

Limitations of Automatic Relevance Assessments with Large Language Models for Fair and Reliable Retrieval Evaluation

David Otero
david.otero.freijeiro@udc.es
IRlab, CITIC
Universidade da Coruña
A Coruña, Spain

Javier Parapar
javier.parapar@udc.es
IRlab, CITIC
Universidade da Coruña
A Coruña, Spain

Álvaro Barreiro
barreiro@udc.es
IRlab, CITIC
Universidade da Coruña
A Coruña, Spain

Abstract

Offline evaluation of search systems depends on test collections. These benchmarks provide the researchers with a corpus of documents, topics and relevance judgements indicating which documents are relevant for each topic. While test collections are an integral part of Information Retrieval (IR) research, their creation involves significant efforts in manual annotation. Large language models (LLMs) are gaining much attention as tools for automatic relevance assessment. Recent research has shown that LLM-based assessments yield high systems ranking correlation with human-made judgements. These correlations are helpful in large-scale experiments but less informative if we want to focus on top-performing systems. Moreover, these correlations ignore whether and how LLM-based judgements impact the statistically significant differences among systems with respect to human assessments. In this work, we look at how LLM-generated judgements preserve ranking differences among top-performing systems and also how they preserve pairwise significance evaluation as human judgements. Our results show that LLM-based judgements are unfair at ranking top-performing systems. Moreover, we observe an exceedingly high rate of false positives regarding statistical differences.

CCS Concepts

• **Information systems** → **Test collections; Relevance assessment; Evaluation of retrieval results; Information retrieval.**

Keywords

Information Retrieval; Test Collections; LLMs; Relevance Assessments

ACM Reference Format:

David Otero, Javier Parapar, and Álvaro Barreiro. 2025. Limitations of Automatic Relevance Assessments with Large Language Models for Fair and Reliable Retrieval Evaluation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 Introduction

Test collections allow for easy evaluation and comparison of Information Retrieval (IR) systems. A test collection consists of a set

of relevance judgements (also called labels, assessments or qrels throughout this paper), which are annotations establishing the documents systems should retrieve for each user's need (also known as *topics*). Collections built in venues like the Text REtrieval Conference (TREC) have enabled tremendous advancements in IR for many years. However, making the relevance judgements for a TREC-like test collection can be very expensive and labour-intensive. Low-budget collections could lead to biased and unstable evaluations.

Recently, several works explored the possibility of using large language models (LLMs) to assess the relevance of documents [1, 4, 7, 12, 15, 16, 17, 18], since they are considerable cheaper than humans and show strong text-understanding capabilities. All these works use ranking correlations, such as Kendall's τ [5], to validate the reliability of judgements generated with an LLM for IR evaluation. The rationale is that if judgements generated by an LLM can provide a similar ranking of systems to the ranking inferred from human judgements, the LLM could replace the human in terms of capturing run-level effectiveness [17]. Thus, the size of future collections in terms of judgements could increase drastically since the budget would no longer be an issue. The problem is that Kendall's τ is not very informative for evaluating agreement on the top-performing systems. Usually, the goal of developing a new system is surpassing the top of the leaderboard. Therefore, it is vital to distinguish meaningful differences between the best systems. Furthermore, some recent works have demonstrated that rank correlation is not enough [6, 10]. Although judgements generated by an LLM may give good results in terms of system rankings, we know nothing about whether the results of a statistical test between any two systems would yield similar results. In addition, recent studies have highlighted that LLM-based judgements may suffer from other issues [2].

In this paper, we aim to contribute to fill this gap. First, we have a closer look at how LLM-generated qrels are able to preserve the same ranking of systems as inferred from human judgements regarding top-performing systems. In particular, we use top-heavy rank correlations that penalize more misses at the top. Second, we evaluate how fair LLM-generated qrels are on a per-run basis. Last, we evaluate if significance decisions regarding the differences between evaluated systems under LLM judgements are similar to those observed under human judgements.

The main highlights of our results are the following. Correlations are not as good as they seem when we focus on the top-performing systems and degrade as we focus more and more on the first few systems. Additionally, there are many runs along the ranking that drop many positions when evaluated with LLM-based



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '25, Padua, Italy*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

Table 1: Statistics of the datasets used for experimentation. Note that $T_H \cap T_{LLM} = \emptyset$.

