

Report on the 1st Workshop on Large Language Model for Evaluation in Information Retrieval (LLM4Eval 2024) at SIGIR 2024

Hossein A. Rahmani

University College London
London, UK

hossein.rahmani.22@ucl.ac.uk

Clemencia Siro

University of Amsterdam
Amsterdam, The Netherlands

c.n.siro@uva.nl

Mohammad Aliannejadi

University of Amsterdam
Amsterdam, The Netherlands

m.aliannejadi@uva.nl

Nick Craswell

Microsoft
Seattle, US

nickcr@microsoft.com

Charles L. A. Clarke

University of Waterloo
Ontario, Canada

claclark@gmail.com

Guglielmo Faggioli

University of Padua
Padua, Italy

guglielmo.faggioli@unipd.it

Bhaskar Mitra

Microsoft
Montréal, Canada

bmitra@microsoft.com

Paul Thomas

Microsoft
Adelaide, Australia

pathom@microsoft.com

Emine Yilmaz

University College London
London, UK

emine.yilmaz@ucl.ac.uk

Abstract

The first edition of the workshop on Large Language Model for Evaluation in Information Retrieval (LLM4Eval 2024) took place in July 2024, co-located with the ACM SIGIR Conference 2024 in the USA (SIGIR 2024). The aim was to bring information retrieval researchers together around the topic of LLMs for evaluation in information retrieval that gathered attention with the advancement of large language models and generative AI. Given the novelty of the topic, the workshop was focused around multi-sided discussions, namely panels and poster sessions of the accepted proceedings papers.

Date: 18 July 2024.

Website: <https://llm4eval.github.io>.

1 Introduction

Large language models (LLMs), such as ChatGPT,¹ have demonstrated increasing effectiveness, with larger models performing well on tasks where smaller models are insufficient. Recently, LLMs have been actively explored for various evaluation tasks, among others.

¹<https://chatgpt.com>

In information retrieval (IR), among other applications, LLMs are actively explored for estimating query-document relevance, both for ranking and for label generation [Rahmani et al., 2024a; Craswell et al., 2024]. The latter can then be used for training and evaluating other less powerful but more efficient rankers. LLMs are employed for relevance labeling in industry [Thomas et al., 2023]. The evaluation methodologies apply a wider range of LLMs and prompts to the labeling problem, potentially addressing a broader range of quality issues.

Motivated by these observations, we believed that a workshop on evaluation strategies in the context of LLMs would question whether IR and NLP were truly facing a paradigm shift in evaluation strategies. Therefore, we organized this workshop to provide a fresh perspective on LLM-based evaluation through an information retrieval lens. The workshop also provided an opportunity to reflect on the benefits and challenges of LLM-based evaluation in academia and industry. Finally, we encouraged submissions and discussions on further evaluation topics and models, where existing literature is scarce, such as recommender systems, learning to rank, and diffusion models.

This paper is an event report of our own LLM4Eval [Rahmani et al., 2024b] event, the first workshop on Large Language Model for Evaluation in Information Retrieval (LLM4Eval 2024), held in conjunction with SIGIR 2024. The workshop had a poster session with accepted papers and a panel discussion. We report on how we organized the workshop (Section 2), provide a descriptive account of what happened at the workshop (Section 3), and report on what we learned from the LLMJudge challenge (Section 4).

2 Workshop Overview

This section provides a descriptive account of the paper review process and how we organized the workshop and panel session. We begin by defining what topics the workshop was mainly focused on among many others.

2.1 Topics

The workshop focused on models, techniques, data collections, and methodologies for information retrieval evaluation in the era of LLMs. These include but are not limited to:

- LLM-based evaluation metrics for traditional IR and generative IR
- Agreement between human and LLM labels
- Effectiveness and/or efficiency of LLMs to produce robust relevance labels
- Investigating LLM-based relevance estimators for potential systemic biases
- Automated evaluation of text generation systems
- End-to-end evaluation of Retrieval Augmented Generation systems
- Trustworthiness in the world of LLMs evaluation
- Prompt engineering in LLMs evaluation
- Effectiveness and/or efficiency of LLMs as ranking models

