

LLM-based Relevance Assessment for Web-Scale Search Evaluation at Pinterest

Han Wang*
Pinterest
San Francisco, CA, USA
hanwang@pinterest.com

Alex Whitworth*
Pinterest
San Francisco, CA, USA
awhitworth@pinterest.com

Pak Ming Cheung
Pinterest
San Francisco, CA, USA
pcheung@pinterest.com

Zhenjie Zhang
Pinterest
San Francisco, CA, USA
zhenjiezhang@pinterest.com

Krishna Kamath
Pinterest
San Francisco, CA, USA
kkamath@pinterest.com

Abstract

Relevance evaluation plays a crucial role in personalized search systems to ensure that search results align with a user’s queries and intent. While human annotation is the traditional method for relevance evaluation, its high cost and long turnaround time limit its scalability. In this work, we present our approach at Pinterest Search to automate relevance evaluation for online A/B experiments using fine-tuned LLMs. We rigorously validate the alignment between LLM-generated judgments and human annotations, demonstrating that LLMs can provide reliable relevance measurement for experiments while greatly improving the evaluation efficiency. Leveraging LLM-based labeling further unlocks the opportunities to expand the query set, optimize sampling design, and efficiently assess a wider range of search experiences at scale. This approach leads to higher-quality relevance metrics and significantly reduces the Minimum Detectable Effect (MDE) in online experiment measurements.

CCS Concepts

• **Information systems** → **Retrieval models and ranking; Web searching and information discovery.**

Keywords

Search Recommendation Systems, LLMs, Relevance Measurement

ACM Reference Format:

Han Wang, Alex Whitworth, Pak Ming Cheung, Zhenjie Zhang, and Krishna Kamath. 2025. LLM-based Relevance Assessment for Web-Scale Search Evaluation at Pinterest. In *RecSys 2025 EARL Workshop, September 22–26, 2025, Prague, Czech Republic*. ACM, New York, NY, USA, 6 pages. <https://doi.org/XXXXXXX.XXXXXXX>

*Equal contribution

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys '25, Prague, Czech Republic

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Search relevance measures how well search results align with user’s search query. For personalized search systems, evaluating relevance is crucial to ensure that displayed contents are pertinent to the user’s information needs, rather than over-relying on user’s past engagement. Online A/B experiments are ubiquitous in industry recommendation systems to measure the impact on user metrics. For search systems, it is essential to track changes in whole-page search relevance during experimentation.

Relevance evaluation typically relies on human annotations. However, this approach suffers from high costs, long turnaround time, and limited scalability, which can ultimately compromise the quality of relevance metrics. In recent years, with the advancement of Large Language Models (LLMs), there is a growing interest in automating relevance assessment in the Information Retrieval (IR) community. There have been several works exploring the feasibility of using LLMs to supplement laborious human labeling efforts [8, 25, 28].

At Pinterest, search is one of the key surfaces where users discover inspiring content that meets their information needs. We track whole-page relevance to ensure a high-quality user experience. In this paper, we present our methodology at Pinterest Search to leverage LLMs for assessing semantic relevance in online A/B experiments. We fine-tune open-source LLMs on relevance prediction tasks using human-annotated labels, then utilize the fine-tuned LLMs to evaluate the ranking results across experimental groups in online A/B experiments. This approach not only significantly reduces labeling costs and improves evaluation efficiency, but also unlocks opportunities to further improve metric quality by scaling up the query sets and refining the sampling design. We hope that this work is useful to other industry practitioners looking to improve the efficiency of relevance measurement.

We highlight our contributions as follows:

- We present the application of LLM-based relevance assessment at Pinterest Search, and outline the practical considerations and validation steps performed prior to adopting LLM-based automated relevance labeling.
- We demonstrate that fine-tuning small-sized LLMs on human-annotated relevance labels produces highly effective relevance models, with generated metrics closely aligning with those derived from human judgments. The average error in

query-level $sDCG@K$ metric (defined in Section 3.4) remains within 0.02.

- We highlight that LLM-based relevance assessment allows for expanding the query set and improving the query sampling design, which leads to higher-quality relevance metrics and substantially lowers the Minimum Detectable Effect (MDE) for online experiment evaluation.

