

CoRank: LLM-Based Compact Reranking with Document Features for Scientific Retrieval

Runchu Tian*
Siebel School of Computing and Data
Science, University of Illinois
Urbana-Champaign
Urbana, United States
runchut2@illinois.edu

Xueqiang Xu*
Siebel School of Computing and Data
Science, University of Illinois
Urbana-Champaign
Urbana, United States
xx19@illinois.edu

Bowen Jin
Siebel School of Computing and Data
Science, University of Illinois
Urbana-Champaign
Urbana, United States
bowenj4@illinois.edu

SeongKu Kang
Department of Computer Science and
Engineering, Korea University
Seoul, South Korea
seongkukang@korea.ac.kr

Jiawei Han[†]
Siebel School of Computing and Data
Science, University of Illinois
Urbana-Champaign
Urbana, United States
hanj@illinois.edu

Abstract

Scientific retrieval is essential for advancing scientific knowledge discovery. Within this process, document reranking plays a critical role in refining first-stage retrieval results. Recently, LLMs have become popular reranking backbones for their strong domain knowledge. However, standard LLM listwise reranking faces challenges in the scientific domain. First-stage retrieval is often suboptimal in the scientific domain, so relevant documents are ranked lower. Meanwhile, conventional listwise reranking places the full text of candidates into the context window, limiting the number of candidates that can be considered. As a result, many relevant documents are excluded before reranking, constraining overall retrieval performance. To address these challenges, we explore semantic-feature-based compact document representations (e.g., categories, sections, and keywords) and propose CoRANK, a *training-free, model-agnostic* reranking framework for scientific retrieval. It presents a three-stage solution: (i) offline extraction of document features, (ii) coarse-grained reranking using these compact representations, and (iii) fine-grained reranking on full texts of the top candidates from (ii). This integrated process addresses suboptimal first-stage retrieval: Compact representations allow more documents to fit within the context window, improving candidate set coverage, while the final fine-grained ranking ensures a more accurate ordering. Experiments on 5 academic retrieval datasets show that CoRANK significantly improves reranking performance across different LLM backbones (average nDCG@10 from 50.6 to 55.5). Overall, these results underscore the interaction between information extraction and information retrieval, demonstrating how

structured semantic features can enhance reranking in the scientific domain. Code and other reproducibility materials are [available](#).

CCS Concepts

• **Information systems** → **Retrieval models and ranking**; *Document representation*; *Information extraction*; • **Computing methodologies** → *Natural language processing*.

Keywords

Scientific Information Retrieval, Document Reranking, Listwise Ranking, Large Language Models

ACM Reference Format:

Runchu Tian, Xueqiang Xu, Bowen Jin, SeongKu Kang, and Jiawei Han. 2026. CoRank: LLM-Based Compact Reranking with Document Features for Scientific Retrieval. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770855.3818844>

1 Introduction

Scientific retrieval [27, 81] is the foundational infrastructure for scientific activities. While current retrievers are effective at retrieving broadly relevant scientific papers, they often struggle to differentiate among documents covering similar topics [37, 65], failing to make fine-grained relevance estimation. Therefore, the reranking stage [8, 26] is particularly important [13], as it refines the first-stage retrieval to better distinguish between closely ranked documents.

Recently, Large Language Models (LLMs) [15, 47] have significantly advanced the field of document reranking, particularly through their application in *listwise reranking* [39, 70]. In this setting, LLMs jointly assess a set of retrieved candidates within their context window and generate a reordered ranking based on their relevance. Previous works [11, 36, 52, 58] show that LLM-based listwise rerankers outperform prior embedding-based approaches [44, 46] on multiple evaluation datasets [73].

The standard practice in LLM listwise reranking [11, 36, 52, 53, 70] is to input the full text of each candidate document into the context window. In scientific retrieval, however, this approach faces

*Both authors contributed equally to this research.

[†]Corresponding author.



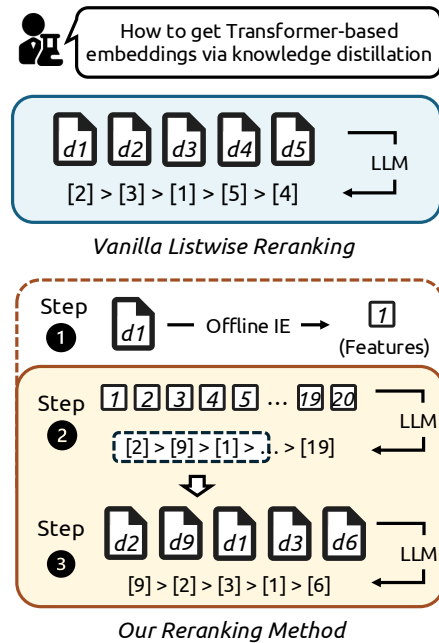


Figure 1: We extract semantic features, rerank a larger candidate pool with them, and finally refine the shortlist with full text instead of directly reranking with full text.

key limitations. In the scientific domain, first-stage retrievers often show limited performance [21, 23, 25], so truly relevant documents may not rank high enough in the first-stage retrieval. Moreover, due to the substantial token overhead of full text representation and finite context length, rerankers can only operate on a limited number (typically 20 per prompt) [39, 70] of candidates. Consequently, when the first-stage retrieval is suboptimal, many relevant documents are excluded from reranking consideration, which inherently constrains overall performance [55].

A straightforward solution is to expand the context window of LLM rerankers to handle more candidates. This can be done using long-context LMs [18, 51] or test-time scaling strategies [69, 83], such as multi-round reranking [39, 70]. However, these methods have notable drawbacks: long-context LMs require additional training [36], suffer from “loss in the middle” problem [17, 33], and incur high computational costs [38]; sliding windows similarly increase inference latency and overhead [39, 70]. To better address the limitations, we explore compressing each document’s representation by abstracting high-level features (e.g., categories, keywords), enabling more documents to fit within a fixed context window. This achieves a similar effect to window expansion but with greater efficiency.

We propose CoRANK, a *zero-shot, model-agnostic* reranking framework for scientific retrieval as exhibited in Figure 1. CoRANK consists of three stages: (i) **Offline Preprocessing**: extract high-level semantic features like categories and keywords from unstructured scientific documents; (ii) **Coarse-grained Reranking**: use these compact representations for an initial ranking and select a subset of top candidates; (iii) **Fine-grained Reranking**: rerank *the top candidates* using full scientific documents to recover the potentially missing details during information extraction. This integrated

reranking framework aims to address the suboptimal first-stage retrieval. Compact representations allow more documents to fit within the context window, improving candidate set coverage, while the final fine-grained ranking ensures a more accurate ordering.

We comprehensively evaluate CoRANK on five diverse scientific retrieval datasets spanning computer science [1, 41], biomedical [6], and clinical [78, 79] domains. Our experiments cover a diversity of LLM backbones, including both proprietary [14, 47] and open-source [15, 20, 54] models. Results show that CoRANK consistently outperforms existing LLM-based zero-shot reranking baselines, achieving an average absolute improvement of +4.9 nDCG@10. Notably, our method delivers strong gains even in challenging settings with suboptimal first-stage retrieval, while also maintaining high efficiency by reducing token consumption compared to the widely adopted sliding window strategy [70]. Extensive ablation studies and further analyses validate the contribution of each component and demonstrate the robustness of our method.

- # 1 We reveal the limitations of LLM listwise reranking in scientific retrieval.
- # 2 We explore semantic features for document representation in reranking and propose CoRANK, a novel zero-shot reranking framework.
- # 3 We show that CoRANK significantly improves reranking performance of LLM-based listwise approach in scientific retrieval.

These findings underscore how semantic information extraction and retrieval interact, with pre-analysis of document features enabling more effective reranking in scientific search.

2 Related Work

2.1 Classical Document Reranking

Document reranking [24, 57, 59] is a critical component in information retrieval (IR) [3, 68], especially when high retrieval precision is required and the first-stage retriever alone is insufficient. Early reranking methods were lexical and probabilistic [61, 64]. Like early sparse retrieval methods [3, 60], they relied on term overlap between the query and document to adjust relevance scores. However, unlike initial retrieval, rerankers *revisit* a smaller set of candidate documents returned by the first stage and apply more sophisticated weighting or adjustment techniques to refine the ranking. Nevertheless, these approaches were limited by their reliance on exact matching and failed to capture in-depth semantics [12, 24].

