

SIGIR 2025 Workshop on eCommerce (ECOM25)

From Research to Product: Challenges, Lessons, and Opportunities in eCommerce Search and Recommendations

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

Pallavi Gudipati
Upwork
San Francisco, California, USA
pallavigudipati@upwork.com

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

Aditya Chichani
Walmart
Sunnyvale, California, USA
aditya_chichani@berkeley.edu

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

Abstract

The eCommerce search and recommendations space is a unique and dynamic domain within information retrieval (IR), characterized by its 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 (e.g. Alibaba, Amazon, eBay, Etsy, Flipkart, Walmart), the data available for retrieval and ranking differs significantly, as do the success signals (e.g. adding items to a cart, purchasing). Our focus for 2025 is on fostering deeper engagement through interactive discussions, exploring crucial topics such as navigating irreproducibility in research-to-product pipelines, 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.

CCS Concepts

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

Keywords

e-commerce, search, recommendations, irreproducibility, AI productionization

ACM Reference Format:

Yubin Kim, Tracy Holloway King, Aditya Chichani, Pallavi Gudipati, and Andrew Trotman. 2025. SIGIR 2025 Workshop on eCommerce (ECOM25): From Research to Product: Challenges, Lessons, and Opportunities in eCommerce Search and Recommendations. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3726302.3730359>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '25, Padua, Italy

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1592-1/2025/07
<https://doi.org/10.1145/3726302.3730359>

SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25), July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 3 pages.
<https://doi.org/10.1145/3726302.3730359>

1 Motivation

The eCommerce search and recommendations space is a unique and dynamic domain within information retrieval (IR), characterized by its 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 (e.g. Alibaba, Amazon, eBay, Etsy, Flipkart, Walmart), 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, popularity, revenue). This complex interplay of data presents unique research challenges, such as developing recall (matching) and ranking functions that balance trade-offs across facets in response to user queries. The features for constructing click models in eCommerce, such as queries, hover time, clicks, browse time, add-to-cart actions, purchases, side-by-side comparisons, remove-from-cart actions, and product returns, are richer and more complex than those in web search. Additionally, eCommerce introduces intricacies like personalized promotions and pricing, a highly dynamic inventory with frequent changes, and a long-tail query distribution.

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 organizations race to adopt new technology, the knowledge required to tackle the practical challenges of productionization risks becoming siloed. Avoiding this requires dialogue between IR communities in academia and industry and across 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 ECOM25 is *From Research to Product: Challenges, Lessons, and Opportunities in eCommerce Search and Recommendations*, and we have prepared a workshop format that heavily emphasizes conversations and collaborations through structured facilitation.

Our focus for 2025 is on fostering deeper engagement through interactive discussions, exploring crucial topics such as navigating irreproducibility in research-to-product pipelines, 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 the audience in a limited question period. Furthermore, we will capture the outcomes of the workshop and publish a SIGIR Forum paper as a “five years later” follow-up to the 2020 paper on challenges and opportunities in eCommerce search and recommendations [1], authored by previous ECOM workshop organizers.

The special theme of this year (*From Research to Product: Challenges, Lessons, and Opportunities in eCommerce Search and Recommendations*) aligns with both long-standing and current interests of the SIGIR community. Simultaneously, the special theme is complementary to the main conference, as we hope to capture practical techniques and learnings that are often difficult to publish in standard academic venues.

2 Theme and Purpose

The primary purpose of ECOM25 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. The workshop relates to all aspects of eCommerce search and recommendation. Research topics and challenges that are frequently encountered in this domain include:

- Ranking and Whole Page Relevance (WPR)
 - Optimization for IR and business metrics
 - Diversity in product search and recommendations
 - Relevance models for multi-faceted entities
 - Relevance vs. revenue
 - Deterministic sorts (e.g. price low to high)
 - 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
 - Reviews and sentiment analysis
- Recommendation and Personalization
 - Personalization & contextualization
 - Privacy, bias and ethics in eCommerce IR
 - Blending recommendations, sponsored products and search results
- Representations and Data
 - Semantic representation of products, queries, and customers
 - Construction and use of knowledge graphs for eCommerce
- IR Fundamentals for eCommerce
 - Unified and universal search and recommendations
 - Cross-lingual search and machine translation
 - Indexing and search in rapidly changing environments (e.g. auction sites)

- Experimentation techniques
- Other challenges
 - Trust, transparency, and fairness in eCommerce
 - UX for eCommerce
 - Question answering and chat bots for eCommerce

Theme: From Research to Product: Challenges, Lessons, and Opportunities in eCommerce Search and Recommendations

There is a major gap between state-of-the-art research results and effectively implementing them in eCommerce production. This year’s special theme addresses the challenges, lessons, and opportunities encountered while transitioning cutting-edge research into practical applications within eCommerce. The workshop will include invited speakers and panelists expert in this area and position papers from academia and industry, emphasizing:

- The practical challenges of translating research into scalable products.
- Lessons learned from tackling irreproducibility while implementing key research papers into industrial eCommerce search and recommendation systems.
- Emerging trends in eCommerce IR, such as multimodality, customer lifetime value, and online and offline usage of LLMs.

We believe that this is a particularly timely special theme. Over the past few years, many organizations have invested in productionizing new technology (e.g. LLMs, generative AI) to solve business problems. Although productionizing research papers requires overcoming substantial challenges with creative problem solving, the associated learnings can be difficult to publish in standard academic venues because they lack the right kind of novelty, e.g. clever heuristics, knowledge of common pitfalls, or negative results. However, that “one weird trick” that fixes a common issue or knowing what “obviously good” idea *doesn’t* work and why, can save precious research time. ECOM25 will provide a platform to share these invaluable from-the-trenches experiences across organization boundaries to enrich our community and reduce repetitive effort. We will capture the discussions in a SIGIR Forum paper as a 5 year follow-up to the eCommerce challenges paper [1].

3 Workshop Format

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, and economics.

ECOM25 is formatted to prioritize and encourage deep discussion. Through well-structured facilitation, we seek to move beyond the presenter/audience dynamic and generate discussion between *all* participants. This will be a highly interactive workshop, not a mini-conference with: Two invited talks; Moderated panel discussion on the workshop theme; Selected long contributed talks with facilitated discussions; Poster boaster presentations followed by the

poster session; Playfully themed lunch tables for casual networking; "From Research to Product" breakout group challenge.

Panel Discussion The panel will reflect the theme *From Research to Product*. Panelists from industry and academia will share lessons learned from applying state-of-the-art research and technology to eCommerce, with a special focus on irreproducibility. Given that negative results are under-reported in publications, we find value in providing a forum to discuss what does *not* work.

Themed Lunch Tables During lunch, we will set up playfully themed lunch 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 and/or job roles, e.g. "The Multiverse" (Multimodality); "Treasure Cove" (Sponsored Search); "The Launchpad" (Start-ups); "Mt. Everest" (Large Scale).

Breakout Group Challenge Participants will break into groups to collaborate on a "From Research to Product" guide for a topic of their choice and produce the kind of guidebook they wished they had. Organizers will provide a list of potential topics (e.g. hybridizing sparse and dense retrieval, image-text multimodal search, LLMs for synthetic data generation), but attendees are encouraged to propose topics. Each group will be assigned a facilitator and will present their guidebook to the attendees. The slides will be publicly shared and will be summarized in our SIGIR Forum paper.

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.