NYU – New York, US
- Jeremy Fineman
Georgetown University – Washington, DC, US
- Pawel Gawrychowski
University of Wroclaw, PL
- Michael Goodrich
University of California – Irvine, US
- Joachim Gudmundsson
The University of Sydney, AU
- Inge Li Gørtz
Technical University of Denmark – Lyngby, DK
- John Iacono
ULB – Brussels, BE
- Rob Johnson
Broadcom – San Jose, US
- Valerie King
University of Victoria, CA
- Tomasz Kociumaka
MPI für Informatik – Saarbrücken, DE
- Hanna Komlós
NYU – New York, US
- László Kozma
FU Berlin, DE
- William Kuszmaul
Carnegie Mellon University – Pittsburgh, US
- Jingxun Liang
Carnegie Mellon University – Pittsburgh, US
- Quanquan C. Liu
Yale University – New Haven, US
- Ulrich Carsten Meyer
Goethe University – Frankfurt am Main, DE
- Ian Munro
University of Waterloo, CA
- Gonzalo Navarro
University of Chile – Santiago de Chile, CL
- Eva Rotenberg
Technical University of Denmark – Lyngby, DK
- Robert Sedgewick
Princeton University, US
- Marek Sokolowski
MPI für Informatik – Saarbrücken, DE
- Teresa Steiner
Technical University of Denmark – Lyngby, DK
- Sabine Storandt
Universität Konstanz, DE
- Robert Endre Tarjan
Princeton University, US
- Ivor van der Hoog
Technical University of Denmark – Lyngby, DK
- Stefan Walzer
KIT – Karlsruher Institut für Technologie, DE
- Nicole Wein
University of Michigan – Ann Arbor, US
- Huacheng Yu
Princeton University, US
- Or Zamir
Tel Aviv University, IL
- Renfei Zhou
Carnegie Mellon University – Pittsburgh, US

Report from Dagstuhl Seminar 25192

# AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation

**Giuseppe De Giacomo*[1], Marlon Dumas*[2], Fabiana Fournier*[3], Timotheus Kampik*[4], and Lior Limonad*[5]**

1 **University of Oxford, GB.** giuseppe.degiacomo@cs.ox.ac.uk
2 **University of Tartu, EE.** marlon.dumas@ut.ee
3 **IBM Research Israel – Haifa, IL.** fabiana@il.ibm.com
4 **SAP Berlin, DE & Umeå University, SE.** timotheus.kampik@sap.com
5 **IBM Research Israel – Haifa, IL.** liorli@il.ibm.com

**Abstract**

Advances in AI are enabling the shift toward Autonomous Business Processes (ABPs), where systems not only suggest actions but also take proactive steps within defined constraints. This concept was introduced in the AI-Augmented Business Process Management Systems (ABPMSs) manifesto, which outlines their lifecycle, features, and research challenges. The "AutoBiz" 25192 Dagstuhl Seminar brought together experts from AI and BPM to collaborate on advancing this vision. The seminar's main goal was to define a research agenda for the realization of ABP systems.

**Seminar** May 4–9, 2025 – https://www.dagstuhl.de/25192

**2012 ACM Subject Classification** Computing methodologies → Artificial intelligence; Information systems → Information systems applications; Applied computing → Business process management; Computer systems organization → Self-organizing autonomic computing; General and reference → Surveys and overviews

**Keywords and phrases** AutoBiz, Artificial Intelligence, Business Process Management, Autonomous Business Processes, Dagstuhl Seminar

**Digital Object Identifier** 10.4230/DagRep.15.5.21

## 1 Executive Summary

*Fabiana Fournier (IBM Research Israel – Haifa, IL, fabiana@il.ibm.com)*



Advances in AI make it possible to push the boundaries of automation into the realm of **Autonomous Business Processes (ABPs)**. In ABPs, AI-based systems not only recommend predefined interventions, as in prescriptive process execution, but they also proactively trigger interventions to respond to unforeseen changes within user-defined constraints. The initial vision towards ABPs was recently introduced in the "*AI-Augmented Business Process Management Systems: A Research Manifesto*" paper. This manifesto coined the concept of **AI-Augmented Business Process Management Systems (ABPMSs)**, their lifecycle, core characteristics, and the research challenges they present. An ABPMS enhances the execution of business processes with the aim of making these processes more: Adaptable, proactive, explainable, and context-sensitive.

* Editor / Organizer


AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation, *Dagstuhl Reports*, Vol. 15, Issue 5, pp. 21–63
Editors: Giuseppe De Giacomo, Marlon Dumas, Fabiana Fournier, Timotheus Kampik, and Lior Limonad
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The most advanced form of ABPs is manifested by **ABP systems**. This Dagstuhl Seminar brought together leading academic and industrial researchers from the AI and BPM communities to collaboratively advance the vision of ABP systems and address the challenges outlined in the manifesto. These include framed autonomy, situation-aware explainability, automated process adaptation, and actionable conversations. The seminar was structured into corresponding working groups, each focusing on one of these aspects.

The discussions in these groups led to the development of four dedicated publications, presented at the PMAI workshop at ECAI 2025. These publications build upon the original manifesto by elaborating on the concrete challenges and technical foundations required to realize different characteristics of ABPMSs. The seminar thus marks a significant step toward a comprehensive research agenda for Agentic BPM systems—systems that integrate autonomous agents capable of reasoning, learning, and acting within framed constraints.

The results of this seminar, as elaborated in this report, form a bold call to action for realizing the vision of Agentic BPM systems. These systems embody the architectural principles and core ideas initially sketched in the manifesto and now further developed through collaborative research and discussion.

For further details, readers are encouraged to consult the full report.

## 2 Table of Contents

**Executive Summary**
*Fabiana Fournier* . . . . . 21

**Overview of Talks**

AI-Augmented BPM – The manifesto and its key characteristics
*Marlon Dumas* . . . . . 24

Framed Autonomy in AI-Augmented Business Process Management Systems
*Marco Montali* . . . . . 24

Uncertainty Quantification and its role in AI-Augmented BPM
*Niek Tax* . . . . . 25

AI-augmented Process Mining
*Stefanie Rinderle-Ma* . . . . . 26

Causal business processes: A new paradigm for agentic observability
*Lior Limonad* . . . . . 26

Generative AI for process explainability
*Fabiana Fournier* . . . . . 27

AI-Assisted Prescriptive Business Process Monitoring
*Andreas Metzger* . . . . . 28

Foundations of Agentic AI: GenAI meets Strategic Reasoning and Planning
*Giuseppe De Giacomo* . . . . . 28

**Working Groups**

Working Group on Framed Autonomy
*Giuseppe De Giacomo, Andrea Marrella, Yves Lesperance, Andrea Matta, Timotheus Kampik, and Diego Calvanese* . . . . . 29

Working Group on Adaptive / Uncertainty Quantification
*Estefania Serral, Achiya Elyasaf, Andreas Metzger, Niek Tax, Sebastian Sardina, and Arik Senderovich* . . . . . 36

Working Group on Conversational Actionability
*Daniel Amyot, Marco Comuzzi, Marlon Dumas, Marco Montali, and Irene Teinemaa* 45

Working Group on Explainability
*Peter Fettke, Fabiana Fournier, Lior Limonad, Andreas Metzger, Stefanie Rinderle-Ma, and Barbara Weber* . . . . . 51

**AUTOBIZ Dissemination**

Embedding of ABPMS-related Training into University Curricula . . . . . 61

Related Publications and Future Venues . . . . . 62

**Participants** . . . . . 63

# 3 Overview of Talks

## 3.1 AI-Augmented BPM – The manifesto and its key characteristics

*Marlon Dumas (University of Tartu, EE, marlon.dumas@ut.ee)*



AI-Augmented Business Process Management Systems (ABPMSs) are an emerging class of process-aware information systems, empowered by trustworthy AI technology. An ABPMS enhances the execution of business processes with the aim of making these processes more adaptable, proactive, explainable and context-sensitive. This talk presents a vision for ABPMSs and discusses research challenges that ought to be surmounted to realize this vision. To this end, we define the concept of ABPMS, we outline the lifecycle of processes within an ABPMS, we discuss core characteristics of an ABPMS, and we derive a set of challenges to realize systems with these characteristics.

Details are given in [1] and [2].

**References**

1 Marlon Dumas, Fabiana Fournier, Lior Limonad, Andrea Marrella, Marco Montali, Jana-Rebecca Rehse, Rafael Accorsi, Diego Calvanese, Giuseppe De Giacomo, Dirk Fahland, Avigdor Gal, Marcello La Rosa, Hagen Völzer, and Ingo Weber. Ai-augmented business process management systems: A research manifesto. *ACM Trans. Manag. Inf. Syst.*, 14(1):11:1–11:19, 2023.
2 David Chapela-Campa and Marlon Dumas. From process mining to augmented process execution. *Softw. Syst. Model.*, 22(6):1977–1986, 2023.

## 3.2 Framed Autonomy in AI-Augmented Business Process Management Systems

*Marco Montali (Free University of Bozen-Bolzano, IT, montali@inf.unibz.it)*



Process framing is about defining suitable boundaries of execution within which AI agents and humans can operate and collaborate to enact one or more work processes. The talk summarized a long-standing line of research aiming at providing an explicitly formal description of the frames, which can be effectively employed in computational terms, e.g., for reasoning, verification, execution, and analysis. Temporal logics on finite traces, together with their automata-theoretic characterization, are the basis of this approach, which has been widely explored within process science to provide a declarative way for process specification and management (cf., in particular the Declare language). The talk started by showing how this approach can be effectively used for process framing using declarative process mining tools for frame elicitation and deviation analysis. It then innovatively expanded the approach to deal with:

- Frames that can be broken – providing anticipatory monitoring techniques to promptly detect deviations at runtime
- Uncertain frames, where “a portion” of traces may violate some constraints
- Data-aware frames, dealing with data attributes or more complex objects and their mutual relationships

All these settings come with suitable extensions of temporal logics on finite traces, with corresponding foundational and practical results on computation. Finally, how agents operating within the frame can employ the repertoire of techniques to operate and enact processes, discussing several open challenges within and across artificial intelligence and process science.

## 3.3 Uncertainty Quantification and its role in AI-Augmented BPM

*Niek Tax (Meta – London, GB, niek@meta.com)*



Uncertainty Quantification (UQ) plays a pivotal role in enhancing the reliability, safety, and adaptability of AI-augmented Business Process Management (BPM) systems. This presentation explores the connection between UQ and several key BPM tasks. Various forms of UQ are discussed, including probability calibration and the modeling of the full posterior predictive distribution (beyond point estimates) and its applications in risk-aware decision making, ensuring fairness, and enabling adaptability through guided exploration (e.g., active learning). The talk also presents recent advances in UQ from the Meta Central Applied Science team, including:

**InfoShap [1]** A SHAP-based method for explaining uncertainty estimates.
**AdaptiveWeightSampling [2]** A theoretically grounded active learning method that scale to industrial applications.
**TCE [3]** A calibration error metric suitable for imbalanced domains.
**MCE [4]** A metric to quantify the extent to which predictions are simultaneously calibrated across multiple subpopulations (i.e., *multicalibrated* [5]).

### References

1 D. Watson, J. O'Hara, N. Tax, R. Mudd, and I. Guy. Explaining predictive uncertainty with information theoretic shapley values. In *Advances in Neural Information Processing Systems*, volume 36, pages 7330–7350, 2023.
2 Daniel Haimovich, Dima Karamshuk, Fridolin Linder, Niek Tax, and Milan Vojnovic. On the convergence of loss and uncertainty-based active learning algorithms. *Advances in Neural Information Processing Systems*, 37:122770–122810, 2024.
3 Takuo Matsubara, Niek Tax, Richard Mudd, and Ido Guy. Tce: A test-based approach to measuring calibration error. In *Uncertainty in Artificial Intelligence*, pages 1390–1400. PMLR, 2023.
4 Ido Guy, Daniel Haimovich, Fridolin Linder, Nastaran Okati, Lorenzo Perini, Niek Tax, and Mark Tygert. Measuring multi-calibration. *arXiv preprint arXiv:2506.11251*, 2025.
5 Ursula Hébert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pages 1939–1948. PMLR, 2018.

## 3.4 AI-augmented Process Mining

*Stefanie Rinderle-Ma (TU Munich, DE, stefanie.rinderle-ma@tum.de)*



Process mining and automation have emerged as technological megatrends, largely driven by recent advances in Artificial Intelligence (AI). Among these, generative AI and particularly large Language Models (LLMs) hold significant promise for transforming process discovery from textual data and enabling the creation of process models by domain experts in conversation with a chatbot. Furthermore, process technologies facilitate the collection of contextualized data, especially event log data augmented with IoT data, which can be leveraged to optimize the prediction of process behaviour by, e.g., including sensor data, targeting, for example, enhanced compliance checks.

The presentation introduces the concepts of conversational process mining and redesign, showcases the collection of contextualized data, and outlines key research directions.

## 3.5 Causal business processes: A new paradigm for agentic observability

*Lior Limonad (IBM Research Israel – Haifa, IL, liorli@il.ibm.com)*


**Joint work of** Lior Limonad, Fabiana Fournier

"A rooster's crow does not cause the sun to rise, even though it always precedes it." (The book of Why, Mackenzie & Pearl, 2018). This quote captures a common fallacy – post hoc ergo propter hoc – our tendency to infer causation from mere sequences: after this, therefore because of this.

In the presentation, I argued that time-correlated events do not necessarily imply a causal execution relationship. As an example, we presented a simple mortgage application process in which the acceptance decision is followed by informing the customer and archiving the application. While conventional process mining might depict the latter two activities as a direct sequence, this does not imply that informing the customer causes archiving, or vice versa.

Understanding true causal execution dependencies is essential for meaningful process improvement. It helps avoid costly empirical interventions by identifying which outcomes stem from specific changes. Our presented method introduces a novel technique for analyzing activity execution times in business processes, based on logs. We extend recent causal discovery methods – originally designed for continuous variables – and adopt them to timestamped events, while aligning with realistic assumptions about process execution. This enables a new paradigm: **Causal Business Processes**.

The talk presented our foundational work on discovering causal execution dependencies from process event logs and modeling them in a unified framework. This approach augments traditional process mining by enabling predictive insights about the effect of interventions, critical for informed process optimization.

Further details about this work is available in [1] and [2].

**References**

1 F. Fournier, L. Limonad, I. Skarbovsky, and Y. David. The why in business processes: Discovery of causal execution dependencies. *Künstliche Intelligenz*, 1 2025.

2 Yuval David, Fabiana Fournier, Lior Limonad, and Inna Skarbovsky. The why in business processes: Unification of causal process models. In *BPM Forum in BPM Conference (to appear)*, 9 2025.

## 3.6 Generative AI for process explainability

*Fabiana Fournier (IBM Research Israel –Haifa, - IL, fabiana@il.ibm.com)*


**Joint work of** Fabiana Fournier, Lior Limonad

This talk tackles the interpretability aspect of SAX explanations and explore the question on how effectively large language models (LLMs) can explain outcomes and decisions in business processes (e.g., loan rejection, parking fines) as perceived by users. In the SAX framework, a set of knowledge extractors analyses an input log (a historical record of process executions within an organization) and generates key perspectives – or "views" – of the process. These three views are then combined into a prompt for an LLM, along with the user's inquiry (e.g., why my loan application was rejected?). The LLM uses this input to generate an explanation tailored to the user's specific question. To assess this, we conducted a user study with 50 participants with the goal to assess both perceptions of fidelity and interpretability, while also considering the roles of trust and curiosity as experienced by the users. Our results show that adding different perspectives does, in fact, improve the perceived fidelity of the generated explanations. However, our study concluded that adding knowledge views improves perceived fidelity, but such improvement can actually come at the cost of perceived interpretability. Once we had developed a scale to assess the quality of explanations generated by LLMs, we began exploring if we can leverage the developed scale for LLM selection and refinement. To address this challenge, we launched another study – this time focused on evaluating how our scale could be used to quantify user perceptions and enable a meaningful comparison between different LLMs. The experiment centered on a tax refund process and included 128 participants, who were asked to rate explanations for various tax refund decisions using our evaluation scale.

Further details about this work is available in [1] and [2].

**References**

1 Lior Limonad, Fabiana Fournier, Hadar Mulian, George Manias, Spiros Borotis, and Danai Kyrkou. Selecting the right llm for egov explanations, 2025.

2 Dirk Fahland, Fabiana Fournier, Lior Limonad, Inna Skarbovsky, and Ava J. E. Swevels. How well can large language models explain business processes as perceived by users? *Data & Knowledge Engineering*, 157:102416, 2 2025.

## 3.7 AI-Assisted Prescriptive Business Process Monitoring

*Andreas Metzger (University of Duisburg-Essen, DE, andreas.metzger@paluno.uni-due.de)*


**Main reference** Andreas Metzger, Tristan Kley, Aristide Rothweiler, Klaus Pohl: "Automatically reconciling the trade-off between prediction accuracy and earliness in prescriptive business process monitoring", Inf. Syst., Vol. 118, p. 102254, 2023.
**URL** http://dx.doi.org/10.1016/J.IS.2023.102254

Prescriptive business process monitoring aims to guide process managers on adapting processes to avoid negative outcomes.

A key challenge is to balance prediction accuracy with tardiness: earlier predictions provide more time for adaptations, but are often less reliable.

This talk explores existing methods for managing this trade-off and compares their performance by using real-world datasets. By evaluating their cost savings, this talk identifies factors influencing effectiveness and provides practical recommendations for selecting appropriate approaches.

Based on these insights, the talk explores directions & opportunities for future research, involving emerging topics, such as explainable AI and Large Language Models (LLMs).

See: `https://doi.org/10.1016/j.is.2023.102254`

## 3.8 Foundations of Agentic AI: GenAI meets Strategic Reasoning and Planning

*Giuseppe De Giacomo (University of Oxford, GB, giuseppe.degiacomo@cs.ox.ac.uk)*



We are entering an era where businesses adopt AI agents to transform operations, drive impact across functions, and accelerate value realization. AI agents are software entities with goal-directed autonomy, capable of choosing and executing actions using AI techniques, see e.g., the 2025 IBM white paper "AI agents: Opportunities, risks, and mitigations". AI agents have long been studied in Reasoning about Actions (KR), Planning (ICAPS), and Autonomous Agents (AAMAS). Recently, LLM-based agents have emerged, combining language models with agency and showing strong potential of bringing about "open-endedness" and "common sense", as suggested by John McCarthy in his 1959 paper "Programs with Common Sense". Agentic AI systems combine AI agents with tools, planners, memory, and data to pursue goals autonomously, often including hardware control. In this talk, we discuss various opportunities and challenges that Agentic AI could deliver.

# 4 Working Groups

Seminar facilitation methodology adhered to an adapted design-thinking style. Overall, the agenda included four major phases: talks-&-ideation, working groups, presentations, and writ-a-thon (see Figure 1).

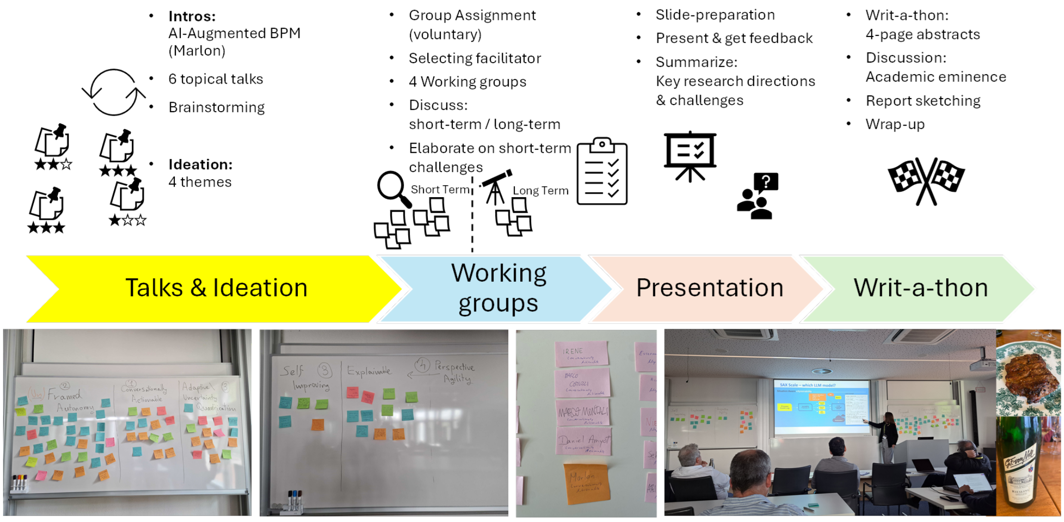


**Figure 1** “Design-thinking”-like seminar structure.

## 4.1 Working Group on Framed Autonomy

*Giuseppe De Giacomo (University of Oxford, GB, giuseppe.degiacomo@cs.ox.ac.uk)*
*Andrea Marrella (Sapienza Università di Roma, IT, marrella@diag.uniroma1.it)*
*Yves Lesperance (York University – Toronto, CA, lesperan@eecs.yorku.ca)*
*Andrea Matta (Politecnico di Milano, IT, andrea.matta@polimi.it)*
*Timotheus Kampik (SAP Berlin, DE & Umeå University, SE, timotheus.kampik@sap.com)*
*Diego Calvanese (Free University of Bozen-Bolzano, IT, diego.calvanese@unibz.it)*



Framed autonomy refers to agents acting and interacting with maximal flexibility within potentially dynamic *frames* consisting of rules, restrictions, and regulations. With the increased deployment of AI-based technologies – recently and most notably large language models – framing the autonomy of agents that enact business processes can be expected to be a key challenge. In this working group report, we sketch problem scenarios and provide a conceptual architecture for framing autonomy in business processes. We highlight a list of practical challenges for the framing of autonomous business process behaviors and conclude with the sketch of a research roadmap.

### 4.1.1 Introduction

Software systems providing the operational backbone of organizations are becoming increasingly autonomous [1], partially driven by advances in deep learning-based technologies such as Large Language Models (LLMs) [2]. While in this context, *autonomy* may pertain to an overall and potentially tightly coupled system, the distributed and complex nature of large organizations requires intelligence at the level of autonomous submodules, reflecting how intelligent business decisions are made by humans. In order to deploy autonomous software agents safely and effectively, one must ensure that they comply with normative requirements, while still utilizing their substantial degrees of autonomy to accomplish their goals to the best possible extent [3].

As abstractions for managing guardrails, we propose the notion of (normative) *frames* that – in contrast to the more operational notions of declarative or procedural business processes and rules – focus only on deontic requirements of how organizations should run. Frame representation and reasoning can draw from a wealth of research on deontic logic [4, 5], temporal reasoning [6], planning [7, 8], and normative multi-agent systems [9]. We provide informal definitions of frames and position them in the context of related abstractions, and sketch scenario types describing how frames can be applied to agents enacting business processes. These partially subsymbolic AI agents must then be augmented with symbolic capabilities for synthesizing plans that guarantee frame compliance, reasoning about their own, others', and process-level goals in order to maximize objectives *within* the frames. Accordingly, on a fundamental level these agents require capabilities for plan and behavior synthesis [10, 11, 12, 13, 14], as well as for goal reasoning [15]. We also highlight a list of practical challenges that require solving to (better) utilize frames in large organizations. Considering these challenges, we outline a research roadmap for framing autonomous business process execution.

As our motivating example, we introduce a simple excerpt of a (fictional) order-to-cash process (Figure 2). A customer sends a natural language wishlist to a retailer, which then generates a symbolic order proposal from it. If all items in the proposal are available, the proposal is sent to the customer. If some items are not available, a gift is added to the proposal. The retailer chooses between chocolate or wine as the gift and sends the gift-augmented proposal to the customer. Upon receiving the answer, the retailer assesses whether the answer is positive or negative. In the former case, the retailer acknowledges the acceptance of the proposal; in the latter, the retailer generates a new order proposal. However, in total not more than three proposals will be generated for a given incoming request. The process – i.e., the *retailer* pool – is enacted by one or several agents that can autonomously generate proposals (including alternative proposals) and decide on the type of gift to add to a proposal. The objective of the process is to maximize the margins of order proposals that are accepted by customers, also considering the cost of gifts.

Some frame constraints for this example are:

- If the customer is underaged, do not add wine as the gift (i.e., a gift containing alcohol). This constraint can be easily specified using linear temporal logic and verified on finite traces ($LTL_f$).
- Once an item has been added to an order proposal, the item's price must not increase in a subsequent proposal. This constraint is harder to verify, as it requires reasoning over quantities over time.

Notice that we may have a single agent taking these decisions, or multiple agents, with different agents taking decisions in various moments during the execution of a single process instance. This gives rise to different scenarios that are discussed in the next section.

### 4.1.2 Framed Autonomy in Business Processes

Framed autonomy requires that the autonomous system operates within its current frame. Intuitively, a frame is a set of rules, restrictions, and regulations, which may evolve over time. The frame establishes the boundaries within which the system may operate with maximal flexibility, making autonomous decisions. In BPM, frames may exist – at least – on *agent type*, *process*, and *organization* level (as well as potentially across organizations).

More analytically, frames are *normative*, i.e., they specify deontic requirements to the process; in contrast, classical process specification languages, such as BPMN and DECLARE are *operational*, i.e., they specify behavior required to accomplish a business goal. In terms

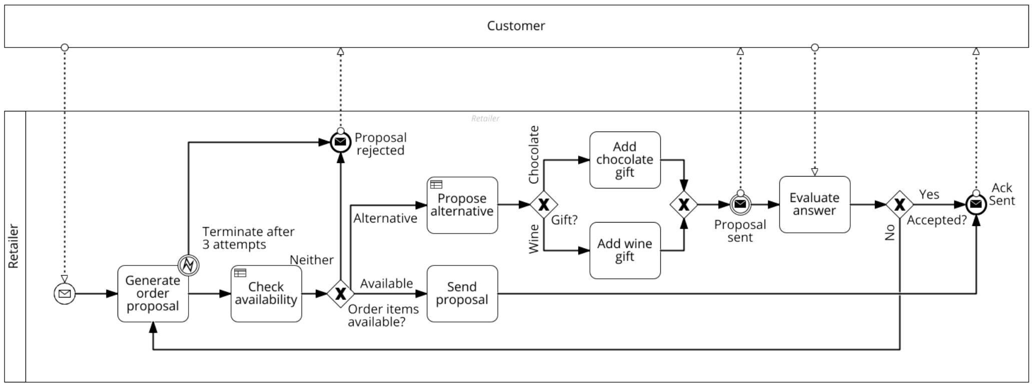


**Figure 2** An excerpt of an order-to-cash process, in which a retailer turns a natural language wishlist into an order proposal.

of strategies, operational specifications require choosing a strategy to achieve a goal, while the frame requires, in principle, identifying the entire set of strategies that are compatible with the norm and then ensuring that the strategy chosen for the goal is one of them.

Notice that sometimes the operational specifications have been called *frames* as well [1]. Indeed, they can be considered a sort of operational frame. In this document, however, we our focus of "frames" is on the normative specification. When we need to distinguish, we call the two frames *normative frame* and *operational frame*, respectively.

Observe that if there are no choices to be made (no autonomous decision-makers), then the normative frame is just an additional condition over the operational frame; but if decision-making is possible then the operational frame requires finding a strategy to satisfy the objective, whereas the normative frame requires choosing a strategy that remains within what is allowed (with respect to the frame).

Strategies for achieving goals under framed autonomy are associated with decision-makers, including software agents. This gives rise to several problem setups, reflecting centralized as well as distributed intelligence.

**Centralized intelligence.** We consider the "AI agents" as a single entity that orchestrates the *process* that is executed in a mutually fully observable and coordinated manner. The *environment* may be stochastic and not fully observable. The frame is over the process. The single entity may have active or passive responsibility for the frame. If we have multiple agents we may break down the problem into several of the above scenarios.

**Distributed intelligence.** We consider AI agents as distributed entities that enact the process as resources. This has wide-ranging implications:
- A resource has partial observability of what other resources do.
- Coordination may be effortful and not always possible.
- Agents may have individual goals that may not be consistent with process-level goals.
- We can frame resources or the entire process.
- We need to assign responsibility to individual agents or groups thereof, and there may be strategic interactions affecting responsibility.

From these problem setups, we can derive three different blueprint scenarios for framed autonomy in business processes (Figure 3):
1. We have a single decision-maker and place a frame on process behavior.
2. We have multiple decision-makers and place frames on individual decision-makers.
3. We have multiple decision-makers and place frame(s) on process behavior or parts thereof.

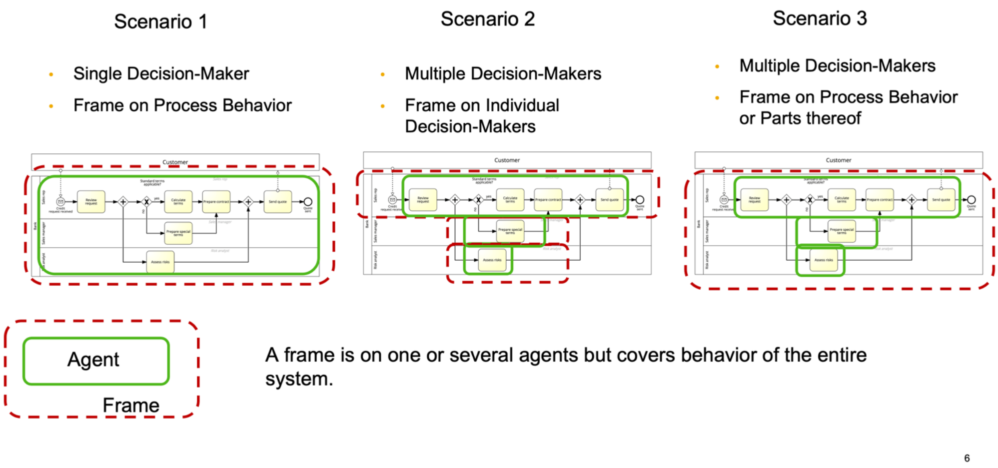


**Figure 3** Different scenarios of framed autonomy in business processes.

In practice, there may be additional variance to the scenarios. For example, normative frames may be partially represented within operational process specifications, restricting overall agent autonomy. An example is a purchasing process where purchase orders can only be created and paid through a central IT system that enforces normative rules, e.g. regarding four-eyes approval policies. Other parts of the global normative frame can potentially be projected to local agent-level norms. For example, overall spending limits may apply on the global level, but could be operationalized locally. Also, from a process perspective, operationalizing some frame constraints (such as the spending limits from above) may require very broad case notions; e.g., a case may be *all purchases executed by a specific business unit within a given month*.

### 4.1.3 Practical Challenges

Achieving framed autonomy in business processes comes with practical challenges. Below, we list (and briefly discuss) three that we consider of particular importance.

1. *What is a pragmatic notion of an agent in the context of business process execution?* Before the broad adoption of LLMs, the notion of an agent did not play a major role in the engineering of business information systems and the processes that run them. Consequently, practitioners cannot be expected to be familiar with the depth and sophistication of agent-related abstractions. To the contrary, a practitioner may consider as an agent a software tool that makes use of an LLM, without much thought about further properties. Defining a more precise and robust notion of an agent that is still intuitively understandable by business process practitioners can thus be considered a key prerequisite. See also Working Group Report 4.2.
2. *How to elicit and specify frames?* Next, approaches for eliciting and specifying frames need to be devised. This requires a meta-model for frames, and one or several specification languages. To this end, existing specification languages can be reused; potentially several languages and their underlying concepts can be combined. For example, declarative

approaches to process specification – such as DECLARE [16] and in more practical contexts business rule and query languages with temporal reasoning capabilities [17] – can be augmented with deontic notions in order to promote normativity to a first-class abstraction. For elicitation, both symbolic and subsymbolic approaches can be used and fused. LLMs can generate frames or parts thereof from natural language text, whereas rule mining approaches can be applied to infer normative constraints from the traces of well-behaved agents and multi-agent systems. See also Working Group Report 4.3.

3. *How to operationalize frames on real-world symbolic data?* Once specified, frames need to be integrated with business information systems, to ensure systems' frame-compliance during runtime. A short- to mid-term prerequisite is the operationalization of frames using technologies that do in fact run in large organizations. Here, explainability is a necessity, considering the practical intricacy of normative requirements, as well as the scale of real-world symbolic queries and data. See also Working Group Report 4.4.

### 4.1.4 Research Roadmap

Given our conceptual architecture of frames for autonomous business process execution, as well as the practical challenges outlined above, we define a research roadmap. The roadmap is divided into four broad categories, each of which comes with a set of simply phrased "how to?" research questions.

#### 4.1.4.1 Frame Representation & Reasoning

The key prerequisite for applying frames to autonomous business process execution is devising ways to represent them and reason about them. Accordingly, our questions are the following:

- *How to combine declarative and procedural paradigms when specifying frames?*
- *How to decide what to model in a trace view versus what to model in a transition system (with choice points)?*
- *How to model agents and integrate agent models with process- and organization-level models?*
- *How to operationalize frames in real-world information systems?*
- *How to ensure responsibility and accountability with respect to outcomes as well as to deontic notions such as obligations, permissions, and prohibitions?*

#### 4.1.4.2 LLMs and Framed Autonomy

The emergence of LLMs as widely applicable natural language processing tools has led to the re-emergence of agents as mainstream abstractions in information systems engineering. Accordingly, a key problem is the assurance of normative compliant LLM agent behavior (frames for LLMs), as well the application of LLMs to generate both normative frames and frame-compliant operational behavior specifications (LLMs for frames). We consider the former direction more relevant than the latter, reflecting the overall objective of frames to maximize autonomy while still ensuring compliance. However, both are covered by the questions below:

- *How to implement LLM agents that can comply with frames?*
- *How to symbolically augment LLM agents to ensure compliance with frames?*
- *How to generate frames from sources in several modalities using LLMs, as well as symbolic approaches?*
- *How to elicit goals in business processes with LLMs?*
- *How to leverage LLMs as tools that help us reason about frames?*

#### 4.1.4.3 Goal Reasoning

While goals are central to informal definitions of business process notions, they are typically not treated as first-class abstractions: the goal of a process or instance thereof is assumed to be implicitly defined by the behavioral specification. However, considering that agents enact processes autonomously, behavioral specification alone is insufficient, as these specifications are necessarily at least partially synthesised from agent and process goals. Accordingly, a range of research questions about goal reasoning in business processes emerge, such as:

- *How to elicit goals in business processes?*
- *How to align agent-level and process-level goals?*
- *How to represent goals in business processes?*
- *How to manage goals (e.g., instantiate, drop, revise, prioritize) in business processes?*
- *How to anticipate future goals?*
- *How to maintain most options open in anticipation of future goals when strategizing and acting?*

#### 4.1.4.4 Meta Frames and Reframing

Finally, we are interested in representing and reasoning about multiple frames, as well as as frames that change:

- *How to compose and decompose different process frames?*
- *How to adopt meta-level frames, e.g., as provided by third-party organizations?*
- *How to navigate through different frames?*
- *How to manage conflicts between frames?*
- *How to revise frames over time?*
- *How to verify that the entity revising the frame has the authority to do so?*

### 4.1.5 Conclusions

When autonomy is included in a business process execution system, the notion of *normative frame* becomes essential to guardrail autonomous decision-making. Normative frames have a deontic nature and are concerned with the sets of strategies that an agent can choose from while satisfying the frame. Accordingly, when goal-oriented agents synthesize their operational strategies, these strategies are implicitly mapped to those at the normative level and checked against the frame. AI agents – whether based on symbolic or subsymbolic methods – that enact business processes must be able to synthesize such strategies so that frame compliance can be guaranteed and exceptional violations can be justified.

**References**

1 Marlon Dumas, Fabiana Fournier, Lior Limonad, Andrea Marrella, Marco Montali, Jana-Rebecca Rehse, Rafael Accorsi, Diego Calvanese, Giuseppe De Giacomo, Dirk Fahland, Avigdor Gal, Marcello La Rosa, Hagen Völzer, and Ingo Weber. Ai-augmented business process management systems: A research manifesto. *ACM Trans. Manag. Inf. Syst.*, 14(1):11:1–11:19, 2023.

2 Timotheus Kampik, Christian Warmuth, Adrian Rebmann, Ron Agam, Lukas Egger, Andreas Gerber, Johannes Hoffart, Jonas Kolk, Philipp Herzig, Gero Decker, Han van der Aa, Artem Polyvyanyy, Stefanie Rinderle-Ma, Ingo Weber, and Matthias Weidlich. Large process models: A vision for business process management in the age of generative ai. *KI - Künstliche Intelligenz*, 2024.

**3** Timotheus Kampik, Adnane Mansour, Olivier Boissier, Sabrina Kirrane, Julian Padget, Terry R. Payne, Munindar P. Singh, Valentina Tamma, and Antoine Zimmermann. Governance of autonomous agents on the web: Challenges and opportunities. *ACM Trans. Internet Technol.*, 22(4), 2022.

**4** Dov Gabbay, John Horty, Xavier Parent, Ron Van der Meyden, Leendert van der Torre, et al. *Handbook of Deontic Logic and Normative Systems, Volume 1.* College Publications, 2013.

**5** Dov Gabbay, John Horty, Xavier Parent, Ron Van der Meyden, Leendert van der Torre, et al. *Handbook of Deontic Logic and Normative Systems, Volume 2.* College Publications, 2021.

**6** Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI*, pages 854–860, 2013.

**7** Giuseppe De Giacomo and Sasha Rubin. Automata-theoretic foundations of FOND planning for ltlf and ldlf goals. In *IJCAI*, pages 4729–4735. ijcai.org, 2018.

**8** Malik Ghallab, Dana S. Nau, and Paolo Traverso. *Acting, Planning and Learning.* Cambridge University Press, 2025.

**9** A. Chopra, L. van der Torre, and H. Verhagen. *Handbook of Normative Multiagent Systems.* College Publications, 2018.

**10** Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.

**11** Bernd Finkbeiner, Felix Klein, and Niklas Metzger. Live synthesis. *Innov. Syst. Softw. Eng.*, 18(3):443–454, 2022.

**12** Giuseppe De Giacomo and Moshe Y. Vardi. Synthesis for LTL and LDL on finite traces. In *IJCAI*, pages 1558–1564, 2015.

**13** Shufang Zhu and Giuseppe De Giacomo. Act for your duties but maintain your rights. In Gabriele Kern-Isberner, Gerhard Lakemeyer, and Thomas Meyer, editors, *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel, July 31 - August 5, 2022*, 2022.

**14** Shufang Zhu and Giuseppe De Giacomo. Synthesis of maximally permissive strategies for $\text{LTL}_f$ specifications. In *IJCAI*, pages 2783–2789. ijcai.org, 2022.

**15** David W. Aha. Goal reasoning: Foundations, emerging applications, and prospects. *AI Mag.*, 39(2):3–24, 2018.

**16** Claudio Di Ciccio and Marco Montali. Declarative process specifications: Reasoning, discovery, monitoring. In Wil M. P. van der Aalst and Josep Carmona, editors, *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*, pages 108–152. Springer, 2022.

**17** Timotheus Kampik and Cem Okulmus. Expressive power and complexity results for signal, an industry-scale process query language. In Andrea Marrella, Manuel Resinas, Mieke Jans, and Michael Rosemann, editors, *Business Process Management Forum - BPM 2024 Forum, Krakow, Poland, September 1-6, 2024, Proceedings*, volume 526 of *Lecture Notes in Business Information Processing*, pages 3–19. Springer, 2024.

## 4.2 Working Group on Adaptive / Uncertainty Quantification

*Estefania Serral (KU Leuven, BE, estefania.serralasensio@kuleuven.be)*
*Achiya Elyasaf (Ben-Gurion University of the Negev – Beer Sheva, IL, achiya@bgu.ac.il)*
*Andreas Metzger (paluno – The Ruhr Institute for Software Technology, University of Duisburg-Essen, DE, andreas.metzger@paluno.uni-due.de)*
*Niek Tax (Meta – London, GB, niek@meta.com)*
*Sebastian Sardina (RMIT University – Melbourne, AU, sebastian.sardina@rmit.edu.au)*
*Arik Senderovich (York University, Toronto, CA, sariks@yorku.ca)*



### 4.2.1 Introduction

In today's dynamic digital landscape, business processes (BPs) are expected to operate autonomously and adapt their structure and behavior to evolving goals, conditions, or constraints [4]. Such systems are known as *autonomous business process systems* (ABPs), whose key capability is self-modification.

Modifications in ABPs fall into two main categories [21]:

- **Adaptation**: Short-term, instance-specific changes to handle unforeseen conditions without altering the overall process schema. These real-time adjustments ensure continuity and resilience, for example by rerouting workflows, reallocating resources, or integrating new data sources during execution.
- **Evolution**: Long-term changes to process logic or models, affecting all future instances. These deliberate updates are driven by recurring issues or strategic shifts and may involve redesigning decision logic or updating policies.

Without these capabilities, ABPs risk brittleness and reduced responsiveness. Thus, engineering self-modifying capabilities is essential for robust, flexible process management.

The importance of adaptive and evolving BPM systems has long been recognized [3, 21, 11], with prior work exploring both adaptation [16, 20, 10, 6, 23, 26] and evolution [12, 28, 27].

In this report, we provide:

- A definition of self-modification in ABPs, distinguishing adaptation and evolution.
- A structured view of the dimensions, goals, and triggers of modifications in ABPs.
- An outline of key challenges in governance, uncertainty, and continuous learning.

### 4.2.2 Running Example: Automated Warehouse

To illustrate self-modifying ABPs, consider an automated warehouse where fleets of robots transport shelves to human pickers across a high-throughput environment.

If a robot fails during the busy holiday season, nearby robots reroute around the blockage, and the order is reassigned. Workers receive updated instructions, ensuring minimal disruption. This short-term response (*adaptation*) resolves the immediate issue without changing the overall process model.

The system also logs the failure and, after noticing similar breakdowns after about 1,000 picks, introduces a maintenance rule requiring inspection every 900 operations. This mid-term, model-level change (*evolution*) refines policies to prevent recurrence.

Together, these examples show how ABPs span from localized adaptations to deliberate evolution, implemented with varying degrees of automation.

### 4.2.3 Types of Modifications in ABPs

Modifications in ABPs can be classified along several dimensions:

#### D1: Adaptation vs. Evolution

We introduced this central dimension already in Section 1, referring to adaptations as short-term, instance-specific modifications, and evolution as long-term, multi-instance modifications of the process logic and/or model itself [22].