With advances in deep learning, neural reranking methods [16, 76], particularly those based on learning-to-rank [7, 35] frameworks, became prevalent. These models overcome previous limitations, capturing semantic similarity beyond surface-level term matches. The emergence of pre-trained language models (PLMs) [10, 56, 77] further transformed document reranking [45, 86]. Cross-encoder-based rerankers [32], in particular, have demonstrated strong performance by modeling query-document relevance scores and optimizing pointwise, pairwise, or listwise loss [91]. Unlike bi-encoder architectures used in dense retrieval [24], cross-encoders jointly encode both the query and each candidate document, allowing the model to directly capture rich dependencies between queries and documents from first-stage retrieval.

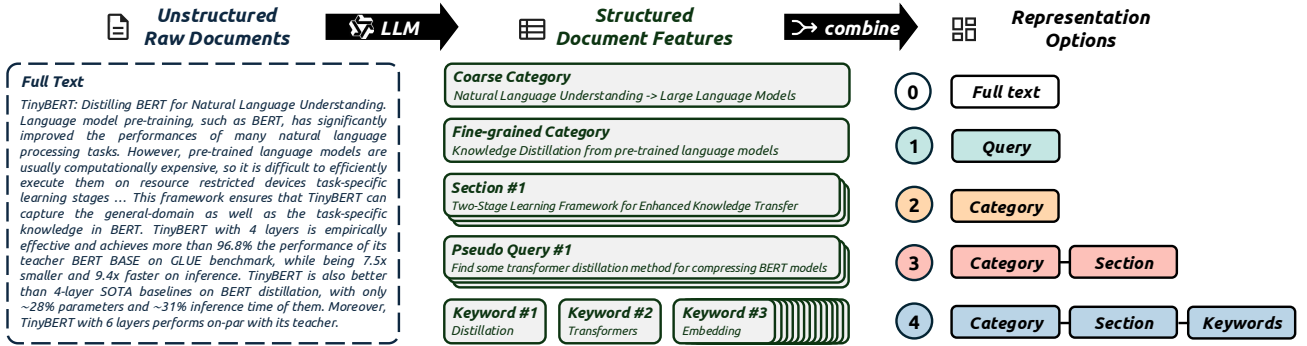


Figure 2: Overview of our feature extraction pipeline: from unstructured documents, we apply zero-shot LLM information extraction to obtain document features including categories, sections, pseudo queries, and keywords, which are then combined into compact representations.

2.2 LLM-Based Listwise Document Reranking

Recently, large language models (LLMs) [49, 75] have revolutionized the field of NLP. With strong general knowledge and instruction-following [50] capabilities, LLMs offer a new paradigm for listwise document reranking. A classic proposal [70] involves placing the query and a list of candidate documents into the context window and prompting the LLM to generate a reordered ranking. This zero-shot strategy has been shown to outperform the traditional reranking baselines [39, 70].

After the introduction of LLM-based listwise reranking, research has rapidly extended in several directions. For *model development*, works such as RankVicuna [52] and RankZephyr [53] pioneered open-source reranking LLMs, while RankMistral [36] leveraged long-context training to adapt LLMs for larger document lists in a single input. For *context-length adaptation*, Sliding window strategies [39, 70] partition candidate lists into overlapping chunks for broader coverage, addressing the limited context windows of LLMs. For *inference efficiency*, FIRST [11] improves the speed of listwise reranking by restricting the process to first-token decoding. Liu et al. [34] and Li et al. [31] use passage embeddings and chunked text blocks to represent documents reducing the token cost per document. From the perspective of *self-consistency*, ScaLR [58] introduces self-calibration to improve consistency across windows. Meanwhile, Tang et al. [72] enhance reranking quality through permutation consistency.

Since our method operates on document representations in a *plug-and-play* fashion, it is orthogonal to most previous baselines and can be applied jointly with them (we illustrate this with the sliding window strategy in Section 5). Notably, two prior studies [31, 34] have also identified limitations of full-document representations, but their approaches are fundamentally different from ours: one [34] adopts non-natural language embeddings, while the other [31] merely relies on document chunking. Furthermore, their analyses focus on the general domain, whereas ours is conducted in the scientific domain.

Overall, our work highlights the limitations of full-document representations for listwise reranking in scientific retrieval. And we leverage IE-based features as document representations for

reranking. This makes our contribution novel and significant within the reranking literature.

3 Preliminary Analysis

In this section, we first define the task of listwise reranking, then show the limitations of current full-document-based methods and explore compact and informative document representation options for reranking in the scientific domain.

3.1 Problem Definition

Reranking Input. Given a query q and a large corpus \mathcal{P}_n of n scientific papers, a first-stage retriever selects a ranked list of m candidate documents $C = \{p_1, p_2, \dots, p_m\} \subset \mathcal{P}_n$. Typically we have $m \ll n$ due to the scale of the corpus and the finite capacity of rerankers. This candidate list C is then passed to a reranking model, which aims to reorder the documents such that more relevant ones are ranked higher.

Reranking Objective. The goal of document reranking is to find a permutation $\sigma : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ whose application

$$C' = [p_{\sigma(1)}, p_{\sigma(2)}, \dots, p_{\sigma(m)}]$$

orders documents in more accurate *descending* relevance to q . The quality is often [11, 36, 70] evaluated using top- k metrics such as nDCG, with small k values that stress accuracy at the top of the ranked list.

LLM Listwise Reranking. Large language models can inspect m candidate documents in their context window and generate a permutation that reflects their relevance ordering. Formally, the model acts as follows:

$$\sigma_{\text{LLM}} = \text{LLM}(q, C)$$

Applying this permutation to the candidate list yields

$$C' = [p_{\sigma_{\text{LLM}}(1)}, p_{\sigma_{\text{LLM}}(2)}, \dots, p_{\sigma_{\text{LLM}}(m)}].$$

3.2 Current Limitations

The standard approach in LLM-based listwise reranking [11, 36, 52, 53, 70] feeds the *full text* of each candidate into the model's context window. While this allows LLMs to capture the complete content of each document, it faces some challenges in the scientific domain.




Figure 3: Token efficiency and performance comparison across different document representations. (a) Per-document token length distribution. (b) nDCG@10 scores for different numbers of documents in the context window of single LLM input (c) Context token overhead to reach equal reranking performance.

Suboptimal First-Stage Retrieval. Unlike in general-domain scenarios, first-stage retrievers—whether sparse [45, 60] or dense [19, 80]—struggle to generalize to the scientific domain [5, 73]. This is due to the complexity of long-tail concepts [22, 23] and the lack of large-scale supervised training data [5, 29]. Consequently, their retrieval performance is limited, and truly relevant documents are often ranked much lower [1].

Token Consumption. Full text representation also introduces significant token overhead, with individual scientific papers often consuming hundreds or even thousands of tokens [1, 41, 73]. At the same time, LLMs have a limited effective context length [9, 84] and are known to suffer from issues such as positional bias for long inputs [33, 74]. As a result, full text rerankers are constrained to operate on a small number of candidates, typically 20 documents per input prompt [39, 52, 53, 70].

When first-stage retrieval is suboptimal and reranking operates over a narrow candidate set, the overall reranking performance is significantly constrained.

3.3 Semantic Features as Alternatives

3.3.1 Semantic Feature Construction. Given the limitations of full text representations in scientific-domain reranking, we explore an alternative approach, which is based on document-level semantic features such as categories and keywords. The intuition is that these features are both significantly more concise and capable of preserving the core semantics. Therefore, more candidate documents can be effectively considered during reranking. Specifically, we examine the following four types of document semantic features:

Category. Categories [71, 89] offer a high-level topical overview of scientific documents. Each document is assigned a three-level hierarchical category path in the format $\{Category\} \rightarrow \{Subcategory\} \rightarrow \{Subsubcategory\}$, which provides a broad-to-specific classification on its topic.

