Report on the 48th European Conference on Information Retrieval (ECIR 2026)

Avishek Anand, Julián Urbano
Delft University of Technology
Delft, The Netherlands
\{avishek.anand,j.urbano\}@tudelft.nl

Zhaochun Ren, Suzan Verberne
Leiden University
Leiden, The Netherlands
\{z.ren,s.verberne\}@liacs.leidenuniv.nl

Maria Heuss
University of Amsterdam
Amsterdam, The Netherlands
m.c.heuss@uva.nl

Bhaskar Mitra
Independent Researcher
Tiohtià:ke / Mooniyang / Montréal, Canada
bhaskar.mitra@acm.org

Maik Fröbe, Jan Heinrich Merker
Friedrich-Schiller-Universität Jena
Jena, Germany
\{maik.froebe,heinrich.merker\}@uni-jena.de

Harrisen Scells
University of Tübingen
Tübingen, Germany
harrisen.scells@uni-tuebingen.de

Abstract

The 48th European Conference on Information Retrieval (ECIR 2026) was held in Delft, the Netherlands, during 29 March – 2 April 2026. The conference took place at Lijm & Cultuur, a former glue factory centrally located between the historic city centre and the TU Delft campus; the tutorial day was hosted on the TU Delft campus. ECIR 2026 was an exclusively in-person event, with authors unable to travel supported by proxy presenters, and attracted a record number of 476 participants from academia and industry making it the most attended ECIR till date. Building on recent editions, ECIR 2026 introduced several novelties, most notably the promotion of the IR-for-Good track to a parallel core track of the main conference anchored around the theme *What is IR-for-Good?*, the addition of a resource paper track, and the integration of IRRJ paper presentations into the main programme.

Date: 29 March – 2 April 2026.

Website: https://ecir2026.eu.

1 Overview

ECIR 2026 took place in Delft, The Netherlands, and was a joint organisation of general chairs from Delft and Leiden. We had to diverge from our original bid, which was based on a venue in The Hague. Between the bid and the start of the organisation, however, prices went up substantially, and we had to move to a cheaper location: Lijm & Cultuur, a former glue factory in Delft, centrally

Role	Name(s)
General Chairs	Avishek Anand, Zhaochun Ren, Suzan Verberne
Full paper Track Chairs	Adam Jatowt, Ricardo Campos, Yanyan Lan
Short paper Track Chairs	Sean MacAvaney, Mohammad Aliannejadi, Christine Bauer
IR-for-Good Track Chairs	Bhaskar Mitra, Maria Heuss
Reproducibility Track Chairs	Andrew Yates, Venktesh Viswanathan
Resource Track Chairs	Petra Galuščáková, Panagiotis Eustratiadis
Demonstration Track Chairs	Yue Feng, Sandipan Sikdar
Industry Day Chairs	Benjamin Piwowarski, Vinay Setty
Doctoral Consortium Chairs	Aldo Lipani, Julián Urbano
CLEF Labs Chairs	Julia Maria Struß, Sean MacAvaney
Workshop Chairs	Negar Arabzadeh, Franco Maria Nardini
Tutorial Chairs	Faegheh Hasibi, Manish Gupta
Collab-a-thon Chairs	Maik Fröbe, Jan Heinrich Merker, Harrisen Scells
Best Paper Awards Chair	Craig Macdonald
Sponsorship Chairs	Ujwal Gadiraju, Edgar Meij
Proceedings Chairs	Alisa Rieger, Johannes Kiesel
Local Organisation Chairs	Masoud Mansoury, Nan Bai
Publicity Chairs	Yifei Yuan, David Graus

Table 1. ECIR 2026 organising team.

located between the city centre and the campus of Delft University of Technology. The tutorial day on Sunday took place at the Delft University campus.

As ECIR chairs, we had the privilege to build on a long history of successful conferences. From the recent previous editions of ECIR we adopted a number of successful initiatives: presentations of CLEF Tracks, the Collab-a-thon, a reproducibility paper track, and a Findings track for the full papers. We expanded on that with a few improvements and extensions: We included the IR for Good Track as a parallel track in the main conference (instead of a separate day), we added a session with presentations of IRRJ papers (collaboration with the open access journal), and we added a resource paper track.

We set some specific ambitions for the program. We did our very best for a high-quality review process, by having meta-reviewers (SPC members) for each track, including the smaller ones, and assigning 4 reviewers to each submission to make sure we had 3 good-quality reviews. We explicitly aimed for a relatively high acceptance rate for short papers, to give ongoing work a chance to be presented at the conference. This required from the short paper chairs that they held discussions with the PC and SPC members to make sure that papers were not rejected too lightly. The acceptance rates of all tracks are listed in Table 2.

We decided that ECIR was an exclusively in-person conference. Authors who were not able to travel could appoint a proxy presenter for their paper. We connected some authors to proxy

Category	Submissions	Accepted	Acceptance rate
Full papers	200	46	23.0%
+Findings		+10	+5.0%
Short papers	108	37	34.3%
Resource papers	35	17	48.6%
Reproducibility papers	17	9	52.9%
IR4Good papers	56	19	33.9%
Industry Day	19	10	52.6%
Demo papers	24	13	54.2%
Doctoral Consortium	20	14	70.0%
Workshops	13	11	84.6%
Tutorials	10	7	70.0%

Table 2. ECIR 2026 submission and acceptance statistics. Findings papers were selected from the pool of rejected full-paper submissions, yielding a combined Main + Findings acceptance rate of 28.0%.

presenter volunteers. This all worked smoothly and we did not have any no-shows in the main conference.