Data	DL-19	DL-20	DL-21	DL-22	DL-23
Topics with human judgements ($ T_H $)	43	54	53	76	82
Topics with LLM judgements ($ T_{LLM} $)	157	146	424	424	618
Average judgements per topic in T_H	215	211	204	5084	272
Average judgements per topic in T_{LLM}	70	128	157	133	73
Submissions	36	59	63	100	35
Pairs of submissions	630	1711	1953	4950	595

qrels, which suggests that they are not entirely fair. Finally, we observe that statistical tests made with LLM-based judgements yield an exceedingly high rate of false positives. Overall, our experiments suggest that LLMs are not yet ready to replace humans for making gold relevance assessments in IR offline evaluation.

2 Differences among Top Systems

In this section, we have a closer look at how LLM-based qrels are able to preserve the same ranking of systems among top-performing ones. We used the data of the SynDL dataset [13] to perform our experiments. This dataset is based on the passage ranking task of the Deep Learning Track, from years 2019 to 2023. Each year, track organizers created a set of queries from which some received human judgements. For example, in DL19, only 43 out of the 200 queries received expert assessments. The creators of SynDL used GPT-4 to provide labels for the rest of the queries that did not receive human judgements. For more details about how the LLM-based labels have been generated, we refer the reader to the paper [13]. Table 1 summarises the data we have used. It is important to note that the set of topics that have human judgements and the set of topics that have LLM judgements are disjoint.

To evaluate to which extent LLM-based qrels can preserve the ordering among top-performing systems, we have computed two different top-weighted correlations, namely τ_{AP} [23] and Rank-Biased Overlap (RBO)¹ [21]. The main idea is that, when comparing two rankings of systems, discrepancies among those systems with high positions are penalized more than those with low positions. Table 2 shows the top-weighted systems ranking correlations, along with Kendall’s τ , between official TREC judgements and LLM-based judgements. As already demonstrated in previous work [4, 13, 17], LLM-based qrels show competitive values regarding Kendall’s τ . However, we can observe that the correlations that penalize more for failures in higher positions give significantly lower values, which highlights the idea that the judgements generated with an LLM are not able to identify the best systems.

2.1 Fairness

These correlations are useful for showing overall performance on large-scale experiments, but they hide the fact that some runs may be treated differently in the two evaluations. That is, it could be the case that, although correlations are actually high, some runs are losing (or gaining) many positions in the ranking.

¹We used the parameter $p = 0.7$ to focus on the first third or four runs [15, Section 4.5]

Table 2: Systems Ranking correlations between official TREC human-made judgements and LLM-based judgements.

Metric		Dataset				
		DL-2019	DL-2020	DL-2021	DL-2022	DL-2023
AP	τ	0.62	0.85	0.80	0.75	0.79
	τ_{AP}	0.49	0.70	0.72	0.63	0.66
	RBO ($p = 0.7$)	0.11	0.43	0.52	0.51	0.33
NDCG	τ	0.80	0.90	0.87	0.85	0.88
	τ_{AP}	0.79	0.84	0.85	0.77	0.83
	RBO ($p = 0.7$)	0.93	0.54	0.97	0.58	0.69

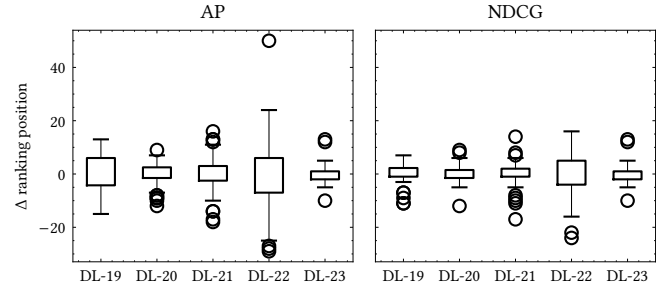


Figure 1: Results of the fairness experiment. For each year and metric, we plot the distribution of the changes in ranking positions that runs undergo when evaluated with LLM-generated assessments. A positive change means that the run is ranked at a higher position under LLM judgements, while a negative value means that the run is ranked at a lower position.

Following a common approach in past work [8, 20], we evaluate the **fairness**. When evaluated with LLM-generated judgements, we compute the change that each run suffers with respect to the ranking of systems inferred from human judgements. We show the results of this experiment in Figure 1. This figure depicts the distributions, as boxplots, of the changes in ranking positions that runs undergo when evaluated with LLM-based assessments. A positive value means that the run is at a higher position with LLM qrels, while a negative value means that the run is at a lower position. Overall, we observe that almost all runs suffer some change in position when evaluated with LLM judgments. More worrying is the fact that there are, in all cases, regardless of the metric or the collection, runs that suffer changes of 5 or more positions, even reaching changes of more than 40 positions.

