

AgentIR: 2nd Workshop on Agent-based Information Retrieval

Pengyue Jia
City University of Hong Kong
jia.pengyue@my.cityu.edu.hk

Qingpeng Cai
Kuaishou Technology
cqpcurry@gmail.com

Xiangyu Zhao
City University of Hong Kong
xianzhao@cityu.edu.hk

Ling Pan
HKUST
lingpan@ust.hk

Xin Xin
Shandong University
xinxin@sdu.edu.cn

Jin Huang
University of Amsterdam
j.huang2@uva.nl

Weinan Zhang
Shanghai Jiao Tong University
wnzhang@sjtu.edu.cn

Li Zhao
Microsoft Research
lizo@microsoft.com

Dawei Yin
Baidu
yindawei@acm.org

Grace Hui Yang
Georgetown University
huiyang@cs.georgetown.edu

ABSTRACT

Information retrieval (IR) systems are essential in modern society, aiding users to efficiently locate relevant information through query expansion, document retrieval, ranking, and re-ranking. User feedback from ranked outputs forms a dynamic interaction loop with IR systems, which can be modeled as either one-time or sequential decision-making problems. Over the past decade, deep reinforcement learning (DRL) has emerged as a promising approach to decision-making, leveraging the high model capacity of deep learning for complex tasks. While significant research has explored the application of DRL to IR tasks, several fundamental challenges remain underexplored, including the underlying information theory in DRL settings, the limitations of reinforcement learning methods for industrial IR applications, and the simulation of DRL-based IR systems. Concurrently, the advent of large language models (LLMs) has introduced new opportunities for optimizing and simulating IR systems. Building on the success of the Agent-based IR Workshop at SIGIR 2024, we propose hosting the second Agent-based IR Workshop at SIGIR 2025. This workshop will continue to provide a platform for researchers and practitioners from academia and industry to present cutting-edge advances in DRL-based and LLM-based IR systems from an agent-based perspective. By building on the foundation laid in the first workshop, the 2025 edition aims to delve deeper into emerging research challenges, foster collaborations, and explore innovative applications. Through engaging discussions and insightful presentations, the workshop seeks to further expand the boundaries of IR research and solidify its role as a premier venue for advancing agent-based IR systems.

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.
Conference'17, July 2017, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

CCS CONCEPTS

• **Information systems** → **Information retrieval**; • **Computing methodologies** → **Reinforcement learning**.

KEYWORDS

DRL, LLM, Agent-based Information Retrieval

ACM Reference Format:

Pengyue Jia, Qingpeng Cai, Xiangyu Zhao, Ling Pan, Xin Xin, Jin Huang, Weinan Zhang, Li Zhao, Dawei Yin, and Grace Hui Yang. 2025. AgentIR: 2nd Workshop on Agent-based Information Retrieval. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 WORKSHOP THEME AND TOPICS

Information retrieval (IR) systems are indispensable in modern society, facilitating users in finding relevant information through processes such as query expansion, item recall, ranking, and re-ranking [2, 39]. User feedback, provided through actions like browsing and clicking, further refines these systems [4, 39]. Such user-item interactions can extend across multiple search sessions over longer timescales [8, 9, 11, 12, 29, 39]. These interactions, whether in single-step or sequential settings, can be modeled as decision-making processes where the IR system acts as the agent and the user as the environment [36, 39]. Decision-making in this context is inherently counterfactual: user feedback may differ if presented with alternative ranked lists under the same context [1, 39]. Consequently, naive supervised learning approaches, which simply regress to logged user behavior data, fail to achieve optimal performance.