2 Scientific program

The ECIR 2026 program [Campos et al., 2026a,b,c,d] comprised papers across multiple tracks, reflecting the breadth and diversity of contemporary information retrieval research. The final program included 46 full papers, 10 findings papers, 9 reproducibility papers, 17 resource papers, 19 IR-for-Good papers, as well as 37 short papers, 13 demonstration papers, 14 doctoral consortium papers, 10 industry track papers, 16 invited CLEF papers, and 7 invited FDIA papers (see Table 2). In addition, one session in the program was devoted to papers published in the first two issues of the Information Retrieval Research journal (IRRJ).[1] This initiative reflects ECIR's ongoing effort to strengthen the connection between archival journal research and the conference community. We focus on the three notable new elements in the ECIR 2026 program – The Resource Track, the IRRJ session, and the IR4Good Track(covered in Section 4).

2.1 Resource Track

ECIR 2026 introduced a dedicated Resource Track for the first time. The motivation for the track was to give first-class visibility, peer review, and citation credit to the kinds of contributions on which much modern IR research depends: datasets, test collections, benchmarks, annotation schemes, evaluation toolkits, open-source software, pretrained models, and other reusable artefacts. Although papers describing such resources are frequently submitted to the main track, they

[1] https://irrj.org/

are often disadvantaged when reviewed under the same criteria as algorithmic or systems contributions. A separate track allowed reviewers to be calibrated for the specific notions of novelty, scope, documentation, reusability, and long-term maintenance that resource papers require.

The track was chaired by Petra Galuščáková and Panagiotis Eustratiadis. Submissions were reviewed by a dedicated programme committee with experience in building and maintaining IR resources. Reviewers were instructed to evaluate not only the technical quality of a resource but also its likely impact on the community: the rigour of its construction, the breadth of tasks it can support, the clarity of its documentation, and the credibility of the authors' plan for its hosting and long-term availability.

The track received **35 submissions**, of which **17 were accepted**. The accepted papers were presented in regular sessions alongside the full and reproducibility paper sessions of the main conference. We feel that the track filled a real gap in the IR publication landscape, and we recommend that it become a permanent fixture of ECIR going forward.

2.2 Information Retrieval Research Journal (IRRJ) Session

The *Information Retrieval Research Journal* (IRRJ)[2] is a recently launched, fully open-access journal dedicated to information retrieval research. The journal complements the existing conference-driven publication culture of the IR community by providing an archival, peer-reviewed venue with no article-processing charges for authors and no access barriers for readers. As a community-owned journal, IRRJ aims to support the long-form work – comprehensive evaluations, reproducibility studies, systematic surveys, and significant extensions of conference papers – that the conference page-limit format does not always accommodate well.

To strengthen the connection between the ECIR community and the journal, ECIR 2026 introduced for the first time a dedicated IRRJ session in the main conference programme. Authors of papers published in the first two issues of IRRJ were invited to present their work to the conference audience. Each paper received a regular oral slot, with the same audience and Q&A format as other main-conference sessions, so that journal contributions were given equal scientific visibility alongside conference papers. The session attracted strong attendance and stimulated lively discussion, in particular around papers reporting larger-scale empirical studies than is typical at ECIR.

In addition, a special issue of IRRJ is currently in preparation to host extended versions of selected ECIR 2026 papers from three different tracks. Authors of these papers were encouraged to develop the conference work into a fuller treatment – adding experiments, comparisons, analyses, or extensions that the conference page limit precluded – to be peer reviewed and archived in the journal. We view the IRRJ session and the planned special issue as a first step towards a more sustainable, multiple-track publication model for the IR community, in which conferences and the community-owned journal reinforce rather than compete with each other. We strongly encourage future editions of ECIR to continue and expand this collaboration.

[2] https://irrj.org/

Figure 1. ECIR 2026 keynote speakers. From the left: Katja Hofmann (Microsoft Research Cambridge), Madeleine Daepp (Public Democracy America and Microsoft Research), and Tetsuya Sakai (Waseda University), winner of the 2026 Keith van Rijsbergen Award.

3 Keynote talks

The keynote speakers included Katja Hofmann (Microsoft Research Cambridge) and Madeleine I. G. Daepp (Public Democracy America and Microsoft Research), as well as the Keith van Rijsbergen Award winner, Tetsuya Sakai. We warmly congratulate Dr. Sakai on this well-deserved recognition.

3.1 Main Conference Keynote Presentation: Katja Hofmann

Katja Hofmann is a Partner Research Manager at Microsoft Research Cambridge, where she co-leads the People-Centric AI research area. Her work focuses on generative AI, interactive media, and game intelligence, combining advances in machine learning with human-computer interaction, design, and social science. With her team she aims to create AI systems that empower people through collaboration, creativity, and play – unlocking new forms of interaction and addressing complex real-world challenges.

Generative AI for Connection and Creativity

In the rapidly evolving landscape of AI models and tooling innovation, a central question guides our work: How do we build technologies that genuinely support people in pursuing what matters to them? This keynote reflects on recent research from the People-Centric AI area at Microsoft Research Cambridge, highlighting progress toward AI systems that understand human context, foster equitable access and representation, and act as generative "connective tissue" for collaboration and creativity. Through examples spanning world-model research, inclusive data practices, and emerging tools that help people express ideas, explore possibilities, and work together more effectively, the talk outlines best practices for developing AI that amplifies human potential and strengthens the connections between us.