As a follow-up analysis, in Figure 2 we include the change suffered by each of the runs that participated in the DL-2020. The x-axis shows the ranking of the runs inferred from human judgements, while the y-axis shows the change that each run suffers, in terms of ranking position, when evaluated with LLM judgements. Overall, we observe that most of the runs suffer a change. Only 10 and 13 runs stay at the same position when evaluated with AP and NDCG, respectively. In other words, 83% of runs with AP, and 78% with NDCG, get a different position when evaluated with LLM-based qrels.

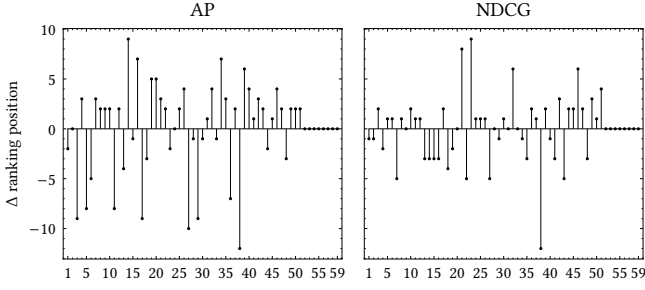


Figure 2: Per-run drops in DL-2020. The x-axis shows the position of each run in the ranking inferred from human judgements. The y-axis shows the change in ranking position when evaluated with LLM-based judgements. A positive change means that the run is ranked at a higher position under LLM judgements, while a negative value means that the run is ranked at a lower position.

3 Preservation of Statistical Significances

In the previous section, we evaluated biases exhibited by LLM-based relevance assessments in ranking correlations and ranking swaps, especially among top-performing systems. In this section, we study the reliability of LLM-generated judgements to evaluate retrieval systems from the perspective of statistical tests. Assuming that human judgements are the ground truth, we evaluate whether the LLM labels can provide a similar pairwise significance evaluation, i.e., whether decisions regarding the statistical difference between a pair of systems are the same using human labels and LLM labels.

Let Q_H be the judgements made by human assessors over a set T_H of topics, and let Q_{LLM} be the judgements made by an LLM over a set T_{LLM} of topics. We do not assume any overlap between T_H and T_{LLM} , in fact, they may be completely disjoint sets and have different sizes. Let R be the set of runs under evaluation, where each run includes results for all topics $T_H \cup T_{LLM}$. Using an evaluation measure of choice (e.g. Average Precision), we compute per-topic scores for each run in R using both sets of judgements independently. Thus, we have two arrays of per-topic scores for each run, one obtained using Q_H and the other obtained with Q_{LLM} . With these per-topic scores, we compute a p -value for each pair of runs,² for both sets of qrels. So, in the end, we have, on the one hand, a p -value for each pair of runs computed from the scores obtained with Q_H and, on the other hand, a different p -value for each pair of runs computed from the scores obtained with Q_{LLM} . Thus, for each pair of runs, given a significance level α , we know if they are significantly different (or not) under Q_H and under Q_{LLM} . Now, assuming that the significance decisions derived from the Q_H qrels are the ground truth, we evaluate how the decisions derived from Q_{LLM} match the former. In particular, we compute the true/false positives/negatives:

- **True Positives (TP):** these are the cases where both set of judgements marked a significant result.
- **False Negatives (FN):** these are the cases where a pair of runs is marked as significantly different under Q_H , but not under Q_{LLM} .

²We use a two-sided test.

Table 3: Classification results observed using the synthetic qrels on different datasets with AP and NDCG.

Metric		Dataset				
		DL-2019	DL-2020	DL-2021	DL-2022	DL-2023
AP	TP	80%	94%	96%	95%	97%
	FN	20%	6%	4%	5%	3%
	TN	29%	42%	17%	20%	11%
	FP	71%	58%	83%	80%	89%
NDCG	TP	92%	98%	98%	97%	98%
	FN	8%	2%	2%	3%	2%
	TN	33%	46%	28%	26%	6%
	FP	67%	54%	72%	74%	94%

- **True Negatives (TN):** these are the cases where both set of judgements marked a non-significant result.
- **False Positives (FP):** these are the cases where a pair of runs is marked as non-significant under Q_H , but marked as significant under Q_{LLM} .