The concept of adaptation in ABPs strongly connects to work in the software engineering community on self-adaptive software systems [32]. Here, the MAPE-K loop has established itself as a widely adopted conceptual model [1, 2]. This model is depicted in Figure 4.

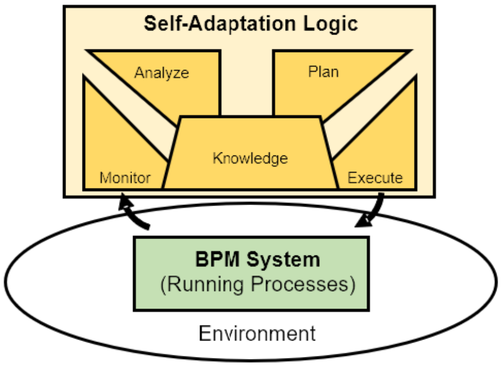


**Figure 4** Conceptual model of self-adaptation.

The self-adaptation logic is structured into four main activities (MAPE) that leverage a shared knowledge base (K). The knowledge base includes adaptation goals (requirements), strategies, and rules. The four activities are:

- **Monitoring**: Observing the system and its environment via sensors.
- **Analysis**: Interpreting monitoring data to determine the need for adaptation.
- **Planning**: Deciding on adaptation actions.
- **Execution**: Implementing adaptation actions via actuators.

The concept of evolution in ABPs, in contrast, focuses on systematic, long-term modifications that affect the process model and logic across multiple future instances. Rather than responding to immediate runtime conditions, evolution is driven by aggregated insights from monitoring and analysis over time – such as recurring performance issues, changes in strategic direction, or compliance updates. These insights inform deliberate revisions of process design, typically through:

- Assessment of existing outcomes.
- Planning of structural or behavioral improvements.
- Implementation of changes.
- Validation to ensure alignment with long-term goals.

While adaptation enables resilience and responsiveness, evolution ensures strategic alignment, sustainability, and optimization.

### D2: Task vs. Flow vs. Process

Another critical dimension concerns what is modified [22]:
- **Task-level changes**: Modify how individual tasks are performed or configured (e.g., adjusting duration, logic, input/output data, or resource assignments).
- **Control Flow-level changes**: Adjust control flow or routing of tasks (e.g., rerouting orders, skipping steps).
- **Process-level changes**: Alter the structure or central resources of the entire process (e.g., introducing new roles, shifting coordination logic, or replacing subsystems).

These levels often interact. For example, rerouting in the warehouse involves flow-level change, while a new rule for reassigning malfunctioning robots might affect task performance. Revising the maintenance schedule constitutes a process-level change. This dimension also affects implementation complexity: task-level changes can often be handled locally, while process-level changes typically require coordination across components.

### D3: Reactive vs. Proactive

Another important distinction concerns the trigger of modifications:
- **Reactive**: Occur in response to specific events or failures, such as a robot malfunction triggering rerouting.
- **Proactive**: Initiated based on forecasts or insights from predictive analytics. Examples include predictive and prescriptive business process monitoring [17], such as scheduling preventive maintenance or restricting robot access during peak loads.

### D4: Human-Driven vs. Autonomous

Changes may be:
- **Human-driven**: Users detect, decide, and implement modifications.
- **Autonomous**: The system identifies issues and enacts changes independently.

In the warehouse example, creating a maintenance rule might initially be human-driven, but an advanced ABPS could learn similar patterns and implement them autonomously.

### D5: Planned vs. Emergent

- **Planned**: Result from deliberate, top-down redesign (e.g., rolling out a maintenance policy across sites).
- **Emergent**: Arise bottom-up, as patterns learned from execution lead to adaptations or eventual model evolution.

An ABPS that generalizes from local failure logs to propose global improvements exemplifies emergent capability.

### Characterizing the Running Example

All key dimensions play out in the warehouse scenario:
- Rerouting around a malfunctioning robot is a short-term, instance-level *adaptation*.
- Introducing a maintenance rule is a deliberate, model-level *evolution*.
- Rerouting is *reactive* when triggered by an event but can be *proactive* if anticipated via predictive monitoring.
- Initially, the maintenance policy is *human-driven* and *planned*.
- Over time, an advanced ABPS could autonomously propose and implement similar rules, making the change *emergent*.

### 4.2.4 Toward Self-modifying ABPS

Today's business process systems typically operate at an augmented level of autonomy, where intelligent components assist human workers but do not independently drive process execution or change. To move toward truly autonomous business process systems (ABPS), we propose a structured roadmap of autonomy levels, inspired by the SAE J3016 standard for driving automation [24]. While Sheridan's Levels of Automation [29] offer an alternative, they focus on isolated task automation rather than holistic process-level behavior, making them less suitable for ABPS.

We define autonomy levels as follows:

- **Level 0: No Automation** – Execution and orchestration of all tasks are fully manual.
- **Level 1: Process Assistance** – The system provides recommendations or highlights anomalies to human workers (e.g., predictive monitoring [18]).
- **Level 2: Partial Autonomy** – The system independently executes isolated tasks (e.g., call routing, task assignment) within predefined boundaries, but without contextual or adaptive behavior.
- **Level 3: Contextual Autonomy** – The system autonomously performs most tasks and orchestrates flows in a context-aware manner, requiring human intervention only in exceptional cases.

While most current systems lie at Level 1 or 2, we envision Level 3 as the target state for ABPS. Achieving this level necessitates the development of self-modifying capabilities. Specifically, a Level 3 ABPS must:

1. Detect changes in the operating environment (e.g., concept drift detection [8, 14, 25]).
2. Decide whether the detected change requires adaptation or evolution.
3. Select an appropriate modification strategy based on goals, context, and system history.
4. Learn from prior adaptations and generalize successful strategies.
5. Communicate its decisions, rationale, and uncertainty to human stakeholders.

#### Goals and Capabilities Across Autonomy Levels

The extent and nature of these capabilities depend on the object of modification: task, flow, or process. For instance, at Level 1, task-level modifications might involve recommending better parameter configurations; at Level 3, the system may autonomously reassign or skip tasks. Similarly, flow-level autonomy ranges from highlighting bottlenecks (Level 1) to real-time rerouting (Level 3), while process-level autonomy progresses from alerting on coordination issues to full reconfiguration of goals, roles, or policies.

In the warehouse scenario, task-level changes may involve altering how picking or navigation tasks are performed; flow-level changes may include rerouting due to blocked paths; and process-level autonomy could entail adjusting global maintenance policies. Thus, progressing toward Level 3 autonomy requires integrating sensing, reasoning, learning, and communication across abstraction levels while minimizing reliance on human oversight.

### 4.2.5 Challenges in Enabling Autonomous Modifications

In this section, we discuss three challenges related to governance and human oversight, continual learning for adaptation management, and uncertainty quantification and communication.

**Table 1** Goals and capabilities for autonomy at different levels and modification objects.

| Object of Modification | Level 1: Process Assistance | Level 2: Partial Autonomy | Level 3: Contextual Autonomy |
|---|---|---|---|
| **Task** | Recommending task performers or configurations (e.g., duration, cost). | Automating task assignment or execution in narrow scopes. | ABPS decides and executes full task reconfiguration autonomously, with human fallback only in edge cases. |
| **Flow** | Suggesting alternative paths or detecting bottlenecks. | Automating routing based on real-time conditions or rules. | Rerouting and dynamically altering execution paths with learned policies under uncertainty. |
| **Process** | Flagging process-wide issues (e.g., coordination delays). | Automating subprocesses, such as resource pooling or exception handling. | Reconfiguring processes, modifying policies or goals, and coordinating across stakeholders with minimal human oversight. |

### Challenge 1: Governance, Oversight, and Human Interaction

To operate safely, ABPs must embed governance and human oversight. A central question is when to refrain from automation and return control to humans, especially in ambiguous or high-risk situations. AI planning and runtime monitoring can help define thresholds for escalation, while prediction with a reject option [7] offers approaches for "learning to defer."

Determining how users validate proposed changes requires explainable adaptation, where the system justifies its modifications. Methods from explainable AI [5] or simulations of expected outcomes can support transparency.

**Table 2** Challenge mapping across modification levels.

| Object of Modification | Governance & Oversight | Continuous Learning & Adaptation Mgmt. | Modeling & Uncertainty |
|---|---|---|---|
| **Task** | When should execution shift from autonomous to manual? | How can the system learn improved task performance? | How confident is the system in altering task strategies? |
| **Flow** | Who approves new routing decisions? | Can routing rules generalize without overfitting? | What risks arise when rerouting or skipping tasks? |
| **Process** | When must humans approve structural redesign? | How to record and evolve process-level modifications? | How is uncertainty modeled across processes? |

Aligning ABP goals with those of human stakeholders remains complex. Multi-objective optimization [15] can help balance performance, cost, compliance, and satisfaction, but resolving conflicting objectives is difficult.

In fully autonomous scenarios, ABPs need internal mechanisms to evaluate whether adaptations succeed. Techniques such as anomaly detection, causal reasoning, and formal verification are promising. LLMs may also support human audits by generating decision justifications [33].

Human input for validation is often costly, motivating efficient solutions like active testing [13]. Lastly, systems must be able to assess whether they have sufficient context to act safely or should defer decisions.

**Core Research Questions:**

- When should control shift between autonomous processes and humans in self-modifying systems?
- How can user validation be incorporated in real-time adaptation without bottlenecking autonomy?
- How can systems optimize multiple objectives while remaining within formal and ethical constraints?
- How can ABPs evaluate the success or failure of modifications when human validation is unavailable?
- What are the data quality and coverage thresholds for safe, autonomous decision-making?

### Challenge 2: Continuous Learning and Adaptation Management

A core capability of self-modifying ABPs is learning from experience to improve behavior over time. This requires continuously recording adaptations – reactive and proactive – and assessing their impact both per instance and across instances. Such meta-knowledge enables feedback loops where effective modifications are reinforced and ineffective ones discarded. AI planning and reinforcement learning [20], especially when combined with causal modeling, are promising for managing these loops.

Evaluating generalization poses another challenge. An adaptation that works once may fail under different conditions. Systems must infer when learned strategies apply and calibrate their confidence accordingly. For example, if a task is often skipped under heavy load, the ABP may learn to skip it when similar patterns arise, but must avoid overgeneralizing. Formal guarantees and validation frameworks are needed to ensure safe transfer of learned policies. This also relates to interpretability, as stakeholders expect explanations for behavior shifts.

Finally, ABPs operating over long periods or at high volume face limits of memory and context. They must build bounded knowledge representations – such as compressed summaries, predictive abstractions, or sliding windows of relevant history. Techniques like stream reasoning, process mining, and transformer-based models can help. Deciding what to remember, forget, or query is essential for scalable, continuous learning.

**Core Research Questions:**

- How can ABPs continuously record adaptations and assess their effectiveness over time?
- What metrics and techniques enable safe generalization of learned behavior across varying contexts?
- How can bounded knowledge representations be maintained for long-running or high-frequency processes?

#### Challenge 3: Modeling and Measuring Uncertainty

ABPs must handle both aleatoric (inherent randomness) and epistemic (lack of knowledge) uncertainty [9]. Aleatoric uncertainty can be modeled with probabilistic distributions over durations and outcomes, while epistemic uncertainty often requires exploration or additional sensing to reduce ambiguity. Identifying which type is present guides whether to gather more data or hedge decisions probabilistically. Probabilistic modeling, Bayesian inference, and fuzzy logic are key tools for representing uncertainty.

Quantification combines qualitative thresholds set by experts and quantitative measures like Bayesian models or ensemble predictions [16]. ABPs must integrate both, relying on statistical models when data are sufficient and heuristics when information is limited. Hybrid approaches that combine fuzzy logic and ensemble learning can improve robustness in uncertain domains.

Uncertainty modeling links closely to Challenges 1 and 2. For Challenge 1, quantifying uncertainty informs when to defer to humans [7]. For Challenge 2, accurate epistemic estimates are essential for learning adaptive policies efficiently [19, 30].

Finally, communicating uncertainty effectively is critical. Operators must grasp what the system plans to do and how confident it is. Explainable AI [31], visual confidence intervals, and natural language outputs can help convey this information, along with potential impacts on risks and service quality.

**Core Research Questions:**

- How can ABPs differentiate between epistemic and aleatoric uncertainty during execution?
- How can qualitative and quantitative uncertainty metrics be combined for robust decision-making?
- What are effective representations for modeling uncertainty using probabilistic or fuzzy paradigms?
- How should ABPs communicate uncertainty and associated risk to users in a transparent and actionable way?

### 4.2.6 Conclusion and Outlook

In this paper, we have laid the conceptual groundwork for understanding and engineering self-modifying capabilities in autonomous business process systems (ABPS). We defined what it means for an ABPS to self-modify, distinguished between the short-term reactivity of adaptation and the long-term reconfiguration of evolution, and introduced a structured framework for levels of business process autonomy. We further mapped these autonomy levels across different objects of modification – task, flow, and process – and connected them to system goals and capabilities.

As ABPS strive toward higher levels of autonomy, we identified three core research challenges that must be addressed to enable safe, intelligent, and human-aware self-modification:

1. Establishing robust governance mechanisms and human oversight.
2. Managing continuous learning to support sustainable and generalizable adaptations.
3. Modeling and communicating uncertainty in ways that foster trust and informed decision-making.

Together, these challenges highlight a critical insight: autonomy is not simply about replacing humans but about redesigning systems that can responsibly decide when to act, when to adapt, and when to defer. The path forward requires integrating techniques from AI planning, machine learning, explainable AI, causal inference, process mining, and human-computer interaction into coherent architectures that balance autonomy with accountability.

We envision ABPS of the future not as black-box automation engines but as collaborative agents capable of engaging with their environments and human counterparts in a transparent, contextual, and goal-aligned manner. To that end, we encourage the community to build benchmarks, share evaluation frameworks, and develop modular toolkits that bring us closer to truly self-modifying, trustworthy, and adaptive business processes.

## References


1. Y. Brun, G. Di Marzo Serugendo, C. Gacek, H. Giese, H. Kienle, M. Litoiu, H. Müller, M. Pezzè, and M. Shaw. Engineering self-adaptive systems through feedback loops. In *Software Engineering for Self-Adaptive Systems*, pages 48–70. Springer, 2009.
2. IBM Computing. An architectural blueprint for autonomic computing. Technical Report 31, IBM White Paper, 2006.
3. P. Dadam and M. Reichert. The adept project: a decade of research and development for robust and flexible process support: challenges and achievements. *Computer Science - Research and Development*, 23:81–97, 2009.
4. M. Dumas, F. Fournier, L. Limonad, A. Marrella, M. Montali, J.-R. Rehse, R. Accorsi, D. Calvanese, G. De Giacomo, and D. Fahland. Ai-augmented business process management systems: a research manifesto. *ACM Transactions on Management Information Systems*, 14:1–19, 2023.
5. R. Dwivedi, D. Dave, H. Naik, S. Singhal, R. Omer, P. Patel, B. Qian, Z. Wen, T. Shah, G. Morgan, et al. Explainable ai (xai): Core ideas, techniques, and solutions. *ACM Computing Surveys*, 55:1–33, 2023.
6. D. Greenwood and R. Ghizzioli. Goal-oriented autonomic business process modelling and execution. In *Multiagent Systems*. IntechOpen, 2009.
7. K. Hendrickx, L. Perini, D. Van der Plas, W. Meert, and J. Davis. Machine learning with a reject option: A survey. *Machine Learning*, 113:3073–3110, 2024.
8. F. Hinder, V. Vaquet, and B. Hammer. One or two things we know about concept drift – a survey on monitoring in evolving environments. part a: detecting concept drift. *Frontiers in Artificial Intelligence*, 7:1330257, 2024.
9. E. Hüllermeier and W. Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110:457–506, 2021.
10. K. Jander, L. Braubach, A. Pokahr, W. Lamersdorf, and K.-J. Wack. Goal-oriented processes with gpmn. *International Journal on Artificial Intelligence Tools*, 20:1021–1041, 2011.
11. N. R. Jennings, P. Faratin, M. Johnson, T. J. Norman, P. O'Brien, and M. E. Wiegand. Agent-based business process management. *International Journal of Cooperative Information Systems*, 5:105–130, 1996.
12. H. Kir and N. Erdogan. A knowledge-intensive adaptive business process management framework. *Information Systems*, 95:101639, 2021.
13. J. Kossen, S. Farquhar, Y. Gal, and T. Rainforth. Active testing: Sample-efficient model evaluation. In *International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5753–5763, 2021.
14. J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang. Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31:2346–2363, 2018.
15. R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26:369–395, 2004.
16. A. Metzger, T. Kley, and A. Palm. Triggering proactive business process adaptations via online reinforcement learning. In *Business Process Management - 18th International Conference, BPM 2020*, volume 12168 of *Lecture Notes in Computer Science*, pages 273–290. Springer, 2020.

17 A. Metzger, T. Kley, A. Rothweiler, and K. Pohl. Automatically reconciling the trade-off between prediction accuracy and earliness in prescriptive business process monitoring. *Information Systems*, 118:102254, 2023.
18 A. E. Márquez-Chamorro, M. Resinas, and A. Ruiz-Cortés. Predictive monitoring of business processes: a survey. *IEEE Transactions on Services Computing*, 11:962–977, 2017.
19 V.-L. Nguyen, M. H. Shaker, and E. Hüllermeier. How to measure uncertainty in uncertainty sampling for active learning. *Machine Learning*, 111:89–122, 2022.
20 A. Palm, A. Metzger, and K. Pohl. Online reinforcement learning for self-adaptive information systems. In *Advanced Information Systems Engineering - 32nd International Conference, CAiSE 2020*, volume 12127 of *Lecture Notes in Computer Science*, pages 169–184. Springer, 2020.
21 M. Reichert and B. Weber. *Enabling Flexibility in Process-Aware Information Systems: Challenges, Methods, Technologies*, volume 54. Springer, 2012.
22 H. A. Reijers. Process design and redesign. In *Process-Aware Information Systems: Bridging People and Software through Process Technology*, pages 205–234. Wiley, 2005.
23 L. Sabatucci, C. Lodato, S. Lopes, and M. Cossentino. Towards self-adaptation and evolution in business process. In *AIBP@AI*IA*, pages 1–10, 2013.
24 SAE International. Taxonomy and definitions for driving automation systems for on-road motor vehicles. Technical report, SAE Recommended Practice J3016, 2021.
25 D. M. V. Sato, S. C. De Freitas, J. P. Barddal, and E. E. Scalabrin. A survey on concept drift in process mining. *ACM Computing Surveys*, 54:1–38, 2021.
26 E. Serral, J. De Smedt, M. Snoeck, and J. Vanthienen. Context-adaptive petri nets: Supporting adaptation for the execution context. *Expert Systems with Applications*, 42:9307–9317, 2015.
27 E. Serral, P. Valderas, and V. Pelechano. Supporting runtime system evolution to adapt to user behaviour. In *Advanced Information Systems Engineering: 22nd International Conference, CAiSE 2010*, volume 6051 of *Lecture Notes in Computer Science*, pages 378–392. Springer, 2010.
28 E. Serral, P. Valderas, and V. Pelechano. Addressing the evolution of automated user behaviour patterns by runtime model interpretation. *Software & Systems Modeling*, 14:1387–1420, 2015.
29 T. B. Sheridan. Adaptive automation, level of automation, allocation authority, supervisory control, and adaptive control: Distinctions and modes of adaptation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 41:662–667, 2011.
30 A. Tharwat and W. Schenck. A survey on active learning: State-of-the-art, practical challenges and research directions. *Mathematics*, 11:820, 2023.
31 D. Watson, J. O'Hara, N. Tax, R. Mudd, and I. Guy. Explaining predictive uncertainty with information theoretic shapley values. In *Advances in Neural Information Processing Systems*, volume 36, pages 7330–7350, 2023.
32 D. Weyns. *An Introduction to Self-Adaptive Systems: A Contemporary Software Engineering Perspective*. John Wiley & Sons, 2020.
33 L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623, 2023.

## 4.3 Working Group on Conversational Actionability

*Daniel Amyot (University of Ottawa, CA, damyot@uottawa.ca)*
*Marco Comuzzi (Ulsan National Institute of Science and Technology, KR, mcomuzzi@unist.ac.kr)*
*Marlon Dumas (University of Tartu, EE, marlon.dumas@ut.ee)*
*Marco Montali (Free University of Bozen-Bolzano, IT, marco.montali@unibz.it)*
*Irene Teinemaa (Google DeepMind – London, GB, iteinemaa@google.com)*



The discussions in this working group focused on the use cases that drive the requirement of *conversational actionability* in AI-Augmented Business Process Management Systems (ABPMS), the key functions an architecture of a conversationally actionable ABPMS, and various concerns and challenges related to the realization of such systems.

### 4.3.1 Definitions and Use Cases

An ABPMS is *conversationally actionable* if it can interact with users or external agents using a *conversational interface* to support, trigger, or guide actions related to business processes. Depending on the user and use case, this conversational interface can support communications using natural language, textual abstractions of process related artifacts, images and graphical models, or specific agent communication protocols.

We identify four principal use cases in a conversationally actionable ABPMS:

1. *Creation*. This use case is concerned with the elicitation of knowledge leading to the creation of business process related artifacts, such as imperative or declarative business process models, constraints, etc. The ABPMS creates the artifacts via several means, including from conversations with domain experts [1], from manually-created process documentation and maps, through discovery from event logs [2], or a combination thereof.
2. *Query*. This use case is concerned with answering questions about the past and current states of business processes and their executions. Besides process monitoring information, e.g., KPI dashboards, constraint violations, and related explanations, queries may also address design time concerns, e.g. process model updates and explanations of process changes. The queries can be issued by users using natural language and other communication means or by agents using some mutually agreed upon protocol. A query can return answers in natural language, as a graphical output (dashboard, annotated process model, etc.) or using a specific agent communication protocol.
3. *Recommendation*. This use case is concerned with generating recommendations for process instance adaptation and process evolution (future states). The ABPMS can provide such recommendations based on internal parametric knowledge or it can combine its knowledge with output retrieved from external *components*, such as simulators, optimizers, and solvers. Recommendations can be issued proactively by the ABPMS or requested by users or agents through a conversational interface. A recommendation can be presented to users in natural language or sent directly to an automation component to be directly implemented.
4. *Automation*. This use case is concerned with business process execution and evolution, which can be triggered by the user through natural language, automatically by an agent after the creation of a process-related artefact, or as the implementation of a

recommendation. Depending on the type of action to be taken, automation is delegated to a suitable component, e.g., an RPA bot for automating the execution of an activity or an agent updating the business rules of an ERP system.

Figure 5 maps the use cases to the phases of the traditional BPM lifecycle.

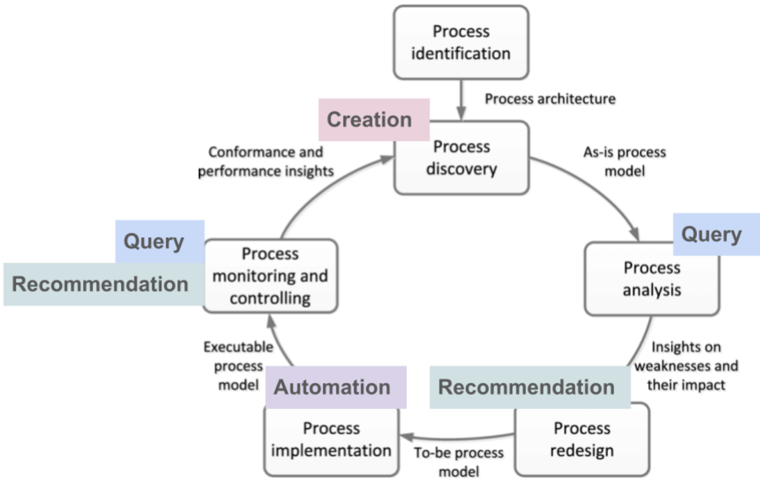


**Figure 5** BPM lifecycle annotated with use cases of a conversationally actionable ABPMS.

### 4.3.2 Architecture

We envision an architecture for a conversationally actionable ABPMS consisting of four subsystems (see Figure 6): a data layer, an intelligence layer, an action layer, and a conversational fabric.

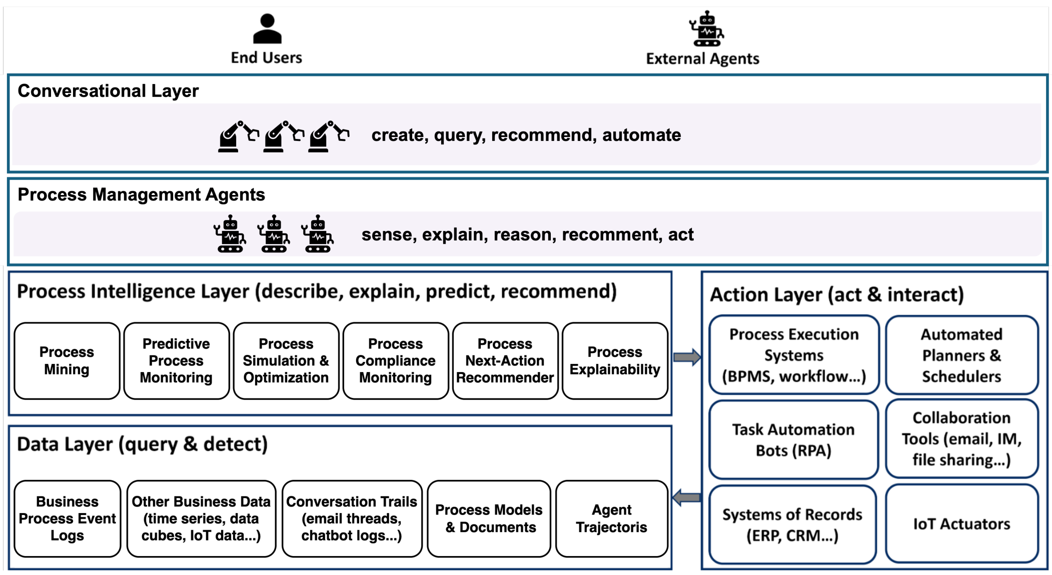


**Figure 6** Architecture of a conversationally actionable ABPMS.

The data layer brings together structured and unstructured data about a business operation, including event logs of business processes, other business data (financial data, time series), business process documentation (including text documents and process models), as well as data from physical systems relayed via IoT sensors. This layer allows the ABPMS to sense the current state and evolution of the business process and its environment. It provides a querying interface allowing other layers to leverage business process data and metadata from a variety of sources.

On top of the data layer sits a process intelligence layer, which provides capabilities for process discovery and performance analysis (process mining), predictive capabilities (e.g. based on simulation and machine learning), as well as prescriptive capabilities (recommending intervention). In other words, the process intelligence layer provides capability to describe and explain the current state of the process and to predict future states of the process under different scenarios.

Next to the process intelligence layer, the action layer provides capabilities for triggering actions affecting one or more business processes, or interactions with external actors (suppliers, customers, etc.). Example of actions include: (1) creating or altering the state of a case in a business process orchestrated by a Business Process Management System (BPMS); (2) triggering a software bot; (3) sending notifications via a communication platform; (4) updating records in a CRM, ERP or other System of Records.

The fourth component is the *conversational fabric*. This layer makes available the capabilities of the data layer, process intelligence layer, and action layer to different types of agents. The capabilities of the lower layers are exposed via *tools*. These tools are exposed to agents via a *Model Context Protocol* (MCP) [3] layer, which provides semantically rich descriptions of the tools for consumption by agents

An agent may leverage some of these tools to, for example, detect degradations in the performance of a process that may lead to a violation of a Service Level Agreement (SLA). Having detected this risk, the agent may leverage the process intelligence tools to determine interventions that may be triggered to prevent this SLA violation, and it may then leverage the tools coming from the action layer to trigger actions or notifications.

The agents in the conversational fabric receive instructions and goals from end users, directly via for example a chat interface, or indirectly via agents operating in other systems, such as an agent running in a CRM platform.

Some agents rely on general-purpose LLMs, others are based on fine-tuned or domain-specific models, and some possess explicit planning or reasoning capabilities. Each agent operates autonomously but may collaborate with other agents by passing tasks, sharing context, or requesting specific operations.

#### 4.3.3 On Actions and their Composition

Conversations do not only happen when external (human or artificial) agents interact with the ABPMS to obtain information, insights, and explanations about the current state or the history of the ABPM itself. They are also central to enable *actionability*, that is, to support actions in a broad sense. Actions may pertain to:

- The very execution of the process, e.g., approving a purchase order request, or sending a warehouse replenishment request.
- Process (re)framing, e.g., adding or removing constraints, and (re)modelling, e.g., evolving a model due to changing requirements.
- Decisions and interventions for process improvement, e.g., hiring a new resource to improve cycle time, or deciding to cancel an order because of economic considerations)

In this context, conversating is essential to gather information and insights before taking – or deciding whether to take – an action. For example, an agent would need to ponder the impact of hiring a new resource before committing to do so, or may decide to cancel an order depending on the corresponding penalty. This can only be achieved with good quality guarantees if the agent can conversate with other agents, as well as it can invoke the tools needed for the task at hand.

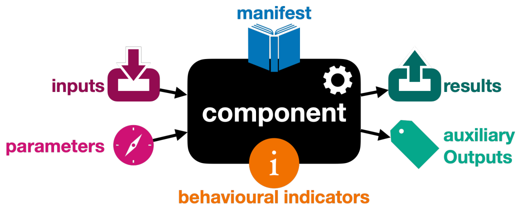


**Figure 7** A component with its inputs, outputs, and metadata.

In general terms, a *component* is conceived as a deterministic, verifiable unit of software that realises a certain task and/or produces some result given some inputs and preconditions. A tool may be a manually crafted software or it may rely on generative AI, but in any case, it is deemed to have gone through sufficient quality control to be trusted.

As shown in Figure 7, to enable agents in using components and interpreting the obtained results, some key aspects must be described:

- *Input* – an input object (possibly with some preconditions) necessary to invoke the component; for example, the process model required by a simulator.
- *Parameter* – an auxiliary input used to tune the component; for example, hyperparameters used to configure the simulator.
- *Result* – an output object produced by the component; for example, the process cycle time produced by the simulator.
- *Auxiliary output* – an auxiliary output produced by the component to help interpreting the output; for example, quantified uncertainty associated to a prediction.
- *Manifest* – a description of the component, its behaviour, and its inputs and outputs; the manifest relates to MCP, as its main purpose is to enable agents to discover and properly consume the component.
- *Behavioral indicator* – a quantitative or qualitative indicator characterising the (functioning of the) component along a key functional or non-functional concern (cf. Section 4.3.4), such whether the component produces analytical or approximate results, or whether it guarantees data privacy, whether it requires a rigid description of the input, and the like; relevant aspects that cannot be clearly described as indicators could still be included in the manifest of the component.

An agent may jointly employ multiple tools through different usage patterns, like concatenating two components to obtain a new functionality, or invoking two components in parallel and then aggregating their results. In addition, as pointed out before, it may conversate with other agents. This could lead to an autonomous decision taken by the agent relying on the result of the conversation and of the employed software components, or in an even more complex setting, resulting from collective decision making (akin to multi-agent systems [4]).

A crucial aspect of this architecture is *trust*. Agents are only deemed trustworthy if they can transparently explain their outputs. This is achieved by enabling agents to trace and articulate how they invoked specific components, what input they provided, and how the resulting outputs contributed to their decisions or responses have been used. This traceability, which relates to the more general concept of ABPMS explainability, ensures that agent outputs are not opaque, but grounded in reproducible and verifiable actions.

In other words, trust propagates from tools to agents if the agents are able to trace their own outputs to outputs obtained from the components. Actionability is in the components, while conversationality is provided by the agents. This trust propagation approach from tools to agents is sketched in Figure 8.

In the context of an ABPMS, tools provide access to descriptions about the current state of the process, diagnosis for issues, predictions, explanations, and recommendations, among others. Other components may provide actions, e.g. executed by rule-based bots, or by bots driven by automated planners, which perform actions on top of systems that have an impact on the "real world" such as transactional systems (which maintain ground-truth of reality) or communication systems (e.g. sending an email).

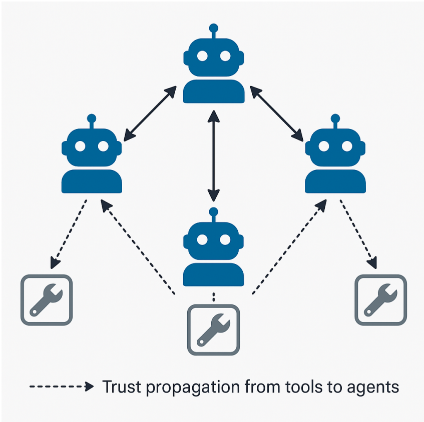


**Figure 8** Trust propagates from tools to agents, which are connected in a network.


**Figure 9** Important conversationally actionable ABPMS concerns.

### 4.3.4 Key Concerns and Challenges

We have identified a non-exhaustive list of key concerns (mostly non-functional in nature) closely relevant to the conversational actionability of ABPMS (Figure 9 and Table 3). These concerns (except Training and probably Strategy Flexibility and Knowledge Representation) are expected to be specified, tracked, and composed using measurable indicators (e.g., in components), as has been done in other areas such as Service-Oriented Architecture [5].

**Table 3** Key concerns for conversationally actionable ABPMS.

| Concern | Description |
|---|---|
| Performance | The extent to which the ABPMS meets time, throughput, and capacity requirements. |
| Cost | The financial impact of invoking external tools and services (including LLMs) on the architecture (e.g., smaller but specialized LLMs) and operation of the ABPMS. |
| Trust | The degree to which users and agents believe that the ABPMS outputs (e.g., models, recommendations, automations) are true, accurate, and dependable. |
| Training | The level of effort and education required for users to effectively learn and utilize the ABPMS and its conversational capabilities. |
| Usability | The degree to which the ABPMS interface and interaction mechanisms enable users to accomplish their goals efficiently and satisfactorily. |
| Helpfulness | The extent to which the outputs of the ABPMS support users in completing or improving their tasks and decision-making. |
| Quality | The accuracy, correctness, and completeness of the results provided by the ABPMS in support of user and agent tasks. |
| Transparency | The extent to which the ABPMS provides understandable and traceable explanations for its outputs, actions, or reasoning. |
| Strategy Flexibility | The degree to which the ABPMS can support the tailoring of human requests to generic or specific strategies (e.g., by specifying which external tools to use and combine). |
| Knowledge Elicitation | The ability of the ABPMS to query and integrate external or internal knowledge sources (e.g., web, databases, ontologies) to enhance its reasoning or task performance. |
| Uncertainty | The extent to which the ABPMS can represent, communicate, and act upon confidence levels of its outputs and outputs. |
| Privacy | The degree to which the ABPMS respects and enforces privacy constraints arising from its frame (laws, policies, or user preferences) during its operations and model fine-tuning. |

Stemming from these concerns and the previous discussion, we identify important challenges for the development of conversationally actionable ABPMS. Most are short-term, but the last two are longer-term.

1. How are actionable conversations defined and evolved?
2. How can conversation needs and mechanisms (e.g., MCP) be elicitated and specified to enable an ABPMS to use discover, understand, and invoke the capabilities of other tools and services (e.g., simulators or AI-based agents)?
3. How can the concerns identified in Table 3, and especially trust, be understood and increased in an ABPMS?

4. How can indicators for these concerns be specified, measured, and aggregated within the ABPMS but also within component contexts (e.g., with MCP)?
5. How can unreliable AI-based components and certified tool components be composed dynamically to contribute towards satisfying the goals of an actionable conversation while addressing the identified concerns?
6. How can an ABPMS itself become a component (MCP-based or other) that can be used by larger ecosystems?
7. (Long-term) How can we support the automated generation of automations (RPA, AI agents, others) from within an APBMS?
8. (Long-term) How can frames related to conversations and to interactions with external services/tools be flexibly managed?

#### References

1 Nataliia Klievtsova, Timotheus Kampik, Juergen Mangler, and Stefanie Rinderle-Ma. Conversationally actionable process model creation. In Marco Comuzzi, Daniela Grigori, Mohamed Sellami, and Zhangbing Zhou, editors, *Cooperative Information Systems - 30th International Conference, CoopIS 2024, Porto, Portugal, November 19-21, 2024, Proceedings*, volume 15506 of *Lecture Notes in Computer Science*, pages 39–55. Springer, 2024.
2 Wil M. P. van der Aalst and Josep Carmona, editors. *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*. Springer, 2022.
3 Xinyi Hou, Yanjie Zhao, Shenao Wang, and Haoyu Wang. Model context protocol (MCP): landscape, security threats, and future research directions. *CoRR*, abs/2503.23278, 2025.
4 Ali Dorri, Salil S. Kanhere, and Raja Jurdak. Multi-agent systems: A survey. *IEEE Access*, 6:28573–28593, 2018.
5 Hanane Becha and Daniel Amyot. Consented consumer-centric non functional property description and composition for SOA-based applications. *Int. J. Web Eng. Technol.*, 10(4):355–392, 2015.

## 4.4 Working Group on Explainability

*Peter Fettke (Saarland University – Saarbrücken, DE & German Research Center for Artificial Intelligence (DFKI) – Saarbrücken, DE, peter.fettke@dfki.de)*
*Fabiana Fournier (IBM Research Israel – Haifa, IL, fabiana@il.ibm.com)*
*Lior Limonad (IBM Research Israel –Haifa, IL, liorli@il.ibm.com)*
*Andreas Metzger (University of Duisburg-Essen, DE, andreas.metzger@paluno.uni-due.de)*
*Stefanie Rinderle-Ma (TU Munich, DE, stefanie.rinderle-ma@tum.de)*
*Barbara Weber (University of St. Gallen, CH, barbara.weber@unisg.ch)*



### 4.4.1 Introduction

An autonomous business process (ABP) is the next generation of AI-Augmented Business Process Management System (ABPMS) [1], which is a self-executing ABPMS that leverages advanced technologies such as Artificial Intelligence (AI) and Machine Learning (ML) to operate with minimal to no human intervention. ABPs can sense and respond to various inputs, reason, make decisions, and adapt to changing circumstances in real time, all without relying on manual triggers or continuous oversight. Think of it like a self-driving car for your business operations. Instead of a human driver controlling all aspects, the system uses sensors, data analysis, and intelligent algorithms to navigate and achieve its objectives.

The notion of ABPs was recently introduced at the AutoBiz Dagstuhl seminar[1], during which the core material for this paper was jointly developed.

One reason for realizing ABPs is to improve operational efficiency, reduce errors, lower costs, improve response times, and free human workers for more strategic and creative work. However, despite the promises that ABPs offer, their characteristics can lead to particular concerns in the context of BPM:

- ABPs may erode *trust* among stakeholders – including process owners, business analysts, end-users, and customers – who may be hesitant to rely on or adopt AI-based process recommendations or automated decisions if they cannot understand the rationale behind them.
- The opacity of ABPs may make it difficult to *debug* process models, as well as identify potential failures, or understand why a process might be under-performing.
- Using ABPs may hinder *accountability*; if an ABP leads to a failure or an unfair outcome, the inability to explain its underlying decisions makes it challenging to assign responsibility or implement corrective actions.
- ABPs may perpetuate hidden *biases* of their underlying AI and ML components. Such biases may lead to discriminatory or unfair process outcomes, which can be difficult to detect and mitigate.
- Demonstrating the *compliance* of ABPs with regulatory frameworks, such as the EU's GDPR and AI Act, requires an increasing level of transparency, particularly in high-risk domains like finance, healthcare, and human resources, which are common areas for BPM applications.

We argue that *explainability* will be a key characteristic of ABP systems to address the aforementioned concerns [1, 2], leading to the notion of **eXplainable ABPs (XABPs)**.

XABPs are particularly relevant when ABPs are realized in the form of *Agentic BPM* systems. An Agentic BPM system is an advanced approach to managing and automating complex business workflows by integrating autonomous AI agents. Unlike traditional BPM or Robotic Process Automation (RPA) systems that follow rigid, predefined rules and workflows, agentic BPM leverages AI to enable systems to make independent decisions, adapt to changing conditions, and learn from experience with minimal human intervention. Here, explainability offers a central mechanism through which agents can articulate the rationale behind their behavior. As such, explainability becomes a first-class citizen in the realization of Agentic BPM systems, supporting agent autonomy from two perspectives:

- Enabling agents to independently resolve misalignments in other agents' behavior.
- Reducing human intervention by making agent behavior understandable and transparent.

Employing state-of-the-art explainable AI (XAI) techniques [3] for XABPs pose several limitations:

1. Inability to express business process model constraints [4].
2. Failure to capture the richness of contextual situations that affect process outcomes [5].

[1] See https://www.dagstuhl.de/25192. We express our gratitude to the Scientific Directorate and staff of Schloss Dagstuhl for their invaluable support. We also thank our fellow participants for their engaging discussions.

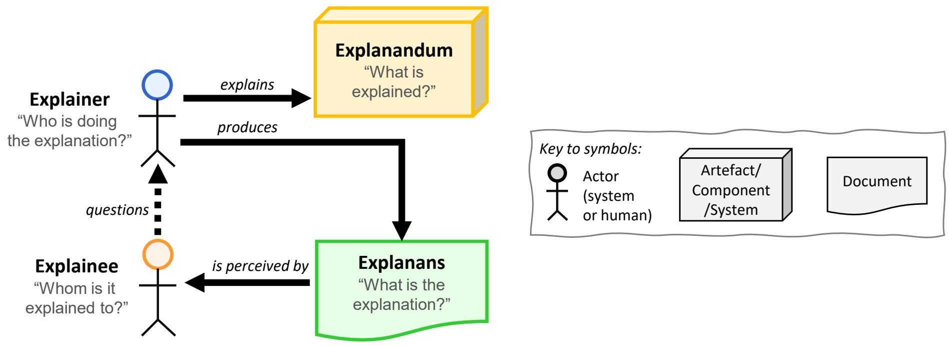


**Figure 10** Explainability Concepts.

3. Inability to reflect causal execution dependencies among activities in the business process [6].
4. Explanations are often nonsensical or not interpretable for human users [7].

### 4.4.2 Characterization and Needs of XABPs

We start with a generic conceptualization of explainability and then refine this to particular concerns in the BPM setting.

#### 4.4.2.1 Fundamental Explainability Concepts

Figure 10 illustrates the key explainability concepts.

