

Report on the 34th ACM Conference on Information and Knowledge Management (CIKM 2025)

Carl Yang Emory University USA	Senjuti Basu Roy NJ Inst. Tech. USA	Zhenhui Jessie Li Yunqi Acad. Eng. China	Jaap Kamps University of Amsterdam The Netherlands kamps@uva.nl
Meeyoung Cha MPI-SP, Germany KAIST, South Korea	Chanyoung Park KAIST South Korea	Noseong Park KAIST South Korea	

Abstract

This paper reports on the 34th ACM Conference on Information and Knowledge Management (CIKM 2025), held in Seoul, South Korea, during November 10–14, 2025. The conference was the largest CIKM ever and brought together researchers from all over the world. For those who like numbers: First, we had 1,436 in-person registrations from 46 countries. Second, we received 2,890 submissions in total (including 1,627 full, 615 short, and 317 applied research papers), excluding further workshop submissions. Third, the technical program committee consisted of 2,845 reviewers in total (including 11 chairs, 29 ACs, 493 SPCs, 1,888 PCs, and 452 sub-reviewers), many of whom served on multiple tracks. Fourth, the proceedings include 870 papers totaling 7,121 pages.

Date: 10–14 November 2025.

Website: <https://cikm2025.org>.

1 Introduction

The 34th ACM Conference on Information and Knowledge Management (CIKM 2025) was held in Seoul, South Korea, during November 10–14, 2025. The conference was organized by three general chairs from South Korea and Germany, in cooperation with the ACM’s Specialist Interest Groups on Information Retrieval (SIGIR) and on Hypertext and the Web (SIGWEB). CIKM is a premier ACM conference in the areas of information retrieval, knowledge management, and databases. Since 1992, it has successfully brought together leading researchers and practitioners from these three communities to identify challenging problems in the development of advanced knowledge and information systems and to shape future research directions through the publication of high-quality, applied, and theoretical research findings.

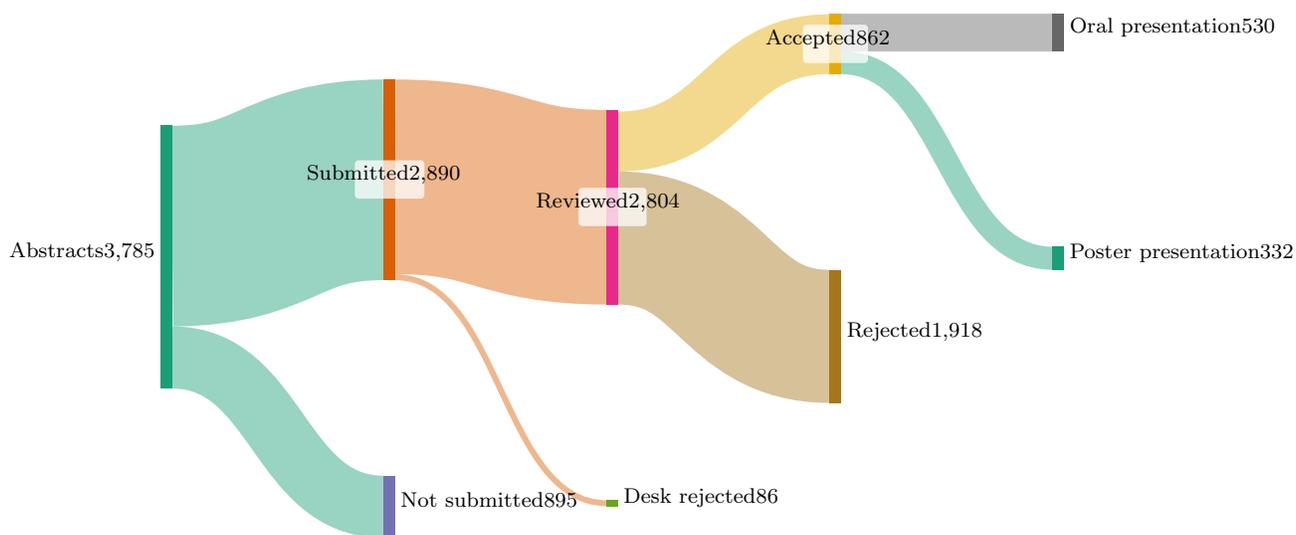


Figure 1. CIKM 2025 submissions.

CIKM 2025 has emerged as the largest edition to date, with a record-breaking number of 2,890 submissions (see Figure 1). CIKM 2025 offers an outstanding technical program, including oral presentation of full research papers and applied research papers, poster sessions for short papers and resource papers, and demonstration sessions. In total, we have received 1,627 submissions of full research papers, 615 short research papers, 317 applied research papers, 118 resource papers, and 95 demo papers. These papers were submitted from 78 countries. For the first time this year, every submitted paper required a mandatory “GenAI Usage Disclosure” section, detailing any use of generative AI tools in research or writing, in accordance with the ACM Authorship Policy.

Organizing CIKM 2025 was a monumental effort by the very large team of organizers shown in Table 1, and countless other volunteers. All full research paper submissions were evaluated by the largest research program committee in CIKM’s history—1,888 PC and 493 SPC members—who conducted rigorous reviews, meta-reviews, and indepth discussions. For the first time, 29 ACs elected from responsive and voluntary SPCs helped the PC Chairs to chase missing reviews, enhance discussions, and suggest final decisions. Final decisions were made with careful consideration of AC recommendations, SPC meta-reviews, PC reviews, and discussions. The full research paper track accepted 443 full papers (27% acceptance rate), whereas the short research paper track accepted 185 papers (31%), the applied research paper track accepted 93 papers (29%), the resource paper track accepted 52 papers (44%), and the demo paper track accepted 37 papers (39%). The overall acceptance rate is 30%. The acceptance rates across tracks reflect CIKM’s commitment to balancing high selectivity with inclusiveness across research, applied, resource, and demo contributions.

