

LLM-based Relevance Assessment Still Can't Replace Human Relevance Assessment

Charles L. A. Clarke
University of Waterloo
Canada

Laura Dietz
University of New Hampshire
USA

Abstract

The use of large language models (LLMs) for relevance assessment in information retrieval has gained significant attention, with recent studies suggesting that LLM-based judgments provide comparable evaluations to human judgments. Notably, based on TREC 2024 data, Upadhyay et al. [11] make a bold claim that LLM-based relevance assessments, such as those generated by the UMBRELA system, can fully replace traditional human relevance assessments in TREC-style evaluations. This paper critically examines this claim, highlighting practical and theoretical limitations that undermine the validity of this conclusion.

First, we question whether the evidence provided by Upadhyay et al. [11] genuinely supports their claim, particularly when the test collection is intended to serve as a benchmark for future research innovations. Second, we submit a system deliberately crafted to exploit automatic evaluation metrics, demonstrating that it can achieve artificially inflated scores without truly improving retrieval quality. Third, we simulate the consequences of circularity by analyzing Kendall's tau correlations under the hypothetical scenario in which all systems adopt UMBRELA as a final-stage re-ranker, illustrating how reliance on LLM-based assessments can distort system rankings. Theoretical challenges — including the inherent narcissism of LLMs, the risk of overfitting to LLM-based metrics, and the potential degradation of future LLM performance — that must be addressed before LLM-based relevance assessments can be considered a viable replacement for human judgments.

1 Introduction

In early 2023, Faggioli et al. [6] applied GPT-3.5 (text-davinci-003) to fully reassess runs submitted to the TREC 2021 Deep Learning track. They reported Kendall's $\tau = 0.86$ for human vs. LLM-based assessment on NDCG@10. A natural conclusion might be that LLMs could now replace humans for routine relevance assessment. Instead, Faggioli et al. [6] issue a warning. While recognizing the potential of LLMs to improve ranking and acknowledging their value as part of the relevance assessment process, they argue strongly against abandoning human assessment. They raise concerns about the potential for unknown biases that LLM-based assessments might introduce. They highlight the issue of circularity, where LLMs evaluate the outputs of other LLMs. Most importantly, their primary concern is that “LLMs are not people.” Since information retrieval systems are designed to serve human needs, their evaluation must ultimately reflect human judgment and preferences.

Recently, Upadhyay et al. [11] analyze the retrieval task results from the TREC 2024 RAG track. This retrieval task (or “R task”) mirrors a traditional TREC ad hoc retrieval task. Participating systems were tasked with executing 301 queries over the MS MARCO Segment V2.1 collection, producing a ranked set of 100 passages for

each query (a “run”). They compare four procedures for assessing runs: 1) fully automatic assessments using the UMBRELA LLM-based relevance assessment tool [12]; 2) fully manual assessments using established TREC evaluation protocols for human judgments; 3) a hybrid method where UMBRELA filtered the set of passages to be judged; and 4) a hybrid method where humans refined UMBRELA's assessments. In this paper, we focus on the first two procedures: fully automatic and fully manual assessments. Based on their analysis, Upadhyay et al. [11] conclude:

*Our results suggest that **automatically generated UMBRELA judgments can replace fully manual judgments** to accurately capture run-level effectiveness. Surprisingly, we find that LLM assistance does not appear to increase correlation with fully manual assessments, suggesting that **costs associated with human-in-the-loop processes do not bring obvious tangible benefits**.... Our work validates the use of LLMs in academic TREC-style evaluations and provides the foundation for future studies.*

We disagree. Not only does their reported correlation fail to provide stronger evidence than that of Faggioli et al. [6], but additional evidence from the track directly contradicts their conclusion. This evidence includes runs submitted by team WaterlooClarke, which were explicitly designed to subvert LLM-based relevance judgments by employing LLM-generated judgments as a final-stage ranker.

Faggioli et al. [6] already demonstrated a strong empirical correlation between manual judgments and LLM judgments, both in terms of inter-annotator agreement and leaderboard correlation, and many other work also observed this empirical correlation. However, after a detailed consideration of competing views, Faggioli et al. [6] concluded that there are too many theoretical concerns before human judgments can be replaced. These concerns, which remain critical to the discussion, **have neither been addressed nor refuted** in the work of Upadhyay et al. [11, 12].

