

# SIGIR 2024 Workshop on eCommerce (ECOM24)

Surya Kallumadi

Lowe's Companies  
Atlanta, GA, USA  
surya@ksu.edu

Yubin Kim

Vody  
Newport, CA, USA  
yubink.cs@gmail.com

Tracy Holloway King

Adobe  
San Jose, CA, USA  
tking@adobe.com

Maarten de Rijke

University of Amsterdam  
Amsterdam, Netherlands  
m.derijke@uva.nl

Vamsi Salaka

Amazon  
Palo Alto, CA, USA  
vsalaka@amazon.com

## ABSTRACT

ECOM24 brings together researchers and practitioners from academia and industry to identify and discuss core research problems in eCommerce search and recommendation. The workshop aims to foster collaboration, to attract research funding, and to introduce IR researchers and postgraduate students to eCommerce product discovery. The workshop features a special theme of eCommerce search in the age of Generative AI and LLMs and a data challenge in collaboration with TREC on how end-to-end retrieval systems can be built and evaluated given a large set of products.

## CCS CONCEPTS

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

## KEYWORDS

eCommerce information retrieval, product search, recommender systems, metrics, evaluation

### ACM Reference Format:

Surya Kallumadi, Yubin Kim, Tracy Holloway King, Maarten de Rijke, and Vamsi Salaka. 2024. SIGIR 2024 Workshop on eCommerce (ECOM24). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3626772.3657988>

## 1 MOTIVATION FOR THE WORKSHOP

Search and recommendation have applications ranging from traditional web search to document collections to vertical search systems. This workshop explores approaches to search and recommendation of products in eCommerce. Although the basic search task (i.e., fulfill a user's information need) is the same as web search, the way in which this is achieved is different. On eCommerce sites (e.g.,

Alibaba, Amazon, eBay, Etsy, Flipkart, Walmart), the data available for retrieval and ranking are different as are the signals of success (e.g., adding items to a cart, purchasing).

For product search, the entities to be discovered are combinations of unstructured text (e.g., titles, descriptions, reviews), images, and structured data (e.g., price, brand, ratings, popularity, revenue). This complex combination of data raises interesting research challenges including recall (matching) and ranking functions that take into account the trade-offs across facets with respect to the input query. The features available for building click models are different, and often stronger, in eCommerce than in web search. As well as queries, hover time, clicks, and browse time, eCommerce sites also have signals from add-to-cart, purchase, side-by-side comparison, remove-from-cart, return of goods, etc. When incorporating promotions and personalization such as individual pricing, the click models are more complex than in web search. eCommerce is also characterized by a dynamic inventory with a high rate of change and turnover, and a very long tail of the query distribution.

This workshop brings together researchers and practitioners from academia and industry to identify and discuss core research problems in eCommerce search and recommendation. The workshop aims to foster collaboration by bringing the community together in a way that rarely happens, to attract research funding to this increasing important domain, and to introduce IR researchers and postgraduate students to eCommerce and product discovery. Finally, it helps broaden the definition of IR at research venues such as SIGIR. To support these goals, the workshop features a special theme as well as a data challenge.

- **Special theme:** eCommerce search in the age of Generative AI and LLMs (Section 2.2)
- **Data challenge:** The workshop supports data availability for eCommerce IR research. In 2018, 2019, 2020, 2021, 2022 we released data from Rakuten, eBay, Coveo and Farfetch. This year, we collaborated with TREC on a product search data challenge (Section 2.3).

## 2 THEME AND PURPOSE

The primary theme of the workshop is eCommerce search and recommendation. The purpose of the workshop is to provide a venue for discussion and publication of IR research as it pertains to eCommerce. The workshop brings together practitioners, researchers, and applied researchers from academia and industry to discuss the challenges and approaches to eCommerce search

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](mailto:permissions@acm.org).

SIGIR '24, July 14–18, 2024, Washington, DC, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0431-4/24/07

<https://doi.org/10.1145/3626772.3657988>

and recommendation. We aim to foster collaboration and discussion with the broader IR community and to raise awareness within the academic community of the unique challenges faced by the eCommerce domain.