The official record of the conference is the ACM DL proceedings (Cha et al., 2025). We are deeply grateful to everyone who made this program possible. First, we thank all authors who submitted their excellent work—the conference could not exist without your contributions. Second, we acknowledge the tireless efforts of our senior program committee members, program committee members, and external reviewers, who devoted significant time and expertise to ensure

Table 1. CIKM 2025 Organization Team

Chairs	Name	Affiliation
General Chairs:	Meeyoung Cha	MPI-SP, Germany & KAIST, South Korea
	Chanyoung Park	KAIST, South Korea
	Noseong Park	KAIST, South Korea
PC Chairs:	Carl Yang	Emory University, USA
	Senjuti Basu Roy	New Jersey Institute of Technology, USA
	Jessie Li	Yunqi Academy of Engineering, China
	Jaap Kamps	University of Amsterdam, Netherlands
Short Paper:	Kijung Shin	KAIST, South Korea
	Bryan Hooi	National University of Singapore, Singapore
	Lifang He	Lehigh University, USA
Applied Research:	Yong Li	Tsinghua University, China
	Julian McAuley	University of California, San Diego, USA
	Yang Song	Kuaishou Technology, China
	Sungchul Kim	Adobe, USA
Resources/Demo:	Jundong Li	University of Virginia, USA
	Cheng-Te Li	National Cheng Kung University, Taiwan
Workshop:	Xiangnan He	University of Science and Technology of China, China
	Qingyun Wu	The Pennsylvania State University, USA
Tutorials:	Matteo Riondato	Amherst College, USA
	Ee-Peng Lim	Singapore Management University, Singapore
	Giulia Preti	CENTAI, Italy
AnalytiCup:	Vachik Dave	Walmart Labs, USA
	Wei Jin	Emory University, USA
PhD Symposium:	Qiaoyu Tan	NYU Shanghai, China
	Isaac Johnson	Wikimedia, USA
Industry Day:	Jingren Zhou	Alibaba Group, China
	Soonmin Bae	Korea Telecom, South Korea
	Xianfeng Tang	Amazon, USA
	Afra Mashhadi	University of Washington Bothell, USA
Publicity:	Diego Saez Trumper	Wikimedia, USA
	Gi-Soo Kim	UNIST, South Korea
Proceedings:	Jian Kang	MBZUAI, UAE
	Eun-Sol Kim	Hanyang University, South Korea
Student Travel:	Dongwon Lee	The Pennsylvania State University, USA
	Dawei Zhou	Virginia Tech, USA
Local:	Yeon-Chang Lee	UNIST, South Korea
	Buru Chang	Korea University, South Korea
	Soo Kyung Kim	Ewha Womans University, South Korea
Web:	Jeongwhan Choi	KAIST, South Korea



Figure 2. Scenes from the CIKM 2025 Banquet in COEX.

a fair and rigorous selection process. We also thank our keynote speakers for enriching the program with their inspiring perspectives.

2 Social Events

CIKM visited South Korea for the very first time, and was held in Seoul, the country’s over 600-year-old capital, and a vibrant hub of economic, technological, and cultural activity at the heart of Asia. The conference was a showcase of Korean hospitality, held at the massive COEX convention and exhibition center in Seoul’s popular Gangnam district. Of particular notice was the view of the Bongeunsa Temple and Bongeun Park across the road during the extensive coffee and lunch breaks.

The conference featured a welcome reception with a buffet and drinks, along with a poster and demo session. The quality of the food and drinks was only surpassed by the quality of the 274 short, resource, and demonstration papers presented in parallel. Given the record size of CIKM 2025, organizing a conference banquet is challenging. The seated banquet in the immense COEX Grand Ballroom was impressive not only for its scale. The banquet was a grand celebration, including culture (see Figure 2) and nurture with great food and wine. The banquet included a spectacular best paper award ceremony, discussed below in Section 4.

The shared experience of CIKM in Seoul made a deep impression on every participant, and those who were fortunate to attend experienced the Korean *jeong*.¹

3 Keynote Talks

Three invited keynote talks were given by Haesun Park (Georgia Institute of Technology), Sihem Amer-Yahia (CNRS / Univ. Grenoble Alpes), and Yong-Yeol Ahn (University of Virginia).

¹A concept with no English equivalent that cannot be defined but can be experienced...

3.1 Numerical Linear Algebraic Foundations for Large-Scale Unsupervised Learning

Haesun Park (Georgia Institute of Technology) gave a keynote on “Numerical Linear Algebraic Foundations for Large-Scale Unsupervised Learning”.

Haesun presented numerical linear algebra as a foundation for large-scale unsupervised learning. She showed that a low-rank approximation framework supported tasks such as clustering, topic modeling, community detection, embedding, and representation learning. She reviewed SVD, LSI, PCA, and constrained LRA methods, including NMF, SymNMF, and JointNMF. She argued that these methods fit into one framework. She then extended it to co-clustering, co-embedding, multi-view, and multi-granularity data, and semi-supervised settings with prior knowledge. She reported that the resulting algorithms achieved scalability, efficiency, and effectiveness, with experiments on synthetic and real data showing substantial benefits. See Park (2025) for more details.

