

Revisiting Human-vs-LLM judgments using the TREC Podcast Track

Watheq Mansour¹[0000-0002-9463-595X], J. Shane Culpepper¹[0000-0002-9463-595X], Joel Mackenzie¹[0000-0001-7992-4633], and Andrew Yates²[0000-0002-5970-880X]

¹ The University of Queensland, Brisbane, Australia

² HLTCOE, Johns Hopkins University, Baltimore, Maryland, USA

Abstract. Using large language models (LLMs) to annotate relevance is an increasingly important technique in the information retrieval community. While some studies demonstrate that LLMs can achieve high agreement with ground truth (human) judgments, other studies have argued for the opposite conclusion. To the best of our knowledge, these studies have primarily focused on classic ad-hoc text search scenarios. In this paper, we conduct an analysis on user agreement between LLM and human experts, and explore the impact disagreement has on system rankings. In contrast to prior studies, we focus on a collection composed of audio files that are transcribed into two-minute segments – the TREC 2020 and 2021 Podcast Track. We employ five different LLMs to re-assess all of the query-segment pairs, which were originally annotated by TREC assessors. Furthermore, we re-assess a small subset of pairs where LLM and TREC assessors have the highest disagreement, and find that the human experts tend to agree with LLMs more than with the TREC assessors. Our results reinforce the previous insights of Sormunen in 2002 – that relying on a single assessor leads to lower user agreement – and emphasize the subjectivity of relevance assessment tasks.

Keywords: Large language models · Relevance judgment

1 Introduction

As researchers increasingly embrace the use of large language models (LLMs), the Information Retrieval (IR) community is now vested in understanding how reliable LLMs are at labeling data [1, 3, 5, 12, 20, 25]. Much of this work has been empirical in nature, where LLM assessors judge query-document pairs that have ground truth judgments created by human assessors, allowing us to compare LLM versus human performance [11, 29, 30, 31]. These studies typically consider ad-hoc passage or document retrieval tasks such as the TREC Deep Learning track, where the MSMARCO corpus is used as the document collection [4]. We are therefore motivated to consider other related tasks, in particular, collections with different properties, to understand how both the collection and retrieval task influence LLM assessments. To this end, we explore the use of

LLMs to assess relevance using the TREC Podcast Track, where the data is more challenging for both human and LLM judges since the collections contain transcription errors, and there is a loss of context across podcast segments. Our main contributions are as follows:

(1) We re-assess 18,284 pairs from the Podcast Track 2020/2021 pools using five different LLMs, resulting in a total of 91,420 judged pairs. (2) We exhaustively study assessor agreement and the impact on system ordering. (3) We identify a set of pairs with the highest disagreement between the original TREC assessors and the LLMs, and re-assess a subset of these topics using three senior IR research experts to better understand why LLMs and TREC judges disagree.

We find that the system ordering from the 2020 collection is relatively stable when using LLM assessments, with the top 3 system orderings remaining the same; however, the 2021 system ordering changes substantially. Surprisingly, we find that assessments from the TREC assessors systematically disagree with both the independent human experts and the LLM assessments. Initial analysis of the system descriptions suggests that LLMs favor lexical-based systems more than dense systems [2]. These findings reveal important issues in the original 2021 assessments and illustrate how challenging assessing relevance with podcast data can be, highlighting the often ambiguous nature of relevance assessment [27].

2 Related Work

We briefly outline the key work from this emerging area of research; for a more comprehensive overview, we refer the reader to the report from the first *LLM4Eval* workshop [24] and to recent perspective papers [6, 12, 26].

In 2023, Faggioli et al. [11] explored the advantages and disadvantages of using an LLM to generate relevance judgments with an IR test collection, observing highly correlated *system orderings* despite exhibiting only modest levels of judgment agreement between the two approaches. Upadhyay et al. [31] reproduced Thomas et al’s work [29] using OpenAI’s GPT-4o, and released an open-source toolkit called *UMBRELA*. Experiments using the TREC Deep Learning Tracks from 2019–2023 demonstrate that system rankings created with LLM labels are highly correlated to the ordering produced by human ground-truth judgments. Interestingly, Upadhyay et al. show several corner cases where LLM judgments were actually more accurate than the corresponding human judgments, which they attributed to the unreliability of human assessments, or a lack of a clear description of the user’s information need.