We report these metrics as percentages over the gold counts of significant and non-significant. TP will be the percentage of correctly detected positives over the total of gold positives, FN will be the ratio of incorrectly detected negatives over the total of positives, and so on. Thus, $TP + FN = 100\%$ and $TN + FP = 100\%$.

3.1 Experimental Setup

We used both Average Precision (AP) and NDCG (Normalized Discounted Cumulative Gain) with a 1000 cutoff for scoring the runs. To compute the p -values, we employed a two-sided Wilcoxon Signed-Rank test as recommended in previous literature [9, 11], and set $\alpha = 0.05$.

3.2 Results

We show the results of our experiments in Table 3. This table reports the TP/TN/FP/FN percentages both when scoring the runs with AP (upper half) and NDCG (lower half).

From this table, we highlight several interesting results. First, we see systematically better figures with NDCG than MAP—higher positives, lower negatives—, suggesting that NDCG is more robust for evaluating systems with LLM-qrels. We observe that LLM-generated labels detect most of the existing significant differences, since TP ratios are near 100% almost in every case. This result holds independently of using AP or NDCG, although TP figures obtained with NDCG are systematically higher. Interestingly, we observe a high number of false positives for every year and metric. These cases represent situations where the LLM judgements Q_{LLM} marked a significant difference when there was none. In other words, some pairs of systems are not significantly different under Q_H but they are under Q_{LLM} . It is a concerning result, since judgements created with an LLM would lead a researcher to incorrectly conclude and present a supposedly significant result. However, we hypothesize that this might be happening because of the difference in sample sizes between both evaluations. As observed in Table 1, the number

Table 4: Classification results observed using the SynDL synthetic qrels on different datasets with AP and NDCG. We randomly undersample T_{LLM} topics such that $|T_{LLM}| = |T_H|$.

Metric		Dataset				
		DL-2019	DL-2020	DL-2021	DL-2022	DL-2023
AP	TP	68%	88%	88%	89%	91%
	FN	32%	12%	12%	11%	9%
	TN	57%	62%	50%	51%	28%
	FP	43%	38%	50%	49%	72%
NDCG	TP	83%	94%	91%	92%	93%
	FN	17%	6%	9%	8%	7%
	TN	58%	63%	64%	55%	33%
	FP	42%	37%	36%	45%	67%

of topics that include LLM-generated labels is much higher than the topics that have human judgements. It is known that more topics provide more power when performing a statistical test [3, 14, 22]. Therefore, these false positives could be happening for this reason.

To check this hypothesis, we sampled random subsets of T_{LLM} so that $|T_H| = |T_{LLM}|$. In this way, both evaluations have an equal sample size. We sampled 200 times, and computed the average of ratios over these iterations. We present the results in Table 4.

These results show that, indeed, under equal sample size conditions, the number of false positives is lower, supporting our previous hypothesis, although they are still concerning. We also observe lower values of true positives. Now some differences that are real under Q_H are being ignored under Q_{LLM} . Leaving some real advancements behind can also hinder progress in the field. Regarding the differences between true negatives and true positives, we see that LLM-generated judgements are equally effective at detecting true significances and non-significances. In this case, we also observe better figures with NDCG, a result in line with previous research in related fields [19]. Overall, the results of this experiment showcase that LLMs are not yet capable of producing IR relevance labels that preserve the same statistical differences as human labels.

Similar to the previous section, we are also interested in whether a run has more or less significant differences when evaluated under the LLM-based qrels. For every run, we count the significant differences observed under human and LLM judgements against this run. Then, a *drop* happens when the number is higher in the first case, meaning that a run is losing some significances under the LLM labels. We also undersampled T_{LLM} 50 times and averaged the results. We plot the distribution of these drops in Figure 1.

This figure shows the distribution of the per-run significance drops, averaged over 50 undersampling iterations, for the different year and effectiveness measures we have used. The ideal case here would be observing narrow boxes near 0 without outliers. This would mean that, on average, no run is treated differently in terms of significances. However, we observe that, to a greater or lesser degree, there is always some run treated differently, regardless of the year and the effectiveness measure. This suggests that the false negatives we observed in a previous experiment (see Table 4) are concentrated within just a few runs. In other words, the false

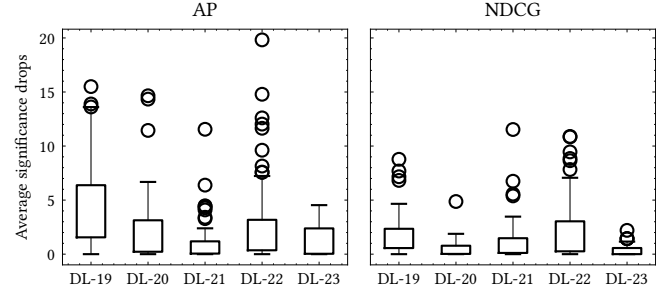


Figure 3: Distribution of the drops in significance counts of the runs when evaluated with LLM labels. Each value is the average over 50 iterations of undersampling T_{LLM} .

negatives appear because the LLM qrels consistently fail to correctly evaluate the same few runs.

