

Report on the 3rd ACM SIGIR Conference on Information Retrieval in the Asia Pacific (SIGIR-AP 2025)

Keping Bi

Institute of Computing Technology, Chinese Academy of Sciences
China
bikeping@ict.ac.cn

Abstract

The 3rd ACM SIGIR Conference on Information Retrieval in the Asia Pacific (SIGIR-AP 2025) was held from 7 to 10 December 2025 in Xi'an, China. Hosted by Xidian University in cooperation with the SIGIR Beijing Chapter, the hybrid conference welcomed 160 participants from 12 countries and regions, representing 61 institutions across academia, research institutes, and industry. The program featured two keynote addresses, four tutorials, two workshops, two industry talks, oral and poster sessions, best paper and outstanding reviewer recognitions, and a social program designed to support community building. This report summarizes the conference organization, venue, technical program, keynotes, tutorials, workshops, awards, social events, sponsorship, and student and volunteer engagement, and reflects on the role of SIGIR-AP in strengthening information retrieval research across the Asia-Pacific region and the broader SIGIR community.

Date: 7–10 December 2025.

Website: <http://www.sigir-ap.org/sigir-ap-2025/>.

1 Introduction

SIGIR-AP 2025 was the third edition of the ACM SIGIR Conference on Information Retrieval in the Asia Pacific. Following the first two editions in Beijing in 2023 and Tokyo in 2024, the 2025 conference was held in Xi'an, China, from 7 to 10 December 2025. The conference continued SIGIR-AP's mission as an ACM-sponsored international forum for information retrieval (IR) research and development across Asia, Australasia, and the islands of the Pacific and Indian Oceans.

The conference was organized as a hybrid event. Participants could join in person in Xi'an, at the Brisbane Satellite venue at the University of Queensland, or remotely from elsewhere in the world. This format allowed SIGIR-AP 2025 to support both regional engagement and wider international participation. In total, the conference attracted 160 participants from 12 countries and regions and 61 institutions. There were 111 in-person participants, including 38



Figure 1. SIGIR-AP 2025 in Xi'an, China.

from international institutions; 30 participants joined through the Brisbane satellite venue; and 19 participated fully online, 15 of whom were international.

2 The Role and Significance of SIGIR-AP

SIGIR-AP was established to promote and disseminate IR research within the Asia-Pacific region while maintaining strong connections with the global SIGIR community. This regional focus is increasingly important. The Asia-Pacific research community is large, multilingual, and methodologically diverse, with strong contributions across web search, recommendation, retrieval-augmented generation (RAG), evaluation, trustworthy IR, conversational search, and user-centered information access. SIGIR-AP provides a venue where these contributions can be presented, discussed, and connected across geographic and institutional boundaries.

The 2025 edition was held at a moment when generative AI is reshaping the design and evaluation of IR systems. The conference theme and program reflected this shift: many sessions addressed RAG, large language models (LLMs), ranking, recommendation, bias and fairness, privacy, auditing, and human-system interaction. At the same time, SIGIR-AP retained a broader view of IR by including work on resources, reproducibility, applications, and evaluation. The conference's Revise-and-Resubmit (R&R) track also remains a distinctive feature: authors of

papers submitted to but not accepted by SIGIR or ICTIR could revise their work in response to reviewer feedback and submit it to SIGIR-AP with a detailed response.

This combination of regional engagement, hybrid access, and a technical program aligned with current research challenges makes SIGIR-AP an important part of the broader IR ecosystem. It gives researchers in the Asia-Pacific region a strong platform for scholarly exchange and gives the global IR community a regular opportunity to engage with work, students, and institutions across the region.

3 Conference Organization

SIGIR-AP 2025 was hosted by Xidian University and organized in cooperation with the SIGIR Beijing Chapter. The General Chairs were Keping Bi (Chinese Academy of Sciences, China) and Qingshan Li (Xidian University, China). The Program Chairs were Evangelos Kanoulas (University of Amsterdam, Netherlands), Yadan Luo (The University of Queensland, Australia), and Bhaskar Mitra (Independent Researcher, Canada).

