

# JUDGERANK: LEVERAGING LARGE LANGUAGE MODELS FOR REASONING-INTENSIVE RERANKING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

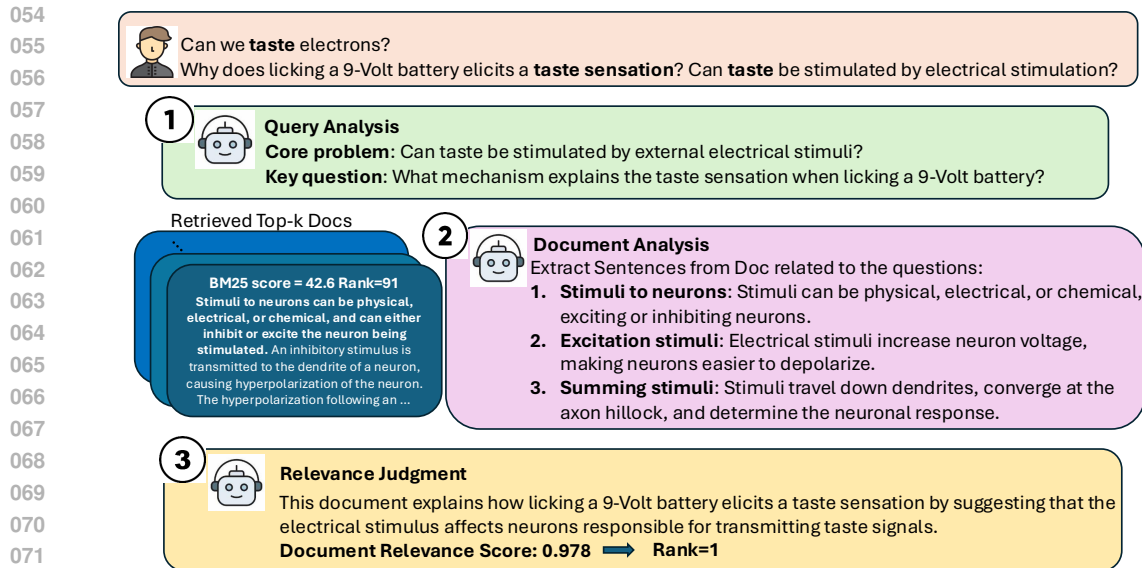
Accurate document retrieval is crucial for the success of retrieval-augmented generation (RAG) applications, including open-domain question answering and code completion. While large language models (LLMs) have been employed as dense encoders or listwise rerankers in RAG systems, they often struggle with reasoning-intensive tasks because they lack nuanced analysis when judging document relevance. To address this limitation, we introduce JUDGERANK,<sup>1</sup> a novel agentic reranker that emulates human cognitive processes when assessing document relevance. Our approach consists of three key steps: (1) *query analysis* to identify the core problem, (2) *document analysis* to extract a query-aware summary, and (3) *relevance judgment* to provide a concise assessment of document relevance. We evaluate JUDGERANK on the reasoning-intensive BRIGHT benchmark, demonstrating substantial performance improvements over first-stage retrieval methods and outperforming other popular reranking approaches. In addition, JUDGERANK performs on par with fine-tuned state-of-the-art rerankers on the popular BEIR benchmark, validating its zero-shot generalization capability. Through comprehensive ablation studies, we demonstrate that JUDGERANK’s performance generalizes well across LLMs of various sizes while ensembling them yields even more accurate reranking than individual models.

## 1 INTRODUCTION

Passage reranking is a critical component in modern information retrieval systems, designed to refine results obtained from efficient first-stage retrieval methods such as BM25 (Robertson et al., 1995; 2009). By narrowing down the pool of candidate documents, reranking substantially improves the quality of downstream tasks, such as retrieval-augmented generation or RAG (Lewis et al., 2020). Two primary approaches have emerged to address the reranking task. The first category comprises encoding-based approaches (Nogueira & Cho, 2019; Gao et al., 2021), which encode queries and documents into fixed-size embedding vectors. These methods use either cosine similarity as a score function or directly output a score from the model (Nogueira et al., 2020; Zhuang et al., 2023). While highly efficient, these approaches face several limitations. One major challenge is their inflexibility in defining relevance, making it difficult to accommodate diverse retrieval objectives (e.g., finding *supporting* vs. *refuting* evidence). Moreover, encoding-based models heavily rely on manual supervision signals due to the discrepancy between LLM pretraining and reranking objectives, limiting their ability to generalize to new domains or models (Nguyen, 2016; Izacard et al., 2022).

Most recently, utilizing Large Language Models (LLMs) for document reranking has led to promising progress in addressing some of these challenges, owing to their superior capabilities in language understanding, generation, interaction, and reasoning (Ouyang et al., 2022). These approaches utilize an LLM either as a pointwise judge (Ma et al., 2024) or a listwise reranker (Sun et al., 2023; Zhuang et al., 2024). While these approaches allow for flexible definition of document relevance and support zero-shot operation, they still require the model to make decisions without intermediate analyses. Consequently, they fall short in scenarios requiring complex reasoning (Su et al., 2024), hampering both performance and interpretability. Moreover, listwise rerankers face significant computational challenges due to context length constraints, often compromising on individual document length when processing multiple documents simultaneously.

<sup>1</sup>We plan to release our code upon acceptance.



073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083

Figure 1: A step-by-step illustration of how JUDGERANK arrives at the final judgment through query and document analyses. The query analysis identifies the core problem being asked, while the document analysis extracts relevant sentences from the document based on the query. This is a real example from the Biology task in the BRIGHT evaluation benchmark.

084  
085  
086  
087

To bridge this gap, we propose JUDGERANK, a novel zero-shot pointwise reranker tailored for reasoning-intensive text retrieval tasks. Inspired by Chain-of-Thought (Wei et al., 2022) and LLM-as-a-Judge (Zheng et al., 2023) methods, JUDGERANK utilizes highly generalizable prompts to guide instruction-tuned LLMs through explicit reasoning steps before arriving at a final judgment.