2 Related Works

2.1 Search Relevance Modeling

For LLM-based search relevance modeling, there are two types of architectures commonly used in practice: bi-encoders and cross-encoders. Bi-encoder models encode queries and documents separately into a common dense space and then score their relevance using vector dot-product or cosine similarity. These models are typically trained by minimizing the contrastive loss with in-batch negative sampling [13, 15]. While bi-encoder models are efficient for retrieval tasks, they are limited in their ability to capture complex interactions between queries and documents. In contrast, cross-encoder models encode queries and documents jointly, hence can better capture the interaction between queries and documents. Various loss functions are used in the literature for training cross-encoder relevance models, such as pointwise classification loss [20], pairwise or listwise ranking loss [32], and contrastive loss [15]. In this work, we adopt a cross-encoder architecture and fine-tune the LLMs by minimizing the point-wise multi-class classification loss. The fine-tuned LLMs are then applied to whole-page relevance assessment at Pinterest Search.

2.2 LLMs for Relevance Judgments

Recent studies have shown the potential of using LLMs for automating relevance assessments in Information Retrieval (IR), as part of the broader “LLM-as-a-judge” paradigm [10, 31]. Initial explorations by Faggioli et al. [8] examined the application of LLMs to relevance judgment, providing insights into human-LLM collaboration. Subsequent work by Thomas et al. [25] and Upadhyay et al. [28] further demonstrated the capability of LLMs to accurately predict searchers’ preferences and generate human-quality relevance labels. Most current research investigates various prompting techniques for LLM-based relevance labeling, including zero-shot [3, 27, 28] and few-shot [16, 21] approaches, with a primary focus on evaluating their alignment with human labeling. Later, Meng et al. [17] and Abbasiantaeb et al. [1] showed that fine-tuning open-source LLMs on human-labeled relevance judgments is crucial for obtaining more reliable relevance predictions, suggesting that fine-tuning much smaller LLMs can yield more effective relevance prediction than few-shot prompting with significantly larger models.

Unlike most existing work that primarily compares different LLMs and prompting techniques [1, 17] or explores human-LLM collaboration [24], our work presents a real-world application of LLM-based relevance assessment to measure relevance change in online experiments at Pinterest Search. We discuss the practical considerations and validation steps, and highlight the opportunities for improving relevance measurement driven by this approach.

3 Methodology

3.1 Problem Statement

At Pinterest, we measure the semantic relevance between queries and Pins¹ using a 5-point guideline: Highly Relevant (L5), Relevant (L4), Marginally Relevant (L3), Irrelevant (L2), and Highly Irrelevant (L1). We use this guideline to measure the whole-page relevance for our search system. For online experiments with new search ranking models, assessing relevance across experimental groups is essential for monitoring the relevance impact. The relevance measurement procedure consists of the following components:

- Data collection: define and collect a representative sample.
- Relevance Labeling: rate each element of the sample.
- Metrics calculation: aggregate rated elements to get the overall relevance metrics.

These measurements have historically been limited by the low availability of human labels and the high marginal cost associated with generating them. This led to measurement designs and sample sizes that could only detect significant topline metric movements, but were insufficient to measure heterogeneous treatment effects or small topline effects. Next, we describe how we leverage LLMs to scale the labeling capabilities and address these bottlenecks.

3.2 Fine-tuned LLMs as Relevance Model

A cross-encoder model architecture is utilized to predict a Pin’s relevance to a given query, as illustrated in Figure 1. We fine-tune open-source LLMs on human-annotated data to optimize their performance on relevance prediction task. The training dataset contains approximately 2.6 million human-annotated query and Pin pairs, and model evaluation is performed on a held-out test split to prevent overfitting. To support search queries and Pins across multiple languages, we leverage multilingual LLMs to take advantage of their cross-lingual transfer capabilities. We formalize the relevance prediction as a multiclass classification problem based on the 5-point relevance guideline, minimizing the point-wise cross-entropy loss during training.




Figure 1: The cross-encoder architecture in the relevance teacher model. Take the encoder language models (e.g., BERT-based models) for illustration.

¹Pins on Pinterest are rich multimedia entities that feature images, videos and other content, often linked to external webpages or blogs.

