SIGIR 2026 Workshop on eCommerce (ECOM26)

Dean E. Alvarez
University of Illinois Urbana-Champaign
Urbana, Illinois, USA
deana3@illinois.edu

Aditya Chichani
Meta
Menlo Park, California, USA
aditya_chichani@berkeley.edu

Surya Kallumadi
Coursera
Atlanta, Georgia, USA
skallumadi@coursera.org

Yubin Kim
Vody
New York, NY, USA
yubin@vody.com

Tracy Holloway King
Adobe
San Jose, California, USA
tking@adobe.com

Andrew Trotman
University of Otago
Dunedin, New Zealand
andrew.trotman@otago.ac.nz

ABSTRACT

The eCommerce search and recommendations space is a unique, dynamic domain within information retrieval (IR), characterized by multimodality and industry-driven challenges. While the basic task of fulfilling a user's information need aligns with web search, the methodologies employed are distinct. On eCommerce platforms, the data available for retrieval and ranking differs significantly, as do the success signals (e.g. adding items to a cart, purchasing). The special theme of ECOM26 is *User Interaction and Experience: Agentic-driven Trends*. Our focus for 2026 is on fostering deeper engagement through interactive discussions, exploring crucial topics such as shifts in user interaction paradigms, and addressing emerging topics such as evaluation metrics for LLMs, multimodality, and the interplay between organic and sponsored search. With our discussion-heavy format and structured facilitation, we aim to spark conversation among all participants, beyond that of the usual interactions between presenters and audience questions.

CCS CONCEPTS

• **Applied computing** → **Electronic commerce**; • **Information systems** → **Information retrieval**.

KEYWORDS

eCommerce, search, recommendations, user experience, AI productionization

1 MOTIVATION

The eCommerce search and recommendations space is a unique, dynamic domain within information retrieval (IR), characterized by multimodality and industry-driven challenges. While the basic task of fulfilling a user's information need aligns with web search, the methodologies employed are distinct. On eCommerce platforms, the data available for retrieval and ranking differs significantly, as do the success signals (e.g. adding items to a cart, purchasing).

For product search, the entities to be discovered are multimodal combinations of unstructured text (e.g. titles, descriptions, reviews), images and videos, and structured data (e.g. price, brand, ratings). This complex interplay of data presents unique research challenges, such as developing recall and ranking functions that balance trade-offs across facets in response to queries. The features for constructing click models in eCommerce, such as queries, hover time, clicks, browse time, add-to-cart actions, purchases, and product returns, are richer than those in web search. Additionally, eCommerce introduces intricacies like personalized promotions, a highly dynamic inventory with frequent changes, and a long-tail query distribution.

The uniqueness of the eCommerce domain necessitates a venue to discuss domain-specific challenges and solutions. Previous iterations of this SIGIR ECOM workshops have consistently attracted strong participation from academic and industry practitioners, and from students to professors and senior applied scientists.

Recently, the rapid evolution of emerging technologies, including large language models (LLMs) and generative AI, has introduced fresh opportunities and challenges for eCommerce search and recommendations. As many organizations race to adopt new technology, the knowledge required to tackle the practical challenges risks becoming siloed. Avoiding this requires continued dialogue between IR academia and industry, and across different eCommerce organizations. The goal of our workshop is to bridge these gaps and foster collaboration for this crucial exchange of information.

To support this goal, the special theme of ECOM26 is *User Interaction and Experience: Agentic-driven Trends* (Section 2.2), and we have prepared a workshop format that heavily emphasizes conversations and collaborations through structured facilitation (Section 3). Our focus for 2026 is on fostering deeper engagement through interactive discussions, exploring crucial topics such as shifts in user interaction paradigms, and addressing emerging topics such as evaluation metrics for LLMs, multimodality, and the interplay between organic and sponsored search. With our discussion-heavy format and structured facilitation, we aim to spark conversation among all participants, beyond that of the usual interactions between presenters and audience questions.

2 THEME AND PURPOSE

The primary purpose of ECOM26 is to provide a venue for discussion and publication of IR research as it pertains to eCommerce, and to bring together practitioners from across academia and industry to produce thoughtful discussion and lasting collaborations.

