

Open Scholarly Information Systems: Status Quo, Challenges, Opportunities

Hannah Bast^{*1}, Guillaume Cabanac^{*2}, Paolo Manghi^{*3}, Jian Wu^{*4}, and Marcel R. Ackermann^{*5}

- 1 Universität Freiburg, DE. bast@informatik.uni-freiburg.de
- 2 University of Toulouse, FR. guillaume.cabanac@univ-tlse3.fr
- 3 Institute of Information Science and Technologies – CNR – Pisa, IT.
paolo.manghi@isti.cnr.it
- 4 Old Dominion University – Norfolk, US. fanchyna@gmail.com
- 5 Schloss Dagstuhl – Trier, DE. marcel.r.ackermann@dagstuhl.de

Abstract

Over the past 30 years, a rich ecosystem of scholarly information systems has developed that openly provide their services to the scientific community. These systems include aggregators of bibliographic metadata (e.g., DBLP, OpenCitations, OpenAIRE Graph, OpenAlex, ORKG, Semantic Scholar, CiteSeerX, and CORE); publication, data, and software repositories (e.g., Arxiv.org, Figshare, Zenodo, Software Heritage, and Dataverse); and PID authorities (e.g., ORCID, ROR, Crossref, and DataCite). This interdisciplinary Dagstuhl Seminar “Open Scholarly Information Systems: Status Quo, Challenges, Opportunities” (25381) was the first of its kind to bring together practitioners from this ecosystem, as well as researchers investigating related questions or relying on these systems in their own research. It provided a unique opportunity for dialogue, sharing insights, building new networks, and fostering collaboration.

Seminar September 14–19, 2025 – <https://www.dagstuhl.de/25381>

2012 ACM Subject Classification Information systems → Digital libraries and archives; Information systems → Information retrieval; Computing methodologies → Machine learning

Keywords and phrases artificial intelligence, knowledge graphs, open infrastructures, scholarly big data, scholarly information systems, semantic search

Digital Object Identifier 10.4230/DagRep.15.9.38

1 Executive Summary

Hannah Bast (Universität Freiburg, DE)

Guillaume Cabanac (University of Toulouse, FR)

Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT)

Jian Wu (Old Dominion University – Norfolk, US)

License  Creative Commons BY 4.0 International license
© Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu

This report presents the outcomes and strategic next steps derived from Dagstuhl Seminar “Open Scholarly Information Systems: Status Quo, Challenges, Opportunities” (25381), held in September 2025. The seminar brought together an international group of experts to address the critical challenges facing Open Scholarly Infrastructure (OSI), the evolution of Scholarly Knowledge Graphs (SKGs), and the transformative impact of Agentic AI on the research lifecycle.

* Editor / Organizer



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Open Scholarly Information Systems: Status Quo, Challenges, Opportunities, *Dagstuhl Reports*, Vol. 15, Issue 9, pp. 38–57

Editors: Hannah Bast, Guillaume Cabanac, Paolo Manghi, Jian Wu, and Marcel R. Ackermann



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Purpose and Context

The primary objective of the seminar was to foster collaboration among diverse infrastructure “owners” and stakeholders to ensure the long-term sustainability and interoperability of the systems that support global research. Participants engaged in high-level plenary discussions and focused working groups to tackle technical, ethical, and economic hurdles within the scholarly ecosystem.

Key Themes and Discussion Pillars

The results detailed in this report are organised around several core pillars:

- *Sustainability and Digital Sovereignty*: A central theme was the urgent need for nations to retain control over research outputs to avoid dependency on commercial monopolies. The group explored the Barcelona Declaration as a framework for promoting open research information and discussed strategies to convince institutions to dedicate a portion of their budgets to open infrastructure.
- *Metadata Excellence and Interoperability*: Discussions focused on harmonising metadata across platforms like DBLP, Wikidata, and OpenReview. The “COMET” approach was proposed to align decentralised metadata enrichment efforts, reducing redundancy and enhancing data quality.
- *The Rise of Agentic AI*: The seminar examined how autonomous AI agents might reshape discovery, writing, and peer review. A critical concern was maintaining human agency and accountability to safeguard scientific integrity, even as AI accelerates baseline tasks like replication.
- *Reforming Research Assessment*: Participants challenged the current “perverse system” of publication metrics and rankings. The consensus emphasized that research quality cannot be measured by numbers alone and that data providers should focus on providing comprehensive data rather than automated rankings.

Strategic Objectives

The outcomes represent a collective effort to move from “building” to “using” and “sustaining” open systems. Key outputs include a draft manifesto on the economic and social value of OSI, technical roadmaps for migrating services like Scholia to QLever, and position papers on the future of human-AI collaboration in science. This document serves as a record of these discussions and a roadmap for the community to ensure that scholarly information remains a transparent, trustworthy, and shared global resource.