The *explainer* provides an explanation of the *explanandum* (explanation subject) by offering one or more *explanans/explanantia* (explanation/explanations). The explanation is generated by the explainer using a specific explanation mechanism at a defined generation time, and is delivered to the *explainee* in a particular presentation format – typically visual or textual. In its simplest form, the explanantia produced by the explainer should provide information about the causes of the explained phenomenon (explanandum) [8]. The content of the explanation must align with both the nature of the explanandum and the needs of the explainee. Furthermore, interaction of the explainee with the explanation can follow different modes, ranging from one-shot explanations to conversational or multi-round interactions.

#### 4.4.2.2 Explanandum: Explanation Subjects *"what is explained?"*

Figure 11 shows the key types of aspects of an explanandum that may be explained, as elaborated below.

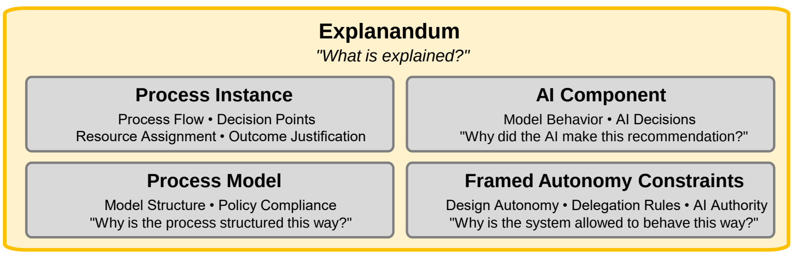


**Figure 11** Explainability Concepts.

**Process Instance Explanation:** ***"Why did this specific process execution take the path and produce the result it did?"***

- *Process Flow* - The sequence of activities, decisions, and events in the business process. Example: "Why did the invoice approval take 5 days?"
- *Decision Points* – Why certain paths or outcomes were chosen during process execution. Example: "Why was a customer's request escalated instead of being resolved at Tier 1?"
- *Resource Assignment* – Why specific tasks were assigned to certain roles or individuals. Example: "Why was this case handled by the senior team?"
- *Outcome Justification* Why a specific result occurred. Example: "Why was this loan application rejected by the process?"

**Process Model Explanation:** ***"Why is the process structured the way it is?"***

- *Model Structure:* Why are certain activities, decisions, or flows included? E.g., "Why do we have a credit history check as a decision point?"
- *Policy Compliance:* Whether and how policies shaped the model or its execution E.g., "Was the data retention policy followed?"

**AI Component Explanation:** ***"Why did an AI component make this recommendation or decision?"***

Note, that this here very much refers to explainable AI (XAI). In more detail:

- *AI Model Outcome:* "Why did the AI component predict deviations or prescribe proactive adaptations?" [9] E.g., "Why was an alarm raised for process event $e_j$?"
- *AI Model Behavior:* "Why does the AI model have certain characteristics or properties?" E.g., "Why does the LSTM prediction model have a Mean Absolute Error (MAE) of only .35 for the given process domain?"

**Framed Autonomy Explanation:** ***"Why is the system or process allowed to behave as it does?"***

- *Design Autonomy:* "Why can the process bypass manual review?"
- *Delegation Rules:* "Why do Tier 1 agents have approval authority?"
- *AI Authority:* "Why can the AI act without a human in the loop?"
- *Escalation Thresholds:* "Why is escalation triggered only after 3 attempts?"
- *Compliance Limits:* "Why is this exception allowed under the GDPR?"

#### 4.4.2.3 Explainer

XABPs involve a range of human and system actors that either generate or consume explanations. Some actors – especially autonomous systems or agents – may fulfill both roles, such as generating explanations for others while also using explanations for self-reflection or system adaptation.

Figure 12a shows the key aspects of the explainer, elaborated below.

**Table 4** Explanation Providers in Business Processes (Explainers).

**(a)** System Explainers: *systems providing or formalizing explanations*

| **Actor Type** | **Role in Ecosystem** | **Explanation About** | **Mechanism or Method** |
|---|---|---|---|
| AI Agents | Intelligent assistants, bots | Predictions, task outcomes, alerts | SHAP, LIME, rule-based reasoning, counterfactuals |
| Monitoring Components | Process mining engines, workflow monitors | Process events, performance, exceptions | Temporal rules, KPI tracking, log analysis |
| Connected Systems | External APIs or services | State changes, synchronization info | Metadata contracts, semantic logging |

**(b)** Human Explainers: *humans providing or formalizing explanations*

| **Actor Type** | **Role in Ecosystem** | **Explanation About** | **Communication Mechanism** |
|---|---|---|---|
| Domain Experts | Analyst specifying business rules | Process logic, decision criteria, exception handling | Process documentation, model annotations, verbal explanation |
| Supervisors / Managers | Operations or compliance manager | Justification for overrides, escalations, or decisions | Reports, notes, emails, verbal feedback |
| Trainers / Annotators | Labelers or human-in-the-loop operators | Ground truth or feedback to train explainable systems | Annotation tools, structured forms, chat interfaces |

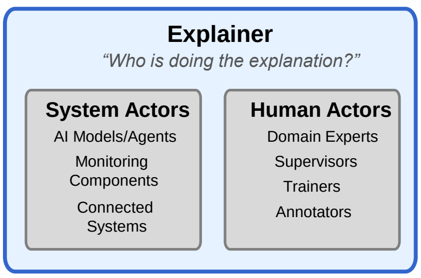


**(a)** Aspects of Explainer

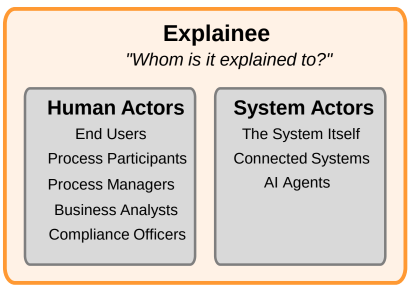


**(b)** Aspects of Explainee

**Figure 12** Overview of explainer and explainee aspects.

#### 4.4.2.4 Explainee

Figure 12b shows the key aspects of the explainee, elaborated below.

**Table 5** Explanation Recipients in Business Processes (Explainees).

**(a)** Human Explainees: *humans consuming explanations.*

| Actor Type | Example Role | Needs from Explanations | Preferred Explanation Style |
|---|---|---|---|
| End Users | Customer applying for a loan | Understand decisions about them (e.g., rejections, delays) | Simple, outcome-focused, natural language |
| Process Participants | Agent handling loan verification | Know what task to do next and why | Step-by-step task rationale, alerts, real-time updates |
| Process Managers | Operations lead, shift manager | Monitor KPIs, react to anomalies, adapt resources | Dashboards, alerts, summaries, what-if analysis |
| Business Analysts / Domain Experts | Person modeling the process | Improve efficiency, detect bottlenecks, validate rule logic | Process mining results, causal analysis, counterfactuals |
| Compliance Officers / Auditors | Internal or external auditors | Ensure traceability, legality, policy adherence | Audit trails, rule execution logs, exception reports |

**(b)** System Explainees: *self-reflective systems consuming explanations.*

| Actor Type | Role in the Ecosystem | Needs from Explanations | Communication Mechanism |
|---|---|---|---|
| The System Itself | Autonomous BPM or AI component | Self-monitoring, internal diagnosis, reconfiguration | Logs, symbolic reasoning, anomaly detection |
| Connected Systems | CRM, ERP, or DMS components | Data or process synchronization with semantic clarity | API contracts, structured events, semantic metadata |

#### 4.4.2.5 Explanans

Figure 13 shows the key aspects of an explanans, elaborated below.

**Explanation Mechanism: *"How is the explanation generated?"***
The explanation mechanism refers to the approach employed by the explainer to generate an explanation – such as attributing feature importance, selecting representative examples, deriving symbolic rules, constructing interpretable approximations, identifying counterfactuals, or visualizing model behavior.

- *Feature Attribution:* Assigns contribution (credit or blame) to input features. *Examples:* SHAP, LIME, Saliency Maps
- *Example-Based:* Uses similar or contrasting examples to justify a decision. *Examples:* k-NN, Prototypes, Counterfactuals

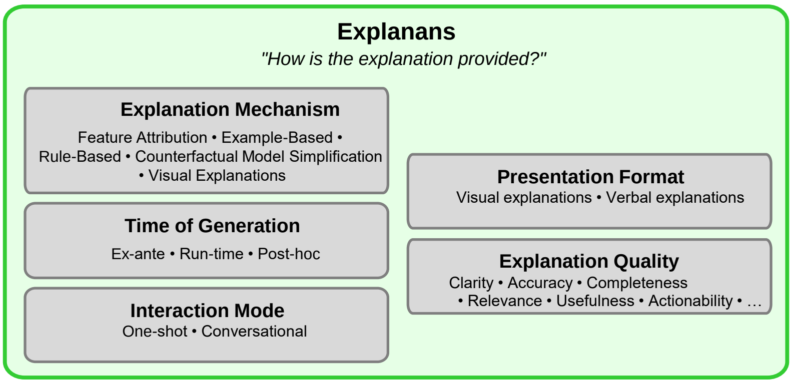


**Figure 13** Aspects of an Explanans.

- *Rule-Based* Derives symbolic or logical rules from data or models. *Examples:* Decision Trees, Rule Lists, Association Rules
- *Model Simplification:* Approximates complex models with interpretable surrogates. *Examples:* Surrogate Decision Trees, Linear Proxies
- *Counterfactual:* Explains what minimal input change would alter the outcome [10]. *Example:* "If income were $5,000 higher, the outcome would have changed."
- *Visual Explanations:* Uses visual indicators to represent decision logic or model behavior. *Examples:* Heatmaps, Partial Dependence Plots

**Time of Explanation Generation: *"When is the explanation produced?"***
The timing of an explanation determines its role in the lifecycle of decision-making systems. Explanations may be generated before, during, or after system execution:

- *Ex-ante Explanations (Before Execution):* Provided before the system executes or makes a decision to validate models or justify decisions before deployment.
- *Run-time Explanations (During Execution):* Delivered while the process is running to support human-in-the-loop oversight or adaptive user feedback.
- *Post-hoc Explanations (After Execution):* Generated after the process completes its actions or decisions in order to audit, debug, or help users understand outcomes.

**Presentation Format of Explanation: *"How is the explanation presented to the user?"***
The chosen presentation method has a direct effect on user comprehension and, therefore, on the success of the explanations [11].
- *Visual explanations*: Heatmaps, charts, dashboards, saliency maps
- *Verbal explanations*: Natural-language output, written rules, factual/counterfactual statements

**Interaction with Explainee: *"How does the user interact with the explanation?"***
Interaction of the explainee with the explanation refers to the mode and extent of user involvement in the explanation process:
- *One-shot explanations:* Explanation provided once, passively
- *Query-based explanations:* Explanation provided on-demand, actively
- *Multi-round / Conversational:* Interactive, iterative, potentially adaptive dialogue [12]

**Explanation Quality:** ***"How to assess the quality of explanations?"***

Explanation quality may be assessed along two complementary dimensions:

- *Technical quality:* This relates to the technical properties of the explanation method itself. Examples include fidelity aka. faithfulness aka. soundness (which measures how accurately the explanation reflects the reasoning or behavior of the explanans), and stability (an explainer should provide similar explanations for similar input or minor perturbations of the input).
- *User-centric quality:* This relates to how the explanation is perceived by humans (in the role of explainees). Examples include usefulness (which quantifies how well it helps the explainee to solve a problem, understand a concept better, or apply the knowledge in a new situation) and meaningfulness (explanation is relevant to the specific explainee and the question or topic at hand and avoids unnecessary tangents or irrelevant information that could confuse the explainee).

### 4.4.3 Challenges for Explainable ABPs

We structure the challenges along the four main explainability concepts as well as along overarching concerns.

#### 4.4.3.1 Challenges Related to Explainee

**Challenge 1: How to specify preferences regarding explanations?** Specifying preferences for explanations presents multifaceted challenges. First, input mechanisms must effectively capture preferences through various channels, whether explicitly declared upfront, interactively elicited through dialogue, or implicitly inferred from user behavior. Systems must accommodate both static preferences that remain consistent and those that dynamically adapt to changing contexts, while supporting the natural evolution of preferences as the explainee's understanding develops. Second, inevitable preference conflicts need to be navigated. This involves carefully balancing competing dimensions, such as detail versus conciseness and speed versus accuracy. This requires finding trade-offs without sacrificing critical explanatory qualities.

#### 4.4.3.2 Challenges Related to Explanandum

**Challenge 2: What explanation subjects are needed for ABPMs?** Explainability related to AI components are rather well understood - not so for ABPM. From our understanding, explanation subjects such as process instance, process models, and framed autonomy constraints are interesting and relevant. However, a more mature taxonomy of explanation types might evolve.

#### 4.4.3.3 Challenges Related to Explainer

**Challenge 3: Which techniques are needed for the explainer to generate explanations?** In the broader field of XAI, several techniques for XAI are known. In addition, in the field of BPM, several techniques are emerging. On this foundation, new techniques should be developed, e.g., what-if analysis, process outcome analysis. All these techniques should take causality into account. In the future, it should be clarified how existing BPM techniques can be integrated, e.g., visualization, and be exploited for creating explanations. Explainability of frames is a nearly unexplored field.

#### 4.4.3.4 Challenges Related to Explanans / Explanantia

**Challenge 4: How may one articulate actionable explanations (e.g., to other agents) to preserve autonomy?** While the explanans is constructed to make an explanation informative about the circumstances that may led to the situation being inquired (i.e., the explanandum), in the context of XABPs, the explanans may also adopt an actionable style – indicating to the explainee which corrective or mitigating actions could be taken to alter the state of the explanandum, particularly without escalating the situation to any external agent. In this way, the explainee may be able to autonomously act upon the condition at hand. However, further work is needed to devise a systematic approach that enables the explainer to determine the most effective content for the explainee, to elicit such corrective action – taking into account both the explanandum and the behavioral intentions of the explainee.

**Challenge 5: When to generate explanations (generation time) and how long to preserve them [13]?** The question here is whether explanations should be generated up-front, whenever possible, or whether we can or should be more conscious about the generation time of the explanation. Another question is when to discard outdated explanations.

**Challenge 6: How can explanations automatically adapt their form to suit the identity of the explainee?** The question is how the explanation can be presented in a way that is easily understandable for the explainee, e.g., leads to low cognitive load for human explainees, and answers the explanation needs of the explainee. This could also be motivated by organizational motivation and goals.

**Challenge 7: How can we accommodate explanations that consider (why) certain behaviors did not occur?** Explaining non-occurring behavior is more challenging than explaining occurring behavior and requires capturing or acquiring knowledge about non-occurring behavior. Causal analysis might be helpful here.

**Challenge 8: How may we synthesize a variety of perspectives (e.g., data, contextual, exogenous) into the explanation?** The first challenge here is to collect and create data sets that cover different perspectives and are of sufficient quality. It is essential to be able to link the synthesized data to process instances. Moreover, providing explanations on synthesized data might also require selecting and filtering the data adequately again to provide adequate explanations.

**Challenge 9: How to identify causal explanations?** Causality vs. correlation: Not every correlation between two variables has a causal explanation. It is therefore important to distinguish between spurious correlations and causality. This classical distinction is well known, but must also be observed in the context of explainable ABPMN. The explainer can provide the explainee with information about the degree of certainty of the explanation offered.

**Challenge 10: How do explanations evolve over time based on feedback or changing context?** Explanations might have to be updated based on changing context and feedback, e.g., if sensor data starts to deviate. The first question is how to detect that an explanation that (partly) takes into account the sensor data has to be updated? Another question is when to present the updated explanation to the user, i.e., directly after a changing context was detected or at another, possibly better fitting moment? This question is related to the question of explanation update frequency. Here, the challenge is to find the sweet spot between keeping explanations up to date and not confusing the explainee. Finally, we have to think about when and how to provide full versus incremental explanation updates.

**Challenge 11: How does the realization of the "frame" in ABPMSs affect the one of explainability?** i.e., with autonomous agents, it may be the means for the agents to share with other agents the rationale for their own behavior.

#### 4.4.3.5 Overarching Challenges

**Challenge 12: How may one assess the quality of the explanations?** Evaluating explanation quality presents a fundamental challenge requiring both empirical and theoretical approaches. From an empirical perspective the question needs to be answered how can we effectively measure explanation quality when objective and subjective dimensions must both be considered? Objective measures include, for example, factual accuracy and completeness, while user-centered aspects cover, for example, comprehensibility, usefulness, effectiveness and efficiency (e.g., see [14]), as well as being actionable.

**Challenge 13: Which kind of datasets are needed to serve as explainability benchmarks?** Benchmarking is a typical approach to evaluating system performance. We expect that such an idea can also promote the development of the field of accountability. However, benchmarking typically relies on adequate benchmark data. In principle, such data can be generated in a laboratory setting. But adequate data are also needed for benchmarking explainability systems in the field.

**Challenge 14: How to ensure that explanations do not reveal information that may be privacy-sensitive, reveal business-critical IPR, or make it easier to undermine the security of the system?** In essence, the challenge lies in providing enough information to satisfy the need for explainability without compromising other crucial aspects of the business, such as data privacy, competitive advantage, and system security.

### 4.4.4 Conclusion

An autonomous business process (ABP) represents a paradigm shift towards self-executing workflows driven by AI and ML Yet ABPs introduce challenges related to trust, transparency, accountability, bias, and regulatory compliance within BPM. To address these issues, this paper introduced the notion of explainable ABPs (XABPs), which can articulate the rationale behind their actions and underlying models. Current explainable AI (XAI) techniques fall short in capturing the complexities of the BPM setting. We therefore introduced a set of challenges to stimulate further research on XABPs.

#### References


1. Marlon Dumas, Fabiana Fournier, Lior Limonad, Andrea Marrella, Marco Montali, Jana-Rebecca Rehse, Rafael Accorsi, Diego Calvanese, Giuseppe De Giacomo, Dirk Fahland, Avigdor Gal, Marcello La Rosa, Hagen Völzer, and Ingo Weber. AI-augmented business process management systems: A research manifesto. *ACM Trans. Manag. Inf. Syst.*, 14(1):11:1–11:19, 2023.
2. Nijat Mehdiyev, Maxim Majlatow, and Peter Fettke. Interpretable and explainable machine learning methods for predictive process monitoring: A systematic literature review, 2023.
3. Amina Adadi and Mohammed Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access*, 6, 2018.
4. Guy Amit, Fabiana Fournier, Shlomit Gur, and Lior Limonad. Model-informed lime extension for business process explainability. In *PMAI@IJCAI'22*. CEUR, 2022.

5 Guy Amit, Fabiana Fournier, Lior Limonad, and Inna Skarbovsky. Situation-aware explainability for business processes enabled by complex events. In *BPM-W 2022*, volume 460 of *LNBIP*, pages 45–57. Springer, 2022.
6 F. Fournier, L. Limonad, I. Skarbovsky, and Y. David. The why in business processes: Discovery of causal execution dependencies. *Künstliche Intelligenz*, 1 2025.
7 Dirk Fahland, Fabiana Fournier, Lior Limonad, Inna Skarbovsky, and Ava J. E. Swevels. How well can large language models explain business processes as perceived by users? *Data & Knowledge Engineering*, 157:102416, 2 2025.
8 Peter Lipton. *Causation and Explanation*, pages 619–631. Oxford University Press, 1 2010.
9 Andreas Metzger, Tristan Kley, Aristide Rothweiler, and Klaus Pohl. Automatically reconciling the trade-off between prediction accuracy and earliness in prescriptive business process monitoring. *Inf. Syst.*, 118:102254, 2023.
10 Tsung-Hao Huang, Andreas Metzger, and Klaus Pohl. Counterfactual explanations for predictive business process monitoring. In Marinos Themistocleous and Maria Papadaki, editors, *EMCIS 2021*, volume 437 of *LNBIP*, pages 399–413. Springer, 2021.
11 Lorenzo Malandri, Fabio Mercorio, Mario Mezzanzanica, and Navid Nobani. Convxai: a system for multimodal interaction with any black-box explainer. *Cogn. Comput.*, 15(2):613–644, 2023.
12 Andreas Metzger, Jone Bartel, and Jan Laufer. An AI chatbot for explaining deep reinforcement learning decisions of service-oriented systems. In *ICSOC 2023*, volume 14419 of *LNCS*, pages 323–338. Springer, 2023.
13 Nijat Mehdiyev and Peter Fettke. Explainable artificial intelligence for process mining: A general overview and application of a novel local explanation approach for predictive process monitoring. In *Interpretable Artificial Intelligence: A Perspective of Granular Computing*, pages 1–18. Springer, 2021.
14 Andreas Metzger, Jan Laufer, Felix Feit, and Klaus Pohl. A user study on explainable online reinforcement learning for adaptive systems. *ACM Trans. Auton. Adapt. Syst.*, 19(3):15:1–15:44, 2024.


## 5 AUTOBIZ Dissemination

This section gives a brief overview of how existing and new knowledge of AUTOBIZ is and will be disseminated. In the former case, the focus is on education (covered by Subsection 5.1, whereas in the latter case, the focus is on venues for the dissemination of novel results, i.e., workshops, journals, and conferences (Subsection 5.2).

### 5.1 Embedding of ABPMS-related Training into University Curricula

With the rise of "agentic AI" in enterprise software, autonomous business process execution has reached the mainstream. However, there is substantial uncertainty on how to achieve autonomy in a reliable and sustainable manner, with analysts such as Gartner using the term "agent washing" to describe pseudo-autonomous enterprise subsystems and predicting that 40% of agent deployment projects will fail by 2027[2]. In order to provide an education foundation for effective, efficient, and sustainable AUTOBIZ, it is crucial to adjust both business process

[2] Cf. `http://s.cs.umu.se/ni2riq` (Gartner press release, URL shortened), *accessed at 2025-07-18.*

management and applied AI curricula as soon as possible. This should happen on both undergraduate and graduate levels; just as important are continuous education offerings for the professionals who are exposed to autonomous process execution systems already now. Many of the seminar participants are prolific educators, having, for example, co-authored widely used textbooks on business process management and related topics. Accordingly, the expectation is that the seminar content will be integrated not only into local education offerings at participants' universities but also into materials that are more widely distributed. Materials are expected to be tailored to different educational needs, based on the educational level as well as technical and practical skills and experience of several target audiences (e.g., business students vs. engineering students vs. seasoned practitioners). Specific first steps for developing course materials have been taken. For example, at Umeå University (Sweden), funding has been secured to develop a course on "Sustainable AI Transformation".

## 5.2 Related Publications and Future Venues

The discussions of the working groups, as presented in this report, have matured into four accepted papers to appear in the CEUR proceedings of the PMAI workshop, held at the ECAI 2025 conference. In addition, a tutorial at BPM 2025 is planned for presentation by Prof. Andreas Metzger on the topic of "AI-Assisted Business Process Monitoring", including insights drawn directly from the discussions in this seminar. For novel and strong technical results that participants and other researchers in the BPM and AI communities produce on the topic, a special issue in the *Information Systems* journal – the prime journal venue for technical BPM research results – has been announced, guest-edited by the seminar organizers.

Given the prevalence of AI in BPM research, further steps may be taken to facilitate the continued dissemination of AUTOBIZ-related research results. For example, discussions have started regarding the establishment of a new *AIXBPM forum* at the Business Process Management conference, the prime conference venue of BPM research.

Future research on AUTOBIZ is expected to enjoy substantial funding success. Notably, seminar co-organizer Marlon Dumas has secured an ERC grant for the *AI-Assisted Optimization of Business Processes*[3]. Also, the organizations of many of the industry participants are expected to continue to fund and drive AUTOBIZ research and its transfer into practice at scale.

[3] `https://cs.ut.ee/en/news/university-tartu-researcher-was-awarded-erc-grant-turn-frontier-science-practical-solutions`, *accessed at 2025-07-18.*

## Participants

- Daniel Amyot
University of Ottawa, CA
- Diego Calvanese
Free University of Bozen-Bolzano, IT
- Marco Comuzzi
Ulsan National Institute of Science and Technology, KR
- Giuseppe De Giacomo
University of Oxford, GB
- Marlon Dumas
University of Tartu, EE
- Achiya Elyasaf
Ben Gurion University – Beer Sheva, IL
- Peter Fettke
DFKI – Saarbrücken, DE
- Fabiana Fournier
IBM Israel – Haifa, IL
- Timotheus Kampik
SAP Berlin, DE & Umeå University, SE
- Yves Lesperance
York University – Toronto, CA
- Lior Limonad
IBM Israel – Haifa, IL
- Andrea Marrella
Sapienza University of Rome, IT
- Andrea Matta
Polytechnic University of Milan, IT
- Andreas Metzger
Universität Duisburg – Essen, DE
- Marco Montali
Free University of Bozen-Bolzano, IT
- Stefanie Rinderle-Ma
TU München, DE
- Sebastian Sardiña
RMIT University – Melbourne, AU
- Arik Senderovich
York University – Toronto, CA
- Estefania Serral Asensio
KU Leuven, BE
- Niek Tax
Meta – London, GB
- Irene Teinemaa
Google DeepMind – London, GB
- Barbara Weber
Universität St. Gallen, CH

Report from Dagstuhl Seminar 25201

# Computational Geometry

**Maarten Löffler*[1], Eunjin Oh*[2], Jeff M. Phillips*[3], and Alexandra Weinberger†[4]**

1 **Utrecht University, NL.** m.loffler@uu.nl
2 **POSTECH – Pohang, KR.** eunjin.oh@postech.ac.kr
3 **University of Utah – Salt Lake City, US.** jeffp@cs.utah.edu
4 **FernUniversität in Hagen, DE.** alexandra.weinberger@fernuni-hagen.de

**Abstract**

This report documents the program and the outcomes of Dagstuhl Seminar 25201 "Computational Geometry". The seminar program spanned the days from 11th May to 16th May 2025, and 39 participants from various countries were on site. Recent advances in computational geometry were presented and discussed, and new challenges were identified, in particular in relation to the two themes "parameterized complexity" and "the interplay between theory and implementation". This report collects the abstracts of the talks and the open problems presented at the seminar, an excerpt from the panel discussion, and partial progress from the active working groups.

**Seminar** May 11–16, 2025 – https://www.dagstuhl.de/25201
**2012 ACM Subject Classification** Theory of computation → Computational geometry; Theory of computation → Design and analysis of algorithms; Mathematics of computing → Discrete mathematics
**Keywords and phrases** algorithms, combinatorics, complexity, geometric computing, implementation
**Digital Object Identifier** 10.4230/DagRep.15.5.64

## 1 Executive Summary

*Maarten Löffler (Utrecht University, NL)*
*Eunjin Oh (POSTECH – Pohang, KR*
*Jeff M. Phillips (University of Utah – Salt Lake City, US)*



This Dagstuhl Seminar constituted a biennial gathering of computational geometers and topologists at the Dagstuhl venue to share recent results, and further research on some of the most important problems of the time in this field. This year the seminar focused on two central challenges within computational geometry and topology: (1) the emergence of parameterized algorithms, and (2) the interaction between implementation and theory in geometry. Within the parameterized algorithms theme, two enlightening overview talks were given. The first was on parameterized techniques in computational geometry, focusing on clustering, covering, and geometric intersection graphs. The second was on parameterized complexity and its role in geometric topology, focusing on triangulated topological manifolds. The second theme on the interplay between implementation and theory covered a broader set of topics from motion planning to persistence homology to geometric software engineering to geometric topology software to non-metric high-dimensional geometries to using Haskell for

* Editor / Organizer
† Editorial Assistant / Collector

low-dimensional geometry. A highlight was a panel discussion on the state of the challenges of promoting implementation-driven research in a field dominated by primarily theoretical results. Finally, an open problem session promoted a variety of intriguing unsolved issues in the field; many are stated within this report. These problems occupied many of the participants throughout the week as they were interrogated in small working groups. Some research resolution during the week, and we expect others to lead to publications soon after.

## Changes to the Lottery

The Computational Geometry Dagstuhl Seminar series is the longest-running series of Dagstuhl seminars, starting in 1990 with Dagstuhl Seminar 9041 (see `https://www.dagstuhl.de/9041`) and the current seminar 25201 being its 19th iteration.

In 2013, the organizers of Dagstuhl Seminar 13101 introduced a *lottery* to select participants. The reason for this was that a steadily growing number of established names in the community being invited each time limited access for more junior researchers to this series. While this is not necessarily a bad thing and is indeed a sign of a healthy and expanding community, the organizers felt this conflicted with Dagstuhl's unique opportunity for junior and senior researchers to interact. They reported in 2013:

> We believe that the lottery created space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar was much "younger" than in the past, and certainly more "family-friendly." Five young children roaming the premises created an even cosier atmosphere than we are used in Dagstuhl. Without decreasing the quality of the seminar, we had a more balanced attendance than in the past. Feedback from both seminar participants and from researchers who were not selected was uniformly positive.

Since 2013, the lottery has been used each time to select most of the participants. Over time, though, the task of maintaining an accurate list of "the community" has proven challenging, and through its lack of transparency the series also acquired some mystery, specifically surrounding the questions how one enters the lottery.

This year, we have for the first time not used the hand-curated list, and instead automatically generated a list of names as input to the lottery. Here we report on how the lottery was run this year.

A large Dagstuhl Seminar has space for about about 45 participants. The organizers reserved 10 seats for hand-selected invitees based on the chosen themes for the seminar. The remaining 35 seats were selected through the lottery. The organizers have used DBLP's SparQL interface `https://sparql.dblp.org/` to generate a list of names to enter the lottery. The query used this year was:

> "all people who have at least 3 publications in SoCG, of which at least 1 appeared in the last 5 years".

From this list, a weighted random selection was made, taking into account Dagstuhl's requirements on geographical and gender diversity, the themes of the seminar, and invitations to previous Dagstuhl Seminars.

## 2 Table of Contents

**Executive Summary**
*Maarten Löffler, Eunjin Oh, and Jeff M. Phillips* . . . 64

**Overview of Talks**

Parametrized Algorithms in Geometry
*Sayan Bandyapadhyay* . . . 68

Lessons Learned in Implementation
*Benjamin Burton* . . . 68

Better Late than Never: The Complexity of Arrangements of Polyhedra
*Otfried Cheong* . . . 69

Fast Persistence for 1D Functions (In Gudhi)
*Marc Glisse* . . . 69

The Practice and Theory of Sampling-Based Motion Planning
*Dan Halperin* . . . 70

Maximal 2-Planar Graphs: A Computational Proof for a Graph Drawing Problem
*Michael Hoffmann* . . . 70

Implementing (Geometric) Algorithms as a Less-Initiated Software Engineer
*Gert Meijer* . . . 71

Motion Planning for Introverts: Social Distancing in Various [F/N]orms
*Valentin Polishchuk* . . . 71

The Geometry of the Set of Equivalent Linear Neural Networks
*Jonathan Shewchuk* . . . 72

Parameterised Complexity in Geometric Topology
*Jonathan Spreer* . . . 72

hgeometry: Experience Report Implementing CG Algorithms
*Frank Staals* . . . 73

Google, Balls, and Technology
*Hubert Wagner* . . . 73

A Free Lunch: Manifolds of Positive Reach Can Be Smoothed Without Decreasing the Reach
*Mathijs Wintraecken* . . . 73

**Panel Discussion: On the Interplay Between Implementing Geometric Algorithms and Theory** . . . 74

**Open Problems and Working Groups** . . . 75

Covering by Double Disks
*Sayan Bandyapadhyay* . . . 75

Is Finding Positive Euler Characteristic Normal Surfaces FPT?
*Benjamin Burton* . . . 76

Turning Planar Triangulations Inside Out
*Jeff Erickson and Birgit Vogtenhuber* . . . 77

Morphing Long-Geodesic Embeddings on the Sphere
*Jeff Erickson* . . . . . 78

Approximate Simplification of Flag Filtrations
*Marc Glisse, Jeff M. Phillips* . . . . . 79

Sub-Alpha-Complex in High Dimension
*Marc Glisse, Jeff M. Phillips* . . . . . 80

Treewidth of 3-Manifolds
*Benjamin Burton, Kristóf Huszár, Jonathan Spreer, Yelena Yuditsky* . . . . . 81

WSPD for Low Density Graphs
*Joachim Gudmundsson* . . . . . 84

Consistent System of Shortest Paths
*Francis Lazarus, Jeff Erickson, Raimund Seidel, Ling Zhou* . . . . . 85

Graph Tiles
*Oswin Aichholzer, Robert Ganian, Phillip Keldenich, Maarten Löffler, Gert Meijer, Alexandra Weinberger, Carola Wenk* . . . . . 87

Approximate Minlink Paths in Subquadratic Time
*Valentin Polishchuk, Sayan Bandyapadhyay, Mark de Berg, Joachim Gudmundsson, André Nusser, Frank Staals* . . . . . 89

Alexander Theorem for Knotted Spheres
*Mathijs Wintraecken and Kristóf Huszár* . . . . . 91

Universal Point Sets
*Alexander Wolff, Stefan Felsner, Michael Hoffmann, Maarten Löffler, Fabrizio Montecchiani, Yoshio Okamoto, Jonathan Spreer* . . . . . 92

**Participants** . . . . . 95

# 3 Overview of Talks

This section contains abstracts of talks that were given during the seminar. Sayan Bandyapadhyay and Jonathan Spreer were invited to speak on the topic of parameterized computational geometry and topology, respectively, one of the two themes of the seminar. For the theme of theoretical questions arising in implementing geometric Benjamin Burton, Marc Glisse, Danny Halperin, Gert Meijer, Frank Staals, and Hubert Wagner were invited to give a talk. Additionally, Ottfried Cheong, Michael Hoffmann, Valentin Polishchuk, Jonathan Shewchuk, Mathijs Wintraecken, gave talks in the area of computational geometry.

## 3.1 Parametrized Algorithms in Geometry

*Sayan Bandyapadhyay (Portland State University, US)*



The design of parameterized algorithms for geometric problems has become a well-established area of research, driven in part by the many natural parameters that arise in these settings. While the field has a rich history, earlier efforts were often fragmented, with only a few problems receiving systematic attention. In contrast, recent work has adopted a more unified approach, focusing on the development of general frameworks that can address a broader class of problems. In this talk, we explore three central topics – clustering, covering, and geometric intersection graphs – where parameterized complexity has played a significant role. We review the literature in these areas, highlight key techniques, and pose several open questions for future investigation.

## 3.2 Lessons Learned in Implementation

*Benjamin Burton (University of Queensland – Brisbane, AU)*



We talk through some of the lessons learned through 26 years of working on the topological software package Regina, and from major computational projects including the tabulation of prime knots up to 20 crossings. Regina implements exact algorithms for 3-manifolds, 4-manifolds, knots and links, and is freely available from regina-normal.github.io.

## 3.3 Better Late than Never: The Complexity of Arrangements of Polyhedra

*Otfried Cheong (Scalgo, DK)*


**Joint work of** Boris Aronov, Sang Won Bae, Sergio Cabello, Otfried Cheong, David Eppstein, Christian Knauer, Raimund Seidel
**Main reference** Boris Aronov, Sang Won Bae, Sergio Cabello, Otfried Cheong, David Eppstein, Christian Knauer, Raimund Seidel: "Better Late than Never: the Complexity of Arrangements of Polyhedra", CoRR, Vol. abs/2506.03960, 2025.
**URL** https://doi.org//10.48550/ARXIV.2506.03960

Let $\mathcal{A}$ be the subdivision of $\mathbb{R}^d$ induced by $m$ convex polyhedra having $n$ facets in total. We prove that $\mathcal{A}$ has combinatorial complexity $O(m^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor})$ and that this bound is tight. The bound is mentioned several times in the literature, but no proof for arbitrary dimension has been published before.

An earlier version of the paper appeared in EuroCG 2025, with a proof for general position only.

## 3.4 Fast Persistence for 1D Functions (In Gudhi)

*Marc Glisse (INRIA – Orsay, FR)*


**Main reference** Marc Glisse: "Fast persistent homology computation for functions on R", CoRR, Vol. abs/2301.04745, 2023.
**URL** https://doi.org//10.48550/ARXIV.2301.04745

0-dimensional persistent homology is known, from a computational point of view, as the easy case. Indeed, given a list of $n$ edges in non-decreasing order of filtration value, one only needs a union-find data structure to keep track of the connected components and we get the persistence diagram in time $O(n\alpha(n))$. The running time is thus usually dominated by sorting the edges in $\Theta(n \log(n))$. A little-known fact is that, in the particularly simple case of studying the sublevel sets of a piecewise-linear function on $\mathbb{R}$ or $\mathbb{S}^1$, persistence can actually be computed in linear time. This talk presents a simple algorithm that achieves this complexity and an extension to image persistence. An implementation is available in Gudhi.

## 3.5 The Practice and Theory of Sampling-Based Motion Planning

*Dan Halperin (Tel Aviv University, IL)*


**Main reference** Oren Salzman: "Sampling-based robot motion planning", Commun. ACM, Vol. 62(10), pp. 54–63, 2019.
**URL** https://doi.org//10.1145/3318164
**Main reference** Kiril Solovey, Dan Halperin: "Sampling-Based Bottleneck Pathfinding with Applications to Fréchet Matching", in Proc. of the 24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark, LIPIcs, Vol. 57, pp. 76:1–76:16, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.
**URL** https://doi.org//10.4230/LIPICS.ESA.2016.76
**Main reference** Kiril Solovey, Dan Halperin: "Efficient sampling-based bottleneck pathfinding over cost maps", in Proc. of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017, pp. 2003–2009, IEEE, 2017.
**URL** https://doi.org//10.1109/IROS.2017.8206019

Since their introduction in the 1990s, sampling-based (SB) planners have become a central tool for solving motion planning problems, both in robotics and beyond. In this talk, I will give a brief introduction to SB planning, highlighting major milestones in its development and the types of theoretical guarantees these methods often provide. I will then discuss the application of SB planning to Fréchet-type problems, as well as its implementation in our DiscoPygal platform – a suite of tools designed for rapid development and experimentation with motion planning in 2D environments. DiscoPygal is built using Python bindings over CGAL. I will conclude with several open problems in SB planning.

## 3.6 Maximal 2-Planar Graphs: A Computational Proof for a Graph Drawing Problem

*Michael Hoffmann (ETH Zürich, CH)*


**Joint work of** Michael Hoffmann, Meghana Mallik Reddy, Shengzhe Wang
**Main reference** Michael Hoffmann, Meghana M. Reddy: "The Number of Edges in Maximal 2-Planar Graphs", in Proc. of the 39th International Symposium on Computational Geometry, SoCG 2023, June 12-15, 2023, Dallas, Texas, USA, LIPIcs, Vol. 258, pp. 39:1–39:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
**URL** https://doi.org//10.4230/LIPICS.SOCG.2023.39

One of the main meta questions in the study of beyond-planar graphs is: Which of the many nice properties of planar graphs carry over to a more general setting? One of these properties is sparsity and one of the most natural and well-studied classes of beyond-planar graphs is the family of $k$-planar graphs. A graph is *$k$-planar* if it can be drawn in the plane such that every edge has at most $k$ crossings. It is known that every $k$-planar graph on $n$ vertices has at most $c_k n$ edges, for some constant $c_k \leq 3.81\sqrt{k}$. However, several other nice properties of planar graphs do not carry over to $k$-planar graphs. For instance, every *maximal* planar graph (adding any edge yields a nonplanar graph) on $n$ vertices has exactly $3n-6$ edges. Such a correspondence does not hold for maximal $k$-planar graphs, with $k \geq 1$.

We studied the case $k = 2$ specifically, and could show that every maximal 2-planar graph on $n \geq 5$ vertices has at least $2n$ edges. Compared to the upper bound of $5n - 10$ from the literature, and also compared to the $3n - 6$ bound for planar graphs, this may not sound too impressive. Thus, we also tried to find examples of maximal 2-planar graphs with few edges. Indeed, it is not too difficult to come up with some plausible candidates. Certifying 2-planarity is easy: Just describe a corresponding drawing. But proving maximality turns out to be a challenge. The reason is that – in contrast to planarity – testing if a given graph is $k$-planar, for $k \geq 1$, is known to be NP-complete.

Our attempts to prove maximality using pencil and paper quickly turned into a mess of case distinctions. Thus – in the spirit of the focus theme about *the interplay between theory and implementation* – it seemed much more reasonable to let a computer check all these cases. But also for a computer-based proof it is not really obvious how to proceed: Our graphs form a conceptually infinite family, where already initial slices are quite large for an exhaustive computation. In this talk I briefly discuss how we overcame these challenges so as to obtain a computational proof of the following statement: For every $n$ there exists a maximal 2-planar graph on $n$ vertices with $2n + c$ edges, where $c \leq 350$ is an absolute constant, independent of $n$.

## 3.7 Implementing (Geometric) Algorithms as a Less-Initiated Software Engineer

*Gert Meijer (NHL Stenden Hogeschool – Leeuwarden, NL)*



This talk gives an overview of my experience implementing theoretical algorithms in an applied software engineering context. This potential audience has a lower mathematical background. I present the challenges we face when trying to translate papers into 'mortal' language, as well as the problems that arise during the implementation phase. I advocate for more accessible papers and end with a comparison of software engineering trends. I hope these insights contribute to making theoretical work more accessible and impactful.

## 3.8 Motion Planning for Introverts: Social Distancing in Various [F/N]orms

*Valentin Polishchuk (Linköping University, SE)*


**Joint work of** Reilly Browne, Kevin A. Buchin, Omrit Filtser, Mayank Goswami, Mart Hagedoorn, Prahlad Narasimham Kasthurirangan, Joseph S. B. Mitchell, Valentin Polishchuk, Roman Voronov

Maintaining separation/invisibility for static/moving agents is the goal in various applications, motivated by security/privacy and similar considerations. I will present several recent results on the topic, ranging from introduction of a new path (dis)similarity measure to gender-aware science to resolving long-standing open problems.

Based on FoCS23, WADS25, ITCS23, SoCG20, FUN18 work with R. Browne, P. Kasthurirangan, J. Mitchell, M. Hagedoorn, O. Filtser, M. Goswami, K. Buchin, L. Sedov, R. Voronov.

