

Document Version

Final published version

Licence

CC BY

Citation (APA)

Yoon, S., Kim, J., Kwon, D., Anand, A., & Hwang, S. W. (2026). On Listwise Reranking for Corpus Feedback. In *WSDM 2026 - Proceedings of the 19th ACM International Conference on Web Search and Data Mining* (pp. 1273-1277). (WSDM 2026 - Proceedings of the 19th ACM International Conference on Web Search and Data Mining). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3773966.3779404>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



 Latest updates: <https://dl.acm.org/doi/10.1145/3773966.3779404>

SHORT-PAPER

On Listwise Reranking for Corpus Feedback

SOYOUNG YOON, Seoul National University, Seoul, South Korea

JONGHO KIM, Seoul National University, Seoul, South Korea

DAEYONG KWON, Seoul National University, Seoul, South Korea

AVISHEK ANAND, Delft University of Technology, Delft, Zuid-Holland, Netherlands

SEUNG-WON HWANG, Seoul National University, Seoul, South Korea

Open Access Support provided by:

Seoul National University

Delft University of Technology



PDF Download
3773966.3779404.pdf
07 April 2026
Total Citations: 0
Total Downloads: 257

Published: 21 February 2026

Citation in BibTeX format

WSDM '26: The Nineteenth ACM
International Conference on Web Search
and Data Mining
February 22 - 26, 2026
ID, Boise, USA

Conference Sponsors:

SIGKDD
SIGWEB
SIGIR
SIGMOD

On Listwise Reranking for Corpus Feedback

Soyoung Yoon
soyoung.yoon@snu.ac.kr
Seoul National University
Seoul, South Korea

Jongho Kim
jongh97@snu.ac.kr
Seoul National University
Seoul, South Korea

Daeyong Kwon
gwondaeyong63@gmail.com
Seoul National University
Seoul, South Korea

Avishek Anand
Avishek.Anand@tudelft.nl
Delft University of Technology
Delft, Netherlands

Seung-won Hwang*
seungwonh@snu.ac.kr
Seoul National University
Seoul, South Korea

Abstract

Reranker improves retrieval performance by capturing document interactions. At one extreme, graph-aware adaptive retrieval (GAR) represents an information-rich regime, requiring a pre-computed document similarity graph in reranking. However, as such graphs are often unavailable, or incur quadratic memory costs even when available, graph-free rerankers leverage large language model (LLM) calls to achieve competitive performance. We introduce L2G, a novel framework that implicitly induces document graphs from listwise reranker logs. By converting reranker signals into a graph structure, L2G enables scalable graph-based retrieval without the overhead of explicit graph computation. Results on the TREC-DL and BEIR subset show that L2G matches the effectiveness of oracle-based graph methods, while incurring zero additional LLM calls.¹

CCS Concepts

• **Information systems** → **Information retrieval**; *Retrieval models and ranking*; Language models; Learning to rank.

Keywords

listwise reranker, adaptive retrieval, corpus graph

ACM Reference Format:

Soyoung Yoon, Jongho Kim, Daeyong Kwon, Avishek Anand, and Seung-won Hwang. 2026. On Listwise Reranking for Corpus Feedback. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3773966.3779404>

1 Introduction

Reranking captures document interactions as a form of corpus feedback to improve retrieval performance [17, 19, 20]. To illustrate, listwise rerankers reflect local interactions by reranking documents within a window. For a large input, this window shifts forward to compare high-ranking documents with the next unseen documents, until the window reaches the other end of the list. However, such

*Corresponding author

¹<https://github.com/soyoung97/l2g-submit>



This work is licensed under a Creative Commons Attribution 4.0 International License. *WSDM '26, Boise, ID, USA*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2292-9/2026/02
<https://doi.org/10.1145/3773966.3779404>