4 Conclusions

Making relevance judgements is the most resource-intensive task when creating new collections for IR evaluation. Recent research has explored the possibility of using LLMs to assess the relevance of documents, showing promising results: LLMs can produce accurate labels compared to human assessments and reliable rankings of systems. However, little work has evaluated the validity of these judgements in terms of distinguishing the best top-performing systems and preserving statistical significance stability. In this work, we looked at these two issues and evaluated the reliability of LLM labels for IR evaluation from different perspectives. First, we evaluated if LLM-based qrels were able to correctly rank top-performing systems by employing top-heavy ranking correlation. Results showed that correlations degrade as we focus more and more on the best systems. We also evaluated whether significance decisions regarding the differences between evaluated systems are similar to those observed under human judgements. Our results showed that LLM-generated judgements yield an exceedingly high rate of false positives. This result suggests they are biased to incorrectly mark a significant difference when there is none. We also looked at whether any of the runs suffered great changes in their ranking position when evaluated with LLM judgements. In this experiment, we observed that there are runs that suffer drops of up to 50 positions, which suggests that these LLM-generated judgements are not entirely fair for evaluating IR systems.

Acknowledgments

The authors thank the financial support supplied by the grant PID2022-137061OB-C21 funded by MICIU/AEI/10.13039/501100011033 and by “ERDF/EU”. The authors also thank the funding supplied by the Consellería de Cultura, Educación, Formación Profesional e Universidades (accreditations ED431G 2023/01 and ED431C 2025/49) and the European Regional Development Fund, which acknowledges the CITIC Research Center as a Center of Excellence and recognizes it as a Member of the CIGUS Network for the period 2024-2027.