Deep reinforcement learning (DRL) and multi-agent reinforcement learning have emerged as promising approaches for decision-making over the past decade [17]. By integrating the high capacity of deep learning with reinforcement learning paradigms, these methods excel in handling complex scenarios and optimizing long-term rewards [18, 21–23, 26, 30, 39, 46]. Despite their potential, reinforcement learning methods often face limitations in the diversity of solutions they discover. Generative flow networks (GFlowNets) have addressed this challenge by achieving

high rewards while maintaining diversity in discovered solutions [3, 24, 25, 27, 28, 37, 38]. Recently, there has been growing interest in applying DRL to various IR tasks, including query reformulation [15], relevance feedback [19], learning to rank [13, 45], and recommendation [5, 6, 9, 10, 16, 35, 41, 42, 44]. However, research on DRL for IR remains underdeveloped, partly due to the following challenges: (i) A lack of clarity regarding the fundamental information theory underpinning decision-making in IR; (ii) Inconsistent evaluation metrics and baselines across studies; (iii) A reliance on off-the-shelf DRL methods without sufficient adaptation to the unique demands of industrial IR, such as the challenges of applying DRL to real-world systems.

LLMs [20] have attracted widespread attention from academia and industry. LLM-based IR systems [32] employ approaches such as directly recalling and ranking items, pre-training ID embeddings, and orchestrating recommendation modules. ChatGPT, powered by reinforcement learning from human feedback (RLHF), has further inspired research [34] into integrating human preferences into recommender systems. To mitigate hallucination in LLM outputs, retrieval-augmented generation (RAG) combines external knowledge retrieval with generation to produce grounded, factual outputs. Recent studies [14, 31] explore reinforcement learning to enhance RAG's retrieval and generation performance. However, the foundations and applications of LLM-based IR remain unclear and insufficiently validated.

We propose the concept of agent-based information retrieval (AgentIR), aiming to develop intelligent agents systems capable of autonomously performing IR tasks by leveraging techniques from LLMs, RL, and memory. LLMs enable the agent to enhance retrieval efficiency and accuracy through language understanding, while RL allows for continuous optimization of decision-making across diverse retrieval tasks. Advanced memory designs further empower the agent to manage and utilize information more effectively, boosting both efficiency and quality. We believe AgentIR represents a highly promising research direction. This workshop will focus on exploring fundamental principles for formulating LLM-based and DRL-based IR tasks, alongside evaluating practical experimental designs. These discussions aim to foster innovative research, uncover intriguing findings, and inspire new applications.

In previous years, we successfully organized workshops at SIGIR'20 [40], '21 [39], '22 [43], and '24 [7], as well as CIKM'23 [33]. These workshops were among the most popular at their respective conferences, consistently attracting over 200 attendees annually. This year, we aim to delve deeper into the fundamental aspects of topics such as DRL, LLMs, and agents. The workshop will feature keynote speeches, paper presentations, and a panel discussion to foster engaging and insightful exchanges. We anticipate an even larger turnout this year, expecting around 300 participants.

2 PROGRAM SKETCH

2.1 Workshop Format

The workshop is designed to be a full-day hybrid event, featuring two keynote speeches, four invited talks, and six oral research presentations. The keynote speakers will be distinguished professors or researchers with expertise in the field. The invited talks and peer-reviewed oral presentations will focus on two main themes:

(i) academic discussions on foundational reinforcement learning research with potential applications in IR, and (ii) industry-oriented presentations showcasing the design or application of DRL techniques in real-world IR scenarios. Each presentation will be delivered as a lecture accompanied by slides, followed by a dedicated Q&A session to encourage interaction and discussion.

2.2 Tentative Workshop Schedule

The workshop schedule is planned with two half-day sessions.