SIGIR '26, July 20–24, 2026, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-2599-9/2026/07
https://doi.org/10.1145/3805712.3808650

2.1 Scope

The workshop relates to all aspects of eCommerce search and recommendation. Research topics and challenges that are frequently encountered in this domain include:

- **User Interaction and Experience: Agentic-driven Trends (2026 Special Theme)**: Agentic search, browsing, and discovery experiences in eCommerce; Trade-offs between accuracy, cost, latency, and scalability in agentic pipelines; User interaction patterns and UX considerations for agentic and conversational systems; Offline and online evaluation of agentic systems; Mobile vs. desktop experiences for agentic and conversational interfaces
- **IR Fundamentals for eCommerce**: Personalization and contextualization; Unified and universal search and recommendations; Indexing and search in rapidly changing environments (e.g. auction sites); Experimentation techniques Small Language Models (SLMs) for efficient and scalable eCommerce IR
- **Ads and Sponsored Products in eCommerce**: Sponsored search and recommendations; Joint optimization of organic and sponsored results; Auctions, bidding, and ranking for sponsored products; Ads impact on relevance and user experience
- **Ranking and Whole Page Relevance (WPR)**: Optimization for IR and business metrics; Diversity in product search and recommendations; Relevance models for multi-faceted entities; Temporal dynamics and seasonality
- **Query and Document Understanding**: Query intent, query suggestions, and auto-completion; Strategies for resolving low or zero recall queries; Cross-modal search (e.g. text, structured data, images); Categorization and facets
- **Representations and Data**: Semantic representation of products, queries, and customers; Construction and use of knowledge graphs for eCommerce; Large foundational recommendation models
- **Fairness and transparency at the system and ecosystem level**

2.2 Special theme: User Interaction and Experience — Agentic-driven Trends

We will also solicit participation on our 2026 theme. The success of eCommerce search and recommendation features hinges not only on the quality of the data and algorithms but crucially on the quality of how users interact with the overall search and recommendation experience. eCommerce saw major user experience shifts with the rise of mobile device usage for eCommerce. This led to new UXs for eCommerce and to new algorithms which focused on precision and diversity among the top few results and which accounted for user preferences across device types. A new shift is occurring with the introduction of agents in eCommerce workflows. ECOM26's workshop theme of user interaction and experience will promote discussion on where we see eCommerce search and recommendations shifting with the advent of agentic experiences. Questions we expect to discuss include:

- How will we guide users from current eCommerce UX to more agentic experiences?
- What will browse and discovery experiences, which are core to the long term success of eCommerce, look like?
- AB testing is a cornerstone of eCommerce, how does AB testing need to be adapted for new UI paradigms?
- How can the UX be adapted to provide a personalized experience while maintaining user trust and a fair marketplace?

In addition to featuring invited speakers and panelists who are experts in this area, we will solicit papers from academia and industry focused on this special theme. We welcome papers on all aspects of this theme from classic HCI research and how to adapt eCommerce algorithms to best practices for eCommerce UX experimentation and position papers on eCommerce user interaction.

To ensure that these learnings are widely shared beyond our workshop, we will publish the proceedings through CEUR Workshop Proceedings archive.[1]

3 FORMAT

ECOM26 is a **full day workshop** and formatted to encourage deep discussion. We limit contributed talks to a select few top papers, with most accepted papers presented as posters. Through well-structured facilitation, we seek to generate discussion between *all* participants. This will be a highly interactive workshop, not a mini-conference. Our workshop format will include:

- Two invited talks
- Moderated panel discussion on the workshop theme
- Selected long contributed talks with facilitated discussions
- Poster boaster presentations followed by a poster session
- Playfully themed lunch tables for casual on-topic networking
- *Agents Anonymous* group activity

3.1 Panel Discussion

The panel session will reflect the special theme *User Interaction and Experience: Agentic-driven Trends*. Panelists across industry and academia will share lessons learned from major UX paradigm shifts in the past and from current experiences with integrating agentic-like interactions in their UX. The panel session will begin with introductions and 2–3 questions from the moderator, which the panel will answer in round robin. The session will then move to an audience Q&A, followed by a final closing question from the moderator. This panel format was well-received at prior workshops and resulted in active audience participation and discussion.

3.2 Data Challenge

ECOM26 will host a data challenge by building on the Search task from the TREC Product Search and Recommendations track on query reformulation for task oriented eCommerce queries.[2] The task targets conversational search where users express a high level goal rather than a specific item and ask the system to generate reformulated queries that make implicit product requirements explicit so that the IR system can retrieve more useful products. It supports an automatic variant that produces one reformulation per query and an interactive variant that produces up to four reformulations per query, scored by the best performing option. This connects directly to agentic search and discovery where an agent can propose multiple interpretations, surface choices to the user,

[1] https://ceur-ws.org/
[2] https://trec-product-search.github.io/search

and use lightweight feedback to refine the search trajectory. The challenge can leverage the released product metadata collection, a BM25-based Search API for testing, and training and development data with human annotated reformulations and relevance judgments. Evaluation centers on task completion oriented measures such as Task Completion nDCG and product recall at K. By grounding the ECOM26 challenge in the TREC benchmark while explicitly encouraging agent mediated interaction patterns, we aim to support reproducible experimentation and stimulate new approaches to user centered product discovery.