2 Table of Contents

Executive Summary

<i>Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu</i>	38
--	----

Overview of Talks

Curating the DBLP Computer Science Bibliography <i>Marcel R. Ackermann</i>	42
Citations in the world's largest encyclopedia – and their future <i>Phoebe Ayers and Lydia Pintscher</i>	42
Narrative Information Access in Digital Libraries <i>Wolf-Tilo Balke</i>	42
Knowledge Graphs and QLever <i>Hannah Bast</i>	43
Zenodo and InvenioRDM: Cross-domain digital repositories for the long tail of research <i>Martin Fenner</i>	43
FAIR Digital Objects as Scholarly Infrastructure Metadata Middleware <i>Carole Goble</i>	44
Please meet AI, our dear new colleague: Are we becoming obsolete? <i>Iryna Gurevych</i>	44
Research Fast and Slow <i>Min-Yen Kan</i>	44
Overview of COnnecting REpositories (CORE) <i>Petr Knuth</i>	45
Openness aspects of scholarly information systems important for adoption, transparency and interoperability <i>Bianca Kramer</i>	45
Scholarly Knowledge Graphs: No community, no fun <i>Paolo Manghi</i>	45
Collaborative curation of bibliographic metadata <i>Daniel Mietchen</i>	46
Google Dataset Search <i>Natasha Noy</i>	47
OpenCitations and its new IT infrastructure <i>Mario Petrella</i>	47
The AIDA Dashboard <i>Angelo Salatino</i>	47
Navigating Scholarly Knowledge <i>Tilahun Abedissa Taffa</i>	48
TIB AI Assistant for Research <i>Sahar Vahdati</i>	48

Digital Science & the Research Data Ecosystem
Kathryn Weber-Boer 48

CiteSeerX and NDLTD
Jian Wu 49

Challenges and opportunities in arXiv
Ramin Zabih 49

Working groups

Theme: Metadata Excellence and Interoperability
Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu 50

Theme: Reforming Research Assessment
Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu 51

Theme: Sustainability and Digital Sovereignty
Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu 53

Theme: The Rise of Agentic AI
Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu 54

Participants 57

3 Overview of Talks

3.1 Curating the DBLP Computer Science Bibliography

Marcel R. Ackermann (Schloss Dagstuhl – Trier, DE)

License  Creative Commons BY 4.0 International license
© Marcel R. Ackermann

The DBLP Computer Science Bibliography is one of the most comprehensive scholarly metadata collections and knowledge graphs in computer science. Correctly identifying the true named entities in publication metadata is one of the biggest challenges in maintaining such a scholarly information system. In this talk, we discuss the unique approach of the dblp computer science bibliography to this curation problem, and how we set up a system that uses automated heuristics based on domain knowledge, while keeping the human editors in the loop.

3.2 Citations in the world's largest encyclopedia – and their future

Phoebe Ayers (MIT Libraries – Cambridge, US) and Lydia Pintscher (Wikimedia – Germany, DE)

License  Creative Commons BY 4.0 International license
© Phoebe Ayers and Lydia Pintscher

Wikipedia encompasses millions of articles in hundreds of language editions, written collaboratively by people who do not know each other. Rather than relying on known expert editors, Wikipedia's reliability rests on its citations to other reliable sources. Taken together, this body of references represents a huge, human curated and open collection of reference metadata and bibliographies on every topic imaginable. But what do these references in Wikipedia actually include, how do editors decide what to include, and what do we know about them? And what could citations be, and what could we know about them? We'll show what citations in Wikipedia look like now, and will talk about some ideas for a more structured approach to citations, "Wikicite", that has been in development in the Wikimedia community for several years. We will also discuss Wikidata and bibliographic metadata in Wikidata as well as the current efforts to scale Wikidata, especially in the area of bibliographic data.

3.3 Narrative Information Access in Digital Libraries

Wolf-Tilo Balke (TU Braunschweig, DE)

License  Creative Commons BY 4.0 International license
© Wolf-Tilo Balke

From early on, narratives have been used as an essential means to convey information and knowledge in a form that is close to human communication and sense making. Moreover, references to archetypical narratives, such as David vs. Goliath or Don Quijote, can also transport a set of connotations beyond the actual story allowing for a framing of information. Facing today's flood of data, data-driven narratives are thus an ideal way to make complex

topics comprehensible, to make sense of certain events, or to assess the plausibility of given narratives. However, these features are rarely used in Digital Libraries today. In particular, most of the current work on narratives is limited to representing structural properties such as story or plot graphs, event chains, or representations of entities and events without exploiting the deeper meaning of narratives. In this talk, we will explore narratives in the sense of logical overlays over heterogeneous knowledge repositories in Digital Libraries, such as knowledge graphs, linked open data sources, document collections, or even concrete datasets. In its simplest form, a narrative then is a directed graph consisting of entities, events, and literals as nodes. Narrative edges describe the flow of the modeled events, i.e. on the one hand the semantic interaction between events and entities and on the other hand the respective types of interaction by suitable edge labels (e.g., in the causal or temporal sense). Essential for the expressive power of this overlay model is that edges of a narrative must always be bound against underlying knowledge repositories. In particular, this allows the plausibility of each edge to be evaluated against a given set of trusted repositories. Of course, this also means that the information in the underlying repositories needs to be carefully extracted with respect to classical dimensions of data quality, such as correctness, completeness, or validity.