- Morning Session
 - 08:30 - 08:40: Welcome and Opening
 - 08:40 - 09:10: Keynote 1
 - 09:10 - 09:35: Academic Invited Talk 1
 - 09:35 - 10:00: Industrial Invited Talk 1
 - 10:00 - 10:20: Coffee Break
 - 10:20 - 12:00: Oral Research Talks
- Afternoon Session
 - 13:30 - 14:10: Keynote 2
 - 14:10 - 14:35: Academic Invited Talk 2
 - 14:35 - 15:00: Industrial Invited Talk 2
 - 15:00 - 15:20: Coffee Break
 - 15:20 - 16:20: Oral Research Talks
 - 16:20 - 17:00: Panel Discussion

3 RELATED WORKSHOPS

List of related workshops:

- Agent-based Information Retrieval - (SIGIR2024)
- Deep Reinforcement Learning for Information Retrieval - (SIGIR 2020-2022, CIKM 2023)
- Reinforcement Learning for Knowledge Discovery - (KDD 2019¹)
- Deep Reinforcement Learning Workshop - (NeurIPS 2015-2020²)
- Deep Reinforcement Learning: Frontiers and Challenges - (IJCAI 2016³)
- Deep Reinforcement Learning Meets Structured Prediction - (ICLR 2019⁴)

The DRL Workshop at NeurIPS (2015–2020) and IJCAI (2016) explored integrating neural networks with reinforcement learning in domains such as robotics, strategy games, and multi-agent systems. Similarly, the Deep Reinforcement Learning Meets Structured Prediction workshop at ICLR (2019) focused on applying reinforcement learning paradigms to structured prediction tasks. The workshop at KDD (2019) highlighted a diverse range of real-world reinforcement learning applications. The DRL4IR workshops at SIGIR (2020, 2021, 2022), CIKM (2023) (*the first* series dedicated to DRL for IR), and AgentIR at SIGIR (2024) (*the first* workshop on agent-based IR) were highly successful and well-received by attendees. Building on this foundation, the proposed workshop will emphasize recent advancements in DRL and LLMs for IR. This workshop aims to bring together experts in IR, reinforcement learning, and LLMs. The program will feature keynote speeches, invited talks, paper presentations, a poster session, and a panel discussion, offering

¹<http://www.cse.msu.edu/~zhaoxi35/DRL4KDD/>

²<https://sites.google.com/view/deep-rl-workshop-neurips2020/home>

³<https://sites.google.com/site/deeprijcai16/>

⁴<https://sites.google.com/view/iclr2019-drlstructpred/>

attendees a comprehensive overview of the current state-of-the-art and future research directions. To support hands-on learning, real-world datasets, and code resources will also be available for attendees to explore and practice beyond the workshop.

4 ORGANIZERS INFORMATION AND QUALIFICATION

Mr. Pengyue Jia is a third-year PhD student at City University of Hong Kong (CityU). His research interest lies in Recommender Systems, Information Retrieval, Large Language Models, and GeoAI. He has published several papers in his areas of interest (e.g., KDD, AAAI, NeurIPS, and NAACL). He achieved second place globally in the Amazon KDD CUP 2024: A Multi-task Online Shopping Challenge for Large Language Models. He leads the WWW'25 tutorial "Joint Modeling in Deep Recommender Systems" and the IJCAI'23 tutorial "Multi-task and Multi-scenario Recommendations".

Dr. Qingpeng Cai is a Senior Staff Algorithm Engineer in KuaiShou Technology. He is responsible for business optimization and technical management. During KuaiShou Technology, he lead the application and evolution of reinforcement learning in recommendation and advertising. He obtained a PhD degree in Computer Science from Tsinghua University. His research interests include reinforcement learning, large language models and recommender system. Currently, he has published more than 30 papers in top international conferences. He was the invited speaker of DRL4IR@SIGIR'22 and VideoRecSys@ACM RecSys'23 with the topic "Reinforcement Learning for Short Video Recommender Systems".

Dr. Xiangyu Zhao is an assistant professor at City University of Hong Kong (CityU). His current research interests include data mining and machine learning, especially (1) Personalization, Recommender System, Online Advertising, Search Engine, and Information Retrieval; and (2) Large Language Models and Reinforcement Learning. He has published more than 100 papers in top conferences and journals. His research received ICDM'22 and 21 Best-ranked Papers, Global Top 100 Chinese New Stars in AI, CCF-Tencent Open Fund, Criteo Research Award, and Bytedance Research Award. He serves as the organizers of DRL4KDD@KDD'19, DRL4IR@SIGIR'20, 2nd DRL4KD@WWW'21, 2nd DRL4IR@SIGIR'21, and a lead tutor at WWW'21/22 and IJCAI'21.

