

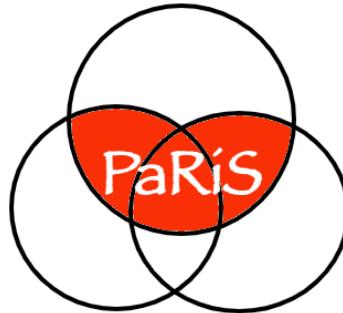
# Third Workshop on Personalization and Recommendations in Search (PaRiS)

Sudarshan Lamkhede  
slamkhede@netflix.com  
Netflix Research  
United States

Moumita Bhattacharya  
mbhattacharya@netflix.com  
Netflix Research  
United States

Hamed Zamani  
zamani@cs.umass.edu  
University of Massachusetts Amherst  
United States

Hongning Wang  
hw5x@virginia.edu  
Dept. of Computer Science, University of Virginia  
United States



## Personalization and Recommendations in Search

### ABSTRACT

We would like to propose a full day workshop on the topic of Personalization and Recommendations in Search (**PaRiS**) to be held in conjunction with SIGIR 2024. This will be the third instance of this workshop. We held two very successful instances of this workshop at the WebConf 2023 and WSDM 2022. This year we plan to especially focus on applications of LLMs and Generative AI to enable personalization and recommendations in the context of search, for example, conversational assistants, while continuing to use this workshop for discussing other advances and applications in the context of personalized search and recommendations in the context of search. Note that all members of the organizing committee plan to attend the conference in person.

### ACM Reference Format:

Sudarshan Lamkhede, Hamed Zamani, Moumita Bhattacharya, and Hongning Wang. 2024. Third Workshop on Personalization and Recommendations in Search (PaRiS). In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

### 1 INTRODUCTION

It has been long argued that personalizing search can be very helpful. In recent years, with proliferation of personal computing devices and large number of logged-in experiences, search has evolved to a stage with many different product scenarios where personalization plays a crucial role for relevance quality and user satisfaction. Though search context plays a big role in determining the relevance of a given result, the utility of a search system for its users can be further enhanced by providing personalized results as well as recommendations within the search context. Moreover, with recent developments in Large Language Models and Generative AI, personalization in the context of search has become even more relevant. For example, searching via conversational systems to find something relevant to shop. A variety of solutions have been developed for search engines in e-commerce systems, streaming/media content providers, social network systems and even in web search systems for such tasks. However, the research discussions around personalization and recommendation for search remain fragmented across different conferences and workshops. We feel that there is a strong need for bringing together researchers and practitioners working on these problems for a robust discussion and sharing of ideas.

---

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 ACM 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).

*Conference'17, July 2017, Washington, DC, USA*  
© 2024 Association for Computing Machinery.  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

This workshop aims for researchers and practitioners from both academia and industry to engage in the discussions of algorithmic and system challenges in search personalization and effectively recommending in search context. It will include but not limited to the topics such as evaluation, query assistance, retrieval, ranking, context modeling, benchmark data and system efficiency for search personalization and recommendations within search contexts, for which more effective and efficient solutions can be shared and discussed. We expect the workshop to be of interest to large audiences in the research community of information retrieval and machine learning.

## 2 THEME AND PURPOSE OF THE WORKSHOP

Personalization has been a prevalent concept in the context of recommender systems for many years now, however, in recent years we have seen a significant increase in the use of personalizations in search. Examples of personalized search include e-commerce sites like Etsy.com, showing different users' different relevant listings for the same query, while web search such as Google, incorporates user intents and context to show different ranked lists of results for different users for the same query. In this workshop, we want to bring all such applications and approaches developed to incorporate personalization and recommendation in search algorithms. This workshop welcomes submissions from academic and industry researchers and practitioners to submit their work related to personalization and recommendations in search. We also welcome work that can highlight the challenges faced in developing a search algorithm with personalization in real production systems. Specifically, the topics of interest are (but are not limited to) listed below:

- Use of Generative AI and LLMs for personalization of search systems
- Personalizations and Recommendations for Conversational Assistants
- Personalized query interpretation and intent disambiguation
- Incorporating user context in real-time or in-session adaptation for personalization in search
- Using the user's search session history for personalization and recommendation
- Personalized evidence generation and explainability of search results
- Fairness and privacy in personalized search
- Latency, caching, and other infrastructural considerations for real-world personalized search systems
- Learnings and challenges for developing large-scale personalized search systems
- Leveraging different modalities of data, including text, image, and audio for the personalization of search
- New learning to rank (LTR) approaches for search personalization as well as beyond traditional LTR approaches, such as reinforcement learning for search personalization
- Joint modeling of search and recommendation
- New applications for personalization and recommendations for search(e.g. in healthcare)
- Evaluation metrics for personalization and recommendation in search

- Optimization of delayed user behavioral rewards for personalization and recommendations in the context of search
- Human-computer interaction considerations in designing personalized Search user interfaces and interaction
- Datasets or design of simulator for personalized retrieval
- Infrastructure for sparse and dense retrieval for a personalized Search and Recommendations experience

Note that Personal Search (e.g. email search, desktop search, etc.) is *not* within the scope of the workshop.

## 3 RATIONALE FOR HOLDING SUCH A WORKSHOP DURING SIGIR

SIGIR is one of the premier conferences that brings together researchers as well as practitioners from academia and industry on web-inspired research involving search and information retrieval.

Search applications, especially those that potentially stand to benefit from incorporating personalization and recommendations, require multidisciplinary research spanning various topics of interest for the SIGIR attendees - machine learning, information retrieval, human-computer interaction, consumer research, distributed systems, data analytics, software engineering, etc. Hence we think that SIGIR would provide a good platform for facilitating a robust discussion around the problem.

While many of the topics could also find a presence at the main conference, we believe that a separate gathering focused on the specific subarea is advantageous because: (1) It would encourage discussion on work in progress that may not be yet ready for the conference (2) Relevant keynotes and panel discussion for subarea maybe difficult to host within the program of the main conference given time and venue constraints while having SIGIR as a single track conference. (3) It would make it easier to find and network with other researchers who are also working on the personalization and recommendations in search.

In our limited survey of our collaborators and ex-colleagues, we heard enthusiastic support for the idea of organizing a workshop on the topic of Personalization and Recommendation in Search. Moreover, we have so far held two very successful instances of this workshop and both times this workshop was one of the most attended one among all the other workshops that we held in the corresponding conference. We think that this enthusiasm will be shared by the attendees of SIGIR, thus widening the appeal of SIGIR 2024.

## 4 PREVIOUS INSTANCES OF THIS WORKSHOP

We have so far held two very successful instances of this workshop. A few details about both are listed below:

### 4.1 WebConf 2023 - Second Workshop

This was the second instance of this workshop ([website](#)), and it was a full-day in-person workshop held in conjunction with WebConf 2023. The program included three keynote speakers, a few invited talks, a few accepted paper presentations and a panel discussion. For the keynote talks, we had Ed Chi from Google Brain, Nick Craswell from Microsoft and Johanne Trippas from the School of Computing Technologies at RMIT University, Australia. Each of the keynotes was very well received, and well attended as well as there was quite

a lot of audience participation. The invited talks included a talk from Netflix on Personalized Recommendations in Netflix Search and a talk from DoorDash on Balancing Relevance and Personalization in Multi-vertical Search at DoorDash. The panel discussion was on the role that large language models and generative AI models will play for personalization and recommendations in the context of search. The workshop was one of the most attended ones in WebConf 2023 with a total of 50 attendees.

## 4.2 WSDM 2022 - First Workshop

The first instance of this workshop was held in conjunction with WSDM 2022 ([website](#)) and was a half-day virtual workshop. Even this year we had three keynote speakers, a few invited talks and a few accepted papers that were presented as a part of the workshop. This year we had the following keynote speakers: Prof. Julian McAuley from UCSD, David Carmel from Amazon and Prof. Hamed Zamani from UMass Amherst. The accepted papers included work from Meta, Walmart Labs, and Ebay. This year it was the second most attended workshop of the eight, with a total of over 47 attendees at WSDM 2022.