The broader organizing team included Makoto Kato (University of Tsukuba, Japan) as Workshop Chair, Krisztian Balog (University of Stavanger, Norway) as Tutorial Chair, Rui Li, Lu Wang, and Yueshen Xu (Xidian University, China) as Local Chairs, Shuai Wang (University of Queensland, Australia) as Queensland Satellite Chair, Zhijing Wu (Beijing Institute of Technology, China) as Registration Chair, Yue Feng (University of Birmingham, United Kingdom) as Proceedings Chair, Fei Sun (Chinese Academy of Sciences, China) as Sponsorship Chair, Shiyu Ni (Chinese Academy of Sciences, China) as Web Master, Razieh Negin Rahimi (University of Massachusetts Amherst, USA) and Zhaochun Ren (Leiden University, Netherlands) as Publicity Chairs, and Faegheh Hasibi (Radboud University, Netherlands), Sean MacAvaney (University of Glasgow, United Kingdom), and Alistair Moffat (The University of Melbourne, Australia) as Best Paper Chairs.

4 Venue

The main conference venue was the Wyndham Grand Xi'an South, located at No. 208 Ci'en East Road, Qujiang New District, Xi'an, Shaanxi. The hotel is situated in Datang Everbright City near the Giant Wild Goose Pagoda, Tang City Wall Ruins Park, Datang Furong Garden, and Qujiang Pool. This setting placed the conference within walking distance of cultural sites, restaurants, shopping areas, and public transportation, including Metro Line 4.

The venue also supported the practical needs of the conference. The Grand Ballroom hosted the main sessions, keynotes, industry talks, reception, and workshops, while function rooms supported tutorials and parallel activities. The location helped combine a focused technical program with opportunities for attendees to experience Xi'an's local history and culture.

5 Technical Program

The technical program was curated by the Program Chairs with support from senior program committee members, reviewers, and meta-reviewers. SIGIR-AP 2025 invited original research, resource, reproducibility, industry, and perspective papers, and also accepted R&R submissions from SIGIR 2025 and ICTIR 2025. Review assignments were based on reviewer expertise and conflict-of-interest information, checked by the program committee chairs, and adjusted where necessary. Each reviewed submission received at least three reviews, followed by discussion under the guidance of senior program committee members.

SIGIR-AP 2025 received 112 abstract submissions. After desk rejections, 98 submissions entered the peer-review process; 46 papers were accepted, yielding an overall acceptance rate of 46.9%. The regular track accepted 35 of 83 reviewed submissions, while the R&R track accepted 11 of 15 reviewed submissions. The proceedings are available in the ACM Digital Library.¹

Table 1. SIGIR-AP 2025 acceptance statistics.

Venue or Track	#Reviewed	#Accepted	Acceptance Rate
Regular submissions	83	35	42.2%
Revise-and-Resubmit submissions	15	11	73.3%
Total	98	46	46.9%




Figure 2. Geographical distribution of accepted papers.

The main conference program included 51 oral presentations: 46 SIGIR-AP papers including one presentation from the Brisbane Satellite venue, and five TOIS papers. The conference also

¹<https://dl.acm.org/doi/proceedings/10.1145/3767695>

included 52 poster presentations including 1 postly-only TOIS presentation. Technical sessions covered Retrieval-Augmented Generation, Bias, Fairness, Privacy, and Auditing, Applications, Ranking, Representation Learning and Retrieval, Recommender Systems, Language Modeling, and User-System Interactions and Benchmarking. A dedicated poster session and coffee break gave attendees additional time for direct discussion with authors.

6 Keynotes

SIGIR-AP 2025 featured two keynote speakers who addressed timely questions about evaluation and human-centered information access in the LLM era.

6.1 Keynote 1: Fernando Diaz

Title: Preference-Based Evaluation. Fernando Diaz is an Associate Professor at Carnegie Mellon University’s Language Technologies Institute and a Research Scientist at Google. His work spans quantitative evaluation of AI systems, retrieval-enhanced AI, and the cultural impact of AI.