3.3 Contributed Talks and Posters

Select top papers will be given 10–15 min oral presentation slots followed by 15 mins of group discussions. Rather than a typical Q&A period, this discussion period will be facilitated by an organizer to encourage participants to contribute ideas and talk to each other. The discussion will be seeded with open-ended questions such as: "How were your experiences similar/dissimilar?", "What adaptations do you think you may need for your problem?" The discussion will be moderated by the facilitator using the progressive stack,[3] a system that prioritizes underrepresented voices. First time speakers will be prioritized. In addition, based on the diversity of existing contributions, the facilitator may solicit certain groups to "step forward", e.g. if most contributions have been from large tech companies, start-ups and academics will be encouraged.

All accepted papers will have poster presentation. Papers which are not part of the long presentations will have a lightning, 2 min "poster boaster" presentation before the poster session, which will give all authors a presentation opportunity and generate interest for their work.

3.4 Themed Lunch Tables

During lunch, we will set up playfully themed tables to help participants meet others working on similar problems and encourage conversations around shared points of interest. The tables will be organized around broad problem categories, e.g. "The Multiverse" (Multimodality); "Treasure Cove" (Sponsored Search); "The Launchpad" (Start-ups); "Mt. Everest" (Large Scale); "Just Do It" (Agents).

3.5 Agents Anonymous

Agents Anonymous is an interactive, "group therapy"-style activity designed to facilitate candid discussion of hard-earned lessons, open questions, and frustrations of building real-world systems. Prior to the workshop, we will invite participants to submit 1-minute problem pitches describing a challenge they are currently grappling with. We will select 4–5 cases. Based on the selected cases, we will invite a panel of experts to provide quality feedback. During the workshop, each problem will be discussed in a structured, time-boxed format that includes a problem pitch, clarifying questions, expert feedback, open group brainstorming, and facilitated synthesis. The emphasis is on sharing patterns, heuristics, and battle-tested experience to create space for insights that are rarely discussed in formal academic settings and create a sense of shared community.

[3] https://en.wikipedia.org/wiki/Progressive_stack

4 DISTINCTION FROM MAIN CONFERENCE

The SIGIR eCom workshop remains timely because its focus is complementary to the main conference program and targets questions that often receive limited coverage in standard research tracks. Our special theme emphasizes user interaction and experience in agentic driven eCommerce, including conversational and multimodal shopping journeys, user control and transparency, trust and calibration, and feedback loops that connect interaction design to retrieval and recommendation quality. These topics sit at the intersection of information retrieval, human computer interaction, behavioral science, and product design, and they raise evaluation and measurement challenges that differ from typical offline benchmarks. The workshop also provides a practical venue for discussing deployment realities such as latency, safety, policy constraints, privacy, and experimentation at scale, which are essential to understanding real user impact but are not always foregrounded in the main program. Finally, the workshop creates space for sharing ideas and lessons that can be difficult to publish in full detail due to industry restrictions on data access, proprietary systems, and limits on what can be publicly disclosed, enabling candid discussion, anonymized case studies, and constructive industry-academia collaboration.

5 ORGANIZERS

Dean Alvarez is a third-year PhD student at the University of Illinois, Urbana-Champaign, where he is advised by Professor ChengXiang Zhai. His research focuses on information discovery and organization, with a specific emphasis on graph representations of text data. He earned his undergraduate degree in Computer Science from Cornell University. In addition to his research, he served as a co-organizer for the 2026 TREC Product Search and Recommendations track.

Aditya Chichani is a Machine Learning Engineer at Meta focusing on product-centric ads ranking. Previously, he worked as a Senior Data Scientist at Walmart focusing on search ranking. Aditya completed his Master's from University of California Berkeley with a Fung Excellence Scholarship, currently serves on the program committee of RecSys and has organized workshops at top conferences such as CIKM, SIGIR and ICDM.

Dr. Surya Kallumadi is a Director of Machine Learning at Coursera, with a focus on core search, recommendations, and personalization. In the past he has worked on eCommerce discovery at Home Depot, Lowe's, Flipkart and eBay. Surya has served as an organizer for multiple workshops at conferences such as SIGIR, ACL, and IJCAI.

