

Why Large Language Models can Secretly Outperform Embedding Similarity in Information Retrieval

Matei Benescu
University of Groningen
matei.benescu22@gmail.com

Ivo Pascal de Jong
University of Groningen
ivo.de.jong@rug.nl

Abstract

With the emergence of Large Language Models (LLMs), new methods in Information Retrieval are available in which relevance is estimated directly through language understanding and reasoning, instead of embedding similarity. We argue that similarity is a *short-sighted* interpretation of relevance, and that LLM-Based Relevance Judgment Systems (LLM-RJS) (with reasoning) have potential to outperform Neural Embedding Retrieval Systems (NERS) by overcoming this limitation. Using the TREC-DL 2019 passage retrieval dataset, we compare various LLM-RJS with NERS, but observe no noticeable improvement. Subsequently, we analyze the impact of reasoning by comparing LLM-RJS with and without reasoning. We find that human annotations also suffer from short-sightedness, and that false-positives in the reasoning LLM-RJS are primarily mistakes in annotations due to short-sightedness. We conclude that LLM-RJS do have the ability to address the short-sightedness limitation in NERS, but that this cannot be evaluated with standard annotated relevance datasets.

1 Introduction

Information Retrieval (IR) is the task of finding relevant information in a pool of numerous documents, where a document can be any retrievable unit like a webpage or text fragment. IR methods have constantly been evolving (Kazi et al., 2025). Vector Space Models (Salton, 1965; Salton et al., 1975) use sparse vector representations of term weights, where each weight represents the frequency of the term in the given document. Documents are ordered by relevance to the query, by the means of an ordering function, such as the BM25 (Robertson et al., 2009, Chapter 3.4). Modern IR methods rely on dense vector representations of text, known as “embeddings”, which are learned by Neural Networks (Huang et al., 2013). We will

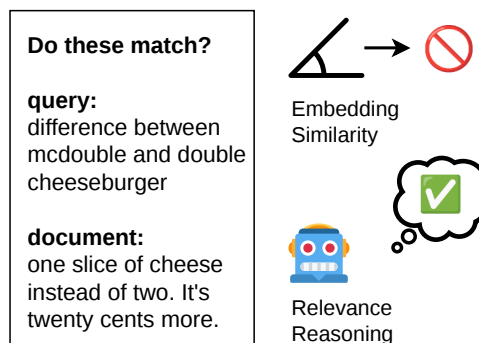


Figure 1: Motivation for LLM-RJS. The query and document are not similar, so embedding similarity is too short-sighted for IR. The document is the correct answer, but requires reasoning about relevance to identify as a match.

refer to these methods as Neural Embeddings Retrieval Systems (NERS). In NERS, documents are ordered by their similarity with the query, which is given by a distance metric such as Euclidean Distance, or Cosine Similarity (Wang and Dong, 2020). Karpukhin et al., 2020 showed that NERS outperform sparse vector representation approaches.

More recently, Large Language Models (LLMs) have been proposed as a different approach to judging how relevant a document is to a query (MacAvaney and Soldaini, 2023; Faggioli et al., 2023; Abbasiantaeb et al., 2024). Alaofi et al., 2024 showed that LLMs’ relevance judgments are similar to those of humans. This has the potential to outperform NERS models, as NERS models are limited by retrieving documents that are *similar* to the query. Similarity tests whether the semantics of a document and query match, but dissimilar contents may still be relevant with some extra thought. An example of NERS being short-sighted is shown in Figure 1.

For the scope of this paper, we will call a system that uses an LLM’s generated relevance score

for matching an *LLM-Based Relevance Judgment System* (LLM-RJS). While LLM-RJS have been shown to be effective (Alaofi et al., 2024), to the best of our knowledge a 1-to-1 comparison against NERS is missing. Additionally, IR applications of reasoning models that are equipped with Chain of Thought (CoT) capabilities, such as GPT-o1 (Jaech et al., 2024) remain unexplored. As reasoning capabilities become cheaper and available in open models (Guo et al., 2025), examining its impact on document relevance judgments becomes increasingly important and broadly applicable.

With this, we make the following contributions:

- We evaluate NERS and multiple LLM-RJS models on the TREC-DL 2019 passage retrieval benchmark to compare their performances in relevance ranking. We see that top LLMs perform on par with NERS, but do not yield demonstrably better rankings.
- We evaluate the impact of “reasoning effort” for LLM-RJS at directly estimating the relevance scores. We propose that this can help to overcome the short-sightedness of NERS. We find that adding reasoning results in higher relevance ratings than the human relevance scores, and that the LLM-RJS overall shows higher relevance ratings.
- We explain the difference in relevance ratings as short-sightedness in the ground-truth annotations, showing examples of documents that have low relevance ratings in the human annotation, but that are highly relevant. From the 94 false positives, 89 are actually mistakes in the annotation because the passages are relevant but not similar.
- We conclude that the empirical capabilities of zero-shot methods for Information Retrieval may be underestimated due to short-sighted annotations, making short-sighted similarity models appear better than more capable reasoning models. LLM-RJS have the ability to overcome the short-sighted limitation of NERS, but this does not lead to better metrics when evaluated on human annotations.

2 Related Work

We describe previous research on NERS and LLM-RJS models. While NERS stem from a long line of existing research in text embeddings (Mikolov et al., 2013; Pennington et al., 2014), we specifically focus on embeddings from Neural Networks.