Dr. Ling Pan is an assistant professor of the Department of Electronic and Computer Engineering at the Hong Kong University of Science and Technology (HKUST). Her research interests mainly include theoretical understanding, algorithmic improvements and practical application of generative flow networks (GFlowNets), reinforcement learning and multi-agent systems. She was also a recipient of Microsoft Research Ph.D. Fellowship (Asia). She has published several high-quality papers on top-tiered machine learning and artificial intelligence conferences and journals including ICML/NeurIPS/ICLR/AAAI/IJCAI/AMAS, and her works have been selected as spotlight presentations for multiple times. She served as the organizer of the Mila GFlowNet Workshop 2023.

Dr. Xin Xin is an assistant professor at Shandong University (SDU). Before joining SDU, he got the Ph.D. degree in University of Glasgow. His current research interests include informational retrieval and machine learning, especially (1) Personalized Recommendation, and (2) Reinforcement Learning. He has published more

than 20 papers in top conferences (e.g., WWW, SIGIR, ICDM, CIKM, WSDM, IJCAI, ACL) and journals (e.g., TKDE). He was an invited speaker of DRL4IR@SIGIR'21 with the topic of offline reinforcement learning for recommendation.

Dr. Jin Huang is a Postdoctoral Researcher at the qurAI group at the University of Amsterdam (UvA), working on trustworthy AI-based systems for sustainable growth. She did her PhD at the Information Retrieval Lab at UvA. Before joining UvA, she got her BE and ME degree in the Renmin University of China. Her current research interests lie in the development of trustworthy, responsible, and socially-aware AI-power advising-giving systems for critical domains, including healthcare, news recommender systems, and public services. She has published several research papers in top-tier conferences, including SIGIR, WSDM, RecSys, and ECIR.

Dr. Weinan Zhang is currently a tenure-track associate professor in Shanghai Jiao Tong University. His research interests include machine learning and big data mining, particularly, deep learning and reinforcement learning techniques for real-world data mining scenarios, such as computational advertising, recommender systems, text mining, web search and knowledge graphs. He has published over 100 papers on first-tier international conferences and journals, including KDD, SIGIR, ICML, ICLR, JMLR, IJCAI, AAAI, WSDM, CIKM etc. He won the Best Paper Honorable Mention Award in SIGIR 2017, the Best Paper Award in DLP Workshop at KDD 2019, the Best System Paper Award at CoRL 2020, ACM Rising Star Award, Alibaba DAMO Young Scholar Award etc. Weinan has organized workshops and tutorials in SIGIR, KDD, CIKM and ECIR etc.

Dr. Li Zhao is currently a Senior Researcher in Machine Learning Group, Microsoft Research Asia (MSRA). Her research interests mainly lie in deep learning and reinforcement learning, and their applications for text mining, recommendation, finance and games. She obtained her Ph.D. degree majoring in Computer Science in July, 2016, from Tsinghua University, supervised by Professor Xiaoyan Zhu. During her Ph.D. studies, she has conducted research on sentiment extraction, text mining and weakly supervised learning. She published several research papers in top conferences, including NeurIPS, KDD, IJCAI, AAAI, EMNLP and CIKM.