Upadhyay et al. [30] compared LLM and human judgments using three different configurations to measure the effect of including humans during the labeling process. The key observation is that *LLM judgments could replace human assessments when using many common IR effectiveness metrics, when the overall effectiveness ordering at the system run level is being measured*. When comparing agreement at the judgment level, they found that human assessors apply more stringent relevance criteria than LLMs currently do – meaning that LLMs tend to *over-rate* relevance compared to humans. In contrast, both Clarke and Dietz

[6], and Soboroff [26], argue against the replacement of humans with LLMs, and provide counterexamples to demonstrate the pitfalls of doing so. The authors argue that there is no clear line between an LLM judge and an LLM re-ranker.

3 Experimental Setup

Collection and Queries. The Spotify podcast corpus provides the “documents” used in the 2020 and 2021 TREC Podcast collections [7, 16, 17]. This dataset consists of 100,000 English podcasts published between 2019 and 2020 on the Spotify platform; the episodes constitute around 60,000 hours of audio. The audio collection was transcribed using Google’s Speech-to-Text API and then partitioned into two-minute segments (each with a one-minute overlap) to form the final text collection, producing around 3.4 million text segments (documents). Note that this automatic transcription process differentiates the podcast corpus from ad-hoc text corpora, as it often contains errors inherent to audio transcription [21], and segments do not necessarily align with a single context like a passage-based collection does (a segment can start halfway through a sentence, for example, which can be problematic to both humans and LLMs).

The 2020 and 2021 TREC Podcast Tracks each contain 50 topics. As is typical with TREC topics, each is accompanied by a short “title” query, and a longer description of the user information need. The relevance judgments were generated by *one assessor per topic*. NIST assessors had access to both the ASR transcript (including text before and after the text of the two-minute segment) and the corresponding audio segment.

Re-Assessing the Judgment Pool. We employ multiple open-source and proprietary LLMs to ensure that our findings are consistent. We use OpenAI’s GPT-4o as a proprietary model, as it has achieved the highest agreement with human judgments in recent studies [2, 30]. For our open-source LLMs, we use the following four models: (1) *Mistral*, the *Mistral-Small-Instruct-2409-Q6_K_L* model from the wider *Mistral* family [15]; (2) *Qwen*, the *Qwen2.5-14B-Instruct-Q8_0* model [32]; (3) *Llama3*, Meta’s *Meta-Llama-3.1-8B-Instruct-Q8_0* model [19]; and (4) *Gemma2*, Google’s *gemma-2-9b-it-Q8_0* [13]. All models are quantized to 8 bits with the exception of *Mistral* which is a 6-bit model. We use `llama-cpp-python`³ to support more efficient inference at scale. All of the open source models are publicly available.⁴

LLM Prompting. Before running the LLMs on the assessment pool, we select an instruction prompt using a 10% stratified random sample of TREC query-segment pairs. In our initial prompt, we included the TREC judgment guidelines in the *Description, Narrative, Aspects* (DNA) prompting style as it achieves the best performance on relevance assessment tasks [29]. The descriptions and narratives assist LLM in understanding the topic intent, and the retrieved segment that should be assessed, whereas the aspects guide the thinking process in a

³ <https://github.com/abetlen/llama-cpp-python>

⁴ <https://huggingface.co/collections/bartowski/>

step-by-step manner. The output is restricted to a JSON object that contains the relevance score and a short justification [28]. We experiment using three variants of the original prompt: a vanilla zero-shot DNA prompt; a prompt that asks LLMs to be more strict in its assessments, inspired by Upadhyay et al. [30] who show that LLMs tend to overestimate relevance compared to humans; and a prompt using in-context learning. We evaluated the quality of the prompts using Cohen Kappa’s score relative to the ground-truth judgments from TREC, and then chose the prompt that had the highest agreement. We used the best-performing prompt (the second variant) in all subsequent experiments.