3.2 AI Planning for Data Exploration

Sihem Amer-Yahia (CNRS / Univ. Grenoble Alpes) gave a keynote on “AI Planning for Data Exploration”.

Sihem presented AI planning for data exploration as an incremental, conversational process. She noted that reinforcement learning had automated parts of this process and that several solutions existed. With the rise of large language models and their sequential reasoning, she asked whether AI planning and LLMs outperformed customized RL policies. She also examined whether LLMs reduced retraining for new tasks and balanced specificity with generality. The talk aimed to answer these questions by reviewing RL training and policy reusability for data exploration. See Amer-Yahia (2025) for more details.

3.3 The Geometry of Knowledge and Computational Discovery

Yong-Yeol Ahn (University of Virginia) gave a keynote on “The Geometry of Knowledge and Computational Discovery”.

Yong-Yeol argued that modern neural networks map large datasets into geometric embedding spaces where semantic relations are encoded as spatial structure, and that the key to harnessing these models is to make those geometries interpretable. He introduced a simple theory of embedding interpretability that clarifies how to read and manipulate these spaces, enabling new ways to analyze data, design tailored metrics, and model system dynamics. This aims to shift from opaque black-box predictions to transparent, data-driven insights. See Ahn (2025) for more details.

4 Award Winners

CIKM 2025’s outstanding research quality was recognized with 10 awards, as well as additional CIKM 2025 AnalytiCup winners.



Figure 3. CIKM 2025 Best Paper Award.

4.1 Full Papers

The *CIKM 2025 Best Full Paper Award* was presented to: Weishuo Ma, Yanbo Wang, Xiyuan Wang, and Muhan Zhang (2025). “Reconsidering the Performance of GAE in Link Prediction”. In: CIKM ’25, pp. 2052–2062. DOI: [10.1145/3746252.3761108](https://doi.org/10.1145/3746252.3761108) (see Figure 3).

Ma et al. (2025) examine recent advances in GNNs for link prediction and warn that outdated baselines inflate perceived gains. They systematically explore Graph Autoencoders using model-agnostic tricks and hyperparameter tuning. A well-tuned GAE matches the performance of complex models while being more efficient on standard benchmarks. The approach works best when structure dominates and features are limited. It achieves an SOTA score on ogbl-ppa. It analyzes the impact of each trick to explain the success and to guide future designs. It emphasizes the need to update baselines for accurate progress assessment.

The *CIKM 2025 Best Student Full Paper Award* was presented to: Simone Merlo, Stefano Marchesin, Guglielmo Faggioli, and Nicola Ferro (2025). “A Cost-Effective Framework to Evaluate LLM-Generated Relevance Judgements”. In: CIKM ’25, pp. 2115–2126. DOI: [10.1145/3746252.3761200](https://doi.org/10.1145/3746252.3761200).

Merlo et al. (2025) show that LLMs reshape IR and already support tasks like query rewriting and retrieval-augmented generation. The community explores using LLMs to help or replace humans in generating relevance judgments and to expand evaluation collections, especially in complicated settings. Validation compares LLM-generated labels with human labels using inter-assessor agreement. To address this need, they introduce a framework that estimates the quality of LLM-generated judgments with statistical guarantees while minimizing human effort. The framework supports either fixed-confidence estimation with minimal assessment or fixed-budget estimation with bounds. Experiments on three IR collections with multiple LLM assessors indicate that reviewing a small fraction of LLM judgments suffices to estimate performance with high confidence.

The *CIKM 2025 Best Full Paper Runner Up Award* was presented to (in alphabetical order):

- Yufei Peng, Cheng Yang, Zhengjie Fan, and Chuan Shi (2025). “Data-centric Prompt Tuning for Dynamic Graphs”. In: *CIKM '25*, pp. 2336–2345. DOI: [10.1145/3746252.3761161](https://doi.org/10.1145/3746252.3761161).
- Zheng Zhang and Liang Zhao (2025). “Transferable Deep Clustering Model”. In: *CIKM '25*, pp. 4273–4282. DOI: [10.1145/3746252.3761188](https://doi.org/10.1145/3746252.3761188).

4.2 Applied Research Papers

The *CIKM 2025 Best Applied Research Paper Award* was presented to: Songpei Xu, Shijia Wang, Da Guo, Xianwen Guo, Qiang Xiao, Bin Huang, Guanlin Wu, and Chuanjiang Luo (2025). “Climber: Toward Efficient Scaling Laws for Large Recommendation Models”. In: *CIKM '25*, pp. 6193–6200. DOI: [10.1145/3746252.3761561](https://doi.org/10.1145/3746252.3761561).

Xu et al. (2025) observe that transformer-based generative models succeed across domains but face recommendation-specific issues. Scaling is not ideal due to heterogeneity across multiple sources and structural mismatches. Latency constraints of tens of milliseconds worsen as behavior sequences lengthen. To address this, the authors propose Climber, a recommendation framework that unites an efficient architecture with co-designed accelerations. Multi-scale sequence extraction reduces time complexity by a constant factor and scales better with sequence length. Dynamic temperature modulation adapts attention to diverse scenarios and behaviors. With single-user, multiple-item batching and memory-efficient key–value caching, throughput rises significantly without performance loss. Offline experiments across datasets show a more ideal scaling curve. Controlled scaling yields continuous online metric growth, with a significant overall lift at reasonable resource cost. Deployment on Netease Cloud Music serves tens of millions of users daily.