2.2 Format

The workshop was a full-day in-person workshop held in Washington D.C., US on the 18th of July 2024. The day was organized as follows:

Time	Agenda
Morning	
9:00 – 9:15	Opening Remarks
9:15 – 10:00	Keynote 1: Ian Soboroff, NIST
10:00 – 10:30	Booster Talks 1
10:30 – 11:00	Coffee Break
11:00 – 11:30	Booster Talks 2
11:30 – 12:30	Poster Session
12:30 – 13:30	Lunch
Afternoon	
13:30 – 14:15	Keynote 2: Donald Metzler, Google
14:15 – 14:30	LLMJudge Presentation
14:30 – 15:00	Discussion on the results of LLMJudge
15:00 – 15:30	Coffee Break
15:30 – 16:55	Panel Discussion

2.3 Program Committee

LLM4Eval exists thanks to the dedication of 24 researchers who volunteered their time to review the submissions. We are deeply grateful to each member for their commitment to the workshop. Below is a list of the program committee members:

- Zahra Abbasiantaeb, University of Amsterdam
- Mofetoluwa Adeyemi, University of Waterloo
- Marwah Alaofi, RMIT University
- Negar Arabzadeh, University of Waterloo
- Shivangi Bithel, IIT Delhi
- Francesco Luigi De Faveri, University of Padua
- Yashar Deldjoo, Polytechnic University of Bari
- Gianluca Demartini, The University of Queensland
- Laura Dietz, University of New Hampshire
- Yue Feng, UCL
- Claudia Hauff, Spotify
- Bhawesh Kumar, Verily Life Sciences
- Yiqun Liu, Tsinghua University
- Sean MacAvaney, University of Glasgow
- James Mayfield, Johns Hopkins University

-
- Chuan Meng, University of Amsterdam
 - Ipsita Mohanty, Carnegie Mellon University
 - Mohammadmehdi Naghiaei, University of Southern California
 - Pranoy Panda, Fujitsu Research
 - Lu Wang, Microsoft
 - Xi Wang, University of Sheffield
 - Orion Weller, Johns Hopkins University
 - Jheng-Hong Yang, University of Waterloo
 - Oleg Zendel, RMIT University

3 Workshop Program

In this section, we present an overview of the LLM4Eval workshop, encompassing details about its participants, accepted papers, poster sessions, panel discussion.

3.1 In Numbers

As the first workshop on LLMs for evaluation in information retrieval, LLM4Eval 2024 has attracted significant interest from the community. The workshop welcomed more than 50 in-person participants, reflecting the growing curiosity and engagement around the evolving role of LLMs in information retrieval evaluation.

3.2 Keynotes

LLM4Eval featured two invited keynote talks. We present the title and abstract of each talk below along with the name of each speaker.

3.2.1 Keynote 1: A Brief History of Automatic Evaluation in IR *by Ian Soboroff, NIST*

Abstract. The ability of large language models such as GPT4 to respond to natural language instructions with flowing, grammatical text that reflects world knowledge has generated (sorry) significant interest in IR, as it has everywhere, and specifically in the area of IR evaluation. It seems that just as we “prompt” a human assessor to provide a relevance judgment, we can do the same thing with an LLM. Researchers are very excited because the fluent, concise, informed, and perhaps even grounded responses from the LLM feel like interacting with a person, and so we guess they might have some of the same capabilities beyond producing fluent textual responses to prompts. In IR we are always complaining about the costs of human assessments, so perhaps this is solved. I would like to point out, although it is not the main thrust of this talk, that if the above is true, IR is solved and we don’t need to have research about it any more. The computer understands the document and the user information need to the degree that it can accurately predict if the document meets the need, and that is what IR systems are supposed to do. Scaling current LLM capabilities to where it can run on your wristwatch is just engineering. The actual

thrust of this talk will be to review some of the history and literature on automatic evaluation methods. This is not automatic evaluation’s first rodeo, as they say. My arrival at NIST was accompanied by a SIGIR paper proposing that relevant documents could be picked using random sampling, and from that point the race was on. Along the way we have reinforced some things we already knew, like relevance feedback is good, and found some new things we did not know.