Normalizing Relevance Grades. According to the TREC judgment guidelines, there are five relevance grades (0–4), and grade 4 was used only for “known item” and “refinding” topic types (and not for the “topical” category). However, upon examining the judgment pairs from both 2020 and 2021, we found that grade 4 was applied on every topic type, making it unclear how to differentiate between grades 3 and 4 based on the categorical description of each. In addition, perfect relevance (grade 4) is not reproducible by anyone but the topic creator. Therefore, we remap all such pairs with a grade of 4 to 3 to provide a more stable testing framework. Thus, we use a four-point relevance scale, which aligns with the graded relevance range used in previous TREC ranking tasks, such as the Deep Learning passage and document ranking tasks [8, 9, 10].

Reproducibility. In the interest of reproducibility, we make the used prompts and LLMs’ judgments publicly available.

 <https://github.com/Watheq9/ecir2026-podcast-judge>

4 Results

After exhaustively reassessing the entire set of query-document pairs derived from the TREC 2020 and 2021 Podcast Tracks using five LLMs, we evaluate how these judgments compare to the ground truth judgments from TREC. In the experiments presented below, all systems are ordered according to the mean RBP $\phi = 0.95$ score [23]; the same trends were observed using NDCG@10 [14].

System Ranking Evaluation. Table 1 reports the system ranking correlations using both an unweighted Kendall’s τ [18], and the top-weighted *Rank-Biased Alignment* (RBA) [22]. Two different settings for RBA are used: (1) a shallow version ($\phi = 0.8$, representing an expected depth of 5), focusing the weight of the comparison at the top of the ranking; and (2) a deeper version ($\phi = 0.9$, representing an expected depth of 10), spreading the weight more uniformly across all system rankings.

In the 2020 comparison, the results are quite stable, with high agreement between the system orderings produced using human relevance assessments compared to the LLM assessments, regardless of the metric or LLM being applied. In particular, observe that Kendall’s τ values are as high as 0.84, and that the

Table 1. Run order correlations for the TREC 2020 pool (28 systems, left) and the TREC 2021 pool (27 systems, right) when comparing human assessments to various LLMs.

Model	TREC 2020			TREC 2021		
	Kendall’s τ	RBA ϕ		Kendall’s τ	RBA ϕ	
		0.8	0.9		0.8	0.9
GPT-4o	0.81	0.96	0.94	0.66	0.94	0.92
Mistral	0.84	0.97	0.94	0.67	0.94	0.92
Qwen	0.83	0.97	0.94	0.62	0.94	0.92
Llama3	0.81	0.96	0.94	0.72	0.95	0.93
Gemma2	0.84	0.97	0.94	0.60	0.92	0.91

top-weighted RBA metric always returns results greater than 0.94, indicating that the top ranking system ordering is being preserved between the human and LLM judges. However, the 2021 judgments present a much different story, with Kendall’s τ values as low as 0.60, and lower RBA values in all of our comparisons. This indicates that the system ordering is much more volatile than in the 2020 data, including the *top* ranking systems.

To illustrate the system ordering volatility as the LLM assessor is changed, we plot the changes in rank position for each system, compared to the human judgments – Figure 1 shows the results for both 2020 (left) and 2021 (right), which align with Table 1. In 2020, it is clear that the top ranking system ordering is largely preserved, with small perturbations occurring after the third-best system. However, the systems’ rank changes are remarkably larger in 2021. The third- and fourth-ranked systems drop between one and five positions depending on the LLM used to create the judgments, the fifth-best system moves two positions up, and the eighth system drops up to *twelve* positions. Even more surprising is that the system that is originally ranked at position 13 moves into the top six systems, with similar large positive deltas observed as deep as rank 19. Initial analysis of the runs in Figure 1 suggests that LLM judgments may favor lexical systems (including hybrids or re-rankers with lexical first stages) as compared to strictly dense systems [2]. For example, the *QL* baseline jumped from rank 18 to 16, and from 22 to 15 in 2020 and 2021, respectively (See Gemma2 in Figure 1). However, more research is required to completely understand the instability of the 2021 data.