3.4 Knowledge Graphs and QLever

Hannah Bast (Universität Freiburg, DE)

License  Creative Commons BY 4.0 International license
© Hannah Bast

I will give an introduction to knowledge graphs, RDF, and SPARQL, with many examples and demos. I hope that a lively discussion will ensue.

3.5 Zenodo and InvenioRDM: Cross-domain digital repositories for the long tail of research

Martin Fenner (Front Matter – Münster, DE)

License  Creative Commons BY 4.0 International license
© Martin Fenner

InvenioRDM is an open source repository platform started at CERN that not only powers the Zenodo repository, but also an increasing number of other digital repositories developed and hosted by the InvenioRDM community. InvenioRDM integrates with DataCite, ORCID, ROR, Crossref, OpenAIRE and other Scholarly Information Systems and is a good example of the challenges and opportunities in building scholarly information systems.

3.6 FAIR Digital Objects as Scholarly Infrastructure Metadata Middleware

Carole Goble (University of Manchester, GB)

License  Creative Commons BY 4.0 International license
© Carole Goble

The scientific enterprise depends on finding, exchanging, understanding, validating, reproducing, integrating, reusing, archiving, and citing research entities across a dispersed community of researchers and an ecosystem of research and scholarly communication platforms. We are not just talking about papers, but also the mix of data, software, protocols, models and so on that are research outputs that together form the compendium of knowledge and the means for reuse and reproducibility of experimental outcomes. What is the unit of Knowledge Exchange for research? If actionable metadata needs to accompany our research entities (data and software) on their journeys to enable FAIR (Findable, Accessible, Interoperable, Reusable) research, how do we do that? If metadata is “a love note to the future”, then what’s the envelope? RO-Crates, that’s what! Platform independent standards-based metadata framework for bundling resources with context into citable reproducible packages.

3.7 Please meet AI, our dear new colleague: Are we becoming obsolete?

Iryna Gurevych (TU Darmstadt, DE)

License  Creative Commons BY 4.0 International license
© Iryna Gurevych

How can AI and LLMs facilitate the work of scientists in different stages of the research process? Can technology even make scientists obsolete? The role of AI and Large Language Models (LLMs) in science as the target application domain has recently been rapidly growing. This includes assessing the impact of scientific work, facilitating writing and revising manuscripts as well as intelligent support for manuscript quality assessment, peer-review and scientific discussions. The talk will illustrate such methods and models using several tasks from the scientific domain. We argue that while AI and LLMs can effectively support and augment specific steps of the research process, expert-AI collaboration may be a more promising mode for complex research tasks.

3.8 Research Fast and Slow

Min-Yen Kan (National University of Singapore, SG)

License  Creative Commons BY 4.0 International license
© Min-Yen Kan

I will argue that while today’s AI-driven, toolkit-rich environment enables rapid, incremental “fast” research – boosting citations and confidence – it also risks short-termism, stress, and crowding around the same problems. Drawing on Kahneman’s Thinking, Fast and Slow and Friedman’s “age of acceleration,” I will highlight the importance of cultivating “slow” research: asking the right questions, seeking interdisciplinary perspectives, tolerating inconvenience, and developing long-term vision.

3.9 Overview of COnnecting REpositories (CORE)

Petr Knuth (The Open University – Milton Keynes, GB)

License  Creative Commons BY 4.0 International license
© Petr Knuth

CORE (Connecting Repositories) is an open scholarly infrastructure that indexes millions of open access research papers and metadata from repositories and journals worldwide. Its goal is to improve the discoverability and reuse of research outputs and support machine access to scholarly content in line with open access and open science principles. This talk will provide an overview of CORE and its services for repositories, including compliance monitoring, metadata validation, and tools to improve interoperability and discoverability. It will also present research by the Big Scientific Data and Text Analytics Group (BSDTAG), showcasing recent innovations such as CORE-GPT, a system for trustworthy question answering over scholarly literature; SDG: Classify, which maps research papers to UN Sustainable Development Goals; and SoFAIR, which addresses reproducibility and research software management. Finally, the talk will discuss how CORE enables external research and innovation in areas such as training large language models, plagiarism detection, library discovery, and the construction of scholarly graphs, fostering a globally connected and machine-readable open research ecosystem.

3.10 Openness aspects of scholarly information systems important for adoption, transparency and interoperability

Bianca Kramer (Sesame Open Science – Utrecht, NL)

License  Creative Commons BY 4.0 International license
© Bianca Kramer

In this presentation, I will discuss various aspects of openness, and their relevance for development and maintenance of scholarly information systems, as well as for decisions around adoption and (financial) support. Taking from various existing frameworks (including the Principles of Open Scholarly Infrastructure, the CoARA working group Open Infrastructures for Responsible Research Assessment and the Barcelona Declaration on Open Research Information) I will put a number of openness criteria up for discussion, and solicit opinions from seminar participants how relevant each criterion is for them, and why. In the second part of the presentation, we will collectively look at a number of infrastructures represented in the room and how they are positioned against a limited set of openness criteria.