3.2.2 Keynote 2: LLMs as Rankers, Raters, and Rewarders *by* Donald Metzler, Google DeepMind

Abstract. In this talk, I will discuss recent advancements in the application of large language models (LLMs) to ranking, rating, and reward modeling, particularly in the context of information retrieval tasks. I will emphasize the fundamental similarities among these problems, highlighting that they essentially address the same underlying issue but through different approaches. Based on this observation, I propose several research questions that offer promising avenues for future exploration.

3.3 Papers

The workshop received 21 paper submissions, each of which was reviewed through EasyChair² in a double-blind process by at least three reviewers from the list in Section 2.3. Reviewers rated papers as reject, weak reject, weak accept, or accept, with no option for a neutral (borderline) stance. Papers with mixed reviews were evaluated further by the organizers, who acted as meta-reviewers. We encouraged authors to include code and reproducibility efforts in their submissions.

All accepted papers are hosted on our website³ in a non-archival format and were presented in a poster session. 7 of these papers were selected for presentation and publication in the CEUR-WS volume “*Proceedings of the 1st Workshop on Large Language Models for Evaluation in Information Retrieval (LLM4Eval)*”.⁴ Other 11 papers were accepted for presentation only. Additionally, the workshop received 5 already published works which, being on topic, were accepted for presentation to inspire group discussions. Although the submissions varied in perspectives, they all focused on evaluation topics. Below, we present the titles and authors of the accepted papers.

3.3.1 Accepted Papers

1. One-Shot Labeling for Automatic Relevance Estimation
[MacAvaney and Soldaini \[2023\]](#)
2. Evaluating Cross-modal Generative Models Using Retrieval Task
[Bithel and Bedathur \[2023\]](#)
3. A Comparison of Methods for Evaluating Generative IR
[Arabzadeh and Clarke \[2024\]](#)
4. A Novel Evaluation Framework for Image2Text Generation
[Huang, Zhu, Shen, Rudinac, Paccas, and Kanoulas \[2024\]](#)

²<https://easychair.org/>

³<https://llm4eval.github.io/papers/>

⁴<https://ceur-ws.org/Vol-3752/>

-
5. Using LLMs to Investigate Correlations of Conversational Follow-up Queries with User Satisfaction
[Kim, Choi, Yang, Lee, Park, Lee, Kim, and Kim \[2024\]](#)
 6. EXAM++: LLM-based Answerability Metrics for IR Evaluation
[Farzi and Dietz \[2024\]](#)
 7. Context Does Matter: Implications for Crowdsourced Evaluation Labels in Task-Oriented Dialogue Systems
[Siro, Aliannejadi, and de Rijke \[2024\]](#)
 8. On the Evaluation of Machine-Generated Reports
[Mayfield, Yang, Lawrie, MacAvaney, McNamee, Oard, Soldaini, Soboroff, Weller, Kayi, et al. \[2024\]](#)
 9. Toward Automatic Relevance Judgment using Vision–Language Models for Image–Text Retrieval Evaluation
[Yang and Lin \[2024\]](#)
 10. Reliable Confidence Intervals for Information Retrieval Evaluation Using Generative A.I.
[Oosterhuis, Jagerman, Qin, Wang, and Bendersky \[2024\]](#)
 11. Selective Fine-tuning on LLM-labeled Data May Reduce Reliance on Human Annotation: A Case Study Using Schedule-of-Event Table Detection
[Kumar, Amar, Yang, Li, and Jia \[2024\]](#)
 12. FollowIR: Evaluating and Teaching Information Retrieval Models to Follow Instructions
[Weller, Chang, MacAvaney, Lo, Cohan, Van Durme, Lawrie, and Soldaini \[2024\]](#)
 13. The Challenges of Evaluating LLM Applications: An Analysis of Automated, Human, and LLM-Based Approaches
[Abeyasinghe and Circi \[2024\]](#)
 14. Can We Use Large Language Models to Fill Relevance Judgment Holes?
[Abbasiantaeb, Meng, Azzopardi, and Aliannejadi \[2024\]](#)
 15. Query Performance Prediction using Relevance Judgments Generated by Large Language Model
[Meng, Arabzadeh, Askari, Aliannejadi, and de Rijke \[2024\]](#)
 16. Large Language Models for Relevance Judgment in Product Search
[Mehrdad, Mohapatra, Bagdouri, Chandran, Magnani, Cai, Puthenputhuserry, Yadav, Lee, Zhai, et al. \[2024\]](#)
 17. Evaluating the Retrieval Component in LLM-Based Question Answering Systems
[Alinejad, Kumar, and Vahdat \[2024\]](#)
 18. Exploring Large Language Models for Relevance Judgments in Tetun
[de Jesus and Nunes \[2024\]](#)
 19. A Comparative Analysis of Faithfulness Metrics and Humans in Citation Evaluation
[Zhang, Aliannejadi, Yuan, Pei, Huang, and Kanoulas \[2024\]](#)
 20. Evaluating RAG-Fusion with RAGElo: an Automated Elo-based Framework
[Rackauckas, Câmara, and Zavrel \[2024\]](#)
 21. Enhancing Demographic Diversity in Test Collections Using LLMs
Alaofi, Ferro, Thomas, Scholer and Sanderson
 22. GPT-4 Relevance Labelling can be Fooled by Query Keyword Stuffing
[Alaofi, Thomas, Scholer, and Sanderson \[2024\]](#)