Section. Sections [90] consist of multiple subtitle-style strings, each summarizing a major part of the scientific document. They

capture the internal structure of the document and serve as mid-level semantic signals that enhance category-level summaries.

Keyword. Keywords [28, 62] are terms or entities that represent fine-grained lexical concepts within a document. They provide the most specific information among different granularities.

Pseudo Query. Pseudo queries [22, 63] are synthetic user questions based on the content. This feature offers a unique query-aligned perspective of the document, simulating how the document might be retrieved in real-world scenarios.

These four features represent the most common types of document-level information extraction (IE) [43, 82], varying in both granularity and style. All semantic features are extracted using LLM-based *information extraction* methods [87, 88]. We adopt zero-shot prompting with instruction-tuned models such as GPT-4.1 [48] and the Qwen model family [54]. This approach leverages the inherent knowledge of pretrained language models [30, 40] to extract information without requiring task-specific training [87], while demonstrating strong generalization across domains and schema variations [85].

Among these features, *sections*, *keywords*, and *pseudo queries* have multiple elements for each document. To reduce noise and focus on the most relevant content, we apply an adaptive selection strategy during inference time. We compute dense embedding similarities between the query and each extracted element and retain only the most relevant: 5 keywords (from 30), 1 pseudo query (from 20), and 1 section (from 3). This improves content relevance while minimizing token overhead.

However, relying on a single feature often lacks representational power. For example, using only the category tends to be too coarse-grained and offers limited discriminative ability, while keywords alone may lack sufficient context or background information. To address this, we represent each document using *combinations* of features. Specifically, we explore the following four configurations of different overall granularities.

Form 1. Pseudo Query

Form 2. Category

Form 3. *Category + Section*

Form 4. *Category + Section + Keywords*

Ablation studies on the effect of each specific component can further be found in Section 5.2.2. The overview of representation construction is presented in Figure 2.

3.3.2 Empirical Validation. To evaluate the effectiveness of our feature-based representations, we conduct a series of experiments on the LitSearch [1], using GPT-4.1-mini [48] as the reranking backbone. Semantic features are extracted using Qwen3-8B [54]. Following Gangi Reddy et al. [11], we use Contriever [19] as the first-stage retriever and the similarity encoder for adaptive selection.

Token Length Per Document. We first evaluate the token cost of different representations. As shown in Figure 3 (a), full text inputs consume around 200 tokens per document on average, with some exceeding 5,000 tokens, limiting the number of documents that can fit into the context window. In contrast, semantic-feature-based representations (Forms 1-4) are much more compact, each averaging between 10 and 50 tokens.

Number of Documents in a Single Input. We then evaluate how different representations perform as the number of candidate documents increases (up to 200) within a 32k-token context window. As shown in Figure 3 (b), all representations improve as more documents are included. In the range of 0 to 80 candidates, full text representations perform better than feature-based ones. This is expected, since full text retains complete document information, though at a much higher token cost. However, beyond 80 candidates, full text inputs frequently exceed the context length limit. In contrast, feature-based methods, especially Form 4, remain compact and can handle more candidates without issue, while still achieving comparable performance.

Token Cost for Equal Reranking Performance. Finally, we examine how many context tokens are needed to achieve the same reranking performance across different representations. Using results from the previous two experiments, we estimate the total token cost by combining the number of documents required with the average token count per document. We take full text reranking with 20 documents as the 100% performance baseline, and compare how many tokens each method needs to reach 70%, 85%, 100%, 115%, and 130% of that level. As shown in Figure 3 (c), compact feature-based representations can match full text performance with significantly fewer tokens. For instance, reaching 100% performance requires around 3,500 tokens for full text, but only 1,000 or even a few hundred tokens for feature-based inputs.

3.3.3 Discussion. The results above demonstrate that feature-based document representations offer clear advantages over full text inputs in reranking. They can express the core content of a document using far fewer tokens, allowing significantly more retrieved candidates to be included within the same token overhead. This is especially valuable for recovering relevant documents that were assigned lower scores by the first-stage retriever. Considering both token efficiency and reranking performance, we finally select Form 4 as our feature-based representation.

However, in the second set of experiments where the number of documents is held constant, we also observe that feature-based

representations may underperform full text. After all, offline information extraction is *not* perfect. The process may omit subtle but important cues present in the full text.

Therefore, these results motivate a balanced reranking strategy: using semantic features in the early stage to cover a broader range of candidates, followed by a refinement stage that reranks the top candidates using full text inputs for more in-depth comparison.

4 Methodology

Based on the preliminary analysis, we propose a new three-stage reranking framework in the scientific domain: (i) offline extraction of structured semantic features; (ii) coarse-grained reranking using the features to select a subset of top candidates; and (iii) fine-grained reranking over the subset with full text inputs.

4.1 Compact and Informative Representation

To construct compact, feature-based representations for each scientific document, we first perform document-level information extraction offline. As described in Section 3.3, we use LLMs to obtain target semantic features: *category*, *section*, and *keyword*. Formally, for a given document p_i , we obtain:

$$\text{LLM}(p_i) = \text{Category}_i, [\text{Section}_i^j], [\text{Keyword}_i^j]$$

These features are designed to capture the core semantics of scientific documents in a token-efficient manner. The extraction is performed offline, and the results are cached and reused during inference time to avoid any additional runtime delay.

4.2 Coarse-grained Reranking w. Compact Features

Given a query q , we obtain a document relevance ranking $C = [p_1, p_2, \dots, p_m]$ from a first-stage retriever. We then replace each document with a compact semantic representation and apply LLM-based listwise reranking to identify a high-quality subset.

For sections and keywords, we apply *adaptive selection* to select only those most relevant to the query. Specifically, for each unit (single section or keyword) u , we compute the cosine similarity with text embedding and select the top-5 keywords and the most relevant section. Analysis on the number of keywords used can be found in Section 5.4. The final feature-based representation r_i is formed as:

$$r_i = \text{Concat} [\text{Category}_i, \text{Section}_i^*, \text{Keywords}_i^*]$$

where $*$ indicates adaptively selected elements.

With the feature-based representations $\mathcal{R} = [r_1, r_2, \dots, r_m]$, we then use the LLM to perform listwise reranking, producing a permutation $\sigma_{\text{feat}} = \text{LLM}(q, \mathcal{R}) \in S_m$. Applying this permutation to the original candidate list $C = [p_1, p_2, \dots, p_m]$ yields the reranked output for the coarse-grained reranking:

$$C_{\text{feat}} = [p_{\sigma_{\text{feat}}(1)}, p_{\sigma_{\text{feat}}(2)}, \dots, p_{\sigma_{\text{feat}}(m)}]$$

We then keep the top- k documents from this list to form the seed set for full text reranking:

$$C_{\text{seed}} = [p_{\sigma_{\text{feat}}(1)}, \dots, p_{\sigma_{\text{feat}}(k)}], \quad k < m$$

Coarse-grained reranking with compact document representations greatly expands the number of candidates that can be considered in the same LLM input. This broader coverage helps recover scientific documents that were assigned low scores by the first-stage retriever.

4.3 Fine-grained Reranking w. Full Text

In the second reranking stage, we refine the ranking over the seed set of candidates using full text representation. Specifically, we take the seed set $C_{\text{seed}} = [p'_1, p'_2, \dots, p'_k]$, (we use ' to distinguish them from the first-stage input) obtained from the previous stage, and replace each compact representation with its original document.

Let $\mathcal{T} = [t_1, t_2, \dots, t_k]$ denote the full text of the selected documents, where t_i corresponds to the full text content of passage p'_i . We then once again use LLM to perform listwise reranking:

$$\sigma_{\text{text}} = \text{LLM}(q, \mathcal{T}) \in S_k$$

With the permutation σ_{text} we get the final ranking:

$$C_{\text{final}} = [p'_{\sigma_{\text{text}}(1)}, \dots, p'_{\sigma_{\text{text}}(k)}]$$

Fine-grained reranking with full text inputs recovers details that may be lost during the information extraction process. Since the candidate set has already been narrowed down, full text representation can now be used without exceeding the LLM’s context limit.