## 3.9 The Geometry of the Set of Equivalent Linear Neural Networks

*Jonathan Shewchuk (University of California – Berkeley, US)*


**Joint work of** Sagnik Bhattacharya, Jonathan Shewchuk
**Main reference** Jonathan Richard Shewchuk, Sagnik Bhattacharya: "The Geometry of the Set of Equivalent Linear Neural Networks", CoRR, Vol. abs/2404.14855, 2024.
**URL** https://doi.org//10.48550/ARXIV.2404.14855

A linear neural network computes a linear transformation of its input vector. Given a fully-connected linear network, the set of all weight vectors for which the network computes the same linear transformation is an algebraic variety in weight space, called a fiber under the matrix multiplication map. Sometimes this variety a manifold, but usually not. The rank stratification of a fiber is a natural partition of the fiber into manifolds of various dimensions called strata. We characterize how these strata are connected to each other. They satisfy the frontier condition: if a stratum intersects the closure of another stratum, then the former stratum is a subset of the closure of the latter stratum. This subset relationship can be expressed as a partial order with a single minimal element, a stratum that is included in the closure of every other stratum. Our main result describes the relationship between this partial order and the ranks of certain matrices in the network. Each stratum represents a different pattern of information flow through the network, and moves up the hierarchy may offer neural network training algorithms more opportunities to succeed. To support this construction, we propose a candidate for the "fundamental subspaces of a linear neural network", analogous to the four fundamental subspaces of a matrix.

## 3.10 Parameterised Complexity in Geometric Topology

*Jonathan Spreer (University of Sydney, AU)*



A major goal in the field of geometric topology is the algorithmic study of topological manifolds, typically represented by triangulations. Specifically, we want to answer topological questions about manifolds by running discrete algorithms on their triangulations.

Methods from parameterised complexity play an important role in the development of these algorithms: On the one hand, parameterised hardness results precisely state how complicated triangulations can obstruct efficient solutions to algorithmic problems; on the other hand, a host of fixed-parameter tractable algorithms exist to solve hard topological problems efficiently for well-behaved manifolds and/or triangulations.

In my talk I will survey hardness results as well as practical algorithms. Some of these are implemented in the low-dimensional topology software Regina. The algorithms are fixed-parameter tractable in the first Betti number of the underlying manifold, or the treewidth of the dual graph of the input triangulation. I will then go on to survey theoretical results about how topological properties of a manifold influence the treewidth of the dual graph of its triangulations.

## 3.11 hgeometry: Experience Report Implementing CG Algorithms

*Frank Staals (Utrecht University, NL)*



hgeometry is a pure Haskell library implementing some textbook computational geometry algorithms "from scratch". In the talk I will discuss some of the lessons learned while implementing these algorithms. The main points are:

- the easy things are hard (or at least annoying),
- degeneracies are a real problem,
- implementing the algorithm itself was only 1/3rd of the work,
- property testing was very useful.

Two challenges in property testing with which the theory community could help (i) are generating a good distribution of test data (that includes corner cases and degeneracies) and (ii) developing approaches for efficiently shrinking (simplifying) a data set containing a bug.

## 3.12 Google, Balls, and Technology

*Hubert Wagner (University of Florida – Gainesville, US)*


**Joint work of** Herbert Edelsbrunner, Hana Dal Poz Kouřimská, Tuyen Pham, Žiga Virk, Hubert Wagner

I start from a computational geometry problem arising from my work on the google alerts service (`https://www.google.com/alerts`). Specifically, nearest neighbour and ball search in spaces that are *not* metric spaces. One example is the Kullback–Leibler (KL) divergence which is used to compare pairs of discrete probability distributions (in applications, e.g. a text document is often represented as a discrete probability distribution.) This is mostly a pretext to describe certain nice properties of Bregman divergences, whose one member is the KL divergence. I show several standard geometric constructions (Delaunay triangulation, Čech filtration, k-means clustering, kd-trees) that naturally extend to the Bregman setting. Additionally, this talk serves as a starting point to a discussion about the interplay between practical applications and theory of computational geometry – which is relevant to the ongoing addition of a second, practical track to the SoCG conference.

## 3.13 A Free Lunch: Manifolds of Positive Reach Can Be Smoothed Without Decreasing the Reach

*Mathijs Wintraecken (Centre Inria d'Université Côte d'Azur, FR)*


**Joint work of** André Lieutier, Hana Dal Poz Kouřimská, Mathijs Wintraecken

Assumptions on the reach are crucial for ensuring the correctness of many geometric and topological algorithms, including triangulation, manifold reconstruction and learning, homotopy reconstruction, and methods for estimating curvature or reach. However, these assumptions are often coupled with the requirement that the manifold be smooth, typically at least $C^2$. In this paper, we prove that any manifold with positive reach can be approximated arbitrarily well by a $C^\infty$ manifold without significantly reducing the reach, by employing techniques from differential topology – partitions of unity and smoothing using convolution kernels. This result implies that nearly all theorems established for $C^2$ or manifolds with a certain reach naturally extend to manifolds with the same reach, even if they are not $C^2$, for free!

## 4 Panel Discussion: On the Interplay Between Implementing Geometric Algorithms and Theory

Moderator: Jeff M. Phillips
Panelists: Dan Halperin, Jonathan Shewchuk, Carola Wenk, Benjamin Burton

The panel was convened to discuss the challenges in implementing geometric algorithms, and how that led to developments in the theoretical advancements and understanding. While this was in planning for over a year before the Dagstuhl Seminar, it became very timely after a large discussion on Implementation and Application focused papers at the 2024 International Symposium on Computational Geometry (SoCG), which led to a Task Force on the topic, which released a report a month before seminar.

The panel started with discussion on successes in and surprising outcomes from implementing geometric algorithms. Where did these outcomes differ from what was expected. Halperin and Burton referred to their earlier talks about how methods in robotics or enumerative topology changed the course of the field. In robotics, it was how simple sample-based motion planning made complicated geometry analysis overly detailed, and in the topology space, it was how by reimplementing (and optimizing!) worst-case expensive methods (eventually) led to really fast approaches in practice. Wenk talked about the surprise in the complexity in working with real-world GPS data, which led to eventually enough material to write a book. And Shewchuk talked about how certain lingering challenges in constrained triangulations took over a decade until only recently what seems like the “right” solution was published.

The panel was also asked about the interplay between experimental and theoretical explorations of the same topic led to advancements in both. Halperin described how the theory communities formalization of the motion planning problem paved the way of the simpler sampling-based methods. Burton mentioned that parameterized algorithms were linked to the topology problems to explain how they could run so fast, and formalize the sort of optimizations that would be possible in practice.

The discussion turned to the challenges implementation-driven work faces, how it is valued, and where it should be published. The challenges of publishing papers on the influential CGAL library were brought up. There were no definitive answers, but the proposed second track at SoCG was discussed, and those working on implementation complements to geometric work were encouraged to continue, and to submit their work for publication.

Finally, the panel was asked what lessons could be learned from their experience on the interaction between implementation and theoretical approaches to geometric problems. While Shewchuk argued that more should be taken from these interactions, Halperin ended on the note that despite it being a struggle, these different approaches had done a lot to move the field forward.

# 5 Open Problems and Working Groups

This section contains open problems that have been proposed on the start of the seminar. Many of the open problems have then be discussed in more depth during the week, two even resolved in full. We expect the working groups that have formed to continue their joined research on the problems beyond the seminar and potentially further collaboration will form still afterwards. A short overview of some of the ideas and discussions happening during the week has also been added to this section (to the respective open problem).

## 5.1 Covering by Double Disks

*Sayan Bandyapadhyay (Portland State University, US)*



We are given $m$ access points and $n$ client points. At each access point, one can place a disk of radius 1 or 2. The goal is to pick a radius (i.e., a disk) for each access point so that (1) the chosen disks cover each client point, and (2) the number of client points covered by the radius 1 disks is maximized.

This problem is known to be NP-hard, and there are two polynomial-time approximation algorithms:

- 4-approximations [1]
- 2.5-approximations [2], FPT in sparsity and the number of radius 2 disks in an optimal solution, also W[1]-hard in the number of radius 2 disks.

**Questions.**

1. What is the best approximation possible in polynomial time?
2. Is the problem FPT in the number of radius 1 disks in an optimal solution?

### 5.1.1 Progress

This problem appears to be equivalent to discrete unit disk cover (given a set of points and a set of equal-radius disks, choose the smallest set of disks that covers all the points) by the following observations:

- All clients that are within distance at most 1 from at least 1 access point will be covered anyway, so we might as well throw them away.
- All access points that are at distance larger than 2 from any remaining client point will never be a radius 2 disk, so we might as well throw them away.
- For the remaining access points, placing radius 1 disks will not cover any more clients, so all remaining clients must be covered by radius 2 disks.
- Minimizing the number of radius 2 disks now is the same as maximizing the number of radius 1 disks in the original question.

However, we are now approximating or parameterizing w.r.t. the number of non-chosen disks rather than the number of chosen disks, so known results for DUDC do not necessarily carry over.

**References**

1. Sayan Bandyapadhyay, Anil Maheshwari, Sasanka Roy, Michiel Smid, and Kasturi R. Varadarajan. Geometric covering via extraction theorem. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287 of *LIPIcs*, pages 7:1–7:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ITCS.2024.7.
2. Sayan Bandyapadhyay and Eli Mitchell. Approximation and parameterized algorithms for covering with disks of two types of radii. In Pat Morin and Eunjin Oh, editors, *19th Algorithms and Data Structures Symposium (WADS)*, 2025. To appear.


## 5.2 Is Finding Positive Euler Characteristic Normal Surfaces FPT?

*Benjamin Burton (University of Queensland – Brisbane, AU)*



### 5.2.1 Problem Statement

The setting is a 3-manifold triangulation. A *normal surface* is a properly embedded surface that meets each tetrahedron in some collection of triangles and/or quadrilaterals (e.g., no tubes, no bubbles, no octagons, etc.).

The main bottleneck for unknot recognition, 3-sphere recognition, prime decomposition and some other core 3-manifold algorithms is this: find a non-trivial connected normal surface with positive Euler characteristic, or determine that no such surface exists. Here *non-trivial* means that the surface is not just a collection of triangles that build a sphere around a vertex (in other words, there is at least one quadrilateral).

Currently the best known running time is $O(3^n \cdot \text{poly}(n))$, where $n$ is the number of tetrahedra. However, it seems plausible that this could be FPT in the treewidth of the dual graph of the triangulation. If so, this would give immediate FPT results for several important topological algorithms (e.g., those listed earlier).

Some further notes: normal surfaces are described by vectors in $\mathbb{R}^{7n}$ (each entry counts the number of triangles or quadrilaterals in some position in some tetrahedron). These vectors must satisfy some linear equations, linear inequalities, and combinatorial constraints. The inequalities and combinatorial constraints are local to each tetrahedron, and the equations come from gluings between adjacent tetrahedra. In this sense, the structure of the equations mirrors the structure of the dual graph, which is one reason such an FPT result might be plausible.

It is also an open problem to prove that this same problem is NP-hard.

## 5.3 Turning Planar Triangulations Inside Out

*Jeff Erickson (University of Illinois Urbana-Champaign, US) and Birgit Vogtenhuber (TU Graz, AT)*



### 5.3.1 Problem Statement

Let $T$ be an arbitrary planar straight-line triangulation, with a triangular outer face, such that the origin $(0,0)$ lies in the interior of some bounded face, and no two vertices lie on the same line through the origin. An *eversion* of $T$ is another planar straight-line triangulation satisfying three conditions:

- $T$ and $T'$ are embeddings of the same maximal planar graph $G$.
- Every vertex of $G$ lies on the same ray from the origin in both $T$ and $T'$.
- For every ray $r$ from the origin, the sequence of edges of $T$ that $r$ intersects is the reversal of the sequence of edges of $T'$ that $r$ intersects. Equivalently, every face of $T$, except the outer face and the face containing the origin, has the opposite orientation of the corresponding face of $T'$. (The outer face of $T$ must be the face containing the origin in $T'$, and vice versa.)

**Prove or disprove.** Every planar straight-line triangulation has an eversion.

We know how to construct an eversion of a given $n$-vertex triangulation $T$, *if one exists*, in $O(n^{\omega/2})$ time, by solving a certain system of linear equations. (The variables in this system are reciprocals of distances of vertices from the origin.) This is also the fastest algorithm we know to decide whether an eversion of $T$ exists.

This problem can be viewed as a radial variant of the *hierarchical planar graph drawing* problem, which asks whether a given planar graph has a straight-line embedding where each vertex lies on a prescribed y-coordinate. This problem has a long history, culminating in a linear-time algorithm by Klemz [1].

This problem comes out of work by Erickson and Howard [3] on spherical morphing. The problem can be equivalently formulated in terms of spherical embeddings as follows: Given a shortest-path triangulation $T$ in the open northern hemisphere, is there an isomorphic shortest-path triangulation $T'$ in the open southern hemisphere, such that every vertex lies on the same longitude in both $T$ and $T'$? (In the language of our paper: Is every northern triangulation $T$, is there a $\theta$-equivalent southern triangulation $T'$?)

### 5.3.2 Progress

A structure that might cause complications are “rotors” or “lens-shutters”. Essentially, this is a sequence of $e_0, \dots, e_k$ of disjoint line segments that cyclically overlap, such that for each $i$ along the ray through of one endpoint of $e_i$, the intersection with $e_{i+1 \bmod k}$ lies closer to the origin than the endpoint. I'll call this endpoint the “far endpoint” $f_i$ of $e_i$. Equivalently, every edge of a rotor also has a “near endpoint” $c_i$. Less formally, all near endpoints are visible from the origin, but no far endpoints.

Rotors are the simplest sets of segments that do not have a *radial shelling order*. A radial shelling order indexes the segments $s_1, s_2, \dots, s_n$, so that if any line segment from the origin to any segment $s_i$ crosses another segment $s_j$, then $j < i$. Arguments of of Awartani and Henderson [2, 3] imply that any set of segments that is radially shellable has an eversion.

An argument of Pallazi and Snoeyink [4] implies that if any ray from the origin misses all the segments, then the segments are radially shellable. Thus, rotors must surround the origin.

Further, we asked whether it might be that a rotor can be drawn with straight edges but needs a very specific shape of the polygon $F = f_0, \ldots, f_k$ formed by the far endpoints. We think that we can prove the following: If there is a realization in which $F$ is a star-shaped polygon, then there is also a realization where the vertices of $F$ lie on a circle, or any other *convex* polygon. Likewise, there should be one for any star-shaped polygon $F$ (with the vertices obeying the given angle condition).

It is still unclear whether a proof that every rotor can be everted would imply that any planar straight-line graph, or even every set of disjoint segments, can be everted.

### References


1. Boris Klemz. Convex drawings of hierarchical graphs in linear time, with applications to planar graph morphing. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPIcs*, pages 57:1–57:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPICS.ESA.2021.57`.
2. Marwan Awartani and David W. Henderson. Spaces of geodesic triangulations of the sphere. *Transactions of the Americal Mathematical Society*, 304(2):721–732, 1987. `doi:10.2307/2000738`.
3. Jeff Erickson and Christian Howard. Shelling and sinking graphs on the sphere. In Oswin Aichholzer and Haitao Wang, editors, *41st International Symposium on Computational Geometry (SoCG)*, volume 332 of *LIPIcs*, pages 47:1–47:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. `doi:10.4230/LIPIcs.SoCG.2025.47`.
4. Larry Palazzi and Jack Snoeyink. Counting and reporting red/blue segment intersections. *Graphical Models and Image Processing*, 56(4):304–310, 1994. `doi:10.1006/cgip.1994.1027`.


## 5.4 Morphing Long-Geodesic Embeddings on the Sphere

*Jeff Erickson (University of Illinois Urbana-Champaign, US)*



### 5.4.1 Problem Statement

A *geodesic* on the sphere is any arc of a great circle. A geodesic is *long* if it contains two antipodal points and *short* otherwise; every short geodesic is the unique shortest path between its endpoints. (Classically, the word "geodesic" refers only to shortest paths on the sphere; here I'm using the more general definition from Riemannian geometry.)

A *geodesic embedding* of a graph $G$ on the sphere maps the vertices of $G$ to distinct points and the edges of $G$ to interior-disjoint geodesics. A geodesic embedding is *short* if each of its edges is short, and *long* otherwise.

**Prove or disprove.** Any two isomorphic geodesic embeddings of the same planar graph on the sphere are connected by a continuous family of geodesic embeddings.

In 1944, Cairns [1] proved that any two *short* geodesic *triangulations* of the sphere are connected by a continuous family of *short* geodesic triangulations. Cairns's proof can be generalized to short geodesic embeddings of 3-connected planar graphs. Using ad-hoc arguments, we can show that the space of all geodesic embeddings of $K_4$ is connected, but nothing else appears to be known.

One significant complication is that if an edge of $G$ is short in one embedding but long in the other, the endpoints of that edge must be *exactly* antipodal at some moment during the morph. Another complication is that specifying the motion of the vertices is not sufficient; the same vertex coordinates and rotation system can be consistent with multiple geodesic embeddings, even if no two vertices are antipodal.

Morphs are known to exist between isomorphic straight-line planar embeddings (starting with Cairns), between isotopic geodesic embeddings on the flat torus [2], and between isotopic geodesic embeddings on any surface with negative curvature [3]. In fact, for all of those surfaces, the space of all geodesic embeddings of a given graph, within each isotopy class of embeddings, modulo rigid motions on the surface, is simply connected. The only orientable surface for which this problem is open is the sphere!

### References


1. Stewart S. Cairns. Isotopic deformations of geodesic complexes on the 2-sphere and on the plane. *Annals of Mathematics*, 45(2):207–217, 1944. URL: `http://www.jstor.org/stable/1969263`.
2. Erin Wolf Chambers, Jeff Erickson, Patrick Lin, and Salman Parsa. How to morph graphs on the torus. In Dániel Marx, editor, *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2759–2778, 2021. `doi:10.1137/1.9781611976465.164`.
3. Yanwen Luo, Tianqi Wu, and Xiaoping Zhu. The deformation space of geodesic triangulations and generalized Tutte's embedding theorem. *Geometry & Topology*, 7:3361–3385, 2023. `doi:10.2140/gt.2023.27.3361`.


## 5.5 Approximate Simplification of Flag Filtrations

*Marc Glisse (INRIA – Orsay, FR), Jeff M. Phillips (University of Utah – Salt Lake City, US)*



### 5.5.1 Problem Statement

Computing persistent homology usually takes time linear in the number of cells of the filtered complex involved. To speed things up, one then needs a small complex. The first (geometric) approach is to define a small complex to begin with. The second (combinatorial) approach, the one I am interested in today, is to simplify a filtration before computing its persistent homology. This only makes sense if we can simplify the complex without actually listing its cells explicitly (same cost as computing PH). A common case is flag filtrations (like Vietoris-Rips), which can be represented by a graph with a filtration value for each vertex and edge, and where higher dimensional cells are implicitly defined as the cliques of the graph. The goal of the simplification is then to replace the graph with a smaller graph before listing its simplices to compute PH (so there are fewer simplices to list).

Glisse and Pritam [1] provide an approach for simplifying the graph while exactly preserving its PH. However, in many cases, we would be happy with an approximation.

**Question.** How can we simplify the graph, more than in the exact case, but still with guarantees on the error?

### 5.5.2 Progress

The following approaches and questions arose during the seminar.

- Forcibly reintroduce some geometry, embedding in dim  log n.
- Could rounding to a grid be more efficient than what the paper does, sometimes, because it gives some edges the same value and thus helps jump over them? Needs experiments; but initial experiments showed this did not provide an improvement.
- Simplification with Morse matching (not at the graph level, just a first step to understand what is possible?)

**References**

1 Marc Glisse and Siddharth Pritam. Swap, shift and trim to edge collapse a filtration. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG)*, volume 224 of *LIPIcs*, pages 44:1–44:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.SOCG.2022.44.

## 5.6 Sub-Alpha-Complex in High Dimension

*Marc Glisse (INRIA – Orsay, FR), Jeff M. Phillips (University of Utah – Salt Lake City, US)*



### 5.6.1 Problem Statement

**Goal.** compute (persistent) homology of low dimension of a union of balls in high dimension.

**Usual approaches.**

1. Cech (or Rips) of size $n^{(k+1)}$.
2. Alpha-complex which "requires" computing the full Delaunay triangulation of size (up to) $n^{(d/2)}$. This is only better than Approach 1 for small $d$.
3. Handle each simplex using linear/quadratic programming [1].

**Question.** Can we improve on it?

- we don't specifically need the alpha-complex, something closer to the wrap complex would be even better
- more output sensitive?
- use more LP before QP?
- etc

### 5.6.2 Progress

It was also suggested to maybe we use function approximation coresets to get some sort of approximation relevant for persistence (diagram) approximation using function approximation coresets. For instance in the line of work by Langberg and Schulmann [2] or and more recently (and perhaps more generally) by Alishahi and Phillips [3]. Constructions like Growing Neural Gas [4] could be viewed as an approximation of what we want, and could be further generalized. Funke and Ramos work on smooth-surface reconstruction [5] might also be helpful.

**References**

1 Erik Gunnar Carlsson and John Carlsson. Computing the alpha complex using dual active set quadratic programming. arXiv report, 2023. URL: https://arxiv.org/abs/2310.00536.
2 Michael Langberg and Leonard J. Schulman. Universal epsilon-approximators for integrals. In Moses Charikar, editor, *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 598–607, 2010. doi:10.1137/1.9781611973075.50.
3 Meysam Alishahi and Jeff M. Phillips. No dimensional sampling coresets for classification. In *41st International Conference on Machine Learning (ICML)*. OpenReview.net, 2024. URL: https://openreview.net/forum?id=jS3CMHtYJD.
4 Bernd Fritzke. A growing neural gas network learns topologies. In Gerald Tesauro, David S. Touretzky, and Todd K. Leen, editors, *Advances in Neural Information Processing Systems 7, [NIPS Conference, Denver, Colorado, USA, 1994]*, pages 625–632. MIT Press, 1994. URL: https://proceedings.neurips.cc/paper_files/paper/1994/hash/d56b9fc4b0f1be8871f5e1c40c0067e7-Abstract.html.
5 Stefan Funke and Edgar A. Ramos. Smooth-surface reconstruction in near-linear time. In David Eppstein, editor, *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 781–790, 2002. URL: http://dl.acm.org/citation.cfm?id=545381.545484.

## 5.7 Treewidth of 3-Manifolds

*Benjamin Burton (University of Queensland – Brisbane, AU), Kristóf Huszár (TU Graz, AT), Jonathan Spreer (University of Sydney, AU), Yelena Yuditsky (ULB-Brussels, BE)*



### 5.7.1 Problem Statement

The *treewidth* $\mathrm{tw}(M)$ of a compact 3-manifold $M$ is defined as the smallest treewidth the dual graph $\Gamma(T)$ of any generalized triangulation $T$ of $M$ may have. In other words

$$\mathrm{tw}(M) = \min\{\mathrm{tw}(\Gamma(T)) : T \text{ is a triangulation of } M\}.$$

The study of $\mathrm{tw}(M)$ is motivated by existing fixed-parameter tractable (FPT) algorithms to efficiently compute various 3-manifold invariants from triangulations with dual graph of bounded treewidth [1, 2, 3, 4, 5] – several of which are known to be NP-hard in general. The treewidth has been related to various other properties of 3-manifolds, such as Heegaard genus [7, 9], hyperbolic volume [10] (cf. [6]), or JSJ decompositions [8], but several fundamental questions remain to be answered.

**Question 1.** What is the complexity of computing the treewidth of a 3-manifold?

It is natural to ask, what is the computational complexity of the treewidth of a 3-manifold given by a triangulation. We believe that it is at least as hard as computing the treewidth of a graph, which is known to be NP-hard. We know that any algorithm to compute the treewidth of a 3-manifold yields (via a polynomial-time reduction) a constant-factor approximation algorithm for the treewidth of bounded-degree graphs, see [8, Remark 15].

**Question 2.** What is the treewidth of the 3-torus?

The 3-torus $\mathbb{T}^3 = \mathbb{S} \times \mathbb{S} \times \mathbb{S}$ is defined as the Cartesian product of three circles. Perhaps surprisingly, the exact value of $\mathrm{tw}(\mathbb{T}^3)$ is unknown. It is known that it is at least two (since

3-manifolds with treewidth at most one have been classified [7, Theorem 2]) and at most four (as shown by its minimal triangulation, which has treewidth four [12, Closed Census]).

**Question 3.** Can we classify closed 3-manifolds of treewidth two?

Given the classification of closed 3-manifolds with treewidth one, it is natural to ask for a similar result regarding 3-manifolds with treewidth two. We know that all Seifert fibered spaces over the 2-sphere or a non-orientable surface have treewidth of at most two [7, Section 5]. Answering Question 3 would also help us get closer to answering Question 2.

**Question 4.** What is the relationship between minimal triangulations of 3-manifolds and triangulations with minimal treewidth?

We know that minimal triangulations of 3-manifolds (as measured by the number of tetrahedra) are not always of minimal treewidth: the minimal triangulation of the *Poincaré homology sphere* consists of five tetrahedra and its dual graph is $K_5$, whose treewidth is four. At the same time, the Poincaré homology sphere is a Seifert fibered space over the 2-sphere, and as such it has a (larger) triangulation with a dual graph of treewidth two [7, Corollary 20]. This is the only example of this kind that we know of.

**Question 5.** Is the Homeomorphism Problem of triangulated 3-manifolds FPT in the parameter $k = \max\{\mathrm{tw}(\Gamma(T_1)), \mathrm{tw}(\Gamma(T_2))\}$?

The Homeomorphism Problem is a fundamental decision problem in computational topology. It asks whether two given triangulated $d$-manifolds $T_1$ and $T_2$ are homeomorphic. For $d = 2$, the problem can be solved in polynomial time, but for $d \geq 4$ it is undecidable. In dimension $d = 3$ it is known to be at least as hard as the Graph Isomorphism Problem, but is elementary recursive. See [11, Section 3.1] for context and details.

### 5.7.2 Progress

During the seminar, we focused primarily on Questions 2 and 3.

**Question 2.** With the help of `Regina` [12], we ran through all $9,944,398$ triangulations of the 3-torus that can be accessed from its minimal triangulation using Pachner moves (bistellar flips) never exceeding 14 tetrahedra at any step. For each triangulation, we used a greedy algorithm to compute a small-width tree decomposition of its dual graph. This greedy algorithm *probably* finds the real treewidth for such small graphs, though this is not guaranteed. We did not find anything with width less than 4 in this list. The following code was used to perform this task in `regina-python`.

```
def test(sig, tri):
    print(TreeDecomposition(tri).width(), sig)
    return False

Example3.threeTorus().retriangulate(8, 1, test)
```

The first parameter `8` above means at most eight extra tetrahedra at any stage (so at most 14 in total, since the minimal triangulation has 6 tetrahedra). The second parameter `1` instructs the program to run in single-thread mode.

**Question 3.** In regard to the classification of closed 3-manifolds with treewidth two the following result may serve as a starting point.

▶ **Theorem 1** (folklore; adapted from [13, Theorem 3(i)]). *If a multigraph $G = (V, E)$ has treewidth at most two and is not the empty graph, then $G$ has a vertex of degree at most two. (Here the degree of a vertex $v \in V$ denotes the number of other vertices in $V$ adjacent to $v$.)*

Based on Theorem 1 we started exploring the following approach. Given a triangulation $T$ of a closed 3-manifold with $\mathrm{tw}(\Gamma(T)) = 2$, pick a vertex $v$ in $\Gamma(T)$ of degree at most two. Observe that, for the 1-neighborhood of $v$ in $\Gamma(T)$ there are five possibilities (Figure 1). Each of these five cases corresponds to a finite number of admissible substructures within a valid triangulation of a 3-manifold. Using `Regina`, we enumerated these substructures. Aiming for a recursive classification method, our next goal is to understand how the topology of the underlying manifold changes when we "reduce" the triangulation $T$ around the tetrahedron corresponding to the vertex $v$.

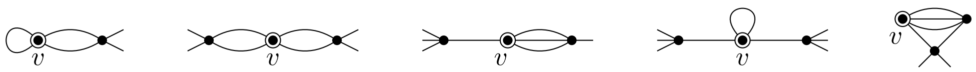


**Figure 1** The five possible local pictures around a vertex $v \in \Gamma(T)$ of degree at most two.

## References


1 Benjamin A. Burton and Rodney G. Downey. Courcelle's theorem for triangulations. *J. Comb. Theory A*, 146:264–294, 2017. doi:10.1016/J.JCTA.2016.10.001.

2 Benjamin A. Burton, Thomas Lewiner, João Paixão, and Jonathan Spreer. Parameterized complexity of discrete Morse theory. *ACM Trans. Math. Softw.*, 42(1):6:1–6:24, 2016. doi:10.1145/2738034.

3 Benjamin A. Burton, Clément Maria, and Jonathan Spreer. Algorithms and complexity for Turaev–Viro invariants. *J. Appl. Comput. Topol.*, 2(1-2):33–53, 2018. doi:10.1007/S41468-018-0016-2.

4 Benjamin A. Burton and William Pettersson. Fixed parameter tractable algorithms in combinatorial topology. In Zhipeng Cai, Alex Zelikovsky, and Anu G. Bourgeois, editors, *20th International Conference Computing and Combinatorics (COCOON)*, volume 8591 of *LNCS*, pages 300–311. Springer, 2014. doi:10.1007/978-3-319-08783-2_26.

5 Benjamin A. Burton and Jonathan Spreer. The complexity of detecting taut angle structures on triangulations. In Sanjeev Khanna, editor, *Proceedings 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 168–183, 2013. doi:10.1137/1.9781611973105.13.

6 Kristóf Huszár. On the pathwidth of hyperbolic 3-manifolds. *Comput. Geom. Topol.*, 1(1):1:1–1:19, 2022. URL: https://www.cgt-journal.org/index.php/cgt/article/view/4.

7 Kristóf Huszár and Jonathan Spreer. 3-manifold triangulations with small treewidth. In Gill Barequet and Yusu Wang, editors, *35th International Symposium on Computational Geometry (SoCG)*, volume 129 of *LIPIcs*, pages 44:1–44:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.SOCG.2019.44.

8 Kristóf Huszár and Jonathan Spreer. On the width of complicated JSJ decompositions. In Erin W. Chambers and Joachim Gudmundsson, editors, *39th International Symposium on Computational Geometry (SoCG)*, volume 258 of *LIPIcs*, pages 42:1–42:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.SOCG.2023.42.

9 Kristóf Huszár, Jonathan Spreer, and Uli Wagner. On the treewidth of triangulated 3-manifolds. *J. Comput. Geom.*, 10(2):70–98, 2019. URL: https://doi.org/10.20382/jogc.v10i2a5, doi:10.20382/JOGC.V10I2A5.

10 Clément Maria and Jessica S. Purcell. Treewidth, crushing and hyperbolic volume. *Algebr. Geom. Topol.*, 19(5):2625–2652, 2019. doi:10.2140/agt.2019.19.2625.

**11** Marc Lackenby. Algorithms in 3-manifold theory. In *Surveys in differential geometry 2020. Surveys in 3-manifold topology and geometry*, volume 25 of *Surv. Differ. Geom.*, pages 163–213. Int. Press, Boston, MA, 2022.

**12** Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. `http://regina-normal.github.io/`, 1999–2023.

**13** Hans L. Bodlaender, Benjamin A. Burton, Fedor V. Fomin, and Alexander Grigoriev. Knot diagrams of treewidth two. In Isolde Adler and Haiko Müller, editors, *46th International Workshop Graph-Theoretic Concepts in Computer Science (WG)*, volume 12301 of *LNCS*, pages 80–91. Springer, 2020. `doi:10.1007/978-3-030-60440-0\_7`.


## 5.8 WSPD for Low Density Graphs

*Joachim Gudmundsson (The University of Sidney, AU)*



### 5.8.1 Problem Statement

Given a geometric graph $G = (V, E)$ where the vertices are points in the plane and the edge weights is the Euclidean distance between its endpoints, the main goal is to construct a $(1+\varepsilon)$-approximate distance oracle ($O(1)$) query time) using linear space and $O(m + n \log n)$ preprocessing.

Of course this is not possible for arbitrary geometric graphs, instead we want to investigate restricted graph classes. In particular, we want to investigate graph classes that include road networks (think about queries in Google maps). Such a graph class should allow for edge intersections and include naturally occurring graphs such as grid graphs or "cul-de-sac" graphs. There are a number of "realistic" graph classes, e.g. planar, low-dilaton, $c$-packed, small highway dimension, small skeleton dimension, and small doubling dimensions that allow for approximate distance oracles, but these graph classes do not include road networks.

There are a few graph classes that include road networks, e.g. disk neighbourhood, crossing degeneracy, and bounded growth graphs but they do not have efficient algorithmic tool sets. Our favourite natural graph class that includes road networks is $\lambda$-low density graphs or the more general class of $\tau$-lanky graphs. A geometric graph is $\lambda$-low density if for every $p$ and $r$ the ball $B(p, r)$ with centre at $p$ and radius $r$ intersects at most $\lambda$ edges of $E$ of length at least $r$.

In a recent paper [1] we showed that $O(1)$-low-density graphs have a $(1+\varepsilon)$-approximate distance oracle using $O(n \log n)$ space and $O(n^{3/2} polylogn)$ preprocessing. The key insight used to obtain this result is a Well-Separated Pair Decomposition (WSPD) of $V$ in the graph metric of $O(n \log n)$ size (which gives the space bound). There are two natural techniques that could be used to improve this result:

1. Improve the size of the WSPDs for low density graphs. No non-trivial lower bound is known. Can we close the gap between $\Omega(n)$ and $O(n \log n)$? Unit disk graphs (UDG) have the same bounds. Even subgraphs of a grid graph would be interesting (or even a tree in a grid graph). In dimensions greater than 2 the lower bound for the WSPD for UDG is $\Omega(n^{2-2/d})$.
2. Find a tree cover $T_1, \ldots, T_k$ of $G$ of constant size and $(1+\varepsilon)$ dilation. A tree cover is a set of trees s.t. for every pair of points $p$ and $q$ (a) there exists a tree $T_i$ with $d_{T_i}(p, q) \leq (1+\varepsilon) \cdot d_G(p, q)$ and (b) for every tree $T_i$ $d_{T_i}(p, q) \geq d_G(p, q)$. If such a tree cover exists we would probably get a $(1+\varepsilon)$-approximate distance oracle of linear space.

Any results for 1 or 2 would give improved results.

### References

1 Joachim Gudmundsson and Sampson Wong. A well-separated pair decomposition for low density graphs. *CoRR*, abs/2411.08204, 2024. doi:10.48550/ARXIV.2411.08204.

## 5.9 Consistent System of Shortest Paths

*Francis Lazarus (CNRS, Grenoble, FR), Jeff Erickson (University of Illinois – Urbana-Champaign, US), Raimund Seidel (Universität des Saarlandes – Saarbrücken, DE), and Ling Zhou (Duke University, Durham, US)*



This problem has been resolved during the seminar. More precisely, the problem was already solved by several authors, starting with Hartvigsen and Mardon [6].

### 5.9.1 Problem Statement

Given a graph $G = (V, E)$ with a weight function $w : E \mapsto \mathbb{R}_+$, the problem is to compute a *consistent* system $\mathcal{P}$ of shortest paths w.r.t. to $w$, so that for every $u, v \in V$ there is a shortest path from $u$ to $v$ in $\mathcal{P}$ and for every paths $p, q \in \mathcal{P}$ their intersection is connected (possibly empty), i.e. is a commmon subpath.

If shortest paths are unique for the weight function, then the set of shortest paths is consistent. There are various perturbation schemes that ensure uniqueness of shortest paths, either non-deterministic or deterministic but with an extra overhead in complexity (see discussion below). The goal here is to produce a consistent system of shortest paths in a deterministic way with no overhead cost compared to the traditional all-pair-shortest-paths algorithms. The hope is to compute a consistent system in a more efficient way, without enforcing uniqueness of shortest paths...

Computing a consistent system of shortest paths appears to be equivalent to compute a *consistent system of rooted shortest path trees.* Here, consistent means that for every $u, v \in V$, the (shortest) path from $u$ to $v$ in the tree $T_v$ rooted at $v$ is the reverse of the path from $v$ to $u$ in the tree $T_u$ rooted at $u$, i.e. $T_v(u) = (T_u(v))^{-1}$

This last characterization is itself equivalent to the following: for every $u, v \in V$

1. $T_v(u)$ and $T_u(v)$ have the same number of edges, and
2. denoting by $\pi_u$ the parent pointer in $T_u$, either $\pi_v(u) = v$ or $\pi_v(u) = \pi_{\pi_u(v)}(u)$.

The first condition is easily enforced by adding a constant perturbation to the edge weights. The second is "local" in the sense that it only involves pointer comparisons (as opposed to compare the paths $T_v(u)$ and $T_u(v)$).

### 5.9.2 Progress

A simple randomized perturbation of the edge weights provides uniqueness of shortest paths with high probability. The perturbations are obtained by drawing independently and uniformly at random an integer in $[1, n^4]$ for every of the $n$ edges. See [1, Sec 6.1] for a full description of this non-deterministic perturbation scheme based on the Isolating Lemma of Mulmeley, Vazirani and Vazirani [2].

There are two known deterministic approaches to compute weight perturbations to guarantee uniqueness of shortest paths for graphs embedded on surfaces.

- One uses lexicographic perturbation, incurring a $O(\log n)$ overhead [1]. This perturbation scheme assigns a unique index to each edge in the graph, and then breaks ties between shortest paths by comparing the lexicographically smallest index where the two paths differ. The scheme uses a dynamic-forest data structure to maintain a shortest-path tree during its construction via Dijkstra's algorithm. In addition to uniqueness, this scheme guarantees that shortest paths are symmetric: The reversal of any shortest path is also a shortest path.
- The other scheme uses homology and incurs a $O(g)$-factor overhead [3]. This "holiest" perturbation scheme uses a vector of $2g+1$ perturbation terms for each edge, $2g$ of which are homology signatures, and the last of which stores the number of vertices in a fixed dual spanning tree on one side of the edge. Path lengths are vectors of length $2g+2$, which are added as vectors and compared lexicographically. This is a generalization of the leftmost-shortest path tie-breaking rule for planar graphs, originally proposed by Ford and Fulkerson. Like the leftmost shortest paths in planar graphs, holiest shortest paths are not necessarily symmetric. Consider two opposite corner vertices in a square unweighted grid graph.

It should be emphasized that the main problem here is not making individual edge weights distinct, but making certain sums of edge weights distinct.

Raimund Seidel points out that the Floyd-Warshall all-pairs shortest path algorithm computes a consistent set of shortest paths. Recall that the $i$th iteration of Floyd-Warshall computes the shortest path from each vertex $u$ to each vertex $v$ whose intermediate vertices have index at most $i$. It follows that FW computes shortest paths that lexicographically minimize the sorted vector of vertex indices.

Unfortunately, Floyd-Warshall runs in $O(n^3)$ time, which is more than we want to spend for most surface-graph algorithms.

It appears that Wullf-Nilsen already described a similar perturbation scheme for arbitrary graphs that allows Dijkstra to compute paths that minimize length, then the number of edges, and finally the sorted vector of vertex indices, all in $O(n \log n)$ time [4] for a single shortest path tree. By repeating Dijkstra from every vertex we thus obtain a consistent system of shortest paths in $O(n^2 \log n)$ time. Note that Wullf-Nilsen uses $O(\log n)$ extra pointers per vertex, so that running a single (modified) Dijkstra actually uses $O(n \log n)$ space instead of the usual linear space implementation. In fact, Hartvigsen and Mardon [6] formerly described an implementation that runs all the Dijsktra's in parallel in $O(n^2 \log n)$ time without using extra pointers.

See also Borradaile, Sankowski, and Wulff-Nilson [5, Sec. 7] for an implementation using top trees.

This is the basis for the $O(\log n)$-factor lexicographic perturbation scheme described above; the $O(\log n)$ overhead does not apply to Dijkstra's algorithm, but rather to the pivots inside our multiple-source shortest paths algorithm. Similarly, Borradaile, Sankowski, and Wulff-Nilson show that more modern shortest-path algorithms based on dense-distance graphs incur a $O(\log^2 n)$ overhead.

## References

1 Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM Journal on Computing*, 42(4):1542–1571, 2013. https://arxiv.org/pdf/1202.0314.

2 Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
3 Jeff Erickson, Kyle Fox, and Luvsandondov Lkhamsuren. Holiest minimum-cost paths and flows in surface graphs. In *50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 2018. URL: `https://arxiv.org/pdf/1804.01045`, `doi:10.1145/3188745.3188904`.
4 Christian Wulff-Nilsen. Minimum cycle basis and all-pairs min cut of a planar graph in subquadratic time. arXiv report, 2009. `doi:10.48550/arXiv.0912.1208`.
5 Glencora Borradaile, Piotr Sankowski, and Christian Wulff-Nilsen. Min st-cut oracle for planar graphs with near-linear preprocessing time. *ACM Transactions on Algorithms*, 11(3):1–29, 2015. `doi:10.1145/2684068`.
6 David Hartvigsen and Russell Mardon. The all-pairs min cut problem and the minimum cycle basis problem on planar graphs. *SIAM Journal on Discrete Mathematics*, 7(3):403–418, May 1994. `doi:10.1137/s0895480190177042`.

## 5.10 Graph Tiles

*Oswin Aichholzer (TU Graz, AT), Robert Ganian (TU Wien, AT), Phillip Keldenich (TU Braunschweig, DE), Maarten Löffler (Utrecht University, NL), Gert Meijer (NHL Stenden Hogeschool – Leeuwarden, NL), Alexandra Weinberger (FernUniversität in Hagen, DE), Carola Wenk (Tulane University – New Orleans, US)*



### 5.10.1 Problem Statement

In a *graph tiling* problem, we are given as input a set of square tiles with pieces of graph drawn (embedded) on them, and the output is an arrangement of those tiles into a larger area, like in Figure 2.

Now, we are interested in output arrangements that satisfy certain properties.

- Tiles should "match up" on the sides.
- We may want that the resulting graph is connected.
- We may want to minimize the length of longest path between any two vertices in the resulting graph.
- We may want to minimize the area of the largest face, or maximize the area of the smallest face.

In general, even satisfying the first property alone is already NP-hard. However, this is not the case when there is only a single type of tile. Can we parametrize this problem by the number and/or complexity of the tiles? And for really simple sets of tiles, how many of these properties can we still test for efficiently?

The following variants and limitations might also be interesting to consider, and may influence to complexity of the problem.