Diaz’s keynote argued that recent advances in AI have exposed limits in traditional evaluation metrics and benchmarks. As modern retrieval models saturate high-precision metrics, evaluation methods need to capture more meaningful differences in system behavior. The talk presented preference-based evaluation as a framework that treats evaluation as an ordering over system behaviors rather than only as the calculation of numeric scores. Drawing on IR evaluation and social choice theory, the keynote introduced preference-based counterparts of precision, recall, and average precision, and showed how these methods can improve sensitivity across retrieval and recommendation tasks [Diaz, 2025].

6.2 Keynote 2: Min Zhang

Title: LLMs vs. Humans in Information Access Tasks: Performances, Behaviors, and Learning Abilities. Min Zhang is a Professor in the Department of Computer Science and Technology at Tsinghua University and Director of the AI Lab. Her research covers web search, recommendation, and user modeling; she is an ACM Distinguished Member, a SIGIR Academy Fellow, and Editor-in-Chief of ACM Transactions on Information Systems.

Zhang’s keynote examined how LLMs compare with human users in information access tasks. The talk approached this question from a user-centric perspective, asking how the performance of large models should be evaluated, how their conversational behaviors differ from human behaviors in IR tasks, and how their test-time learning from conversational reasoning experiences compares with human learning. By bringing together performance evaluation, behavioral analysis, and reasoning-based learning, the keynote offered a framework for studying LLMs not only as tools for retrieval, but also as actors whose behavior must be understood in relation to human information seeking [Zhang, 2025].

7 Tutorials

Four half-day tutorials were offered on the first day of the conference.

1. **Retrieval and Ranking with LLMs (R²LLMs)** (Guido Zuccon, Shengyao Zhuang, Xueguang Ma, and Bevan Koopman). This tutorial introduced the use of LLMs for retrieval and ranking, covering LLM-based retrievers, LLM-based rankers, training and prompting strategies, effectiveness and efficiency considerations, and open research challenges in deploying LLMs within retrieval pipelines [Zuccon et al., 2025].
2. **Dynamic and Parametric Retrieval-Augmented Generation** (Weihang Su, Qian Dong, Qingyao Ai, and Yiqun Liu). This tutorial focused on two complementary directions in RAG. Dynamic RAG studies when, what, and how to retrieve during generation, while parametric RAG examines how retrieved knowledge can be integrated into model parameters and reasoning processes. Together, these directions address both adaptive retrieval behavior and deeper forms of knowledge integration [Su et al., 2025].
3. **Trustworthy Information Retrieval in the LLM Era: Bias, Unfairness, and Hallucination** (Sunhao Dai, Chen Xu, Shicheng Xu, Zhongxiang Sun, Liang Pang, Zhenhua Dong, and Jun Xu). This tutorial addressed trustworthiness challenges that arise when LLMs are integrated into IR systems, including bias, unfairness, and hallucination. It organized these challenges around distributional mismatch and discussed mitigation strategies for building more reliable and responsible IR systems [Dai et al., 2025].
4. **Conversational Agents: From RAG to LTM** (Dell Zhang, Yue Feng, Haiming Liu, Changzhi Sun, Jixiang Luo, Xiangyu Chen, and Xuelong Li). This tutorial connected RAG-based conversational systems with long-term memory (LTM), covering trustworthy conversational agents, tool use, multimodal conversational search, and memory mechanisms including textual, graph-based, and parametric memory [Zhang et al., 2025].

8 Workshops

Two workshops were held with SIGIR-AP 2025, both addressing the rapidly developing research agenda around RAG.

1. **R³AG 2025**.² The Second Workshop on Refined and Reliable Retrieval-Augmented Generation was organized by Haitao Yu, Yubo Fang, Xuri Ge, Xin Xin, Zihan Wang, Junchen Fu, Joemon M. Jose, Weizhi Ma, and Zhaochun Ren. The workshop focused on refined and reliable RAG, bringing together work on fundamental challenges, implementation strategies, evaluation, and robustness for retrieval-augmented generation systems [Yu et al., 2025].
2. **BREV-RAG**.³ Beyond Relevance-based Evaluation of RAG Systems was organized by Tetsuya Sakai, Sijie Tao, Zhicheng Dou, Junjie Wang, Haoxiang Shi, Nuo Chen, and Atsushi Keyaki. This workshop examined RAG evaluation beyond relevance, correctness, and groundedness. Topics included fairness, calibrated confidence, and broader evaluation criteria for systems that combine retrieval and generation [Sakai et al., 2025].