3.2 IR-for-Good Keynote Presentation: Madeleine I. G. Daepp

Dr. Madeleine I. G. Daepp is Visiting Director of Civic Innovation at Public Democracy America, where she advises nonprofits on navigating the AI transition, and senior researcher at Microsoft Research studying the impacts of generative AI on global democracies. A civic technologist with a Ph.D. from MIT, her work centers on collaborating with communities and leveraging novel technologies to solve problems in shared public spaces. Dr. Daepp has written invited perspectives on the societal impacts of AI for The Economist and Nature Computational Science (forthcoming) and published in peer-reviewed venues across computer science, public health, and urban planning.

Gemini Hegemony

For some people AI is a megaphone. But many people around the world are experiencing its harms, from cultural erasure to expanding scam and influence operations to scaled surveillance states. In this keynote, I draw on fieldwork conducted in Taiwan and India during the 2024 mega-election year to characterize key emerging challenges that generative AI poses for democracies around the world. Based on interviews with more than 70 frontlines actors in both defender and creator roles, I identify a set of recurring challenges as generative AI systems are deployed at scale. I then show how these findings shaped my work to expand red-teaming efforts across languages and cultures, and map the challenges I observed to key insights and opportunities in IR research. We are in a critical moment to avoid repeating social media's missteps, and the choices made by builders – from development to interaction design – will be key in ensuring a democratic future with AI.

3.3 Keith van Rijsbergen Award 2026: Tetsuya Sakai

The KvR Award recognises outstanding contributions to the field of information retrieval, with a particular emphasis on scientific excellence, innovation, and lasting impact on the research community. The award is sponsored by The University of Glasgow School of Computing Science and BCS IRSG. The committee for the award was composed of a representative of the University of Glasgow, a representative of BCS IRSG, one of the full paper chairs, and the three general chairs.

The KvR Award 2026 has been awarded to Tetsuya Sakai, Professor in the Department of Computer Science and Engineering at Waseda University and Dean of the Center for Data Science at Waseda University. As part of the award, Tetsuya gave a keynote at ECIR 2026.

Did You Know? Eleven Nerdy Trivia about Evaluation Measures

Evaluation measures are often treated as fixed mathematical tools for benchmarking retrieval systems. Yet behind every metric lies a set of assumptions about users: what they value, how patiently they browse, when they stop searching, and what kinds of errors matter most. In this keynote, Tetsuya Sakai revisits more than two decades of research on information retrieval evaluation through a collection of "nerdy trivia" about classic and modern evaluation metrics. This talk shows how many ideas considered modern today, e.g., cascade browsing models, gain-based evaluation, and utility-oriented metrics, can already be found in much earlier work, while also highlighting limitations of widely used measures. Through examples spanning graded relevance,

diversity, fairness, calibration, and nugget-based evaluation, this keynote argues that evaluation metrics should not merely generate scores, but should faithfully model user experience and distinguish meaningful system differences.

Tetsuya Sakai is internationally recognised for his influential work on evaluation metrics, test collections, and experimental methodologies. His research has fundamentally shaped how information retrieval systems are assessed and compared, contributing to more robust and reproducible research practices across the discipline.

4 IR-for-Good track

The IR-for-Good Track was introduced at ECIR for the first time in 2024 to encourage increased focus on societal issues in IR and to make progress towards a responsible IR agenda incorporating interdisciplinary perspectives McDonald et al. [2024]. For the first couple of years, the track ran on the final day of the conference in parallel to other workshops. This year, the third edition of this special track continued at ECIR and was promoted to a core conference track that ran parallel to other conference paper presentation sessions.

This year, the track received a total of 56 complete submissions. Each submission was reviewed by three to four PC members and meta-reviewed by one senior PC member. Subsequently, the track chairs jointly discussed the review feedback and meta-review recommendation for each submission and made the final accept/reject decisions. 19 submissions were eventually accepted, setting the acceptance rate for the track this year to 34.5%.

This year, the PC consisted of 54 reviewers and 10 meta-reviewers (senior PCs). We want to express our deepest appreciation to the PC for their thoughtful feedback on each submission, their active engagement in the review discussion stage, and for adapting to new requirements introduced this year, such as the theory of change section.

4.1 Theme: What is IR for Good?

The IR-for-Good track aims to affect positive societal impact through IR research. To realize material societal benefits, the community must have a clear and shared understanding of what constitutes societal good, and appraise our progress against those goals. However, the topic of societal good is largely undertheorized in IR [Mitra, 2026]. If our normative position is that all of IR should contribute positively to society, then it further raises questions about why we need a special track on this topic separate from the main conference and what we intend to achieve with the track [Mitra and Heuss, 2025]. At ECIR 2026, the IR-for-Good track chairs foregrounded these questions by adopting the prompt "*What is IR-for-Good?*" as this year's theme for the conference track.