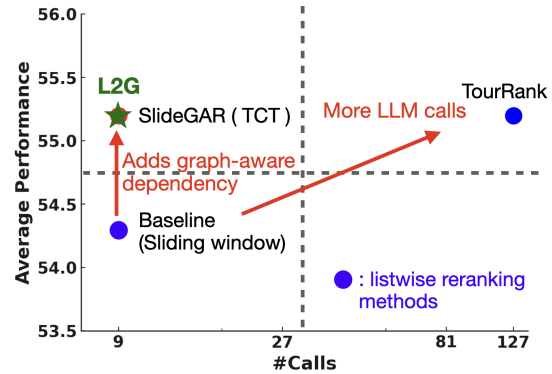


Figure 1: A cost-performance trade-off of rerankers (nDCG@10; see arxiv version of our paper (link) full results).

baselines, due to serialized window shifts, fail to capture interaction feedback across non-adjacent windows.

One extreme of enriching corpus feedback is SlideGAR [20], which assumes a graph capturing all pairwise document interactions, then leverages the neighborhood information beyond the window from the graphs. While effective, this approach requires full corpus graph access to construct them, and incurs substantial cost from repeated bi-encoder calls (e.g., TCT [12]), as well as a prohibitive $O(N^2)$ memory footprint for a corpus of size N . In another extreme, graph-free rerankers such as TourRank [2] conduct multiple rounds of tournaments across windows. Despite delivering a strong performance with SlideGAR without requiring explicit graphs, TourRank incurs heavy computation due to the large number of LLM calls required to remain competitive.

Figure 1 summarizes this tradeoff landscape by plotting average effectiveness against LLM calls. SlideGAR lies in the upper-left: strong performance with sliding-window baseline budget (9 calls), but requires additional compute and memory to construct the corpus graph. Another extreme in the upper-right is tournament-style listwise reranker (e.g., TourRank) that is graph-free, improves performance but requires far more calls (about 127).

Our contribution is achieving high performance at low call budgets (upper-left quadrant). Observing that listwise methods already expose local relations, we ask: can we reuse these signals across queries to approximate a document graph, without any additional LLM calls or requiring a doc-doc graph or retriever?

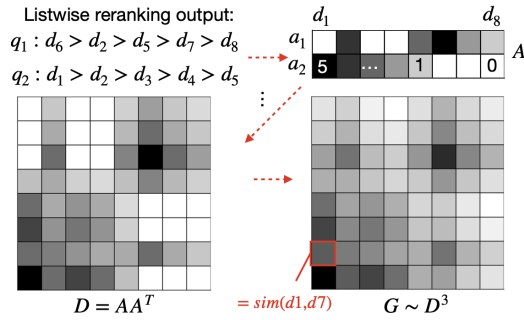


Figure 2: Illustration of building G , a proxy of doc-doc graph, from listwise signals for L2G. Darker cells indicate higher pairwise relevance, and vice versa.

This motivates us to propose Listwise-to-Graph (L2G), requiring *no* additional LLM or retriever calls while restoring SlideGAR performance. Our contributions are threefold: (i) we provide a yet effective perspective that frames listwise reranking as a source of implicit corpus feedback; (ii) we introduce L2G, which reconstructs an explicit doc-doc graph directly from ranking outputs (no separate doc-doc retriever); and (iii) we demonstrate SlideGAR performance with lower cost, making graph-free adaptive retrieval practical even in dynamic or resource-constrained settings.

2 Related Works

Graph-based reranker Graph-aware methods such as SlideGAR [20] exploit *explicit* doc-doc relations to guide adaptive reranking. While they achieve strong quality, they requiring pre-built corpus graphs with huge memory and pre-processing costs. Our work keeps the call budget low without requiring a doc-doc retriever by *reconstructing* a sparse interaction graph directly from past listwise rankings and reusing it in SlideGAR.