3.11 Scholarly Knowledge Graphs: No community, no fun

Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT)

License  Creative Commons BY 4.0 International license
© Paolo Manghi

Scholarly Knowledge Graphs (SKGs) currently play a crucial role in enabling open data as a means of fostering transparent and trustworthy research assessment. However, the emergence of Open Science and its associated publishing workflows introduce a range of challenges

that cannot be addressed by SKGs alone. Drawing from the experience of developing the OpenAIRE Graph, this presentation will highlight some of these challenges and illustrate how the interaction among researchers, publishing data sources and venues, and SKGs creates a synergistic ecosystem. In this ecosystem, SKGs serve as essential tools to enhance and streamline scholarly workflows.

3.12 Collaborative curation of bibliographic metadata

Daniel Mietchen (FIZ Karlsruhe – Berlin, DE)

License  Creative Commons BY 4.0 International license
© Daniel Mietchen

This talk sketches out several facets of a collaborative approach to the curation of scholarly information. It starts by outlining two use cases from the Wikipedia ecosystem: The first is that when media files from open-access sources are uploaded to Wikimedia Commons, then bibliographic metadata about the open-access source need to be curated with each file on Wikimedia Commons. The second is that if Wikipedia articles are translated from one language to another, then the bibliographic metadata of the cited references will have to be curated for each language version of the article. Such use cases would benefit from a shared repository of bibliographic information, which is what the WikiCite initiative is about that uses Wikidata as a hub to collect and curate bibliographic information about scholarly sources cited on Wikimedia projects. Next, the Scholia tool is introduced, which uses the information from Wikidata to provides some dozens of scholarly profile types (e.g. author, work, topic, organization, award, taxon, gene or location). For each profile type, it generates a HTML page that contains a number of visualizations that are based on preformulated Wikidata SPARQL queries parametrized by the entity to be profiled. This way, Scholia users do not need to have SPARQL knowledge but if they do, they can use it to finetune or modify the queries to explore the data further. Scholia has a range of features that are relevant to the discussion of open scholarly infrastructures. For instance, many Scholia profiles have an “improve data” button that leads to a dedicated curation page listing known data quality issues, e.g. author name strings yet to be disambiguated, along with links to tooling that helps address these issues. Some of its visualizations provide impact proxies. The profile for an individual work include a panel listing statements supported by a given work, or Wikipedia mentions of it, or information about relevant retractions. The profile for this very Dagstuhl Seminar includes co-author networks and a list of recent publications by seminar participants. Next, the platform zbMATH Open is introduced, which covers the mathematical research literature in a way similar to how the DBLP platform covers the computer science literature. zbMATH Open profiles authors and publications and, through its sister platform swMATH, also curates links between mathematical publications and the associated software. Finally, nanopublications are introduced as miniature knowledge graphs that can represent individual units of curation. They can be queried, shared and aggregated and are functionally similar to individual wiki edits or git commits, yet they are structured and machine actionable in a way that makes them compatible with larger knowledge graphs, for which they might serve as a vehicle to exchange curation information.

3.13 Google Dataset Search

Natasha Noy (Google – Mountain View, US)

License  Creative Commons BY 4.0 International license
© Natasha Noy

Datasets constitute a key part of scientific knowledge: they are described in publications and referenced in them; datasets themselves connect with one another in intricate ways. I will build on our experience with Google Dataset Search and discuss the different types of these connections, focusing on datasets. I will discuss the complementary roles of community and tooling to build the connections.

3.14 OpenCitations and its new IT infrastructure

Mario Petrella (University of Bologna, IT)

License  Creative Commons BY 4.0 International license
© Mario Petrella

This presentation will examine OpenCitations IT infrastructure, showing how our Kubernetes-based microservices architecture effectively implements the 'Living Will' principle by allowing complete replication of our open scholarly services in just a few hours. This technical approach ensures OpenCitations resilience and ability to continue operating independently of its current supporting institutions.

3.15 The AIDA Dashboard

Angelo Salatino (The Open University – Milton Keynes, GB)

License  Creative Commons BY 4.0 International license
© Angelo Salatino

The AIDA Dashboard is a powerful system for analysing and comparing scientific journals and conferences. By leveraging a large-scale Knowledge Graph that integrates billions of data points about research from multiple sources, it offers unique, sophisticated analytics and rankings. This tool gives researchers and other key stakeholders unique insights into the evolution of different venues and helps them make crucial decisions. In my talk, I will focus on two main challenges. These include a lack of conference data information outside of the Computer Science field (where we are so lucky to have DBLP), as well as the absence of fine-grained representation of research topics in several disciplines. These are critical to enable appropriate categorisation and management of information beyond computer science.

3.16 Navigating Scholarly Knowledge

Tilahun Abedissa Taffa (Leuphana Universität Lüneburg, DE)

License  Creative Commons BY 4.0 International license
© Tilahun Abedissa Taffa

Scholarly knowledge navigation tools are essential for finding relevant information from bibliographic data sources. This presentation focuses on complementary advances: ASK-DBLP, a user-in-the-loop KGQA system that converts natural language questions to editable SPARQL and mitigates schema drift; and a RAG-based scholarly QA layer on the NFDI4DataScience Gateway that enables conversational, federated access to diverse scientific databases.