Overall, our strategy effectively addresses the challenges of LLM-based listwise reranking in the scientific domain. The coarse reranking stage expands the reranking pool, improving robustness to suboptimal first-stage retrieval in the scientific domain. The fine-grained reranking stage, in turn, preserves sufficient detail for precise relevance comparisons.

5 Experiments

In this section, we first outline the experimental setup, then present the main reranking performance results, followed by ablation studies and further analysis.

5.1 Experimental Setup

We describe the experimental setup for evaluating CoRANK and prior zero-shot LLM reranking baselines on multiple scientific retrieval benchmarks. We keep the evaluation protocol and reranking parameters consistent across methods for fair comparison.

Datasets & Metric. We evaluate different reranking methods on five high-quality scientific retrieval benchmarks: LitSearch [1], CS-FCube [41], NFCorpus [6], SciFact [79] and Trec-Covid [78], which collectively cover diverse scientific domains including computer science, biomedical and clinical information retrieval. Following standard practice [11, 36, 39, 70], we report nDCG and Recall as our evaluation metrics for top-10 document.

Models. For the reranking backbones, we use Qwen3-32B [54] as the representative open-source model, Gemini 2.0 Flash [14] and GPT-4.1-mini [48] as proprietary examples. For generative parameters, we use a temperature of 1.0 and a fixed random seed of 42 for reproducibility.

For semantic feature extraction, we find that small open-source models are sufficiently effective; we therefore adopt Qwen3-8B [54]

for document feature extraction. For both first-stage retrieval and semantic similarity scoring for feature filtering, we use Contriever [19] following previous works [11, 55]. To ensure the robustness against different first-stage retrievers, we also evaluate the performance of our method with diverse retrievers in Section 5.3.2.

Baselines. Our method is a *zero-shot, model-agnostic* approach. Therefore, we primarily compare against prior reranking baselines under the same setting, most notably vanilla listwise reranking [70] with zero-shot instruction-following. We also consider the sliding window strategy [39, 70], a test-time scaling technique [83] designed to expand the reranking context. Results are reported both with and without the sliding window strategy. We conduct a focused comparison between our method and the sliding window strategy in Section 5.3.1.

We also include some reranking methods from the literature that represent the state-of-the-art supervised LLM rerankers like RankVicuna [52], RankZephyr [53], and RankMistral [36], which are trained on large-scale general-domain datasets such as MS MARCO [4]. While these models differ from ours in both model size and supervision conditions, we report them for reference to illustrate the overall performance distribution.

For a fairer comparison with supervised baselines, we evaluate our method using smaller-parameter LLMs (7B–8B) in Section 5.3.3, allowing us to isolate the effect of model size from supervision.

Reranking Parameters. We follow the standardized setup [11, 39, 52, 53, 70] from prior work to unify reranking parameters. For vanilla reranking, we include 20 full text documents within the context. An exception is RankMistral [36], which benefits from long-context training and is able to process 100 full text documents in a single input. When using the sliding window strategy, we rank a total of 100 full text documents by applying a window size of 20 with a step size of 10. For CoRANK, we include 200 compact representations in the coarse reranking stage, followed by 20 full text documents in the fine-grained reranking stage.

5.2 Main Results

5.2.1 Reranking Performance. We begin by evaluating the performance of different reranking methods on the target academic retrieval benchmarks, as shown in Table 1.

First, we observe that all reranking methods outperform the first-stage dense retriever baseline, confirming the critical role of reranking in scientific document retrieval. However, while supervised models such as RankVicuna [52] do yield noticeable improvements, their gains remain limited, partly because these models are relatively small (7B–8B parameters) and thus less competitive despite having supervision. Meanwhile, current general LLMs already demonstrate strong instruction following capabilities for zero-shot reranking. Furthermore, our proposed method, CoRANK, which is model-agnostic, and training-free, delivers even stronger and consistent gains across different LLM backbones. Without using sliding windows, CoRANK improves the average nDCG@10 from 47.2 to 54.9, and when combined with sliding window inputs, it still provides an average gain of +2.0 nDCG@10 (from 54.1 to 56.2).

We also find that both CoRANK and the sliding window strategy independently lead to substantial performance gains (ours being

¹Low because each query in Trec-Covid has far more than 10 relevant documents.

Table 1: Reranking performance of different methods on LitSearch [1], CSFCube [41], NFCorpus [6], SciFact [79] and Trec-Covid [78] with nDCG@10 (N@10) and Recall@10 (R@10) as metrics. Higher scores under the same setting are bold. All values are shown as percentage absolute scores, with improvement (blue) and degradation (red) relative to first-stage retrieval.

Reranking Model	Reranking Strategy	LitSearch		CSFCube		NFCorpus		SciFact		Trec-Covid ¹	
		N@10	R@10	N@10	R@10	N@10	R@10	N@10	R@10	N@10	R@10
None	First-stage retrieval only	14.2	21.2	23.6	16.4	32.8	15.5	67.7	81.3	59.6	1.6
RankMistral	Vanilla Reranking <i>Full</i>	20.1	24.3	21.9	13.3	22.9	9.5	16.9	27.3	52.3	1.3
RankVicuna	Vanilla Reranking <i>Sliding</i>	26.5	32.4	28.7	17.5	26.3	11.6	29.4	40.7	64.1	1.7
RankZephyr	Vanilla Reranking <i>Sliding</i>	34.1	39.0	34.1	22.7	28.6	15.1	44.4	80.0	59.5	1.6
Qwen3-32B	Vanilla Reranking <i>Full</i>	25.7	27.3	28.4	20.2	36.8	17.2	77.6	85.7	67.0	1.7
	CoRANK <i>Full</i>	39.6	41.9	31.8	23.4	39.8	18.5	80.8	89.1	76.4	2.0
Qwen3-32B	Vanilla Reranking <i>Sliding</i>	39.8	43.5	31.2	22.9	38.8	18.0	80.0	89.8	76.7	2.1
	CoRANK <i>Sliding</i>	42.2	46.2	35.5	26.6	39.8	18.7	81.7	90.3	78.1	2.1
Gemini 2.0 Flash	Vanilla Reranking <i>Full</i>	26.1	27.3	28.5	18.8	35.9	16.8	75.6	85.4	68.9	1.8
	CoRANK <i>Full</i>	40.7	44.2	35.1	27.8	38.0	17.5	77.6	88.3	76.1	2.0
Gemini 2.0 Flash	Vanilla Reranking <i>Sliding</i>	40.8	44.6	32.2	23.2	38.3	17.9	77.2	88.8	79.7	2.2
	CoRANK <i>Sliding</i>	43.3	48.9	38.1	31.0	39.0	18.6	78.0	91.1	79.2	2.2
GPT-4.1-mini	Vanilla Reranking <i>Full</i>	26.1	27.2	28.7	19.7	37.1	16.8	76.7	84.9	68.4	1.8
	CoRANK <i>Full</i>	46.3	49.4	39.0	30.5	41.4	19.1	79.6	89.8	80.7	2.2
GPT-4.1-mini	Vanilla Reranking <i>Sliding</i>	41.9	44.4	34.9	25.3	40.0	17.9	79.3	90.5	81.0	2.2
	CoRANK <i>Sliding</i>	45.8	49.1	39.4	30.4	41.4	19.2	80.4	91.6	80.5	2.2

larger and more efficient; see Section 5.3.1). This is likely because both methods expand the reranking scope, which is particularly valuable in the scientific domains considering the suboptimal first-stage retrieval. These findings also reinforce our earlier analysis of current limitations presented in Section 3.2.

5.2.2 Ablation Studies. To evaluate the contribution of each component in our design, we conduct ablation studies on the LitSearch [1] dataset using GPT-4.1-mini [48] as the backbone.

Specifically, we remove individual semantic features—*category*, *section*, and *keywords*—from the compact representation to assess their relative importance. In addition, we ablate two key components of our pipeline: *adaptive selection* and *fine-grained reranking*, to understand their impact on overall performance.