3.4 Poster Session

In light of the acceptance of 23 papers, we organized a dynamic poster session to facilitate the dissemination of their findings. All the presentations were held in person. Additionally, we integrated our poster session with two other SIGIR workshops, namely IR-RAG⁵ and ReNeuIR,⁶ fostering a collaborative environment for sharing insights.

3.5 Panel Discussion

The workshop included a panel discussion on relevant topics raised by the audience concerning the use of the LLMs for evaluation. The invited panellists were Charlie Clarke (University of Waterloo), Laura Dietz (University of New Hampshire), Michael D. Ekstrand (Drexel University), and Ian Soboroff (National Institute of Standards and Technology (NIST)). The moderator was Bhaskar Mitra (Microsoft Research).

3.5.1 Evaluation Validity

A large part of the discussion focused on the validity of the evaluation using LLMs. One thing that we should address if we envision the use of LLMs as assessors is the circularity of the evaluation. While it is true that, based on the TREC paradigm, some specific IR models are used to construct the pool of documents to be annotated, it was also demonstrated that the TREC-style evaluation does not introduce any form of bias towards such models. On the contrary, if we were to use an LLM both as an assessor and as a ranker, we could expect such a model to be favoured over other evaluated models. This might also impair the development of new LLMs if we were to evaluate them on judgements constructed using a simpler LLM. If we assume a similar evaluation protocol, an LLM would be considered perfect if it behaves exactly as the LLM used for the annotations, which might be worse and therefore suboptimal.

3.5.2 Intrinsic Randomness of the LLMs

A second element discussed during the panel concerned the intrinsic randomness of these models. Indeed, some operations that are becoming more and more common when operating with an LLM, such as prompt engineering or parameter tuning, induce randomness in the generation: it is impossible to know beforehand what the output will be, given a certain prompt. To address this limitation, one of the proposed solutions involved the development of repositories of baseline prompts for a series of tasks that should be as minimal as possible. In this regard, there was a consensus on the fact that some “tricks” that are known to work in practice should be avoided to build a solid evaluation strategy. Examples of such tricks involve using special characters, “threatening” or “flattering” sentences and other word sequences that might work in practice in specific use cases but for which we are not able to devise a mathematical model describing why and how they work. To address the randomness intrinsic to the LLMs, a point raised by the audience concerned the possibility of exploiting it to build distributions of answers. For example, it is possible to envision a scenario in which, instead of interrogating a single LLM with a single

⁵<https://coda.io/@rstless-group/ir-rag-sigir24>

⁶<https://reneuיר.org/>

prompt once, we could interrogate multiple models multiple times using multiple prompts, to construct a distribution of probability over the answers/relevance judgements that can be used to summarize the LLMs’ opinion on the topic. A downside of this approach, as highlighted by the panellists, is the consumption of the LLMs. Indeed, LLMs are not only expensive from an economic perspective, but they have also an environmental impact. This side should be taken into consideration when using these models, especially if we consider resource-intensive procedures as the sampling.