088  
089  
090  
091  
092  
093  
094  
095  
096  
097

Figure 1 illustrates a real example of how our model works on the Biology dataset in the BRIGHT (Benchmark for Reasoning-Intensive Generative Retrieval Tasks) benchmark (Su et al., 2024). Specifically, our reranker first prompts the LLM to identify the core problem in the query, allowing it to focus on the central question while filtering out irrelevant context. Next, the model produces an extractive summary for each of the candidate documents and explains how it addresses the query. Finally, the model makes a relevance judgment based on the previous analyses. This process closely mimics how humans approach questions: by first skimming the document, identifying relevant parts, and then carefully reading these parts to obtain an answer (Masson, 1983). This structured pipeline enables JUDGERANK to transcend surface-level lexical matching, leveraging deeper semantic understanding to improve reranking accuracy.

098  
099  
100  
101  
102  
103  
104  
105  
106  
107

We evaluate JUDGERANK on the recently constructed BRIGHT benchmark, widely regarded as one of the most challenging retrieval evaluation datasets. Despite the poor performance of state-of-the-art text embedding models and rerankers on this benchmark, our method achieves significant improvements over all existing baselines and secures the top position on the BRIGHT benchmark leaderboard among 89 models, surpassing the previous best model by a significant margin (9 points). Additionally, we demonstrate that JUDGERANK readily generalizes to other popular retrieval benchmarks such as BEIR and performs competitively with state-of-the-art rerankers. We also analyze the complementarity of models at different scales by investigating the alignment of their ranking decisions. Contrary to our expectations, we observe that models of different sizes demonstrate a surprisingly orthogonal behavior on their relevance judgements. Leveraging this key finding, we propose a simple ensembling strategy that allows us to combine multiple models flexibly and demonstrate that it translates to considerable performance gains on the final ranking.

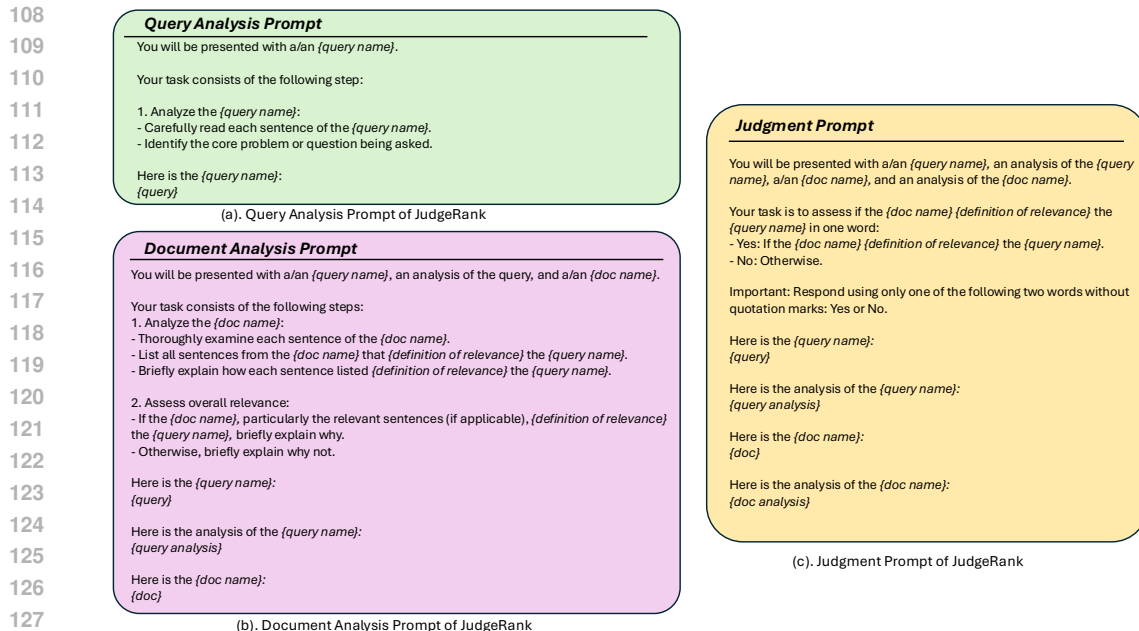


Figure 2: (a) Prompt to analyze query, where  $\{query\ name\}$  (e.g., “Biology post”) and  $\{query\}$  are placeholders for the query type and content. (b) Prompt for analyzing a document, where  $\{doc\ name\}$  (e.g., “document”) and  $\{doc\}$  are placeholders for the document type and content. (c) Prompt for making the final one-word relevance judgment.

## 2 METHOD

### 2.1 AGENTIC STEPS

Mimicing human cognitive process, JUDGERANK consists of three main steps: *Query Analysis*, *Document Analysis*, and *Relevance Judgment*. The prompt templates for these steps are illustrated in Figure 2.

**Query analysis** The query analysis prompt (Figure 2 (a).) directs the LLM to analyze the query by identifying the core problem being asked. Note that this prompt only depends on the query so that we can generate the query analyses separately and store them. Since the number of queries  $n_q$  is usually much smaller than that of documents  $n_d$  ( $n_q \ll n_d$ ), we can afford to use a more expensive LLM (e.g., GPT-4) to handle this important step, and leave the other steps to relatively smaller LMs.

**Document analysis** The document analysis prompt (Figure 2 (b).) asks the LLM to output an extractive summary of the document that helps answer the query, and assess the overall relevance of the document based on the summary.

**Relevance judgment** The judgment prompt (Figure 2 (c).) asks the model to make a one-word judgment, either “Yes” or “No”. We isolate this step to make it easier to ensemble with different judgment prompts or models.