The *CIKM 2025 Best Student Applied Research Paper Award* was presented to: Lidi Zhang, Yinfeng Xiang, Wenjun Lyu, Zhiqing Hong, Haotian Wang, Desheng Zhang, Yunhuai Liu, and Tian He (2025). “D3-TR: Data-driven Daily Delivery Task Rescheduling for Cost-effective Last-mile Delivery”. In: *CIKM '25*, pp. 6266–6273. DOI: [10.1145/3746252.3761550](https://doi.org/10.1145/3746252.3761550).

L. Zhang et al. (2025) study last-mile logistics where couriers work in fixed zones and delays arise from leave or high-priority tasks. Historical analysis shows that daily rescheduling among on-duty couriers is a cost-effective alternative to temporary hires. The problem requires accurate workload assessment and fair task assignment. Existing methods lack optimization-aware workload models and often sacrifice fairness for efficiency. To address this, the authors propose D3-TR, which combines a consistency-guided workload predictor with a workload-aware genetic algorithm for task allocation. Experiments show superior results over baselines. Real-world deployment on millions of orders yields a significant improvement in on-time delivery.

4.3 Short Papers

The *CIKM 2025 Best Short Paper Award* was presented to: Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro (2025). “DP-COMET: A Differential Privacy Contextual Obfuscation MEchanism for Texts in Natural Language Processing”. In: *CIKM '25*, pp. 4700–4705. DOI: [10.1145/3746252.3760888](https://doi.org/10.1145/3746252.3760888).

De Faveri, Faggioli, and Ferro (2025) note that protecting sensitive text depends on context, while many epsilon-DP obfuscation methods operate term by term without context. The authors introduce DP-COMET, an epsilon-DP mechanism that evaluates a text’s context before obfuscation. The technique builds a representation of the entire text and adds noise to this representation, controlled by epsilon, to produce a private version. Experiments on NLP and IR tasks show performance comparable to term-by-term approaches. The obfuscated outputs are also less similar to the originals, which strengthens privacy.

4.4 Resource and Demo Papers

The *CIKM 2025 Best Resource Paper Award* was presented to: Clark Mingxuan Ju, Liam Collins, Leonardo Neves, Bhuvesh Kumar, Louis Yufeng Wang, Tong Zhao, and Neil Shah (2025). “Generative Recommendation with Semantic IDs: A Practitioner’s Handbook”. In: *CIKM ’25*, pp. 6420–6425. DOI: [10.1145/3746252.3761612](https://doi.org/10.1145/3746252.3761612).

Ju et al. (2025) note that generative recommendation shows strong performance, aided by semantic IDs that map continuous semantics to discrete sequences. The authors note that heterogeneous modeling choices, hyperparameters, and setups hinder fair comparison. They introduce GRID, an open-source, modular framework for GR with SIDs that supports easy component swapping and fast iteration. Using GRID, they run systematic experiments and ablations on public benchmarks. Results show that overlooked architectural components in SID-based GR can strongly affect performance. The work provides new insights and demonstrates the value of a unified platform for robust benchmarking and faster research progress.

The *CIKM 2025 Best Demo Paper Award* was presented to: Chengshuai Zhao, Riccardo De Maria, Tharindu Kumarage, Kumar Satvik Chaudhary, Garima Agrawal, Yiwen Li, Jongchan Park, Ying-Chih Chen, Yuli Deng, and Huan Liu (2025). “CyberBOT: Ontology-Grounded Retrieval Augmented Generation for Reliable Cybersecurity Education”. In: *CIKM ’25*, pp. 6752–6756. DOI: [10.1145/3746252.3761478](https://doi.org/10.1145/3746252.3761478).

C. Zhao et al. (2025) highlight LLM advances that enable inquiry-based learning tools in technical domains. In cybersecurity education, accuracy and safety are crucial, so systems must deliver trustworthy, domain-appropriate answers. The authors introduce CyberBOT, a question-answering chatbot that uses a retrieval-augmented generation pipeline to pull context from course materials and validate responses with a domain-specific cybersecurity ontology. The ontology serves as a structured reasoning layer that constrains and verifies LLM outputs, reducing the risk of misleading or unsafe guidance. CyberBOT is deployed in a large graduate-level course at Arizona State University.

4.5 Text of Time Award

The *CIKM Test of Time Award* is awarded to honor a paper published at *CIKM* at least 10 years ago that has had an extensive impact.

The *CIKM 2025 Test-of-Time Award* was presented to: Ee-Peng Lim, Viet-An Nguyen, Nitin Jindal, Bing Liu, and Hady Wirawan Lauw (2010). “Detecting product review spammers using rating behaviors”. In: *CIKM ’10*, pp. 939–948. DOI: [10.1145/1871437.1871557](https://doi.org/10.1145/1871437.1871557).

E.-P. Lim et al. (2010) aim to detect review spammers by modeling characteristic behaviors. The approach captures targeting specific products or groups, as well as deviations from other reviewers' ratings. They propose scoring methods to measure each reviewer's spam degree and apply them to Amazon reviews. A subset of highly suspicious reviewers is selected for human scrutiny using a web-based evaluation tool. Results show that ranking and supervised methods effectively identify spammers and outperform a baseline based on helpfulness votes. Moreover, detected spammers have a greater impact on product ratings than unhelpful reviewers.

The impact of E.-P. Lim et al. (2010) has been significant. At the time of writing, the ACM DL lists 461 citations and 2,871 downloads, and Google Scholar lists 1,030 citations.