3.5.3 Replicability and Reproducibility

Another important issue raised during the discussion concerned the replicability of the experiments that involve LLMs as assessors. The community should agree on policies and guidelines concerning proprietary models that cannot be reimplemented and replicated autonomously by the research community. We should foresee, address, and prevent possible scenarios in which changes in a proprietary model impact the scientific conclusions and findings of the papers that rely on such a model for the evaluation and empirical validation of the hypotheses. In this regard, the Diversity, Equity and Inclusion aspects should also be taken into consideration: often, proprietary models are subject to costs that might not be sustainable for research groups with fewer economic resources. In this sense, our future evaluation protocols based on LLMs should be applicable regardless of the resources available to the different research groups. A counter-argument that was raised in this regard, concerns the Cranfield paradigm and TREC-style evaluation. Akin to LLM-based evaluation, during its initial stages, a part of the research community considered TREC-style evaluation to be expensive, hard to replicate, and not deterministic, due to the partial annotation of the topics. With the development and refinement of the protocols, as well as the increasing familiarity of the researchers with this type of evaluation, TREC-style evaluation has become the de facto standard procedure. It is possible to envision a similar path also for LLMs-based evaluation.

3.5.4 The Parallelism Between Human and LLMs Assessment

Finally, an open issue concerns the parallelism between human and LLM annotations. One observation that was made is that we are somehow used to feeding “prompts” to “black-box operators”. Indeed, this is for example what happens with the assessor guidelines commonly used by both the research and industry communities. For humans, it is common to “experience” the act of searching: the annotation process in this sense can be described as a form of generalization of the act of finding relevant information, which is indeed something the assessors have experienced in their lives. This is certainly not the case for the current LLMs who do not have empirical experience in the real world and therefore are not capable of generalizing something they cannot have experienced.

4 LLMJudge Challenge

The goal of the challenge was to attract the attention of the community towards using LMs for evaluation and to release datasets that could later be used to enhance research in this area.




Figure 1. Scatter plot of Cohen’s κ and Kendall’s τ for submitted labelers

The LLMJudge challenge reused the MS MARCO datasets [Nguyen et al., 2016] as the primary benchmark. The test queries were a mix of previous years’ TREC 2023 Deep Learning Track (TREC DL ’23) test sets, which were released along with a development set for fine-tuning or in-context learning purposes. Participants were given a set of $\langle \text{query}, \text{document} \rangle$ pairs and were asked to generate a relevance label.

Participants needed to submit their exact prompt together with the predicted labels for the documents. When submitting prompts, participants were also able to indicate the exact LLM model and parameters they employed to generate the run, which could be used to reproduce it. By allowing participants to submit their prompts, we could further analyze how these prompts might work across a variety of different LLM models.

In order to evaluate the quality of the generated labels, we used Cohen’s κ to see the labeler’s agreement with LLMJudge test data at query-document level and the Kendall’s τ to check the labeler’s agreement with LLMJudge test data on system ordering, i.e., the runs that submitted to TREC DL 2023. In total, we had 39 submissions (i.e., the 39 labelers) from 7 groups from National Institute of Standards and Technology (NIST), RMIT University, The University of Melbourne, University of New Hampshire, University of Waterloo, Included Health, and University of Amsterdam.

Figure 1 shows the performance of submitted labelers on LLMJudge test set. The x-axis indicates Cohen’s κ while the y-axis shows the labeler’s agreement on system ordering. It can be seen that labelers have a low variability for Kendall’s τ but a larger for Cohen’s κ . Most of the labelers are clustered in a narrow range of τ values, indicating that while they agree well on the ordering of systems, there is more variation in their inter-rater reliability as measured by Cohen’s κ . This suggests that while the labelers tend to rank systems similarly, there is less consistency in their exact labeling, leading to variability in Cohen’s κ values.