3.17 TIB AI Assistant for Research

Sahar Vahdati (TIB – Hannover, DE)

License  Creative Commons BY 4.0 International license
© Sahar Vahdati

The TIB AI Assistant for Research is an AI-supported, modular assistant designed to help scholars across the research lifecycle from ideation and question formulation to literature exploration and structured synthesis. It combines large language models with semantic/vector search and knowledge graphs (notably the Open Research Knowledge Graph), enabling natural-language querying and producing concise, structured outputs such as tabular evidence summaries. Emphasizing openness and reproducibility, components like ORKG Ask are open-source and operate over tens of millions of open-access publications. Together, these capabilities accelerate rigorous, transparent discovery while aligning with TIB's open-science mission.

3.18 Digital Science & the Research Data Ecosystem

Kathryn Weber-Boer (Digital Science – London, GB)

License  Creative Commons BY 4.0 International license
© Kathryn Weber-Boer

This presentation will describe a range of Digital Science software solutions (with a focus primarily on Figshare and Dimensions), in the context of their relationships to each other, the way they connect to, draw on, and—in turn—feed other data sources. Figshare is an open repository where users share datasets, figures, and other research-related output. Dimensions is a platform where open systems (Crossref, DataCite, ORCID) are used, their data are transformed and combined with proprietary data (IFI CLAIMS and ReadCube), and the results are then fed back into the open space (e.g., ROR, Covid-19 Dataset). The presentation will address the delicate relationship between industry and the open science movement, and the vision that drives our understanding of that relationship.

3.19 CiteSeerX and NDLTD

Jian Wu (Old Dominion University – Norfolk, US)

License © Creative Commons BY 4.0 International license
© Jian Wu

CiteSeerX is one of the World's earliest digital library search engines serving 15 million academic documents crawled over the Web. CiteSeerX uses AI techniques for crawling, information classification, and information extraction. Over the past 20+ years, CiteSeerX has overcome a lot of challenges in an academic setting and is still indexed by Google Scholar. In this talk, we show the upcoming challenges of this legacy and discuss an emergent question about how to make CiteSeerX sustainable and continue to serve the academic communities.

The Networked Digital Library of Theses and Dissertations (NDLTD) is an international organization that promotes the creation, dissemination, and preservation of electronic theses and dissertations (ETDs). Established in 1987, NDLTD has organized an annual symposium since 1998, on ETD-related research and development. It also maintains a union catalog containing metadata for more than 6.5 million ETDs worldwide and publishes a dedicated journal. In recent years, NDLTD has faced significant infrastructure and organizational challenges and is seeking collaborative efforts, particularly with the United States Electronic Thesis and Dissertation Association (USETDA), to address these issues.

3.20 Challenges and opportunities in arXiv

Ramin Zabih (Cornell Tech – New York, US)

License © Creative Commons BY 4.0 International license
© Ramin Zabih

arXiv has been a key open access resource since 1991. It has become a dominant force in many areas of science, particularly in computer science, math and physics. Since 2002 arXiv has been hosted at Cornell University, which has provided a stable home and has helped it develop a broad funding base. While from the outside arXiv looks like a smoothly functioning machine, it actually is facing a wide range of difficult issues. I will cover some of the main challenges and opportunities, and give some insights into how the arXiv model has proven sustainable for over 3 decades.

4 Working groups

4.1 Theme: Metadata Excellence and Interoperability

Hannah Bast (Universität Freiburg, DE), Guillaume Cabanac (University of Toulouse, FR), Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT), and Jian Wu (Old Dominion University – Norfolk, US)

License  Creative Commons BY 4.0 International license
© Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu

The outcomes of the seminar regarding Metadata Excellence and Interoperability focus on harmonizing fragmented systems, adopting advanced querying technologies, and establishing collaborative models to reduce redundant manual work. The participants highlight that while many organizations enrich metadata, these efforts are often independent, leading to inconsistencies across platforms.

4.1.1 The COMET Approach to Collaborative Metadata

A primary outcome was the discussion of the *COMET (Collaborative Metadata)* model, which aims to align decentralized curation efforts with one another and with authoritative sources.

- *Reducing Redundancy:* By promoting shared principles and technical standards, COMET seeks to facilitate the deduplication of effort, enabling more efficient data exchange.
- *Cross-System Comparison:* Next steps involve using Knowledge Graphs (KGs) to automatically detect agreement or disagreement in metadata, such as author attribution, across different systems.
- *Affiliation Parsing:* The group plans to compare results on affiliation parsing from sources like arXiv to test the interoperability of current curation workflows.

4.1.2 Transitioning to Knowledge Graphs and SPARQL

The seminar identified a significant opportunity to move from “scripts and JSON blobs” to formal Scholarly Knowledge Graphs to simplify complex data tasks.

- *OpenReview and DBLP Integration:* Participants agreed that OpenReview’s current REST API and MongoDB storage could be complemented by an RDF/SPARQL pipeline. This would allow conference organisers to run complex queries – such as finding authors with specific paper counts across multiple venues – without writing custom scripts.
- *Wikidata Scalability with QLever:* To address the scalability issues of the Wikidata Query Service (currently using Blazegraph), the group successfully tested QLever as a new backend. Scholia was identified as the primary test case for this migration, with promising results in rewriting SPARQL queries to be more standard-compliant.