- One can consider the tiles to be arranged into different shapes.
  - A square
  - A rectangle
  - A polygon without holes
  - A polygon with holes
  - . . .

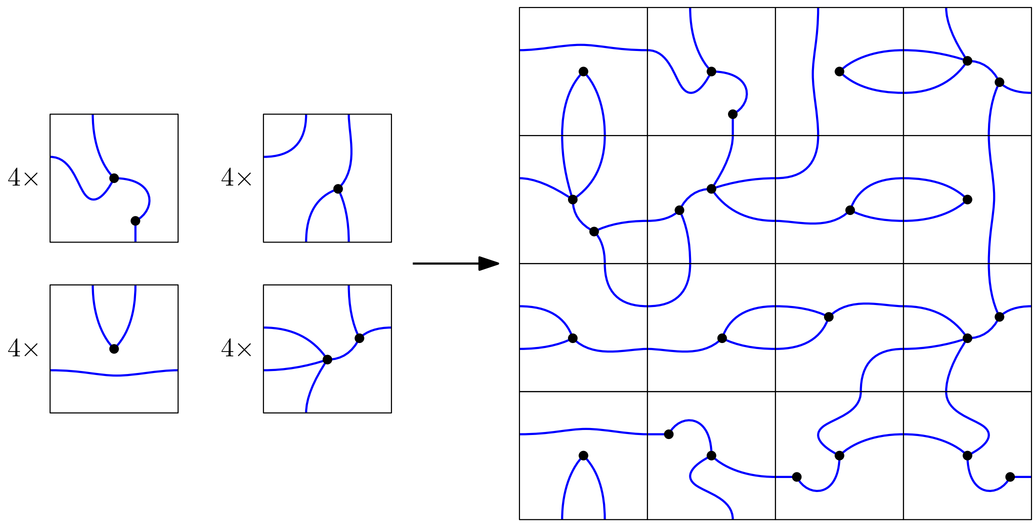


**Figure 2** A graph tiling.

- For the boundary of the form (where tiles do not meet other tiles and so are not restricted by having to match another tile) we can …
  - allow everything.
  - prescribe something (natural would be that all are empty or all have a path going in or there is a prescription for each).
  - avoid having a boundary by putting the tiles on a torus.
  - …
- Restrictions on the graph other than (or in addition to) connectedness could also be made. For example, one could require the graph to be a tree.

### 5.10.2 Progress

During the seminar, we considered the following setting: Given $n^2$ tiles that have constantly many different types. Can all the tiles be arranged into one square such that the sides of the tiles match and the boundary is arbitrary?

When we restrict the tile types to the six types of tiles from Figure 3 and allow rotation, then if all tiles are of the same type, it is always possible to arrange them into a valid square. If tiles are of exactly two types, then we have polynomial time algorithms to decide if they can be arranged into a square. (For most combinations, the answer is always yes or always no.) For three types most cases still give easy polynomial time algorithms, but some cases are still open.

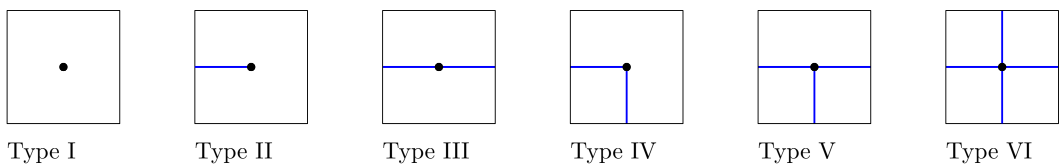


**Figure 3** Six types of tiles.

## 5.11 Approximate Minlink Paths in Subquadratic Time

*Valentin Polishchuk (Linköping University, SE), Sayan Bandyapadhyay (Portland State University, US), Mark de Berg (TU Eindhoven, NL), Joachim Gudmundsson (The University of Sydney, AU), André Nusser (CNRS, FR), Frank Staals (Utrecht University, NL)*



### 5.11.1 Problem Statement

Let P be a polygonal domain with n vertices. A minlink path between points s,t in P is an s-t path with fewest edges (links). Computing the path is 3SUM-hard, and an $O(\sqrt{n})$ approximation can be found in subquadratic time [1]. Can a better approximation be computed in subaquatic time?

### 5.11.2 Progress

The following are sketches of ideas that arose during the seminar.

#### $O(n \log n)$ space, $O(\log n)$ query, $+2$ approx

Build a balanced hierarchical subdivision. For the bottom endpoint; construct a SPM with respect to the link distance (which uses $O(n)$ space). Recurse.

**Space:** $O(n \log n)$

**Query:** For $LCA(s,t)$, and all its ancestors $v_1, \ldots, v_k$, query the min-link path from $s$ to $v_i$ and from $v_i$ to $t$. Report the cheapest path.Beschreibung

▶ Claim 1. This is a +2 approx.

**Proof.** Look at min-link path from $s$ to $t$: If $s, t$ are on opposite sides of chord through $v_i$, the approx ratio is +2. If $s, t$ are on same side of $v_i$, the approx is +1. ◀

#### $O(n^2)$ space, $O(\log n)$ query

Let v be the first vertex on the geodesic shortest path from $s$ to $t$.

▶ Claim 2. There exists a min-link path from s to t that passes through v.

Consider the (geodesic) SPM of v, and consider the min-link diagram $LPM(v, \phi)v$, in which the first link has fixed orientation $\phi$.

▶ Claim 3. Any edge of $LPM(v, \phi)$ is contained in a triangle of $SPM(v)$.

**Proof sketch.** Assume by contradiction that $e = wq$ of $SPM(v)$ intersects an edge of $LPM(v, \phi)$ in some point $p$. Now consider two points $p_1$ and $p_2$ on $e$ on opposite sides of $p$. This mean the min-link distance from $v$ (with orientation $\rho$) to $p_1$ is $i$, and the min-link distance from $v$ to $p_2$ is $j \neq i$.

However, by Claim 2 there exists a min-link path from $p_1$ to $v$ or $s$ that goes through $w$. Similarly, there exists a min-link path from $p_2$ to $v$ or $s$ that goes through $w$. They must have the same length. Contradiction. ◀

(We may want to argue that an edge of $LMD(s)$ cannot intersect an edge of $SPM(v)$.)

So, for any triangle of $SPM(v)$, there is at most one edge $e(\phi)$ of $LPM(v, \phi)$ in it.

▶ Claim 4. $e(\phi)$ rotates monotonically as a function of $\phi$ (or it is fixed).

So, for any triangle; we store the interval $(-\infty, \phi_1]$ on which the triangle is simply contained in a "length $i+1$" LPM region, and the interval $[\phi_2, \infty)$ in which it is "length $i+1$". In the interval $[\phi_1, \phi_2]$ we store the function describing how $e(\phi)$ moves.

▶ Claim 5. For any point $q$, the link-distance from $v$ to $q$ whose first link has orientation $\phi$ is either $i$ or $i+1$ (over all $\phi$s.)

Hence, for every vertex $v$ we have some $O(n)$ space structure. (The SPM, and per triangle the description of one $LPM(v, phi)$ edge).

So $O(n^2)$ in total. In addition we store the Guibas Hershberger structure. Space remains $O(n^2)$.

Query: We locate v using the Guibas Hershberger structure. This also gives us the orientation $\phi$. We now point locate $t$ in the $SPM(v)$; we then compare $\phi$ against the interval $[phi_1, phi_2]$ and report either $i$ or $i+1$.

### $O(n)$ space $O(k \log n)$ query

**Main idea.** We simulate the min-link propagation using the Guibas Hershberger structure.

Store the Guibas Hershberger structure: Compute the first vertex v on geodesic shortest path from $s$ to $t$. There exists a min-link path through $v$.

Consider the window $w_1 = vq$ on the ray from $s$ through $v$. Consider the weak visibility polygon $V$ of $w_1$. We essentially want to compute which pocket of $P \setminus V$ contains $t$. Let $w_2$ be the window (of $V$) separating $s$ from $t$.

Query the funnel from $t$ to $v$ and $q$ (that has apex $a$).

▶ Claim 6. $a$ is a vertex of $w_2$.

**Main idea.** We can use the Guibas Hershberger structure to actually find $a$ in $O(\log n)$ time. We can then use a ray shooting query to compute the other endpoint of $w_2$ (i.e. by extending the "middle" triangle of the funnel in the other direction"). So, we simulate one step of the visibility propagation in $O(\log n)$ time; this gives us $O(k \log n)$ query time.

### $O(n^2)$ space, $O(\log n)$ query, $+4$ approx

For every pair of vertices, store the minimum link distance.

Store a triangulation of the free space.

Locate the triangles containing $s$ and $t$. Let $v$ and $w$ be vertices incident to these triangles. Report 2+ length min-link path $(v, w)$.

▶ Claim 7. This is a +4 approx for $d(s, t)$.

### $O(n^4)$ space, $O(\log^c n)$ query, $+2$ approx

Rather than picking some arbitrary first vertex v and last vertex $w$, we actually find the best pair.

Let $v, w$ be the pair of vertices ($v$ visible from $s$, $w$ visible from $t$) for which min-link distance from $v$ to $w$ is minimum. We will report the path $s$ to $v$ to $\ldots w$ to $t$.

▶ Claim 8. This is a +2 approx.

**Main idea.** We build some three-level structure to efficiently find $v$ and $w$.

For every vertex, compute its visibility polygon, store it in a DS for stabbing queries, so that we can report the polygons stabbed by $s$ as $O(\log n)$ canonical subsets. For each such canonical subsets we again store a structure for stabbing queries with $t$ as $\log n$ canonical subsets. As a third level structure: simply store the link distance.

(This is very similar to the $2 - pt$ geodesic query structure Sarita de Berg, Tillmann Miltzow, and Frank Staals [2] constructed before.)

#### $O(n^8)$ space, $O(\log n)$ query, +1 approx

The vague idea is the following: We do the same as above, except that for every vertex $v$, we compute its minimum link map (which has $O(n^4)$ complexity or so); we throw all these regions into our stabbing structure. (This may also require us to pick the right angle from which we leave $v$.)

#### $O(n^{1+1/k})$ space, $O(\log n)$ query, $*2$ approx

**Main idea.** Same as our $O(n^2)$ space structure, except that we store approximate distances between every pair of vertices.

There exists an $O(n^{1+1/k})$ space structure to store distances in an arbitrary graph [3]. We use this on the complete graph whose edge lengths are the link-distances.

**References**


1 Joseph S. B. Mitchell, Valentin Polishchuk, and Mikko Sysikaski. Minimum-link paths revisited. *Comput. Geom.*, 47(6):651–667, 2014. URL: `https://doi.org/10.1016/j.comgeo.2013.12.005`, `doi:10.1016/J.COMGEO.2013.12.005`.
2 Sarita de Berg, Tillmann Miltzow, and Frank Staals. Towards space efficient two-point shortest path queries in a polygonal domain. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG)*, volume 293 of *LIPIcs*, pages 17:1–17:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. `doi:10.4230/LIPICS.SOCG.2024.17`.
3 Shiri Chechik. Approximate distance oracles with constant query time. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 – June 03, 2014*, pages 654–663. ACM, 2014. `doi:10.1145/2591796.2591801`.


## 5.12 Alexander Theorem for Knotted Spheres

*Mathijs Wintraecken (Centre Inria d'Université Côte d'Azur, FR) and Kristóf Huszár (TU Graz, AT)*



### 5.12.1 Problem Statement

Alexander [1] proved that every knot or even link can be represented as a closed braid, that is, a braid with ends identified. This result is known as the Alexander theorem. The question is if this has a generalisation to knotted spheres in the following sense: Can every knotted 2-sphere in $\mathbb{R}^4$ be embedded sufficiently close to the standard symmetrically embedded 2-sphere $\mathbb{S}^2 \subset \mathbb{R}^3 \subset \mathbb{R}^4$ such that the closest point projection onto the symmetrically embedded 2-sphere is a local diffeomorphism? Has this question been answered before?

### 5.12.2 Progress

With the help of ChatGPT, Kristóf Huszár found the appropriate literature references, of which [2, Chapter 1] already contains this result with a pedagogical explanation.

#### References

1 James W. Alexander. A lemma on systems of knotted curves. *Proceedings of the National Academy of Sciences*, 9(3):93–95, 1923. doi:10.1073/pnas.9.3.93.

2 Scott Carter, Seiichi Kamada, and Masahico Saito. *Surfaces in 4-Space*. Springer Berlin Heidelberg, 2004. doi:10.1007/978-3-662-10162-9.


## 5.13 Universal Point Sets

*Alexander Wolff (Universität Würzburg, DE), Stefan Felsner (TU Berlin, DE), Michael Hoffmann (ETH Zürich, CH), Maarten Löffler (Utrecht University, NL), Fabrizio Montecchiani (University of Perugia, IT), Yoshio Okamoto (The University of Electro-Communications, JP), Jonathan Spreer (University of Sydney, AU)*



### 5.13.1 Problem Statement

A set $S$ of points in the plane is called *universal* with respect to a given drawing style and an integer $n$ if every graph on $n$ vertices admits a drawing in the given style such that the vertices are mapped to the points of $S$. For planar graphs, for example, the drawing style could require crossing-free straight-line edges. When considering drawings where the edges are circular arcs [1] or polygonal lines [4], the encoding of the edges may also need to be considered; for instance, a third point on an arc or the bends of a polygonal line. Also the number of edges can be a relevant parameter, which leads to the notion of *universal geometric (or topological) graphs* [2]. Another direction is to study universal point sets for restricted classes of planar graphs, such as planar 3-trees [3].

### 5.13.2 Progress

**Higher dimensions.** One possible way to transfer the problem of universal point sets into higher dimensions is the following: Can we find a universal point set in 3-space for all $n$-vertex triangulations of the 2-sphere?

By a classical result due to Steinitz, every triangulated 2-sphere can be embedded into 3-space in convex position (in particular, as a polyhedral embedding – that is, with straight-line edges and flat triangles). Moreover, we know that all of these embeddings can be chosen to have integer coordinates.

By a result of André Schulz [5], the maximum coordinate of such an integer embedding is $O(2^{7.55n})$, where $n$ is the number of vertices. This trivially yields a universal point set of size $\left(c \cdot 2^{7.55n}\right)^3$. This result holds for arbitrary polytopes. The hope is that the bound for simplicial polytopes (triangulations of the 2-sphere) would be much smaller. There might be hints in the proof of the above result that can help us.

This touches on a famous problem in discrete geometry: The polyhedral realization problem for triangulated (or polyhedral) surfaces. Not every triangulation of an orientable surface has a polyhedral embedding into 3-space. In fact, an $n$-vertex triangulation of an orientable surface can have genus quadratic in $n$, but the highest genus surfaces that are known to be polyhedrally embeddable have genus $n \log(n)$. Even more, deciding polyhedral embeddability for a given triangulation is likely a hard problem in general. See [6] for a nice overview of the polyhedral realization problem.

**Central vertices in Schnyder woods.** To construct universal point sets in the plane, it seems natural to find a central vertex in the given planar graph, map it to a central point of the point set, split the graph into, say, three regions that meet in the central vertex, split the point set accordingly, and recurse. In fact, in a Schnyder wood of a 3-connected plane graph $G$, every vertex does have three associated regions, which form a partition of the inner faces of $G$; see the regions $R_r$, $R_g$, and $R_b$ of vertex $v$ in Figure 4c. A vertex is $\delta$-central if each of its regions contains at most $(1-\delta)f_0$ faces of $G$, where $f_0$ is the number of inner faces of $G$. The centrality of the outer vertices is 0. Here we show that every Schnyder wood has a 1/3-central vertex.

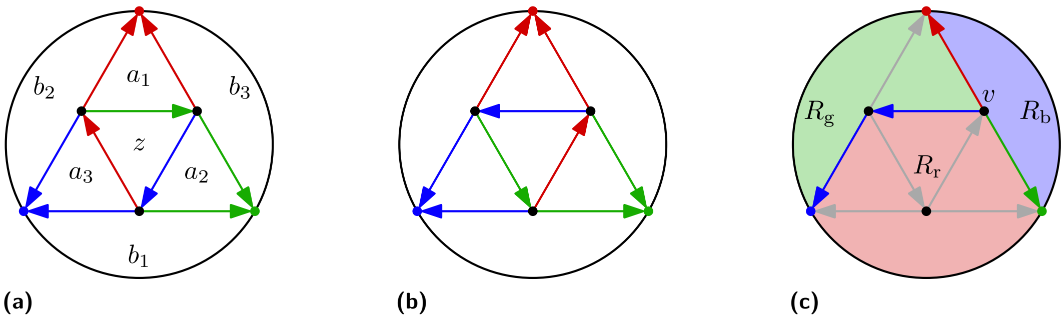


**Figure 4** The two Schnyder woods of the octahedron (a,b) and the three regions of a vertex $v$ (c).

▶ Proposition 1. If $S$ is a Schnyder wood of a plane 3-connected graph $G$ with $f_0$ bounded faces, then there is a vertex $c$ that is 1/3-central with respect to $S$.

The proof, which we omit here, uses Sperner's Lemma (see "Proofs from the book" by Aigner and Ziegler). Next we show that the above centrality bound is tight.

▶ Proposition 2. For every $\varepsilon > 0$, there exists a planar triangulation $T$ with $f_0$ bounded faces such that, in every Schnyder wood of $T$, no vertex has centrality greater than $(1/3+\varepsilon)$.

**Proof.** Consider the skeleton of the octahedron; see Figure 4. We first consider a weighted version of the partition problem where we assign a weight $w_f$ to each face $f$ among the seven inner faces; namely $w(z) = w(b_1) = w(b_2) = w(b_3) = 1$ and $w(a_1) = w(a_2) = w(a_3) = \alpha$ for some real $\alpha > 0$. The total weight of the faces is $3\alpha + 4$. Each inner vertex has a region of weight $2\alpha + 2$. Hence, with increasing $\alpha$, the weight of a region of each vertex can get arbitrarily close to 2/3 of the total weight of the faces of the graph.

We can subdivide the faces of the octahedron and use the weights to count the number of (sub)faces in each original octahedron face. To this end, we place a triangulation with $k$ inner vertices into each of the faces $a_1$, $a_2$, and $a_3$, and set $\alpha = 2k + 5$. In every Schnyder wood of the new triangulation with $3k + 6$ vertices, the edges of the original octahedron are still colored and oriented in one of the two ways shown in Figures 4a and 4b. Let $i \in \{1, 2, 3\}$. If $u$ is a vertex of the triangulation that refines the octahedron face $a_i$, then the largest region of $u$ contains $z$, $b_i$, and all the faces in the triangulations subdividing $a_{i-1}$ and $a_{i+1}$. Hence $u$ has a region whose weight exceeds $2\alpha + 2$. ◀

## References

1 Patrizio Angelini, David Eppstein, Fabrizio Frati, Michael Kaufmann, Silvain Lazard, Tamara Mchedlidze, Monique Teillaud, and Alexander Wolff. Universal point sets for drawing planar graphs with circular arcs. *J. Graph Algorithms Appl.*, 18(3):313–324, 2014. doi:10.7155/jgaa.00324.

2 Fabrizio Frati, Michael Hoffmann, and Csaba D. Tóth. Universal geometric graphs. *Combinatorics, Probability and Computing*, 32(5):742–761, 2023. doi:10.1017/S0963548323000135.

3 Radoslav Fulek and Csaba D. Tóth. Universal point sets for planar three-trees. *J. Discrete Algorithms*, 30:101–112, 2015. doi:10.1016/J.JDA.2014.12.005.

4 Maarten Löffler and Csaba D. Tóth. Linear-size universal point sets for one-bend drawings. In Emilio Di Giacomo and Anna Lubiw, editors, *23rd Int. Symp. Graph Drawing & Network Vis. (GD)*, volume 9411 of *LNCS*, pages 423–429. Springer, 2015. doi:10.1007/978-3-319-27261-0_35.

5 André Schulz. *Lifting planar graphs to realize integral 3-polytopes and topics in pseudo-triangulations*. PhD thesis, Freie Universität Berlin, 2008. doi:10.17169/refubium-16284.

6 Günter M. Ziegler. *Polyhedral Surfaces of High Genus*, pages 191–213. Birkhäuser Basel, 2008. doi:10.1007/978-3-7643-8621-4_10.

## Participants

- Oswin Aichholzer
TU Graz, AT
- Sayan Bandyapadhyay
Portland State University, US
- Benjamin Burton
University of Queensland – Brisbane, AU
- Otfried Cheong
Scalgo, DK
- Mark de Berg
TU Eindhoven, NL
- Jeff Erickson
University of Illinois Urbana-Champaign, US
- Fedor V. Fomin
University of Bergen, NO
- Robert Ganian
TU Wien, AT
- Joachim Gudmundsson
The University of Sidney, AU
- Dan Halperin
Tel Aviv University, IL
- Michael Hoffmann
ETH Zürich, CH
- Kristóf Huszár
TU Graz, AT
- Phillip Keldenich
TU Braunschweig, DE
- David G. Kirkpatrick
University of British Columbia – Vancouver, CA
- Francis Lazarus
CNRS – Grenoble, FR
- Jim Little
University of British Columbia – Vancouver, CA
- Maarten Löffler
Utrecht University, NL
- Gert Meijer
NHL Stenden Hogeschool – Leeuwarden, NL
- Fabrizio Montecchiani
University of Perugia, IT
- Ian Munro
University of Waterloo, CA
- André Nusser
CNRS – Sophia Antipolis, FR
- Eunjin Oh
POSTECH – Pohang, KR
- Yoshio Okamoto
The University of Electro-Communications – Tokyo, JP
- Jeff M. Phillips
University of Utah – Salt Lake City, US
- Valentin Polishchuk
Linköping University, SE
- Raimund Seidel
Universität des Saarlandes – Saarbrücken, DE
- Jonathan Shewchuk
University of California – Berkeley, US
- Diane Souvaine
Tufts University – Medford, US
- Jonathan Spreer
University of Sydney, AU
- Frank Staals
Utrecht University, NL
- Birgit Vogtenhuber
TU Graz, AT
- Hubert Wagner
University of Florida – Gainesville, US
- Alexandra Weinberger
FernUniversität in Hagen, DE
- Carola Wenk
Tulane University – New Orleans, US
- Mathijs Wintraecken
Centre Inria d'Université Côte d'Azur – Sophia Antipolis, FR
- Alexander Wolff
Universität Würzburg, DE
- Yelena Yuditsky
ULB-Brussels, BE
- Ling Zhou
Duke University – Durham, US

Report from Dagstuhl Seminar 25202

# Generative Models for 3D Vision

**Laura Neschen*[1], Bernhard Egger†[2], Adam Kortylewski†[3], William Smith†[4], and Stefanie Wuhrer†[5]**

1 **INRIA Rhône-Alpes, FR.** laura.neschen@inria.fr
2 **Friedrich-Alexander-Universität Erlangen-Nürnberg, DE.** egger.bernhard@gmail.com
3 **Universität Freiburg, DE and MPI für Informatik – Saarbrücken, DE.** akortyle@mpi-inf.mpg.de
4 **University of York, GB.** william.smith@york.ac.uk
5 **INRIA – Grenoble, FR.** stefanie.wuhrer@inria.fr

**Abstract**

Generative models that allow synthesis of realistic 3D models have been of interest in computer vision and graphics for over 2 decades. While traditional methods use morphable models for this task, more recent works have adopted powerful tools from the 2D image domain such as generative adversarial networks, neural fields and diffusion models, and have achieved impressive results. The question of which tools are most suitable for applications such as reconstructing 3D geometry from partial data, and creating digital 3D content remains open. This report documents the program and outcomes of Dagstuhl Seminar 25202 titled "Generative Models for 3D Vision". This meeting of 25 researchers covered a variety of topics such as generative models and priors for 2D tasks, medical applications, and digital representations of humans, including how to evaluate and benchmark different methods. We summarise the discussions, presentations, and results of this seminar.

**Seminar** May 11–16, 2025 – http://www.dagstuhl.de/25202
**2012 ACM Subject Classification** Computing methodologies → 3D imaging; Computing methodologies → Computer vision; Computing methodologies → Computer graphics
**Keywords and phrases** 3D Computer Vision, Computer Graphics, Generative Models, Implicit Representations, Neural Rendering, Statistical Modelling
**Digital Object Identifier** 10.4230/DagRep.15.5.96

## 1 Executive Summary

*Bernhard Egger (Friedrich-Alexander-Universität Erlangen-Nürnberg, DE)*
*Adam Kortylewski (Universität Freiburg, DE and MPI für Informatik – Saarbrücken, DE)*
*Laura Neschen (INRIA Rhône-Alpes, FR)*
*William Smith (University of York, GB)*
*Stefanie Wuhrer (INRIA – Grenoble, FR)*



The rise of purely data-driven generative models, in particular generative adversarial networks, auto-regressive models, neural fields and diffusion models, has led to a step change in image synthesis quality. It is now possible to create photorealistic images with high level semantic control and solve many desirable use cases such as 2D inpainting. Whilst prior models were

* Editorial Assistant / Collector
† Editor / Organizer

object specific (e.g. 3D Morphable Models of Faces), we now have generative models for images and videos that can represent various object classes and generate a huge variety of objects and scenes, even in different styles. The drawback of purely data-driven approaches is that the control and explainability provided by 3D and physically-based parameters is lost. It is also difficult (and perhaps prohibitively inefficient) to learn 3D consistent representations without prior models purely from 2D data alone.

For this seminar, a total of 58 researchers were invited, and 25 of them attended. Participants came from both academia and industry and at varying stages of their careers. Thirteen participants presented their work in around 15-30 minute presentations, and an abstract of each presentation is included in this report. We started the seminar with a short introduction of each participant. Everyone was given one slide to introduce themselves and asked to prepare a question, challenge or goal to discuss during the seminar.

In addition to traditional presentations, multiple slots were left for research discussions with the full group or sub-groups of the participants. The first set of these slots was filled with topics that participants proposed before the start of the seminar. Five participants led research discussions of about 1 hour each about a topic or a problem they considered important. Some of these discussions were led with the full group, while others were discussed in sub-groups, and the resulting conclusions were shared with the full group afterwards. Additionally, two 2 hour discussion slots were initially reserved to be filled with suggestions that came up during the seminar. These two long discussions concerned research questions that were identified as being important for the research community in the course of the seminar, namely the topics of metrics and capture, and hard problems in the research community that merit being studied more. All proposed topics led to lively discussions about vaious problems around generative models. Summaries of the results of these flexible sessions are contained in this report. In addition to these organized discussions, there were numerous informal discussions during both the Wednesday outing and free time slots that are not summarized in this report.

## 2 Table of Contents

**Executive Summary**
*Bernhard Egger, Adam Kortylewski, Laura Neschen, William Smith, and Stefanie Wuhrer* . . . 96

**Overview of Talks**

Synthetic Data for Generative AI
*Thabo Beeler* . . . 100

Learning to Infer Parametric Representations for 3D Plants from Scans
*Samara Ghrer* . . . 100

Statistical Approaches to Internal Anatomy Prediction
*Marilyn Keller* . . . 101

On Learning to Reconstruct Shape using World Priors, Handling Topological Inconsistencies, and Edge Integration as a Model of Choice
*Ron Kimmel* . . . 101

Towards Photorealistic 3D Head Avatars
*Tobias Kirschstein* . . . 103

Understanding Self-Supervised Learning
*Adam Kortylewski* . . . 103

Leveraging 3D Representations for Multi-view Synthesis and Editing with 2D Generative Models
*Or Patashnik* . . . 104

Scaling Digital Humans
*Shunsuke Saito* . . . 104

Clever Data Curation for Smart Supervision
*William Smith* . . . 104

Generating 3D Human Motion with Language
*Gül Varol* . . . 106

3D Foundation Models for Enhanced Geometry in Gaussian Splatting
*Yaniv Wolf* . . . 106

D-Garment: Physics-Conditioned Latent Diffusion for Dynamic Garment Deformations
*Stefanie Wuhrer* . . . 107

**Open problems**

Interpolation of Shapes with Topological and Structural Variability
*Andreea Ardelean and Samara Ghrer* . . . 107

Exploring problems with infinite data
*Timotei Ardelean* . . . 109

Benchmarking Monocular Reconstruction Discussion
*James Gardner* . . . 109

Capturing useful datasets for 3D vision tasks
*Samara Ghrer* . . . 111

Believe Propagation over Spatial Domains with Generative Models
*Jan Eric Lenssen* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 111

Hard problems in computer vision and the necessity of modeling
*Yaniv Wolf* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 112

**Participants** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 113

# 3 Overview of Talks

## 3.1 Synthetic Data for Generative AI

*Thabo Beeler (Google Research – Zürich, CH)*



Inspired by the seminal "Fake It Til You Make It" paper published in 2021, people have successfully trained ML models on synthetic data for research and product alike, demonstrating that the infamous domain gap can be overcome, even though the generated data still falls short in visual quality. Research has produced models trained on synthetic data that not only achieve parity with models trained on real data, but by now even surpass models trained on real data due to the full control (i.e. multiview data), additional modalities (i.e. depth and normals) and pixel perfect annotations (i.e. 3D keypoints). Research such as the seminal "Sapiens" work has further demonstrated success by finetuning foundational models trained on real data with synthetic data, to enable generalization despite limited amounts of synthetic assets.

All of those models are discriminative – they analyze real images to derive higher order semantics from it, such as keypoints, segmentations, or surface normals. On the generative side people have been considerably less successful, due to the remaining domain gap in the data generated, causing domain shifts as visible for example in "Rodin". Recent works, such as "Cafca" or "SynShot", have proposed to overcome this domain shift by fine-tuning on real data, typically following a three-step approach: 1) pretrain a prior on synthetic data, 2) invert sparse views into that prior, 3) employ a fine-tuning strategy to go off model. While this can produce results without domain shift, it's of course not a fully generative model anymore, and as such has its own limitations – for example one cannot unconditionally sample from it. Looking forward we predict the development to trend towards fusing high quality 3D synthetic data with the emerging large scale image or video foundational models to overcome the domain gap for generative 3D models.

## 3.2 Learning to Infer Parametric Representations for 3D Plants from Scans

*Samara Ghrer (University of Grenoble, FR)*


**Joint work of** Samara Ghrer, Christophe Godin, Stefanie Wuhrer
**Main reference** Samara Ghrer, Christophe Godin, Stefanie Wuhrer: "Learning to Infer Parameterized Representations of Plants from 3D Scans", CoRR, Vol. abs/2505.22337, 2025.
**URL** https://doi.org//10.48550/ARXIV.2505.22337

Reconstructing faithfully the 3D architecture of plants from unstructured observations is a challenging task. Plants frequently contain numerous organs, organized in branching systems in more or less complex spatial networks, leading to specific computational issues due to self-occlusion or spatial proximity between organs. Existing works either consider inverse modeling where the aim is to recover the procedural rules that allow to simulate virtual plants, or focus on specific tasks such as segmentation or skeletonization. We propose a unified approach that, given a 3D scan of a plant, allows to infer a parameterized representation of the plant. This representation describes the plant's branching structure, contains parametric information for each plant organ, and can therefore be used directly in a variety of tasks.

## 3.3 Statistical Approaches to Internal Anatomy Prediction

*Marilyn Keller (MPI für Intelligente Systeme – Tübingen, DE)*


**Joint work of** Marilyn Keller, Keenon Werling, Soyong Shin, Scott L. Delp, Sergi Pujades, C. Karen Liu, Michael J. Black, Silvia Zuffi, Vaibhav Arora, Abdelmouttaleb Dakri, Shivam Chandhok, Jürgen Machann, Andreas Fritsche
**Main reference** Marilyn Keller, Keenon Werling, Soyong Shin, Scott L. Delp, Sergi Pujades, C. Karen Liu, Michael J. Black: "From Skin to Skeleton: Towards Biomechanically Accurate 3D Digital Humans", ACM Trans. Graph., Vol. 42(6), pp. 253:1–253:12, 2023.
**URL** https://doi.org//10.1145/3618381
**Main reference** Marilyn Keller, Silvia Zuffi, Michael J. Black, Sergi Pujades: "OSSO: Obtaining Skeletal Shape from Outside", in Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022, pp. 20460–20469, IEEE, 2022.
**URL** https://doi.org//10.1109/CVPR52688.2022.01984
**Main reference** Marilyn Keller, Vaibhav Arora, Abdelmouttaleb Dakri, Shivam Chandhok, Jürgen Machann, Andreas Fritsche, Michael J. Black, Sergi Pujades: "HIT: Estimating Internal Human Implicit Tissues from the Body Surface", in Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024, pp. 3480–3490, IEEE, 2024.
**URL** https://doi.org//10.1109/CVPR52733.2024.00334

Personalized anatomical digital twins are essential for medicine, computer graphics, and biomechanics, but serving internal anatomy usually requires expensive medical imaging. Instead, we can leverage the correlation between external body shape and internal structures to predict the anatomy directly from a subject's appearance. Learning this correlation raises three key challenges: building datasets with paired observations of body shapes and internal anatomy, annotating these datasets, and learning models that capture the correlation between external and internal features. In this talk, I will showcase how we became able to predict skeleton geometry, bone location, and soft tissue distribution solely from external body shape.

## 3.4 On Learning to Reconstruct Shape using World Priors, Handling Topological Inconsistencies, and Edge Integration as a Model of Choice

*Ron Kimmel (Technion – Haifa, IL)*



Deep learning is a disruptive line of research that continues to reshape how computational problems are formulated and solved. While remarkable advances have been made in tasks involving classification, segmentation, and reconstruction, certain fundamental limitations remain – especially when handling geometric data not suited to traditional convolutional structures. This note outlines our ongoing efforts to address shape reconstruction and analysis in such scenarios, drawing inspiration from differential geometry, spectral theory, and human vision.

Deep learning architectures are designed around the notion of shift-invariance and low-dimensional latent structures. These assumptions empower convolutional neural networks (CNNs) to generalize well across a range of problems. However, for geometric structures where no natural shift-invariant coordinate system exists, standard methods often fall short. In our group, we have been exploring new avenues that combine classical geometry with modern machine learning to reconstruct and analyze shapes – even when those shapes are partially occluded, topologically noisy, or nonrigid. These include tools for shape matching, geodesic measurement, and the design of semi-supervised learning techniques rooted in axiomatic geometric principles.

One of geometry's grand challenges is explaining the world we live in. Shape comparison, particularly of nonrigid objects, exemplifies this. Unlike rigid objects, nonrigid shapes lack a universal parameterization, complicating comparison. Over the past two decades, we have developed methods for analyzing and matching such shapes using the Gromov Hausdorff (GH) distance, which quantifies discrepancies between metric spaces. We have advanced this from theoretical abstraction to practical approximation using spectral embeddings.

Traditionally, computer vision (CV) aimed to extract geometry from images, while computer graphics (CG) sought to generate images from geometry. Today, these domains converge. A fundamental question we address is: Can neural networks help discover invariants or derive parameterizations of geometric data? To this end, we showed that the Gromov distance between metric spaces offers a powerful framework for non-rigid shape analysis.

Classical computer vision focused on "shape from X" problems: shape from shading (often solved using eikonal solvers), shape from stereo (requiring correspondence resolution), and shape from auto-stereograms. A compelling example comes from stereoscopic vision. Consider looking through a cardboard tube with one eye, while the other eye is blocked by one hand – this creates a visual illusion of a "hole" in the hand, as the brain attempts to force a false correspondence between views. Such hallucination phenomena motivated the study of auto-stereograms generation, see [2], and geometric reconstruction from such data using methods described in [1]. It motivates the concept of prior based geometry reconstruction as recently adopted and adapted in [3]. It demonstrates how prior knowledge of image formation models captured by neural network trained to solve the shape from stereo problem, can significantly improve reconstruction accuracy integrating Gaussian splatting with synthetic stereo pairs.

Recently, our attention has turned to shape invariants, especially under partial observation. We introduced the *Wormhole Loss* [4], designed for robust partial matching by emphasizing topological consistency. This follows our earlier work on partial shape correspondence using functional maps [5], which outperforms previous efforts such as [6, 7] by incorporating both local and global topology sensitive invariants into the learning objective.

Finally, another classical idea that should be re-examined through the lens of deep learning is edge detection. The variational model for optimal edge integration via regularized Laplacian zero-crossings [8] raises an intriguing question: can such elegant mathematical formulations be embedded into a neural network framework? If so, which applications would benefit most, and how could these formulations serve as loss functions or regularizers in modern computer vision applications?

### References


1 Ron Kimmel. 3D Shape Reconstruction from Autostereograms and Stereo. *Journal of Visual Communication and Image Representation*, 13:324-333, 2002.

2 A. M. Bruckstein, R. Onn, and T. J. Richardson. Improving the vision of magic eyes: A guide to better autostereograms. In *Advances in Image Understanding, A Festschrift for Azriel Rosenfeld*, IEEE Comput. Sci., 1996.

3 Y. Wolf, A. Bracha, and R. Kimmel. Gs2mesh: Surface reconstruction from Gaussian splatting via novel stereo views. *European Conference on Computer Vision*, 2024.

4 A. Bracha, T. Dagès, and R. Kimmel, "Wormhole Loss for Partial Shape Matching," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [Online]. Available: `https://arxiv.org/abs/2410.22899`

5 A. Bracha, T. Dagès, and R. Kimmel, "On Partial Shape Correspondence and Functional Maps," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023. [Online]. Available: `https://arxiv.org/abs/2310.14692`

6 G. Rosman, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Nonlinear dimensionality reduction by topologically constrained isometric embedding. *International Journal of Computer Vision*, 2010.
7 A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Topology-invariant similarity of nonrigid shapes. *International Journal of Computer Vision*, 2009.
8 R. Kimmel and A. M. Bruckstein. Regularized Laplacian zero crossings as optimal edge integrators. *International Journal of Computer Vision*, 2003.

## 3.5 Towards Photorealistic 3D Head Avatars

*Tobias Kirschstein (TU München – Garching, DE)*



In recent years, 3D head avatars have become ever more realistic, leading to their integration into first commercial products. In this talk, we will explore some of the key developments that drive research in this area, including datasets, methods, and efforts to standardize evaluation. Studio-level multi-view video datasets such as the NeRSemble dataset enable creating high quality animatable reconstructions of human heads. For example, the current state-of-the-art method, Neural Parametric Gaussian Avatars, combines multi-view supervision with a detailed 3D face model to produce lifelike digital doubles. In the future, the field is moving toward learning larger generative 3D priors which could help tackle under-constrained scenarios such as reconstructing 3D head avatars from single images. Ultimately, we will have to answer the question whether the future lies in recording more data in controlled studio environments or exploring new paradigms, like 3D GANs, that can infer 3D priors from 2D data alone.

## 3.6 Understanding Self-Supervised Learning

*Adam Kortylewski (Universität Freiburg, DE and MPI für Informatik – Saarbrücken, DE)*



SSL methods gave us general-purpose visual features like DinoV2 that perform well as backbone for a variety of tasks, like classification, segmentation and even low-level 3D tasks like depth or normal estimation. Interestingly, these features emerge despite being trained with pre-text objectives that are rather unrelated to actual (3D) vision tasks, like making image crops of the same image similar to each other. In this talk, I will discuss some initial ideas about the connection between SSL of visual features and visual correspondence tasks. I will discuss experiments that show how current SSL features are not just good at 2D vision but also 3D vision tasks, and how we could possibly develop pre-text objectives that explicitly optimize semantic correspondence capabilities of visual features, ultimately leading to the emergence of more advanced visual features.

## 3.7 Leveraging 3D Representations for Multi-view Synthesis and Editing with 2D Generative Models

*Or Patashnik (Tel Aviv University, IL)*



While 2D generative models have achieved remarkable success in image synthesis and editing, extending their power to multi-view and 3D tasks often leads to inconsistencies and limited control. In this talk, I will discuss a line of work showing how incorporating 3D representations, ranging from explicit 3D shapes to emergent 3D feature fields, can significantly enhance the capabilities of 2D generative models for multi-view generation and editing. I will first demonstrate how explicit geometric structures, such as rough meshes or low-fidelity 3D models, can guide and stabilize the generation process, enabling better control and consistency. I will then discuss how, even without an explicit shape, constructing a 3D-aware feature field during generation can substantially improve multi-view coherence. Together, these approaches show that introducing 3D representations, whether provided upfront or learned during the process, fundamentally amplifies the capabilities of 2D diffusion models and enables high-quality, controllable 3D content creation.

## 3.8 Scaling Digital Humans

*Shunsuke Saito (Codec Avatars Lab – Pittsburgh, US)*



In this talk, we discuss how to put the modeling of 3D Digital Humans on the table of large-scale training. In particular, we share our recent experiments to see the potential benefit and risk of pre/post-training regime. Self-supervised learning with large scale data and then finetuning on small but clean data leads to generalization and high-quality at scale. This shows surprising generalization to the domain outside the post-training data distribution. We also discuss the importance of 3D in the era of powerful 2D generative models. By presenting our recent work on efficient 3D Gaussian avatar decoding, we illustrate how efficiency plays a critical role in content creation with 3D-based approaches.

## 3.9 Clever Data Curation for Smart Supervision

*William Smith (University of York, GB)*



This talk summarised work from the past 10 years that opportunistically repurposed data collected for one purpose to provide supervision for a 3D vision task. The purpose of the talk was to: 1. highlight that there are likely many more datasets available that have not yet been tapped for such purposes, 2. to try to categorise previous methods and 3. to stimulate discussion around ideas for problems or datasets that could benefit from or be used in this way. The proposed initial categorisation along with representative methods was:

- **Multiview curation at training time for single view task supervision** [1, 2, 3, 4]
- **Two view curation at training time for two view task supervision** [5, 6, 7]
- **Compositing** [8]
- **Wide FOV to supervise beyond cropped narrow FOV** [9]
- **Diffusion-based editing** [10]

In the discussion that followed, many sources of interesting supervision were proposed. Aging or fitness journey timelapse videos can be used to understand changes in appearance with a fixed identity. Zipline or rollercoaster videos provide smooth fly throughs that are repeated in many different conditions with almost the same camera trajectory. Ice bucket challenge videos provide examples of the same clothes dry and wet. Flash photography, for example at red carpet events, provides photometric stereo cues. Google streetview provides many lighting conditions for the same scene. Object press videos show how different materials deform. Slow TV such as railway journeys provide revisits of slowly changing landscapes.

## References


1 Zhengqi Li and Noah Snavely. *MegaDepth: Learning Single-View Depth Prediction from Internet Photos*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2041-2050, 2018.