In line with this theme, we took several initiatives this year to drive towards more clarity on these questions and make room for the IR community to critically reflect on and develop its own shared understanding of these topics. These conversations are already happening in different parts of the IR community—*e.g.*, the call from the fourth Strategic Workshop in Information Retrieval in Lorne (SWIRL) to recenter IR research on societal, democratic, and emancipatory values [Trippas et al., 2025]. We build on these ongoing conversations to adopt the following operative definition of IR-for-Good for this year's conference:

IR-for-Good refers to IR research and practices that contribute towards realizing more equitable, emancipatory, and sustainable futures.

We consider this definition to be neither fixed nor complete. It is finally up to the IR community to iterate, extend, and further explicate their societal goals over time. But we hope this definition provides a reasonable starting point for those conversations. Based on this definition, we updated the scope and the list of relevant topics for the track in our call-for-papers[3] that we detail in Section 4.1.1. This year, we also encouraged authors to explicitly include a section on their theories of change, as described in Section 4.1.2.

Lastly, on the question of the role of the IR-for-Good special track, we concluded that the track should be a space where we can explore, experiment with, and develop new community practices and norms to promote more societally-beneficial IR research. We should subsequently contribute the identified best practices back to the broader IR community in an effort to ensure that all of IR research is IR-for-Good.

4.1.1 Scope and Topics

This year, we defined the scope of IR-for-Good track to include IR research that:

1. Explicitly concerns with new research directions and system design to achieve specific societally beneficial outcomes,
2. Develops new fairness, privacy, transparency, accessibility, sustainability, and other similar societally-motivated interventions, and/or
3. Identifies and critiques the ways in which existing IR methods and systems and how we do IR research may contribute to systemic harm or impede social progress.

Within the above specified scope, we invited contributions to the track that explore new positions, critiques, tools, methods, resources, and interventions for IR-for-Good.

Starting from our operative definition of "IR-for-Good", we enumerated relevant topics of interest for this track to include **how IR intersects with and/or can support**: (i) Accessibility and disability justice, (ii) art, culture, and representation, (iii) crisis and disaster management, (iv) decolonization and racial justice, (v) emancipation, anti-oppression, and social justice, (vi) gender and sexuality justice, (vii) informed citizenry, democracy, and collective decision making, (viii) law and restorative justice, (ix) literacy and knowledge production, (x) privacy and dignity, (xi) public health and community health, (xii) social, political, and economic equity, (xiii) sustainability and environmental justice, and (xiv) worker rights and labor movements. We made an intentional choice to center these topics on societally-beneficial outcomes (*e.g.*, equity, emancipation, justice, and sustainability) rather than on the approaches that may help us progress towards those outcomes (*e.g.*, procedural fairness, interpretability, and transparency). While we also welcomed submissions focusing on different approaches in the track, our motivation for centering outcomes over approaches was to encourage exploration of a broader space of diverse sociotechnical methods as well as to hold ourselves accountable to the ultimate goal of affecting positive societal impact.

[3] https://ecir2026.eu/calls/call-for-ir-for-good-papers

4.1.2 Theories of Change

If we want to encourage more critical scholarly discourse within the IR community on how specific research directions may contribute towards desired societal outcomes, we need to make our theories of change explicit in our scholarship. With this motivation, this year, we encouraged every submission to the IR-for-Good track that propose new IR tools, methods, resources, and interventions to explicitly include a separate section elaborating how their work contributes towards desired societal outcomes. Position papers and critiques were exempted from this requirement as these arguments should anyway be a core contribution of those submissions.

We recommended that the "Theory of Change" section should explicitly state:

1. What is the identified societal need / problem, and how are the core contributions from this current work expected to address them?
2. What preconditions are necessary or what assumptions need to hold for this work to have its desired effect, and how likely are they to hold true in practice?
3. What are possible negative externalities of this approach and is it plausible that this may lead to new or different harms?

Authors were encouraged to include any additional discussions that they may deem relevant in this section. Contributions focusing on algorithmic bias, fairness, transparency, interpretability, explainability, trustworthiness, misinformation, disinformation, hate speech, replicability, transferability, robustness, uncertainty, security, ethics, and other related topics were also required to explicitly articulate how the work contributes towards positive societal outcomes and not implicitly assume that all research on these topics contribute to societal good. As a corollary, certain IR topics that may not have historically been seen as societally focused (e.g., designing distributed information access platforms or developing more effective ranking models without the use of user behavior data) were also welcome to the track if they appropriately argued how the work is likely to contribute to societal good, *e.g.*, by making platforms more robust to authoritarian capture or disincentivizing mass ubiquitous user surveillance, respectively.

We strongly emphasized that this section should not be an afterthought. Instead it should be a critical part of the key motivations for the work and as important as any other core sections of the respective papers. We encouraged authors and reviewers to critically engage with this section while acknowledging the real uncertainty of how any well-intentioned research may impact society in practice. We cautioned authors against inflating their claims of societal impact and instead encouraged them to rigorously deliberate on their sociotechnical assumptions and to thoroughly enumerate the necessary preconditions for the work to have its desired impact and potential negative externalities.

Having explicit theories of change in our publications further opens up the opportunity for future scholarship to analyze, critique, validate, and improve upon these theories of change. It also creates the possibility for scholars from non-IR disciplines to engage with and critically analyze the emerging theories of change within the IR community. And we hope that over time this practice also encourages IR researchers to more actively reach beyond their disciplinary boundaries to work with other scholars, experts, practitioners, policymakers, civil rights advocates, activists, and movement organizers pushing for social justice and sustainability.