5 Tutorials and Workshops

CIKM 2025 hosted 10 tutorials and 14 workshops.

5.1 Tutorials

The following 10 tutorial were offered:

- Hou, Zhang, Sheng, Wu, Wang, Chua, and McAuley (2025) on “*Towards Large Generative Recommendation: A Tokenization Perspective*”.
- Huang, Chen, Cheng, Kailkhura, Chawla, and Zhang (2025) on “*Socially Responsible and Trustworthy Generative Foundation Models: Principles, Challenges, and Practices*”.
- Sunwoo Kim, Lee, Gao, Antelmi, Polato, and Shin (2025) on “*A Tutorial on Hypergraph Neural Networks: An In-Depth and Step-By-Step Guide*”.
- Li, Huang, Li, Zhou, Zhang, and Liu (2025) on “*Generative Models for Synthetic Data: Transforming Data Mining in the GenAI Era*”.
- Oh, D. Lim, and Sungil Kim (2025) on “*Neural Differential Equations for Continuous-Time Analysis*”.
- Roy, Chakrabarti, and De (2025) on “*Retrieval of Graph Structured Objects: Theory and Applications*”.
- Saeedi and Fani (2025) on “*Neural Shifts in Collaborative Team Recommendation*”.
- Wang, Palikhe, Yin, and Zhang (2025) on “*Fairness in Language Models: A Tutorial*”.
- Yin, Wang, Palikhe, and Zhang (2025) on “*Uncertain Boundaries: A Tutorial on Copyright Challenges and Cross-Disciplinary Solutions for Generative AI*”.
- Yoo, Kang, and Tong (2025) on “*Continual Recommender Systems*”.

5.2 Workshops

The following 14 workshops were held:

- Tiropanis, Roussos, Bahrani, and Ragab (2025) on “*DESERE: The 2nd Workshop on Decentralized Search and Recommendation*”.
- Bin, Li, Ma, Zhang, Wenjie, Ma, Yang, and Chua (2025) on “*International Workshop on Multimodal Generative Search and Recommendation (MMGenSR@CIKM 2025)*”.

-
- Chatterjee, Wang, Zhang, Ebrahimi, Ren, Ganguly, Jones, Yilmaz, and Zamani (2025) on “*ProActLLM: Proactive Conversational Information Seeking with Large Language Models*”.
 - Choi, Han, Koo, Bae, Yoo, Woo, and Samek (2025) on “*Human-Centric AI: From Explainability and Trustworthiness to Actionable Ethics*”.
 - Di Teodoro, Guarrasi, Siciliano, and Silvestri (2025) on “*Advances in Medical Knowledge Systems: LLMs, RAG and Foundation Models*”.
 - Galletta, Taheri, Di Modica, and Ficara (2025) on “*SIoTEc 2025 - 6th edition of ACM Workshop on Secure IoT, Edge and Cloud systems*”.
 - E. Kim, Han, Lim, Selcuk, Back, Han, and Lee (2025) on “*The 1st International Workshop on Retrieval-driven Generative AI & ScienceON AI Challenge: RDGENAI 2025*”.
 - Lee, Mehra, Choi, Chen, Mehta, Zohren, Kim, Lee, Lee, and Oh (2025) on “*Advances in Financial AI: Innovations, Risk, and Responsibility in the Era of LLMs*”.
 - Merra, Skračić, Malitesta, Golebiowski, and Minervini (2025) on “*SmaLLEXT: 1st Workshop on Small and Efficient Large Language Models for Knowledge Extraction*”.
 - Miao, Zhao, Liang, Yang, Zheng, and Jensen (2025) on “*The International Workshop on Spatio-Temporal Data Intelligence and Foundation Models*”.
 - Rostami, Felfernig, Wörndl, Oussalah, Anand, Jalili, and Banerjee (2025) on “*Recommender Systems for Sustainable Development through Responsible Nudging*”.
 - Sun, Zhang, Fu, Song, Li, and Yu (2025) on “*Frontiers in Graph Machine Learning for the Large Model Era*”.
 - Wu, Wang, Chen, Zhao, Yu, Deldjoo, and Lian (2025) on “*Trustworthy Knowledge Discovery and Data Mining (TrustKDD)*”.
 - Yuan, Zhou, Xu, Pang, Su, Zhang, Xiao, Xu, Ren, and Chen (2025) on “*The 1st Workshop on LLM Agents for Social Simulation*”.

6 Envoi

This concludes our report on CIKM 2025. In addition to the paper tracks discussed in this report, CIKM also hosted an Industry Day and a PhD Symposium.

We want to give two special thanks in the main text rather than hide them in an acknowledgment section, as they were crucial to making the conference a great success. First, we thank all the sponsors supporting CIKM. Their generous support helped more students (and seniors) to attend CIKM and improved the social events. But we also want to thank them for their participation in the technical program. Second, we were fortunate to be supported by a massive program committee that helped the crucial peer-review process to create a strong technical program for the conference. We were fortunate to work with a combine program committee of 2,845 fellow researchers in total, apart from the chairs, we worked with 29 Associate Chairs, 439 Senior PC members, 1,888 PC members, and 452 subreviewers. We want to thank each one of them, for offer their time and expertise, to make fair recommendations resulting in a very strong CIKM program.²

The record-breaking conference in Seoul demonstrated the growth and vitality of the field. CIKM 2025 had the largest number of submissions and accepted papers, the largest PC and the

²For the complete list of names, we refer to the front matter of the CIKM 2025 proceedings (Cha et al., 2025).