### 2.2 GENERALIZABILITY OF THE PROMPTS

All three steps consume natural language as input and generate a response, making it more flexible to transfer and stack across different LLMs. The templates also show that the prompts are highly generalizable: to adapt them to a new reranking task, one only needs to replace the query name, the document name, and the relation between them. Leveraging LLMs in a zero-shot setting allows us to flexibly define the relation between the query and the document. This flexibility is important because the user may define either “document that supports a query” or “document that refutes a query” as

the relation, which are opposites of each other. Encoding-based models usually cannot achieve such behavior zero-shot because most of them use cosine similarity or metrics alike to represent “relevance”. One way encoding-based models could achieve such flexibility is through extensive fine-tuning. However, this requires additional training data and introduces new model parameters, potentially causing an unintended distribution shift. Similarly, models like RANKZEPHYR also finetunes the LLM to output document ids, thus suffering from a distribution shift as well.

### 2.3 METHODOLOGY OF SCORING DOCUMENTS

**Discrete version** The discrete version creates a binary partition between accepted (when the model outputs a “Yes”) and rejected (when the model outputs a “No”) documents, maintaining the first-stage retrieval ranking within each category. More specifically, let  $D = \{d_1, d_2, \dots, d_k\}$  be an ordered list of top- $k$  documents ranked by the first-stage retrieval model. Let  $D_y$  and  $D_n$  be a partition of  $D$ , such that

$$D_y \cup D_n = D \quad (1)$$

$$D_y \cap D_n = \emptyset \quad (2)$$

, where  $D_y$  is the set of documents that the reranker judged as relevant and  $D_n$  is the set of documents that the reranker judged as non-relevant. Let  $R$  be the reranking function which maps each document  $d$  to its rank (lower rank means “more relevant”), then

$$\forall d \in D_y \text{ and } d' \in D_n, R(d) < R(d') \quad (3)$$

, and for the relative ranking within each partition,

$$\forall d_i, d_j \in D_y, R(d_i) < R(d_j) \iff R_0(d_i) < R_0(d_j) \quad (4)$$

, where  $R(d)$  is the final rank of document  $d$  and  $R_0(d)$  is the rank of  $d$  from the first-stage retrieval. The same applies to  $D_n$ .

While straightforward, this approach is sensitive to prompt wording and relies heavily on first-stage retrieval performance.

**Continuous version** The continuous version addresses the limitations of the discrete version by using the probability of the “Yes” judgment  $p_y$  and the probability of the “No” judgment  $p_n$  to construct a complete ranking. This is similar to the relevance generation approach proposed in Liang et al. (2023). The scoring function  $S$  is defined by normalizing the probabilities between  $p_y$  and  $p_n$  as follows:

$$S(d) = \frac{p_y}{p_y + p_n} \quad (5)$$

This normalization step is necessary because the combined probability mass allocated to  $p_y$  and  $p_n$  is not always a fixed value across different documents. Without normalization, the  $p_y$  values for different documents would not be directly comparable.

The final ranking  $D_R$  is then defined as  $D_R = \{d_1, d_2, \dots, d_k\}$  such that

$$\forall i, j \in \{1, 2, \dots, k\}, i < j \iff S(d_i) > S(d_j) \quad (6)$$

This continuous version provides a more fine-grained ranking compared to the binary partition, as it utilizes the full range of probabilities output by the LLM.

**Hybrid version** Additionally, we explore a variant of the continuous version where the final score is computed by taking a weighted sum of the probability score  $S_{\text{prob}}$  and the BM25 score  $S_{\text{BM25}}$ . More specifically, the final score is computed by

$$S = \alpha S_{\text{prob}} + S_{\text{BM25}} \quad (7)$$

, where  $S_{\text{BM25}}$  is the score provided by BM25 in the first-stage retrieval, and  $\alpha$  is the relative weight of the probability score. We set  $\alpha = 100$  in this work to bring  $S_{\text{prob}}$  to the same scale as  $S_{\text{BM25}}$ . This version leverages model ensembling to consider both reasoning and surface-level matching, thus marrying the benefits of both approaches. Unless otherwise specified, we use this setting to compute the final score throughout the paper. In the ablation studies, we compare these three settings and show their relative performances.

### 3 EXPERIMENTAL SETUP

#### 3.1 DATASETS

**BRIGHT** We use the BRIGHT benchmark (Su et al., 2024) to assess the performance of our reranker. BRIGHT is specifically designed to evaluate text retrieval systems on complex, reasoning-intensive queries that go beyond simple keyword matching. The benchmark comprises 1,398 real-world queries spanning diverse domains, including economics, psychology, robotics, math, and software engineering. These queries are carefully curated to represent challenging scenarios that require deep understanding and reasoning to identify relevant documents. We use this dataset to evaluate our approach because unlike traditional benchmarks that focus on simple information-seeking tasks, BRIGHT queries require complex reasoning to determine document relevance, making it an excellent tool for evaluating advanced retrieval systems in realistic scenarios. The benchmark has also been validated to be robust against potential data leakage, maintaining its effectiveness even when benchmark documents have been included in model training data.

Because of its challenging nature, state-of-the-art retrieval models have shown significantly lower performance on BRIGHT compared to other benchmarks (Su et al., 2024). For example, the leading model on the MTEB leaderboard (Muennighoff et al., 2022) achieves an nDCG@10 of only 18.0 on BRIGHT, compared to 59.0 on other benchmarks. The GPT-4 listwise reranker also only improves around 2 points on nDCG@10 on top of the BM25 first-stage retrieval, while Gemini (Team et al., 2023) features less improvement than that. The cross-encoder reranker MiniLM (Wang et al., 2020) even significantly underperforms the BM25 baseline.