5 Conclusion

The LLM4Eval 2024 workshop was designed as a platform to foster collaboration between academia and industry researchers from diverse backgrounds, united by a shared interest in the concept, development, and application of large language models for evaluation in information retrieval. This commitment to inclusivity is reflected in our workshop program, which features a panel comprising 4 researchers, a poster session showcasing 22 accepted papers, and a roundtable discussion. The immense potential of large language models in information retrieval and their subsequent applications in downstream services is widely recognized.

Acknowledgments

We would like to thank ACM and SIGIR for hosting this workshop and extend our appreciation to the SIGIR 2024 workshop chairs, Zhuyun Dai, Hussein Suleman, and Andrew Trotman as well as our exceptional panellists, program committee members, paper authors, and participants. This research is supported by the Engineering and Physical Sciences Research Council [EP/S021566/1], the EPSRC Fellowship titled “Task Based Information Retrieval” [EP/P024289/1].

References

- Zahra Abbasiantaeb, Chuan Meng, Leif Azzopardi, and Mohammad Aliannejadi. Can we use large language models to fill relevance judgment holes? *arXiv preprint arXiv:2405.05600*, 2024. URL <https://arxiv.org/abs/2405.05600>.
- Bhashithe Abeysinghe and Ruhan Circi. The challenges of evaluating llm applications: An analysis of automated, human, and llm-based approaches. *arXiv preprint arXiv:2406.03339*, 2024. URL <https://arxiv.org/abs/2406.03339>.
- Marwah Alaofi, Paul Thomas, Falk Scholer, and Mark Sanderson. Llms can be fooled into labelling a document as relevant (best café near me; this paper is perfectly relevant). In *International ACM SIGIR Conference on Information Retrieval in the Asia Pacific*, December 2024. URL <https://dl.acm.org/doi/abs/10.1145/3673791.3698431>.
- Ashkan Alinejad, Krtin Kumar, and Ali Vahdat. Evaluating the retrieval component in llm-based question answering systems. *arXiv preprint arXiv:2406.06458*, 2024. URL <https://arxiv.org/abs/2406.06458>.
- Negar Arabzadeh and Charles LA Clarke. A comparison of methods for evaluating generative ir. *arXiv preprint arXiv:2404.04044*, 2024. URL <https://arxiv.org/abs/2404.04044>.
- Shivangi Bithel and Srikanta Bedathur. Evaluating cross-modal generative models using retrieval task. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1960–1965, 2023. URL <https://dl.acm.org/doi/abs/10.1145/3539618.3591979>.

-
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Hossein A. Rahmani, Daniel Campos, Jimmy Lin, Ellen M. Voorhees, and Ian Soboroff. Overview of the trec 2023 deep learning track. In *Text REtrieval Conference (TREC)*. NIST, TREC, February 2024. URL <https://www.microsoft.com/en-us/research/publication/overview-of-the-trec-2023-deep-learning-track/>.
- Gabriel de Jesus and Sérgio Nunes. Exploring large language models for relevance judgments in tetun. *arXiv preprint arXiv:2406.07299*, 2024. URL <https://arxiv.org/abs/2406.07299v1>.
- Naghmeh Farzi and Laura Dietz. Exam++: Llm-based answerability metrics for ir evaluation. In *Proceedings of LLM4Eval: The First Workshop on Large Language Models for Evaluation in Information Retrieval*, 2024. URL <https://ceur-ws.org/Vol-3752/paper3.pdf>.
- Jia-Hong Huang, Hongyi Zhu, Yixian Shen, Stevan Rudinac, Alessio M. Paccas, and Evangelos Kanoulas. A novel evaluation framework for image2text generation, 2024. URL <https://arxiv.org/abs/2408.01723>.
- Hyunwoo Kim, Yoonseo Choi, Taehyun Yang, Honggu Lee, Chaneon Park, Yongju Lee, Jin Young Kim, and Juho Kim. Using llms to investigate correlations of conversational follow-up queries with user satisfaction. *arXiv preprint arXiv:2407.13166*, 2024. URL <https://ceur-ws.org/Vol-3752/paper5.pdf>.
- Bhawesh Kumar, Jonathan Amar, Eric Yang, Nan Li, and Yugang Jia. Selective fine-tuning on llm-labeled data may reduce reliance on human annotation: A case study using schedule-of-event table detection. *arXiv preprint arXiv:2405.06093*, 2024. URL <https://www.arxiv.org/abs/2405.06093>.
- Sean MacAvaney and Luca Soldaini. One-shot labeling for automatic relevance estimation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2230–2235, 2023. URL <https://arxiv.org/abs/2302.11266>.
- James Mayfield, Eugene Yang, Dawn Lawrie, Sean MacAvaney, Paul McNamee, Douglas W Oard, Luca Soldaini, Ian Soboroff, Orion Weller, Efsun Kayi, et al. On the evaluation of machine-generated reports. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1904–1915, 2024. URL <https://arxiv.org/abs/2405.00982>.
- Navid Mehrdad, Hrushikesh Mohapatra, Mossaab Bagdouri, Prijith Chandran, Alessandro Magnani, Xunfan Cai, Ajit Puthenpuhussery, Sachin Yadav, Tony Lee, ChengXiang Zhai, et al. Large language models for relevance judgment in product search. *arXiv preprint arXiv:2406.00247*, 2024. URL <https://arxiv.org/abs/2406.00247>.
- Chuan Meng, Negar Arabzadeh, Arian Askari, Mohammad Aliannejadi, and Maarten de Rijke. Query performance prediction using relevance judgments generated by large language models. *arXiv preprint arXiv:2404.01012*, 2024. URL <https://arxiv.org/abs/2404.01012>.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. *choice*, 2640: 660, 2016.