4.1.3 Harmonising Subject Classification and Ontologies

The participants note a “perverse trend” where the proliferation of new standards actually hinders interoperability.

- *Reducing Standard Proliferation:* The strategic goal is to decrease the number of generic subject area classifications used by open infrastructures.

- *Integration of Ontologies:* Plans are underway to link DBLP records with the Computer Science Ontology (CSO), a taxonomy of 14,000 research topics, to provide better subject-level discoverability for venues and authors.
- *Automatic Classification:* Experiments are being conducted using the Leiden/OpenAlex model on Hugging Face to automatically classify tens of thousands of records within the InvenioRDM platform.

4.1.4 Advanced Researcher Disambiguation

Addressing the “author identity” challenge, the seminar explored strategies to combine automated and manual curation.

- *Evidence-Based Identification:* Disambiguation efforts will focus on combining various indicators, including co-authorship, subject area, and email domains.
- *Chinese Name Disambiguation:* A specific initiative involves using AMiner to improve the disambiguation of Chinese names, as traditional ASCII transliteration often loses critical variations present in Hanzi characters.
- *Researcher Autonomy:* The group emphasised that systems must respect researchers who choose to maintain multiple ORCID profiles for reasons of identity change or personal safety in different political contexts.

4.1.5 Expanding Metadata Scope

The seminar also looked toward “giving back” to the ecosystem by capturing information currently missing from major graphs:

- *Software and Acknowledgments:* Efforts are starting to index software mentions in DBLP, syncing workflows with zbMATH and swMATH. Additionally, nanopublications are being explored as a mechanism to track acknowledgments for infrastructures like Dagstuhl itself.
- *Fake and Spam Metadata:* A dedicated working group discussed methods for identifying and flagging “fake” or “spam” conferences to prevent them from polluting open metadata records.
- *Scalability of Metadata Access:* Moving beyond API calls, the group discussed the necessity of using regular data dumps in formats like Parquet or JSONL for large-scale analysis, as API-based retrieval does not scale for multi-million record collections.

4.2 Theme: Reforming Research Assessment

Hannah Bast (Universität Freiburg, DE), Guillaume Cabanac (University of Toulouse, FR), Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT), and Jian Wu (Old Dominion University – Norfolk, US)

License © Creative Commons BY 4.0 International license
© Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu

The results from Dagstuhl Seminar 25381 regarding Reforming Research Assessment center on a collective agreement to move away from what participants described as a “perverse system” of publication metrics and automated rankings. The consensus among the experts was that scientific quality and impact are multi-dimensional and cannot be reduced to a single numerical score.

4.2.1 Rejection of Automated Rankings

The seminar participants, including representatives from major data providers, reached a consensus on several fundamental points regarding the presentation of research data:

- *Quality vs. Numbers:* There was unanimous agreement that the quality of scientific research cannot be measured by numbers alone.
- *The Responsibility of Infrastructure:* Databases such as DBLP, Dimensions, and Open-Review should not automatically deliver publication statistics for the purpose of ranking or estimating expertise. Instead, these platforms should “live their values” by choosing how – and if – they display indicators.
- *Data over Rankings:* The primary mission of data providers should be to curate and provide comprehensive, transparent data, leaving the interpretation and derivation of statistics to the users for specific, contextual purposes.

4.2.2 Strategic “Refusal” of Metrics

A significant outcome was the discussion on how specific infrastructures intentionally limit the metrics they provide to prevent misuse:

- *Withholding H-indices:* It was noted that platforms like Dimensions and DBLP have explicitly chosen to refuse to provide H-indices or Impact Factors (IF) on their sites, even though the raw data to calculate them is available.
- *Contextualisation:* By default, search results in these systems are often sorted by date rather than “impact,” forcing users to engage with the research itself rather than a pre-sorted list of “top” authors.
- *Transparency and “Data Cards”:* To combat the “metric frenzy,” the group proposed the use of “data cards” to make the assumptions behind research data clear, ensuring that any derived rankings are transparent and accountable.

4.2.3 Addressing the “Perverse” Incentives

The working groups explored why flawed metrics persist and how to transition to a more equitable system:

- *The Seniority Problem:* Participants discussed how the current system is maintained by tenured, senior colleagues who were selected by these very metrics. There is an urgent need for these “overloaded seniors” to drive change.
- *Information Overload:* It was acknowledged that rankings often persist because they are an “easy” solution to manage information overload and time pressure during hiring or grant review processes.
- *Multi-Criteria Assessment:* The group advocated for moving toward narrative-based assessments where researchers “name their top five” contributions, and evaluations include broader criteria such as teaching, collaboration styles, and institutional contributions.

4.2.4 Alignment with International Initiatives

The seminar’s outcomes are designed to support and fill “actionable gaps” in existing international frameworks:

- *COARA and DORA:* The discussions were closely aligned with the Coalition for Advancing Research Assessment (CoARA) and the San Francisco Declaration on Research Assessment (DORA).