²<https://ii-research.github.io/R3AG/2025/>

³<http://sakailab.com/brev-rag/>

9 Industry Talks

Two industry talks complemented the research program. Daoxin Zhang presented **Application of Multimodal LLM in Xiaohongshu Search**, highlighting how multimodal large language models are being applied in a large-scale search setting. Ruiming Tang presented **Recent Technical Advances in Recommender System of Kuaishou**, discussing recent recommender-system developments from an industrial perspective.

10 Awards and Recognition

SIGIR-AP 2025 recognized outstanding research contributions and reviewing service through best paper awards and outstanding reviewer recognitions.

10.1 Paper Awards

Best Paper Award:

- **ScaleFormer: Span Representation Cumulation for Long-Context Transformer**, by Jiangshu Du, Wenpeng Yin, and Philip Yu [Du et al., 2025].

Best Paper Runner-Up Award:

- **Adversarial Attacks against Neural Ranking Models via In-Context Learning**, by Amin Bigdeli, Negar Arabzadeh, Ebrahim Bagheri, and Charles L. A. Clarke [Bigdeli et al., 2025].



Figure 3. Best paper and outstanding reviewer award presentations.

10.2 Outstanding Reviewers

The conference also recognized reviewers who made exceptional contributions to the peer-review process:

- Florian Boudin, Université de Nantes, France.

- Matthias Hagen, Friedrich-Schiller-Universität Jena, Germany.
- Wei Huang, Chinese Academy of Sciences, China.
- Makoto P. Kato, University of Tsukuba, Japan.
- Luis Martínez, University of Jaén, Spain.
- Alistair Moffat, The University of Melbourne, Australia.
- Iadh Ounis, University of Glasgow, United Kingdom.
- Virgil Pavlu, Northeastern University, USA.
- Maya Ramanath, IIT Delhi, India.
- Tetsuya Sakai, Waseda University, Japan.
- Paul Thomas, Microsoft, Australia.
- Julián Urbano, Delft University of Technology, Netherlands.
- Takehiro Yamamoto, University of Hyogo, Japan.

Figure 3 shows the presentation of the best paper and outstanding reviewer awards during the conference banquet.

11 Social Events

SIGIR-AP 2025 featured a social program designed to support informal exchange and community building. The reception was held on the evening of the tutorial day, giving attendees an early opportunity to meet, reconnect, and continue conversations beyond the technical sessions, shown in Figure 4.



Figure 4. Reception at SIGIR-AP 2025.

The banquet introduced participants to interactive cultural activities, including calligraphy, fan lacquering, and Touhu(a traditional pitching game), shown in Figure 5. These activities gave attendees a relaxed setting in which to experience local traditions while continuing discussions from the conference program.

The banquet also included cultural performances: Tang Dynasty drums, Bian Lian face-changing, silk-string melodies by an ensemble of traditional Chinese instruments, and the dance performance *A Thousand Miles of Rivers and Mountains: Only Green*. A lucky draw added to the festive atmosphere, with one first prize, a shadow puppet art piece; three second prizes, grape-



Figure 5. Interactive cultural activities before the banquet.

and-bird pattern sachets; and five third prizes, Tang Dynasty mural photo frames. Figure 6 shows the cultural performances and lucky draw at the banquet.

12 Sponsorship

SIGIR-AP 2025 was supported by ACM SIGIR, with additional financial support from Xiaohongshu as Platinum Sponsor, Kuaishou as Gold Sponsor, and Weixin Search as Silver Sponsor. The organizers gratefully acknowledge this support, which helped reduce the financial burden of conference participation and strengthened exchange between academia and industry within the IR community.