Graph-free reranker However, a corpus or graph is not always available. Then, advanced reranking for ranked results can be used to overcome the gap, albeit at an additional cost [21]. We focus on listwise rerankers [19] that jointly reason over candidate sets. They often rely on sliding windows or tournament-style comparisons. RankZephyr [19] is a strong fine-tuned listwise ranker; ReaRank [25] emphasizes reasoning ability with more inference time; TourRank [2] improves effectiveness at the cost of more LLM calls. Contemporary advanced ranking approaches [13–15] optimize in ICL/demo selection, collaborate small/large rankers with order adjustment, or expand context windows. Our distinction is that we frame listwise methods as *implicitly* inducing local interaction graphs among candidates with zero LLM-call overhead.

Query-stream reuse Transforming an entire corpus into a graph is a costly process, while previous work [1] argued is unnecessary. Our work shares this motivation, leveraging batch and stream processing [4–6, 16] to demonstrate gains from shared computation and intermediate results across queries. Our incremental L2G updates operationalize this principle for reranking: listwise signals accumulate into a sparse graph that stays effective under order perturbations.

3 Method: Listwise signals to Graph (L2G)

We study treating reranker outputs as document relationship signals and aggregate document pairs and their co-occurrence with queries, which induces a first-order similarity matrix that serves as a proxy for an implicit graph over documents. Fig. 2 visualizes the building process. Let $Q = \{q_1, \dots, q_m\}$ be a stream of user queries. For each q_i a listwise reranker receives a candidate set $C_i = \{d_{i,1}, \dots, d_{i,k}\}$ and returns a total order $\pi_i : C_i \rightarrow [k]$, where $\pi_i(d)$ is the rank of document d for query q_i (one-dimensional). We denote by $\mathcal{D} = \bigcup_i C_i$ the (ever-growing) global document pool.

3.1 From Rankings to a First-order Graph

First, we introduce a method for converting each ranked list for q_i into a score vector. When accumulated, this forms a vectorized matrix for multiple queries, which becomes the basis for the underlying doc-doc graph structures. We map each ranked list into a dense *document score vector* $\mathbf{a}_i \in \mathbb{R}^{|\mathcal{D}|}$ with $[\mathbf{a}_i]_d = (k - \pi_i(d) + 1)$ if $d \in C_i$, and 0 otherwise. In other words, the scores get assigned from highest to lowest depending on their rank for each query. We approximate pairwise document affinity as $\mathbf{D}^{(1)} = \mathbf{A}\mathbf{A}^T$, where $\mathbf{D}_{d_1, d_2}^{(1)} = \sum_i [\mathbf{a}_i]_{d_1} [\mathbf{a}_i]_{d_2}$ accumulates evidence over all queries in which *both* documents co-occur, and treat this $\mathbf{D}^{(1)}$ as an approximation of the pairwise document similarity matrix. Related works suggest that greater query diversity leads to more accurate approximations [9, 24].

3.2 Handling Sparsity and Higher-Order Connectivity

The first-order similarity $\mathbf{D}^{(1)}$ is often sparse due to the limited overlap between candidate sets across queries. However, multi-hop connections—e.g., two documents may not co-occur but may each co-occur with a shared third document—can capture higher-order similarity structure. Formally, we extend the first-order graph by propagating affinity via k -hop random walks: $\mathbf{D}^{(k)} = \mathbf{D}\mathbf{D} \cdots \mathbf{D}$ (k times), analogous to power iteration in PageRank [18]. Larger k increases recall but risks *rank collapse* towards a uniform stationary distribution. We therefore (i) renormalize rows to unit k_1 norm after each multiplication and (ii) cap $k \leq 3$ in practice.

Query-conditioned locality for robustness. While our L2G graph is defined over the full corpus, its induced edges are intentionally *sparse* compared to fully precomputed pairwise graphs. To convert this sparsity into a precision advantage and avoid spurious hub expansion, L2G queries the graph *locally*: relevant candidates are restricted to the first-stage retriever’s top- c pool (we use $c \in \{100, 1000\}$). This query-conditioned locality reduces sensitivity to noisy long-range links when applying k -hop propagation, and improves by constraining exploration to a calibrated, high-precision pool.