Dr. Dawei Yin is Engineering Director at Baidu inc.. He is managing the search science team at Baidu, leading Baidu's science efforts of web search, question answering, video search, etc.. Previously, he was Senior Director, managing the recommendation engineering team at JD.com between 2016 and 2020. Prior to JD.com, he was Senior Research Manager at Yahoo Labs, leading relevance science team and in charge of Core Search Relevance of Yahoo Search. He obtained Ph.D. (2013), M.S. (2010) from Lehigh University and B.S. (2006) from Shandong University. From 2007 to 2008, he was an M.Phil. student in The University of Hong Kong. His research interests include data mining, applied machine learning, information retrieval and recommender system. He published more than 80 research papers in premium conferences and journals, and was the recipients of WSDM2016 Best Paper Award, KDD2016 Best Paper Award, WSDM2018 Best Student Paper Award, and ICHI 2019 Best Paper Honorable Mention.

Dr. Grace Hui Yang is Associate Professor in Computer Science at Georgetown University, Washington D.C. Dr. Yang's current research interests include deep reinforcement learning, interactive agents, search engines, and privacy-preserving information

retrieval. Prior to this, she has conducted research on question answering, automatic ontology construction, near-duplicate detection, multimedia information retrieval, and opinion and sentiment detection. Dr. Yang led the effort for the Text Retrieval Conference (TREC) Dynamic Domain Tracks from 2015 to 2017 and SIGIR privacy-preserving information retrieval workshops from 2014 to 2016 and co-organized the SIGIR Deep Reinforcement Learning Workshops since 2020. Dr. Yang is associate editor for ACM Transactions on Information Systems and served on the editorial board of Information Retrieval Journal from 2014 to 2017. She has actively served as an organizing or program committee member in many top-tier international conferences such as SIGIR, ECIR, ACL, AAAI, ICTIR, CIKM, WSDM, and WWW, and the general co-chair of SIGIR 2024. She is a recipient of the NSF Faculty Early Career Development Program (CAREER) Award and a co-author of the 2016 book "Dynamic Information Retrieval Modeling."