Figure 4. Scenes from the CIKM 2025 Opening Session.

largest attendance, in the history of CIKM. This growth reflects the status of CIKM and its broad scope including SIGIR, SIGWEB, SIGMOD, SIGKDD, SIGAI, etc. With the advent of current AI models, traditional boundaries among these areas are diminishing, and CIKM's broad scope is a natural home for AI. As noted during the opening session (see Figure 4), CIKM has traditionally been regarded as the #2 conference in many of these areas, but it has now surpassed the #1 conferences in size.

The growth of the conference in 2025 is promising a bright future for CIKM, and we are proud to transfer the conference to the CIKM 2026 organizers. CIKM 2026 will be held in Rome, Italy, with Nicola Ferro (University of Padova) and Antonella Poggi (Sapienza University of Rome) acting as the general chairs. We look forward to attending CIKM 2026 in Rome.

7 Additional Authors

This report and the conference itself would never happen without the concerted effort of the following co-organizers: Kijung Shin, Bryan Hooi, and Lifang He (Short Paper Program Chairs); Yong Li, Julian McAuley, Yang Song, and Sungchul Kim (Applied Research Track Chairs); Jundong Li and Cheng-Te Li (Resources and Demo Chairs); Xiangnan He and Qingyun Wu (Workshop Chairs); Matteo Riondato, Ee-Peng Lim, and Giulia Preti (Tutorial Chairs); Vachik Dave and Wei Jin (AnalytiCup Chairs); Jingren Zhou, Soonmin Bae, and Xianfeng Tang (Industry Day Chairs); Qiaoyu Tan and Isaac Johnson (PhD Symposium Chairs); Afra Mashhadi (Diversity and Inclusion Chair); Diego Saez Trumper and Gi-Soo Kim (Publicity Chair); Jian Kang and Eun-Sol Kim (Proceedings Chairs); Dongwon Lee and Dawei Zhou (Student Travel Award Chairs); Yeon-Chang Lee, Buru Chang, and Soo Kyung Kim (Local Chairs); and Jeongwhan Choi (Web Chair).

8 Acknowledgments

It was a distinct honor and privilege to work with the CIKM Steering Committee chair, Prof. Alistair Moffat, and the rest of the CIKM SC. The level of support and the wealth of knowledge and experience were key to realizing a conference of this size and complexity. We thank SIGIR and SIGWEB for sponsoring the CIKM Conference.

The success of CIKM 2025 would not have been possible without all the help from the team of volunteers and reviewers. We wish to thank all the reviewers and meta-reviewers who helped to ensure the high quality of the program. We also wish to thank the entire organization team (see Table 1).

For further information, we refer to the conference web pages <https://cikm2025.org/> and the proceedings. The CIKM'25 proceedings are published in a single volume (Cha et al., 2025). Several workshops also published their own proceedings, see the pages of each workshop <https://cikm2025.org/program/workshops>.

References

- Yong-Yeol Ahn (2025). “The Geometry of Knowledge and Computational Discovery”. In: CIKM '25, p. 1. DOI: [10.1145/3746252.3763787](https://doi.org/10.1145/3746252.3763787).
- Sihe Amer-Yahia (2025). “AI Planning for Data Exploration”. In: CIKM '25, p. 2. DOI: [10.1145/3746252.3763786](https://doi.org/10.1145/3746252.3763786).
- Yi Bin, Haoxuan Li, Haokai Ma, Yang Zhang, Wang Wenjie, Yunshan Ma, Yang Yang, and Tat-Seng Chua (2025). “International Workshop on Multimodal Generative Search and Recommendation (MMGenSR@CIKM 2025)”. In: CIKM '25, pp. 6890–6893. DOI: [10.1145/3746252.3761597](https://doi.org/10.1145/3746252.3761597).
- Meeyoung Cha, Chanyoung Park, Noseong Park, Carl Yang, Senjuti Basu Roy, Jessie Li, Jaap Kamps, Kijung Shin, Bryan Hooi, and Lifang He, eds. (2025). *CIKM '25: Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. DOI: [10.1145/3746252](https://doi.org/10.1145/3746252).
- Shubham Chatterjee, Xi Wang, Shuo Zhang, Sajad Ebrahimi, Zhaochun Ren, Debasis Ganguly, Gareth Jones, Emine Yilmaz, and Hamed Zamani (2025). “ProActLLM: Proactive Conversational Information Seeking with Large Language Models”. In: CIKM '25, pp. 6894–6897. DOI: [10.1145/3746252.3761593](https://doi.org/10.1145/3746252.3761593).
- Jaesik Choi, Bohyung Han, Myoung-Wan Koo, Kyungman Bae, Chang D. Yoo, Simon S Woo, and Wojciech Samek (2025). “Human-Centric AI: From Explainability and Trustworthiness to Actionable Ethics”. In: CIKM '25, pp. 6898–6900. DOI: [10.1145/3746252.3761599](https://doi.org/10.1145/3746252.3761599).
- Francesco Luigi De Faveri, Guglielmo Faggioli, and Nicola Ferro (2025). “DP-COMET: A Differential Privacy Contextual Obfuscation MEchanism for Texts in Natural Language Processing”. In: CIKM '25, pp. 4700–4705. DOI: [10.1145/3746252.3760888](https://doi.org/10.1145/3746252.3760888).
- Giulia Di Teodoro, Valerio Guarrasi, Federico Siciliano, and Fabrizio Silvestri (2025). “Advances in Medical Knowledge Systems: LLMs, RAG and Foundation Models”. In: CIKM '25, pp. 6901–6904. DOI: [10.1145/3746252.3761592](https://doi.org/10.1145/3746252.3761592).