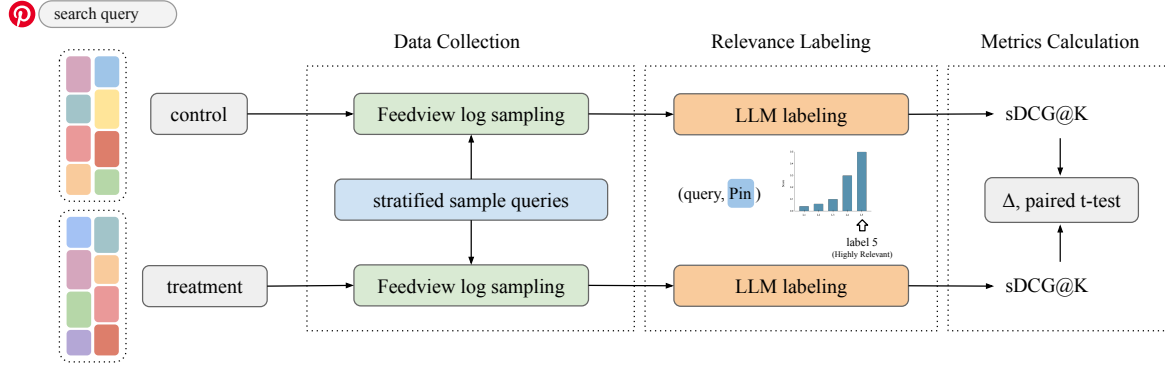


Figure 2: Components of LLM-based relevance measurement at Pinterest Search.

To effectively represent each Pin for relevance prediction, we leverage a comprehensive set of textual features, including Pin titles and descriptions, BLIP image captions [14], linked page titles and descriptions, user-curated board titles where the Pin has been saved, and highly-engaged queries associated with the Pin. These features together form a robust text representation crucial for accurate relevance assessment.

We experiment with various pre-trained language models, including multilingual BERT_{base} [7], T5_{base} [22], mDeBERTaV3_{base} [11], XLM-RoBERTa_{large} [6], and Llama-3-8B [2]. The comparative performance for these models and ablation studies on Pin text features can be found in Wang et al. [29]. We then use the fine-tuned model to generate 5-dimensional relevance scores and use the label corresponding to the highest score (argmax) for relevance assessment.

3.3 Stratified Sampling Design

LLM labeling significantly reduces relevance labeling costs as well as labeling time, which enables much larger sampling designs. Therefore, we propose a stratified query sampling design that enables measurement of heterogeneous treatment effects and reduces minimum detectable effects (MDEs) by an order of magnitude. Prior to LLM labeling, stratified query sampling with human annotations was impractical, as it required a large number of queries to adequately represent each fine-grained stratum.

Stratification plays an important role in sampling-based measurement. First, stratification ensures the sample population is representative of the whole population. In addition, if the strata are chosen such that each stratum is relatively homogeneous, variance reduction can be achieved [18]. From a statistical perspective, consider the whole-page relevance score as a random variable Y with mean μ and variance σ^2 . The variance of the sample mean \bar{Y} is given by $V(\bar{Y}) = \sigma^2/N$, where N is the sample size. If K strata are chosen such that each stratum has a different mean, $V(\bar{Y})$ can be decomposed into within-strata and between-strata variance (see Equation 1), the latter of which can be removed due to stratified

sampling [30],

$$V(\bar{Y}) = \sum_{k=1}^K \frac{n_k}{N} \sigma_k^2 + \sum_{k=1}^K \frac{n_k}{N} (\mu_k - \mu)^2 \geq \sum_{k=1}^K \frac{n_k}{N} \sigma_k^2 = V(\bar{Y}_{\text{strata}}), \quad (1)$$

where n_k , μ_k , σ_k^2 denote the sample size, mean, and variance of each stratum. Therefore, by using a stratified sampling design, we can achieve reduced variance and, as a result, reduced MDEs.

To determine the query strata, we evaluated multiple strata choices. While BERTopic [9] produced the largest variance reduction, we adopted the in-house query-to-interest model based on DistilBERT [23], due to ease of integration with existing systems and scalability.

3.4 Relevance Measurement with LLMs

To measure the relevance impact of an A/B experiment on search ranking, we take a stratified sample of paired search queries from control and treatment experiment groups, ensuring that the sample is representative of overall user usage [26]. The use of paired samples blocks between-query differences, an important source of variation in experiment measurement.

For each query in our paired sample, we retain the top K search results and generate LLM-based relevance labels. We then compute $sDCG@K$ for each query and aggregate query-level metrics to derive topline experiment metrics. The $sDCG@K$ metric is a variant of the standard $nDCG@K$, where we assume an infinite supply of highly relevant (L5) documents (see Equation 2). We use $K = 25$ throughout our evaluation.

$$sDCG@K = \frac{\sum_{k=1}^K L / \log_2(1+k)}{\sum_{k=1}^K 5 / \log_2(1+k)}, \quad L \in \{1, 2, 3, 4, 5\}. \quad (2)$$

Lastly, we calculate heterogeneous effects by query popularity and query interest (e.g. beauty, women’s fashion, art, etc), utilizing a Benjamini-Hochberg procedure [4] to control the false discovery rate. The LLM-based relevance measurement procedure at Pinterest Search is illustrated in Figure 2.