**BEIR** To test the generalizability of our approach, we also evaluate on the BEIR benchmark (Thakur et al., 2021), a robust and heterogeneous evaluation benchmark for information retrieval. We evaluate on all tasks in this benchmark that are publicly available (Kamalloo et al., 2023). For all datasets we use the test set, except for MSMARCO where we follow BEIR convention to evaluate on the dev set. For *cqadupstack* we follow BEIR convention and evaluate on all sub-datasets and compute their average. Because BEIR is a large benchmark, and the largest dataset has more than 13K queries, we only generate query analysis to evaluate on this dataset. This almost adds no overhead to the generation because the query analysis generation does not depend on the document. Since our model only generates a single token “Yes” or “No”, its latency is almost the same as encoding both the query and the document with an encoding-based retrieval model, making it a highly efficient alternative.

#### 3.2 FIRST-STAGE RETRIEVAL

For both benchmarks we evaluate on, we follow common settings from previous work to rerank the top-100 documents from the first-stage retrieval and use nDCG@10 score as the evaluation metric. This metric assesses the quality of the retrieved documents, taking into account both their relevance and ranking position.

**BRIGHT** The original BRIGHT paper explores using LLMs to generate Chain-of-Thought (Wei et al., 2022) reasoning steps as queries (Su et al., 2024), resulting in up to 12.2 point improvements

on average. We thus build on top of this best first-retrieval model on the leaderboard, which achieves an nDCG@10 score of 26.5 with BM25 and reasoning chains generated by GPT-4-0125-preview.

**BEIR** We follow the original BEIR paper (Thakur et al., 2021) and use ElasticSearch BM25 (elasticsearch, 2015) as the first-stage retriever.

### 3.3 BASE MODEL

Our main model builds on top of Llama-3.1-70B-instruct (AI@Meta, 2024). We choose the Llama-3.1 model family because its ROPE scaling Liu et al. (2024b) allows longer context length up to 128K, which is essential in handling long documents. To speed up experiments, we evaluate on a quantized version of this model, namely Llama-3.1-70B-instruct-awq-int4.<sup>2</sup> We also perform ablation studies where we evaluate on Llama-3.1-8B<sup>3</sup> and Llama-3.1-405B-instruct-awq-int4.<sup>4</sup> Note that for the 8B model we do not use the quantized version because it can already fit on a single A100 GPU, while the other two bigger models require quantization to save computational cost. More specifically, the 70B version requires at least 2 x A100 GPUs, while the 405B version requires 8.

### 3.4 BASELINE RERANKERS

We reproduce RANKLLAMA (Ma et al., 2024) and RANKZEPHYR (Pradeep et al., 2023), two state-of-the-art rerankers as evaluated by the BEIR benchmark. RankLlama is a pointwise reranker that directly outputs a score. This model is trained on the MS MARCO passage ranking dataset (Bajaj et al., 2016). RANKZEPHYR is a listwise reranker that takes a query and a list of documents together as input and outputs a ranking. This model uses the queries sourced by Sun et al. (2023) from the MS MARCO dataset to distill GPT-3.5 and GPT-4 in sequence. We use the RERANKERS library (Clavié, 2024), a lightweight unified API that allows users to run diverse reranking models out-of-the-box.<sup>5</sup>

### 3.5 EFFICIENCY AND OPTIMIZATION

To make the encoding and generation more efficient, we use vLLM (Kwon et al., 2023). This library leverages paged attention, which improves the throughput of popular LLMs by 2 - 4x with the same level of latency. Importantly, when designing the prompts used in our approach, we append the query and the document at the very end of each prompt to make the best use of Automatic Prefix Caching (Gim et al., 2024), which temporarily stores the KV cache of existing inputs so that a new input can directly reuse the KV cache if it shares the same prefix with one of the existing ones. This allows the new inputs to skip the computation of the shared part. This design greatly improves the efficiency of our experiments.

## 4 RESULTS AND DISCUSSION

### 4.1 MAIN RESULTS

**BRIGHT** As shown in Table 1, JUDGERANK achieves state-of-the-art results on the BRIGHT evaluation benchmark as measured by nDCG@10. Our best performing model improves upon the no-rerank baseline by more than 9 points, while RANKLLAMA underperforms the baseline and RANKZEPHYR stays barely above the baseline. The smaller Llama-3.1-8B-instruct also outperforms the baseline by more than 3 points, showing the generalizability of our approach across different model sizes. Interestingly, increasing model size from 70B to 405B does not bring a significant gain on nDCG@10. We thus select the 70B version as our main model to balance between efficiency and performance. In contrast, the original BRIGHT paper shows that GPT-4 with listwise reranking

<sup>2</sup><https://huggingface.co/hugging-quants/Meta-Llama-3.1-70B-Instruct-AWQ-INT4>

<sup>3</sup><https://huggingface.co/meta-llama/Llama-3.1-8B>

<sup>4</sup><https://huggingface.co/hugging-quants/Meta-Llama-3.1-405B-Instruct-AWQ-INT4>

<sup>5</sup><https://github.com/AnswerDotAI/rerankers>

	BM25	RankLlama	RankZephyr	JudgeRank			
				8B	70B	405B	Ensemble
Biology	53.63	11.05	44.37	54.44	57.59	60.33	<b>60.70</b>
Earth Science	53.65	11.83	35.24	54.62	58.36	55.11	<b>58.72</b>
Economics	24.28	9.56	24.44	28.04	33.01	32.15	<b>35.39</b>
Psychology	38.59	11.38	36.92	42.16	46.59	45.42	<b>47.57</b>
Robotics	18.77	8.53	18.80	24.09	<b>28.30</b>	27.64	28.16
Stack Overflow	22.74	11.40	19.62	27.18	27.47	28.30	<b>29.74</b>
Sustainable Living	25.90	11.75	29.38	30.69	39.55	38.54	<b>41.88</b>
Leetcode	19.27	20.32	<b>24.58</b>	17.49	20.06	22.47	20.23
Pony	17.73	18.88	<b>48.96</b>	22.85	30.82	31.54	32.74
Aops	3.92	3.55	6.98	6.15	8.24	7.74	<b>8.57</b>
TheoremQA-Questions	18.90	11.82	22.34	24.09	23.83	<b>26.93</b>	25.86
TheoremQA-Theorems	20.22	5.63	7.78	29.89	34.55	36.16	<b>36.20</b>
Average	26.47	11.31	27.25	30.14	34.03	34.36	<b>35.48</b>

Table 1: JUDGERANK nDCG@10 results on the BRIGHT evaluation benchmark. Best results on each dataset and the entire benchmark are boldfaced. “Ensemble” stands for model ensembling of JUDGERANK-8B, 70B, and 405B.