Human Assessor Agreement. To better understand when the LLM assessors disagree with the TREC assessors, we randomly sampled 22 out of 826 query-document pairs representing “high disagreements” – where the absolute difference in label between the TREC assessors and the (majority vote) LLM assessors was greater than two. Then, three IR experts independently judged these pairs. The annotation guidelines were the same prompt used when judging with LLMs, with additional information about query types. The pairs were provided in a sheet, where each pair was represented by the query, topic description, and a transcript

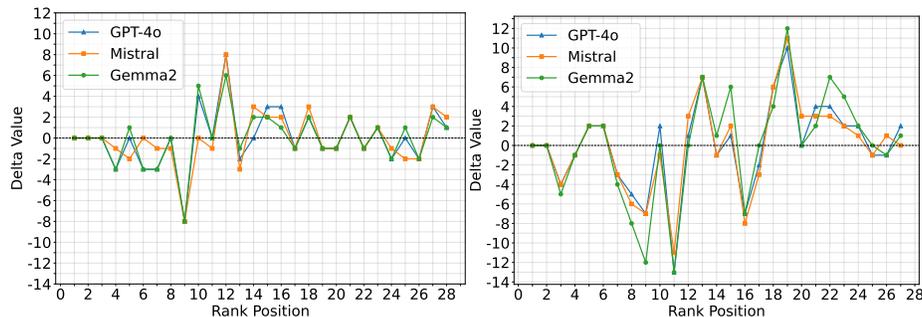


Fig. 1. Order volatility in TREC systems scored using RBP $\phi = 0.95$ in 2020 (left) and 2021 (right), when ordering changes using LLM assessments. The y -axis expresses the difference in rank position compared to the ground truth system ordering.

Table 2. Inter-rater agreement as measured with Krippendorff’s α between assessments of three senior human annotators, LLMs (majority vote), and the TREC judgments across the sampled query-document pairs from the TREC pool.

	Annotator 2	Annotator 3	TREC	LLMs
Annotator 1	0.67	0.73	-0.66	0.71
Annotator 2	-	0.82	-0.77	0.86
Annotator 3	-	-	-0.55	0.77
TREC	-	-	-	-0.76

of a 2-minute audio segment. Table 2 shows the inter-rater agreement between each expert assessor, the TREC assessor, and the LLM assessments. Surprisingly, the agreement between the expert assessors and the LLMs falls in the *tentative* to *reliable* range; on the other hand, there is a systematic disagreement between the TREC assessors and both the human assessors and the LLMs. This supports the earlier findings of Sormunen [27] who also demonstrated that the ambiguity of relevance assessments can result in vastly different outcomes for a query-document pair – and, in this context, it suggests that (many) LLMs may be more reliable than (one) human [31] and are clearly more reliable than other related work suggests. In total, 234 out of 9,386 pairs in 2020 (2.5%), and 1,014 out of 8,897 pairs in 2021 (11.4%) had TREC assessor assign a 0 label, compared to all of the LLM labels that were ≥ 2 . The converse (when the TREC label is a 2 or 3 and the LLM label is a 0) occurs in a much smaller number of disagreements – 11 and 28 pairs in 2020 and 2021, respectively – corroborating the notion that LLMs tend to assign higher relevance to a pair than humans [30], and providing a potential explanation for at least some of the instability observed for the 2021 collection.