13 Student and Volunteer Engagement

SIGIR-AP 2025 placed strong emphasis on early-career engagement and community building. A total of 24 student volunteers supported conference organization and on-site operations, and 9 students received student travel grants. The hybrid format, Queensland Satellite venue, poster sessions, social events, and informal interactions all helped students connect with researchers across institutions and countries. These efforts continue SIGIR-AP’s role in nurturing the next generation of IR researchers in the Asia-Pacific region.




Figure 6. Cultural performances and lucky draw at the banquet.



Figure 7. SIGIR-AP 2025 sponsors: Xiaohongshu (Platinum), Kuaishou (Gold), and Weixin Search (Silver).

14 Looking Ahead

SIGIR-AP 2025 demonstrated the value of a regional conference that is both locally grounded and globally connected. The program showed the field’s rapid movement toward generative AI, RAG, LLM-based search and recommendation, and trustworthy information access, while also maintaining attention to evaluation, representation learning, user interaction, resources, and reproducibility. The strong international participation, hybrid access model, and satellite venue further reinforced SIGIR-AP’s commitment to inclusiveness and regional engagement.

Looking ahead, SIGIR-AP will continue to serve as a platform for advancing IR research and strengthening collaboration across the Asia-Pacific and global SIGIR communities. Those interested in hosting SIGIR-AP in the future may contact Alistair Moffat or Yiqun Liu by email; further information is available on the SIGIR-AP website.⁴

⁴<http://www.sigir-ap.org>

Acknowledgments

The organizers thank the Organizing Committee, Program Committee, reviewers, meta-reviewers, student volunteers, sponsors, authors, keynote speakers, tutorial and workshop organizers, industry speakers, and all participants for their contributions to SIGIR-AP 2025. We also thank Xi'an Kaili Conference and Exhibition Co., Ltd. for supporting conference management and on-site operations.

References

- Amin Bigdeli, Negar Arabzadeh, Ebrahim Bagheri, and Charles LA Clarke. Adversarial attacks against neural ranking models via in-context learning. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 211–220. ACM, 2025.
- Sunhao Dai, Chen Xu, Shicheng Xu, Zhongxiang Sun, Liang Pang, Zhenhua Dong, and Jun Xu. Trustworthy information retrieval in the LLM era: Bias, unfairness, and hallucination. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 442–446. ACM, 2025.
- Fernando Diaz. Preference-based evaluation. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 1–1. ACM, 2025.
- Jiangshu Du, Wenpeng Yin, and Philip Yu. Scaleformer: Span representation cumulation for long-context transformer. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 337–342. ACM, 2025.
- Tetsuya Sakai, Sijie Tao, Zhicheng Dou, Junjie Wang, Haoxiang Shi, Nuo Chen, and Atsushi Keyaki. BREV-RAG: Beyond relevance-based evaluation of RAG systems: A SIGIR-AP 2025 workshop proposal. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 459–460. ACM, 2025.
- Weihsang Su, Qian Dong, Qingyao Ai, and Yiqun Liu. Dynamic and parametric retrieval-augmented generation. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 453–458. ACM, 2025.
- Haitao Yu, Yubo Fang, Xuri Ge, Xin Xin, Zihan Wang, Junchen Fu, Joemon M. Jose, Weizhi Ma, and Zhaochun Ren. R3AG 2025: The second workshop on refined and reliable retrieval-augmented generation. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 461–464. ACM, 2025.

-
- Dell Zhang, Yue Feng, Haiming Liu, Changzhi Sun, Jixiang Luo, Xiangyu Chen, and Xuelong Li. Conversational agents: From RAG to LTM. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 447–452. ACM, 2025. doi: 10.1145/3767695.3769671.
- Min Zhang. LLMs vs. humans in information access tasks: Performances, behaviors, and learning abilities. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 177–177. ACM, 2025.
- Guido Zuccon, Shengyao Zhuang, Xueguang Ma, and Bevan Koopman. SIGIR-AP 2025 tutorial on retrieval and ranking with LLMs (R2LLMs). In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 437–441. ACM, 2025.