2 Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T. Freeman. *Learning the Depths of Moving People by Watching Frozen People*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4521-4530, 2019.

3 Zhengqi Li and Noah Snavely. *Learning Intrinsic Image Decomposition from Watching the World*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9039-9048, 2018.

4 Ye Yu and William A. P. Smith. *InverseRenderNet: Learning single image inverse rendering*. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3155-3164, 2019.

5 Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. *Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow*. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023.

6 Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. *DUSt3R: Geometric 3D Vision Made Easy*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20697-20709, 2024.

7 Linyi Jin, Richard Tucker, Zhengqi Li, David Fouhey, Noah Snavely, and Aleksander Holynski. *Stereo4D: Learning How Things Move in 3D from Internet Stereo Videos*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

8 Finlay G. C. Hudson and William A. P. Smith. *Track Anything Behind Everything: Zero-Shot Amodal Video Object Segmentation*. arXiv preprint arXiv:2411.19210, 2024.

9 Rundong Luo, Matthew Wallingford, Ali Farhadi, Noah Snavely, and Wei-Chiu Ma. *Beyond the Frame: Generating 360° Panoramic Videos from Perspective Videos*. arXiv preprint arXiv:2504.07940, 2025.

10 Alex Trevithick, Ruihan Yang, Rundi Wu, Ruiqi Gao, Ben Poole, Aleksander Holynski, Jonathan T. Barron, Noah Snavely, and Angjoo Kanazawa. *SimVS: Simulating World Inconsistencies for Robust View Synthesis*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

## 3.10 Generating 3D Human Motion with Language

*Gül Varol (ENPC – Marne-la-Vallée, FR)*



Text-driven 3D human motion generation methods have recently evolved to synthesize language-controllable avatars. My talk briefly shows results of motion generation with various levels of granularity and presents recent extensions to motion editing and hand motion generation. An obvious challenge in this domain is the scarcity of data. We will discuss ways to scale up the data to train these models, with a particular focus on noise vs scale trade-off.

## 3.11 3D Foundation Models for Enhanced Geometry in Gaussian Splatting

*Yaniv Wolf (Technion – Haifa, IL)*


**Joint work of** Yaniv Wolf, Amit Bracha, Ron Kimmel
**Main reference** Yaniv Wolf, Amit Bracha, Ron Kimmel: "GS2Mesh: Surface Reconstruction from Gaussian Splatting via Novel Stereo Views", in Proc. of the Computer Vision – ECCV 2024 – 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXXIX, Lecture Notes in Computer Science, Vol. 15147, pp. 207–224, Springer, 2024.
**URL** https://doi.org//10.1007/978-3-031-73024-5_13

In recent years, Gaussian Splatting (GS) has emerged as an efficient and accurate method for scene representation. However, while GS excels in novel view synthesis, it struggles with accurately representing geometry, as appearance and geometry are often conflicting objectives for the Gaussian representation itself (consider a colorful flat wall, or a textureless round surface). Attempts to simultaneously optimize both appearance and geometry often lead to suboptimal compromises. To bridge the gap between appearance and geometry, we introduce GS2Mesh, a novel pipeline that extracts accurate surface geometry from any optimized GS scene without degrading its visual quality. Rather than directly optimizing the Gaussians for both appearance and geometry, we leave the Gaussians optimized for appearance, and leverage a robust 3D data-driven prior, in the form of a pre-trained stereo matching foundation model, for the geometry extraction. Specifically, we render stereo-aligned image pairs from the optimized GS scene and feed them to the pre-trained stereo model, resulting in high-quality disparity maps. These maps are then scaled according to the camera parameters to form multi-view consistent depth maps, which are fused together to reconstruct accurate 3D surfaces. Our approach achieves state-of-the-art performance on popular 3D reconstruction benchmarks, as well as robust performance on in-the-wild scenes captured via smartphones. Due to the modularity of the approach, performance is expected to improve as better GS and stereo matching models emerge.

## 3.12 D-Garment: Physics-Conditioned Latent Diffusion for Dynamic Garment Deformations

*Stefanie Wuhrer (INRIA – Grenoble, FR)*


**Joint work of** Antoine Dumoulin, Adnane Boukhayma, Laurence Boissieux, Bharath Bhushan Damodaran, Pierre Hellier, Stefanie Wuhrer
**Main reference** Antoine Dumoulin, Adnane Boukhayma, Laurence Boissieux, Bharath Bhushan Damodaran, Pierre Hellier, Stefanie Wuhrer: "D-Garment: Physics-Conditioned Latent Diffusion for Dynamic Garment Deformations", CoRR, Vol. abs/2504.03468, 2025.
**URL** https://doi.org//10.48550/ARXIV.2504.03468

Adjusting and deforming 3D garments to body shapes, body motion, and cloth material is an important problem in virtual and augmented reality. Applications are numerous, ranging from virtual change rooms to the entertainment and gaming industry. This problem is challenging as garment dynamics influence geometric details such as wrinkling patterns, which depend on physical input including the wearer's body shape and motion, as well as cloth material features. Existing work studies learning-based modeling techniques to generate garment deformations from example data, and physics-inspired simulators to generate realistic garment dynamics. The presentation covered our recent learning-based approach trained on data generated with a physics-based simulator. Compared to prior work, our 3D generative model learns garment deformations for loose cloth geometry, especially for large deformations and dynamic wrinkles driven by body motion and cloth material. Furthermore, the model can be efficiently fitted to observations captured using vision sensors. We propose to leverage the capability of diffusion models to learn fine-scale detail: we model the 3D garment in a 2D parameter space, and learn a latent diffusion model using this representation independent from the mesh resolution. This allows to condition global and local geometric information with body and material information. The presentation concluded with some open questions that benefitted from discussions in the diverse group of the seminar.

# 4 Open problems

## 4.1 Interpolation of Shapes with Topological and Structural Variability

*Andreea Ardelean (Friedrich-Alexander-Universität Erlangen-Nürnberg, DE) and Samara Ghrer (University of Grenoble, FR)*



This open discussion centers on the problem of representing and understanding the structural topology of growing plants, raising foundational questions about how plant morphology evolves over time and how it can be modeled computationally. Participants explored whether plant topology – particularly the branching structure – changes during growth, and if so, how to model this dynamic nature in a consistent and analyzable way. Analogies to human anatomy (e.g., variability in tooth roots or spinal structure) highlighted similar topological ambiguity in biological systems.

A range of modeling strategies were discussed. Neural implicit representations such as NeuralSDF offer a powerful tool for interpolating between shapes, after which a skeleton can be extracted. However, challenges arise when interpolations generate non-physical or biologically invalid intermediate states – for example, a plant structure with a fractional number of branches, or a tooth with an unrealistic root configuration. In this context, skeleton extraction and branching modeling must be grounded in biological plausibility.

Stochastic modeling was proposed as an effective alternative to rule-based procedural methods. While procedural modeling requires the system to strictly follow hand-crafted rules, stochastic models can learn from data, making them more flexible for capturing the natural variability in plant structures. Parameters of branching processes (e.g., angles, lengths, probabilities of bifurcation) can be learned directly from real plant data, enabling more realistic simulations.

There was a shift in perspective proposed: rather than thinking in purely topological terms, plant structures may be better understood as graph-like entities, where branches and junctions form a structured hierarchy. This aligns better with biological reality and allows for more intuitive comparisons between different plant forms.

To compare plant structures across stages of growth or between individuals, several mathematical tools were proposed. Diffusion distances were suggested as a more robust alternative to geodesic distances, offering smoother and more meaningful metrics for shape comparison. In cases where the structure contains holes or discontinuities, Fourier-based approaches might offer better representations. Still, the core challenge remains: how to define meaningful correspondences between topologically diverse structures, especially when interpolations between them don't represent real physical states (e.g., morphing between different glyph forms of the letter "g").

From an application-driven standpoint, several motivations for understanding plant topology were discussed. These include:

- Monitoring and diagnosing plant health (e.g., detecting disease via structural anomalies).
- Understanding growth patterns of individual plants versus species-level trends (nature vs. nurture).
- Predicting future states or motions of plants based on current observations.

The discussion emphasized the importance of discrete vs. continuous modeling, particularly in cases like the transition from flower to fruit, which represents a discrete structural change atop a generally continuous growth process. Mixture models or hybrid representations (continuous shapes with discrete categorical states) were considered useful for capturing this dual nature.

Drawing parallels to 3D morphable models (3DMMs) used in face reconstruction, the conversation highlighted the trade-off between model complexity and usability. While complex models may better capture biological variance, they are harder to fit and interpret. 3DMMs benefit from disentanglement and controllability, features that plant models currently lack due to the high diversity and complexity of plant morphologies.

Throughout the discussion, there was a recurring theme: the purpose of the model should drive its design. For instance, determining a person's eye color is more efficiently done through observation than genetic analysis – a metaphor for focusing modeling efforts on efficient, goal-aligned representations. Similarly, there's a risk in over-parameterizing models, leading to poor generalization, or under-parameterizing, resulting in loss of essential structure.

Ultimately, effective plant topology modeling must balance generative and discriminative approaches, provide biologically grounded constraints, and respect the discrete-continuous duality of natural structures. The goal is not just to simulate growth but to create meaningful, applicable models that align with real-world use cases such as agriculture, ecology, and evolutionary biology.

## 4.2 Exploring problems with infinite data

*Timotei Ardelean (Friedrich-Alexander-Universität Erlangen-Nürnberg, DE)*



In order to explore the interplay between data, compute, and algorithms, it was proposed to analyze our current methods and training methodologies on constrained, well-defined tasks where perfect training data can be easily generated. By generating synthetic data without restrictions, the effect of scaling w.r.t to data and training time could more easily be disentangled, and algorithmic limitations could be identified. The discussions produced several suggestions of Computer Vision problems where shortcomings of the current models could be investigated, including maze solving, the n-queens puzzle, and Sudoku (as inspired by the previous talk of this seminar: Believe Propagation over Spatial Domains with Generative Models). A very simple example we considered is Conway's Game of Life, which we analyzed with a quick hands-on experiment, by training a model to predict the update rule of the game. The model quickly learned the rules of the game using synthetic randomly generated boards, supervised by the true update rule. We also tested for generalization by excluding certain patterns from training. For the simple cases tested, the model was able to generalize to these unseen patterns, suggesting the network learns more than a simple look-up table operation.

## 4.3 Benchmarking Monocular Reconstruction Discussion

*James Gardner (University of York, GB)*



### The Limitations of Current Benchmarking

It was noted that existing metrics, especially for generative tasks like single-view 3D reconstruction, are fundamentally flawed. An analogy was drawn to image compression, where metrics such as PSNR failed to capture perceptual artifacts, this led to the development of new perceptual metrics, which themselves became targets for optimization, not always resulting in genuinely better qualitative outcomes. A similar pattern is evident in LLM evaluation, where comprehensive benchmark suites are still critiqued for fostering optimization towards specific metrics over broader user satisfaction. This phenomenon, often described by Goodhart's Law – where a metric ceases to be effective once it becomes a target – is a significant point of concern. Consequently, reviewers often find it necessary to look beyond numerical scores, emphasizing qualitative aspects and the soundness of the proposed methods.

### Proposed Directions for Better Benchmarks

Several approaches for enhancing benchmarks were put forward:
Focus on Failure Modes and Use Cases: It was proposed that new metrics should not be designed merely to show a method's superiority, but should instead arise from an understanding of how and why current methods fail. If a particular artifact, such as object interpenetration, is identified as problematic, metrics should be formulated to quantify it. This, however, introduces complexities, including the definition of object separability (e.g.,

determining if cushions are distinct from a sofa). The specific application or use case was highlighted as critically important, metrics must align with the intended outcome, whether for applications like 3D printing (where plausibility might be prioritized over exact scale) or interactive scene manipulation (where object decomposition is vital).

Plausibility and Perceptual Evaluation: For ill-posed problems that involve substantial hallucination or completion, the plausibility of the generated 3D geometry or novel views is of paramount importance. Proposals included:

- More extensive user studies.
- The leveraging of other sophisticated models, such as Video Models or Vision-Language Models (VLMs), to evaluate the consistency and "common sense" of reconstructed scenes (e.g., how objects might interact or react to physical phenomena like an earthquake). The underlying principle is that if novel views are plausible and internally consistent, the geometry might be deemed adequate for many applications, even if precise geometric accuracy remains elusive. Embracing Ambiguity and Distributions: Given that many 3D vision tasks are ill-posed and possess multiple valid solutions (e.g., completing an occluded part of an object), it was argued that models should predict a distribution of potential solutions. Evaluation methodologies would then need to assess the likelihood of the ground truth under this predicted distribution. This approach necessitates considerably larger and more varied test datasets to adequately probe ambiguous scenarios.

Identifying the "Next Frontier": It was stated that effective benchmarks should delineate the "next significant problem." When current methodologies begin to saturate existing benchmarks, this is considered a signal to identify new, more demanding tasks that will cause the performance of existing state-of-the-art methods to decline, thereby stimulating further advancements.

### Contentious Metrics and Community Influence

The discussion also briefly addressed alternative, non-traditional indicators of a method's utility, such as the number of GitHub stars. It was argued by some that this reflects real-world adoption, the quality of the accompanying code, and overall impact. However, strong counter arguments were presented, noting that such indicators lack scientific rigor, can be disproportionately influenced by "hype," may not accurately reflect true scientific merit or superior performance (particularly for less mainstream research areas), and could disadvantage robust methods that have less polished codebases or are not aligned with current popular trends.

In summary, there was an agreement that current benchmarking practices within 3D vision are insufficient. The proposed way forward involves the development of metrics that are more intimately aligned with specific use cases, effectively capture perceptual plausibility, appropriately account for the inherent ambiguity present in many 3D tasks, and are derived from an understanding of current model limitations rather than focusing on simple numerical rankings. The intrinsic difficulty is in the creation of metrics that are simultaneously robust, meaningful, and resistant to superficial optimization.

## 4.4 Capturing useful datasets for 3D vision tasks

*Samara Ghrer (University of Grenoble, FR)*



Dataset propositions and ideas to be collaboratively captured:

- 4D human motion dataset with a large number of markers to capture motions that cannot be captured by cameras, or by other motion capture datasets with lower number of markers that lead to smoothed motions.
- Large volumetric dataset of humans in motion with multi-people, self and object interactions.
- Multi view face dataset in both indoor and outdoor lighting conditions with marker alignment (invisible markers on the back of the head, or using sunscreen).
- MRI + volumetric dynamic or multi-pose captures.
- Large and diverse dataset that is captured across different labs with an agreed setting.
- DagMark: A 3D vision (different tasks) benchmarking challenge where people can submit datasets, evaluation schemes and metrics. Once submitted, the benchmarks are automatically compared to existing ones and added to the collection. This challenge can be hosted on a website for worldwide contributions.

Raised questions:

- Is it possible to simultaneously capture MRI scans and real-life images using a combination of MRI scanners and cameras?
- Plausibility vs accuracy: Can we design a dataset that supports both ends?

## 4.5 Believe Propagation over Spatial Domains with Generative Models

*Jan Eric Lenssen (MPI für Informatik – Saarbrücken, DE)*


**Joint work of** Christopher Wewer, Bartlomiej Pogodzinski, Bernt Schiele, Jan Eric Lenssen
**Main reference** Christopher Wewer, Bartlomiej Pogodzinski, Bernt Schiele, Jan Eric Lenssen: "Spatial Reasoning with Denoising Models", in Proc. of the International Conference on Machine Learning (ICML), 2025.
**URL** https://geometric-rl.mpi-inf.mpg.de/srm/

Video and Multi-View Generation Models are taking over (3D) Computer Vision. We are expecting them to model extremely complex, high-dimensional distributions. However, in many cases they spectacularly fail and collapse to hallucinations. I'd like to pose the question: is scaling up data and compute the best or only solution to solve this problem? And is it a good idea to keep up with exponential growth of distribution complexity? Bayesian inference came up with intelligent ways to simplify distribution complexity in the past. It might be a good time to explore similar ideas for continuous generative models in spatial domains. In this talk, I would introduce these discussion points and present some preliminary results we obtained in recent investigations.

## 4.6 Hard problems in computer vision and the necessity of modeling

*Yaniv Wolf (Technion – Haifa, IL)*



The group discussed what defines a "hard problem" in computer vision and how models should be designed to solve them. While some problems benefit from large-scale "big hammer" solutions, others require specialized tools. In that context, neural networks were recognized as convenient, but also problematic in post-data regimes due to heavy data dependence, limited explainability, and robustness issues.

The discussion focused on two main issues:

- The first issue discussed was whether a "hard problem" is defined by the amount of data required to train a model to solve it, and whether any problem can be solved without the need for generalization if enough data is available. Several examples were presented, such as training NeRFs with synthetic/real data, learning geometric consistency in video models, understanding grammar in LLMs, and solving complex puzzles such as 3D mazes and Sudoku.
- The second issue discussed was the impact of explainability in models. Some participants argued that explainability is essential for humans, as the designers of these models, to understand what they are designing. Others believed that introducing parametric models might negatively affect model performance. An example discussed was self-driving cars, and whether they should adhere strictly to theoretical traffic laws or learn from actual human driving behavior.

Finally, caution was expressed regarding overly complex solutions ("big hammers") that risk memorization instead of genuine generalization.

## Participants

- Andreea Ardelean
Friedrich-Alexander-Universität Erlangen-Nürnberg, DE
- Timotei Ardelean
Friedrich-Alexander-Universität Erlangen-Nürnberg, DE
- Thabo Beeler
Google Research – Zürich, CH
- Timo Bolkart
Google Research – Zürich, CH
- Neill Campbell
University of Bath, GB
- Rishabh Dabral
MPI für Informatik – Saarbrücken, DE
- Olaf Dünkel
MPI für Informatik – Saarbrücken, DE
- Bernhard Egger
Friedrich-Alexander-Universität Erlangen-Nürnberg, DE
- James Gardner
University of York, GB
- Samara Ghrer
University of Grenoble, FR
- Marilyn Keller
MPI für Intelligente Systeme – Tübingen, DE
- Ron Kimmel
Technion – Haifa, IL
- Tobias Kirschstein
TU München – Garching, DE
- Adam Kortylewski
Universität Freiburg, DE & MPI für Informatik – Saarbrücken, DE
- Jan Eric Lenssen
MPI für Informatik – Saarbrücken, DE
- Ruoshi Liu
Columbia University – New York, US
- Laura Neschen
INRIA Rhône-Alpes, FR
- Or Patashnik
Tel Aviv University, IL
- Ryan Po
Stanford University, US
- Shunsuke Saito
Codec Avatars Lab – Pittsburgh, US
- William Smith
University of York, GB
- Christian Theobalt
MPI für Informatik – Saarbrücken, DE
- Gül Varol
ENPC – Marne-la-Vallée, FR
- Yaniv Wolf
Technion – Haifa, IL
- Stefanie Wuhrer
INRIA – Grenoble, FR

Report from Dagstuhl Seminar 25211

# The Constraint Satisfaction Problem: Complexity and Approximability

**Manuel Bodirsky*[1], Venkatesan Guruswami*[2], Dániel Marx*[3], Stanislav Živný*[4], and Žaneta Semanišinová†[5]**

1 **TU Dresden, DE.** manuel.bodirsky@tu-dresden.de
2 **University of California – Berkeley, US.** venkatg@berkeley.edu
3 **CISPA – Saarbrücken, DE.** marx@cispa.de
4 **University of Oxford, GB.** standa.zivny@cs.ox.ac.uk
5 **TU Dresden, DE.** zaneta.semanisinova@tuwien.ac.at

**Abstract**

Constraint satisfaction has always played a central role in computational complexity theory; appropriate versions of CSPs are classical complete problems for most standard complexity classes. CSPs constitute a very rich and yet sufficiently manageable class of problems to give a good perspective on general computational phenomena. For instance, they help to understand which mathematical properties make a computational problem tractable (in a wide sense, e.g., polynomial-time solvable, non-trivially approximable, etc.). One of the most striking features of this research direction is the variety of different branches of mathematics (including algebra and logic, combinatorics and graph theory, probability theory and mathematical programming, and most recently topology) that are used to achieve deep insights in the study of the CSP, and this seminar will contribute towards further synergy in the area. In the last 20 years, research activity in this area has significantly intensified and hugely impressive progress was made. The Dagstuhl Seminar 25211 "The Constraint Satisfaction Problem: Complexity and Approximability" was aimed at bringing together researchers using all the different techniques in the study of the CSP so that they can share their insights obtained during the past four years. This report documents the material presented during the course of the seminar.

**Seminar** May 18–23, 2025 – https://www.dagstuhl.de/25211

**2012 ACM Subject Classification** Theory of computation → Problems, reductions and completeness

**Keywords and phrases** computational complexity, constraint satisfaction problem, hardness of approximation, parameterized complexity, semidefinite programming

**Digital Object Identifier** 10.4230/DagRep.15.5.114

## 1 Executive summary

*Manuel Bodirsky (TU Dresden, DE)*
*Venkatesan Guruswami (University of California – Berkeley, US)*
*Dániel Marx (CISPA – Saarbrücken, DE)*
*Stanislav Živný (University of Oxford, GB)*

The *constraint satisfaction problem*, or CSP in short, provides a unifying framework in which it is possible to express, in a natural way, a wide variety of computational problems dealing with mappings and assignments, including satisfiability, graph colorability, and systems of equations. The CSP framework originated over 40 years ago independently in artificial

* Editor / Organizer
† Editorial Assistant / Collector



The Constraint Satisfaction Problem: Complexity and Approximability, *Dagstuhl Reports*, Vol. 15, Issue 5, pp. 114–133
Editors: Manuel Bodirsky, Venkatesan Guruswami, Dániel Marx, Stanislav Živný, and Žaneta Semanišinová
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

intelligence, database theory, and graph theory, under three different guises, and it was realised only in the late 1990s that these are in fact different faces of the same fundamental problem. Nowadays, the CSP is extensively used in theoretical computer science, being a mathematical object with very rich structure that provides an excellent laboratory both for classification methods and for algorithmic techniques, while in AI and more applied areas of computer science this framework is widely regarded as a versatile and efficient way of modelling and solving a variety of real-world problems, such as planning and scheduling, software verification and natural language comprehension, to name just a few. An instance of CSP consists of a set of variables, a set of values for the variables, and a set of constraints that restrict the combinations of values that certain subsets of variables may take. Given such an instance, the possible questions include (a) deciding whether there is an assignment of values to the variables so that every constraint is satisfied, or optimising such assignments in various ways; e.g., (b) finding an assignment satisfying as many constraints as possible, or (c) finding an assignment that violates as few constraints as possible. There are many important modifications and extensions of this basic framework, e.g., those that deal with counting assignments or involve soft or global constraints.

Constraint satisfaction has always played a central role in computational complexity theory; appropriate versions of CSPs are classical complete problems for most standard complexity classes. CSPs constitute a very rich and yet sufficiently manageable class of problems to give a good perspective on general computational phenomena. For instance, they help to understand which mathematical properties make a computational problem tractable (in a wide sense, e.g. polynomial-time solvable or non-trivially approximable, fixed-parameter tractable or definable in a weak logic). One of the most striking features of this research direction is the variety of different branches of mathematics (including universal algebra and logic, combinatorics and graph theory, probability theory and mathematical programming, and most recently topology) that are used to achieve deep insights in the study of the CSP, and the proposed seminar will contribute towards further synergy in the area.

After about 20 years of intense research activity and hugely impressive progress, the culmination of the algebraic-approach to fixed-template CSPs was the resolution of the Feder-Vardi conjecture independently by Bulatov and Zhuk in 2017. While some fundamental questions (such as a fine-grained understanding of tractable CSPs) remain open, new research directions on generalizations of CSPs started emerging. One such direction is the fixed-template Min-CSP studied with respect to fixed-parameter tractability. Another is the fixed-template promise CSP (PCSP), which is motivated by the goal to better understand the approximability of problems with perfect completeness. PCSPs are a vast generalization of CSPs. A prime and well-known example is the approximate graph coloring problem: distinguish $k$-colorable graphs from graphs that are not even $c$-colorable, for some fixed $c \geq k$. The main topic of this Dagstuhl Seminar is PCSPs, a highly ambitious research directions with intriguing connections to both old open problems (such the approximate graph coloring problem) and new research directions (such as new algorithms for CSPs) and its connections to gems of computational complexity (such as Håstad's optimal inapproximability results) and fixed-parameter tractability of Min-CSPs, which has recently seen some remarkable progress.

The recent flurry of activity on the topic of the seminar is witnessed by six previous Dagstuhl Seminars, titled "Complexity of constraints" (06401) and "The CSP: complexity and approximability" (09441, 12541, 15301, 18231, 22201), that were held in 2006, 2009, 2012, 2015, 2018, and 2022, respectively. This seminar was a follow-up to the 2009, 2012, 2015, 2018, and 2022 seminars. Indeed, the exchange of ideas at the 2009, 2012, 2015, 2018, and 2022 seminars has led to ambitious new research projects and to establishing regular communication channels. There is clearly the potential for further systematic interaction that will keep on cross-fertilising the areas and opening new research directions. The 2025 seminar brought together 47 researchers from different highly advanced areas of constraint satisfaction and involved many specialists who use universal-algebraic, combinatorial, geometric, and probabilistic techniques to study CSP-related algorithmic problems. The participants presented, in 24 talks and 2 tutorial series, their recent results on a number of important questions concerning the topic of the seminar. One particular feature of this seminar is a significant increase in the number of talks involving multiple subareas and approaches within its research direction – a definite sign of the growing synergy, which is one of the main goals of this series of seminars.

**Concluding remarks and future plans.** The seminar was well received as witnessed by the very high rate of accepted invitations (many more researchers would have liked to participate, but we were unable to accommodate them) and the great degree of involvement by the participants. Because of a multitude of impressive results reported during the seminar and active discussions between researchers with different expertise areas, the organisers regard this seminar as a great success. With steadily increasing interactions between such researchers, we foresee another seminar focusing on the interplay between different approaches to studying the complexity and approximability of the CSP. Finally, the organisers wish to express their gratitude to the Scientific Directors of the Dagstuhl Centre for their support of the seminar.

## Description of the topics of the seminar

Following the two independent proofs of the finite-domain CSP dichotomy in 2017 by Bulatov and Zhuk, the research in the area of constraint satisfaction problems focuses on various generalizations of the CSP framework and on uniformizing the approaches to these variants of CSPs. We list the concrete topics of the talks in the seminar below.

### Uniform algorithms for finite-domain CSPs

On the topic of understanding better the polynomial-time algorithms for finite-domain CSPs with the hope of getting a uniform algorithm, Dmitriy Zhuk gave an overview talk. He discussed the algorithms for finite-domain CSPs arising from the combination of consistency checking with basic linear programming and affine integer programming, and their singleton variants. He also provided polymorphism conditions for these algorithms to be correct, some of which were not known before.

### Infinite-domain CSPs

One of the most natural and intensely studied generalizations of finite-domain CSPs are CSPs with infinite-domain $\omega$-categorical templates. The workshop contained a 2-hour tutorial on such CSPs given by Antoine Mottet.

- In the first part, Mottet presented a template-free approach, showing the similarity between the finite- and infinite-domain CSPs.
- In the second part, he outlined more recent results in the area, in particular the theory of smooth approximations and the connections of infinite-domain CSPs to promise CSPs.

There were 5 contributed talks on recent results in the area of infinite-domain CSPs and related topics.

- Jakub Rydval talked about fundamental questions in infinite-domain constraint satisfaction such as the structural assumptions that can be imposed on the templates in the scope of the Bodirsky-Pinsker conjecture.
- Michał Wrona gave a talk about the complexity of infinite-domain CSPs with bounded strict width in the scope of Bodirsky-Pinsker conjecture.
- Johanna Brunar presented a result on pseudo loops in smooth digraphs.
- Michael Pinsker talked about a necessary polymorphism condition for tractability of CSPs that applies to a large class of finite and $\omega$-categorical templates.
- Maximilian Hadek gave a talk about the connection between König's tree lemma and Ramsey's theorem, two important tools for infinite-domain CSPs.

### Promise CSPs

Following the trend of the previous workshops, the topic of promise CSPs, which can be viewed as a framework for qualitative approximation in CSPs, was addressed in 3 contributed talks and touched in several more talks.

- Andrei Krokhin presented a proof that $\mathrm{PCSP}(\mathrm{LO}_2, \mathrm{LO}_3)$ is NP-hard.
- Santiago Guzmán Pro presented his result showing that if the tractability of a finite-template PCSP is explained by a GMSNP sandwich, then it is also explained by a finite sandwich.
- Per Austrin introduced the concept of usefulness for PCSPs, focusing on Boolean PCSPs.

### Topological methods for CSPs and PCSPs

Continuing the newly emerging topological approach to classifying complexity of CSPs and PCSPs, there were 2 contributed talks concerning this topic.

- Jakub Opršal gave an introduction to homomorphism complexes and their use in classifying the complexity of CSPs.
- Uli Wagner presented a topological proof that it is NP-hard to distinguish between graphs that are $G$-colorable and graphs that are not even 4-colorable, where $G$ is any non-bipartite, 4-colorable graph.

### Approximability of CSPs

Another prominent topic of the seminar is the quantitative approximability of CSPs, sometimes combined with the setting of promise CSPs to obtain, in some sense, both quantitative and qualitative approximation results. The seminar started with a 3-hour tutorial given by Amey Bhangale.

- In the first part, he presented approximation algorithms and known results on hardness of approximation.
- The second part of the tutorial was focused on Raghavendra's characterization of the approximability of almost satisfiable Max-CSPs.
- In the third part, he discussed the recent work on the approximability of satisfiable CSPs including the hybrid algorithm that combines SDP rounding and Gaussian elimination algorithms.

There were 8 contributed talks devoted to this topic:

- Silvia Butti presented a result on inapproximability of promise equations over finite groups.
- Libor Barto talked about the framework of valued PCSPs and reduction between them based on the minor condition problem.

- Konstantin Makarychev gave a talk on approximability of phylogenetic CSPs.
- George Osipov presented results on optimal approximation factors for Boolean MinCSPs.
- Euiwoong Lee presented results on approximability of MinCSPs on complete instances.
- Tamio-Vesa Nakajima introduced a new approximation framework called maxPCSPs and presented some complexity results for such problems.
- Santhoshini Velusamy talked about approximability of constraint satisfaction problems in the streaming model.
- Anuj Dawar gave a talk about undefinability of approximation in the sense that hardness is captured by undefinability in FPC.

### Other variants of CSPs

There were 3 contributed talks concerning other variants of CSPs than those mentioned above.

- Yury Makarychev presented results on algorithms for CSPs with ML oracle advice.
- Barnaby Martin talked about about a new concept of restricted CSPs on digraphs and a connection to CSPs of digraphs with a restricted class of instances.
- Žaneta Semanišinová gave a talk on valued CSPs and their connection with resilience problems from database theory.

### Applications

Two speakers gave talks on applications of CSPs to other areas in combinatorics and coding theory.

- Pravesh Kothari gave a talk on using a method based on k-XOR CSPs to resolve extremal combinatorial problems, including advances on the hypergraph Moore bound, lower bounds for locally decodable and correctable codes, and other applications.
- Xuandi Ren presented results on inapproximability for the problems of Nearest Code Word and Minimum Distance.

## 2 Table of Contents

**Executive summary**
*Manuel Bodirsky, Venkatesan Guruswami, Dániel Marx, and Stanislav Živný* . . . 114

**Overview of Talks**

On the Usefulness of Promises
*Per Austrin* . . . 121

The Rise of Plurimorphisms: Algebraic Approach to Approximation
*Libor Barto* . . . 121

On Approximability of Satisfiable Max-P-CSPs
*Amey Bhangale* . . . 122

The sorrows of a smooth digraph: lifting finite to infinite
*Johanna Brunar* . . . 122

Optimal Inapproximability of Promise Equations over Finite Groups
*Silvia Butti* . . . 123

Undefinability of Approximations
*Anuj Dawar* . . . 123

GMSNP and finite-template PCSPs
*Santiago Guzmán Pro* . . . 123

Kőnig = Ramsey
*Maximilian Hadek* . . . 124

Spectral Refutations and Their Applications to Algorithms and Combinatorics
*Pravesh K. Kothari* . . . 124

PCSP($LO_2$,$LO_3$) is NP-hard
*Andrei Krokhin* . . . 125

Min-CSPs on Complete Instances
*Euiwoong Lee* . . . 125

Phylogenetic CSPs are Approximation Resistant
*Konstantin Makarychev* . . . 126

Constraint Satisfaction Problems with Advice
*Yury Makarychev* . . . 126

Restricted CSPs and H-free Digraph Algorithmics
*Barnaby Martin* . . . 127

Constraint satisfaction with orbit-finite constraint languages
*Antoine Mottet* . . . 127

An invitation to maxPCSP
*Tamio-Vesa Nakajima* . . . 128

An introduction to homomorphism complexes (in the context of CSPs)
*Jakub Opršal* . . . 128

Optimal FPT-Approximation Factors for some CSPs
*George Osipov* . . . 128

Binary symmetries of tractable non-rigid structures
*Michael Pinsker* . . . . . 129

PCP-free Inapproximability of Nearest Codeword and Minimum Distance
*Xuandi Ren* . . . . . 129

Fundamental Questions in Modern Infinite-Domain Constraint Satisfaction
*Jakub Rydval* . . . . . 130

Valued CSPs and Resilience in Database Theory
*Žaneta Semanišinová* . . . . . 130

Approximability of Constraint Satisfaction Problems in the Streaming Setting
*Santhoshini Velusamy* . . . . . 131

Topology and Hardness of 4-Coloring G-colorable graphs
*Uli Wagner* . . . . . 131

On Bounded Strict Width and NL over Homogeneous Structures with Finite Duality
*Michal Wrona* . . . . . 132

A characterization of singleton algorithms for the Constraint Satisfaction Problem
*Dmitriy Zhuk* . . . . . 132

**Participants** . . . . . 133

# 3 Overview of Talks

## 3.1 On the Usefulness of Promises

*Per Austrin (KTH Royal Institute of Technology - Stockholm, SE)*


**Joint work of** Per Austrin, Johan Håstad, Njörn Martinsson

We introduce the concept of usefulness for PCSPs. A promise A is said to be useful if PCSP(A,B) is tractable for some non-trivial B. We then initiate investigations of this notion, focusing on the setting of folded Boolean PCSPs. This exploration leads us to revisit and generalize the "(2+eps)-Sat" problem.

While we do not obtain complete characterizations, we are able to characterize all but a tiny fraction of promises. Perhaps surprisingly, this is achieved using existing general algorithms and hardness results, and the main technical challenge is to prove that those results are applicable.

Along the way we encounter a number of curious concrete Boolean PCSPs which can serve as challenges for pushing existing algorithms and hardness criteria further.

## 3.2 The Rise of Plurimorphisms: Algebraic Approach to Approximation

*Libor Barto (Charles University - Prague, CZ)*


**Joint work of** Libor Barto, Silvia Butti, Alexandr Kazda, Caterina Viola, Stanislav Zivný
**Main reference** Libor Barto, Silvia Butti, Alexandr Kazda, Caterina Viola, Stanislav Zivný: "Algebraic Approach to Approximation", in Proc. of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2024, Tallinn, Estonia, July 8-11, 2024, pp. 10:1–10:14, ACM, 2024.
**URL** https://doi.org//10.1145/3661814.3662076

Following the success of the so-called algebraic approach to the study of decision constraint satisfaction problems (CSPs), exact optimization of valued CSPs, and most recently promise CSPs, we propose an algebraic framework for valued promise CSPs.

To every valued promise CSP we associate an algebraic object, its so-called valued minion. Our main result shows that the existence of a homomorphism between the associated valued minions implies a polynomial-time reduction between the original CSPs. We also show that this general reduction theorem includes important inapproximability results, for instance, the inapproximability of almost solvable systems of linear equations beyond the random assignment threshold.

## 3.3 On Approximability of Satisfiable Max-P-CSPs

*Amey Bhangale (University of California - Riverside, US)*


**Joint work of** Amey Bhangale, Subhash Khot, Dor Minzer, Yang P. Liu
**Main reference** Amey Bhangale, Subhash Khot, Dor Minzer: "On Approximability of Satisfiable k-CSPs: V", in Proc. of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025, pp. 62–71, ACM, 2025.
**URL** https://doi.org//10.1145/3717823.3718127

This tutorial explores past and recent advances in understanding the approximability of Max-CSPs. In the first tutorial, I will discuss the approximation algorithms and the known hardness of approximation results. In the second tutorial, I will go over Raghavendra's characterization of the approximability of almost satisfiable Max-CSPs. Finally, in the last tutorial, I will discuss the recent work on the approximability of satisfiable CSPs including the hybrid algorithm that combines SDP rounding and Gaussian elimination algorithms in a non-trivial way. Inverse theorems are the theorem statements that find a structure on correlated functions. I will discuss the statement of the inverse theorems for 3-ary correlations and how to use them in the analysis of the hybrid algorithm.

## 3.4 The sorrows of a smooth digraph: lifting finite to infinite

*Johanna Brunar (TU Wien, AT)*


**Main reference** Johanna Brunar, Marcin Kozik, Tomás Nagy, Michael Pinsker: "The sorrows of a smooth digraph: the first hardness criterion for infinite directed graph-colouring problems", CoRR, Vol. abs/2501.17060, 2025.
**URL** https://doi.org//10.48550/ARXIV.2501.17060

As powerful as algebraic methods are for finite-domain CSPs, they exhibit challenges to be generalised for infinite, countably categorical structures. Rather than lifting the algebraic notions tailored for finite algebras, it is often easier to extend combinatorial constructions to the infinite setting. The broader goal, therefore, is to recast known finite structural dichotomies as combinatorial proofs that can then be lifted to the countably categorical case. Pursuing this programme, we overcome previous obstacles to lifting structural results for digraphs in this context from the finite to the omega-categorical realm – the strongest lifting results hitherto not going beyond a generalisation of the Hell-Nešetřil theorem for undirected graphs.

## 3.5 Optimal Inapproximability of Promise Equations over Finite Groups

*Silvia Butti (University of Oxford, GB)*


**Joint work of** Silvia Butti, Alberto Larrauri, Stanislav Zivný
**Main reference** Silvia Butti, Alberto Larrauri, Stanislav Zivný: "Optimal Inapproximability of Promise Equations over Finite Groups", CoRR, Vol. abs/2411.01630, 2024.
**URL** https://doi.org//10.48550/ARXIV.2411.01630

A celebrated result of Håstad established that, for any constant $\varepsilon > 0$, it is NP-hard to find an assignment satisfying a $(1/|G| + \varepsilon)$-fraction of the constraints of a given 3-LIN instance over an Abelian group G even if one is promised that an assignment satisfying a $(1 - \varepsilon)$-fraction of the constraints exists. Engebretsen, Holmerin, and Russell showed the same result for 3-LIN instances over any finite (not necessarily Abelian) group. In other words, for almost-satisfiable instances of 3-LIN the random assignment achieves an optimal approximation guarantee. We prove that the random assignment algorithm is still best possible under a stronger promise that the 3-LIN instance is almost satisfiable over an arbitrarily more restrictive group.

## 3.6 Undefinability of Approximations

*Anuj Dawar (University of Cambridge, GB)*


**Joint work of** Anuj Dawar, Bálint Molnár
**Main reference** Anuj Dawar, Bálint Molnár: "Undefinability of Approximation of 2-To-2 Games", in Proc. of the 33rd EACSL Annual Conference on Computer Science Logic, CSL 2025, February 10-14, 2025, Amsterdam, Netherlands, LIPIcs, Vol. 326, pp. 16:1–16:21, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.CSL.2025.16

In this talk, I review work on the hardness of defining the approximation of optimization problems in a logic such as fixed-point logic with counting (FPC), which corresponds to a natural symmetric fragment of P. Decision problems in FPC include those solvable by linear and semidefinite programming, yet we are able to show unconditional lower bounds against it. In recent work with Bálint Molnár, we establish the undefinability of approximation of 2-to-2 games by showing that no FPC class separates the satisfiable instances from those that are no more than epsilon-satisfiable. This has significant implications for promise graph colouring problems

## 3.7 GMSNP and finite-template PCSPs

*Santiago Guzmán Pro (TU Dresden, DE)*


**Main reference** Santiago Guzmán-Pro: "GMSNP and Finite Structures", CoRR, Vol. abs/2406.13529, 2024.
**URL** https://doi.org//10.48550/ARXIV.2406.13529

CSPs expressible in guarded monotone strict NP (GMSNP) remain one of the most prominent open cases within the scope of the Bodirsky-Pinsker conjecture about infinite-domain CSPs. A conjecture of Brakensiek and Guruswami about promise CSPs implies that if G is an

infinite non-bipartite graph with finite chromatic number, then CSP(G) is NP-hard. In this talk we are interested in the interplay between CSPs expressible in GMSNP and promise CSPs. We will we see that for every template S of a CSP expressible in GMSNP, there is a finite-domain "approximation" C of its CSP. Using this approximation we show the following.
1) If G is a non-bipartite graph with finite chromatic number whose CSP is in GMSNP, then CSP(G) is NP-complete.
2) If the tractability of a finite-template PCSP is explained by a GMSNP sandwich, then it is also explained by a finite sandwich.

## 3.8 Kőnig = Ramsey

*Maximilian Hadek (Charles University - Prague, CZ)*



Ramsey's theorem and Kőnig's tree lemma are two famous combinatorial results from the early 20th century, the former talks about colouring subsets of finite sets, while the latter is more akin to a choice principle, allowing us to find infinite paths within trees. At first glance they seem unrelated, though they share one moral similarity: both establish order within seemingly chaotic, infinite objects. We make this precise, by proving that a generalized Kőnigs lemma holds on a category if and only if it has the Ramsey property.

## 3.9 Spectral Refutations and Their Applications to Algorithms and Combinatorics

*Pravesh K. Kothari (Princeton University, US)*



I will present a method to reduce extremal combinatorial problems to establishing the unsatisfiability of k-sparse linear equations mod 2 (aka k-XOR CSP) with a limited amount of randomness. This latter task is then accomplished by bounding the spectral norm of certain "Kikuchi" matrices built from the k-XOR formulas. In these talks, I will discuss a couple of applications of this method from the following list.