-
- Antonino Galletta, Javid Taheri, Giuseppe Di Modica, and Annamaria Ficara (2025). “SIoTEc 2025 - 6th edition of ACM Workshop on Secure IoT, Edge and Cloud systems”. In: CIKM ’25, pp. 6905–6907. DOI: [10.1145/3746252.3761600](https://doi.org/10.1145/3746252.3761600).
- Yupeng Hou, An Zhang, Leheng Sheng, Jiancan Wu, Xiang Wang, Tat-Seng Chua, and Julian McAuley (2025). “Towards Large Generative Recommendation: A Tokenization Perspective”. In: CIKM ’25, pp. 6821–6824. DOI: [10.1145/3746252.3761451](https://doi.org/10.1145/3746252.3761451).
- Yue Huang, Canyu Chen, Lu Cheng, Bhavya Kailkhura, Nitesh Chawla, and Xiangliang Zhang (2025). “Socially Responsible and Trustworthy Generative Foundation Models: Principles, Challenges, and Practices”. In: CIKM ’25, pp. 6825–6828. DOI: [10.1145/3746252.3761448](https://doi.org/10.1145/3746252.3761448).
- Clark Mingxuan Ju, Liam Collins, Leonardo Neves, Bhuvesh Kumar, Louis Yufeng Wang, Tong Zhao, and Neil Shah (2025). “Generative Recommendation with Semantic IDs: A Practitioner’s Handbook”. In: CIKM ’25, pp. 6420–6425. DOI: [10.1145/3746252.3761612](https://doi.org/10.1145/3746252.3761612).
- Eunhui Kim, Caren Caren Han, KeyongTae Lim, Yesim Selcuk, Moonki Back, SangJun Han, and KyongHa Lee (2025). “The 1st International Workshop on Retrieval-driven Generative AI & ScienceON AI Challenge: RDGENAI 2025”. In: CIKM ’25, pp. 6908–6911. DOI: [10.1145/3746252.3761603](https://doi.org/10.1145/3746252.3761603).
- Sunwoo Kim, Soo Yong Lee, Yue Gao, Alessia Antelmi, Mirko Polato, and Kijung Shin (2025). “A Tutorial on Hypergraph Neural Networks: An In-Depth and Step-By-Step Guide”. In: CIKM ’25, pp. 6829–6832. DOI: [10.1145/3746252.3761449](https://doi.org/10.1145/3746252.3761449).
- Yongjae Lee, Nazanin Mehrasa, Chanyeol Choi, Chung-Chi Chen, Dhagash Mehta, Stefan Zohren, Yoon Kim, Chulheum Lee, Yeonhee Lee, and Eunsook Oh (2025). “Advances in Financial AI: Innovations, Risk, and Responsibility in the Era of LLMs”. In: CIKM ’25, pp. 6912–6915. DOI: [10.1145/3746252.3761591](https://doi.org/10.1145/3746252.3761591).
- Dawei Li, Yue Huang, Ming Li, Tianyi Zhou, Xiangliang Zhang, and Huan Liu (2025). “Generative Models for Synthetic Data: Transforming Data Mining in the GenAI Era”. In: CIKM ’25, pp. 6833–6836. DOI: [10.1145/3746252.3761455](https://doi.org/10.1145/3746252.3761455).
- Ee-Peng Lim, Viet-An Nguyen, Nitin Jindal, Bing Liu, and Hady Wirawan Lauw (2010). “Detecting product review spammers using rating behaviors”. In: CIKM ’10, pp. 939–948. DOI: [10.1145/1871437.1871557](https://doi.org/10.1145/1871437.1871557).
- Weishuo Ma, Yanbo Wang, Xiyuan Wang, and Muhan Zhang (2025). “Reconsidering the Performance of GAE in Link Prediction”. In: CIKM ’25, pp. 2052–2062. DOI: [10.1145/3746252.3761108](https://doi.org/10.1145/3746252.3761108).
- Simone Merlo, Stefano Marchesin, Guglielmo Faggioli, and Nicola Ferro (2025). “A Cost-Effective Framework to Evaluate LLM-Generated Relevance Judgements”. In: CIKM ’25, pp. 2115–2126. DOI: [10.1145/3746252.3761200](https://doi.org/10.1145/3746252.3761200).
- Felice Antonio Merra, Kristian Skračić, Daniele Malitesta, Jacek Golebiowski, and Pasquale Minervini (2025). “SmaLLEXT: 1st Workshop on Small and Efficient Large Language Models for Knowledge Extraction”. In: CIKM ’25, pp. 6916–6919. DOI: [10.1145/3746252.3761595](https://doi.org/10.1145/3746252.3761595).
- Hao Miao, Yan Zhao, Yuxuan Liang, Bin Yang, Kai Zheng, and Christian S. Jensen (2025). “The International Workshop on Spatio-Temporal Data Intelligence and Foundation Models”. In: CIKM ’25, pp. 6920–6922. DOI: [10.1145/3746252.3761601](https://doi.org/10.1145/3746252.3761601).
- Yongkyung Oh, Dongyoung Lim, and Sungil Kim (2025). “Neural Differential Equations for Continuous-Time Analysis”. In: CIKM ’25, pp. 6837–6840. DOI: [10.1145/3746252.3761447](https://doi.org/10.1145/3746252.3761447).