As shown in Figure 4, removing any single semantic component (category, section, or keywords) consistently degrades reranking performance, confirming that each feature contributes complementary information. Category captures coarse-grained background context that tends to align more easily with the query but offers less specific detail. In contrast, section and keywords reflect finer-grained semantics that are harder to match but, when relevant, provide highly targeted signals. This balance shows the value of combining semantic features across different levels of abstraction.

Additionally, we find that both adaptive selection and fine-grained reranking contribute to the final performance. Adaptive selection ensures that the most query-relevant features are retained, while the fine-grained reranking stage substantially compensates for potential information loss in compact representations by recovering more nuanced evidence from the full text.




Figure 4: Ablation study results. Foreground bars report performance after removing the labeled component; the translucent background bars indicate the full method score.

5.3 Further Analysis

5.3.1 Comparison with Sliding Window. Sliding window strategy [70] is a widely used [11, 39, 52, 53] method for overcoming the context length limitations of LLMs in listwise reranking. Instead of ranking all candidates in a single prompt, it partitions the candidate list into overlapping windows and reranks each window independently from the bottom up. This method has been shown to notably improve reranking performance [70].

As noted in Section 5.1, we evaluate both the standalone and combined use of sliding windows with our method. While CoRANK and

sliding window are orthogonal and compatible, isolating them in comparison allows us to clearly demonstrate the superior efficiency and effectiveness of our approach when used alone.

Table 2: Comparison of performance, token usage, and API cost between CoRANK and Sliding Window.

Method	N@10	R@10	Token Usage	Cost
Vanilla	26.0	27.3	1.01M	\$0.40
Sliding	40.8	44.2	9.06M	\$3.62
CoRANK	42.2	45.2	3.60M	\$1.44
Both	43.8	48.1	11.65M	\$4.66

We evaluate CoRANK and the sliding window strategy on LitSearch and CSFCube on three tested models, reporting nDCG@10, total token consumption, and API cost for GPT-4.1-mini. As shown in Table 2, CoRANK not only achieves better average reranking quality (nDCG@10 improved by 1.4), but also requires only 40% of the token budget compared to sliding windows. These results demonstrate that CoRANK provides an efficient yet effective alternative for expanding reranking range. Moreover, combining these two methods can further boost performance, offering an advanced solution for complex retrieval scenarios.

5.3.2 Impact of First-Stage Retrievers. Following previous works [11, 55], we choose Contriever [19] as our first-stage retriever in our main experiments. To further evaluate robustness to the choice of the first-stage retriever, we here additionally evaluate the performance of our method with BM25 [60], Specter2 [66], GTR-T5-large [42] as first-stage retrievers. Many of the selected retrievers, such as Contriever and Specter2, are trained on academic corpora. We therefore examine whether LLM-based reranking remains beneficial even when the first-stage retriever is already specialized for scholarly data, under both the vanilla reranking baseline and our method. We run evaluation on LitSearch [1] with Gemini 2.0 Flash [14]. The results are shown in Figure 5.




Figure 5: Comparison of reranking methods (None, Vanilla, CoRANK) across four first-stage retrievers on LitSearch with Gemini 2.0 Flash.

The first-stage retrievers exhibited differences in their standalone retrieval quality. However, our method consistently improved retrieval quality over the vanilla reranking baseline, regardless of

which initial retriever was used. In other words, our gains are not confined to weak initial retrieval: CoRANK adds value on strong first-stage retrievers as well, indicating that our method is robust to the choice of first-stage retriever.

5.3.3 Impact of Reranking Backbones. To further validate the robustness of our reranking method to different reranking backbones, we conducted additional experiments on smaller language models. While our main experiments employed LLMs like Qwen3-32B [54] and Gemini 2.0 Flash [14], it remains a question whether our method generalizes effectively to more compact language models in the 7B–8B parameter range. We test two smaller models, Mistral-7B [20] and Llama-3.1-8B-Instruct [15] as reranking backbones and the results are shown in Table 3.

Table 3: Reranking Performance with different small language models as backbones on LitSearch.

Reranking Method	Model Backbone	N@10	R@10
None	None	14.2	21.2
Vanilla	Mistral-7B	20.9	24.6
CoRANK	Mistral-7B	24.8	28.4
Vanilla	Llama-3.1-8B	22.1	25.3
CoRANK	Llama-3.1-8B	26.0	31.2
Vanilla	RankMistral	20.1	24.3
Vanilla	RankVicuna	26.5	32.4
Vanilla	RankZephyr	34.1	39.0

Our experiments confirmed two main findings. First, CoRANK remained effective on smaller language models, consistently outperforming initial retrieval and vanilla reranking baseline. Second, although it did not surpass the supervised models in absolute performance, this outcome was expected, since the supervised methods require extensive training resources. For example, RankVicuna [52] and RankZephyr [53] use GPT-3.5-distilled training data (1,000 K MS MARCO instances) and significant computational resources (320 GPU-hours on NVIDIA RTX A6000s). CoRANK holds key practical advantages: it requires no additional training data or computational resources, works seamlessly as a plug-and-play method across different models and domains, and ensures data privacy. These strengths underscore CoRANK’s value despite its lower absolute scores compared to supervised models.

5.4 Hyperparameter Study

We investigate the effect of key hyperparameters in CoRANK using the GPT-4.1-mini [48] reranker on the LitSearch [1] dataset, reporting nDCG@10, Recall@10, token usage, and estimated API cost. Our analysis focuses on understanding the *trade-off* between cost and performance, as well as the *robustness* of CoRANK across a wide hyperparameter range.

Number of Keywords. By default, CoRANK extracts 30 keywords from each document and selects the top 5 most relevant ones via cosine similarity in the embedding space, concatenating them into the document representation for *coarse-grained* reranking. We vary the number of concatenated keywords from 0 to 20 and report the immediate results from coarse-grained reranking (Table 4).

Table 4: Effect of the number of selected keywords on coarse-grained reranking: performance, token usage and cost.

# Keyword	N@10	R@10	Token Usage	Cost
0 keyword	38.9	44.2	1.80M	\$0.72
1 keyword	40.4	46.3	1.95M	\$0.78
3 keywords	40.1	44.9	2.26M	\$0.90
5 keywords	41.1	46.7	2.57M	\$1.03
10 keywords	42.4	46.5	3.33M	\$1.33
15 keywords	42.6	48.6	4.10M	\$1.64
20 keywords	42.4	47.4	4.87M	\$1.95

Results show a clear *cost–performance trade-off*: adding more keywords linearly increases token usage while steadily improving coarse-grained reranking quality. However, performance gains diminish beyond 10 keywords, as later keywords contribute weaker relevance signals. Importantly, CoRank achieves strong performance even with substantially fewer keywords, indicating that it is *robust* to this hyperparameter and does not require precise tuning to remain competitive.

Fine-Grained Pool Size. In the main experiments, the fine-grained reranker processes the top 20 candidates from the coarse stage to align with other baselines. Here, we vary the candidate pool from 5 to 100 documents and report the reranking performance of the final fine-grained stage (Table 5).

Table 5: Effect of the fine-grained candidate pool size: performance, token usage and cost.

Pool Size	N@10	R@10	Token Usage	Cost
5 docs	43.0	46.7	0.25M	\$0.10
10 docs	44.3	46.7	0.50M	\$0.20
20 docs	46.3	49.4	1.01M	\$0.40
40 docs	45.4	48.0	2.01M	\$0.81
60 docs	45.0	48.3	3.02M	\$1.21
80 docs	45.3	48.0	4.03M	\$1.61
100 docs	46.0	49.1	5.03M	\$2.01

We again observe a *robust cost–performance trade-off*: increasing the pool size up to 20 yields the most significant improvements, after which performance plateaus while token usage continues to grow. This plateau suggests that CoRank maintains *stable effectiveness* across a broad range of pool sizes, allowing practitioners to adjust this parameter flexibly according to budget without risking notable performance loss.

Summary. These studies confirm that CoRank offers tunable hyperparameters that enable smooth navigation of cost–performance trade-offs, while maintaining consistently strong results. Its effectiveness is *not narrowly dependent* on exact hyperparameter values, alleviating concerns of fragility or heavy manual tuning. For example, varying the number of keywords between 5 and 20 results in only a 3.1% variance in nDCG@10, underscoring the robustness of the method.