Figure 2. IR-for-Good panel discussion. From the left: Maria Heuss (Track Chair and panel moderator), Bhaskar Mitra (Track Chair), Sanne Vrijenhoek, Madeleine I. G. Daepp (Conference Keynote Speaker), and Djoerd Hiemstra. Photograph taken by Arjen P. de Vries.

4.2 Program

The program consisted of three paper sessions, a keynote by Madeleine I. G. Daepp (see Section 3.2), and a panel session distributed over a couple of days of the main conference. We allocated a mix of algorithmic-focused papers and papers that incorporated expertise from other disciplines to each of the paper sessions to encourage more interdisciplinary interactions within the track.

We kicked off the panel session with an invited talk by Djoerd Hiemstra titled "*OpenWebSearch.eu: Towards a shared infrastructure for assembling web search engines*" that challenged the audience to imagine what a transparent, well-regulated, community-based web search engine would look like and to envision a future where a search engine is "assembled" from parts provided by many different institutions, based on public standards.

The panel discussion was chaired by Maria Heuss who was joined by Bhaskar Mitra, Sanne Vrijenhoek, Madeleine I. G. Daepp, and Djoerd Hiemstra on stage. The panel discussion foregrounded this year's theme of the track and invited the panelists—as well as the audience—to engage in thoughtful deliberation over what constitutes IR-for-Good. As such, the discussion centered around what societal good the IR community should aspire for and the need for rigorous theorization of the concept to drive clarity within the community. This was followed by a discussion on how we should do this research differently, the current research practices that we should consider changing, and new practices that we should adopt. Following that, there was also a discussion about areas where we are currently under-invested in IR. Other conversation topics revolved around the OpenWebSearch.eu project and the opportunity to expand that vision towards global public service IR platforms. Panelists also highlighted the need to exercise cau-

tion in normalizing the technologists-as-liberator framing. The panel discussion ended with some concluding thoughts on next steps and where the community should go from here.

5 Awards

5.1 Best paper awards

The awards chair, Craig Macdonald, composed a best paper committee consisting of: Krisztian Balog, Matthias Hagen, Ben He, Donald Metzler, and Gabriella Pasi. For each of the three tracks, the track chairs proposed a shortlist. The committee members ranked the papers independently and anonymously. Ties were broken through discussion. We handed out best paper awards for the full paper track, short paper track, and IR-for-Good track, and a honourable mention for each of these tracks:

- **Best full paper:** *LLM-based Listwise Reranking under the Effect of Positional Bias.* Jingfen Qiao, Jin Huang, Xinyu Ma, Shuaiqiang Wang, Dawei Yin, Evangelos Kanoulas, and Andrew Yates.
- **Full paper honourable mention:** *When Reducing Representations Improves Performance.* Andrea Pasin, Guglielmo Faggioli, Nicola Ferro, Raffaele Perego, and Nicola Tonellotto.
- **Best short paper:** ReFormeR: *Learning and Applying Explicit Query Reformulation Patterns.* Amin Bigdeli, Mert Incesu, Negar Arabzadeh, Charles L. A. Clarke, and Ebrahim Bagheri.
- **Short paper honourable mention:** *Query Performance Prediction using a Child-focused Definition of Relevance.* Hrishita Chakrabarti and Maria Soledad Pera.
- **Best IR-for-Good paper:** *Measuring Political Stance and Consistency in Large Language Models.* Salah Feras Alali, Mohammad Nashat Maasfeh, Mucahid Kutlu, and Şaban Kardaş.
- **IR-for-Good paper honourable mention: Bias in Book Recommendation:** *A Case Study on the Danish Public Libraries.* Savvina Daniil, Søren Højlund Mollerup, and Laura Hollink.

5.2 Reviewer awards

Reviewer awards (PC members and SPC members) were announced during the opening ceremony. To acknowledge the effort of reviewers who contributed high-quality reviews and discussion, we awarded all reviewers who were nominated by an SPC member, and all SPC members who were nominated by the track chairs.

- **PC member awards** went to: *Adam Roegiest, Alberto Veneri, Aleksandr Petrov, Andrea Papenmeier, Antonela Tommasel, Benjamin Kille, Chandan Kumar, Despoina Chatzakou, Elena Garcia-Morato, Ferdinand Schlatt, Fiana Raiber, Francesco Busolin, Giuseppe Spillo, Harrie Oosterhuis, Junichi Tatemura, Ketan Thakkar, Kidist Amde Mekonnen, Ladislav Peska, Maik Fröbe, Marco Alessio, Maria Vlachou, Mrinal Ahlawat, Riccardo Lunardi, Richard A. A. Jonker, Ro Encarnacion, Saber Zerhoudi, Shubham Chatterjee, Takehiro Ya-*

mamoto, Theresia Veronika Rampisela, Thibault Formal, Timo Breuer, Weronika Łajewska, Wiradee Imrattanatrai.

- **Senior PC members awards** went to: *Alessandro Fabris, Arjen de Vries, Carsten Eickhoff, Graham McDonald, Iadh Ounis, Jaime Arguello, Javier Parapar, Jose G Moreno, Julia Neidhardt, Laura Dietz, Ludovico Boratto, Makoto P. Kato, Marc Najork, Martin Halvey, Matthias Hagen, Michael Färber, Raffaele Perego, Sean MacAvaney, Tamer Elsayed, Udo Kruschwitz.*

6 Tutorials and workshops

6.1 Tutorials

ECIR 2026 hosted seven tutorials, all of which were well attended and received positive feedback:

Practical, Efficient, In-Memory Inverted Indexes.[4] Joel Mackenzie, Sean MacAvaney, Antonio Mallia and Michal Siedlaczek.

