eCom'23: The SIGIR 2023 Workshop on eCommerce

Surya Kallumadi
Lowe's Companies
surya@ksu.edu

Yubin Kim
Etsy
ykim@etsy.com

Tracy Holloway King
Adobe Inc.
tking@adobe.com

Shervin Malmasi
Amazon
malmasi@amazon.com

Maarten de Rijke
University of Amsterdam
m.derijke@uva.nl

Jacopo Tagliabue
New York University, Bauplan
jacopo.tagliabue@nyu.edu

ABSTRACT

eCommerce Information Retrieval (IR) is receiving increasing attention in the academic literature and is an essential component of some of the largest web sites (e.g. Airbnb, Alibaba, Amazon, eBay, Facebook, Flipkart, Lowes's, Taobao, Target). SIGIR has for several years seen sponsorship from eCommerce organizations, reflecting the importance of IR research to them. The purpose of this workshop is (1) to bring together researchers and practitioners of eCommerce IR to discuss topics unique to it, (2) to determine how to use eCommerce's unique combination of free text, structured data, and customer behavior data to improve search relevance, and (3) to examine how to build datasets and evaluate algorithms in this domain.

The theme of this year's eCommerce IR workshop is **Foundation Models and Unified Information Access in eCommerce**. The workshop solicits papers on this topic and includes a panel focused on this area (see Section 2). In addition, Lowe's is sponsoring an eCommerce data challenge on Cross-modal and Multi-modal Visual Search for eCommerce. The data challenge reflects themes from the successful SIGIR workshops in 2017, 2018, 2019, 2020, 2021, and 2022. ECOM23 will be held as a full day hybrid workshop to accommodate for diverse participation.

CCS CONCEPTS

• **Information systems** → **Online shopping**; **Specialized information retrieval**;

KEYWORDS

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

ACM Reference Format:
Surya Kallumadi, Yubin Kim, Tracy Holloway King, Shervin Malmasi, Maarten de Rijke, and Jacopo Tagliabue. 2023. eCom'23: The SIGIR 2023 Workshop on eCommerce. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23), July 23–27, 2023, Taipei, Taiwan.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3539618.3591927

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

ACM ISBN 978-1-4503-9408-6/23/07.
https://doi.org/10.1145/3539618.3591927

Search and recommendation have applications ranging from traditional web search to document collections to vertical search systems. In this workshop, we explore approaches to search and recommendation of products in eCommerce IR. Although the basic search task (i.e. fulfilling a user's information need) is the same as in web search, the way in which this is achieved is different [13]. 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).

The entities that need 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 used in ranking 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 increasingly important domain, and to introduce IR researchers and postgraduate students to eCommerce and product discovery. Finally, it will help broaden the definition of IR at research venues such as SIGIR. To support these goals, the workshop features a **data challenge** and a **special theme**.

1 THEME AND PURPOSE

The overarching 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. We will bring together practitioners, researchers, and applied researchers from academia and industry to discuss the challenges and approaches to eCommerce search 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.

1.1 Scope

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

- 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 and 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 (special theme of the workshop for 2023)
  - 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

1.2 Special Theme: Foundation Models and Unified Information Access in eCommerce

For a long time, researchers in the field of information retrieval have tried to develop a universal model that can effectively handle a wide range of information access tasks, such as retrieval, recommendations, and question answering [1]. These tasks are often carried out using different information access systems, as seen on eCommerce websites that expose both search and recommendation functionalities to users. As part of this, specialized models are often developed for each type of information access scenario. However, recent works have suggested that jointly optimizing and modeling search engines and collaborative filtering could lead to better generalization and improve the quality of both search and recommendation results [14, 15]. Moreover, large multi-modal models trained on high-quality eCommerce datasets seem to be able to become domain experts without losing their zero-shot capabilities [2]. As part of the special theme, we will delve into the use of foundation models, as well as the application of large language models, such as transformer-based models, for universal search and recommendation tasks. We will also discuss the challenges and limitations of these approaches, and how they can be overcome.

In addition to having invited speakers who are experts in this area for ECOM'23, we solicit position papers from academia and industry focused on this area.

1.3 Data Challenge

In order to facilitate community building in the eCommerce information retrieval domain, each year we have organized a data challenge 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.[1] These data challenges spanned a wide variety of topics including taxonomy classification, session-oriented personalization, complete-the-look recommendations and high accuracy recall. These data challenges have led to increased interest in the eCommerce domain with participation from diverse groups across industry and academia. As a result, they have helped build a community of researchers interested in this domain. Often, even after the end of the challenge, the dataset made available to the community continues to be frequently used for benchmarking [3], educational and product development purposes.

For the sixth year in a row we have worked with industry researchers to create a new data challenge: the *Cross-modal and Multi-modal Visual Search for eCommerce* data challenge, supported by Lowe's. The goal of this data challenge is to encourage the development of innovative approaches for cross-modal and multi-modal product search using visual and textual information. eCommerce platforms rely heavily on text-based search to allow users to find products they are interested in. However, there are often cases where users may not know the exact wording to describe the product they are looking for or may prefer to use visual information, such as an image of the product, to initiate their search. Moreover the image uploaded by a user often does not look similar to the professional images available in the product catalogue.