6 Limitations and Ethical Considerations

Limitations. CoRank relies on LLM-based extraction to build intermediate features, so performance may vary with the underlying model/prompt and extraction errors can propagate to reranking. It also incurs additional inference cost compared to embedding-only rerankers. Finally, our evaluation targets intermediate outputs and IR results, so conclusions about downstream benefits are limited.

Ethical considerations. We use established benchmarks and follow their licenses and usage terms. We do not introduce human subjects. Therefore, IRB approval is not required. When using proprietary LLM APIs, inputs may be processed by third-party providers, which is not suitable for sensitive or regulated settings. As with any reranker, the system may amplify dataset/model biases; it should not be used as a sole decision maker in high-stakes settings.

7 Conclusion

In this work, we identify two bottlenecks of standard LLM listwise reranking for scientific retrieval: first-stage retrievers often mis-rank relevant papers due to limited domain-specific supervision, while full-text listwise reranking is token-heavy and thus constrained to few candidates. To address this, we propose CoRank, a training-free, model-agnostic framework that first reranks a larger candidate set using compact semantic features, then refines the top results with full-text reranking. Experiments on five scientific datasets show consistent improvements over standard reranking baselines. Overall, our findings suggest that semantic feature extraction provides an effective compact representation for LLM reranking, enabling broader candidate coverage and stronger scientific literature retrieval.

GenAI Disclosure

We used Claude Code [2] for coding assistance and GPT-5 [67] for grammar checking and language polishing. The tools were not used to generate scientific claims, experimental results, or interpretations. All generated code and wording were reviewed, edited, validated, and approved by the authors, who take full responsibility for the implementation, results, and conclusions.

Acknowledgments

Research was supported in part by the AI Institute for Molecular Discovery, Synthetic Strategy, and Manufacturing; Molecule Maker Lab Institute (MMLI), funded by U.S. National Science Foundation under Award 2505932, NSF IIS 25-37827, and the Institute for Geospatial Understanding through an Integrative Discovery Environment (I-GUIDE) by NSF under Award No. 2118329. The research has used the Delta/DeltaAI advanced computing and data resource, supported in part by the University of Illinois Urbana-Champaign and through allocation #250851 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants OAC 2320345, #2138259, #2138286, #2138307, #2137603, and #2138296. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily represent the views, either expressed or implied, of DARPA or the U.S. Government.

References

- [1] Anirudh Ajith, Mengzhou Xia, Alexis Chevalier, Tanya Goyal, Danqi Chen, and Tianyu Gao. 2024. LitSearch: A Retrieval Benchmark for Scientific Literature Search. <https://arxiv.org/abs/2407.18940>
- [2] Anthropic. [n. d.]. Claude Code by Anthropic | AI Coding Agent, Terminal, IDE. <https://claude.com/product/claude-code>. Accessed: 2026-02-08.
- [3] Ricardo Baeza-Yates, Berthier Ribeiro-Neto, et al. 1999. *Modern information retrieval*. Vol. 463. ACM press, New York, NY, USA.
- [4] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. MS MARCO: A Human Generated MACHine Reading COMprehension Dataset. <https://arxiv.org/abs/1611.09268>
- [5] Luiz Henrique Bonifacio, Hugo Queiroz Abonizio, Marzieh Fadaee, and Rodrigo Frassetto Nogueira. 2022. InPars: Unsupervised Dataset Generation for Information Retrieval. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11–15, 2022*, Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (Eds.). ACM, New York, NY, USA, 2387–2392. doi:10.1145/3477495.3531863
- [6] Vera Boteva, Demian Gholipour, Artem Sokolov, and Stefan Riezler. 2016. A Full-text Learning to Rank Dataset for Medical Information Retrieval. In *Proceedings of the European Conference on Information Retrieval (ECIR) (Lecture Notes in Computer Science, Vol. 9626)*. Springer, Cham, 716–722.
- [7] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20–24, 2007 (ACM International Conference Proceeding Series, Vol. 227)*, Zoubin Ghahramani (Ed.). ACM, New York, NY, USA, 129–136. doi:10.1145/1273496.1273513
- [8] J. Carbonell and Jade Goldstein-Stewart. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. *ACM SIGIR Forum* 51 (1998), 209–210. <http://dl.acm.org/citation.cfm?id=3130369>
- [9] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 2978–2988. doi:10.18653/v1/P19-1285
- [10] 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)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. doi:10.18653/v1/N19-1423
- [11] Revanth Gangi Reddy, JaeHyeok Doo, Yifei Xu, Md Arafat Sultan, Deevya Swain, Avirup Sil, and Heng Ji. 2024. FIRST: Faster Improved Listwise Reranking with Single Token Decoding. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 8642–8652. doi:10.18653/v1/2024.emnlp-main.491
- [12] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2021. COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, Online, 3030–3042. doi:10.18653/v1/2021.naacl-main.241
- [13] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2021. Rethink training of BERT rerankers in multi-stage retrieval pipeline. In *European Conference on Information Retrieval*. Springer, Cham, 280–286.
- [14] Google DeepMind. 2025. Gemini 2.5: Our most intelligent AI model. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/> Google Blog.
- [15] Aaron Grattafiori et al. 2024. The Llama 3 Herd of Models. arXiv:2407.21783 [cs.AI] <https://arxiv.org/abs/2407.21783>
- [16] J. Guo, Yixing Fan, Liang Pang, Liu Yang, Qingyao Ai, Hamed Zamani, Chen Wu, W. Bruce Croft, and Xueqi Cheng. 2019. A Deep Look into Neural Ranking Models for Information Retrieval. *Inf. Process. Manag.* 57 (2019), 102067. <https://api.semanticscholar.org/CorpusID:81977235>
- [17] Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekeshe, Fei Jia, Yang Zhang, and Boris Ginsburg. 2024. RULER: What's the Real Context Size of Your Long-Context Language Models? arXiv:2404.06654 [cs.CL] <https://arxiv.org/abs/2404.06654>
- [18] Yumpeng Huang, Jingwei Xu, Junyu Lai, Zixu Jiang, Taolue Chen, Zenan Li, Yuan Yao, Xiaoxing Ma, Lijuan Yang, Hao Chen, et al. 2023. Advancing transformer architecture in long-context large language models: A comprehensive survey.
- [19] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised Dense Information Retrieval with Contrastive Learning. <https://arxiv.org/abs/2112.09118>
- [20] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. arXiv:2310.06825 [cs.CL] <https://arxiv.org/abs/2310.06825>
- [21] SeongKu Kang, Shivam Agarwal, Bowen Jin, Dongha Lee, Hwanjo Yu, and Jiawei Han. 2024. Improving Retrieval in Theme-specific Applications using a Corpus Topical Taxonomy. In *Proceedings of the ACM on Web Conference 2024 (WWW '24), Singapore, May 13–17, 2024*. ACM, New York, NY, USA, 1497–1508. doi:10.1145/3589334.3645512
- [22] SeongKu Kang, Bowen Jin, Wonbin Kweon, Yu Zhang, Dongha Lee, Jiawei Han, and Hwanjo Yu. 2025. Improving Scientific Document Retrieval with Concept Coverage-based Query Set Generation. <https://arxiv.org/abs/2502.11181>
- [23] SeongKu Kang, Yunyi Zhang, Pengcheng Jiang, Dongha Lee, Jiawei Han, and Hwanjo Yu. 2024. Taxonomy-guided Semantic Indexing for Academic Paper Search. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 7169–7184. doi:10.18653/v1/2024.emnlp-main.407
- [24] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 6769–6781. doi:10.18653/v1/2020.emnlp-main.550
- [25] Jihyuk Kim, Minsoo Kim, Joonsuk Park, and Seung-won Hwang. 2023. Relevance-assisted Generation for Robust Zero-shot Retrieval. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Mingxuan Wang and Imed Zitouni (Eds.). Association for Computational Linguistics, Singapore, 723–731. doi:10.18653/v1/2023.emnlp-industry.67
- [26] Oren Kurland and Lillian Lee. 2005. PageRank without hyperlinks: structural re-ranking using links induced by language models. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (Salvador, Brazil) (SIGIR '05)*. Association for Computing Machinery, New York, NY, USA, 306–313. doi:10.1145/1076034.1076087
- [27] Steve Lawrence, Kurt Bollacker, and C. Lee Giles. 1999. Indexing and retrieval of scientific literature. In *Proceedings of the Eighth International Conference on Information and Knowledge Management (Kansas City, Missouri, USA) (CIKM '99)*. Association for Computing Machinery, New York, NY, USA, 139–146. doi:10.1145/319950.319970
- [28] Wanhae Lee, Minki Chun, Hyeonhak Jeong, and Hyunggu Jung. 2023. Toward Keyword Generation through Large Language Models. In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces (Sydney, NSW, Australia) (IUI '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 37–40. doi:10.1145/3581754.3584126
- [29] Haitao Li, Qingyao Ai, Jia Chen, Qian Dong, Yueyue Wu, Yiqun Liu, Chong Chen, and Qi Tian. 2023. SAILER: Structure-aware Pre-trained Language Model for Legal Case Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23–27, 2023*, Hsin-Hsi Chen, Wei-Jou (Edward) Duh, Hen-Hsen Huang, Makoto P. Kato, Josiane Mothe, and Barbara Pobelte (Eds.). ACM, New York, NY, USA, 1035–1044. doi:10.1145/3539618.3591761
- [30] Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024. Pre-trained language models for text generation: A survey. *Comput. Surveys* 56, 9 (2024), 1–39.
- [31] Minghan Li, É. Gaussier, Juntao Li, and Guodong Zhou. 2024. KeyB2: Selecting Key Blocks is Also Important for Long Document Ranking with Large Language Models. <https://arxiv.org/abs/2411.06254>
- [32] Robert Litschko, Ivan Vulić, and Goran Glavaš. 2022. Parameter-Efficient Neural Reranking for Cross-Lingual and Multilingual Retrieval. In *Proceedings of the 29th International Conference on Computational Linguistics*, Nicoletta Calzolari, Chu-Ren Huang, Hansaem Kim, James Pustejovsky, Leo Wanner, Key-Sun Choi, Pum-Mo Ryu, Hsin-Hsi Chen, Lucia Donatelli, Heng Ji, Sadao Kurohashi, Patrizia Paggio, Nianwen Xue, Seokhwan Kim, Younggyun Hahm, Zhong He, Tony Kyungil Lee, Enrico Santus, Francis Bond, and Seung-Hoon Na (Eds.). International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 1071–1082. <https://aclanthology.org/2022.coling-1.90>
- [33] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173. doi:10.1162/tacl_a_00638
- [34] Qi Liu, Bo Wang, Nan Wang, and Jiaxin Mao. 2025. Leveraging Passage Embeddings for Efficient Listwise Reranking with Large Language Models. In *Proceedings of the ACM on Web Conference 2025 (Sydney NSW, Australia) (WWW '25)*. Association for Computing Machinery, New York, NY, USA, 4274–4283. doi:10.1145/3696410.3714554