Conversational Search: Foundations, Large Language Models, and Agents.[5] Chuan Meng, Fengran Mo, Mohammad Aliannejadi, Jeff Dalton and Jian-Yun Nie.

Tutorial on Mechanistic Interpretability.[6] Catherine Chen, Maria Heuss and Carsten Eickhoff.

Reasoning for IR & IR for Reasoning.[7] Mohanna Hoveyda, Panagiotis Eustratiadis, Arjen de Vries and Maarten de Rijke.

Neural Lexical Search with Learned Sparse Retrieval.[8] Andrew Yates, Carlos Lassance, Cosimo Rulli, Eugene Yang, Sean MacAvaney, Siddharth Singh, Thong Nguyen and Yibin Lei.

Economic Perspectives on Fairness in Information Retrieval.[9] Chen Xu, Clara Rus, Yuanna Liu, Marleen de Jonge, Jun Xu and Maarten de Rijke.

Uncertainty Quantification for Large Language Models. Maxim Panov, Artem Shelmanov, Roman Vashurin, Artem Vazhentsev, Ekaterina Fadeeva, Lyudmila Rvanova and Timothy Baldwin.

6.2 Workshops

ECIR hosted eleven workshops, all of which were well attended:

Full-day Workshops:

- The 9th International Workshop on Narrative Extraction from Texts (Text2Story'26).[10] Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, Marina Litvak.

[4] https://pisa-engine.github.io/ecir-2026.html
[5] https://convsearch.github.io/
[6] https://mech-interp-tutorial-ecir26.github.io/
[7] https://reasoning-for-ir.github.io/
[8] https://lsr-tutorial.github.io/
[9] https://economic-fairness-ir.github.io/
[10] https://text2story26.inesctec.pt

- The Third Search Futures Workshop.[11] Leif Azzopardi, Charles Clarke, Claudia Hauff, Yubin Kim, Adam Roegiest, Johanne R. Trippas, Zhaochun Ren, Saber Zerhoudi.
- 4th International Workshop on Geographic Information Extraction from Texts (GeoExT 2026).[12] Xuke Hu, Ludovic Moncla, Jens Kersten, Anna Kruspe, Inhye Kong.
- INFUSE: The First Workshop on Information Access in Uncertainty Scenarios.[13] Alisa Rieger, Ran Yu, Amir Ebrahimi Fard, Nicolas Mattis, Johanne R. Trippas.

Half-day Workshops:

- The First Workshop on Information Retrieval for Accountability and Integrity (IRAI).[14] Chung-Chi Chen, Juyeon Kang, Anaïs Lhuissier, Dittaya Wanvarie, Min-Yuh Day, Hiroya Takamura, Yohei Seki.
- The Second Workshop on Scholarly Information Access (SCOLIA 2026).[15] Ingo Frommholz, Christin Katharina Kreutz, Philipp Mayr, Guillaume Cabanac.
- Reducing Online Misinformation through Credible Information Retrieval (ROMCIR 2026).[16] Marcos Fernández-Pichel, Marinella Petrocchi, Kevin Roitero, Marco Viviani.
- The Third International Workshop on Open Web Search (WOWS).[17] Laura Caspari, Maik Fröbe, Sebastian Heineking, Michael Granitzer, Gijs Hendriksen, Djoerd Hiemstra, Martin Potthast, Arjen de Vries, Saber Zerhoudi.
- LIR: The First Workshop on Late Interaction and Multi-Vector Retrieval.[18] Benjamin Clavié, Xianming Li, Antoine Chaffin, Omar Khattab, Tom Aarsen, Manuel Faysse, Jing Li.
- Workshop on Conversational Search for Complex Information Needs.[19] Roxana Petcu, Mert Yazan, Mohanna Hoveyda, Jirui Qi, Maarten de Rijke.
- Synthetic Data and Simulation Synergy for Information Retrieval.[20] Manel Slokom, Alejandro Bellogin, Mayank Singh, Andrea Barraza-Urbina.

7 Industry day

The ECIR 2026 Industry Day, chaired by Vinay Setty, was held as a full-day programme at the main conference venue and was very well attended. The ECIR 2026 Industry Day highlighted the growing role of IR in agentic AI systems-intelligent agents that plan, act, and retrieve. The focus of the day was generative AI, covering the design, integration, and evaluation of IR components in industrial contexts, and the call for contributions was scoped accordingly. The day was structured around four invited industry keynote talks, interleaved with sessions of peer-reviewed industry track paper presentations, so that talks from leading practitioners framed and contextu-

[11] https://searchfutures.github.io
[12] https://geo-ext.github.io/GeoExT2026
[13] https://sites.google.com/view/infuse-workshop
[14] https://nlpfin.github.io/sites/ECIR2026.html
[15] https://sites.google.com/view/bir-ws/scolia-2026
[16] https://romcir.disco.unimib.it
[17] https://opensearchfoundation.org/events-osf/wows2026
[18] https://lateinteraction.com
[19] https://convsearch-complex-info-needs.github.io
[20] https://synirgy-workshop.github.io

alised the day's research contributions. The four invited speakers were Louis Milliken (Perplexity), who presented *Diffusion-Pretrained Dense and Contextual Embeddings* for web-scale retrieval; Johannes Hoffart (SAP), who outlined a vision for business foundation models in *Unlocking Relational Business Data with In-Context Learning, Moving Toward Operational Knowledge*; Antoine Chaffin (LightOn), who covered open-source encoder and multi-vector retrieval; and Gonzalo Fiz Pontiveros and Roger Zhe Li (Huawei), who shared industrial experience with one-step generative recommendation and multi-modal semantic IDs in *Generative Recommendation & Multi-Modal Semantic IDs: Industrial Insights from Huawei*.