This data challenge aims to address these issues by providing a dataset of images and associated product metadata for participants to use in developing and evaluating cross-modal and multi-modal product search systems. To solve this important use case, a new dataset has been curated, starting from existing large collections of products [7], shopping behavior and search interactions [11]: the final dataset contains $\sim 1,400,000$ products, including multiple images for each SKU. The main goal for the challenge is to make progress in the *image-to-product retrieval task*: participants will be asked to train models that, given an image as input (often taken under real-world conditions), will retrieve relevant SKUs from the full list of products. Successful models will be able to correctly match the input image to its original product, and the evaluation will follow standard Information Retrieval best practices.

[1] For more information, refer to materials published by each data challenge: https://sigir-ecom.github.io/files/Rakuten_Data_Challenge.pdf, https://github.com/eBay/sigir-2019-ecom-challenge, [12].

To run the competition, we will leverage the same tools and processes that proved to be successful in the past. The organizers will maintain an online leaderboard and a slack channel for questions, clarifications, announcements etc. We will invite paper submissions from all participants. The accepted contributions as well as the winners will prerecord a presentation on their submission and will participate in a session dedicated to the data challenge.

1.4 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 to future research. We will capture what we can in the form of a SIGIR Forum workshop report and, as with previous workshops, we will produce a proceedings of the workshop and work with CEUR Workshop Proceedings to ensure they are appropriately archived.

This year's workshop theme of unified information access and the cross-modal and multi-modal data challenge are of increasing importance to eCommerce. We hope that through discussion at the workshop and, more generally, at SIGIR 2023, we can help steer the research community towards these problems and in doing so work towards solutions.

2 WORKSHOP FORMAT

ECOM'22 was held in a hybrid fashion with a format to encourage maximum participation and discussion. In each session, significant time was allotted for discussion and the chairs and workshop organizers acted as facilitators. ECOM'23 will be held as a **full day hybrid workshop** to accommodate for diverse participation.

Overview of format:

- 2 invited talks (1 from academia, 1 from industry)
- Contributed talks (3–4 full oral presentations)
- Discussion sessions: Discussions organized by themes with lightning talks by all the accepted papers
- Data challenge session: Data challenge overview and discussion
- Panel discussion on the workshop theme: Foundation Models and Unified Information Access in eCommerce

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] Patrick Chia, Giuseppe Attanasio, Federico Bianchi, Silvia Terragni, Ana Magalhães, Diogo Goncalves, Ciro Greco, and Jacopo Tagliabue. 2022. Contrastive language and vision learning of general fashion concepts. *Scientific Reports* 12 (11 2022). https://doi.org/10.1038/s41598-022-23052-9

[3] Patrick John Chia, Jacopo Tagliabue, Federico Bianchi, Chloe He, and Brian Ko. 2022. Beyond NDCG: Behavioral Testing of Recommender Systems with RecList. In *Companion Proceedings of the Web Conference 2022 (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 99–104. https://doi.org/10.1145/3487553.3524215

[4] Jon Degenhardt, Giuseppe Di Fabbrizio, Surya Kallumadi, Mohit Kumar, Andrew Trotman, Yiu-Chang Lin, and Huasha Zhao. 2018. *Proceedings of ECOM18: The SIGIR 2018 Workshop on eCommerce*. CEUR.

[5] Jon Degenhardt, Surya Kallumadi, Maarten de Rijke, Luo Si, Andrew Trotman, and Yinghui Xu. 2017. *Proceedings of ECOM17: The SIGIR 2017 Workshop on eCommerce*. CEUR.

[6] Jon Degenhardt, Surya Kallumadi, Utkarsh Porwal, and Andrew Trotman. 2019. *Proceedings of ECOM19: The SIGIR 2019 Workshop on eCommerce*. CEUR.

[7] Ruining He and Julian McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 507–517. https://doi.org/10.1145/2872427.2883037

[8] Dietmar Jannach, Surya Kallumadi, Tracy Holloway King, Weihua Luo, and Shervin Malmasi. 2020. *Proceedings of ECOM20: The SIGIR 2020 Workshop on eCommerce*. CEUR.

[9] Ajinkya Kale, Surya Kallumadi, Tracy Holloway King, Shervin Malmasi, Maarten de Rijke, and Jacopo Tagliabue. 2022. *Proceedings of ECOM22: The SIGIR 2022 Workshop on eCommerce*. CEUR.

[10] Surya Kallumadi, Tracy Holloway King, Shervin Malmasi, and Maarten de Rijke. 2021. *Proceedings of ECOM'21: The SIGIR 2021 Workshop on eCommerce*. CEUR.

[11] Chandan K. Reddy, Lluís Màrquez, Fran Valero, Nikhil Rao, Hugo Zaragoza, Sambaran Bandyopadhyay, Arnab Biswas, Anlu Xing, and Karthik Subbian. 2022. Shopping Queries Dataset: A Large-Scale ESCI Benchmark for Improving Product Search. (2022). arXiv:2206.06588

[12] Jacopo Tagliabue, Ciro Greco, Jean-Francis Roy, Federico Bianchi, Giovanni Cassani, Bingqing Yu, and Patrick John Chia. 2021. SIGIR 2021 E-Commerce Workshop Data Challenge. In *SIGIR eCom 2021*.

[13] 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.

[14] Hamed Zamani. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. ACM, 717–725.

[15] Hamed Zamani and W. Bruce Croft. 2018. Joint Modeling and Optimization of Search and Recommendation. (2018).