- [35] Tie-Yan Liu et al. 2009. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* 3, 3 (2009), 225–331.
- [36] Wenhan Liu, Xinyu Ma, Yutao Zhu, Ziliang Zhao, Shuaiqiang Wang, Dawei Yin, and Zhicheng Dou. 2024. Sliding Windows Are Not the End: Exploring Full Ranking with Long-Context Large Language Models. <https://arxiv.org/abs/2412.14574>
- [37] Ye Liu, Kazuma Hashimoto, Yingbo Zhou, Semih Yavuz, Caiming Xiong, and Philip Yu. 2021. Dense Hierarchical Retrieval for Open-domain Question Answering. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 188–200. doi:10.18653/v1/2021.findings-emnlp.19
- [38] Yijun Liu, Jinzheng Yu, Yang Xu, Zhongyang Li, and Qingfu Zhu. 2025. A Survey on Transformer Context Extension: Approaches and Evaluation. arXiv:2503.13299 [cs.CL] <https://arxiv.org/abs/2503.13299>
- [39] Xueguang Ma, Xinyu Zhang, Ronak Pradeep, and Jimmy Lin. 2023. Zero-Shot Listwise Document Reranking with a Large Language Model. <https://arxiv.org/abs/2305.02156>
- [40] Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veysseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. 2023. Recent advances in natural language processing via large pre-trained language models: A survey. *Comput. Surveys* 56, 2 (2023), 1–40.
- [41] Sheshera Mysore, Tim O’Gorman, Andrew McCallum, and Hamed Zamani. 2021. CSFCube – A Test Collection of Computer Science Research Articles for Faceted Query by Example. <https://arxiv.org/abs/2103.12906>
- [42] Jianmo Ni, Chen Qu, Jing Lu, Zhuoyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, and Yinfei Yang. 2022. Large Dual Encoders Are Generalizable Retrievers. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 9844–9855. doi:10.18653/v1/2022.emnlp-main.669
- [43] Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2018. A Survey on Open Information Extraction. In *Proceedings of the 27th International Conference on Computational Linguistics*, Emily M. Bender, Leon Derczynski, and Pierre Isabelle (Eds.). Association for Computational Linguistics, Santa Fe, New Mexico, USA, 3866–3878. <https://aclanthology.org/C18-1326>
- [44] Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document Ranking with a Pretrained Sequence-to-Sequence Model. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 708–718. doi:10.18653/v1/2020.findings-emnlp.63
- [45] Rodrigo Nogueira and Jimmy Lin. 2019. From doc2query to docTTTT-Tquery. https://cs.uwaterloo.ca/~jimmylin/publications/Nogueira_Lin_2019_docTTTTQuery-v2.pdf. Micropublication; version 2, Dec 7, 2019.
- [46] Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. Multi-Stage Document Ranking with BERT. <https://arxiv.org/abs/1910.14424>
- [47] OpenAI. ; Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, et al. 2024. GPT-4o System Card. arXiv:2410.21276 [cs.CL] <https://arxiv.org/abs/2410.21276>
- [48] OpenAI. 2024. Introducing GPT-4.1 in the API. <https://openai.com/index/gpt-4-1/>. Accessed: 2025-05-19.
- [49] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, et al. 2023. GPT-4 Technical Report. <https://arxiv.org/abs/2303.08774>
- [50] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [51] Saurav Pawar, SM Tonmoy, SM Zaman, Vinija Jain, Aman Chadha, and Amitava Das. 2024. The What, Why, and How of Context Length Extension Techniques in Large Language Models—A Detailed Survey.
- [52] Ronak Pradeep, Sahel Sharifmoghaddam, and Jimmy Lin. 2023. RankVicuna: Zero-Shot Listwise Document Reranking with Open-Source Large Language Models. <https://arxiv.org/abs/2309.15088>
- [53] Ronak Pradeep, Sahel Sharifmoghaddam, and Jimmy Lin. 2023. RankZephyr: Effective and Robust Zero-Shot Listwise Reranking is a Breeze! <https://arxiv.org/abs/2312.02724>
- [54] Qwen. ; An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, et al. 2024. Qwen2.5 Technical Report. <https://arxiv.org/abs/2412.15115>
- [55] Revanth Gangi Reddy, Pradeep Dasigi, Md Arafat Sultan, Arman Cohan, Avirup Sil, Heng Ji, and Hannaneh Hajishirzi. 2023. ReFIT: Relevance Feedback from a Reranker during Inference. <https://arxiv.org/abs/2305.11744>
- [56] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 3982–3992. doi:10.18653/v1/D19-1410
- [57] Ruiyang Ren, Shangwen Lv, Yingqi Qu, Jing Liu, Wayne Xin Zhao, QiaoQiao She, Hua Wu, Haifeng Wang, and Ji-Rong Wen. 2021. PAIR: Leveraging Passage-Centric Similarity Relation for Improving Dense Passage Retrieval. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 2173–2183. doi:10.18653/v1/2021.findings-acl.191
- [58] Ruiyang Ren, Yuhao Wang, Kun Zhou, Wayne Xin Zhao, Wenjie Wang, Jing Liu, Ji-Rong Wen, and Tat-Seng Chua. 2024. Self-Calibrated Listwise Reranking with Large Language Models. <https://arxiv.org/abs/2411.04602>
- [59] Ruiyang Ren, Wayne Xin Zhao, Jing Liu, Hua Wu, Ji-Rong Wen, and Haifeng Wang. 2023. TOME: A Two-stage Approach for Model-based Retrieval. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 6102–6114. doi:10.18653/v1/2023.acl-long.336
- [60] Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval* 3, 4 (2009), 333–389.
- [61] Stephen E. Robertson and Karen Spärck Jones. 1976. Relevance weighting of search terms. *J. Am. Soc. Inf. Sci.* 27 (1976), 129–146. <https://api.semanticscholar.org/CorpusID:45186038>
- [62] Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. Automatic Keyword Extraction from Individual Documents. In *Text Mining: Applications and Theory*, Michael W. Berry and Jacob Kogan (Eds.). John Wiley & Sons, Ltd, Chichester, UK, 1–20. doi:10.1002/9780470689646.ch1
- [63] Devendra Sachan, Mike Lewis, Mandar Joshi, Armen Aghajanyan, Wen-tau Yih, Joelle Pineau, and Luke Zettlemoyer. 2022. Improving Passage Retrieval with Zero-Shot Question Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 3781–3797. doi:10.18653/v1/2022.emnlp-main.249
- [64] Gerard Salton, Anita Wong, and Chung-Shu Yang. 1975. A vector space model for automatic indexing. *Commun. ACM* 18 (1975), 613–620. <https://api.semanticscholar.org/CorpusID:6473756>
- [65] Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. Simple Entity-Centric Questions Challenge Dense Retrievers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6138–6148. doi:10.18653/v1/2021.emnlp-main.496
- [66] Amanpreet Singh, Mike D’Arcy, Arman Cohan, Doug Downey, and Sergey Feldman. 2023. SciRepEval: A Multi-Format Benchmark for Scientific Document Representations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 5548–5566. doi:10.18653/v1/2023.emnlp-main.338
- [67] Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. 2025. OpenAI GPT-5 System Card. arXiv:2601.03267 [cs.CL] <https://arxiv.org/abs/2601.03267>
- [68] Amit Singhal et al. 2001. Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.* 24, 4 (2001), 35–43.
- [69] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters.
- [70] Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. Is ChatGPT Good at Search? Investigating Large Language Models as Re-Ranking Agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 14918–14937. doi:10.18653/v1/2023.emnlp-main.923
- [71] Xiaofei Sun, Xiaoya Li, Jiwei Li, Fei Wu, Shangwei Guo, Tianwei Zhang, and Guoyin Wang. 2023. Text Classification via Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 8990–9005. doi:10.18653/v1/2023.findings-emnlp.603
- [72] Raphael Tang, Crystina Zhang, Xueguang Ma, Jimmy Lin, and Ferhan Ture. 2024. Found in the Middle: Permutation Self-Consistency Improves Listwise Ranking in Large Language Models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 2327–2340. <https://aclanthology.org/2024.naacl-long.129>
- [73] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. <https://arxiv.org/abs/2104.08663>
- [74] Runchu Tian, Yanghao Li, Yuepeng Fu, Siyang Deng, Qinyu Luo, Cheng Qian, Shuo Wang, Xin Cong, Zhong Zhang, Yesai Wu, et al. 2025. Distance between