NERS Neural Embedding Retrieval Systems (NERS) rely on pre-trained language models to create a dense vector representation of text. These language models are Neural Networks trained using objectives such as Masked Language Modeling (Devlin et al., 2019) or Contrastive Learning (Gao et al., 2021). The penultimate layer of the Neural Network therefore has a dense representation of the semantic meaning of the text, capturing properties like synonymy (Karpukhin et al., 2020).

NERS models work by computing the embedding representation of the query and the document, and evaluating the distance between them, typically through cosine similarity. When a document and a query describe similar concepts they get a higher similarity score. Document relevance is therefore ranked by similarity to the query.

We consider this a fundamental limitation of IR systems that rely on embedding similarities. While similarity between query and document is a good heuristic for relevance, it is not exactly the same. Similarity scores assess relevance based on whether the semantics match, but sometimes it is possible to link the document and the query with extra thought, as the example in Figure 1. It is not similar, but with some extra thought can still be highly relevant. Even with ever-improving embedding models (Trablsi et al., 2021), this necessarily gives a limitation of NERS.

LLM-RJS Large Language Model-Based Relevance Judgement Systems (LLM-RJS) instead use the generated answers from an LLM to establish relevance. In this work, the LLM is prompted with query and document, and instructed to quantify how “relevant” the document is in line with the human annotations, though many variations for prompting exist (Arabzadeh and Clarke, 2025).

Faggioli et al., 2023 found that LLMs can perform comparably to humans in identifying irrelevant information but struggled to consistently recognize relevant pieces of information. In their results GPT-3.5 showed a higher false negative rate than humans.

Evaluating the goodness of document relevance assigned by LLMs can be directly compared to human-labeled query-document pairs are used as the “ground truth”. However, humans are not consistent in rating document relevance (Bailey et al., 2008). To evaluate LLMs fairly, Alaofi et al., 2024 posed the issue as “if LLMs produce labels that agree with humans to the same extent that humans

Query	Passage	Rel.
axon terminals or synaptic knob definition	At the end of the axon’s branches and form synapses with the dendrites of other neurons.	0
definition declaratory judgment	That is especially true in situations where a non-practicing entity has sent a licensing offer to a potential infringer. Cases finding declaratory judgment jurisdiction. A patent owner wrote to its licensee alleging that a product sold by the licensee was subject to a license agreement. The licensee paid the requested royalties under protest and filed a declaratory judgment suit.	1
difference between a mcdouble and a double cheeseburger	Although we previously reported that the McDouble would be on the Dollar Menu, in Manhattan the cheeseburger appears on the Dollar Menu while the McDouble is \$1.39 and the double cheeseburger is \$1.49.	2
do goldfish grow	A: The conditions goldfish are kept in plus their diet determine how large they will grow. I have seen goldfish grow ridiculously large in very small containers when their water was changed frequently. Goldfish will not grow if water conditions are poor. Fancy goldfish don’t grow as large as Common goldfish. A good size would be around 5 inches body length for most fancy varieties, 8 inches for Comets and 12 inches for Common Goldfish. These sizes are usually only attained by pond grown fish.	3

Table 1: TREC-DL 2019 examples. Relevance: 0 = Irrelevant, 1 = Related, 2 = Highly relevant, 3 = Perfectly relevant.

agree with each other (human-to-human range), we can conclude that they are sufficiently reliable for use”. Their results show that larger LLMs, such as GPT-4, GPT-4o, and LLaMA70B give results in the human-to-human range. Only GPT-4 and GPT-4o were robust to the prompt variations. However, this perspective cannot show when LLMs might outperform humans in the ability to estimate relevance.

Because LLM-RJS systems are zero-shot models and are not trained as supervised learning systems on the human annotations, we should consider that it has the potential to outperform human annotations. Additionally, with reasoning models (Wei et al., 2022), we can generate reasoning processes that show relevance in a way that has more depth than simple similarity scores, and might have more depth than the human labels.

3 Dataset

TREC (Text REtrieval Conference) datasets are commonly used in the IR field (MacAvaney and Soldaini, 2023; Faggioli et al., 2023; Alaofi et al., 2024). Similarly to Alaofi et al., 2024, we use queries and passages from the test set of the passage retrieval task dataset of the Deep Learning Track of TREC 2019 (TREC-DL 2019) (Voorhees et al., 2020). TREC-DL 2019 contains triples of form (query, passage, relevance score).

The passage corpora of TREC-DL 2019 are derived from the MS MARCO dataset. MS MARCO is a large-scale dataset of real web search queries and web text originally sampled from Bing search logs. Queries in MS MARCO are real user queries

and passages are extracted from web documents obtained via Bing search results. Table 1 gives some examples of (query, passage, relevance) triples.

The relevance scores of TREC-DL 2019 are evaluated by human assessors. The scores (Voorhees et al., 2020, Section 2.2) are integer values using a 4-point scale (0-3) where:

3. **Perfectly relevant:** The passage is dedicated to the query and contains the exact answer.
2. **Highly relevant:** The passage has some answer for the query, but the answer may be a bit unclear, or hidden amongst extraneous information.
1. **Related:** The passage seems related to the query but does not answer it.
0. **Irrelevant:** The passage has nothing to do with the query.

For both experiments, the LLMs are instructed to only output the relevance score of the (query, document) pair. Any response that could not be cast to an integer between 0 and 3 was considered a failed response. All others are successes. For each model, the (query, document) pairs were divided into successes and failures. The intersection of all successful pairs was kept for final analysis.