Conceptually, there is no fundamental difference between an LLM-based relevance assessment and an LLM-based re-ranking method. Both predict an affinity score for a passage to be relevant for a given query. In contrast, human relevance judgments are privileged precisely because they *are* created by humans, and only humans can provide a gold standard for the evaluation of usefulness. While employing LLMs to train and implement rankers can lead to substantial performance gains, these improvements risk being illusory if they fail to reflect human judgments. The observation that LLM-based relevance judgments closely mimic the outcomes of human relevance judgments suggests that these LLM assessments may themselves represent a strong ranking method, rather than a valid evaluation metric.




Figure 1: Scatter plot extracted from Figure 3 of Upadhyay et al. [11] comparing manual assessment and automatic assessment. Each red dot represents the average performance of a run over all queries. The blue dots plot all run-query combination. The inset provides a closer view of the points in the green rectangle.

We acknowledge the value of work by Upadhyay et al. [12]. In particular, their use of a new collection with fresh queries, which guarantees that LLMs were not trained on this collection. Moreover, their work confirms that the correlation seen by Faggioli et al. [6] is not merely an artifact of training on the test collection. However, the warnings issued by Faggioli et al. [6] remain both valid and increasingly urgent, especially given the growing prevalence of LLM-based relevance assessments in information retrieval tasks.

2 Reproduction of Umbrella Results

Kendall’s τ is a suitable metric for showing overall rank correlations on large-scale experiments. Upadhyay et al. [11] report a high overall Kendall’s $\tau = 0.89$ between manual relevance assessments and automatic UMBRELA assessments. However, we note that some submitted systems perform substantially worse than others, making them easy to distinguish by an evaluation. The presence of such under-performing systems can artificially inflate Kendall’s τ scores.

To investigate this effect, we extend the original analysis by excluding the bottom-ranked 15 systems (out of 75) and recomputing Kendall’s τ over the remaining top 60 systems. As shown in Figure 2, this yields a slightly lower but still relatively high Kendall’s τ of 0.84. This corresponds to approximately 8% of system swaps, where the two evaluators disagree on which system performs better. Overall, these findings broadly confirm the results reported by Upadhyay et al. [11]. Nevertheless, we observe some notable outliers. For example, one system achieves a high LLM-based evaluation score (0.81) but a substantially lower manual evaluation score (0.63), as shown in Figure 2. We analyze such discrepancies further in Sections 4.1 and 4.2.

3 Differences among Top-Performing Systems

Demonstrating methodological advancements often involves re-using test collections with the goal of surpassing state-of-the-art systems. Identifying meaningful differences among top-performing




Figure 2: Reproduction of Upadhyay et al. [11]: On top 60 original TREC RAG 24 systems and data, the UMBRELA LLM evaluator correlates highly with manual assessors. Only few submitted retrieval systems included approaches from LLM evaluators. Each system represents one dot. Red dots mark systems provided by team WaterlooClarke, which are known to contain LLM evaluations for re-ranking.

runs is therefore critical for measuring significant progress. In contrast to the overall Kendall’s τ of 0.89 between manual and automatic assessments, the correlation weakens substantially among the highest-scoring systems. While Kendall’s τ among the top 60 systems¹ remains relatively high at 0.85 (corresponding to 8% system swaps), it drops to 0.51 (24% swaps) among the top 20 systems, and to 0.56 (21% swaps) among the top 15.

Thus, when using an LLM-based evaluator to demonstrate improvements over state-of-the-art systems, precise agreement with manual judgments is essential at the top of the leaderboard. A Kendall’s $\tau@15$ of 0.56 suggests that automatic UMBRELA assessments fail to demonstrate strong alignment with manual judgments at the top of the leaderboard. This misalignment undermines the reliability of automatic evaluations for tracking progress at the frontier of retrieval effectiveness.

Closer consideration of results from Upadhyay et al. [11] further highlights specific discrepancies in top-performing system rankings. Figure 1 reproduces a scatter plot extracted from Figure 3 of that paper, showing the performance of submitted runs (red dots) under manual assessment (y-axis) vs. automatic assessment (x-axis).² Focusing on the top-performing systems (inset of Figure 1), discrepancies become evident: for instance, the system ranked highest under automatic evaluation would only place fifth under manual evaluation. Conversely, the top system under manual evaluation ranks sixth under automatic evaluation. Particularly interesting is the case of the run circled in green. While it ranks fifth under automatic evaluation, it drops to 28th under manual evaluation.