-
- Harrie Oosterhuis, Rolf Jagerman, Zhen Qin, Xuanhui Wang, and Michael Bendersky. Reliable confidence intervals for information retrieval evaluation using generative ai. *arXiv preprint arXiv:2407.02464*, 2024. URL <https://arxiv.org/abs/2407.02464>.
- Zackary Rackauckas, Arthur Câmara, and Jakub Zavrel. Evaluating rag-fusion with ragelo: an automated elo-based framework. *arXiv preprint arXiv:2406.14783*, 2024. URL <https://arxiv.org/abs/2406.14783>.
- Hossein A Rahmani, Nick Craswell, Emine Yilmaz, Bhaskar Mitra, and Daniel Campos. Synthetic test collections for retrieval evaluation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2647–2651, 2024a. URL <https://dl.acm.org/doi/10.1145/3626772.3657942>.
- Hossein A Rahmani, Clemencia Siro, Mohammad Aliannejadi, Nick Craswell, Charles LA Clarke, Guglielmo Faggioli, Bhaskar Mitra, Paul Thomas, and Emine Yilmaz. Llm4eval: Large language model for evaluation in ir. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3040–3043, 2024b. URL <https://dl.acm.org/doi/10.1145/3626772.3657992>.
- Clemencia Siro, Mohammad Aliannejadi, and Maarten de Rijke. Context does matter: Implications for crowdsourced evaluation labels in task-oriented dialogue systems. *arXiv preprint arXiv:2404.09980*, 2024. URL <https://arxiv.org/abs/2404.09980>.
- Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. Large language models can accurately predict searcher preferences. *arXiv preprint arXiv:2309.10621*, 2023. URL <https://dl.acm.org/doi/abs/10.1145/3626772.3657707>.
- Orion Weller, Benjamin Chang, Sean MacAvaney, Kyle Lo, Arman Cohan, Benjamin Van Durme, Dawn Lawrie, and Luca Soldaini. Followir: Evaluating and teaching information retrieval models to follow instructions. *arXiv preprint arXiv:2403.15246*, 2024. URL <https://arxiv.org/abs/2403.15246>.
- Jheng-Hong Yang and Jimmy Lin. Toward automatic relevance judgment using vision–language models for image–text retrieval evaluation, 2024. URL <https://arxiv.org/abs/2408.01363>.
- Weijia Zhang, Mohammad Aliannejadi, Yifei Yuan, Jiahuan Pei, Jia-Hong Huang, and Evangelos Kanoulas. Towards fine-grained citation evaluation in generated text: A comparative analysis of faithfulness metrics, 2024. URL <https://arxiv.org/abs/2406.15264>.