REFERENCES

- [1] Aman Agarwal, Kenta Takatsu, Ivan Zaitsev, and Thorsten Joachims. 2019. A general framework for counterfactual learning-to-rank. In *Proc. of SIGIR*.
- [2] Ricardo Baeza-Yates, Berthier Ribeiro-Neto, et al. 1999. *Modern information retrieval*. ACM press New York.
- [3] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. 2021. Flow network based generative models for non-iterative diverse candidate generation. *Proc. of NeurIPS* (2021).
- [4] Alexey Borisov, Martijn Wardenaar, Ilya Markov, and Maarten de Rijke. 2018. A click sequence model for web search. In *Proc. of SIGIR*.
- [5] Qingpeng Cai, Shuchang Liu, Xueliang Wang, Tianyou Zuo, Wentao Xie, Bin Yang, Dong Zheng, Peng Jiang, and Kun Gai. 2023. Reinforcing User Retention in a Billion Scale Short Video Recommender System. *arXiv preprint arXiv:2302.01724* (2023).
- [6] Qingpeng Cai, Zhenghai Xue, Chi Zhang, Wanqi Xue, Shuchang Liu, Ruohan Zhan, Xueliang Wang, Tianyou Zuo, Wentao Xie, Dong Zheng, et al. 2023. Two-Stage Constrained Actor-Critic for Short Video Recommendation. In *Proceedings of the ACM Web Conference 2023*.
- [7] Qingpeng Cai, Xiangyu Zhao, Ling Pan, Xin Xin, Jin Huang, Weinan Zhang, Li Zhao, Dawei Yin, and Grace Hui Yang. 2024. AgentIR: 1st Workshop on Agent-based Information Retrieval. In *Proc. of SIGIR*.
- [8] Haokun Chen, Xinyi Dai, Han Cai, Weinan Zhang, Xuejian Wang, Ruiming Tang, Yuzhou Zhang, and Yong Yu. 2019. Large-scale interactive recommendation with tree-structured policy gradient. In *Proc. of AAAI*.
- [9] Jia Chen, Jiabin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2020. A context-aware click model for web search. In *WSDM*.
- [10] Wenqi Fan, Tyler Derr, Xiangyu Zhao, Yao Ma, Hui Liu, Jianping Wang, Jiliang Tang, and Qing Li. 2020. Attacking Black-box Recommendations via Copying Cross-domain User Profiles. *arXiv preprint arXiv:2005.08147* (2020).
- [11] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. 2021. Towards Long-term Fairness in Recommendation. *arXiv preprint arXiv:2101.03584* (2021).
- [12] Li He, Liang Wang, Kaipeng Liu, Bo Wu, and Weinan Zhang. 2018. Optimizing Sponsored Search Ranking Strategy by Deep Reinforcement Learning. *arXiv preprint arXiv:1803.07347* (2018).
- [13] Yujing Hu, Qing Da, Anxiang Zeng, Yang Yu, and Yinghui Xu. 2018. Reinforcement Learning to Rank in E-Commerce Search Engine: Formalization, Analysis, and Application. *arXiv preprint arXiv:1803.00710* (2018).
- [14] Mandar Kulkarni, Praveen Tangarajan, Kyung Kim, and Anusua Trivedi. 2024. Reinforcement Learning for Optimizing RAG for Domain Chatbots. *arXiv preprint arXiv:2401.06800* (2024).
- [15] Johnny Flame Lee. 2020. *Evaluation of a Deep Reinforcement Learning Method For Query Reformulation*. Ph.D. Dissertation.
- [16] Shuchang Liu, Qingpeng Cai, Bowen Sun, Yuhao Wang, Ji Jiang, Dong Zheng, Peng Jiang, Kun Gai, Xiangyu Zhao, and Yongfeng Zhang. 2023. Exploration and Regularization of the Latent Action Space in Recommendation. In *Proceedings of the ACM Web Conference 2023*.
- [17] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [18] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* (2015).
- [19] Ali MontazerAlghaem, Hamed Zamani, and James Allan. 2020. A Reinforcement Learning Framework for Relevance Feedback. In *Proc. of SIGIR*.
- [20] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Proc. of NeurIPS* (2022).
- [21] Ling Pan, Qingpeng Cai, Zhixuan Fang, Pingzhong Tang, and Longbo Huang. 2019. A deep reinforcement learning framework for rebalancing dockless bike sharing systems. In *Proc. of AAAI*.
- [22] Ling Pan, Qingpeng Cai, and Longbo Huang. 2020. Softmax deep double deterministic policy gradients. *Proc. of NeurIPS* (2020).
- [23] Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu. 2022. Plan better amid conservatism: Offline multi-agent reinforcement learning with actor rectification. In *Proc. of ICML*.
- [24] Ling Pan, Moksh Jain, Kanika Madan, and Yoshua Bengio. 2024. Pre-Training and Fine-Tuning Generative Flow Networks. In *Proc. of ICLR*.