These inconsistencies underscore a fundamental limitation of LLM-based assessments, such as those used by UMBRELA, in reliably

¹Throughout the paper, when we refer to “top-performing systems”, we use the manual relevance judgments as a basis for this assessment.

²At the time of writing, Upadhyay et al. [11] do not provide a full public data release. They have generously provided limited access to their data for the purpose of confirming factual statements in this paper.

LLM-based Relevance Assessment Still Can't Replace Human Relevance Assessment




Figure 3: Reranking with an LLM evaluator (UMBRELA) improves performance under human relevance labels. This plot compares the original and reranked versions of all TREC RAG 24 systems based on manual assessment.

identifying the best-performing systems. As a result, caution is warranted when using these methods for evaluating and validating progress in retrieval tasks. While some differences could be attributed to statistical noise — only 27 queries were manually judged, and no error bars are provided — the available evidence remains insufficient to justify the replacement of human judgments. This concern is particularly important when a collection is intended for re-use: claims of a novel system’s superiority over existing methods must be supported by improvements that are both statistically meaningful and aligned with manual assessments. At present, such confidence is lacking.

4 Subverting Automatic Evaluation

When task relevance labels are generated entirely through a publicly known automatic process, such as UMBRELA, the evaluation metric becomes vulnerable to manipulation. For instance, a participant could aggregate the outputs of many rankers, apply the UMBRELA system to this pooled set, and submit the resulting relevance labels as a new system for evaluation. Such a strategy could, in principle, achieve perfect scores across all metrics. Even if the specific LLM-based relevance assessment process includes undisclosed elements, such as the exact prompt or LLM used, participants could approximate the process enough to subvert the automatic evaluations.

4.1 Empirical Demonstration of the Risk

The run circled in green in Figure 1 exemplifies this vulnerability. Submitted by team WaterlooClarke as run uwc1, this submission was deliberately designed to subvert the automatic evaluation process. Specifically, the team pooled the top 20 documents from 15 preliminary runs, spanning neural and traditional rankers, with and without query expansion. This pool was then judged by GPT-4o with the prompt described by Faggioli et al. [6]. Top-graded passages were subsequently judged pairwise following the procedure of Clarke et al. [4] but substituting LLM-based assessments for crowdsourced human judgments.




Figure 4: Demonstration of the effects of circularity when using UMBRELA as both evaluator and ranker using TREC RAG 24 data. Each submitted retrieval system is first re-ranked with UMBRELA, then evaluated under NDCG with relevance labels from human judges and the UMBRELA evaluator. We see that especially among top ranked systems, the evaluation strategy no longer agrees with human judges on which system is better. Axis ranges are adjusted to display the same top 60 (of 75) systems as in Figure 2.

The final ranking for uwc1 was determined using LLM-based preference judgments as the primary key, LLM-based relevance assessments as the secondary key, and the reciprocal rank fusion [5] of preliminary runs as a tertiary tie-breaker. As intended, this deliberate attempt to manipulate the evaluation process led uwc1 to rank significantly higher under automatic assessment (5th) than under manual assessment (28th).

4.2 LLM-based Relevance Assessment for Re-Ranking

As Soboroff [10] recognizes: “Retrieval and evaluation are the same problem. Asking a computer to decide if a document is relevant is no different than using a computer to retrieve documents and rank them in order of predicted degree of relevance.” In this sense, LLM-based relevance assessment can be viewed as a specific form of LLM-based re-ranking. To employ an LLM-based relevance assessment tool as a LLM-based re-ranker, one starts with an initial ranking and prompts the LLM to assign a score — expressed as a relevance grade — to each of the top- k passages. These passages are then re-ranked according to their assigned relevance grades, preserving their original order among passages with equal grades.

This re-ranking interpretation is further illustrated by another run submitted by team WaterlooClarke, uwc2. For uwc2, the team re-ranked the track’s baseline run using the prompt described by Arabzadeh and Clarke [2]. This re-ranking improves the baseline’s performance from 9th to 4th place under manual assessments, and from 7th to 3rd place under automatic assessments.

4.3 Simulation of Circularity

To explore the potential effects of widespread adoption, we simulate what would have happened if TREC RAG 2024 systems had incorporated UMBRELA as part of their pipeline. Following the methodology described in Section 4.2, we apply UMBRELA as a final-stage re-ranker to the outputs of all submitted retrieval systems. For this experiment, we re-rank using the UMBRELA relevance labels from the track itself. While the track organizers employed UMBRELA for evaluation, the labels it produced could instead have been employed as a final-stage ranker. While we use the labels from the track organizers, any track participant could just as easily have employed UMBRELA as a final-stage re-ranker themselves.