1. Proving hypergraph Moore bound (Feige's 2008 conjecture) – the optimal trade-off between the number of equations in a system of k-sparse linear equations modulo 2 and the size of the smallest linear dependent subset. This theorem generalizes the famous irregular Moore bound of Alon, Hoory and Linial (2002) for graphs (equivalently, 2-sparse linear equations mod 2).
2. Proving a cubic lower bound on 3-query locally decodable codes (LDCs), improving on a quadratic lower bound of Kerenedis and de Wolf (2004) and its generalization to q-query locally decodable codes for all odd q,
3. Proving an exponential lower bound on linear 3-query locally correctable codes (LCCs). This result establishes a sharp separation between 3-query LCCs and 3-query LDCs that are known to admit a construction with a sub-exponential length. It is also the first result to obtain any super-polynomial lower bound for >2-query local codes.

Time permitting, I may also discuss applications to strengthening Szemeredi's theorem, which asks for establishing the minimal size of a random subset of integers S such that every dense subset of integers contains a 3-term arithmetic progression with a common difference from S, and the resolution of Hamada's 1970 conjecture on the algebraic rank of binary 4-designs.

## 3.10 PCSP($LO_2$,$LO_3$) is NP-hard

*Andrei Krokhin (Durham University, GB)*


**Joint work of** Andrei Krokhin, Danny Vagnozzi

We prove NP-hardness of a PCSP with a specific small template, which concerns linearly ordered colourings of 3-uniform hypergraphs. This PCSP was often mentioned as a specific obstacle in the project of classifying the complexity of "approximate 1-in-3-SAT" problems. Our proof uses a well-known sufficient condition for NP-hardness of PCSPs, but exhibits a new interesting behaviour of polymorphisms.

## 3.11 Min-CSPs on Complete Instances

*Euiwoong Lee (University of Michigan - Ann Arbor, US)*


**Joint work of** Aditya Anand, Euiwoong Lee, Jason Li, Yaowei Long, Thatchaphol Saranurak
**Main reference** Aditya Anand, Euiwoong Lee, Jason Li, Yaowei Long, Thatchaphol Saranurak: "Unbreakable Decomposition in Close-to-Linear Time", in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 1464–1493, SIAM, 2025.
**URL** https://doi.org//10.1137/1.9781611978322.46

Given a fixed arity $k \geq 2$, Min-k-CSP on complete instances is the problem whose input consists of a set of n variables V and one (nontrivial) constraint for every k-subset of variables (so there are $\binom{n}{k}$ constraints), and the goal is to find an assignment that minimizes the number of unsatisfied constraints. Unlike Max-k-CSP that admits a PTAS on more general dense or expanding instances, the approximability of Min-k-CSP has not been well understood. Moreover, for some CSPs including Min-k-SAT, there is an approximation-preserving reduction from general instances to dense and expanding instances, leaving complete instances as a unique family that may admit new algorithmic techniques.

In this paper, we initiate the systematic study of Min-CSPs on complete instances. First, we present an O(1)-approximation algorithm for Min-2-SAT on complete instances, the minimization version of Max-2-SAT. Since O(1)-approximation on dense or expanding instances refutes the Unique Games Conjecture, it shows a strict separation between complete and dense/expanding instances.

Then we study the decision versions of CSPs, whose goal is to find an assignment that satisfies all constraints; an algorithm for the decision version is necessary for any nontrivial approximation for the minimization objective. Our second main result is a quasi-polynomial time algorithm for every Boolean k-CSP on complete instances, including Min-k-SAT. We complement this result by giving additional algorithmic and hardness results for CSPs with a larger alphabet, yielding a characterization of (arity, alphabet size) pairs that admit a quasi-polynomial time algorithm on complete instances.

## 3.12 Phylogenetic CSPs are Approximation Resistant

*Konstantin Makarychev (Northwestern University - Evanston, US)*


**Joint work of** Vaggos Chatziafratis, Konstantin Makarychev

In this talk, I will introduce a class of constraint satisfaction problems known as Phylogenetic Constraint Satisfaction Problems. We will consider specific examples from this class, including MaxTriplets and MaxQuarters, which are motivated by applications in computational biology and database theory. We will explore the notions of approximation resistance and biased random assignment. Finally, we will examine the relationship between phylogenetic CSPs and ordering CSPs and show that phylogenetic CSPs are approximation resistant building on the seminal result of Guruswami, Håstad, Manokaran, Raghavendra, and Charikar (2011).

## 3.13 Constraint Satisfaction Problems with Advice

*Yury Makarychev (TTIC - Chicago, US)*


**Joint work of** Suprovat Ghoshal, Konstantin Makarychev, Yury Makarychev
**Main reference** Suprovat Ghoshal, Konstantin Makarychev, Yury Makarychev: "Constraint Satisfaction Problems with Advice", in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 1202–1221, SIAM, 2025.
**URL** https://doi.org//10.1137/1.9781611978322.36

We initiate the study of algorithms for constraint satisfaction problems with ML oracle advice. We introduce two models of advice and then design approximation algorithms for Max Cut, Max 2-Lin, and Max 3-Lin in these models. In particular, we show the following.

1. For Max-Cut and Max 2-Lin, we design an algorithm that yields a near-optimal solution when the average degree is larger than a threshold degree, which only depends on the amount of advice and is independent of the instance size. We also give an algorithm for nearly satisfiable Max 3-Lin instances with quantitatively similar guarantees.

2. Further, we provide impossibility results for algorithms in these models. In particular, under standard complexity assumptions, we show that Max 3-Lin is still $1/2 + \eta$ hard to approximate given access to advice, when there are no assumptions on the instance degree distribution. Additionally, we also show that Max 4-Lin is $1/2 + \eta$ hard to approximate even when the average degree of the instance is linear in the number of variables.

Based on joint work with Suprovat Ghoshal and Konstantin Makarychev.

## 3.14 Restricted CSPs and H-free Digraph Algorithmics

*Barnaby Martin (Durham University, GB)*


**Joint work of** Santiago Guzmán-Pro, Barnaby Martin
**Main reference** Santiago Guzmán-Pro, Barnaby Martin: "Restricted CSPs and F-Free Digraph Algorithmics", in Proc. of the 52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark, LIPIcs, Vol. 334, pp. 158:1–158:21, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.ICALP.2025.158

In recent years, much attention has be placed on the complexity of graph homomorphism problems when the input is restricted to $P_k$-(subgraph)-free graphs. We consider the directed version of this research line, by addressing the question: is it true that digraph homomorphism problems CSP($H$) have a P versus NP-complete dichotomy when the input is restricted to $DP_k$-(subgraph)-free digraphs (where $DP_k$ is the directed path on $k$ vertices)? We build on the theory of constraint satisfaction problems to address this question.

Our first main results are a P versus NP-complete classification of CSPs when the input is restricted to $\mathcal{F}$-homomorphism-free digraphs, or restricted to CSP($H'$) for some finite digraph $H'$. We then use the established connection to constraint satisfaction theory to show our third main result (and partial answer to the question above): if CSP($H$) is NP-complete, then there is a positive integer $N$ such that CSP($H$) remains NP-hard even for $DP_N$-(subgraph)-free digraphs. Moreover, CSP($H$) remains NP-hard for $DP_N$-(subgraph)-free acyclic digraphs, and becomes polynomial-time solvable for $DP_{N-1}$-(subgraph)-free acyclic digraphs.

Another contribution of this work is verifying the question above for digraphs on three vertices and a family of smooth tournaments.

## 3.15 Constraint satisfaction with orbit-finite constraint languages

*Antoine Mottet (TU Hamburg, DE)*



In this tutorial, I give an introduction to the type of problems that are studied in the area of infinite-domain constraint satisfaction with omega-categorical templates. With a template-free approach first, I show that the finite and infinite-domain CSPs are very similar and that infinite-domain CSPs are easily motivated by certain coloring problems for graphs. In a second part, I outline some of the more recent results concerning infinite-domains CSPs, in particular the recent theory of smooth approximations, used to study questions of hardness for infinite-domain CSPs, as well as some of my recent results showing how to think about infinite-domain CSPs in terms of associated promise CSPs with finite templates.

## 3.16 An invitation to maxPCSP

*Tamio-Vesa Nakajima (University of Oxford, GB)*


**Joint work of** Tamio-Vesa Nakajima, Stanislav Zivný
**Main reference** Tamio-Vesa Nakajima, Stanislav Zivný: "Maximum Bipartite vs. Triangle-Free Subgraph", in Proc. of the 52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8-11, 2025, Aarhus, Denmark, LIPIcs, Vol. 334, pp. 121:1–121:16, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.ICALP.2025.121

We present a new framework of approximation problems, combining quantitative and qualitative approximation. We explore several new tractability and hardness results within this framework: maximum dicut vs. cut, maximum dicut vs. acyclic subgraph, maximum k vs. l colouring, and maximum cut vs. triangle-free subgraph.

This is joint work with Standa Živný.

## 3.17 An introduction to homomorphism complexes (in the context of CSPs)

*Jakub Opršal (University of Birmingham, GB)*


**Joint work of** Sebastian Meyer, Jakub Opršal
**Main reference** Sebastian Meyer, Jakub Oprsal: "A topological proof of the Hell-Nešetřil dichotomy", in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 4507–4519, SIAM, 2025.
**URL** https://doi.org//10.1137/1.9781611978322.154

The talk included some topology in an emerging new application of homotopy theory in CSPs and promise CSPs. This approach is build on the question: How does the topology of the solution space of a computational problem influence its complexity? I presented an introduction to a key notion of *homomorphism complexes* used in this line of work, and a brief overview of what have we done with it so far.

## 3.18 Optimal FPT-Approximation Factors for some CSPs

*George Osipov (Linköping University, SE)*


**Joint work of** George Osipov, Konrad K. Dabrowski, Peter Jonsson, Sebastian Ordyniak, Magnus Wahlström

In MinCSP(A), we are given an instance of CSP(A), and the goal is to compute its cost, which is the minimum number of constraints violated by any assignment to its variables. For many choices of A, MinCSP(A) is notoriously hard. For example, even in the Boolean case where A consists solely of binary disequality, the Unique Games Conjecture rules out constant-factor approximation in polynomial time. We study FPT-algorithms for this problem with the natural parameter being the cost.

For every Boolean bijunctive language A, we show that there exists a constant $c = c(A)$ such that: – MinCSP(A) admits a factor-$c$ FPT-approximation, and – MinCSP(A) does not admit a factor – $(c - \varepsilon)$ FPT-approximation for any $\varepsilon > 0$, unless the ETH is false.

Note that the case with $c = 1$ is equivalent to finding an exact solution, so our result refines part of the Boolean MinCSP FPT dichotomy by Kim, Kratsch, Pilipczuk and Wahlström [SODA'2023]. In fact, their FPT-algorithm with a small modification in the last step is sufficient to achieve optimal approximation factors, so our main contribution is proving matching lower bounds under the ETH. Our reduction builds upon the W[1]-hardness proof by [Marx & Razgon, IPL'2009] for, essentially, Boolean MinCSP with constant unary relations and a 4-ary relation $(a \to b) \land (c \to d)$. We use a generalization of k-Clique - the Densest k-Subgraph problem - combined with the recent breakthrough of Guruswami, Lin, Ren, Sun & Wu [STOC'2024] proving Parameterized Inapproximability Hypothesis under ETH and a densification result due to Bafna, Karthik & Minzer [STOC'25]. As a corollary, we obtain that several FPT-approximation algorithms in the literature achieve optimal approximation factors under the ETH.

## 3.19 Binary symmetries of tractable non-rigid structures

*Michael Pinsker (TU Wien, AT)*


**Main reference** Paolo Marimon, Michael Pinsker: "Minimal operations over permutation groups", CoRR, Vol. abs/2410.22060, 2025
**URL** https://arxiv.org/abs/2410.22060

It is well-known that finite or omega-categorical tractable CSP templates must have an essential polymorphism, i.e. a polymorphism that depends on more than one variable. In fact, it is folklore that they must have a ternary such polymorphism. Well-known examples of ternary essential polymorphisms are those satisfying the (ternary) majority or minority identities; these examples also show that in general, one cannot expect binary essential polymorphisms for tractable templates.

We report on a recent result stating that if a template is a finite or omega-categorical core, and its automorphisms do not happen to form a Boolean group acting freely, then tractability of its CSP implies the existence of a binary essential polymorphism. The proof of this result builds on a new generalization/improvement of results of Rosenberg and Bodirsky+Chen on minimal clones.

## 3.20 PCP-free Inapproximability of Nearest Codeword and Minimum Distance

*Xuandi Ren (University of California - Berkeley, US)*


**Joint work of** Vijay Bhattiprolu, Vanketasan Guruswami, Euiwoong Lee, Xuandi Ren

We present simple deterministic gap-producing reductions from the canonical NP-hard problem of solving systems of quadratic equations to the approximate versions of the Nearest Codeword Problem (NCP) and the Minimum Distance Problem (MDP)—achieving inapproximability within arbitrary constant factors. Our reductions and proofs are short and conceptually clean, relying primarily on linear codes in which all codewords have roughly equal weight.

## 3.21 Fundamental Questions in Modern Infinite-Domain Constraint Satisfaction

*Jakub Rydval (TU Wien, AT)*


**Joint work of** Michael Pinsker, Jakub Rydval, Moritz Schöbi, Christoph Spiess
**Main reference** Michael Pinsker, Jakub Rydval, Moritz Schöbi, Christoph Spiess: "Three Fundamental Questions in Modern Infinite-Domain Constraint Satisfaction", in Proc. of the 50th International Symposium on Mathematical Foundations of Computer Science, MFCS 2025, August 25-29, 2025, Warsaw, Poland, LIPIcs, Vol. 345, pp. 83:1–83:20, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
**URL** https://doi.org//10.4230/LIPICS.MFCS.2025.83

The Feder-Vardi dichotomy conjecture for Constraint Satisfaction Problems (CSPs) with finite templates, confirmed independently by Bulatov and Zhuk, has an extension to certain well-behaved infinite templates due to Bodirsky and Pinsker which remains wide open. In this talk, I motivate several fundamental questions on the scope of the Bodirsky-Pinsker conjecture and provide answers to some of them. This concerns, e.g., identifying model-theoretic assumptions that can be added without loss of generality or finding general connections to finite-domain promise CSPs. The presented material stems from joint work with Michael Pinsker, Moritz Schöbi, and Christoph Spiess.

## 3.22 Valued CSPs and Resilience in Database Theory

*Žaneta Semanišinová (TU Dresden, DE)*


**Joint work of** Manuel Bodirsky, Colin Jahel, Carsten Lutz, Žaneta Semanišinová
**Main reference** Manuel Bodirsky, Žaneta Semanišinová, Carsten Lutz: "The Complexity of Resilience Problems via Valued Constraint Satisfaction Problems", in Proc. of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2024, Tallinn, Estonia, July 8-11, 2024, pp. 14:1–14:14, ACM, 2024.
**URL** https://doi.org//10.1145/3661814.3662071

A recent research topic in database theory is the computational complexity of resilience of queries. For a fixed conjunctive query, the problem is to compute the number of facts that need to be removed from a given database so that it does not satisfy the query. In this talk, I will explain how resilience problems can be viewed as valued constraint satisfaction problems (VCSPs) of structures that are finite or at least finite-like (in the sense that they have an oligomorphic automorphism group). Building on the known results about VCSPs on finite domains, we explore how the setting can be generalized to infinite domains and give some general powerful hardness and tractability conditions for the resilience problem. We also present recent results for queries over a signature consisting of a single binary relation, and results on the existence and uniqueness of cores of valued structures stemming from resilience problems. This is based on joint work with Manuel Bodirsky, Colin Jahel, and Carsten Lutz.

## 3.23 Approximability of Constraint Satisfaction Problems in the Streaming Setting

*Santhoshini Velusamy (TTIC - Chicago, US)*


**Joint work of** Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, Santhoshini Velusamy
**Main reference** Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, Santhoshini Velusamy: "Sketching Approximability of All Finite CSPs", J. ACM, Vol. 71(2), pp. 15:1–15:74, 2024.
**URL** https://doi.org//10.1145/3649435

Streaming model is one of the most successful models for solving combinatorial optimization problems involving big data. In this talk, we explore the topic of approximability of constraint satisfaction problems in the streaming model. The main theorem that we cover is a dichotomy theorem from the joint work of Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, and Santhoshini Velusamy, in the dynamic streaming model, where constraints can be added to or deleted from the stream. For any $\gamma > \beta$, consider the promise problem where the Max-CSP value is promised to be either larger than $\gamma$ or smaller than $\beta$. The dichotomy theorem states that the problem of distinguishing the promise instances can either be solved using logarithmic space or requires at least polynomial space, and gives a PSPACE characterization to decide them. This characterization is discussed in detail for the special case of Boolean CSPs and the talk concludes with some interesting open problems in this area.

## 3.24 Topology and Hardness of 4-Coloring G-colorable graphs

*Uli Wagner (IST Austria - Klosterneuburg, AT)*


**Joint work of** Sergey Avvakumov, Marek Filakovský, Jakub Opršal, Gianluca Tasinato, Uli Wagner
**Main reference** Sergey Avvakumov, Marek Filakovský, Jakub Oprsal, Gianluca Tasinato, Uli Wagner: "Hardness of 4-Colouring k-Colourable Graphs", in Proc. of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025, pp. 72–83, ACM, 2025.
**URL** https://doi.org//10.1145/3717823.3718154

Let G be a non-bipartite, 4-colorable graph. We show that it is NP-hard to distinguish between graphs that are G-colorable and graphs that are not even 4-colorable. This is a common generalisation of previous results of Khanna, Linial, and Safra [Comb. 20(3), 393-415 (2000)] and of Krokhin and Opršal [FOCS 2019] and Wrochna and Živný [SODA 2020].

The proof combines the algebraic theory of promise constraint satisfaction problems developed by Barto, Bulín, Krokhin, and Opršal [J. ACM 68(4), 28:1–66 (2021)] with methods of topological combinatorics and equivariant obstruction theory.

## 3.25 On Bounded Strict Width and NL over Homogeneous Structures with Finite Duality

*Michal Wrona (Jagiellonian University - Kraków, PL)*



The celebrated result of Barto, Kozik and Willard states that every finite-domain CSP with bounded strict width is in NL. We prove that the same holds for a large class of structures within Bodirsky-Pinsker conjecture. The class consists of first-order expansions of structures with finite duality with natural restrictions on the transitivity, homogeneity and algebraicity of the underlying group of automorphisms. In particular, the class contains first-order expansions of homogeneous graphs, digraphs and hypergraphs studied in the literature.

## 3.26 A characterization of singleton algorithms for the Constraint Satisfaction Problem

*Dmitriy Zhuk (Charles University - Prague, CZ)*



We consider singleton versions of standard universal algorithms for the (Promise) Constraint Satisfaction Problem, such as Arc Consistency, Basic Linear Programming (BLP), and Affine Integer Programming (AIP). The only difference is that, for every constraint and every tuple (or for every variable and every element of the domain), we fix the tuple, run one of the above algorithms, and rule out the value if the algorithm returns "No". If all tuples are ruled out, we return "No"; otherwise, we return "Yes".

The minions characterizing the power of these algorithms have been known for some time, but they were too complicated to verify the existence of a minion homomorphism. Using the Hales-Jewett theorem, we showed that the existence of a minion homomorphism is equivalent to the existence of palette block functions with certain properties. As a result, we obtained a characterization of the singleton versions of Arc Consistency, BLP, AIP, and their combinations in terms of symmetric polymorphisms of the constraint language. As a side result, we also showed that the two ways of defining singleton - by fixing a tuple of a constraint or by fixing a value for a variable - have the same power.

## Participants

- Per Austrin
KTH Royal Institute of Technology – Stockholm, SE
- Libor Barto
Charles University – Prague, CZ
- Arash Beikmohammadi
Simon Fraser University – Burnaby, CA
- Amey Bhangale
University of California – Riverside, US
- Manuel Bodirsky
TU Dresden, DE
- Zarathustra Brady
Setauket, US
- Johanna Brunar
TU Wien, AT
- Andrei A. Bulatov
Simon Fraser University – Burnaby, CA
- Silvia Butti
University of Oxford, GB
- Karthik C. S.
Rutgers University – New Brunswick, US
- Siu On Chan
Hong Kong, HK
- Victor Dalmau
UPF – Barcelona, ES
- Anuj Dawar
University of Cambridge, GB
- Dmitry Feichtner-Kozlov
Universität Bremen, DE
- Florian Frick
Carnegie Mellon University – Pittsburgh, US
- Venkatesan Guruswami
University of California – Berkeley, US
- Santiago Guzmán Pro
TU Dresden, DE
- Maximilian Hadek
Charles University – Prague, CZ
- Prahladh Harsha
TIFR – Mumbai, IN
- Johan Hastad
KTH Royal Institute of Technology – Stockholm, SE
- Peter Jonsson
Linköping University, SE
- Eun Jung Kim
KAIST – Daejeon, KR & CNRS – Paris, FR
- Pravesh K. Kothari
Princeton University, US
- Marcin Kozik
Jagiellonian University – Kraków, PL
- Andrei Krokhin
Durham University, GB
- Euiwoong Lee
University of Michigan – Ann Arbor, US
- Konstantin Makarychev
Northwestern University – Evanston, US
- Yury Makarychev
TTIC – Chicago, US
- Barnaby Martin
Durham University, GB
- Dániel Marx
CISPA – Saarbrücken, DE
- Antoine Mottet
TU Hamburg, DE
- Tomáš Nagy
Jagiellonian University – Kraków, PL
- Tamio-Vesa Nakajima
University of Oxford, GB
- Jakub Opršal
University of Birmingham, GB
- George Osipov
Linköping University, SE
- Marcin Pilipczuk
University of Warsaw, PL
- Michael Pinsker
TU Wien, AT
- Aaron Potechin
University of Chicago, US
- Xuandi Ren
University of California – Berkeley, US
- Jakub Rydval
TU Wien, AT
- Žaneta Semanišinová
TU Dresden, DE
- Santhoshini Velusamy
TTIC – Chicago, US
- Uli Wagner
IST Austria – Klosterneuburg, AT
- Magnus Wahlström
Royal Holloway, University of London, GB
- Michal Wrona
Jagiellonian University – Kraków, PL
- Dmitriy Zhuk
Charles University – Prague, CZ
- Stanislav Živný
University of Oxford, GB

Report from Dagstuhl Seminar 25212

# Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs


**Sujoy Bhore*[1], Jie Gao*[2], Hung Le*[3], Csaba D. Tóth*[4], and Lazar Milenković†[5]**

1 **Indian Institute of Technology Bombay – Mumbai, IN.** sujoy@cse.iitb.ac.in
2 **Rutgers University – Piscataway, US.** jg1555@cs.rutgers.edu
3 **University of Massachusetts Amherst, US.** hungle@cs.umass.edu
4 **California State University Northridge – Los Angeles, US.** csaba.toth@csun.edu
5 **Tel Aviv University, IL.** milenkovic.lazar@gmail.com



## Abstract

Sketching is a basic technique to handle big data: Compress a big input dataset into a small dataset, called a *sketch*, that (approximately) preserves the important information in the input dataset. A metric space is often given as a distance matrix with $\Omega(n^2)$ entries, and *metric sketching* techniques aim to reduce the space to linear. One goal of this Dagstuhl Seminar was to understand different sketching techniques and metric spaces that admit small sketches. Another common approach to handling big datasets is dynamic algorithms. Typically, large datasets do not arrive in a single batch; instead, they are updated over time in small increments. The objective of dynamic algorithms is to respond to data updates quickly, ideally with an update time that is polylogarithmic in the size of the whole dataset.

In this Dagstuhl Seminar "Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs" (25212), we considered sketching and dynamic algorithms in the context of geometric intersection graphs and topological graphs. Geometric intersection graphs have been used to model many real-world massive graphs, such as wireless networks. Topological graphs, including planar graphs, have been used in applications such as geographic information systems and motion planning. While geometric intersection graphs and topological graphs are seemingly different, they have common structural properties that allow the transfer of algorithmic techniques between the two domains, which was the motivation of this seminar: Uncovering deeper connections between metric sketching, dynamic algorithms, geometric intersection graphs, and topological graphs. More concretely, we studied: (1) the construction of sketching structures, such as spanners, tree covers, distance oracles, and emulators with optimal parameters for various metrics and graphs, including geometric and topological graphs; (2) dynamic problems in geometric intersections graphs, including connectivity, spanners, shortest paths; and (3) dynamic maintenance of metric sketching structures in topological graphs.



**Seminar** May 18–23, 2025 – https://www.dagstuhl.de/25212

**2012 ACM Subject Classification** Theory of computation → Computational geometry; Mathematics of computing → Discrete mathematics; Theory of computation → Design and analysis of algorithms

**Keywords and phrases** geometric spanners, geometric intersection graphs, planar metrics, metric covering, computational geometry

**Digital Object Identifier** 10.4230/DagRep.15.5.134


* Editor / Organizer
† Editorial Assistant / Collector




Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs, *Dagstuhl Reports*, Vol. 15, Issue 5, pp. 134–157

Editors: Sujoy Bhore, Jie Gao, Hung Le, Csaba D. Tóth, and Lazar Milenković

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

# 1 Executive Summary

*Sujoy Bhore (Indian Institute of Technology Bombay – Mumbai, IN)*
*Jie Gao (Rutgers University – Piscataway, US)*
*Hung Le (University of Massachusetts Amherst, US)*
*Csaba D. Tóth (California State University – Northridge, US)*



Our goal was to bring together world-renowned researchers in Computational Geometry, Graph Theory, Data Structures, and Algorithms and collaboratively advance the research agenda for the study of metric sketching, and dynamic algorithms for geometric and topological graphs. The participants identified and investigated fundamental open problems and research directions. In particular, we worked on specific open problems about metric sketching, geometric intersection graphs, geometry of topological graphs, and their applications to dynamic algorithms.

This Dagstuhl Seminar "Metric Sketching and Dynamic Algorithms for Geometric and Topological Graphs" (25212) brought together researchers from four research domains, and harvested the benefits of the significant interplay between fundamental geometric & topological structures, and modern algorithmic paradigms. The general objectives of the seminar were:

1. To advance knowledge on metric sketching and dynamic algorithms for geometric and topological graphs, and investigate fundamental open problems and research directions.
2. To stimulate inter-disciplinary collaboration between Computer Science (computational geometry, data structure and algorithms) and Mathematics (combinatorics and topological graph theory).
3. To catalog the state of the art, and distill major open problems on metric sketching and dynamic algorithms for geometric and topological graphs.

In addition to identifying fundamental research problems on metric sketching and dynamic algorithms for geometric and topological graphs, we shall studied concrete problems chosen by the participants during the seminar. The specific scientific aims of this seminar were:

1. *Structural properties:* Identify key characteristics of classes of geometric and topological graphs, and leverage those understandings to address fundamental challenges of metric sketching.
2. *Trade-offs:* Provide the best possible trade-offs among fundamental aspects for metric sketching mechanisms, e.g., lightness, sparsity, diameter, degree, and others.
3. *Efficient Algorithms:* Develop efficient and robust algorithms for fundamental problems defined on geometric and topological graphs, in both static and dynamic regimes.

Behind the scientific challenges of this seminar lies the pragmatic need for better understanding the complex network systems, such as, banking, education, transportation, healthcare, among others. The scale of the data in any present day network system is huge. Hence, sparsifcation is a natural mechanism of handing such big networks. Moreover, the data is constantly changing. With a huge amount of data, recoumputation is not an option anymore. Recent progress on dynamic/online algorithms led to the development of the current best tools to address some of these problems. Hence, a deeper understanding of geometric data structures helps improve tradeoffs between resource utilization and solution quality.

## 2 Table of Contents

**Executive Summary**
*Sujoy Bhore, Jie Gao, Hung Le, and Csaba D. Tóth* . . . . . . . . . . . . . . . . . 135

**Introduction** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 138

**Overview of Talks** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 142

Metric Embeddings into Trees and its Various Spin-offs
*Arnold Filtser* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 142

Developments in Geometric Intersection Graphs
*Édouard Bonnet* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 142

Dynamic Algorithms in Geometry
*Wolfgang Mulzer* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 142

Constructing and Routing on Plane Constant Spanners
*Prosenjit Bose* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 142

Bicriteria Approximation for Minimum Dilation Graph Augmentation
*Sampson Wong* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 143

Tree Covers for Euclidean and Planar Metrics and Their Applications
*Lazar Milenković* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 143

VC Dimension and Their Applications in Topological and Geometric Graphs
*Da Wei Zheng* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 144

Distance Approximating Structures for Planar and Minor-Free Graphs
*Jonathan Conroy* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 144

**Open Problems** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 145

DAG Covers for Directed Graph Classes
*Nicole Wein* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 145

Plane $(1+\varepsilon)$-spanners in $\mathbb{R}^2$
*Csaba D. Tóth* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 146

Vantage Point Selection
*Jie Gao, Nicole Wein* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 147

Fat Minors
*Michał Pilipczuk* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 148

Planar Spanners
*Prosenjit Bose* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 148

Patterns in Unweighted Planar Graphs
*Da Wei Zheng* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 149

Sparse $(1+\epsilon)$-emulator for Euclidean Point Sets
*Hung Le* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 150

Tree Cover for Unit Ball Graphs
*Jie Gao* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 151

Similarity of Curves (Fréchet Distance)
*Omrit Filtser* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 152

Dynamic Subgraph Connectivity
*Liam Roditty* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 153

Sparse and Light ($k$-)Banyans
*Sándor Kisfaludi-Bak* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 154

Online Spanners in Metric Spaces
*Sujoy Bhore* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 155

Steiner $(1+\varepsilon)$-spanners in $\mathbb{R}^d$
*Csaba D. Tóth* . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 155

**Working Groups** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 156

**Participants** . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 157

# 3 Introduction

*Sujoy Bhore (Indian Institute of Technology Bombay – Mumbai, IN)*
*Jie Gao (Rutgers University – Piscataway, US)*
*Hung Le (University of Massachusetts Amherst, US)*
*Csaba D. Tóth (California State University – Northridge, US)*


In the last several years multiple new techniques have been developed along separate lines of metric sketching (spanners and tree covers), geometry of topological graphs, and geometric intersection graphs. These new techniques have allowed for new ways to solve problems that were deemed challenging previously. For example, geometric tools such as VC-dimension and additive Voronoi diagram have been used to solve problems in planar graphs or graphs of forbidden minors. These ideas then feed back progress on geometric intersection graphs which were studied mainly in the geometry community. We also believe that these new data structures will inspire further development in handling dynamic settings.

Below we briefly review both recent developments in the four topics and also list directions and open problems we tackled.

**Metric Sketching.** The goal of metric sketching is, given a set $P$ of $n$ points in a metric space $(X, \delta_X)$, construct a *compact* (or low space) structure to (approximately) represent the distances between the points in $P$. Ideally, we would like the structure to have a linear space of $O(n)$. There are many sketching structures; here, we highlight the two most important structures: *spanners* and *tree covers*.

A *$t$-spanner* is a graph $G$ whose vertex set is $P$ and edge set contains pairs of points in $P$ with weight to be the metric distance between the endpoints, and furthermore, $\delta_G(x, y) \leq t \cdot \delta_X(x, y)$ for any two points $x, y \in P$. Spanners with $t = 1 + \epsilon$ have been studied extensively for low dimensional Euclidean and doubling metrics. It was known in the 90s that one can construct Euclidean $(1 + \epsilon)$-spanners which are sparse [1], light [2], and have a small diameter [3]. A recent line of work determines exactly the dependency of sparsity [6], lightness [6], and hop bounds [8] on the distance error parameter $\epsilon$ as well as the role of other points in $(X, \delta_X)$ in these constructions [4, 6].

Another important sketching structure is a tree cover. Informally, a tree cover for $P$ is a collection of trees such that the distance between any two points in $P$ is preserved in at least *one of the trees* in the cover. Formally, a $t$-tree cover is a collection $\mathcal{T}$ of trees such that for every two points $x, y \in P$, $\delta_X(x, y) \leq \min_{T \in \mathcal{T}} \delta_T(x, y) \leq t \cdot \delta_X(x, y)$. Tree covers are stronger than spanners in the sense that they are not only compact but also allow querying distances quickly (in constant time); in this sense, tree covers are closer to an ideal metric sketching. On the other hand, constructing a small tree cover is a very difficult problem. Small tree covers were constructed for low-dimensional Euclidean point sets in the 90s [3], but only recently, the number of trees needed was determined [9], and the result was extended to doubling metrics [7]. Remarkably, it was recently shown that *metrics induced by shortest distances in topological graphs* have a tree cover with a small number of trees [10, 11, 12].

In the context of the developments above, there are a number of questions we have explored in this Dagstuhl Seminar with the aim of pushing the frontier in this topic much further. Which metrics admit a compact distance sketching? Can we construct a sketching structure for high-dimensional Euclidean spaces? What is the precise dependency between the number of trees in the tree cover for topological metrics on the distance error parameter $\epsilon$? Could the existing distance sketching structures be dynamized efficiently when points are inserted/deleted?

**Geometric Intersection Graphs.** Given a set $\mathcal{O}$ of $n$ objects in $\mathbb{R}^d$, a geometric intersection graph $G_\times$ for $\mathcal{O}$ has each vertex corresponding to an object in $\mathcal{O}$ and an edge between two vertices if their corresponding objects intersect. Geometric intersection graphs have been successfully used to model many practical networks, such as wireless networks and map labeling. Geometric intersection graphs share certain interesting structures as some other well-studied graph families, such as planar graphs – specifically, planar graphs are geometric intersection graphs of disks by the circle packing theorem (a.k.a. the Koebe–Andreev–Thurston theorem). However, geometric intersection graphs can be dense, which makes them quite different from planar graphs from a computational perspective. To obtain efficient algorithms, typically, one would like to avoid computing all the edges explicitly. Examples of such algorithms include single source shortest paths in such graphs. On the other hand, useful structures such as balanced separators for planar graphs need to be modified/augmented and may lose the nice bound on the small size [13].

Many problems that are easy for point sets in $\mathbb{R}^d$ become extremely hard for $G_\times$ even for simple, uniform, fat objects. One example is the diameter problem: can we compute the diameter of unit disk graphs, i.e., intersection graphs of unit disks, in truly subquadratic time? This problem remains wide open despite many attempts [14, 15]. Related to metric sketching, it is not known if unit disk graphs have a tree cover with a constant number of trees [5] or spanners of small hop diameters [16, 18].

Another interesting perspective of geometric intersection graphs is that the existence of efficient algorithms on such graphs sometimes crucially depends on the shape, size, dimension, and topological properties of the objects. Of particular interest in this family is to understand how far one can push on efficient algorithms with only topological properties of how two objects intersect (e.g., geometric intersection graphs of pseudo-disks versus disk graphs).

We also investigated dynamic algorithms for geometric intersection graphs. In the dynamic settings, objects are inserted and deleted from time to time. The insertion/deletion of an object implicitly induces the insertion/deletion of up to $\Omega(n)$ edges incident to the object. However, the goal of the dynamic algorithm is to handle each update in (amortized) $O(\text{polylog}(n))$ time. The massive number of implicit updates makes the design of dynamic algorithms extremely challenging. We focused on fundamental problems such as dynamic connectivity [17], dynamic shortest paths [19], and dynamic optimization problems, including independent sets [20, 21] and matching [22], and shortest paths [23, 24].

**Geometry of Topological Graphs.** Topological graphs are graphs equipped with topological properties, for example, planar graphs, surface-embedded graphs, and, more generally, minor-free graphs. Research on topological graphs has largely focused on understanding the structural properties of these graphs and hopes to exploit structural insights to develop efficient algorithms. The separator theorem by Lipton and Tarjan [25] for planar graphs and the graph minor project by Robertson and Seymour are representative examples of this line of research. A recent line of research focuses on understanding the *geometry of topological graphs*, specifically the shortest path metrics of these graphs. A representative question is: To what extent does a planar metric resemble the Euclidean plane? Research on this question has produced deep results, including tree cover of constant size [11, 12], VC dimension [26], and low-treewidth embedding [27], with a large number of algorithmic applications. In this seminar, we sought a better understanding of the geometry of topological graphs by drawing connections to research in metric sketching. For example, techniques in constructing spanners were crucial in developing the first linear-space constant query-time distance oracle in planar graphs [28]. A concrete question was to determine the best trade-offs between the distance error parameter $\epsilon$ and the number of trees, the treewidth, or the size of the metric sketches.

**Dynamic Algorithms.** In dynamic algorithms, the input is updated over time, often involving insertions and deletions of new "data points", for example, inserting/deleting points from a point set or edges from a graph. The goal is to quickly update the solution of an algorithm task when the input is updated. Handling dynamic updates brings an extra (and often very challenging) dimension to algorithm design in the three aforementioned topics. We have touched upon designing dynamic algorithms for metric sketching and geometric intersection graphs. Designing dynamic algorithms for topological graphs under edge updates is an emerging line of research that has not been duly explored.

On the negative side it is known that for exact problems, such as exact shortest paths and exact all pairs shortest paths, any dynamic algorithm must spend polynomial update time or query time under popular fine-grained hypotheses [29]. However, these lower bounds do not hold when we seek an approximation, which is often good enough for applications. Basic problems we explored include approximate single source shortest paths, approximate distance oracles, and approximate optimization problems, such as independent sets and matching [30].

## References


1 J. Mark Keil. Approximating the complete Euclidean graph. In *SWAT*, volume 318 of *Lecture Notes in Computer Science*, pages 208–213. Springer, 1988.

2 Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New sparseness results on graph spanners. *Int. J. Comput. Geom. Appl.*, 5:125–144, 1995.

3 Sunil Arya, Gautam Das, David M. Mount, Jeffrey S. Salowe, and Michiel H. M. Smid. Euclidean spanners: Short, thin, and lanky. In *STOC*, pages 489–498. ACM, 1995.

4 Sujoy Bhore and Csaba D. Tóth. Euclidean Steiner spanners: Light and sparse. *SIAM J. Discret. Math.*, 36(3):2411–2444, 2022.

5 Haim Kaplan, Wolfgang Mulzer, Liam Roditty, and Paul Seiferth. Routing in unit disk graphs. *Algorithmica*, 80(3):830–848, 2018.

6 Hung Le and Shay Solomon. Truly optimal Euclidean spanners. In *FOCS*, pages 1078–1100. IEEE Computer Society, 2019.

7 Yair Bartal, Ora Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. *J. Comput. Syst. Sci.*, 130:26–42, 2022.

8 Hung Le, Lazar Milenkovic, and Shay Solomon. Sparse Euclidean spanners with optimal diameter: A general and robust lower bound via a concave inverse-Ackermann function. In *SoCG*, volume 258 of *LIPIcs*, pages 47:1–47:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

9 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Optimal Euclidean tree covers. In *SoCG*, volume 293 of *LIPIcs*, pages 37:1–37:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

10 Sujoy Bhore, Balázs Keszegh, Andrey Kupavskii, Hung Le, Alexandre Louvet, Dömötör Pálvölgyi, and Csaba D. Tóth. Spanners in planar domains via Steiner spanners and non-Steiner tree covers. In *SODA*, pages 4292–4326. SIAM, 2025.

11 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): O(1) trees suffice. In *FOCS*, pages 2231–2261. IEEE, 2023.

12 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Shortcut partitions in minor-free graphs: Steiner point removal, distance oracles, tree covers, and more. In *SODA*, pages 5300–5331. SIAM, 2024.

13 Mark de Berg, Sándor Kisfaludi-Bak, Morteza Monemizadeh, and Leonidas Theocharous. Clique-based separators for geometric intersection graphs. *Algorithmica*, 85(6):1652–1678, 2023.

14 Karl Bringmann, Sándor Kisfaludi-Bak, Marvin Künnemann, André Nusser, and Zahra Parsaeian. Towards sub-quadratic diameter computation in geometric intersection graphs. In *SoCG*, volume 224 of *LIPIcs*, pages 21:1–21:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

15 Hsien-Chih Chang, Jie Gao, and Hung Le. Computing diameter+2 in truly-subquadratic time for unit-disk graphs. In *SoCG*, volume 293 of *LIPIcs*, pages 38:1–38:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

16 Timothy M. Chan and Zhengcheng Huang. Constant-hop spanners for more geometric intersection graphs, with even smaller size. In *SoCG*, volume 258 of *LIPIcs*, pages 23:1–23:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

17 Timothy M. Chan and Zhengcheng Huang. Dynamic geometric connectivity in the plane with constant query time. In *SoCG*, volume 293 of *LIPIcs*, pages 36:1–36:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

18 Jonathan B. Conroy and Csaba D. Tóth. Hop-spanners for geometric intersection graphs. *J. Comput. Geom.*, 14(2):26–64, 2022.

19 Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in geometric intersection graphs. *J. Comput. Geom.*, 10(1):27–41, 2019.

20 Sujoy Bhore, Martin Nöllenburg, Csaba D. Tóth, and Jules Wulms. Fully dynamic maximum independent sets of disks in polylogarithmic update time. In *SoCG*, volume 293 of *LIPIcs*, pages 19:1–19:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

21 Monika Henzinger, Stefan Neumann, and Andreas Wiese. Dynamic approximate maximum independent set of intervals, hypercubes and hyperrectangles. In *SoCG*, volume 164 of *LIPIcs*, pages 51:1–51:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

22 Sujoy Bhore and Timothy M. Chan. Fully dynamic geometric vertex cover and matching. *CoRR*, abs/2402.07441, 2024.

23 Peyman Afshani, Mark de Berg, Kevin Buchin, Jie Gao, Maarten Löffler, Amir Nayyeri, Benjamin Raichel, Rik Sarkar, Haotian Wang, and Hao-Tsung Yang. Approximation algorithms for multi-robot patrol-scheduling with min-max latency. In *WAFR*, volume 17 of *Springer Proceedings in Advanced Robotics*, pages 107–123. Springer, 2021.

24 Jie Gao, Mayank Goswami, Karthik C. S., Meng-Tsung Tsai, Shih-Yu Tsai, and Hao-Tsung Yang. Obtaining approximately optimal and diverse solutions via dispersion. In *LATIN*, volume 13568 of *Lecture Notes in Computer Science*, pages 222–239. Springer, 2022.

25 Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.

26 Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications. In *SODA*, pages 5332–5360. SIAM, 2024.

27 Vincent Cohen-Addad, Hung Le, Marcin Pilipczuk, and Michal Pilipczuk. Planar and minor-free metrics embed into metrics of polylogarithmic treewidth with expected multiplicative distortion arbitrarily close to 1. In *FOCS*, pages 2262–2277. IEEE, 2023.