4 Experiment 1

To evaluate whether LLM-RJS can outperform NERS, and therefore attempt to address the short-sightedness issue of NERS, we compare the performance of NERS with LLM-RJS in document ranking. As a baseline we evaluate the BM25 model. We first describe the implementation details of the

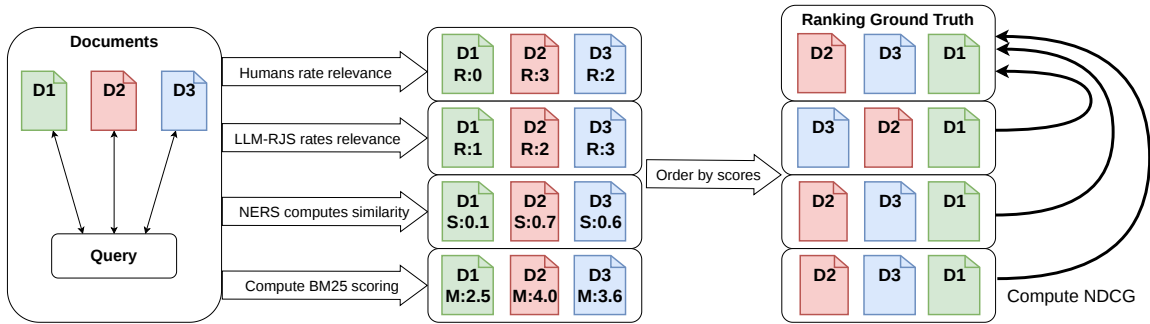


Figure 2: Overview of Experiment 1. Documents are ranked (ordered) based on the “relevance score” with respect to the query assigned by each system. The rankings are compared to the ranking yielded by humans labelers, via NDCG. The mean NDCG is used as final evaluation.

Model	Provider	Open Weights
GPT-5.1	OpenAI	No
Sonnet 4.5	Anthropic	No
Gemini 2.5 Pro	Google	No
gpt-oss:20b	OpenAI	Yes
gpt-oss:120b	OpenAI	Yes
llama3.2:1b	Meta	Yes
llama3.2:3b	Meta	Yes
llama3.1:8b	Meta	Yes

Table 2: LLMs used in Experiment 1 (OpenAI, 2025; Anthropic, 2025; Google, 2025; Ollama, 2024; Meta, 2024b,a).

three models, and then how they are evaluated. The overall setup of Experiment 1 is demonstrated in Figure 2.

LLM-RJS All models were asked to estimate the relevance of (query, passage) pairs on a 4-point scale pair, using the prompt found in Appendix A. The prompt describes the task and provides the relevance score convention. The models were also given examples of (query, document, relevance) triples, as examples are known to enhance LLM performance (Brown et al., 2020). We use both closed state-of-the-art (at the moment of writing), as well as open-weights models. Table 2 provides an overview of the LLM models evaluated. The temperature hyperparameter (Ackley et al., 1985; Peepkorn et al., 2024), controls the amount of “randomness” in LLM token generation. The temperature was set to 0 for all models, to ensure reproducibility and because higher temperatures have been shown to have a moderate correlation with incoherence (Peepkorn et al., 2024). The reasoning efforts of GPT-5.1 was set to “none”. The open-weights models were run using Ollama¹.

¹<https://ollama.com/>

Neural Embedding Retrieval Systems: For NERS, we used a closed model and an open-weights one. The closed model, gemini-embedding001, is the state-of-the-art from Google, showed to outperform most of the current embedding models (Lee et al., 2025). The open model is nomic-embed-text:v1.5². It is an extension of nomic-embed-text:v1 (Nussbaum et al., 2024) that uses Matryoshka Representation Learning (Kusupati et al., 2022). The base model, nomic-embed-text:v1, showed better performance (Nussbaum et al., 2024) than similarly sized embedding models and established models like GloVe (Pennington et al., 2014).

BM25 baseline: BM25 is a probabilistic scoring function that ranks documents based on query term frequencies while balancing the effects of term frequency and document length.

The BM25 scoring function is defined as:

$$\text{BM25}(d, q) = \sum_{t \in q} \text{IDF}(t) \frac{f_{t,d}^{(k_1+1)}}{f_{t,d} + k_1 \left(1 - b + b \frac{|d|}{\text{avgdl}}\right)}$$

where q is the query, $f_{t,d}$ is the frequency of term t in document d , k_1 controls the degree of term-frequency saturation, $|d|$ is the length of document d , avgdl is the average document length in the collection, and $b \in [0, 1]$ controls the strength of document length normalization.

The inverse document frequency (IDF) reduces the weight of common terms while increasing the importance of rarer terms across the collection:

$$\text{IDF}(t) = \log \frac{N - n_t + 0.5}{n_t + 0.5}$$

where N is the total number of documents and n_t is the number of documents containing term t .