5 Conclusion

We have revisited using LLMs as relevance assessors. We found that, although the correlation between the TREC and the LLM assessors was high in the 2020 TREC Podcast collection, it was much more volatile in the 2021 collection, raising doubts about the stability of the gold label assessments. Our analysis suggests that the LLM assessments tend to favor lexical systems, causing them to score much higher in system ranking comparisons. We also had three IR experts independently reassess a subset of pairs where the TREC and LLM judgments had the highest disagreement, and found that the new human judgments have a much higher agreement with the LLM labels than with the original human judgments. This preliminary work corroborates a number of recent findings on LLMs for relevance assessments using two new test collections, and further emphasizes the ambiguous nature of relevance assessment tasks. We plan to continue our analysis to better understand the instability we observed on the 2021 TREC Podcast campaign in future work.

Acknowledgments. We thank the anonymous referees for their feedback and suggestions. The third author was supported by a Google Research Scholar grant.

Disclosure of Interests. The authors have no competing interests of any sort.

Bibliography

- [1] Abbasiantaeb, Z., Meng, C., Azzopardi, L., Aliannejadi, M.: Can we use large language models to fill relevance judgment holes? In: Proc. EMT CIR (2024)
- [2] Alaofi, M., Thomas, P., Scholer, F., Sanderson, M.: LLMs can be fooled into labelling a document as relevant: best café near me; this paper is perfectly relevant. In: Proc. SIGIR-AP, pp. 32–41 (2024)
- [3] Arabzadeh, N., Clarke, C.L.A.: Benchmarking LLM-based relevance judgment methods. In: Proc. SIGIR, pp. 3194–3204 (2025)
- [4] Bajaj, P., Campos, D., Craswell, N., Deng, L., Gao, J., Liu, X., Majumder, R., McNamara, A., Mitra, B., Nguyen, T., Rosenberg, M., Song, X., Stoica, A., Tiwary, S., Wang, T.: MS MARCO: A Human Generated MACHine Reading COmprehension Dataset. arXiv:1611.09268v3 (2018)
- [5] Balog, K., Metzler, D., Qin, Z.: Rankers, judges, and assistants: Towards understanding the interplay of LLMs in information retrieval evaluation. In: Proc. SIGIR, pp. 3865–3875 (2025)
- [6] Clarke, C.L., Dietz, L.: LLM-based relevance assessment still can’t replace human relevance assessment. In: Proc. NTCIR (2025)
- [7] Clifton, A., Reddy, S., Yu, Y., Pappu, A., Rezapour, R., Bonab, H., Eskevich, M., Jones, G., Karlgren, J., Carterette, B., Jones, R.: 100,000 podcasts: A spoken English document corpus. In: Proc. COLING, pp. 5903–5917 (2020)