Dr. Yubin Kim is the Chief Science Officer of Vody, a GenAI startup in the eCommerce space focused on enriching product and buyer representations using multimodal LLMs. She previously worked on large-scale search in eCommerce (Etsy) and healthcare (UPMC). Yubin received her Ph.D. at the Language Technologies Institute in Carnegie Mellon University. She serves on the program committees for top IR conferences and journals, and currently serves on the editorial board of FnTIR. She has organized workshops and scientific events at conferences such as WSDM, CIKM, and SIGIR.

Dr. Tracy Holloway King is a Senior Principal Scientist focusing on multimodal search science in Adobe's Search & Discovery team. After getting her Ph.D. from Stanford, she was an NLP researcher

in PARC for over a decade. She then moved to applied science roles at Microsoft Bing, eBay Search Science, and Amazon Query Understanding before joining Adobe. She has co-organized workshops at ACL, COLING, and SIGIR, and co-organized SIRIP 2023.

Dr. Andrew Trotman is an Associate Professor at the University of Otago in New Zealand. He previously worked on eCommerce at eBay. He has co-chaired 10 SIGIR workshops, co-chaired INEX, was co-chair of SIGIR 2014, was co-PC chair of SIGIR 2025, and has co-chaired and co-PC chaired several ADCSs. Andrew is a senior member of the SIGIR Artifact Evaluation Committee, is a previous member of the SIGIR executive committee (secretary), was on the editorial board of the IRJ, and IP&M.

6 TARGET AUDIENCE

This workshop brings together researchers and practitioners from academia and industry to identify and discuss core research problems in eCommerce search and recommendation. Attendees are expected to include:

- Early-, mid-, and late-career industry professionals from large and small eCommerce companies
- Professors and PhD students of IR, NLP, multi-modal, and economics

7 RELATED MATERIALS AND WORKSHOPS

Past SIGIR ECOM Workshops: Why Another Workshop? Interest in ECOM by the SIGIR community has increased over the years, as reflected by the growth in participation and submissions. This stems from multiple factors. (1) eCommerce is a complex, unsolved search domain involving multimodal documents and business-driven objective functions. (2) The real-world applications of eCommerce search provide many applied science job opportunities for IR graduates and many collaboration opportunities for IR faculty: ECOM allows them to showcase their work and network. (3) Rapid evolution in technologies, especially in LLMs and generative AI, have created new challenges as practitioners adapt these tools to solve problems and unlock new user experiences in eCommerce. The SIGIR ECOM workshop typically receives participation from 100–200 delegates. Prior iterations of this workshop were held at SIGIR 2017–2025; records of previous workshops are available at https://sigir-ecom.github.io/ (except 2017).

We expect many SIGIR participants to be interested in the special theme on user interaction and experience with agentic trends in eCommerce. This rapidly changing area brings new challenges for designing user experiences, collecting and using feedback, and evaluating systems in realistic settings. It also raises practical questions about how to turn these ideas into reliable product experiences, an area not fully covered in the main conference.

eCommerce Search Overview Materials. An eCommerce search overview [1] and the SIGIR 2018 tutorial "Information Discovery in E-Commerce"[4]) provide background to participants.

Related Workshops. The IJCAI-17 Workshop on AI Applications in E-Commerce (WAAE 2017) dealt with some of the research areas that our workshop will cover; however the scope of WAAE was much broader covering all of AI as opposed to search and recommendations. In May 2019 The Web Conference held the "First Workshop on e-Commerce and NLP"[5] which covered a complementary field, NLP, and was followed by two workshops in 2020 at WebConf[6] and ACL[7] and continues to be held. SIGIR eCOM workshops have been successful at SIGIRs since 2017, with the proceedings held by CEUR. The ACM Special Interest Group on E-commerce (SIGecom) runs an annual conference. However, search has never been well represented. SIGIR ECOM addresses this shortfall by examining eCommerce search and recommendation.

REFERENCES

[1] Manos Tsagkias, Tracy Holloway King, Surya Kallumadi, Vanessa Murdock, and Maarten de Rijke. 2020. Challenges and Research Opportunities in eCommerce Search and Recommendations. *SIGIR Forum* 54, 1 (June 2020), 1–23.

[4] https://sites.google.com/view/sigir2018-info-ec/home

[5] https://sites.google.com/view/ecnlp/past-workshops/ecnlp-2019

[6] https://sites.google.com/view/ecnlp/past-workshops/www-2020

[7] https://sites.google.com/view/ecnlp/past-workshops/acl-2020