## 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:

- eCommerce search in the age of Generative AI and LLMs (2024 special theme)
- Ranking and Whole Page Relevance
  - 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
  - Converting across modalities (e.g., text, structured data, images)
  - Categorization and facets
  - Reviews and sentiment analysis
- Recommendation and Personalization
  - Personalization & contextualization, including the use of personal facets such as age, gender, location
  - Privacy, bias and ethics in eCommerce IR
  - Blending recommendations 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 including AB testing and multi-armed bandits
- Other challenges
  - Trust, transparency, and fairness in eCommerce
  - UX for eCommerce
  - The role of search in trust and security for marketplaces
  - Question answering and chat bots for eCommerce

## 2.2 Special Theme: eCommerce search in the age of Generative AI and LLMs

Last year's theme focused on the use of foundation models including large language models (LLMs) and multi-modal models for universal search and recommendation tasks. LLMs bring us closer to the goal of a universal model that can effectively handle a wide range of information access tasks, such as retrieval, recommendations, and question answering [1]. This year we build on these innovations to focus on how generative AI and LLMs are changing eCommerce. These include query-time capabilities, such as accurate answering

of questions about products (i.e., without hallucinations), chatbots for customer support and stateful product search, and improved results for natural language queries. We also consider offline use cases of LLMs, such as improved product understanding to create taxonomies and metadata for left-rail filters, and automated relevance judgements for improved eCommerce ranker training.

The workshop discusses the challenges and limitations of these approaches, and how they can be overcome either by improved leveraging of LLMs or by hybridization with existing IR techniques (e.g., for model number searches).

## 2.3 Data Challenge

To facilitate community building in the eCommerce information retrieval domain, we have been organizing data challenges in conjunction with this workshop. In the past, eCommerce companies such as Rakuten (2018, 2020), eBay (2019), Coveo (2021), and Farfetch (2022) supported this initiative by organizing data challenges. These data challenges spanned a variety of topics including taxonomy classification, session-oriented personalization, and a high accuracy recall task. They have led to increased interest in the eCommerce domain with participation from diverse groups across industry and academia and to building a community of researchers interested in this domain. Even after the end of the challenge, the datasets continue to be downloaded and used for benchmarking, educational and product development purposes.

This year we collaborated with TREC on a product search data challenge (<https://trec-product-search.github.io/index.html>). The over-arching goal is to study how end-to-end retrieval systems can be built and evaluated given a large set of products. The data challenge provides a corpus of products and a set of user intents (queries): the goal is to find the product that suits the user's need.

The TREC 2024 dataset builds on the 2023 TREC Product Search Track [2] which itself is based on the ESCI Challenge for Improving Product Search.<sup>1</sup> In 2024, the focus is on exploring methods for generating synthetic queries via simulation and leveraging large language models (LLMs). This focus aligns with the ECOM24 workshop theme (Section 2.2). Given the challenge product corpus, there are three tasks based on different modes of retrieval: re-ranking, text-only retrieval, and multi-modal retrieval. (1) For re-ranking, given an initial set of 1000 documents for each query extracted using a BM25 baseline, research groups re-rank the existing results using any modeling approach. (2) For text-only retrieval, the task formulation is much like re-ranking but without the top 1000 documents as these must be generated by the research group using the text for each product. (3) For multimodal retrieval, the usage of additional information in the form of images, reviews, and product taxonomy can be used to improve retrieval or ranking performance.

## 2.4 Workshop Outcomes

We believe that the most important outcome of the workshop is the discussion between individual participants at the workshop. It is these discussions that lead to collaboration across institutions, including across academia and industry, and future research. We will capture the discussions in a SIGIR Forum workshop report and

<sup>1</sup><https://amazonkddcup.github.io/>

will produce a proceedings of the workshop and publish it through CEUR Workshop Proceedings for archival<sup>2</sup>.

This year's workshop theme of eCommerce search in the age of Generative AI and LLMs, and the data challenge with text-only and multi-modal data for recall and ranking are of increasing importance to eCommerce. We hope that through discussion at the workshop and, more generally, at SIGIR 2024, we can help steer the research community towards these problems and in doing so work towards solutions.

### 3 FORMAT

To encourage maximum participation and discussion, significant time in each session is allotted for discussion and the chairs and workshop organizers act as facilitators. Format:

- 2 invited talks (1 from academia, 1 from industry)
- Long (15 min) contributed talks
- Lightning (5 min) contributed talks with discussion session
- Data challenge overview and discussion
- Panel discussion on the workshop theme

### 4 TARGET AUDIENCE

The workshop brings together researchers and practitioners from academia and industry to identify and discuss core research problems in eCommerce search and recommendation. It aims to foster collaboration by bringing the community together in a way that rarely happens. Similar to past ECOM workshops, 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

### REFERENCES

- [1] Nicholas J. Belkin and W. Bruce Croft. 1992. Information filtering and information retrieval: two sides of the same coin? *Commun. ACM* 35 (1992), 29–38.
- [2] Daniel Campos, Surya Kallumadi, Corby Rosset, Cheng Xiang Zhai, and Alessandro Magnani. 2023. Overview of the TREC 2023 Product Search Track. (2023). arXiv:cs.IR/2311.07861

---

<sup>2</sup><https://ceur-ws.org/>