3.3 Handling Frequency Bias

Still, the similarity scores in the matrix \mathbf{D} can be skewed by document popularity rather than true semantic relevance. For example, documents that appear frequently across many queries will accumulate high similarity scores simply due to repeated co-occurrence,

not because they are genuinely related. Therefore, we borrow a well-established idea from information retrieval: **inverse document frequency (IDF) re-weighting**. For each document d , we divide its score vector by $\log(1 + \text{df}(d))$, where df counts the number of queries that retrieved document d .

3.4 Online Updates for Unseen Documents

The challenge in maintaining the document matrix is that in practice, the document pool grows dynamically as new queries introduce previously unseen documents, making it impractical to recompute the entire similarity matrix $\mathbf{D}^{(1)}$ from scratch each time. To support pairwise similarity for uncovered or newly added documents, As a solution, we develop an incremental update strategy that efficiently extends the graph structure without full recomputation. Formally, the document pool approximates the document universe unknown *a priori*. As a dynamic approximation, when a new batch of queries adds a set $\Delta\mathcal{D}$ of unseen documents, we update $\mathbf{A}_{\text{new}} = \begin{bmatrix} \mathbf{A} \\ \mathbf{A}_\Delta \end{bmatrix}$ and maintain $\mathbf{D}^{(1)}$ in block form instead of re-computing from scratch: $\mathbf{D}_{\text{new}}^{(1)} = \begin{bmatrix} \mathbf{D}^{(1)} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{bmatrix}$, where $\mathbf{B} = \mathbf{A}\mathbf{A}_\Delta^\top$ and $\mathbf{C} = \mathbf{A}_\Delta\mathbf{A}_\Delta^\top$. This costs $\mathcal{O}(|\mathcal{D}| |\Delta\mathcal{D}|)$ rather than $\mathcal{O}(|\mathcal{D}_{\text{new}}|^2)$. Since $|\Delta\mathcal{D}| \ll |\mathcal{D}|$ in typical streaming scenarios, this provides substantial computational savings.

3.5 Complexity and Memory Footprint

A key advantage of L2G is its elimination of expensive doc-doc similarity computations that plague existing graph-based methods. Unlike approaches such as SlideGAR-TCT that require explicit bi-encoder calls for every document pair, L2G introduces no additional LLM calls for doc-doc similarity and downstream modules simply read the needed relations directly from the on-the-fly graph G .

Operationally, the method requires expansion of D upon the arrival of a new ranked list, and applies three hops of propagation to form $\mathbf{D}^{(3)}$ when a new query’s ranking arrives. In terms of space, L2G stores only the sparse graph G (plus minimal ID maps), never materializing doc–doc matrices or document embeddings while entirely avoiding the need for any GPU-dependent doc-doc retriever. Direct comparison of L2G over SlideGAR-TCT under a matched setup is further discussed at Sec. 4.3.

4 Results

Our empirical study addresses three questions. **1. Effectiveness:** can L2G, built only from past listwise rankings, recover doc-doc structure closely enough to match graph-aware oracles such as doc-doc affinity and full-corpus SlideGAR? **2. Efficiency and cost:** does L2G keep the number of LLM calls fixed at the sliding-window budget while offering favorable latency and memory scaling compared to graph-based baselines? **3. Robustness and generality:** is L2G stable under variations in query order and coverage, and does it transfer across different rerankers and first-stage retrievers?

4.1 Experimental Setup

Datasets & Evaluation. We evaluate on MS MARCO TREC [3] DL’19/’20 (43/54 queries) and a subset of BEIR [23] collections (TREC-COVID, Signal, TREC-NEWS, and Touche) with query size

less than 100, and report $n\text{DCG}@10$ [8] under reranking budget c ($c=100$ or 1000) rounded to the nearest tenth, mostly prepared with the pre-built Pyserini index [10].

Retrieval & Ranking Models. We use BM25 [22] and extend to Contriever [7] as the first-stage retriever. Listwise rerankers include *RankZephyr* [19] and *textitReaRank* [25], operating with sliding windows (size=20, step=10) that output permutations without explicit scores.