²<https://huggingface.co/nomic-ai/nomic-embed-text-v1.5>

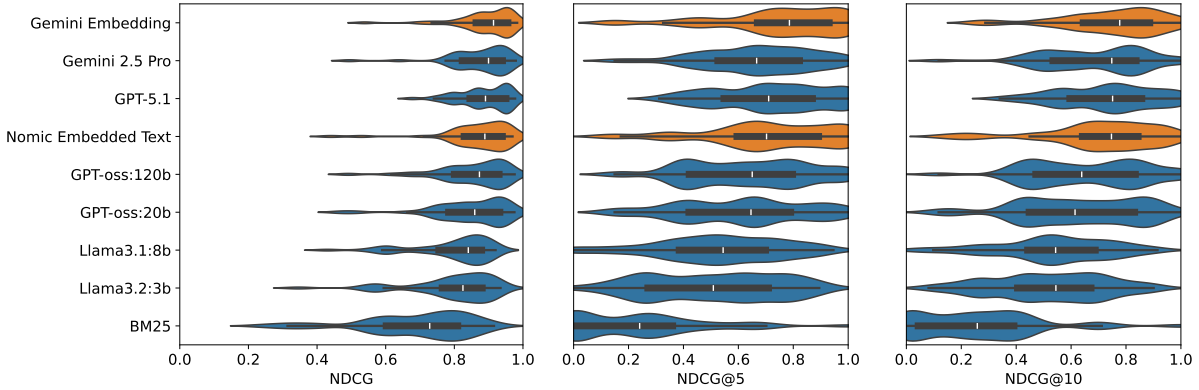


Figure 3: Performance of **LLM-RJS** and **NERS** models in relevance ranking, measured by NDCG@k. Top LLM-RJS models perform on par with NERS models, but do not appear to outperform them. Results are in line with previous research from Hofstätter et al. (2021); Xing et al. (2023), reporting mean NDCG@10 of 0.72 and 0.73 respectively.

Similarly to Voorhees et al., 2020, BM25 with the default parameters ($k_1 = 1.5$, $b = 0.75$) was chosen as a baseline ranking system. The rankings yielded by both LLM-RJS and NERS will be compared to the one of BM25.

Model evaluation: The relevance estimates from NERS and BM25 cannot be directly compared to ground truth relevance scores, due to different representations of relevance. Instead, the document rankings, i.e. the set of documents ordered by their relevance to the query, yielded by the different systems can be directly compared.

The Normalized Discounted Cumulative Gain (NDCG) (Wang et al., 2013) is a measure of ranking quality, which we will use to compare the rankings obtained by different retrieval systems.

NDCG is defined as:

$$\text{NDCG}(A) = \frac{\text{DCG}(A)}{\text{DCG}(I)},$$

$$\text{DCG}(R) = \sum_{i=1}^N \frac{h(d_i)}{\log(i+1)},$$

where I is the ideal ranking and $h(d_i)$ is the relevance of document d at position i in the ranking. NDCG ranges from 0 to 1, with 1 corresponding to the ideal ranking induced by human relevance judgments.

NDCG@k is NDCG evaluated for the sub-ranking that consists of the first k documents in the initial ranking.

Since each query yields a ranking, and the dataset consists of multiple queries, the average NDCG and NDCG@k with $k = 5$ and $k = 10$

over queries will be computed as final evaluation metrics.

4.1 Results and Discussion

For each query, all the passages in relation with it have been ranked, and the NDCG, NDCG@5 and NDCG@10 has been computed for the rank yielded by each model. Figure 3 shows the results of Experiment 1 as a violin plot with NDCG, NDCG@5 and NDCG@10 of all models.

The NERS and LLM-RJS models performed on par with previous models (hofstatter2021efficiently, xing2023query). These previous models are purposely trained dense retrievers, unlike the models analyzed by this paper, which are zero-shot. The presented NDCG@10 of these previous models was reported in the respective papers and not implemented here.

For this experiment, only Llama3.1:8b, Llama3.2:1b and Sonnet 4.5 had any failures, out of which, Llama3.2:1bs and Sonnet 4.5 have been excluded for having failed more than 2/3 of all data points (lack of instruction compliance, see Appendix A. Table 5 provides an example of failure for each model.). The models failed to only output one integer between 0 and 3. Expectedly, one of the excluded models is Llama3.2:1b, the smallest model in the selection, with 8312/9260 failed responses. Surprisingly, the other one was Sonnet 4.5, the state-of-the-art (at the time of writing) of Anthropic, which failed to only give relevance scores for 7473/9260 data points. Llama3.1:8b only failed 84/9260 data points.

Overall, NERS are still most proficient for ranking documents, and perform the task better for only

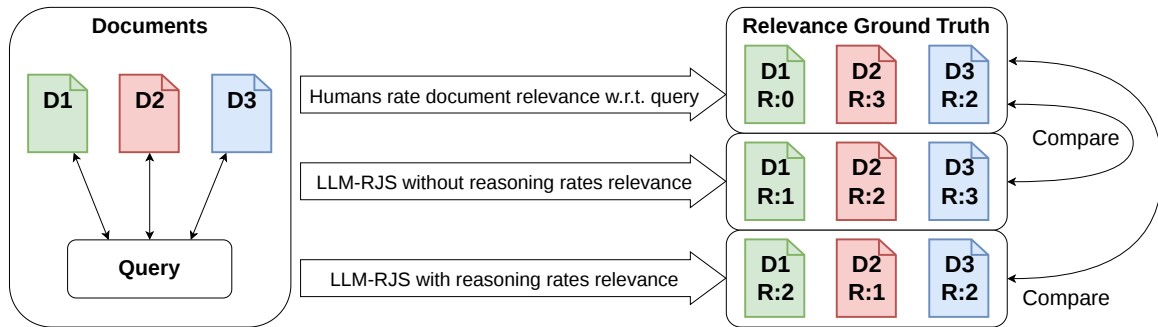


Figure 4: Overview of Experiment 2. The LLM estimates relevance of (query, document) pairs in “high reasoning” and “low reasoning” conditions. The mean absolute error between LLM and human assessment in the two conditions is used as evaluation metric.

a fraction of the cost, and a fraction of the time. The price per 1 million tokens strongly favors NERS, with Gemini Embeddings costing 0.15\$ per 1 million tokens. On the other side, GPT-5.1 costs 1.25\$ per 1 million input tokens and 10\$ per 1 million output tokens. Albeit financially undesirable, LLMs do not leave much on the table, when it comes to ranking (query, document) pairs.