Comparing original and re-ranked systems on manual (i.e., human) judgments in Figure 3, we see that re-ranking with the UMBRELA judgments consistently improves system performance. TREC and similar evaluation experiments generally allow participants to use any automatic re-ranking process for their submissions. Since UMBRELA re-ranking is an entirely automatic process, it immediately loses its value as a measurement tool for those ranking experiments. If any automatic re-ranking process can be used, using the measurement tool itself provides an optimal re-ranking [10]. More generally, if system performance is measured solely by LLM-based tools, system developers are strongly incentivized to incorporate the same tools into their systems.

Once we adopt UMBRELA as both a system component and an evaluation metric, it leads to an invalid **circular evaluation**. This effect is demonstrated in Figure 4, where the UMBRELA-re-ranked runs from Figure 3 are evaluated with the same UMBRELA relevance labels used for re-ranking and compared against evaluations based on human judgments. We observe a substantial increase in disagreement between the two evaluation methods, with discordant system pairs rising to 18% within the top 60 systems. As a result, Kendall’s τ drops sharply to 0.63. This degradation becomes even more pronounced at the top of the leaderboard: Kendall’s τ further decreases to 0.44 among the top 20 systems, 0.49 among the top 15, 0.38 among the top 10, and even turns negative ($\tau = -0.40$) among the top 5 systems. Under UMBRELA-based evaluation, twelve systems now obtain NDCG scores exceeding 0.95, implying near-perfect ranking performance. Yet the same systems achieve manual NDCG scores only between 0.68 and 0.72, illustrating substantial score inflation due to circularity.

For publication in peer-reviewed information retrieval research venues it is often necessary to demonstrate that a proposed system significantly outperforms all strong baselines. Under a circular evaluation, however, such findings would no longer be credible. These results highlight the risks of using LLM-based evaluation pipelines without safeguards against feedback loops, especially when test collections are intended for reuse. Taking manually created relevance labels as the gold standard, we conclude that evaluating systems using the UMBRELA LLM assessor — when those systems internally apply UMBRELA-based re-ranking — results in an invalid circular experimental evaluation.

5 Automatic Judgments are not Gold Standards

The prompts used to elicit relevance grades from LLM-based assessment tools resemble instructions typically given to human assessors. However, this resemblance is superficial and we should not be fooled by it. Such prompts merely represent one of many possible ways an LLM-based re-ranking method might assign scores, akin to an LLM-based point-wise ranker. Despite being commonly referred to as relevance assessments, these scores are not equivalent to the judgments produced by humans.

LLM-based relevance assessments cannot serve as a gold standard because they lack the grounding of a human carrying out an information task necessary to evaluate the usefulness of retrieval systems. A true gold standard must originate from human assessments, as only humans can determine the relevance of information in a way that reflects real-world utility.

Faggioli et al. [6] raised concerns about the potential unknown biases inherent in LLM-based assessments. However, one clear and concerning bias is that LLM-based relevance assessments tend to favor LLM-based ranking systems. Recently, Balog et al. [3] report a detailed evaluation of how LLM-based rankers can influence LLM-based judges, providing the first empirical evidence that LLM judges exhibit “a clear and substantial bias in favor of LLM-based rankers.” This bias has been observed in other contexts as well [8, 9, 13], where LLMs demonstrate a form of “narcissism,” disproportionately favoring outputs generated by similar models. Furthermore, Alaofi et al. [1] show that LLMs can be deceived through well-crafted prompt attacks embedded in content, leading them to incorrectly judge irrelevant text as relevant. These vulnerabilities highlight not only the susceptibility of LLM-based assessments to manipulation but also their inability to objectively evaluate diverse ranking approaches. Such biases and flaws further undermine the reliability of LLM-based assessments as a substitute for human judgments in critical tasks.

6 When Automatic Judgments become Useless

While the uwc1 run demonstrates how a bad actor can strategically subvert an evaluation experiment, it is reasonable to assume that most participants are well-intentioned. These participants are not merely competing to win but are contributing to the creation of reusable test collections that support the development of innovative systems. However, even without malicious intent, the next generation of information systems will likely incorporate the latest advancements in LLMs, including prompting LLMs for relevance. As a result, some ranking methods will inherently embed elements that mirror LLM-based relevance judgments. In a future evaluation experiment, it is plausible that even a well-intentioned participant could inadvertently undermine the evaluation process.