	BM25	JudgeRank	RankLlama	RankZephyr
webis-touche2020	<b>34.71</b>	27.94	32.97	33.34
trec-covid	68.80	83.70	81.08	<b>85.77</b>
scifact	69.06	73.24	75.57	<b>74.68</b>
nfcopus	34.28	37.73	35.92	<b>38.79</b>
dbpedia-entity	32.02	<b>44.30</b>	43.72	44.19
fiqa	25.36	40.35	<b>42.70</b>	40.40
scidocs	16.47	19.47	19.26	<b>19.70</b>
arguana	47.16	<b>62.77</b>	32.08	47.07
nq	32.61	56.87	59.19	<b>60.10</b>
climate-fever	18.61	19.00	17.75	<b>24.26</b>
fever	64.94	69.97	78.55	<b>79.52</b>
msmarco	22.75	34.01	<b>41.40</b>	37.89
hotpotqa	60.22	68.30	70.65	<b>71.16</b>
quora	80.77	<b>84.25</b>	82.73	79.60
cqadupstack	32.53	<b>43.98</b>	42.43	42.68
Average	42.69	51.06	50.40	<b>51.94</b>

Table 2: JUDGERANK nDCG@10 results on the BEIR evaluation benchmark. Best results on each dataset and the entire benchmark are boldfaced.

improves on top of vanilla BM25 baseline by an average of 2.7 points on nDCG@10, a much smaller improvement than our approach despite using a much stronger LLM.

**BEIR** As shown in Table 2, our model delivers competitive results on the BEIR evaluation benchmark despite the fact that RANKLLAMA and RANKZEPHYR are heavily fine-tuned on in-domain data including MS MARCO, which is part of the BEIR benchmark.

## 4.2 DISCUSSION

We pose several research questions to illustrate whether and how our approach works.

**How complementary are LLMs of different scales?** In Table 1, we observe that JudgeRank-70B and JudgeRank-405B performs on par with each other. However, nDCG@10 alone does not reveal the whole picture. One natural question to ask is: do these two models make similar judgments or are complementary to each other? To answer this question, we obtain statistics on the percentage of both models agreeing and disagreeing each other and show them on the left of Figure 3. From the tables we can see that for all three combinations of the models, the majority case is always that both models rejects the documents. This is understandable because only a few out of the top-

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431




Figure 3: **On the left:** judgment alignment studies for models of three sizes: 8B, 70B, and 405B. Percentages are shown for each quadrant. **On the right:** nDCG@10 of each individual model and model ensembling on the BRIGHT evaluation benchmark.

100 documents are supposed to be relevant. The interesting pattern emerges when we inspect the other three cases: each pair of the models spends more time disagreeing with each other than both outputting “Yes”. For the pairs 8B vs 70B and 8B vs 405B, there is a higher difference because the capabilities of the two models differ more. In contrast, for 70B vs 405B there is less disagreement. From these observations, we indeed see that each two models may be complementary to each other.

Motivated by this observation, we take model ensembling one step further. Recall that in Section 2.3, we ensemble the BM25 score with each of the scores output by the Llama models. Here we first take the average score output by all the Llama models, and then perform the weighted sum with the BM25 score. More specifically, let  $S_{8B}$ ,  $S_{70B}$ , and  $S_{405B}$  be the score assigned by each model, respectively, the ensemble score of the three models is computed as  $\alpha(S_{8B} + S_{70B} + S_{405B})/3 + S_{BM25}$ , where again  $\alpha = 100$  and  $S_{BM25}$  is the score given by the BM25 model. The same equation generalizes analogously to two-model ensembles.

We present all ensembling results on the right of Figure 3. We can see that each ensembling performance is better than its individual model performances, with the strongest performance observed when ensembling all three models. This result shows that a salient performance boost can be achieved by ensembling two of the strongest models (70B + 405B), while even the model with lower performance (i.e., 8B) could contribute positively in model ensembling. Intuitively, such ensembling is equivalent to a verification or a majority voting step. The final score is the highest when both models say “Yes”, the score is medium when one of the two says “No”, and the lowest score is observed when both say “No”.

**How does the choice of reranking score impact the final performance?** Recall that to compute the final score, we take a weighted sum of the BM25 score and the probability score from the judgment step. To justify this choice, we compare it with two other settings: the first is binary judgment, and the second only uses the normalized probability to rerank documents (introduced in Section 3). The left part of Figure 4 shows that binary judgment performs the worst among the three settings while using only probability achieves somewhere in between. This is understandable because binary judgments are sensitive to wordings. Imagine that if we change the relation from “substantially helps answer” to “helps answer” or “at least partially helps answer,” the number of “Yes” that the model outputs will keep increasing, thus also increasing the number of false positives. However, the other two settings are not sensitive to such changes.

**How useful are the query and document analysis steps?** To show the effectiveness of the analysis steps, we perform an ablation study on BRIGHT. We remove the two analyses steps and keep the judgment step untouched to compare its performance with the original approach. The right of Figure 4 shows that judging with query and document analyses performs consistently better than the direct judgment approach.