The proprietary models consistently outperform the open models. The bottom line is, if one is looking for the best option, Gemini Embeddings is the winning model. If, aside performance, running the retrieval system locally is also a consideration, nomic-embed-text:v1.5 is a viable option.

5 Experiment 2

Chain of Thought (CoT) reasoning (Wei et al., 2022) has been shown to yield substantial improvement in LLM capabilities on a variety of tasks. It works by letting the LLM generate a longer series of tokens as “reasoning” before giving the final answer. This strategy has been shown to lead to better results in image Information Retrieval (Wu et al., 2025), and is therefore also promising for having LLM-RJS outperform NERS models.

We prompt GPT-5.1 to estimate the relevance of (document, query) pairs with the prompt described in Appendix A. To compare the reasoning against the non-reasoning conditions we change the “reasoning effort”³, which controls the amount of reasoning tokens before the output. A visual overview of the experiment is shown in Figure 4.

The same experiment was also performed with gpt-oss:20b, a smaller open-weights model, with

³<https://platform.openai.com/docs/guides/reasoning#get-started-with-reasoning>

similar results reported in Appendix C, suggesting that the findings are consistent.

Since the LLM-RJS can directly estimate prompt relevance under the same definitions as the human annotations, the relevance estimates are compared directly, instead of by the ranking. Figure 5 shows the confusion matrices of the non-reasoning and reasoning LLM-RJS relevance rating compared to human relevance ratings.

We observe reasonable agreement between the LLM-RJS systems and the human annotations, but find that in many cases the LLM-RJS assigns higher relevance ratings than the humans annotations. This is particularly visible in the large number of cases where the human rating has relevance 0 (irrelevant), while the LLM-RJS has relevance 1 (related). This difference increases for the reasoning model, which identifies yet more (document, query) pairs as related. These higher ratings are also reflected in the average relevance ratings, where the average human rating is 0.79, LLM-RJS without reasoning is 1.03, and with reasoning is 1.23.

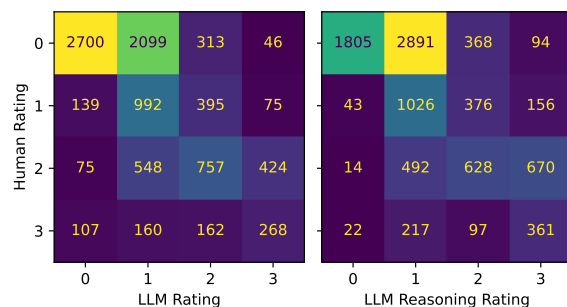


Figure 5: Differences between human and LLM relevance ratings with and without reasoning for GPT 5.1. The LLMs overall are more likely to see relevance in documents, especially with reasoning.

Query	Passage	Human	Reasoning	No Reasoning	NERS
axon terminals or synaptic knob definition	The small bulges at the end of the axon that send messages to other neurons are called	0	3	2	0.48
causes of left ventricular hypertrophy	Heart hypertrophy is a condition where the muscular wall of the heart becomes thickened. [...] It can thicken as a result of high blood pressure [...]	0	3	2	0.52
cost of interior concrete flooring	Polished concrete pricing and cost per square metre for concrete polishing can range from as little as £20 to as much as £150 per metre square [...]	0	3	2	0.46
definition declaratory judgment	n. a judgment of a court which determines the rights of parties without ordering anything be done or awarding damages [...]	0	3	3	0.53
difference between a mcdouble and a double cheeseburger	One slice of cheese instead of two. It's twenty cents more to get a double cheeseburger for \$1.19. [...]	0	3	3	0.43
difference between rn and bsn	Once you graduate, you have your BSN, but to become an RN, you have to take state boards and pass them before you are an RN. [...]	0	3	3	0.57
do goldfish grow	average ones are four to six inches (10 -15 cm) long, but they can grow up to 7 inches [...]	0	3	2	0.35
example of monotonic function	[...] Since then, this proportion has grown steadily and monotonically with each successive election to 43% in 1996, 54% in 2000, 58% in 2004 [...]	0	3	2	0.14
exons definition biology	A portion of DNA that codes for a section of the mature messenger RNA obtained from that DNA [...]	0	3	2	0.55
how long is life cycle of flea	[...] Adults live about 7 to 10 days and deposit eggs on the moist soil surface or in soil cracks [...]	0	3	2	0.38

Table 3: First 10 query-IDs where GPT5.1 gave rating 3 while the humans gave 0. The annotations suffer from short-sightedness, while GPT-5.1 with reasoning correctly identifies relevance. Of the 94 such cases only 5 were genuine false positives (shown in Appendix B). The NERS similarity averages 0.44 for these 94 cases, while it is 0.50 for true positives (higher is more similar). **Bold** highlights key examples of the short-sightedness limitation.

Particularly interesting is that the reasoning configuration, which would be expected to perform better, has double the incidences where human ratings identify a pair as irrelevant (0), while the LLM-RJS identifies the pair as perfectly relevant.