In addition to the invited talks, the day featured ten peer-reviewed industry track papers presented across four sessions. The accepted papers spanned a broad set of industrial IR topics, including query understanding and search suggestions for alphanumeric queries, synthetic data for personalization and cold-start recommendation, and even neural IR over cultural heritage collections. The combination of strong invited talks with concrete industry case studies made for a programme that consistently sparked discussion between researchers and practitioners, and the day closed with broad consensus that the Industry Day remains one of the most valuable components of ECIR.

8 Doctoral consortium

The Doctoral Consortium (DC) aimed to provide PhD students with constructive feedback from senior members of the IR community. After initially receiving only 7 submissions, the deadline was extended and senior colleagues were contacted to encourage submissions from their students. The final number of submissions increased to 20. PC members were then invited based on topic affinity, with the expectation that they would also serve as mentors during the event. Assignments were made according to bids and topic fit, and each submission received two independent reviews. In the end, 14 students were accepted, with decisions guided mainly by the expected benefit they would gain from participating in the DC. To our knowledge, this was the largest ECIR Doctoral Consortium to date.

The mentoring process was designed to maximize relevance and openness. Additional mentors were invited to cover for PC members who could not attend ECIR, and final assignments followed four principles: each student should have two mentors; no mentor in the consortium should come from any accepted student's institution, helping create a safe space for discussion; mentors should come from a different country than the student, making the in-person exchange more valuable; and priority should be given to mentors who had reviewed the submission or bid positively. In total, 22 senior community members participated as PC members and/or mentors. The consortium day featured presentations by all accepted students and two one-hour private mentoring sessions. It was held in TU Delft's Senaatzaal, the university's formal PhD defense room, making it an especially fitting venue for an event centered on doctoral research.

9 Collab-a-thon

The ECIR Collab-a-thon aims to foster new collaborations yielding new collaborative research projects and relationships with like-minded researchers. The first Collab-a-thon ran at ECIR in

Figure 3. Impressions from ECIR 2026. From the left: the main conference venue at Lijm & Cultuur, a former glue factory in Delft; the conference dinner in the historic Arsenaal in the Delft city centre; and the welcome reception that opened the conference.

2024 [MacAvaney et al., 2024]. Historically, each Collab-a-thon session is organized around one ECIR-related research topic so that researchers interested in the topic can meet in an informal way and brainstorm ideas for possible collaborations. Like previous iterations of the Collab-a-thon, we aimed to engage all members of the research community by providing a structured, welcoming event.

This year, the Collab-a-thon session was conducted at the end of each day during or after the last session. We used the sessions of the day to seed discussion groups. A show of hands about each topic was used to split participants into groups. One member of each group was designated as note taker to record the discussions. To ensure that all ECIR attendees that are open for new collaborations know which Collab-a-thon sessions are interesting for them, we coordinated with the organisers of ECIR to hand out leaflets containing instructions on how to participate in the Collab-a-thon. The leaflets also contained places where participants could collect session-specific stickers as a way to incentivise participation. Filling out the entire sticker book allowed participants to collect a special 'ECIR26 Collab-a-thon' sticker.

We held three Collab-a-thon sessions that each had between one and three discussion groups. Overall, 5 discussion groups brainstormed ideas, on (1) Evaluation, (2) LLM-as-a-Judge, (3) Multimodal Search, (4) Data and Infrastructure, and (5) synthetic data and simulations.

10 Social events

The social programme of ECIR 2026 was designed to make the most of Delft's historic city centre and the character of the conference venue. The main conference was hosted at *Lijm & Cultuur*, a former glue factory between the TU Delft campus and the city centre that has been repurposed as a cultural venue. The conference opened with a welcome reception at the Sunday venue (TU Delft Aula) in the evening, providing attendees with a first opportunity to reconnect with colleagues and to meet new members of the IR community in a relaxed setting before the start of the technical programme.

The conference dinner was held at the *Arsenaal*, a historic former armoury in the centre of Delft, within walking distance of the venue and most hotels. Despite a sharp increase in the

Figure 4. Behind and on stage at ECIR 2026. From the left: the organising chairs, the student volunteer team, and the opening ceremony of the conference.

number of guests in the final two weeks before the conference, the event was smoothly served and gave participants the chance to continue discussions over dinner in one of Delft's most distinctive heritage buildings.

11 Conclusion

In all respects, ECIR 2026 was a success. We had a record number of participants (476), many of whom were students and representatives from industry. The industry day was very well attended with talks of high quality. The change of venue from a conference hotel in The Hague to an industrial heritage site in Delft appeared an excellent choice: participants appreciated the 'homey' and hip atmosphere and continuous supply of coffee and tea (see Section 10). The integration of the IR-for-Good track in the main conference programme, the introduction of the resource track, and the integration of IRRJ paper presentations were successful: these sessions were very well attended. A special issue for IRRJ with extended versions of ECIR papers is currently under preparation.