Baselines. We compare to: (i) non-adaptive *Sliding Window*, and 3 different (ii) Graph-based *SlideGAR-TCT* baselines. (1) *full-corpus* assumes oracle access to a pre-computed TCT [11] graph over the entire collection. (2) *doc–doc affinity* is a more practical baseline, which mimics the online graph construction scenario just like ours (Sec. 3.2), limiting the graph neighbors to a *local* graph built only from the query’s retrieved top- c candidates, e.g., top-100. Lastly, (iii) *Random* randomizes neighbor ordering from local top- c candidates, verifying that L2G’s gains are not due to chance.

4.2 Effectiveness

Table 1 summarizes $n\text{DCG}@10$ on MS MARCO DL’19/’20 and BEIR subsets with BM25 first-stage pools ($c \in 100, 1000$). First, we observe that L2G consistently outperforms the sanity-check baseline of a random doc-doc ranker (SlideGAR-Random) across all setups. Under Top-100, L2G effectively ties the graph-based oracles—SlideGAR-TCT with doc–doc affinity and the full-corpus variant—within ≈ 0.1 $n\text{DCG}$ on average, all at the same LLM-call budget as the sliding-window reranker. Moving to Top-1000, L2G even *wins* over full-corpus SlideGAR-TCT while relying only on local interaction signals. We hypothesize this increased effectiveness due to the larger top- c overlap reported in the table, which lets L2G reuse cross-query evidence to enrich its induced graph. This indicates that once overlap is sufficient, listwise-induced graphs recover—and in some cases exceed—the benefits of explicit corpus graphs without building them.

4.3 Efficiency / Cost

We evaluate only the adaptive stage (excluding first-stage retrieval and any LLM calls), reporting per-query latency and peak memory. For L2G, **latency** is measured from the moment a ranked list for query q arrives: we measure the per-query cumulative time to expand the document pool, form $D=AA^\top$, and yield D^3 . We compare L2G with SlideGAR-TCT (doc-affinity), as it is a strong baseline that also constructs a graph online. As summarized in Tab. 2, SlideGAR-TCT averages 0.427 s per-query, whereas L2G requires no precomputation and remains well below this even as the maintained graph grows. On **memory**, SlideGAR-TCT must load a bi-encoder (about 419 MB), store top-100 doc-doc affinity graph dense embeddings (once per query in the first window), and maintain a neighbor dictionary (doc IDs with float scores). L2G avoids the bi-encoder and dense vectors entirely: its footprint is a single sparse graph over documents actually observed along with the minimal docID maps, scaling with unique candidate documents. Full-corpus graphs are not comparable on a per-query basis but are much heavier than the online ones; e.g., for DL’19 the full-corpus TCT dump occupies ~ 3.1 GB on disk, and even the smaller Touche