- Relevant Information Pieces Causes Bias in Long-Context LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 521–533. <https://aclanthology.org/2025.findings-acl.28/>
- [75] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruiti Bhosale, et al. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. <https://arxiv.org/abs/2307.09288>
- [76] Mohamed Ali Trabelsi, Zhiyu Chen, Brian D. Davison, and Jeff Heflin. 2021. Neural ranking models for document retrieval. *Information Retrieval Journal* 24 (2021), 400–444. <https://api.semanticscholar.org/CorpusID:232035744>
- [77] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*. Curran Associates, Inc., Red Hook, NY, USA, 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [78] Ellen Voorhees, Tasmeer Alam, Steven Bedrick, Dina Demner-Fushman, William R. Hersh, Kyle Lo, Kirk Roberts, Ian Soboroff, and Lucy Lu Wang. 2021. TREC-COVID: constructing a pandemic information retrieval test collection. *SIGIR Forum* 54, 1, Article 1 (Feb. 2021), 12 pages. doi:10.1145/3451964.3451965
- [79] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7534–7550. doi:10.18653/v1/2020.emnlp-main.609
- [80] Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text Embeddings by Weakly-Supervised Contrastive Pre-training. <https://arxiv.org/abs/2212.03533>
- [81] Howard D White, H Cooper, LV Hedges, et al. 2009. Scientific communication and literature retrieval. *The handbook of research synthesis and meta-analysis 2* (2009), 51–71.
- [82] Daya C Wimalasuriya and Dejing Dou. 2010. Ontology-based information extraction: An introduction and a survey of current approaches. 306–323 pages.
- [83] Shijie Xia, Yiwei Qin, Xuefeng Li, Yan Ma, Run-Ze Fan, Steffi Chern, Haoyang Zou, Fan Zhou, Xiangkun Hu, Jiahe Jin, et al. 2025. Generative AI Act II: Test Time Scaling Drives Cognition Engineering. <https://arxiv.org/abs/2504.13828>
- [84] Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, et al. 2024. Effective Long-Context Scaling of Foundation Models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 4643–4663. <https://aclanthology.org/2024.naacl-long.260>
- [85] Xueqiang Xu, Jinfeng Xiao, James Barry, Mohab Elkaref, Jiaru Zou, Pengcheng Jiang, Yunyi Zhang, Max Giammona, Geeth de Mel, and Jiawei Han. 2025. Zero-Shot Open-Schema Entity Structure Discovery. arXiv:2506.04458 [cs.CL] <https://arxiv.org/abs/2506.04458>
- [86] Andrew Yates, Rodrigo Frassetto Nogueira, and Jimmy Lin. 2021. Pretrained Transformers for Text Ranking: BERT and Beyond. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11–15, 2021*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.). ACM, New York, NY, USA, 2666–2668. doi:10.1145/3404835.3462812
- [87] Chenhan Yuan, Qianqian Xie, and Sophia Ananiadou. 2023. Zero-shot Temporal Relation Extraction with ChatGPT. In *Proceedings of the 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks*, Dina Demner-Fushman, Sophia Ananiadou, and Kevin Cohen (Eds.). Association for Computational Linguistics, Toronto, Canada, 92–102. doi:10.18653/v1/2023.bionlp-1.7
- [88] Kai Zhang, Bernal Jimenez Gutierrez, and Yu Su. 2023. Aligning Instruction Tasks Unlocks Large Language Models as Zero-Shot Relation Extractors. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 794–812. doi:10.18653/v1/2023.findings-acl.50
- [89] Yunyi Zhang, Ruozhen Yang, Xueqiang Xu, Rui Li, Jinfeng Xiao, Jiaming Shen, and Jiawei Han. 2025. TELEClass: Taxonomy Enrichment and LLM-Enhanced Hierarchical Text Classification with Minimal Supervision. In *Proceedings of the ACM on Web Conference 2025 (Sydney NSW, Australia) (WWW '25)*. Association for Computing Machinery, New York, NY, USA, 2032–2042. doi:10.1145/3696410.3714940
- [90] Xuanhe Zhou, Guoliang Li, and Zhiyuan Liu. 2023. LLM As DBA. arXiv:2308.05481 [cs.DB] <https://arxiv.org/abs/2308.05481>
- [91] Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and Michael Bendersky. 2023. RankT5: Fine-Tuning T5 for Text Ranking with Ranking Losses. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23–27, 2023*, Hsin-Hsi Chen, Wei-Jou (Edward) Duh, Hen-Hsen Huang, Makoto P. Kato, Josiane Mothe, and Barbara Pobleto (Eds.). ACM, New York, NY, USA, 2308–2313. doi:10.1145/3539618.3592047