Acknowledgments

The success of ECIR 2026 would not have been possible without the dedication and effort of a large team of volunteers and reviewers. We wish to thank all reviewers and meta-reviewers for their careful evaluations and constructive feedback, which were essential in ensuring the high quality of the conference program. We are also grateful to all of the track, programme, and organisation chairs whose work shaped every aspect of the event; the full organising team is listed in Table 1.

We would like to thank all student volunteers, whose commitment and hard work were instrumental in ensuring a smooth, welcoming, and memorable experience for all participants and attendees. We are particularly grateful to Bart van Leeuwen, the official conference photographer, whose photographs document the spirit of ECIR 2026 throughout this report. We further acknowledge the support of our sponsors and partners, whose contributions helped make the conference possible.

Finally, we thank all authors and participants for their contributions to ECIR 2026 and for continuing to make ECIR a vibrant, inclusive, and forward-looking forum for information retrieval research.

References

Ricardo Campos, Adam Jatowt, Yanyan Lan, Mohammad Aliannejadi, Christine Bauer, Sean MacAvaney, Avishek Anand, Zhaochun Ren, Suzan Verberne, Nan Bai, and Masoud Mansoury, editors. *Advances in Information Retrieval - 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 - April 2, 2026, Proceedings, Part I*, volume 16483 of *Lecture Notes in Computer Science*, 2026a. Springer. ISBN 978-3-032-21288-7. doi: 10.1007/978-3-032-21289-4. URL https://doi.org/10.1007/978-3-032-21289-4.

Ricardo Campos, Adam Jatowt, Yanyan Lan, Mohammad Aliannejadi, Christine Bauer, Sean MacAvaney, Avishek Anand, Zhaochun Ren, Suzan Verberne, Nan Bai, and Masoud Mansoury, editors. *Advances in Information Retrieval - 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 - April 2, 2026, Proceedings, Part II*, volume 16484 of *Lecture Notes in Computer Science*, 2026b. Springer. ISBN 978-3-032-21299-3. doi: 10.1007/978-3-032-21300-6. URL https://doi.org/10.1007/978-3-032-21300-6.

Ricardo Campos, Adam Jatowt, Yanyan Lan, Mohammad Aliannejadi, Christine Bauer, Sean MacAvaney, Avishek Anand, Zhaochun Ren, Suzan Verberne, Nan Bai, and Masoud Mansoury, editors. *Advances in Information Retrieval - 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 - April 2, 2026, Proceedings, Part III*, volume 16485 of *Lecture Notes in Computer Science*, 2026c. Springer. ISBN 978-3-032-21323-5. doi: 10.1007/978-3-032-21324-2. URL https://doi.org/10.1007/978-3-032-21324-2.

Ricardo Campos, Adam Jatowt, Yanyan Lan, Mohammad Aliannejadi, Christine Bauer, Sean MacAvaney, Avishek Anand, Zhaochun Ren, Suzan Verberne, Nan Bai, and Masoud Mansoury, editors. *Advances in Information Retrieval - 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 - April 2, 2026, Proceedings, Part IV*, volume 16486 of *Lecture Notes in Computer Science*, 2026d. Springer. ISBN 978-3-032-21320-4. doi: 10.1007/978-3-032-21321-1. URL https://doi.org/10.1007/978-3-032-21321-1.

Sean MacAvaney, Adam Roegiest, Aldo Lipani, Andrew Parry, Björn Engelmann, Christin Katharina Kreutz, Chuan Meng, Erlend Frayling, Eugene Yang, Ferdinand Schlatt, Guglielmo Faggioli, Harrisen Scells, Iana Atanassova, Jana Friese, Janek Bevendorff, Javier Sanz-Cruzado, Johanne R. Trippas, Kanaad Pathak, Kaustubh D. Dhole, Leif Azzopardi, Maik Fröbe, Marc Bertin, Nishchal Prasad, Saber Zerhoudi, Shuai Wang, Shubham Chatterjee, Thomas Jänich, Udo Kruschwitz, Xi Wang, and Zijun Long. Report on the collab-a-thon at ECIR 2024. *SIGIR Forum*, 58(1):1–11, 2024. doi: 10.1145/3687273.3687287. URL https://doi.org/10.1145/3687273.3687287.

Graham McDonald, Craig Macdonald, and Iadh Ounis. Report on the 46th european conference on information retrieval (ECIR 2024). *SIGIR Forum*, 58(1):1–18, 2024. doi: 10.1145/3687273.3687283. URL https://doi.org/10.1145/3687273.3687283.

Bhaskar Mitra. Towards Critical IR Theories and Practices, 2026. URL https://bhaskar-mitra.github.io/files/critical-ir.pdf.

Bhaskar Mitra and Maria Heuss. What is IR-for-Good?, 2025. URL https://bhaskar-mitra.github.io/posts/2025/09/01/what-is-ir-for-good/.

Johanne R Trippas, J Shane Culpepper, et al. Report from the fourth strategic workshop on Information Retrieval in Lorne (SWIRL 2025). In *ACM SIGIR Forum*, volume 59, page 68, 2025.