- *Institutional Advocacy*: A key next step is for participants to convince their own institutions to adopt formal strategies for handling publication metrics and to stop over-relying on automated counting.
- *The Barcelona Declaration*: Participants discussed using the Barcelona Declaration to promote open research information as the standard for decision-making, highlighting non-traditional outputs like datasets instead of just highly cited articles.

4.2.5 Summary

The seminar concluded that while statistics can be useful tools for specific goals, they must be contextualised and non-automated to prevent the “rich get richer” or Matthew effect¹ and to safeguard the diversity and integrity of the scholarly ecosystem.

4.3 Theme: Sustainability and Digital Sovereignty

Hannah Bast (Universität Freiburg, DE), Guillaume Cabanac (University of Toulouse, FR), Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT), and Jian Wu (Old Dominion University – Norfolk, US)

License © Creative Commons BY 4.0 International license
© Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu

The outcomes of the Dagstuhl Seminar regarding *Sustainability and Digital Sovereignty* emphasize a critical shift from viewing Open Scholarly Infrastructure (OSI) as a purely academic concern to framing it as a vital national and economic asset. Participants concluded that for open systems to survive, they must align with broader societal priorities, such as *security, AI productivity, and democratic resilience*.

In summary, the seminar participants agreed that the scholarly community must move toward co-ownership and collective responsibility to ensure that research information remains a transparent and shared global resource.

4.3.1 The Economic Imperative for Sustainability

The participants highlight that several established OSIs, including CiteSeerX, NDLTD, and CORE, currently face significant sustainability hurdles due to financial, administrative, and human resource constraints. To address these challenges, the seminar proposed several strategic shifts:

- *Dedicated Institutional Funding*: A primary recommendation is for universities and libraries to dedicate a specific proportion of their budgets – suggested at 5% – to support open scholarly infrastructure, rather than exclusively funding commercial publishers.
- *Aligning with National Priorities*: Future funding proposals should move beyond justifying needs from a researcher’s perspective. Instead, they should demonstrate how OSIs support national priorities such as AI development and digital security.
- *Diverse Business Models*: While philanthropic funding has supported entities like Open-Alex, the group noted that such support is not a guaranteed long-term solution. Instead, they explored membership models (e.g., CORE), donations (e.g., Wikipedia), and ERIC (European Research Infrastructure Consortia) which utilize national subscriptions.

¹ https://en.wikipedia.org/wiki/Matthew_effect

- *Consolidation and Efficiency:* To ensure long-term preservation, the group suggested that under-supported OSIs might need to merge or integrate. For example, work is already underway to integrate CiteSeerX into the Internet Archive Scholar.

4.3.2 Digital Sovereignty and the Knowledge Economy

Digital sovereignty was identified as the need for nations to retain long-term control over their research outputs to avoid dependency on commercial monopolies. Key discussion points included:

- *The Threat of Commercial Monopolies:* Participants noted that institutions are increasingly reliant on commercial Current Research Information Systems (CRIS), which are often non-interoperable and require high annual fees.
- *OSI as the Backbone of AI:* The participants argue that we cannot have trustworthy, evidence-based AI without OSI, as open infrastructures provide the essential reliable corpora needed for AI models.
- *Resilience and Security:* Distributed and interoperable infrastructures provide protection against sanctions, cyber threats, and misinformation. The group emphasized that “OSI saves lives,” citing the role of transparent, reproducible science in developing vaccines and combating medical misinformation.

4.3.3 Strategic Actions and the Barcelona Declaration

The seminar leveraged the *Barcelona Declaration on Open Research Information* as a central framework for these efforts. Notable progress and next steps include:

- *Expanding Signatories:* Efforts are underway to add organizations such as CWI and NWO-I as signatories to the Barcelona Declaration.
- *International Advocacy:* Representatives will take these arguments to high-level policy discussions with the OECD, G7 Open Science Working Group (OSWG), and UKRI.
- *Manifesto for Decision Makers:* The group produced a draft manifesto detailing the economic and social value of OSI to convince decision makers of the return on investment (RoI) provided by open systems.
- *The COMET Approach:* To reduce redundancy and enhance sustainability, the COMET model was proposed to align decentralized metadata curation efforts, ensuring that infrastructure “owners” collaborate rather than duplicate work.

4.4 Theme: The Rise of Agentic AI

Hannah Bast (Universität Freiburg, DE), Guillaume Cabanac (University of Toulouse, FR), Paolo Manghi (Institute of Information Science and Technologies – CNR – Pisa, IT), and Jian Wu (Old Dominion University – Norfolk, US)

License © Creative Commons BY 4.0 International license
© Hannah Bast, Guillaume Cabanac, Paolo Manghi, and Jian Wu

The outcomes regarding The Rise of Agentic AI from Dagstuhl Seminar 25381 focus on the transformative potential of autonomous agents across the research lifecycle, while simultaneously highlighting the risks to scientific integrity and human agency. Participants explored how these systems might evolve from simple assistants to semi-autonomous collaborators.

4.4.1 Levels of Agent Autonomy

A core framework discussed during the seminar defines five levels of autonomy for AI agents in a research context, based on the degree of human involvement:

- *L1 (Operator)*: The human directs all decisions and actions.
- *L2 (Collaborator)*: Human and agent plan and execute tasks together.
- *L3 (Consultant)*: The agent takes the lead but consults the human for preferences or expertise.
- *L4 (Approver)*: The agent operates autonomously and only seeks human intervention for risky or pre-specified scenarios.
- *L5 (Observer)*: The agent operates with full autonomy under human monitoring.