| | BM25 Top-100 | | | | | | | BM25 Top-1000 | | | | | | |
|--|--------------|-------------|-------------|-------------|-------------|-------------|-------------|---------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | Covid | Sig. | News | Tou. | DL'19 | DL'20 | Avg | Covid | Sig. | News | Tou. | DL'19 | DL'20 | Avg |
| Query count | 50 | 97 | 57 | 49 | 43 | 54 | – | 50 | 97 | 57 | 49 | 43 | 54 | – |
| Top-c Overlap (%) | 10.4 | 0.8 | 3.3 | 1.3 | 0.1 | 0.1 | – | 39.7 | 3.1 | 13.6 | 15.3 | 0.5 | 3.1 | – |
| Total corpus size | 171k | 2.9M | 595k | 383k | 8.8M | 8.8M | – | 171k | 2.9M | 595k | 383k | 8.8M | 8.8M | – |
| <i>(0) baseline – without doc2doc graphs</i> | | | | | | | | | | | | | | |
| BM25 | 59.5 | 33.0 | 39.5 | 44.2 | 50.6 | 48.0 | 45.4 | 59.5 | 33.0 | 39.5 | 44.2 | 50.6 | 48.0 | 45.4 |
| Sliding Window | 84.1 | 32.0 | 52.3 | 32.4 | 74.0 | 70.2 | 57.5 | 80.7 | 28.9 | 51.0 | 30.9 | 75.1 | 78.8 | 57.6 |
| <i>(1) with doc2doc graphs</i> | | | | | | | | | | | | | | |
| SlideGAR-TCT (full corpus) | 80.1 | 31.0 | 51.7 | 34.9 | 74.2 | 79.3 | 58.5 | 86.2 | 28.9 | 51.1 | 34.4 | 75.4 | 79.8 | 59.3 |
| SlideGAR-TCT (doc. affinity) | 83.0 | 31.1 | 53.0 | 37.5 | 74.2 | 71.1 | 58.3 | 85.0 | 29.5 | 52.7 | 32.7 | 75.1 | 80.0 | 59.2 |
| SlideGAR-Random | 82.5 | 31.2 | 54.1 | 35.4 | 72.7 | 71.5 | 57.9 | 83.8 | 30.3 | 54.6 | 36.4 | 73.4 | 76.3 | 59.1 |
| L2G | 84.2 | 31.8 | 53.4 | 37.4 | 72.3 | 71.2 | 58.4 | 85.4 | 29.3 | 55.8 | 35.2 | 74.7 | 76.6 | 59.5 |

Table 1: Dataset stats and NDCG@10 (%) with BM25 first-stage under two pool sizes (Top-100 vs. Top-1000), using the RankZephyr-7B model. Benchmarks indicate TREC-Covid, Signal, News, Touche, TREC-DL19, and DL20, respectively. All variants use the same number of LLM calls as the Sliding Window.

| | SlideGAR-TCT (doc. affinity) | L2G (ours) |
|---------------------|---------------------------------|-------------------------|
| Extra store | doc-doc embeddings | No pre-built embeddings |
| Latency / query (s) | 0.427 | 0.103 |
| Peak Storage (MB) | 0.862 | 0.855 |
| Peak VRAM (MB) | 418.73 + α | None |

Table 2: Efficiency comparison between SlideGAR-TCT and L2G, averaged for all reported BEIR subsets. Storage is measured as the in-memory footprint of the data structure (in bytes). Smaller is better; L2G achieves consistently lower latency and memory without prebuilt corpus embeddings.

| Method / Setting | Covid | Sig. | News | Tou. | Avg |
|---|-------|------|------|------|-------------|
| (1) Perturbing query stream (BM25 top-100, RankZephyr) | | | | | |
| L2G (dataset order) | 84.2 | 31.8 | 53.4 | 37.4 | 51.7 |
| L2G (max overlap) | 83.0 | 31.4 | 53.5 | 37.6 | 51.4 |
| L2G (min overlap) | 83.8 | 31.7 | 53.9 | 36.8 | 51.6 |
| (2) Different first-stage retriever (Contriever top-100, RankZephyr) | | | | | |
| Sliding Window | 70.8 | 26.5 | 49.9 | 30.4 | 44.4 |
| SlideGAR (doc. affinity) | 72.1 | 25.0 | 51.9 | 29.7 | 44.7 |
| SlideGAR-Random | 72.0 | 25.0 | 50.1 | 27.1 | 43.6 |
| L2G | 72.0 | 26.8 | 50.9 | 28.9 | 44.7 |
| (3) Different reranking model (Contriever top-100, ReaRank) | | | | | |
| Sliding Window | 71.6 | 27.0 | 50.5 | 24.4 | 43.4 |
| SlideGAR (doc. affinity) | 71.6 | 26.9 | 50.4 | 23.9 | 43.2 |
| SlideGAR-Random | 74.4 | 27.2 | 51.3 | 23.6 | 44.1 |
| L2G | 74.3 | 27.3 | 52.4 | 24.7 | 44.7 |

Table 3: RQ3: Robustness & Generalizability (NDCG@10, %).

collection takes >1 hour to dump on our machine, while many scenarios do not even provide full-corpus access. Taken together, L2G delivers doc-affinity-level effectiveness at a fraction of the runtime and with a smaller footprint. In practice, delaying or batching graph updates (e.g., delay three propagation hops) further reduces constant factors without affecting ranking quality, which we leave as future work.