## 5 TARGET AUDIENCE

Researchers, practitioners and students interested in Search, Personalization Search and Large Language Models applications in Search. In general, we think it will be interesting to most SIGIR attendees regardless of whether they are from industry or academia.

We plan to advertise the workshop via social network such as LinkedIn, sending CFP notifications via group google mailing list (e.g. SIR-IR List, WiML) as well as directly reaching out to relevant teams and labs in our network.

## 6 TENTATIVE WORKSHOP SCHEDULE

We are proposing to have a full-day workshop, held in the Eastern time zone (local time for Washington DC). We plan to have the following in our workshop program:

- **Full day workshop**
- 4 Keynotes
- 4-6 oral papers presentations and posters (this will be combination of invited talks and accepted papers)
- 1 panel discussion

Please see the table 1 for a tentative full-day schedule.

### 6.1 Panel Discussion

One of the following proposed topics will be discussed by the panel.

- Personalization of a conversational recommendation experience
- Applications of generative AI and LLMs for enabling personalization and recommendations in search.
- Will personalization research be primarily on proprietary datasets?

**Table 1: Tentative Workshop Schedule**

Time	Session	People
9:00 am - 9:15 am	Introduction	Organizers
9:15 am - 9:55 am	Keynote # 1	TBD
9:55 am - 10:10 am	Keynote # 1 Q & A -	
10:15 am - 10:30 am	Break	-
10:30 am - 11:10 am	Keynote # 2	TBD
11:10 am - 11:25 am	Keynote # 2 Q & A	-
11:30 am - 12:10 am	Keynote # 3	TBD
12:10 am - 12:25 pm	Keynote # 3 Q & A	-
12:25 pm - 1:00 pm	Break	-
1:00 pm - 1:40 pm	Keynote # 4	TBD
1:40 pm - 1:55 pm	Keynote # 4 Q & A	-
1:55 pm - 2:00 pm	Break	-
2:00 pm - 4:00 pm	Presentations	Authors
4 pm - 5 pm	Panel Discussion	Panel

## 7 RELATED WORKSHOPS

As far as we know no other workshop or conference was organized on the exact topic of the proposed workshop. But below are some relevant conferences and workshops in recent years.

### 7.1 User Modeling

A few previous workshops listed below have been organized in the past that explicitly model user preferences either by adapting to their behavior in-session or by modeling user's preference from their historical usage data. Although there is some overlap of goals of these workshops with ours, ours is broader in scope with respect to how user preference can be leveraged, while narrower with respect to its application to just search systems. For example, topics such as long-term user modeling is a part of a few of these below-mentioned workshops, whereas in our workshop we are proposing topics that specifically use user's search session history for Personalization and incorporate user context in-session for Personalization in search. Similarly, another workshop aims to provide a platform for discussing the challenges and innovative approaches in fusing multi-dimensional information for user modeling, whereas we aim to make this workshop a platform for discussing challenges and corresponding novel-approaches for Personalization in search. As such, although there are some relevant topics in these previously held workshops, none of them capture the overall mission of our proposed gathering, Personalization and

recommendation in search, which is a very relevant area for many large scale practical search systems.

- Workshop on Online Recommender Systems and User Modeling - Recsys
- Conference on User Modeling, Adaptation and Personalization
- Workshop SUM'20: State-based User Modelling - WSDM 2020
- IFUP: Workshop on Multi-dimensional Information Fusion for User Modeling and Personalization - WSDM 2018
- Learning from User Interactions - WSDM 2018

## 7.2 Fairness, Privacy and Socially Responsible

Dedicated workshops have been organized in the past to discuss very important topics such as fairness, accountability, privacy and social responsibility in information retrieval. Since the workshop we are proposing here aims to delve into the aspects of Personalization and Recommendations within search, considering similar topics such as privacy, user data safety, and fairness are very important and we hope to be able to discuss these during our workshop session. However, unlike the aforementioned previously held workshops, ours aim to discuss these topics all in the context of search systems and what should and should not be taken into account while developing Personalization in search.