28 Hung Le and Christian Wulff-Nilsen. Optimal approximate distance oracle for planar graphs. In *FOCS*, pages 363–374. IEEE, 2021.

29 Amir Abboud and Søren Dahlgaard. Popular conjectures as a barrier for dynamic planar graph algorithms. In *FOCS*, pages 477–486. IEEE Computer Society, 2016.

30 Tuukka Korhonen, Wojciech Nadara, Michal Pilipczuk, and Marek Sokolowski. Fully dynamic approximation schemes on planar and apex-minor-free graphs. In *SODA*, pages 296–313. SIAM, 2024.

# 4 Overview of Talks

## 4.1 Metric Embeddings into Trees and its Various Spin-offs

*Arnold Filtser (Computer Science Department, Bar Ilan University – Ramat Gan, IL)*



Given a metric space $(X, d_X)$, a stochastic metric embedding into trees is a distribution over dominating embeddings of the points $X$ into a tree with polylogarithmic expected distortion. We will begin be describing the classic Bartal'96 embedding, and it's construction of padded decompositions. We will then discuss how this approach can be utilized to obtain online metric embedding (in particular deterministic into Euclidean space). We will then present the padded decomposition of Miller, Peng and Xu, and see how it can be used to obtain stochastic metric embedding into spanning trees. We will finish by mentioning some of the recent works of embedding structured graph families into bounded treewidth graphs.

## 4.2 Developments in Geometric Intersection Graphs

*Édouard Bonnet (Laboratoire de l'Informatique du Parallélisme, ENS de Lyon, FR)*



We survey some recent developments in algorithmic graph theory for geometric intersection graphs, and present some conjectures in the novel topic of induced minor theory.

## 4.3 Dynamic Algorithms in Geometry

*Wolfgang Mulzer (Institut für Informatik, Freie Universität Berlin, DE)*



I will present some classic results and recent advances in dynamic algorithms for geometric problems. In particular, I will focus on algorithms for dynamic lower envelopes in two and three dimensions, for lines, planes, and more general objects. I will also discuss applications for dynamic geometric intersection graphs.

## 4.4 Constructing and Routing on Plane Constant Spanners

*Prosenjit Bose (School of Computer Science, Carleton University – Ottawa, CA)*



Given a weighted graph and a real number $t$, a $t$-spanner of $G$ is a spanning subgraph $H$ with the property that for every edge $xy \in G$, there exists a path between $x$ and $y$ in $H$ whose weight is no more than $t$ times the weight of the edge $xy$. An online routing algorithm is an algorithm that finds a short path in a graph without having full knowledge of the graph. We review results and present open problems on different variants of the problem of constructing and routing on plane geometric $t$-spanners.

## 4.5 Bicriteria Approximation for Minimum Dilation Graph Augmentation

*Sampson Wong (University of Copenhagen, DK)*



Spanner constructions focus on the initial design of the network. However, networks tend to improve over time. In this paper, we focus on the improvement step. Given a graph and a budget $k$, which $k$ edges do we add to the graph to minimise its dilation? Gudmundsson and Wong [1] provided the first positive result for this problem, but their approximation factor is linear in $k$.

Our main result is a $(2\sqrt[r]{2}\ k^{1/r}, 2r)$-bicriteria approximation that runs in $O(n^3 \log n)$ time, for all $r \geq 1$. In other words, if $t^*$ is the minimum dilation after adding any $k$ edges to a graph, then our algorithm adds $O(k^{1+1/r})$ edges to the graph to obtain a dilation of $2rt^*$. Moreover, our analysis of the algorithm is tight under the Erdős girth conjecture.

### References

1 Joachim Gudmundsson and Sampson Wong. Improving the dilation of a metric graph by adding edges. *ACM Trans. Algorithms*, 18(3):20:1–20:20, 2022.

## 4.6 Tree Covers for Euclidean and Planar Metrics and Their Applications

*Lazar Milenković (Tel Aviv University, IL)*



I will survey some recent results on tree covers for Euclidean and planar graphs [1, 2, 3]. Due to their structural simplicity, tree cover are suitable for various algorithmic applications. In this talk I will highlight some of the applications, such as compact routing schemes [1, 2] and construction of spanners with optimal tradeoff between hop-diameter and other parameters, including sparsity [4, 5] and lightness [6].

### References

1 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): O(1) trees suffice. In *FOCS*, pages 2231–2261. IEEE, 2023.

2 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Optimal Euclidean tree covers. In *SoCG*, volume 293 of *LIPIcs*, pages 37:1–37:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

3 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Shortcut partitions in minor-free graphs: Steiner point removal, distance oracles, tree covers, and more. In *SODA*, pages 5300–5331. SIAM, 2024.

4 Hung Le, Lazar Milenkovic, and Shay Solomon. Sparse Euclidean spanners with tiny diameter: A tight lower bound. In *SoCG*, volume 224 of *LIPIcs*, pages 54:1–54:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
5 Hung Le, Lazar Milenkovic, and Shay Solomon. Sparse Euclidean spanners with optimal diameter: A general and robust lower bound via a concave inverse-Ackermann function. In *SoCG*, volume 258 of *LIPIcs*, pages 47:1–47:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
6 Sujoy Bhore and Lazar Milenkovic. Light spanners with small hop-diameter. In *ICALP*, volume 334 of *LIPIcs*, pages 30:1–30:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.


## 4.7 VC Dimension and Their Applications in Topological and Geometric Graphs

*Da Wei Zheng (University of Illinois Urbana-Champaign, US)*



Minor-free graphs and geometric intersection graphs are two classes of graphs that share a surprisingly geometric property: they have very natural set systems that have bounded VC-dimension. We can algorithmically exploit these set systems to design algorithms and data structures not possible in general graphs such as subquadratic time algorithms for diameter and subquadratic space data structures for distance oracles.

## 4.8 Distance Approximating Structures for Planar and Minor-Free Graphs

*Jonathan Conroy (Dartmouth College – Hanover, US)*



In this talk, I will discuss several ways in which one can transform planar/minor-free metrics into some simpler object. One helpful tool in this area is the "buffered cop decomposition", a new way of partitioning minor-free graphs which we recently introduced [1]. I will try to convey some intuition about this structural result, and additionally sketch several applications including $O(1)$-distortion Steiner point removal [1] and optimal $O(\log r)$ padded decomposition for $K_r$-minor-free graphs [2].

### References


1 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Optimal Euclidean tree covers. In *SoCG*, volume 293 of *LIPIcs*, pages 37:1–37:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
2 Jonathan Conroy and Arnold Filtser. How to protect yourself from threatening skeletons: Optimal padded decompositions for minor-free graphs. In *STOC*, pages 2281–2292. ACM, 2025.

# 5 Open Problems

## 5.1 DAG Covers for Directed Graph Classes

*Nicole Wein (University of Michigan – Ann Arbor, US)*



This is an open-ended problem based on a recent paper of mine with Sepehr Assadi and Gary Hoppenworth [1]. In this paper we define the notion of a DAG Cover, which is a directed analog of a tree cover. Given a directed graph, the goal is to construct a small collection of DAGs so that for all $s$, $t$: (1) no DAG in the collection underestimates $\text{dist}(s,t)$, and (2) some DAG in the collection gives a good approximation for $\text{dist}(s,t)$.

In the paper we study the problem when the input is a general directed graph, and get some upper and lower bound results that I will describe below. But first, I will state the problem.

**Problem.** Can we get interesting/useful DAG covers for special classes of directed graphs? Specifically, since tree covers have been very useful for classes of geometric graphs, it is natural to wonder whether DAG covers could be useful for geometric classes of directed graphs. As for what classes of graphs would be natural, I haven't thought about this at all!

**Prior Work:** I will outline the results that we get for general graphs, to help contextualize potential questions to ask about special classes of graphs.

First, one could ask: Do the DAGs need to be subgraphs of the original graph? This turns out to be a consequential question. If "yes" then a directed cycle shows that you need $n$ DAGs in the cover even if you just want to preserve reachability. If "no", then there's a simple construction (that I can show you) that achieves only 2 DAGs and preserves exact distances. So, the interesting question is when you allow only a limited number of additional edges. Here are our results.

1. Lower bound: You can't have $O(n^{2-\varepsilon})$ added edges and $O(n^{1-\varepsilon})$ DAGs, and preserve even reachability (i.e. essentially the strongest possible lower bound for when the number of added edges is a function of the number of vertices $n$. Thus, we turn out attention to the number of edges $m$.
2. Upper bound: With $O(m)$ added edges, you can get only 2 DAGs and preserve reachability. (This one is simple and I can show it to you.)
3. Main upper bound: With $\tilde{O}(m)$ added edges, you can get polylog $n$ DAGs and polylog $n$ approximation factor. (Similar to the known tree cover result of log $n$ trees and log $n$ approximation factor.)
4. Lower bound: With $\tilde{O}(m)$ added edges, for preserving exact distances you can't get $n^{1/6-\varepsilon}$ DAGs.

### References

1 Sepehr Assadi, Gary Hoppenworth, and Nicole Wein. Covering approximate shortest paths with DAGs. In *STOC*, pages 2269–2280. ACM, 2025.

## 5.2 Plane $(1 + \varepsilon)$-spanners in $\mathbb{R}^2$

*Csaba D. Tóth (California State University Northridge – Los Angeles, US)*



For $n$ points in Euclidean plane, several constructions are known for $(1 + \varepsilon)$-spanners: Θ-graphs, Yao-graphs, Greedy spanners, and WSPD-based spanners. As $\varepsilon \to 0$ an $(1 + \varepsilon)$-spanner will inevitably have many edge crossings. If we allow Steiner points, we can subdivide the edges at crossing points and obtain a plane $(1+\varepsilon)$-spanner. How many Steiner points do we really need?

Let $s(n, \varepsilon)$ be the minimum integer such that every set of $n$ points in the plane admits a plane $(1 + \varepsilon)$-spanner with at most $s(n, \varepsilon)$ Steiner points.

**Problem 1.** Determine $s(n, \varepsilon)$.

We know that for $n$ points in the plane, there are $(1 + \varepsilon)$-spanners with $O(\varepsilon^{-1}n)$ edges; and Steiner $(1 + \varepsilon)$-spanners with $O(\varepsilon^{-1/2}n)$ edges. However, if the edges pairwise cross (naively), this would give $O(\varepsilon^{-1}n^2)$ crossings, hence $s(n, \varepsilon) \leq O(\varepsilon^{-1} \cdot n^2)$. It is not difficult to improve this naive bound and show that $s(n, \varepsilon)$ is closer to linear in $n$ The goal is to find a bound of the form $s(n, \varepsilon) = O(f(\varepsilon) \cdot n)$ for some function $f(.)$ with the best possible function $f(.)$.

One can also require the Steiner $(1 + \varepsilon)$-spanner to minimize the number of Steiner point and meet other optimization criteria (e.g., small weight or small hop diameter).

**Problem 2.** Determine the minimum number of Steiner points for a Steiner $(1 + \varepsilon)$-spanner of weight $O(\varepsilon^{-1} \cdot w(MST))$ on $n$ points in the plane.

To show that $s(n, \varepsilon)$ is linear in $n$ one can use the planarization of a greedy $(1+\varepsilon)$-spanner. It is known that the crossing graph is $O(f(\varepsilon))$-degenerate [1]. For a nontrivial lower bound, one might be able to use the so-called degenerate crossing number [2, 3] if Steiner points are modeled as degenerate crossings. But perhaps they can be modeled by bundled crossings [4, 5].

### References

1 David Eppstein and Hadi Khodabandeh. On the edge crossings of the greedy spanner. In *SoCG*, volume 189 of *LIPIcs*, pages 33:1–33:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
2 Eyal Ackerman and Rom Pinchasi. On the degenerate crossing number. *Discret. Comput. Geom.*, 49(3):695–702, 2013.
3 János Pach and Géza Tóth. Degenerate crossing numbers. *Discret. Comput. Geom.*, 41(3):376–384, 2009.
4 Md. Jawaherul Alam, Martin Fink, and Sergey Pupyrev. The bundled crossing number. In *GD*, volume 9801 of *Lecture Notes in Computer Science*, pages 399–412. Springer, 2016.
5 Alan Arroyo and Stefan Felsner. Approximating the bundled crossing number. *J. Graph Algorithms Appl.*, 27(6):433–457, 2023.

## 5.3 Vantage Point Selection

*Jie Gao (Rutgers University – Piscataway, US)*
*Nicole Wein (University of Michigan – Ann Arbor, US)*



**Problem Setup:** We have an undirected graph $G$. (You can assume the graph is either weighted or unweighted, whichever is more convenient for what you're trying to prove.) Assume that for every pair of vertices there is one unique shortest path. Each edge also has a unique "capacity" that is hidden from us. We are allowed to choose vertices to "query". When we query a vertex $s$ then for every vertex $t$ the capacity of the minimum capacity edge on the shortest $st$-path is revealed.

Let OPT be the maximum number of edges that could be revealed by performing a single query in $G$. Note that OPT $< n$ since any shortest paths tree has $< n$ edges.

**Question.** How many queries do we need in the worst case to reveal the capacity of $\Omega(\text{OPT})$ edges? It is known that the correct answer is between $\tilde{\Omega}(\sqrt{n})$ and $O(n)$. We think it should be possible to get $O(n^{1-\varepsilon})$

**Alternate version of question.** We don't have a proof that this question is equivalent to the original question, but we would be happy to answer either one. Given an instance of the problem, let $\text{ALG}_1$ be the number of edges we can reveal with just one single query (in expectation if the algorithm is randomized). What is the minimum competitive ratio $\text{OPT}/\text{ALG}_1$ that we can guarantee? The known bounds are the same as the original problem. I.e., if we can reveal $\geq n^{\varepsilon}$ edges in expectation with a single query, then we've achieved the desired $n^{1-\varepsilon}$ upper bound.

**Motivation.** This problem was recently defined by people who work in the theory and practice of networking, in the context of "internet bandwidth estimation". Direct all of your questions about motivation to Jie.

**Prior Work.** Very little is known about this problem. The trivial upper bound is $n$. There is also a simple lower bound of $\tilde{\Omega}(\sqrt{n})$. Consider $\sqrt{n}$ paths of length $\sqrt{n}$ each. One of the paths (we don't know which) has capacities ordered smallest to largest along the path. All of the other paths have capacities in a random order. In this case, OPT$=\sqrt{n}$ by picking the first vertex of the special path. However, if we query a vertex on a non-special path, in expectation we only reveal roughly $\log n$ capacities. Since we don't know which path is special, we need $\Omega(\sqrt{n})$ queries to find it.

**Observations and simplifying assumptions.** Let's say we are trying to achieve an upper bound of $n^{1-\varepsilon}$ for any small constant $\varepsilon$. Notice that if we query a vertex $v$ all of $v$s incident edges are revealed. So, if $G$ has a vertex of degree at least $n^{\varepsilon}$ then we are done. In fact, we may want to assume that every vertex in $G$ has constant degree, since we don't know how to solve that case.

If OPT $= O(n^{1-\varepsilon})$ then we're done because we can pick $O(n^{1-\varepsilon})$ arbitrary vertices to query, each of which reveals at least one edge. As a simplifying assumption, we may want to assume that OPT reveals $\Omega(n)$ capacities.

Another simplifying assumption under which we don't know the answer: Diameter is $\Theta(\log n)$.

## 5.4 Fat Minors

*Michał Pilipczuk (University of Warsaw, PL)*



All graphs are unweighted. For a graph $G$ and a vertex subset $X \subseteq G$ we call $X$ $(k,r)$-coverable if $X$ can be covered by $k$ balls of radius $r$ in $G$. Motivated by the talk of Édouard, we pose the following.

**Conjecture.** For every graph $H$ there exist constants $c, r \in \mathbb{N}$ such that the following holds: Every $n$-vertex graph $G$ that excludes $H$ as an induced minor has a 2/3-balanced separator that is $(c\sqrt{n}, r)$-coverable.

More generally, the suspicion is that exclusion of an induced minor could be relaxed to exclusion of a fat minor. We say that $G$ contains $H$ as a $d$-fat minor, for a parameter $d \in \mathbb{N}$ if there are connected subgraphs $\{B_u : u \in V(H)\}$ and $\{B_e : e \in E(H)\}$ of $G$ such that

- whenever $u$ is an endpoint of $e$ the subgraphs $B_u$ and $B_e$ intersect; and
- except for the above, all the subgraphs in $\{B_f : f \in V(G) \cup E(G)\}$ are pairwise disjoint and at distance more than $d$ in $G$.

Clearly, if $G$ contains $H$ as a $d$-fat minor for any $d \geq 1$ then $G$ contains $H$ as an induced minor. Therefore, exclusion of a fat minor is a weaker assumption than exclusion of an induced minor. We conjecture that the statement above should also hold under the assumption that $G$ excludes $H$ as a $d$-fat minor for some fixed $d \in \mathbb{N}$ then $c$ and $r$ may also depend on $d$.

To give some background, Georgakopoulos and Papasoglu [1] conjectured that if $G$ excludes $H$ as a $d$-fat minor for some constant $d$ then $G$ is quasi-isometric to some graph $G'$ that is $H'$-minor-free, for some $H'$ depending only on $H$ and $d$. (This is not exactly how they stated it, and their original formulation was disproved. See [2] for a discussion.) If the conjecture of Georgakopoulos and Papasoglu was true, then the conjecture above would follow even for fat minors, because the balanced separator of $G'$ could be pulled back through the quasi-isometry. But the conjecture above can be also attacked without trying to verify the conjecture of Georgakopoulous and Papasoglu in full generality (the feeling is that this conjecture may well be just false).

### References

1 Agelos Georgakopoulos and Panos Papasoglu. Graph minors and metric spaces. *Comb.*, 45(3):33, 2025.
2 James Davies, Robert Hickingbotham, Freddie Illingworth, and Rose McCarty. Fat minors cannot be thinned (by quasi-isometries). *CoRR*, abs/2405.09383, 2024.

## 5.5 Planar Spanners

*Prosenjit Bose (Carleton University – Ottawa, CA)*



Given a set $P$ of $n$ points in the plane, a geometric graph $G(P, E)$ is a graph whose vertex set is $P$ whose edges are segments joining pairs of vertices and edges are typically weighted by their length. A geometric graph $H$ is a $t$-spanner of $G$ if $H$ is a subgraph of $G$ and $\forall xy \in E, d_H(x, y) \leq t|xy|$ for some $t \geq 1$, where $d_H(x, y)$ is the length of the shortest path from $x$ to $y$ in $H$. The smallest value of $t$ that satisfies this property is called the *spanning ratio* of $H$. The graph $G$ is the *underlying* graph of the $t$-spanner $H$. In this setting the underlying graph is the complete geometric graph.

An online routing algorithm $A$ is an algorithm that forwards a message from its current vertex $s$ to a destination vertex $t$ where the forwarding decision is made as a function of $s$, $t$, $I(s)$ and $M$ where $s$ is the coordinates of the current vertex, $t$ is the coordinates of the destination vertex, $I(s)$ is information stored with vertex $s$ in $G$ and $M$ is some bits of information stored in $A$. Typically, $I(s)$ is the set of vertices adjacent to $s$ in $G$ but sometimes some additional information is stored. The routing ratio of an algorithm $A$ is the maximum spanning ratio among all $s, t$ pairs of the path followed by $A$ from $s$ to $t$ in $G$. Thus the routing ratio is an upper bound on the spanning ratio.

The following are some questions regarding spanners where $G$ is the complete geometric graph and $H$ is a planar spanning subgraph of $G$.

- The spanning ratio of the standard Delaunay triangulation on a set of points is at most 1.998 and in some cases it is at least 1.59. Can we improve the upper or lower bound?
- For points in convex position, the upper bound is 1.83 and the lower bound remains 1.59. Can we improve either bound?
- Is there a plane spanner with maximum degree 3 and constant spanning ratio? The lowest known degree bound is 4.
- Can we find a bounded degree plane spanner with spanning ratio better than 2.91. Currently, the best known bound is degree 14 and spanning ratio 2.91. The best known upper bound on the spanning ratio for spanners with maximum degree less than 10 is a degree bound of 8 and spanning ratio 4.41.
- For standard Delaunay graphs, the routing ratio is at most 3.5 and at least 1.7. Can these bounds be improved?
- Tight spanning ratios of 2.6 and 2 are known for Delaunay graphs where the empty circle is replaced by a square and hexagon, respectively. However, the routing ratio on these graphs is $\sqrt{10}$ and 6.4 respectively which leaves a large gap with the spanning ratio.

## 5.6 Patterns in Unweighted Planar Graphs

*Da Wei Zheng (University of Illinois – Urbana-Champaign, US)*



Let $G = (V, E)$ be an undirected unweighted planar graph with $n$ vertices, and let $S = \{s_1, \ldots, s_k\}$ be the (ordered clockwise) vertices of one face of $G$. Let $d(\cdot, \cdot)$ denote the shortest path metric of $G$ The pattern $p_v$ of a vertex $v \in V$ is the vector

$$p_v = \langle d(v, s_2) - d(v, s_1), d(v, s_3) - d(v, s_2), \ldots, d(v, s_k) - d(v, s_{k-1}) \rangle.$$

Note that as the graph is unweighted and the vertices $s_1, \ldots, s_k$ are adjacent on one face of $G$ the entries of $p_v$ are in $\{1, 0, -1\}$. Since there are $n$ vertices in $G$ there are $n$ patterns. However, Li and Parter [3] showed (using VC dimension type arguments) that there are only at most $O(k^3)$ many distinct patterns:

**Theorem.** The number of distinct patterns over all vertices of a graph $G$ is $O(k^3)$

An alternative proof is given in [1], along with a $\Omega(k^2)$ lower bound (a grid graph can also give the same lower bound) and the following conjecture:

**Conjecture.** The number of distinct patterns over all vertices of a graph $G$ is $O(k^2)$.

What is the true answer? I suspect the conjecture is true and the answer should be $\Theta(k^2)$

**Motivation and Consequences.** This set system has been generalized to directed minor-free graphs [2, 4] and geometric intersection graphs [5]. Resolving the conjecture either way give use interesting new tools to work with in these settings.

If the conjecture is true, it would imply a simple algorithm (using separators) to compute the diameter of an unweighted planar graph in $O(n^{5/3})$ time (matching the best known algorithm using Voronoi diagrams).


**References**

1 Shay Mozes, Nathan Wallheimer, and Oren Weimann. Improved compression of the Okamura-Seymour metric. In *ISAAC*, volume 248 of *LIPIcs*, pages 27:1–27:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

2 Adam Karczmarz and Da Wei Zheng. Subquadratic algorithms in minor-free digraphs: (weighted) distance oracles, decrementai reachability, and more. In *SODA*, pages 4338–4351. SIAM, 2025.

3 Jason Li and Merav Parter. Planar diameter via metric compression. In *STOC*, pages 152–163. ACM, 2019.

4 Hung Le and Christian Wulff-Nilsen. VC set systems in minor-free (di)graphs and applications. In *SODA*, pages 5332–5360. SIAM, 2024.

5 Hsien-Chih Chang, Jie Gao, and Hung Le. Computing diameter+2 in truly-subquadratic time for unit-disk graphs. In *SoCG*, volume 293 of *LIPIcs*, pages 38:1–38:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.


## 5.7 Sparse $(1+\epsilon)$-emulator for Euclidean Point Sets

*Hung Le (University of Massachusetts Amherst, US)*



Given an Euclidean point set $P$ say in $\mathbb{R}^2$ an emulator for $P$ with stretch $1+\epsilon$ is a graph $G = (V, E)$ where $P \subseteq V$ such that $\|x, y\|_2 \leq d_G(x, y) \leq (1+\epsilon) \cdot \|x, y\|_2$ for every pair of points $x, y \in P$ How many edges does $G$ have to have?

Note that vertices in $V \setminus P$ do not have to be points in $\mathbb{R}^2$.

If $V \setminus P$ are restricted be points in $\mathbb{R}^2$ then we know the answer: $O(n/\sqrt{\epsilon})$ edges are both necessary and sufficient. Nothing is known when $V \setminus P$ are not restricted to be points in $\mathbb{R}^2$.

Here is an embarrassingly simple case: points in $P$ are on two opposite sides of a unit square, and for each side, there are $n = O(1/\sqrt{\epsilon})$ points, where the distance between any two consecutive points is $\sqrt{\epsilon}$. The complete bipartite graph, which only use points in $P$ and has stretch 1 has $O(1/\epsilon) = O(n/\sqrt{\epsilon})$ edges. Could we beat this bound, using points that are not in $P$ (and also not in $\mathbb{R}^2$)?

## 5.8 Tree Cover for Unit Ball Graphs

*Jie Gao (Rutgers University – Piscataway, US)*



A tree cover for a metric space $(X, d)$ is a collection of trees $\mathcal{T}$ such that for any two points $p, q \in X$ the following two conditions are true:

- for every tree $T \in \mathcal{T}$ $d(p, q) \leq d_T(p, q)$;
- there is a tree $T \in \mathcal{T}$ such that $d_T(p, q) \leq \alpha \cdot d(p, q)$.

The tree cover is called to have distortion $\alpha$.

Tree covers are known to exist for many settings:

- $d$-dimensional Euclidean space: $(1 + \epsilon)$-tree cover with $O(\epsilon^{\frac{(-d+1)}{2}} \log(\frac{1}{\epsilon}))$ trees. [1]
- Doubling metrics: $(1 + \epsilon)$-tree cover with $O(\epsilon^{-d} \log(\frac{1}{\epsilon}))$ trees. [2]
- Planar graphs: $(1 + \epsilon)$-tree cover with $O(1)$ trees. [3]

For unit disk graphs, notice that by using the tree cover of planar graphs, one can immediately construct a $O(1)$-tree cover with $O(1)$ trees, by using a constant planar spanner and then results on planar graphs (e.g., [3]). Naturally one can ask, can we find a $(1+\epsilon)$-tree cover for unweighted unit disk graphs? Through personal communication during the seminar, there seems to be already ongoing work towards resolving the aforementioned question. What is still open is whether such a tree cover can be found for unit ball graphs.

A second question to ask is, what is the best stretch factor one can get to approximate a unit disk graph (or a general geometric intersection graph) by a planar graph?

For unit disk graphs, if we want to take a planar spanner graph (using only edges from the input unit disk graph), the known constructions give a stretch factor such as 449 [4] and 341 [5]. Recently we could get a stretch factor of 286 if we allow Steiner vertices. But these stretch factors are still fairly large. Can we do better? What is a lower bound?

### References

1 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Optimal Euclidean tree covers. In *SoCG*, volume 293 of *LIPIcs*, pages 37:1–37:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

2 Yair Bartal, Ora Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. *J. Comput. Syst. Sci.*, 130:26–42, 2022.

3 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): O(1) trees suffice. In *FOCS*, pages 2231–2261. IEEE, 2023.

4 Nicolas Catusse, Victor Chepoi, and Yann Vaxès. Planar hop spanners for unit disk graphs. In *ALGOSENSORS*, volume 6451 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2010.

5 Ahmad Biniaz. Plane hop spanners for unit disk graphs: Simpler and better. *Comput. Geom.*, 89:101622, 2020.

## 5.9 Similarity of Curves (Fréchet Distance)

*Omrit Filtser (The Open University of Israel – Ra'anana, IL)*



A polygonal curve $A : [0, n] \to \mathbb{R}^d$ consists of $n$ line segments $\overline{a_i a_{i+1}}$ for $i = \{0, 1, \ldots, n-1\}$ where $a_i = A(i)$.

The Fréchet distance between two polygonal curves $A : [0, m] \to \mathbb{R}^d$ and $B : [0, n] \to \mathbb{R}^d$ is often described (informally) by the following analogy. Consider a person and a dog connected by a leash, each walking along a curve from its starting point to its end point. Both can control their speed but they are not allowed to backtrack. The Fréchet distance between the two curves is the minimum length of a leash that is sufficient for traversing both curves in this manner (for a formal definition, see below).

The discrete Fréchet distance (DFD for short) is a popular variant in which the two curves are replaced by two sequences of points, and the person and dog are replaced by two frogs, hopping from one point to the next along their respective sequences. The discrete Fréchet distance is defined as the smallest maximum distance between the frogs that can be achieved in such a joint sequence of hops.

**Some background.** The Fréchet distance can be computed in $\tilde{O}(n^2)$ time, and the discrete version in $O(n^2)$ time, via DP approach. For both variants there is a lower bound of $O(n^{2-\varepsilon})$ (conditioned on SETH), even in 1D, and for an approximation factor of up to 3. The best known approximation algorithm is a very recent breakthrough result [1]: a $(7+\varepsilon)$-approximation (w.h.p) in $\tilde{O}(n^{1.99})$ time. Prior to that, the best approximation algorithms presented trade-offs between approximation quality and running time (see e.g. [2, 3]).

**Proximity queries.** In short: Construct an efficient NNS data structure or a distance oracle, under the (discrete) Fréchet distance.

**Nearest Neighbor Search.** Construct a compact data structure over a set $C$ of $n$ input curves, each of length at most $m$ such that given a query curve $Q$ of length $k$ one can efficiently find the curve from $C$ (approximately) closest to $Q$.

**Distance oracle.** Preprocess a curve of length $m$ into a data structure such that given a query curve of length $k < m$ returns its (approximated) distance to the input curve in $\tilde{O}(k)$ time.

For both problems, under DFD, a $(1+\varepsilon)$ approximation and (near) linear query time can be achieved, but with space exponential in $k$ $n \cdot O(\frac{1}{\varepsilon})^{kd}$ for NNS [4], $O(\frac{1}{\varepsilon})^{kd}$ for distance oracle [5]).

**Open questions.** Can we avoid the exponential dependence on $k$ in the space bound, by allowing a constant approximation factor (instead of $1+\varepsilon$) and a slightly larger query time? Or, by allowing $O(\log\log n)$ approximation factor? Is the $O(\frac{1}{\varepsilon})^{kd}$ factor in the space bound necessary if we wish to have a small query time and an approximation factor of $1+\varepsilon$?

For the (continuous) Fréchet distance the situation is much more complicated. Recently there has been some development in 1D case (see [6]).

Exat distance: For $k \in \{1, 2, 3\}$ there are exact distance oracle with near linear size and polylogarithmic query time [7]. For $k = 4$ the query time becomes $O(\sqrt{n})$. Can it be improved? (See also [8] for exact NNS under $L_\infty$ for $k = 2$.)

### References


1 Siu-Wing Cheng, Haoqiang Huang, and Shuo Zhang. Constant approximation of Fréchet distance in strongly subquadratic time. In *STOC*, pages 2329–2340. ACM, 2025.

2 Thijs van der Horst and Tim Ophelders. Faster Fréchet distance approximation through truncated smoothing. In *SoCG*, volume 293 of *LIPIcs*, pages 63:1–63:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

3 Timothy M. Chan and Zahed Rahmati. An improved approximation algorithm for the discrete Fréchet distance. *Inf. Process. Lett.*, 138:72–74, 2018.

4 Arnold Filtser, Omrit Filtser, and Matthew J. Katz. Approximate nearest neighbor for curves: Simple, efficient, and deterministic. *Algorithmica*, 85(5):1490–1519, 2023.

5 Arnold Filtser and Omrit Filtser. Static and streaming data structures for Fréchet distance queries. *ACM Trans. Algorithms*, 19(4):39:1–39:36, 2023.

6 Karl Bringmann, Anne Driemel, André Nusser, and Ioannis Psarros. Tight bounds for approximate near neighbor searching for time series under the Fréchet distance. In *SODA*, pages 517–550. SIAM, 2022.

7 Boris Aronov, Tsuri Farhana, Matthew J. Katz, and Indu Ramesh. Discrete Fréchet distance oracles. In *SoCG*, volume 293 of *LIPIcs*, pages 10:1–10:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

8 Boris Aronov, Omrit Filtser, Michael Horton, Matthew J. Katz, and Khadijeh Sheikhan. Efficient nearest-neighbor query and clustering of planar curves. In *WADS*, volume 11646 of *Lecture Notes in Computer Science*, pages 28–42. Springer, 2019.


## 5.10 Dynamic Subgraph Connectivity

*Liam Roditty (Bar-Ilan University – Ramat Gan, IL)*



Let $G = (V, E)$ be a fixed undirected graph. In the dynamic subgraph connectivity problem each vertex of $G$ is either marked as on or off. An update either turns off a vertex that is on or turns on a vertex that is off. The data structure has to support connectivity queries on the graph induced by the on vertices while vertices are being updated. The dynamic subgraph connectivity problem was introduced in [1]. In [2] the first data structure for general graphs with $O(m^{0.94})$ amortized update time (using FMM) and $\tilde{O}(m^{1/3})$ query time was presented. In [3] the update time of [2] was improved to amortized update time of $\tilde{O}(m^{2/3})$ (without FMM) while keeping the query time $\tilde{O}(m^{1/3})$.

A data structure with $\tilde{O}(m^{4/5})$ worst case update time and $\tilde{O}(m^{1/5})$ query time was presented in [4]. A randomized data structure with $\tilde{O}(m^{3/4})$ worst case update time and $\tilde{O}(m^{1/4})$ query time was presented in [5]. Notice that the product of the update time and the query time is $\tilde{O}(m)$ in all these results.

This is not a coincidence since there is CLB presented by [6] which states that there is no algorithm with polynomial preprocessing time, $O(m^{\gamma-\Omega(1)})$ amortized update time, and $O(m^{1-\gamma-\Omega(1)})$ amortized query time, for any constants $0 \leq \gamma \leq 1$ as long as $m \leq \min(n^{1/\gamma}, n^{1/(1-\gamma)})$.

**Open problem.** Does there exist an algorithm for dynamic subgraph connectivity with worst-case update time of $\tilde{O}(m^{2/3})$ and worst-case query time $\tilde{O}(m^{1/3})$?

**References**

1 Daniele Frigioni and Giuseppe F. Italiano. Dynamically switching vertices in planar graphs. *Algorithmica*, 28(1):76–103, 2000.
2 Timothy M. Chan. Dynamic subgraph connectivity with geometric applications. In *STOC*, pages 7–13. ACM, 2002.
3 Timothy M. Chan, Mihai Pătraşcu, and Liam Roditty. Dynamic connectivity: Connecting to networks and geometry. *SIAM J. Comput.*, 40(2):333–349, 2011.
4 Ran Duan. New data structures for subgraph connectivity. In *ICALP (1)*, volume 6198 of *Lecture Notes in Computer Science*, pages 201–212. Springer, 2010.
5 Ran Duan and Le Zhang. Faster randomized worst-case update time for dynamic subgraph connectivity. In *WADS*, volume 10389 of *Lecture Notes in Computer Science*, pages 337–348. Springer, 2017.
6 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *STOC*, pages 21–30. ACM, 2015.

## 5.11 Sparse and Light ($k$-)Banyans

*Sándor Kisfaludi-Bak (Aalto University, FI)*



A banyan of a point set $P \subset \mathbb{R}^d$ is a geometric graph $G$ where for any terminal set $K \subset P$ we have that $G$ contains a subtree covering $K$ whose weight is at most $1 + \epsilon$ times the weight of the minimum Euclidean Steiner tree of $K$. (The minimum Euclidean Steiner tree of $K$ is the connected geometric graph of minimum weight that contains $K$ in its vertex set.)

A $k$-banyan is as above, but we only require the property for sets $K$ of size at most $k$. With this definition, 2-banyans are simply (Steiner) spanners.

- What is the best trade-off between the banyan stretch and the number of Steiner vertices (or the number of edges) of a ($k-$)banyan for a given dimension $d$ What about 3-banyans?
- What metrics allow light ($k-$)banyans? What is the optimal lightness in $\mathbb{R}^d$?

There is plenty of work on Steiner spanners, so the case $k = 2$ is well-studied and we have good answers for the optimal sparsity and lightness for fixed $d$ in this case. Light banyans have been constructed in a handful of papers cited below, but with the goal of using them for Steiner tree/forest algorithms; I suspect that these constructions are far from optimal in terms of stretch/edge count and stretch/lightness trade-offs.

Starting points are: [1, 2, 3]. The appendix of [4] may also be relevant.

**References**

1 Yair Bartal and Lee-Ad Gottlieb. Near-linear time approximation schemes for Steiner tree and forest in low-dimensional spaces. In *STOC*, pages 1028–1041. ACM, 2021.
2 David Eisenstat, Philip N. Klein, and Claire Mathieu. An efficient polynomial-time approximation scheme for Steiner forest in planar graphs. In *SODA*, pages 626–638. ACM-SIAM, 2012.
3 Satish Rao and Warren D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *STOC*, pages 540–550. ACM, 1998.
4 Sándor Kisfaludi-Bak, Jesper Nederlof, and Karol Wegrzycki. A gap-ETH-tight approximation scheme for Euclidean TSP. In *FOCS*, pages 351–362. IEEE, 2021.

## 5.12 Online Spanners in Metric Spaces

*Sujoy Bhore (Indian Institute of Technology Bombay – Mumbai, IN)*



We are given a sequence of points $(s_1, \ldots, s_n)$ in a metric space, where the points are presented one-by-one, i.e., point $s_i$ is revealed at step $i$ and $S_i = \{s_1, \ldots, s_i\}$ for $i \in \{1, \ldots, n\}$. The objective of an online spanner algorithm is to maintain a $t$spanner $G_i$ for $S_i$ for all $i$. The algorithm is allowed to add edges to the spanner when a new point arrives, however, it is not allowed to remove any edge from the spanner. Moreover, the algorithm does not know the total number of points in advance. The performance of an online $t$spanner algorithm is measured by comparing it with the offline optimum using the standard notion of competitive ratio, which is defined as $\sup_\sigma \frac{ALG(\sigma)}{OPT(\sigma)}$, where the supremum is taken over all input sequences $\sigma$, $OPT(\sigma)$ is the minimum weight of a $t$spanner for the (unordered) set of points in $\sigma$, and $ALG(\sigma)$ denotes the weight of the $t$spanner produced by the algorithm. Note that, to measure the competitive ratio, it is important that $\sigma$ is a finite sequence of points.

**Problem.** Determine the best possible bounds for the competitive ratios for the weight/-lightness, sparsity, and stretch of online $t$-spanners, for $t > 1$.

References: [1, 2, 3, 4].

### References

1 Sujoy Bhore and Csaba D. Tóth. Online Euclidean spanners. *ACM Trans. Algorithms*, 21(1):5:1–5:22, 2025.
2 Sujoy Bhore, Arnold Filtser, Hadi Khodabandeh, and Csaba D. Tóth. Online spanners in metric spaces. *SIAM J. Discret. Math.*, 38(1):1030–1056, 2024.
3 Noga Alon and Yossi Azar. On-line Steiner trees in the Euclidean plane. In *SCG*, pages 337–343. ACM, 1992.
4 Makoto Imase and Bernard M. Waxman. Dynamic Steiner tree problem. *SIAM J. Discret. Math.*, 4(3):369–384, 1991.

## 5.13 Steiner $(1 + \varepsilon)$-spanners in $\mathbb{R}^d$

*Csaba D. Tóth (California State University Northridge – Los Angeles, US)*



For any set $S$ of $n$ points in Euclidean plane and $\varepsilon > 0$ one can construct a Steiner $(1 + \varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$ where the lightness is the ratio between the weight of the spanner and the weight of an MST for the same point set. This bound is the best possible in the plane. In dimensions $d \geq 3$, however, the current best upper and lower bounds do not match: $\Omega(\varepsilon^{-d/2})$ and $\tilde{O}(\varepsilon^{(1-d)/2})$.

**Problem.** Close the gap between the upper and lower bounds for $d \geq 3$ (or, at least for $d = 3$).

References: [1, 2].

### References

1 Sujoy Bhore and Csaba D. Tóth. Euclidean Steiner spanners: Light and sparse. *SIAM J. Discret. Math.*, 36(3):2411–2444, 2022.
2 Hung Le and Shay Solomon. A unified framework for light spanners. In *STOC*, pages 295–308. ACM, 2023.

# 6 Working Groups

- Jonathan Conroy, Liam Roditty, and Hsien-Chih Chang: *DAG Covers for Directed Graph Classes*
- Sujoy Bhore, Sándor Kisfaludi-Bak, Lazar Milenković, Csaba D. Tóth, Karol Wegrzycki, and Sampson Wong: *Plane* $(1+\varepsilon)$*-spanners in* $\mathbb{R}^2$
- Sergio Cabello, Jie Gao, Paloma Lima, and Nicole Wein: *Vantage Point Selection*
- Édouard Bonnet, Hung Le, and Michał Pilipczuk: *Fat Minors*
- Prosenjit Bose, Omrit Filtser, and Linda Kleist: *Planar Spanners*
- Arnold Filtser, Eunjin Oh, Marcin Pilipczuk, and Da Wei Zheng: *Patterns in Unweighted Planar Graphs*

## Participants

- Sujoy Bhore
Indian Institute of Technology Bombay – Mumbai, IN
- Édouard Bonnet
ENS – Lyon, FR
- Prosenjit Bose
Carleton University – Ottawa, CA
- Sergio Cabello
University of Ljubljana, SI
- Hsien-Chih Chang
Dartmouth College – Hanover, US
- Jonathan Conroy
Dartmouth College – Hanover, US
- Arnold Filtser
Bar-Ilan University – Ramat Gan, IL
- Omrit Filtser
The Open University of Israel – Ra'anana, IL
- Jie Gao
Rutgers University – Piscataway, US
- Sándor Kisfaludi-Bak
Aalto University, FI
- Linda Kleist
Universität Potsdam, DE
- Hung Le
University of Massachusetts Amherst, US
- Paloma Lima
IT University of Copenhagen, DK
- Lazar Milenković
Tel Aviv University, IL
- Wolfgang Mulzer
FU Berlin, DE
- Eunjin Oh
POSTECH – Pohang, KR
- Marcin Pilipczuk
University of Warsaw, PL
- Michał Pilipczuk
University of Warsaw, PL
- Liam Roditty
Bar-Ilan University – Ramat Gan, IL
- Csaba D. Tóth
California State University Northridge – Los Angeles, US
- Karol Wegrzycki
MPI für Informatik – Saarbrücken, DE
- Nicole Wein
University of Michigan – Ann Arbor, US
- Sampson Wong
University of Copenhagen, DK
- Da Wei Zheng
University of Illinois – Urbana-Champaign, US