-
- Haesun Park (2025). “Numerical Linear Algebraic Foundations for Large-Scale Unsupervised Learning”. In: CIKM ’25, p. 3. DOI: [10.1145/3746252.3763785](https://doi.org/10.1145/3746252.3763785).
- Yufei Peng, Cheng Yang, Zhengjie Fan, and Chuan Shi (2025). “Data-centric Prompt Tuning for Dynamic Graphs”. In: CIKM ’25, pp. 2336–2345. DOI: [10.1145/3746252.3761161](https://doi.org/10.1145/3746252.3761161).
- Mehrdad Rostami, Alexander Felfernig, Wolfgang Wörndl, Mourad Oussalah, Avishek Anand, Mahdi Jalili, and Ashmi Banerjee (2025). “Recommender Systems for Sustainable Development through Responsible Nudging”. In: CIKM ’25, pp. 6923–6926. DOI: [10.1145/3746252.3761596](https://doi.org/10.1145/3746252.3761596).
- Indradyumna Roy, Soumen Chakrabarti, and Abir De (2025). “Retrieval of Graph Structured Objects: Theory and Applications”. In: CIKM ’25, pp. 6841–6844. DOI: [10.1145/3746252.3761454](https://doi.org/10.1145/3746252.3761454).
- Mahdis Saeedi and Hossein Fani (2025). “Neural Shifts in Collaborative Team Recommendation”. In: CIKM ’25, pp. 6845–6848. DOI: [10.1145/3746252.3761456](https://doi.org/10.1145/3746252.3761456).
- Qingyun Sun, Ziwei Zhang, Xingcheng Fu, Yangqiu Song, Jianxin Li, and Philip S. Yu (2025). “Frontiers in Graph Machine Learning for the Large Model Era”. In: CIKM ’25, pp. 6927–6929. DOI: [10.1145/3746252.3761590](https://doi.org/10.1145/3746252.3761590).
- Thanassis Tiropanis, George Roussos, Mohammad Bahrani, and Mohamed Ragab (2025). “DESERE: The 2nd Workshop on Decentralized Search and Recommendation”. In: CIKM ’25, pp. 6887–6889. DOI: [10.1145/3746252.3761602](https://doi.org/10.1145/3746252.3761602).
- Zichong Wang, Avash Palikhe, Zhipeng Yin, and Wenbin Zhang (2025). “Fairness in Language Models: A Tutorial”. In: CIKM ’25, pp. 6849–6852. DOI: [10.1145/3746252.3761453](https://doi.org/10.1145/3746252.3761453).
- Le Wu, Jindong Wang, Ling Chen, Xiangyu Zhao, Kui Yu, Yashar Deldjoo, and Defu Lian (2025). “Trustworthy Knowledge Discovery and Data Mining (TrustKDD)”. In: CIKM ’25, pp. 6930–6932. DOI: [10.1145/3746252.3761598](https://doi.org/10.1145/3746252.3761598).
- Songpei Xu, Shijia Wang, Da Guo, Xianwen Guo, Qiang Xiao, Bin Huang, Guanlin Wu, and Chuanjiang Luo (2025). “Climber: Toward Efficient Scaling Laws for Large Recommendation Models”. In: CIKM ’25, pp. 6193–6200. DOI: [10.1145/3746252.3761561](https://doi.org/10.1145/3746252.3761561).
- Zhipeng Yin, Zichong Wang, Avash Palikhe, and Wenbin Zhang (2025). “Uncertain Boundaries: A Tutorial on Copyright Challenges and Cross-Disciplinary Solutions for Generative AI”. In: CIKM ’25, pp. 6853–6856. DOI: [10.1145/3746252.3761450](https://doi.org/10.1145/3746252.3761450).
- Hyunsik Yoo, Seongku Kang, and Hanghang Tong (2025). “Continual Recommender Systems”. In: CIKM ’25, pp. 6857–6860. DOI: [10.1145/3746252.3761452](https://doi.org/10.1145/3746252.3761452).
- Yige Yuan, Junkai Zhou, Bingbing Xu, Liang Pang, Du Su, An Zhang, Teng Xiao, Fengli Xu, Zhaochun Ren, and Xu Chen (2025). “The 1st Workshop on LLM Agents for Social Simulation”. In: CIKM ’25, pp. 6933–6936. DOI: [10.1145/3746252.3761594](https://doi.org/10.1145/3746252.3761594).
- Lidi Zhang, Yinfeng Xiang, Wenjun Lyu, Zhiqing Hong, Haotian Wang, Desheng Zhang, Yunhuai Liu, and Tian He (2025). “D3-TR: Data-driven Daily Delivery Task Rescheduling for Cost-effective Last-mile Delivery”. In: CIKM ’25, pp. 6266–6273. DOI: [10.1145/3746252.3761550](https://doi.org/10.1145/3746252.3761550).
- Zheng Zhang and Liang Zhao (2025). “Transferable Deep Clustering Model”. In: CIKM ’25, pp. 4273–4282. DOI: [10.1145/3746252.3761188](https://doi.org/10.1145/3746252.3761188).
- Chengshuai Zhao, Riccardo De Maria, Tharindu Kumarage, Kumar Satvik Chaudhary, Garima Agrawal, Yiwen Li, Jongchan Park, Ying-Chih Chen, Yuli Deng, and Huan Liu (2025). “CyberBOT: Ontology-Grounded Retrieval Augmented Generation for Reliable Cybersecurity Education”. In: CIKM ’25, pp. 6752–6756. DOI: [10.1145/3746252.3761478](https://doi.org/10.1145/3746252.3761478).