- Fairness, Accountability, Confidentiality, Transparency, and Safety - SIGIR 2019
- Workshop on Privacy-Preserving IR - SIGIR 2016
- International Workshop on Social Personalization and Search - ACM Hypertext 2014

## 7.3 Evaluation of Personalization in IR

The WEPiR workshop focuses on one specific, but very relevant topic to our proposed workshop - evaluation of Personalization in Information Retrieval. However, it does not cover topics other than evaluation of Personalization.

- Workshop on Evaluation of Personalization in Information Retrieval - CHIIR 2018-2020

## 8 TENTATIVE KEYNOTE SPEAKERS

As keynote speakers, we plan to invite a few eminent industry leaders and professors who are working on relevant areas. A few tentative keynote speakers are:

- Michael Bendersky - Engineering Director, Google
- Fernando Diaz - Associate Professor, CMU
- Pranam Kolari - VP Engineering, Coupang
- Paul Bennett - Director of Research, Spotify

We also plan to invite speakers from Netflix and Etsy to present their work in this area.

## 9 TENTATIVE PROGRAM COMMITTEE

- Qingyao Ai (Tsinghua University)
- Yongfeng Zhang (Rutgers University)
- Maarten de Rijke (University of Amsterdam)
- Changsang Kang (Walmart)
- Chihoon Lee (Facebook)

- Alex Cozzi (EBay)
- Georges-Eric Dupret (Apple)
- Fabrizio Silvestri (Sapienza Università di Roma)
- Roger Luo (Niantic)
- Liangjie Hong (LinkedIn)
- Edgar Meij (Bloomberg)
- Leo Boytsov (Bosch AI)
- Diane Hu (Etsy)
- Emine Yilmaz (UCL)

## 10 ORGANIZING COMMITTEE

All organizers are planning to attend the workshop in-person.

- **Sudarshan Lamkhede**

*Manager, Machine Learning - Search and Recommendations, Netflix Research.*

Sudarshan leads the applied research team focused on Search and Recommendations algorithms. Prior to Netflix, he led research engineering for various Web Search, Page Optimization and Personalization algorithms at Yahoo! Research. He co-organizes the San Francisco Bay Area Machine Learning symposium ([BayLearn](#)).

- **Hamed Zamani**

*Assistant Professor, University of Massachusetts Amherst.*

Zamani is an Assistant Professor at the University of Massachusetts Amherst, where he also serves as the Associate Director of the Center for Intelligent Information Retrieval (CIIR). Prior to UMass, he was a Researcher at Microsoft working on search and recommendation problems. His research focuses on designing and evaluating statistical and machine learning models with applications to (interactive) information access systems, including search engines, recommender systems, and question answering. His work has led to over 80 refereed publications in the field including a few Best Paper and Honorable Mentions, in addition to a number of open-source research tools. He organized workshops at SIGIR, RecSys, WWW and WSDM conferences.

- **Moumita Bhattacharya**

*Senior Research Scientist - Search and Recommendations, Netflix Research.*

At Netflix Moumita works on developing at-scale machine learning models for Search and Recommendation Systems. Prior to Netflix, she was a Senior Applied Scientist at Etsy, where she was leading a team that developed recommendation systems to show relevant products to Etsy users. Moumita has a PhD in Computer Science with a focus on Machine Learning and its applications. She has been actively serving as Program Committees for WebConf and is a reviewer for conferences such as WebConf, AAAI, and various journals. Moumita is also an adjunct faculty in the Data Science Institute of the University of Delaware.

- **Hongning Wang**

*Associate Professor at University of Virginia, Department of Computer Science.*

He received his PhD degree in computer science at the University of Illinois at Champaign-Urbana in 2014. His research generally lies in the intersection of machine learning and

information retrieval, with a special focus on sequential decision optimization and computational user modeling. His work has generated over 80 research papers in top venues in data mining and information retrieval areas. He is a recipient of 2016 National Science Foundation CAREER Award, 2020 Google Faculty Research Award, and SIGIR'2019 Best Paper Award. He has been serving as Senior Program Committee members for WSDM'2019-21 and CIKM'2020-21.

## ACKNOWLEDGMENTS

This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant number 2143434, and in part by the Office of Naval Research contract number N000142212688. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.