Looking ahead, we anticipate that future retrieval systems will increasingly rely on automatically generated training data to optimize machine learning components. Here, Goodhart’s law serves as a cautionary principle [7]: “When a measure becomes a target, it ceases to be a good measure.” While the current observed correlation between manual and automatic assessment methods is strong, we predict that this correlation will degrade as developers incorporate LLM-based evaluation components into their systems in more refined ways than our simulation in Section 4.3. Over time, these

systems risk becoming disconnected from the human judgments they are intended to serve. If the entire end-to-end experimental pipeline — from query formulation to relevance labeling — is fully automated, the evaluation process devolves into an LLM assessing its own assessments. The circularity feared by Faggioli et al. [6] is no longer a hypothetical concern; it has already begun to manifest in practice.

7 Conclusion

This paper raises serious concerns about the claims made by Upadhyay et al. [11], which presents a preliminary analysis of data from the retrieval (“R”) task of the TREC 2024 RAG Track. The author list of Upadhyay et al. [11] includes some of the most prominent experts in the area of information retrieval evaluation. Despite being preliminary, their conclusions strongly imply that LLM-based relevance assessment can replace human relevance assessment — a claim that does **not** withstand scrutiny. Given the authority of the authors and the strength of their implied conclusions, there is a risk that these findings may gain widespread acceptance within the research community without sufficient critical consideration. Nearly two years ago, Faggioli et al. [6] reached the opposite conclusion based on similar evidence. Their concerns have still not been addressed.

Acknowledgments

We have discussed our concerns with some of the authors of Upadhyay et al. [11], and we appreciate their attention and feedback. They also generously provided us with early access to their data to confirm factual statements in this paper.

This material is based in part upon work supported by the National Science Foundation under Grant No. 1846017. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] 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*. 32–41.
- [2] Negar Arabzadeh and Charles L. A. Clarke. 2024. A Comparison of Methods for Evaluating Generative IR. arXiv:2404.04044 [cs.IR] <https://arxiv.org/abs/2404.04044>
- [3] Krisztian Balog, Donald Metzler, and Zhen Qin. 2025. Rankers, Judges, and Assistants: Towards Understanding the Interplay of LLMs in Information Retrieval Evaluation. arXiv:2503.19092 [cs.IR] <https://arxiv.org/abs/2503.19092>
- [4] Charles L. A. Clarke, Alexandra Vtyurina, and Mark D. Smucker. 2021. Assessing Top-*k* Preferences. *ACM Transactions on Information Systems*, Article 33 (May’ 2021), 21 pages.
- [5] Gordon V. Cormack, Charles L. A. Clarke, and Stefan Buettcher. 2009. Reciprocal Rank Fusion Outperforms Condorcet and Individual Rank Learning Methods. In *32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 758–759.
- [6] Guglielmo Faggioli, Laura Dietz, Charles L. A. Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and Henning Wachsmuth. 2023. Perspectives on Large Language Models for Relevance Judgment. In *SIGIR International Conference on Theory of Information Retrieval*. 39–50.
- [7] Charles Goodhart. 1975. Problems of Monetary Management: The UK experience in papers in monetary economics. *Monetary Economics* 1 (1975).
- [8] Yiqi Liu, Nafise Sadat Moosavi, and Chenghua Lin. 2023. LLMs as narcissistic evaluators: When ego inflates evaluation scores. *arXiv preprint arXiv:2311.09766* (2023).
- [9] Arjun Panickssery, Samuel R Bowman, and Shi Feng. 2024. LLM Evaluators Recognize and Favor Their Own Generations. *arXiv preprint arXiv:2404.13076* (2024).
- [10] Ian Soboroff. 2025. Don’t Use LLMs to Make Relevance Judgments. *Information Retrieval Research* 1, 1 (March 2025), 29–46.
- [11] 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. arXiv:2411.08275 [cs.IR] <https://arxiv.org/abs/2411.08275>
- [12] 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 [cs.IR] <https://arxiv.org/abs/2406.06519>
- [13] Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghui Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. Large Language Models are not Fair Evaluators. *arXiv preprint arXiv:2305.17926* (2023).