4.4 Robustness / Generalizability

Ordering robustness. While document overlap between queries is crucial for L2G’s graph construction as discussed in Sec. 4.2, a potential concern with L2G’s streaming approach is sensitivity to query arrival orders. We address this concern by analyzing L2G’s robustness to query order variations within the same benchmark. (i) a **max overlap** order greedily picks the next query whose top-100 overlaps the most with the pool already processed, and (ii) a **min overlap** order greedily minimizes this cumulative reuse. Tab. 3–(1) further shows that what matters is the overall amount of overlap present in the corpus, not the *specific order* in which overlapping queries arrive; although Table 1 suggested that effectiveness increases with higher corpus-level overlap, results confirm that the arrival order itself is *not* directly related to output performance. **Generalizability.** On Tab. 3–(2) and (3), we observe that L2G dominates over SlideGAR-TCT (doc-affinity) even when swapping first-stage retrievers (BM25 → Contriever) and base listwise rankers (RankZephyr → ReaRank). Overall, L2G behaves as a reusable graph prior that transfers across ranking stacks.

5 Conclusion

This work explores listwise rerankers as an alternative to graph-based adaptive retrieval when a graph is unavailable or costly to build. We demonstrate how the graph structure can be constructed from the listwise reranking results and reused across queries. We validate that our L2G performs on par with graph-based oracles. We are first to frame listwise reranking as an alternative form of adaptive retrieval, bridging the gap between graph-based and graph-free reranking.

Acknowledgments

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [NO.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)], and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2024-00414981),

6 Ethical Considerations

While the L2G is far more efficient than SlideGAR, additional requirements for graph construction may still limit accessibility for resource-constrained organizations. The efficiency of L2G can be further optimized by tuning graph update intervals at the cost of performance, or through implementation of the GPU version of L2G. Our simple implementation leaves additional speed and memory gains as future work, which could help address these accessibility concerns.