4.4.2 Integration with Scholarly Infrastructure

The group debated whether Open Scholarly Information Systems (OSIS) must reinvent themselves to remain relevant in an AI-dominated landscape.

- *Model Context Protocol (MCP)*: There is a strategic discussion regarding whether OSIS services should be offered as MCP services. This would allow AI agents to use scholarly databases as reliable corpora for tasks like Retrieval-Augmented Generation (RAG).
- *The Shift to Bot Users*: Predictions suggest that by 2030, search engines and scholarly systems will be used primarily by AI bots rather than human researchers. This necessitates new protocols for data access and perhaps “micro-credits” to recoup the costs of high-frequency bot requests.

4.4.3 Impact on the Research Lifecycle

The seminar examined specific case scenarios where agentic AI is already active or emerging:

- *Replication of Baselines*: AI agents are being tested for “boring” research tasks, such as replicating baseline experiments. While systems like Sakana’s “AI Scientist” can produce full manuscripts rapidly and at low cost (USD 6–15 per paper), evaluations showed significant flaws, including poor novelty assessments, coding errors in 42% of experiments, and hallucinated results.
- *Peer Review*: While LLMs can support human reviewers by providing evidence or checking submission guidelines, they currently fail to detect faulty research logic. Studies indicate that flaws in internal consistency often go unnoticed by fully automatic review generators.
- *Literature Reviews*: AI can achieve high recall in identifying citations but lacks the nuanced judgment and background knowledge of an expert reviewer.

4.4.4 Safeguarding Human Agency and Education

A recurring concern throughout the sessions was the potential for deskilling among researchers.

- *Critical Thinking*: Participants emphasized that human agency, critical thinking, and accountability must remain at the heart of scholarship.
- *Reforming the PhD*: The role of PhD training may need to shift from technical proficiency to high-level analytical skills and “human-AI co-construction”.
- *The Future of H-AI Collaboration*: The consensus was that the future lies in Human-AI collaboration, where humans remain responsible for the integrity and ethical purpose of the research while leveraging AI to tackle increasingly complex problems.

4.4.5 Strategic Next Steps

The results of these discussions are being compiled into a position paper (currently being drafted on Overleaf) titled “oAsIs”, focusing on the role of open agentic scholarly information systems. Ongoing work will also investigate preference-based cooperation (HAICo2) to ensure AI agents adapt to human expertise and institutional guidelines.

Participants

- Marcel R. Ackermann
Schloss Dagstuhl – Trier, DE
- Phoebe Ayers
MIT Libraries – Cambridge, US
- Wolf-Tilo Balke
TU Braunschweig, DE
- Hannah Bast
Universität Freiburg, DE
- Guillaume Cabanac
University of Toulouse, FR
- A. Seza Dogruöz
Ghent University, BE
- Martin Fenner
Front Matter – Münster, DE
- Ingo Frommholz
MODUL Universität Wien, AT
- Carole Goble
University of Manchester, GB
- Iryna Gurevych
TU Darmstadt, DE
- Lynda Hardman
CWI – Amsterdam, NL
- Holger Hermanns
Universität des Saarlandes –
Saarbrücken, DE
- Min-Yen Kan
National University of
Singapore, SG
- Petr Knoth
The Open University –
Milton Keynes, GB
- Bianca Kramer
Sesame Open Science –
Utrecht, NL
- Christin Kreutz
THM – Gießen, DE
- Cyril Labbé
University of Grenoble, FR
- Michael Ley
Schloss Dagstuhl – Trier, DE
- Paolo Manghi
Institute of Information Science
and Technologies – CNR –
Pisa, IT
- Philipp Mayr
GESIS – Köln, DE
- Daniel Mietchen
FIZ Karlsruhe – Berlin, DE
- Carlos Daniel Mondragón
Chapa
OpenReview – Cambridge, US
- Patrick Neises
Schloss Dagstuhl – Trier, DE
- Natasha Noy
Google – Mountain View, US
- Mario Petrella
University of Bologna, IT
- Lydia Pintscher
Wikimedia – Germany, DE
- Ruzica Piskac
Yale University – New Haven, US
- Rüdiger Reischuk
Universität zu Lübeck, DE
- Angelo Salatino
The Open University –
Milton Keynes, GB
- Ralf Schenkel
Universität Trier, DE
- Ansgar Scherp
Universität Ulm, DE
- Raimund Seidel
Universität des Saarlandes –
Saarbrücken, DE
- Tilahun Abedissa Taffa
Leuphana Universität
Lüneburg, DE
- Sahar Vahdati
TIB – Hannover, DE
- Nees Jan van Eck
Leiden University, NL
- Ruijie Wang
Universität Zürich, CH
- Kathryn Weber-Boer
Digital Science – London, GB
- Jian Wu
Old Dominion University –
Norfolk, US
- Ramin Zabih
Cornell Tech – New York, US