Table 3 shows examples from these cases. Predominantly, the passages are actually highly relevant, but might not be immediately obvious. The “difference between a mcdouble and a double cheeseburger” is indeed one slice of cheese instead of two. This may not be obvious to human annotators or short-sighted NERS models that evaluate similarity, but in these cases the reasoning LLM-RJS outperforms the benchmark. Overall the results from Experiment 2 show that the addition of reasoning allows the model to identify relevance that is not obvious, beyond the human annotations.

6 Conclusion

We proposed that the short-sightedness problem of NERS may be resolved with relevance reasoning LLM-RJS. Similarity is not always sufficient for relevance. In Experiment 1, we evaluated whether these LLM-RJS outperform NERS on a ranking task, and do not find improvement. In Experiment 2 we added reasoning, and directly compared the LLM-RJS relevance ratings to the human ratings. This showed the LLM-RJS was able to identify more documents as relevant than the human annotations. This suggests that the human annotations suffer from similar short-sightedness as the NERS systems, explaining why the LLM-RJS do not reach a noticeable improvement compared to NERS when evaluated on NDCG with short-sighted human annotations as ground truth.

System trade-offs The primary downside to LLM-RJS systems is the added cost compared to

embedding models. The price-per-token of top performing LLMs is roughly 10x that of equivalent embedding models. Additionally, for NERS systems implemented at scale, each document only needs to be embedded once even for multiple queries, and all document embeddings can be stored in a vector database for efficient retrieval. The LLM-RJS system demonstrated here requires LLM inference for each (query, document) pair and therefore does not scale well. A solution may be found in CoT re-ranking (Wu et al., 2025), where an embedding model is used to make a first selection of documents that are similar, and then the reasoning-LLM selects the relevant documents from the subset of similar documents. However, this only addresses the problem where similar documents are not relevant, but would still fail in cases as the example in Figure 1, where the document is relevant but not similar.

Future work The presented results suggest, but do not definitively show that LLM-RJS models can outperform human annotations. To establish the performance difference between human relevance and LLM-RJS relevance re-annotation may be done where trained annotators compare existing labels and LLM-RJS relevance judgements and establish which is correct. This can be supported by additionally providing the reasoning chain from the LLM-RJS, to provide argumentation for the judgment.

We have established in this instance of an Information Retrieval a mismatch between similarity and relevance, and show this as a short-sightedness-bottleneck for embeddings. However, for a variety of different Information Retrieval tasks, a variety of retrieval requirements exist beyond “relevance”. In tasks further conditions apply such as preferring high-quality documents, while in other tasks specifically similarity might be preferred over relevance (e.g. for detecting copyright infringement). With advanced Information Retrieval capabilities of LLM-RJS models, soft requirements can be considered, and there are more possibilities to deviate from ‘similarity’ or requirements that can be quantified. Evaluating LLM-RJS against human annotations still appears to be the standard (Arabzadeh and Clarke, 2025), and efforts are needed to re-specify what the (sub)tasks of Information Retrieval problems are, and how LLM-RJS models can be evaluated.

In conclusion, we found that the short-

sightedness motivation presented in the illustrative Figure 1 is a valid limitation of NERS models that can be addressed by LLM-RJS. However, as shown with the examples of short-sighted annotations in Table 3, human annotations suffer from the same short-sightedness limitation, they present as false-positives for a confusion matrix in Figure 5. Therefore, under a standard annotated benchmark evaluation as shown in Figure 3, they do not demonstrably outperform NERS, while they still could be giving better relevance judgments.

7 Limitations

We have introduced short-sightedness as a fundamental problem of embedding similarity and demonstrate through examples that it exists. However, the concept of short-sightedness lacks a formal definition, making it difficult to identify exactly which false-negatives are due to short-sightedness and which are not. Not all examples in Table 3 are necessarily due to the short-sightedness problem, and there may be short-sighted examples that the LLM-RJS also fails to identify as relevant. We therefore cannot quantify the scale of this problem, and further theoretical analysis of short-sightedness is needed.

We should consider that the short-sightedness in human annotations is not universal for all datasets. We have shown cases of short-sightedness in this dataset, but expert annotators or more effortful annotation may not suffer from the same short-sightedness problem. Since LLM-RJS do not need large training datasets, efforts can be focused on only creating test sets.

Declaration on Generative AI: For this paper, OpenAi’s GPT-5, GPT-5.1 and GPT-5.2 models were used for minor reformulations and spelling checking, and LaTeX formatting.

References

- 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*.
- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. 1985. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169.
- Marwah Alaofi, Paul Thomas, Falk Scholer, and Mark Sanderson. 2024. Llms can be fooled into labelling a document as relevant: best café near me; this paper is