Table 1: LLM vs human labels alignment in query-level sDCG@K for different query popularity segments in US market relevance evaluation.

Query Popularity	Kendall's τ	Spearman's ρ	Query-level Error			
			Mean	P10	Median	P90
Overall	0.652	0.817	0.005	-0.047	0.002	0.059
Head	0.539	0.668	0.007	-0.020	0.001	0.043
Torso	0.600	0.764	0.008	-0.036	0.004	0.057
Tail	0.595	0.775	-0.002	-0.073	0.001	0.063
Single	0.620	0.801	0.004	-0.068	0.003	0.075

Table 2: LLM vs human labels alignment in query-level sDCG@K for different query popularity segments in France (FR) and Germany (DE) markets relevance evaluation.

Country	Query Popularity	Kendall's τ	Spearman's ρ	Query-level Error			
				Mean	P10	Median	P90
FR	Overall	0.483	0.618	0.021	-0.022	0.015	0.079
	Head	0.289	0.366	0.022	-0.009	0.014	0.074
	Torso	0.353	0.466	0.025	-0.011	0.017	0.078
	Tail	0.531	0.666	0.028	-0.007	0.017	0.071
	Single	0.572	0.743	0.015	-0.057	0.014	0.086
DE	Overall	0.470	0.611	0.020	-0.021	0.015	0.078
	Head	0.401	0.513	0.024	-0.010	0.017	0.071
	Torso	0.368	0.483	0.026	-0.011	0.019	0.078
	Tail	0.408	0.531	0.019	-0.026	0.013	0.079
	Single	0.493	0.678	0.009	-0.068	0.008	0.088

4 Results

We use XLM-RoBERTa_{large} as the LLM backbone for our relevance model. The model is lightweight yet delivers high-quality predictions. Inference runs on a single A10G GPU, allowing us to label 150,000 rows within 30 minutes. While the Llama-3-8B model offers slight improvement in accuracy, its inference time and cost increase by 6 times. Therefore, we pick XLM-RoBERTa_{large} as it offers a good balance between prediction quality and inference efficiency. Results are presented to answer the following three research questions (RQs):

- **RQ1:** Can LLM-based relevance assessment provide reliable metrics for online A/B experiment measurements?
- **RQ2:** Does LLM-based relevance assessment offer greater metric sensitivity compared to human labeling?
- **RQ3:** Is multilingual LLM-based relevance assessment effective for non-English queries as well?

4.1 RQ1: Alignment with Human Labels

We conducted a rigorous validation of the metrics derived from LLM labeling by comparing with human labels. The human labels were collected by instructing annotators to assess the relevance of each search query and Pin pair. The exact match rate between LLM and human labels is 73.7%, with 91.7% of ratings differing by no more than one point. These results underscore the high alignment between the relevance labels produced by LLMs and those from human annotators.

Following the established best practices in literature [1, 17] to measure alignment between LLMs and human labels, we compute and report the rank-based correlation Kendall's τ and Spearman's ρ to assess the correlation between the two rankings at query-level sDCG@K metric. To understand the performance on queries with different popularity, we also categorize the queries into 4 popularity segments based on search volume [12]: head, torso, tail, and single². The results are summarized in Table 1. For all query popularity segments, we achieve Kendall's $\tau > 0.5$ and Spearman's $\rho > 0.65$, indicating a strong alignment across all segments.

In addition to Kendall's τ and Spearman's ρ , we also validate the query-level sDCG@K error distribution. Here, the error refers to the difference between the sDCG@K metric derived from LLM labels and human labels. According to Table 1, the overall error is below 0.01, with the 10-th and 90-th percentiles falling within the range of [-0.1, 0.1]. We also visualize the error distribution in Figure 3. The error is tightly centered around 0, indicating its negligible magnitude and that the average bias will approach 0 as the size of the query set grows.

For experimental evaluation, we need to calculate the metric difference between the control and treatment groups. Therefore, we also validate how well these metric differences align in paired comparison. As shown on the right-hand side of Figure 3, the errors in paired differences are even more centered around 0 with lighter

²Single queries are defined as those that receive fewer than 10 searches.