References

- [1] Shengyuan Chen, Chuang Zhou, Zheng Yuan, Qinggang Zhang, Zeyang Cui, Hao Chen, Yilin Xiao, Jiannong Cao, and Xiao Huang. 2025. You Don't Need Pre-built Graphs for RAG: Retrieval Augmented Generation with Adaptive Reasoning Structures. arXiv:2508.06105 [cs.CL] <https://arxiv.org/abs/2508.06105>
- [2] Yiqun Chen, Qi Liu, Yi Zhang, Weiwei Sun, Xinyu Ma, Wei Yang, Daiting Shi, Jiaxin Mao, and Dawei Yin. 2025. TourRank: Utilizing Large Language Models for Documents Ranking with a Tournament-Inspired Strategy. arXiv:2406.11678 [cs.IR] <https://arxiv.org/abs/2406.11678>
- [3] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the TREC 2019 deep learning track. *arXiv preprint arXiv:2003.07820* (2020).
- [4] Shuai Ding, Josh Attenberg, Ricardo Baeza-Yates, and Torsten Suel. 2011. Batch query processing for web search engines. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining (Hong Kong, China) (WSDM '11)*. Association for Computing Machinery, New York, NY, USA, 137–146. doi:10.1145/1935826.1935858
- [5] Mehrad Eslami, Vahid Mahmoodian, Iman Dayarian, Hadi Charkhgard, and Yicheng Tu. 2020. Query batching optimization in database systems. *Computers & Operations Research* 121 (2020), 104983.
- [6] Mehrad Eslami, Yicheng Tu, Hadi Charkhgard, Zichen Xu, and Jiacheng Liu. 2019. PsiDB: A framework for batched query processing and optimization. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 6046–6048.
- [7] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. arXiv:2112.09118 [cs.IR] <https://arxiv.org/abs/2112.09118>
- [8] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated Gain-Based Evaluation of IR Techniques. *ACM Trans. Inf. Syst.* 20, 4 (oct 2002), 422–446. doi:10.1145/582415.582418
- [9] Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. 2012. Sampling methods for the Nyström method. *J. Mach. Learn. Res.* 13, null (April 2012), 981–1006.
- [10] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A Python Toolkit for Reproducible Information Retrieval Research with Sparse and Dense Representations. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*. 2356–2362.
- [11] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2020. Distilling Dense Representations for Ranking using Tightly-Coupled Teachers. arXiv:2010.11386 [cs.IR] <https://arxiv.org/abs/2010.11386>
- [12] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-Batch Negatives for Knowledge Distillation with Tightly-Coupled Teachers for Dense Retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP (ReplANLP-2021)*, Anna Rogers, Iacer Calixto, Ivan Vulic, Naomi Saphra, Nora Kassner, Oana-Maria Camburu, Trapit Bansal, and Vered Shwartz (Eds.). Association for Computational Linguistics, Online, 163–173. doi:10.18653/v1/2021.repl4nlp-1.17
- [13] Wenhan Liu, Xinyu Ma, Yutao Zhu, Lixin Su, Shuaiqiang Wang, Dawei Yin, and Zhicheng Dou. 2025. CoRanking: Collaborative Ranking with Small and Large Ranking Agents. arXiv:2503.23427 [cs.CL] <https://arxiv.org/abs/2503.23427>
- [14] 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. arXiv:2412.14574 [cs.IR] <https://arxiv.org/abs/2412.14574>
- [15] Wenhan Liu, Yutao Zhu, and Zhicheng Dou. 2024. DemoRank: Selecting Effective Demonstrations for Large Language Models in Ranking Task. arXiv:2406.16332 [cs.IR] <https://arxiv.org/abs/2406.16332>
- [16] Joel Mackenzie and Alistair Moffat. 2023. Index-based batch query processing revisited. In *European Conference on Information Retrieval*. Springer, 86–100.
- [17] Shengyu Mao, Yong Jiang, Boli Chen, Xiao Li, Peng Wang, Xinyu Wang, Pengjun Xie, Fei Huang, Huajun Chen, and Ningyu Zhang. 2024. RaFe: Ranking Feedback Improves Query Rewriting for RAG. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 884–901. doi:10.18653/v1/2024.findings-emnlp.49
- [18] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank Citation Ranking : Bringing Order to the Web. In *The Web Conference*.
- [19] Ronak Pradeep, Sahel Sharifymoghaddam, and Jimmy Lin. 2023. RankZephyr: Effective and Robust Zero-Shot Listwise Reranking is a Breeze! arXiv:2312.02724 [cs.IR] <https://arxiv.org/abs/2312.02724>
- [20] Mandeep Rathee, Sean MacAvaney, and Avishek Anand. 2025. Guiding Retrieval using LLM-based Listwise Rerankers. arXiv:2501.09186 [cs.IR] <https://arxiv.org/abs/2501.09186>
- [21] Mandeep Rathee, V Venkatesh, Sean MacAvaney, and Avishek Anand. 2025. Test-time Corpus Feedback: From Retrieval to RAG. arXiv:2508.15437 [cs.IR] <https://arxiv.org/abs/2508.15437>
- [22] 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.
- [23] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. arXiv:2104.08663 [cs.IR] <https://arxiv.org/abs/2104.08663>
- [24] Anders Kirk Uhrenholt, Valentin Charvet, and Bjørn Sand Jensen. 2021. Probabilistic selection of inducing points in sparse Gaussian processes. arXiv:2010.09370 [cs.LG] <https://arxiv.org/abs/2010.09370>
- [25] Le Zhang, Bo Wang, Xipeng Qiu, Siva Reddy, and Aishwarya Agrawal. 2025. REARANK: Reasoning Re-ranking Agent via Reinforcement Learning. arXiv:2505.20046 [cs.IR] <https://arxiv.org/abs/2505.20046>