- perfectly relevant. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 32–41.
- Anthropic. 2025. Claude Sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5>. Accessed: 2026-01.
- 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*, pages 3194–3204.
- Peter Bailey, Nick Craswell, Ian Soboroff, Paul Thomas, Arjen P. de Vries, and Emine Yilmaz. 2008. **Relevance assessment: are judges exchangeable and does it matter**. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, page 667–674, New York, NY, USA. Association for Computing Machinery.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and 1 others. 2023. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 ACM SIGIR international conference on theory of information retrieval*, pages 39–50.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Google. 2025. Gemini 2.5 Pro. <https://docs.cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-5-pro>. Accessed: 2026-01.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently teaching an effective dense retriever with balanced topic aware sampling. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 113–122.
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 2333–2338.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP (1)*, pages 6769–6781.
- Samreen Kazi, Shakeel Khoja, and Ali Daud. 2025. Bridging the gap: A survey of document retrieval techniques for high-resource and low-resource languages. *Computer Science Review*, 57:100756.
- Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, and Ali Farhadi. 2022. **Matyoshka representation learning**. In *Advances in Neural Information Processing Systems*, volume 35, pages 30233–30249. Curran Associates, Inc.
- Jinhyuk Lee, Feiyang Chen, Sahil Dua, Daniel Cer, Madhuri Shanbhogue, Iftekhar Naim, Gustavo Hernández Ábrego, Zhe Li, Kaifeng Chen, Henrique Schechter Vera, and 1 others. 2025. Gemini embedding: Generalizable embeddings from gemini. *arXiv preprint arXiv:2503.07891*.
- 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*, pages 2230–2235.
- Meta. 2024a. LLaMA 3.1 Models. <https://ollama.com/library/llama3.1>. Accessed: 2026-01.
- Meta. 2024b. LLaMA 3.2 Models. <https://ollama.com/library/llama3.2>. Accessed: 2026-01.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

- Zach Nussbaum, John X Morris, Brandon Duderstadt, and Andriy Mulyar. 2024. Nomic embed: Training a reproducible long context text embedder. *arXiv preprint arXiv:2402.01613*.
- Ollama. 2024. GPT-OSS Models. <https://ollama.com/library/gpt-oss>. Accessed: 2026-01.
- OpenAI. 2025. GPT-5.1 Model. <https://openai.com/index/gpt-5-1/>. Accessed: 2026-01.
- Max Peeperkorn, Tom Kouwenhoven, Dan Brown, and Anna Jordanous. 2024. Is temperature the creativity parameter of large language models? *arXiv preprint arXiv:2405.00492*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Stephen Robertson, Hugo Zaragoza, and 1 others. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Gerard Salton. 1965. The evaluation of automatic retrieval procedures—selected test results using the smart system. *American Documentation*, 16(3):209–222.
- Gerard Salton, Anita Wong, and Chung-Shu Yang. 1975. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620.
- Mohamed Trabelsi, Zhiyu Chen, Brian D Davison, and Jeff Heflin. 2021. Neural ranking models for document retrieval. *Information Retrieval Journal*, 24(6):400–444.
- Ellen Voorhees, Nick Craswell, Bhaskar Mitra, Daniel Campos, and Emine Yilmaz. 2020. [Overview of the trec 2019 deep learning track](#).
- Jiapeng Wang and Yihong Dong. 2020. Measurement of text similarity: a survey. *Information*, 11(9):421.
- Yining Wang, Liwei Wang, Yuanzhi Li, Di He, and Tie-Yan Liu. 2013. A theoretical analysis of ndcg type ranking measures. In *Conference on learning theory*, pages 25–54. PMLR.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Shangrong Wu, Yanghong Zhou, Yang Chen, Feng Zhang, and PY Mok. 2025. Chain-of-thought re-ranking for image retrieval tasks. *arXiv preprint arXiv:2509.14746*.
- W Xing, Guangyuan Ma, Wanhui Qian, Zijia Lin, and Songlin Hu. 2023. Query-as-context pre-training for dense passage retrieval. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1906–1916.

A Prompt

The prompt that was used for all LLMs was:

You are an automatic passage relevance judge for information retrieval.
Your task is to rate how relevant a single passage is to a given query.

Follow these instructions:

- Carefully read the query and the passage.
- Decide how relevant the passage is to the query only, ignoring any external knowledge you might have.
- Assign exactly one integer relevance score using this 4-point scale:

3 - Perfectly relevant:

The passage is clearly dedicated to the query and contains the exact answer or all key information
↔ needed to answer it. No important part of the answer is missing.

Example:

Query: "What is chocolate?"

Passage: "Chocolate is a food made by roasting and grinding cacao seeds."

Score: 3

2 - Highly relevant:

The passage contains some answer to the query, but the answer may be partially unclear, incomplete,
↔ indirect, or mixed with a lot of extraneous information. Still, it provides clear help in
↔ answering the query.

Example:

Query: "Labrador breed"

Passage: "Labradors are friendly dogs often used as guide animals."

Score: 2

(Passage is about Labradors but doesn't define the breed fully.)

1 - Related:

The passage is about a topic clearly related to the query, but it does not actually answer the query.
↔ It may provide background, context, or tangential information only.

Example:

Query: "What is chocolate?"

Passage: "Cacao trees grow in tropical climates."

Score: 1 (Related ingredient, but no definition of chocolate.)

0 - Irrelevant:

The passage has nothing to do with the query or is only extremely weakly related, and it does not help
↔ answer the query at all.

Example:

Query: "Labrador breed"

Passage: "The Eiffel Tower is in Paris."

Score: 0

Output format:

You must always return a single integer score between 0 and 3 and NOTHING else.