Figure 3: Query-level $sDCG@K$ error distribution for single group (left) and paired differences (right) in US market relevance evaluation.

tails, indicating that LLM-based labeling provides highly reliable estimates of paired differences for A/B experiment assessment.³

4.2 RQ2: Metrics Sensitivity and MDEs

We evaluate the impact of these changes on experiment sensitivity by evaluating the MDEs for our experimentation system. The MDE is the smallest change in a metric that an experiment can reliably detect given the sample size, statistical power, and significance level chosen for the test. Since the typical experiment for most online platforms has a small effect, achieving small MDEs is a critical factor in team velocity and shipping new features to our users.

Before the introduction of LLM labeling, relevance measurement had large MDEs, typically ranging from 1.3% to 1.5%. These large MDEs were primarily the result of the constraints on our sampling designs imposed by the high cost and time consumption of human labeling. The introduction of LLM labeling enabled us to redesign our sampling approach. We increased our sample sizes, moved from simple random sampling (SRS) to stratified sampling, and now use a stratified sampling estimator [26]. Optimal allocation [19] is used to allocate sample units to strata, which are defined as the intersection of query interest and popularity segment. These changes enabled us to reduce our MDEs to $\leq 0.25\%$.

To attribute the MDE reduction to each source of change, we follow a standard derivation of MDE [5] as a function of the number of queries in the sample (n), metric mean ($\hat{\mu}$), and variance ($\hat{\sigma}^2$) assuming standard values for $\alpha = 0.05$ and $\beta = 0.8$,

$$\text{MDE} = \text{Lift\%Detectable} = \frac{(z_{1-\alpha/2} + z_{\beta}) \times \sqrt{2\hat{\sigma}^2/n}}{\hat{\mu}}. \quad (3)$$

We can therefore express MDE reduction in terms of reduction in $\hat{\sigma}$ and increased sample size. We present these results in Table 3. The vast majority of reduction comes from the variance reduction due to stratification. This is consistent with prior findings at Pinterest that most variance in relevance occurs across queries, largely due to differences in query interest and popularity.

³Note that, despite the low bias, we continue to use human-annotated relevance labels for evaluating online experiments that involve relevance ranking model changes in order to avoid potential bias.

Table 3: Improvement in metric sensitivity (MDE) with proposed sampling design and estimator.

Sample Type	Estimator	n	$\hat{\sigma}$	Reduction in $\hat{\sigma}$
SRS	SRS	2000	0.184	-
Stratified	Stratified	2000	0.011	94%
Stratified	SRS	5000	0.117	36%
Stratified	Stratified	5000	< 0.001	99.6%

4.3 RQ3: Performance on Non-English Queries

We fine-tuned multilingual LLMs on human-annotated data, with the majority of query-Pin pairs in English. As a result, careful validation is required for non-English queries to extend LLM-based relevance assessment to those queries. For this analysis, we focus on France (FR) and Germany (DE) markets.

The query-level metric alignment is summarized in Table 2. The overall Kendall’s τ and Spearman’s ρ are approximately 0.47 and 0.61, respectively. While these rank-based correlations are lower than those observed for English queries, they are still considered strong according to existing literature. The distribution of query-level metric errors is shown in Figure 4. Similar to the results of the US market, the errors are tightly concentrated around 0 for both countries, indicating a low average bias, with an even smaller bias for paired differences. These results provide confidence that the LLM-based relevance assessment is also suitable for non-English queries. Expanding relevance evaluation to countries beyond the US leads to further reductions in labeling costs and improvements in evaluation efficiency.




Figure 4: Query-level $sDCG@K$ error distribution for single group (left) and paired differences (right) in France (top) and Germany (bottom) markets relevance evaluation.

5 Conclusion and Future Work

In this work, we explore the use of LLM-based relevance labeling to generate query-level relevance metrics for online A/B experiments evaluation. We demonstrate that fine-tuned LLMs achieve low bias on query-level $sDCG@K$ metrics and paired differences. Transition to LLM-based relevance assessment enables us to scale up the evaluation query set and redesign the sampling strategy to improve the quality of relevance metrics for online experiment evaluation. We have successfully deployed the LLM-based relevance assessment at Pinterest Search, significantly reducing the manual annotation costs and turnaround time, while achieving an order of magnitude reduction in MDEs for improved detection of relevance shifts. Future work will explore using Visual Language Models (VLMs) to better leverage raw images for relevance prediction. Additionally, the observed performance gap with non-English queries highlights opportunities to further improve the multilingual capabilities of our LLM-based relevance model. We leave it for future work.