- [8] Craswell, N., Mitra, B., Yilmaz, E., Campos, D.: Overview of the TREC 2020 deep learning track. In: Proc. TREC (2021)
- [9] Craswell, N., Mitra, B., Yilmaz, E., Campos, D., Lin, J.: Overview of the TREC 2021 deep learning track. In: Proc. TREC (2021)
- [10] Craswell, N., Mitra, B., Yilmaz, E., Campos, D., Voorhees, E.M.: Overview of the trec 2019 deep learning track. arXiv preprint arXiv:2003.07820 (2019)
- [11] Faggioli, G., Dietz, L., Clarke, C.L., Demartini, G., Hagen, M., Hauff, C., Kando, N., Kanoulas, E., Potthast, M., Stein, B., et al.: Perspectives on large language models for relevance judgment. In: Proc. ICTIR, pp. 39–50 (2023)
- [12] Faggioli, G., Dietz, L., Clarke, C.L.A., Demartini, G., Hagen, M., Hauff, C., Kando, N., Kanoulas, E., Potthast, M., Stein, B., Wachsmuth, H.: Who determines what is relevant? Humans or AI? Why not both? *Comm. ACM* **67**(4), 31–34 (2024)
- [13] Gemma Team: Gemma 2: Improving open language models at a practical size. arXiv:2408.00118 (2024)
- [14] Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Sys.* **20**(4), 422–446 (2002)
- [15] Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., de las Casas, D., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., Lavaud, L.R., Lachaux, M.A., Stock, P., Scao, T.L., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mistral 7b. arXiv:2310.06825 (2023)
- [16] Jones, R., Carterette, B., Clifton, A., Eskevich, M., Jones, G.J., Karlgren, J., Pappu, A., Reddy, S., Yu, Y.: TREC 2020 podcasts track overview. In: Proc. TREC (2020)
- [17] Karlgren, J., Jones, R., Carterette, B., Clifton, A., Tanaka, E., Eskevich, M., Jones, G., Reddy, S.: TREC 2021 podcasts track overview. In: Proc. TREC (2021)
- [18] Kendall, M.G.: A new measure of rank correlation. *Biometrika* **30**(1/2), 81–93 (1938)
- [19] Llama 3 team: The llama 3 herd of models. arXiv:2407.21783 (2024)
- [20] MacAvaney, S., Soldaini, L.: One-shot labeling for automatic relevance estimation. In: Proc. SIGIR, pp. 2230–2235 (2023)
- [21] Mansour, W., Culpepper, J.S., Mackenzie, J.: Examining the impact of transcript variation on podcast search and re-ranking. In: Proc. ECIR, pp. 118–127 (2025)
- [22] Moffat, A., Mackenzie, J., Mallia, A., Petri, M.: Rank-biased quality measurement for sets and rankings. In: Proc. SIGIR-AP, pp. 135–144 (2024)
- [23] Moffat, A., Zobel, J.: Rank-biased precision for measurement of retrieval effectiveness. *ACM Trans. Inf. Sys.* **27**(1) (2008)
- [24] Rahmani, H.A., Siro, C., Aliannejadi, M., Craswell, N., Clarke, C.L.A., Faggioli, G., Mitra, B., Thomas, P., Yilmaz, E.: Report on the 1st workshop on large language model for evaluation in information retrieval (LLM4Eval 2024) at SIGIR 2024. SIGIR Forum (2024)
- [25] Rahmani, H.A., Yilmaz, E., Craswell, N., Mitra, B.: JudgeBlender: Ensembling automatic relevance judgments. In: Proc. WWW, pp. 1268–1272 (2025)

- [26] Soboroff, I.: Don't use LLMs to make relevance judgments. *Inf. Retr. Res.* **1**, 29–46 (2025)
- [27] Sormunen, E.: Liberal relevance criteria of TREC: Counting on negligible documents? In: *Proc. SIGIR*, pp. 324–330 (2002)
- [28] Tam, Z.R., Wu, C.K., Tsai, Y.L., Lin, C.Y., Lee, H.y., Chen, Y.N.: Let me speak freely? A study on the impact of format restrictions on performance of large language models. In: *Proc. EMNLP (Industry Track)* (2024)
- [29] Thomas, P., Spielman, S., Craswell, N., Mitra, B.: Large language models can accurately predict searcher preferences. In: *Proc. SIGIR*, pp. 1930–1940 (2024)
- [30] Upadhyay, S., Pradeep, R., Thakur, N., Campos, D., Craswell, N., Soboroff, I., Lin, J.: A large-scale study of relevance assessments with large language models using UMBRELA. In: *Proc. ICTIR*, pp. 358–368 (2025)
- [31] Upadhyay, S., Pradeep, R., Thakur, N., Craswell, N., Lin, J.: UMBRELA: UMBrela is the (Open-Source Reproduction of the) Bing RElevance Assessor. *arXiv 2406.06519* (2024)
- [32] Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., Lin, H., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Lin, J., Dang, K., Lu, K., Bao, K., Yang, K., Yu, L., Li, M., Xue, M., Zhang, P., Zhu, Q., Men, R., Lin, R., Li, T., Tang, T., Xia, T., Ren, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Wan, Y., Liu, Y., Cui, Z., Zhang, Z., Qiu, Z.: Qwen2.5 technical report. *arXiv:2412.15115* (2025)