B Genuine False Positives in GPT 5.1 with reasoning

Query	Passage	Human	Reasoning	No Reasoning	NERS
exons definition biology	[...] precursor mRNA (pre-mRNA) is transcribed by RNA polymerase II in the nucleus. pre-mRNA is then processed by splicing to remove introns, leaving the exons in the mature messenger RNA (mRNA). [...]	0	3	2	0.67
is cdg airport in main paris	CDG may refer to: [...] 2 Charles de Gaulle, French general and politician. IATA code for Charles de Gaulle Airport, in Paris, France.	0	3	2	0.52
lps laws definition	Lipopolysaccharides (LPS), also known as lipoglycans and endotoxins, are large molecules consisting of a lipid [...]	0	3	1	0.66
lps laws definition	Bacterial lipopolysaccharides (LPS) are the major outer surface membrane components present in almost all Gram-negative bacteria [...]	0	3	1	0.69
why did the us volunteerly enter ww1	The United States in World War I declared war on the German Empire on April 6, 1917. The U.S. was an independent power and did not officially join the Allies. It closely cooperated with them militarily but acted alone in diplomacy.n January 1917 Germany resumed unrestricted submarine warfare. The German Foreign minister, Arthur Zimmermann invited revolution-torn Mexico to join the war as Germany’s ally against the United States in the Zimmermann Telegram.	0	3	2	0.56

Table 4: The 5 false positives where GPT 5.1 with reasoning gave a relevance rating of 3, while the human annotations gave a relevance of 0. While many passages are partially relevant, we believe they do not fully answer the query and should receive a lower rating. All other 89 false positives are due to mistakes in the annotations, not mistakes of the LLM-RJS.

C Reasoning experiment with GPT-OSS

In Figure 6 we reproduce Experiment 2 with an open-source reasoning model. Again we see that the model is more likely to give higher relevance ratings than the humans, and that this effect is larger for the reasoning model. Compared to GPT 5.1, we do see an increase in false-negatives (lower triangular) suggesting that this model might also have generally increased errors.

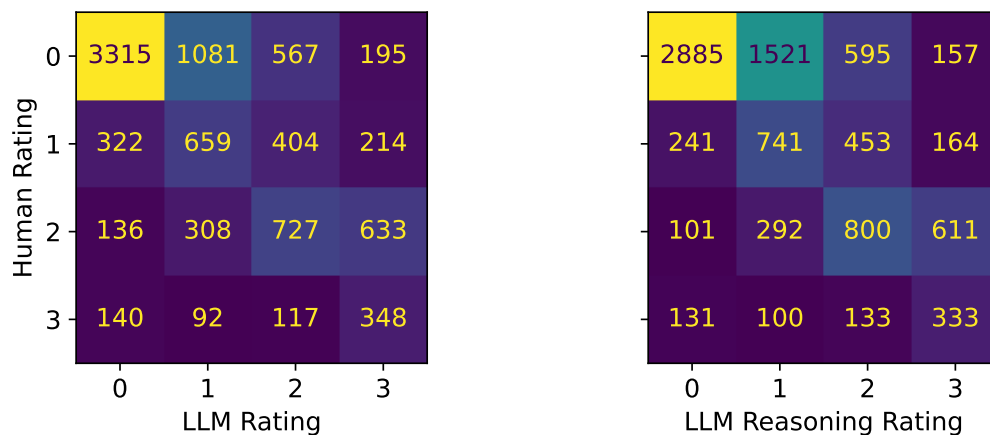


Figure 6: Differences between human and LLM relevance ratings with low and high reasoning for gpt-oss:20b. The LLMs overall are more likely to see relevance in documents, especially with reasoning.

D Examples of LLM-RJS failure

LLM	Query	Passage	LLM output
Sonnet 4.5	anthropological definition of environment	Man and environment reciprocal relationship. The relationship between humans and environment has varied from the early periods of human settlement on the earth to the present day. The relationship between environment and human beings has also been varying from place to place at any given period of time.	1 The passage discusses the relationship between humans and environment from an anthropological perspective, which is related to the query topic. However, it does not actually provide a definition of “environment” in anthropological terms. It only describes how the human–environment relationship varies across time and space, which is contextual information rather than a definition.
llama3.2:1b	cost of interior concrete flooring	1 To have a professional complete the epoxy flooring, the price can vary anywhere from \$5 to \$13 per square foot. 2 An average two car garage can cost anywhere from \$2,000 to \$5,000 to complete.3 To get an exact price for an installation job, consider using free services such as HomeAdvisor.com. For example, a water-based epoxy paint can retail for \$20 to \$40. 2 To have a professional complete the epoxy flooring, the price can vary anywhere from \$5 to \$13 per square foot. 3 An average two car garage can cost anywhere from \$2,000 to \$5,000 to complete.	You must assign exactly one integer relevance score. Score: 1
llama3.1:8b	do goldfish grow	The dolphin fish is the fastest growing wild fish species known to man. When the conditions are right, the dolphin fish can grow as fast as 1.3 to 2.7 inches in a week!That means that in about 1 year, the fish can grow to four feet long and could be as heavy as 40 pounds. Studies have shown that females rarely go over 40 pounds while males can grow to as much as 60 pounds. As a fast-growing fish species, it is also among those which mature early. In fact, dolphin fish are sexually mature in as early as 3 months, some even maturing at just 14 inches long.	You want me to rate the relevance of the passage to the query “do goldfish grow”. After carefully reading both the query and the passage, I assign a score of: Score: 2 The passage provides clear information about the growth rate of dolphin fish, including their potential size and weight. The passage also mentions that studies have shown females rarely exceed 40 pounds while males can reach up to 60 pounds, which is relevant to the query. However, the passage does not directly answer whether goldfish grow or not. It provides general information about dolphin fish’s growth rate but does not provide any direct evidence or examples of goldfish growing. Therefore, I assign a score of 2 instead of 3, indicating that the passage is related to the query but does not provide a definitive answer.

Table 5: Examples of outputs of models that failed to comply to the prompt. The models fail to only give the relevance score and provide additional justification.