Acknowledgments

We would like to thank Kurchi Subhra Hazra for the support throughout this project; Mukuntha Narayanan and Jinfeng Rao for collaboration on developing the LLM-based relevance model; and Maggie Yang, Maria Alejandra Morales Gutierrez, Miguel Madera, Pedro Sanchez, Jorge Amigon, Francisco Navarrete for their assistance with LLM-labeling integration.

References

- [1] Zahra Abbasiantaeb, Chuan Meng, Leif Azzopardi, and Mohammad Aliannejadi. 2024. Can We Use Large Language Models to Fill Relevance Judgment Holes? *arXiv preprint arXiv:2405.05600* (2024).
- [2] AI@Meta. 2024. Llama 3 Model Card. (2024). https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md
- [3] Negar Arabzadeh and Charles LA Clarke. 2025. Benchmarking LLM-based relevance judgment methods. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3194–3204.
- [4] Yoav Benjamini and Yoel Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.
- [5] Howard S Bloom. 1995. Minimum detectable effects: A simple way to report the statistical power of experimental designs. *Evaluation Review* 19, 5 (1995), 547–556.
- [6] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116* (2019).
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [8] Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, 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*. 39–50.
- [9] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).
- [10] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594* (2024).
- [11] Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. *arXiv:2111.09543 [cs.CL]*
- [12] Alpa Jain and Gilad Mishne. 2010. Organizing query completions for web search. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. 1169–1178.
- [13] Vladimir Karpukhin, Barlas Ögüz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906* (2020).
- [14] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*. PMLR, 12888–12900.
- [15] Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. 2023. Fine-tuning llama for multi-stage text retrieval. *arXiv preprint arXiv:2310.08319* (2023).
- [16] 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*. 2230–2235.
- [17] Chuan Meng, Negar Arabzadeh, Arian Askari, Mohammad Aliannejadi, and Maarten de Rijke. 2025. Query performance prediction using relevance judgments generated by large language models. *ACM Transactions on Information Systems* 43, 4 (2025), 1–35.
- [18] Luke W Miratrix, Jasjeet S Sekhon, and Bin Yu. 2013. Adjusting treatment effect estimates by post-stratification in randomized experiments. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 75, 2 (2013), 369–396.
- [19] Jerzy Neyman. 1934. On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society* 97, 4 (1934), 158–606.
- [20] Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. 2020. Document ranking with a pretrained sequence-to-sequence model. *arXiv preprint arXiv:2003.06713* (2020).
- [21] Catarina Pires, Sérgio Nunes, and Luís Filipe Teixeira. 2025. Expanding Relevance Judgments for Medical Case-based Retrieval Task with Multimodal LLMs. *arXiv preprint arXiv:2506.17782* (2025).
- [22] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. <http://jmlr.org/papers/v21/20-074.html>
- [23] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [24] Shreya Shankar, JD Zamfirescu-Pereira, Björn Hartmann, Aditya Parameswaran, and Ian Arawjo. 2024. Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [25] 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*. 1930–1940.
- [26] Steven K Thompson. 2012. *Sampling*. John Wiley & Sons, Chapter 11, 141–170.
- [27] 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 preprint arXiv:2411.08275* (2024).
- [28] Shivani Upadhyay, Ronak Pradeep, Nandan Thakur, Nick Craswell, and Jimmy Lin. 2024. Umbrella: Umbrella is the (open-source reproduction of the) bing relevance assessor. *arXiv preprint arXiv:2406.06519* (2024).
- [29] Han Wang, Mukuntha Narayanan Sundararaman, Onur Gungor, Yu Xu, Krishna Kamath, Rakesh Chalasani, Kurchi Subhra Hazra, and Jinfeng Rao. 2024. Improving Pinterest Search Relevance Using Large Language Models. *arXiv preprint arXiv:2410.17152* (2024).
- [30] Huizhi Xie and Juliette Auriisset. 2016. Improving the sensitivity of online controlled experiments: Case studies at netflix. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 645–654.
- [31] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems* 36 (2023), 46595–46623.
- [32] Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and Michael Bendersky. 2023. Rankt5: Fine-tuning t5 for text ranking with ranking losses. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2308–2313.

Received 04 September 2025