References

- [1] Zahra Abbasiantaeb, Chuan Meng, Leif Azzopardi, and Mohammad Aliannejadi. 2024. Can we use large language models to fill relevance judgment holes? (2024). <https://arxiv.org/abs/2405.05600> arXiv: 2405.05600 [cs. IR] (cited on p. 1).
- [2] Marwah Alaofi, Paul Thomas, Falk Scholer, and Mark Sanderson. 2024. LLMs can be fooled into labelling a document as relevant: best café near me; this paper is perfectly relevant. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region (SIGIR-AP 2024)*. Association for Computing Machinery, Tokyo, Japan, 32–41. ISBN: 9798400707247. doi:10.1145/3673791.3698431 (cited on p. 1).
- [3] Ben Carterette, Virgil Pavlu, Evangelos Kanoulas, Javed A. Aslam, and James Allan. 2008. Evaluation over thousands of queries. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, New York, NY, USA, 651–658. ISBN: 9781605581644. doi:10.1145/1390334.1390445 (cited on p. 4).
- [4] Guglielmo Faggioli et al. 2023. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval (ICTIR '23)*. Association for Computing Machinery, Taipei, Taiwan, 39–50. doi:10.1145/3578337.3605136 (cited on pp. 1, 2).
- [5] Maurice G. Kendall. 1948. *Rank Correlation Methods*. Charles Griffin and Company Limited (cited on p. 1).
- [6] Mucahid Kutlu, Tamer Elsayed, Maram Hasanain, and Matthew Lease. 2018. When rank order isn't enough: new statistical-significance-aware correlation measures. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. Association for Computing Machinery, Torino, Italy, 397–406. ISBN: 9781450360142. doi:10.1145/3269206.3271751 (cited on p. 1).
- [7] Sean MacAvaney and Luca Soldaini. 2023. One-shot labeling for automatic relevance estimation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*. Association for Computing Machinery, Taipei, Taiwan, 2230–2235. doi:10.1145/3539618.3592032 (cited on p. 1).
- [8] David Otero, Javier Parapar, and Álvaro Barreiro. 2023. Relevance feedback for building pooled test collections. *Journal of Information Science*. doi:10.1177/01655515231171085 (cited on p. 2).
- [9] David Otero, Javier Parapar, and Álvaro Barreiro. 2025. Towards reliable testing for multiple information retrieval system comparisons. In *Advances in Information Retrieval*. Springer Nature Switzerland, Cham, 424–439. ISBN: 978-3-031-88711-6 (cited on p. 3).
- [10] David Otero, Javier Parapar, and Nicola Ferro. 2023. How discriminative are your qrels? how to study the statistical significance of document adjudication methods. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*. Association for Computing Machinery, New York, NY, USA, 11 pages. doi:10.1145/3583780.3614916 (cited on p. 1).
- [11] Javier Parapar, David E. Losada, and Álvaro Barreiro. 2021. Testing the tests: simulation of rankings to compare statistical significance tests in information retrieval evaluation. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing (SAC '21)*. Association for Computing Machinery, Virtual Event, Republic of Korea, 655–664. doi:10.1145/3412841.3441945 (cited on p. 3).
- [12] Hossein A. Rahmani, Nick Craswell, Emine Yilmaz, Bhaskar Mitra, and Daniel Campos. 2024. Synthetic test collections for retrieval evaluation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, Washington DC, USA, 2647–2651. doi:10.1145/3626772.3657942 (cited on p. 1).
- [13] Hossein A. Rahmani, Xi Wang, Emine Yilmaz, Nick Craswell, Bhaskar Mitra, and Paul Thomas. 2024. Syndl: a large-scale synthetic test collection for passage retrieval. (2024). <https://arxiv.org/abs/2408.16312> arXiv: 2408.16312 [cs. IR] (cited on p. 2).
- [14] Mark Sanderson and Justin Zobel. 2005. Information retrieval system evaluation: effort, sensitivity, and reliability. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '05)*. Association for Computing Machinery, Salvador, Brazil, 162–169. ISBN: 1595930345. doi:10.1145/1076034.1076064 (cited on p. 4).
- [15] Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. 2024. Large language models can accurately predict searcher preferences. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, Washington DC, USA, 1930–1940. doi:10.1145/3626772.3657707 (cited on pp. 1, 2).
- [16] Shivani Upadhyay, Ehsan Kamaloo, and Jimmy Lin. 2024. LLMs can patch up missing relevance judgments in evaluation. (2024). <https://arxiv.org/abs/2405.04727> [cs. IR] (cited on p. 1).
- [17] Shivani Upadhyay, Ronak Pradeep, Nandan Thakur, Daniel Campos, Nick Craswell, Ian Soboroff, Hoa Trang Dang, and Jimmy Lin. 2024. A large-scale study of relevance assessments with large language models: an initial look. (2024). <https://arxiv.org/abs/2411.08275> arXiv: 2411.08275 [cs. IR] (cited on pp. 1, 2).
- [18] Shivani Upadhyay, Ronak Pradeep, Nandan Thakur, Nick Craswell, and Jimmy Lin. 2024. Umbrela: umbrella is the (open-source reproduction of the) bing relevance assessor. arXiv:2406.06519 (cited on p. 1).
- [19] Daniel Valcarce, Alejandro Bellogin, Javier Parapar, and Pablo Castells. 2018. On the robustness and discriminative power of information retrieval metrics for top-n recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. Association for Computing Machinery, Vancouver, British Columbia, Canada, 260–268. doi:10.1145/3240323.3240347 (cited on p. 4).
- [20] Ellen M. Voorhees. 2018. On building fair and reusable test collections using bandit techniques. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. Association for Computing Machinery, Torino, Italy, 407–416. ISBN: 9781450360142. doi:10.1145/3269206.3271766 (cited on p. 2).
- [21] William Webber, Alistair Moffat, and Justin Zobel. 2010. A similarity measure for indefinite rankings. *ACM Transactions on Information Systems*, 28, 4, Article 20, 38 pages. doi:10.1145/1852102.1852106 (cited on p. 2).
- [22] William Webber, Alistair Moffat, and Justin Zobel. 2008. Statistical power in retrieval experimentation. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (CIKM '08)*. Association for Computing Machinery, Napa Valley, California, USA, 57100580. ISBN: 9781595939913. doi:10.1145/1458082.1458158 (cited on p. 4).
- [23] Emine Yilmaz, Javed A. Aslam, and Stephen Robertson. 2008. A new rank correlation coefficient for information retrieval. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '08)*. Association for Computing Machinery, Singapore, Singapore, 587–594. ISBN: 9781605581644. doi:10.1145/1390334.1390435 (cited on p. 2).