Figure 4: Ablation studies of JudgeRank. **On the left:** Comparison of three scoring settings on the BRIGHT evaluation benchmark. *Binary* stands for binary judgment, *Prob* stands for probability, and *Hybrid* stands for a weighted sum of BM25 and probability scores. **On the right:** Comparison of direct judge and judge with query and document analyses on the BRIGHT evaluation benchmark.

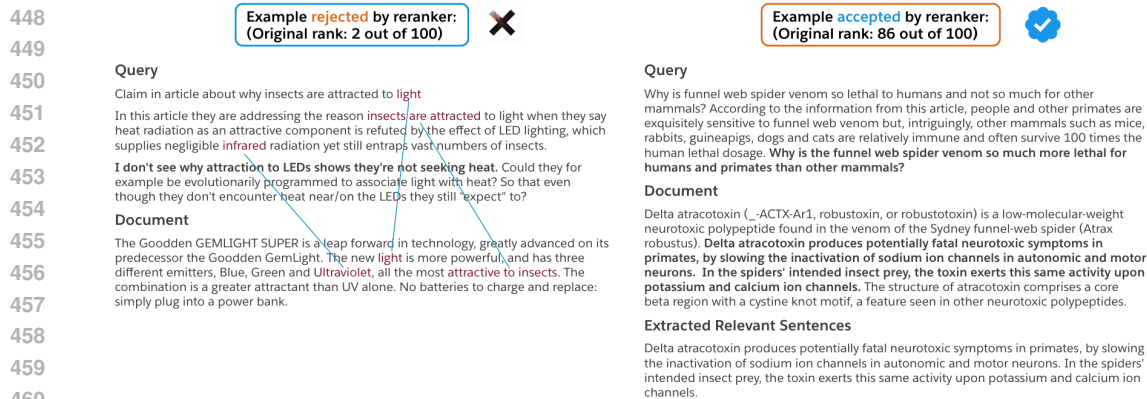


Figure 5: Illustration of how agentic generations of JUDGE RANK help identifying the relevant documents. **On the left,** the document is ranked high by the first-stage retriever but judged as negative by the reranker. **On the right,** the document is ranked low by the first-stage retriever but judged as positive by the reranker because the document analysis prompt helps the LLM to locate the relevant sentences that answer the query.

**Qualitative examples** Figure 5 demonstrates how JUDGE RANK enhances document relevance identification using real examples from the BRIGHT dataset. In the left example, we observe a document initially ranked high by the first-stage retriever due to significant word overlap with the query. However, JUDGE RANK’s deeper analysis correctly judges this document as irrelevant. Despite surface-level similarities, the reranker fails to extract any sentences that help answer the query, revealing that the document is merely an advertisement coincidentally sharing common terminology with the query. The right example presents a contrasting scenario, where a document initially ranked low by the first-stage retriever due to minimal word overlap with the query is accurately identified by JUDGE RANK as highly relevant. In this instance, the document analysis prompt enables the LLM to pinpoint key sentences that elucidate the underlying mechanism of funnel web spider venom’s lethality, precisely addressing the query’s intent. These extracted sentences further inform the LLM to make the final positive judgment, demonstrating JudgeRank’s ability to uncover deeply relevant content that might be overlooked by traditional retrieval methods.

## 5 RELATED WORK

The field of reranking models can be understood through two primary dimensions.

The first dimension distinguishes between encoding-based and LLM-based approaches. Encoding-based models (Nogueira & Cho, 2019; Gao et al., 2021), which have been the standard since the

introduction of BERT, typically require extensive training to adapt to the specific objectives of reranking tasks. In contrast, LLM-based models (Sun et al., 2023; Tang et al., 2024; Qin et al., 2024; Zhuang et al., 2024), particularly those with decoder-only architectures, often demonstrate impressive zero-shot capabilities, allowing them to perform effectively without task-specific fine-tuning. However, it’s worth noting that some researchers have explored pretraining, fine-tuning and distillation techniques for LLM-based rerankers to further enhance their performance (Zhang et al., 2023; Ma et al., 2024; Yu et al., 2024).

The second dimension differentiates between pointwise and listwise reranking strategies. Pointwise rerankers (Ma et al., 2024; Guo et al., 2024) evaluate the relevance of individual query-document pairs in isolation, producing a score for each pair without direct comparison between documents. Listwise rerankers (Sun et al., 2023; Pradeep et al., 2023; Ma et al., 2023; Yoon et al., 2024; Liu et al., 2024a; Tang et al., 2024), on the other hand, consider the entire set of documents for a given query, generating a fully ordered list as output. To address the challenges posed by input length limitations, researchers have developed innovative techniques such as sliding window approaches (Sun et al., 2023), which allow for the ranking of smaller subsets of documents before aggregating them into a comprehensive ranking. This framework encompasses pairwise (Nogueira et al., 2019; Pradeep et al., 2021; Qin et al., 2024) and setwise (Zhuang et al., 2024) rerankers as specific instances of the broader listwise category.

Our contribution is a novel LLM-based, pointwise reranker that leverages the capabilities of instruction-tuned language models and incorporates explicit reasoning steps in the relevance judgment process. This approach sets our work apart from previous efforts by enhancing the model’s ability to handle complex, reasoning-intensive reranking tasks while simultaneously improving the interpretability of its decisions.

## 6 CONCLUSION

In this work, we target document retrieval tasks that require intensive context-based reasoning, which even the strongest retrieval models struggle to achieve satisfactory performance. Through experiments and ablation studies, we show that our agentic reranker can effectively recover low-ranked documents and outperform previous state-of-the-art reranking models, while remaining flexible and efficient. In section 4.2 we have shown the significant benefit of model ensembling on document reranking. Yet we do not have to stop there. There are at least two categories of ensembling that we envision. First, sampling ensembling: For each generation, we sample several generations, each of which could lead to a different judgment. This kind of ensembling is similar to the self-consistency approach (Tang et al., 2024). Second, prompt ensembling: we could leverage paraphrases of the same prompt to perform ensembling. We leave the exploration as future work because such approaches could be generalized to many prompting-related tasks, and thus better be addressed in separate works.