- [25] Ling Pan, Nikolay Malkin, Dinghui Zhang, and Yoshua Bengio. 2023. Better training of gflownets with local credit and incomplete trajectories. In *Proc. of ICML*.
- [26] Ling Pan, Tabish Rashid, Bei Peng, Longbo Huang, and Shimon Whiteson. 2021. Regularized softmax deep multi-agent q-learning. *Proc. of NeurIPS* (2021).
- [27] Ling Pan, Dinghui Zhang, Aaron Courville, Longbo Huang, and Yoshua Bengio. 2022. Generative Augmented Flow Networks. In *Proc. of ICLR*.
- [28] Ling Pan, Dinghui Zhang, Moksh Jain, Longbo Huang, and Yoshua Bengio. 2023. Stochastic generative flow networks. In *Uncertainty in Artificial Intelligence*.
- [29] Raziheh Rahimi and Grace Hui Yang. [n. d.]. Modeling Exploration of Intrinsically Diverse Search Tasks as Markov Decision Processes. [n. d.].
- [30] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. 2020. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research* (2020).
- [31] Zheng Wang, Shu Xian Teo, Jieer Ouyang, Yongjun Xu, and Wei Shi. 2024. M-RAG: Reinforcing Large Language Model Performance through Retrieval-Augmented Generation with Multiple Partitions. *arXiv preprint arXiv:2405.16420* (2024).
- [32] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2023. A Survey on Large Language Models for Recommendation. *arXiv preprint arXiv:2305.19860* (2023).
- [33] Xin Xin, Xiangyu Zhao, Jin Huang, Weinan Zhang, Li Zhao, Dawei Yin, and Grace Hui Yang. 2023. DRL4IR: 4th Workshop on Deep Reinforcement Learning for Information Retrieval. In *Proc. of CIKM*.
- [34] Wanqi Xue, Qingpeng Cai, Zhenghai Xue, Shuo Sun, Shuchang Liu, Dong Zheng, Peng Jiang, Kun Gai, and Bo An. 2022. PrefRec: Recommender Systems with Human Preferences for Reinforcing Long-term User Engagement. *arXiv preprint arXiv:2212.02779* (2022).
- [35] Wanqi Xue, Qingpeng Cai, Ruohan Zhan, Dong Zheng, Peng Jiang, and Bo An. 2023. ResAct: Reinforcing Long-term Engagement in Sequential Recommendation with Residual Actor (ICLR'23). (2023).
- [36] Grace Hui Yang, Marc Sloan, and Jun Wang. 2016. Dynamic information retrieval modeling. *Synthesis lectures on information concepts, retrieval, and services* (2016).
- [37] Dinghui Zhang, Hanjun Dai, Nikolay Malkin, Aaron C Courville, Yoshua Bengio, and Ling Pan. 2024. Let the flows tell: Solving graph combinatorial problems with GFlowNets. *Proc. of NeurIPS* (2024).
- [38] Dinghui Zhang, Ling Pan, Ricky TQ Chen, Aaron Courville, and Yoshua Bengio. 2023. Distributional gflownets with quantile flows. *arXiv preprint arXiv:2302.05793* (2023).
- [39] Weinan Zhang, Xiangyu Zhao, Li Zhao, Dawei Yin, and Grace Hui Yang. 2021. DRL4IR: 2nd Workshop on Deep Reinforcement Learning for Information Retrieval.
- [40] Weinan Zhang, Xiangyu Zhao, Li Zhao, Dawei Yin, Grace Hui Yang, and Alex Beutel. 2020. Deep Reinforcement Learning for Information Retrieval: Fundamentals and Advances. In *Proc. of SIGIR*.
- [41] Kesen Zhao, Shuchang Liu, Qingpeng Cai, Xiangyu Zhao, Ziru Liu, Dong Zheng, Peng Jiang, and Kun Gai. 2023. KuaiSim: A comprehensive simulator for recommender systems. *arXiv preprint arXiv:2309.12645* (2023).
- [42] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep Reinforcement Learning for Page-wise Recommendations. In *Proceedings of the 12th ACM Recommender Systems Conference*.
- [43] Xiangyu Zhao, Xin Xin, Weinan Zhang, Li Zhao, Dawei Yin, and Grace Hui Yang. 2022. DRL4IR: 3rd Workshop on Deep Reinforcement Learning for Information Retrieval. In *Proc. of SIGIR*.
- [44] Xiangyu Zhao, Liang Zhang, Zhuoye Ding, Long Xia, Jiliang Tang, and Dawei Yin. 2018. Recommendations with Negative Feedback via Pairwise Deep Reinforcement Learning. In *Proc. of KDD*.
- [45] Xiangyu Zhao, Xudong Zheng, Xiwang Yang, Xiaobing Liu, and Jiliang Tang. 2020. Jointly learning to recommend and advertise. In *Proc. of KDD*.
- [46] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A Deep Reinforcement Learning Framework for News Recommendation. In *Proc. of WWW*.