## REFERENCES

- AI@Meta. Llama 3 model card. 2024. URL [https://github.com/meta-llama/llama3/blob/main/MODEL\\_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md).
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.
- Benjamin Clavié. rerankers: A lightweight python library to unify ranking methods, 2024. URL <https://arxiv.org/abs/2408.17344>.
- elasticsearch. elasticsearch/elasticsearch, 2015. URL <https://github.com/elasticsearch/elasticsearch>.
- Luyu Gao, Zhuyun Dai, and Jamie Callan. Rethink training of bert rerankers in multi-stage retrieval pipeline. In *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 – April 1, 2021, Proceedings, Part II*, pp. 280–286, Berlin, Heidelberg, 2021. Springer-Verlag. ISBN 978-3-030-72239-5. doi: 10.1007/978-3-030-72240-1\_26. URL [https://doi.org/10.1007/978-3-030-72240-1\\_26](https://doi.org/10.1007/978-3-030-72240-1_26).

- 540 In Gim, Guojun Chen, Seung-seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. Prompt  
541 cache: Modular attention reuse for low-latency inference. In P. Gibbons, G. Pekhimenko,  
542 and C. De Sa (eds.), *Proceedings of Machine Learning and Systems*, volume 6, pp. 325–  
543 338, 2024. URL [https://proceedings.mlsys.org/paper\\_files/paper/2024/  
544 file/a66caal703fe34705a4368c3014c1966-Paper-Conference.pdf](https://proceedings.mlsys.org/paper_files/paper/2024/file/a66caal703fe34705a4368c3014c1966-Paper-Conference.pdf).
- 545 Fang Guo, Wenyu Li, Honglei Zhuang, Yun Luo, Yafu Li, Le Yan, and Yue Zhang. Generating  
546 diverse criteria on-the-fly to improve point-wise llm rankers. *arXiv preprint arXiv:2404.11960*,  
547 2024.
- 548
- 549 Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand  
550 Joulin, and Edouard Grave. Unsupervised dense information retrieval with contrastive learn-  
551 ing. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL [https:  
552 //openreview.net/forum?id=jKN1pXi7b0](https://openreview.net/forum?id=jKN1pXi7b0).
- 553
- 554 Ehsan Kamalloo, Nandan Thakur, Carlos Lassance, Xueguang Ma, Jheng-Hong Yang, and Jimmy  
555 Lin. Resources for brewing beir: Reproducible reference models and an official leaderboard,  
556 2023.
- 557
- 558 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph  
559 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model  
560 serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Prin-  
561 ciples*, pp. 611–626, 2023.
- 562
- 563 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
564 Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe  
565 Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the  
566 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook,  
NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- 567
- 568 Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga,  
569 Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan,  
570 Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re,  
571 Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda  
572 Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng,  
573 Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khat-  
574 tab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar,  
575 Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William  
576 Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. Holistic evaluation of lan-  
577 guage models. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL  
578 <https://openreview.net/forum?id=iO4LZibEqW>. Featured Certification, Expert  
579 Certification.
- 580
- 581 Qi Liu, Bo Wang, Nan Wang, and Jiaxin Mao. Leveraging passage embeddings for efficient listwise  
582 reranking with large language models. *arXiv preprint arXiv:2406.14848*, 2024a.
- 583
- 584 Xiaoran Liu, Hang Yan, Chenxin An, Xipeng Qiu, and Dahua Lin. Scaling laws of roPE-based  
585 extrapolation. In *The Twelfth International Conference on Learning Representations*, 2024b.  
586 URL <https://openreview.net/forum?id=JO7k0SJ5V6>.
- 587
- 588 Xueguang Ma, Xinyu Zhang, Ronak Pradeep, and Jimmy Lin. Zero-shot listwise document rerank-  
589 ing with a large language model. *arXiv preprint arXiv:2305.02156*, 2023.
- 590
- 591 Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. Fine-tuning llama for multi-stage  
592 text retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and  
593 Development in Information Retrieval, SIGIR '24*, pp. 2421–2425, New York, NY, USA, 2024.  
Association for Computing Machinery. ISBN 9798400704314. doi: 10.1145/3626772.3657951.  
URL <https://doi.org/10.1145/3626772.3657951>.
- 594
- 595 Michael EJ Masson. Conceptual processing of text during skimming and rapid sequential reading.  
*Memory & cognition*, 11(3):262–274, 1983.

- 594 Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embed-  
595 ding benchmark. *arXiv preprint arXiv:2210.07316*, 2022. doi: 10.48550/ARXIV.2210.07316.  
596 URL <https://arxiv.org/abs/2210.07316>.  
597
- 598 T Nguyen. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint*  
599 *arXiv:1611.09268*, 2016.
- 600 Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with bert. *arXiv preprint*  
601 *arXiv:1901.04085*, 2019.  
602
- 603 Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. Multi-stage document ranking with  
604 bert. *arXiv preprint arXiv:1910.14424*, 2019.
- 605 Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. Document ranking with a pre-  
606 trained sequence-to-sequence model. In Trevor Cohn, Yulan He, and Yang Liu (eds.), *Findings*  
607 *of the Association for Computational Linguistics: EMNLP 2020*, pp. 708–718, Online, November  
608 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.63.  
609 URL <https://aclanthology.org/2020.findings-emnlp.63>.  
610
- 611 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
612 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kel-  
613 ton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike,  
614 and Ryan Lowe. Training language models to follow instructions with human feedback. In  
615 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in*  
616 *Neural Information Processing Systems*, volume 35, pp. 27730–27744. Curran Associates, Inc.,  
617 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf)  
618 [file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf).
- 619 Ronak Pradeep, Rodrigo Nogueira, and Jimmy Lin. The expando-mono-duo design pattern for text  
620 ranking with pretrained sequence-to-sequence models. *arXiv preprint arXiv:2101.05667*, 2021.
- 621 Ronak Pradeep, Sahel Sharifmoghammad, and Jimmy Lin. RankZephyr: Effective and robust zero-  
622 shot listwise reranking is a breeze! *arXiv:2312.02724*, 2023.  
623
- 624 Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi  
625 Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. Large language mod-  
626 els are effective text rankers with pairwise ranking prompting. In Kevin Duh, Helena Gomez,  
627 and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL*  
628 *2024*, pp. 1504–1518, Mexico City, Mexico, June 2024. Association for Computational Lin-  
629 guistics. doi: 10.18653/v1/2024.findings-naacl.97. URL [https://aclanthology.org/](https://aclanthology.org/2024.findings-naacl.97)  
630 [2024.findings-naacl.97](https://aclanthology.org/2024.findings-naacl.97).
- 631 Stephen Robertson, Hugo Zaragoza, et al. The probabilistic relevance framework: Bm25 and be-  
632 yond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389, 2009.  
633
- 634 Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford,  
635 et al. Okapi at trec-3. *Nist Special Publication Sp*, 109:109, 1995.
- 636 Hongjin Su, Howard Yen, Mengzhou Xia, Weijia Shi, Niklas Muennighoff, Han-yu Wang, Haisu  
637 Liu, Quan Shi, Zachary S Siegel, Michael Tang, et al. Bright: A realistic and challenging bench-  
638 mark for reasoning-intensive retrieval. *arXiv preprint arXiv:2407.12883*, 2024.  
639
- 640 Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin,  
641 and Zhaochun Ren. Is ChatGPT good at search? investigating large language models as re-ranking  
642 agents. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference*  
643 *on Empirical Methods in Natural Language Processing*, pp. 14918–14937, Singapore, December  
644 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.923. URL  
645 <https://aclanthology.org/2023.emnlp-main.923>.
- 646 Raphael Tang, Crystina Zhang, Xueguang Ma, Jimmy Lin, and Ferhan Ture. Found in the mid-  
647 dle: Permutation self-consistency improves listwise ranking in large language models. In  
Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of*

- 648 *the North American Chapter of the Association for Computational Linguistics: Human Lan-*  
649 *guage Technologies (Volume 1: Long Papers)*, pp. 2327–2340, Mexico City, Mexico, June  
650 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.129. URL  
651 <https://aclanthology.org/2024.naacl-long.129>.
- 652 Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu,  
653 Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly  
654 capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- 655 Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych.  
656 Beir: A heterogeneous benchmark for zero-shot evaluation of information retrieval mod-  
657 els. In J. Vanschoren and S. Yeung (eds.), *Proceedings of the Neural Information Pro-*  
658 *cessing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL [https://](https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/65b9eea6e1cc6bb9f0cd2a47751a186f-Paper-round2.pdf)  
659 [datasets-benchmarks-proceedings.neurips.cc/paper\\_files/paper/](https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/65b9eea6e1cc6bb9f0cd2a47751a186f-Paper-round2.pdf)  
660 [2021/file/65b9eea6e1cc6bb9f0cd2a47751a186f-Paper-round2.pdf](https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/65b9eea6e1cc6bb9f0cd2a47751a186f-Paper-round2.pdf).
- 661 Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-  
662 attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neu-*  
663 *ral Information Processing Systems*, 33:5776–5788, 2020.
- 664 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi,  
665 Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language  
666 models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Ad-*  
667 *vances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/](https://openreview.net/forum?id=_VjQlMeSB_J)  
668 [forum?id=\\_VjQlMeSB\\_J](https://openreview.net/forum?id=_VjQlMeSB_J).
- 669 Soyoung Yoon, Eunbi Choi, Jiyeon Kim, Hyeongu Yun, Yireun Kim, and Seung-won Hwang.  
670 ListT5: Listwise reranking with fusion-in-decoder improves zero-shot retrieval. In Lun-Wei  
671 Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting*  
672 *of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2287–2308,  
673 Bangkok, Thailand, August 2024. Association for Computational Linguistics. URL [https://](https://aclanthology.org/2024.acl-long.125)  
674 [aclanthology.org/2024.acl-long.125](https://aclanthology.org/2024.acl-long.125).
- 675 Yue Yu, Wei Ping, Zihan Liu, Boxin Wang, Jiakuan You, Chao Zhang, Mohammad Shoeybi, and  
676 Bryan Catanzaro. Rankrag: Unifying context ranking with retrieval-augmented generation in  
677 llms. *arXiv preprint arXiv:2407.02485*, 2024.
- 678 Longhui Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, Meishan Zhang, and Min Zhang.  
679 Rankingpt: Empowering large language models in text ranking with progressive enhancement.  
680 *arXiv preprint arXiv:2311.16720*, 2023.
- 681 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
682 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.  
683 Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on*  
684 *Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL [https://](https://openreview.net/forum?id=uccHPGDlao)  
685 [openreview.net/forum?id=uccHPGDlao](https://openreview.net/forum?id=uccHPGDlao).
- 686 Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and  
687 Michael Bendersky. Rankt5: Fine-tuning t5 for text ranking with ranking losses. In *Proceedings*  
688 *of the 46th International ACM SIGIR Conference on Research and Development in Information*  
689 *Retrieval*, pp. 2308–2313, 2023.
- 690 Shengyao Zhuang, Honglei Zhuang, Bevan Koopman, and Guido Zuccon. A setwise approach  
691 for effective and highly efficient zero-shot ranking with large language models. In *Proceedings*  
692 *of the 47th International ACM SIGIR Conference on Research and Development in Information*  
693 *Retrieval*, pp. 38–47, 2024.
- 694  
695  
696  
697  
698  
699  
700  
701