

# Top-Down Partitioning for Efficient List-Wise Ranking

Andrew Parry<sup>1,\*</sup>, Sean MacAvaney<sup>1</sup> and Debasis Ganguly<sup>1</sup>

<sup>1</sup>University of Glasgow, Glasgow, UK

## Abstract

Large Language Models (LLMs) have significantly impacted many facets of natural language processing and information retrieval. Unlike previous neural approaches, the enlarged context window of these generative models allows for ranking multiple documents at once, commonly called list-wise ranking. However, there are still limits to the number of documents that can be ranked in a single inference of the model, leading to the broad adoption of a sliding window approach to identify the  $k$  most relevant items in a ranked list. We argue that the sliding window approach is not well-suited for list-wise re-ranking because it (1) cannot be parallelized in its current form, (2) wastes compute repeatedly re-scoring the best set of documents as it works its way up the initial ranking, and (3) prioritizes the lowest-ranked items for scoring rather than the highest-ranked items by taking a bottom-up approach. We investigate the precision of recent list-wise rankers within a single inference and find that when few relevant documents are found, it is preferable that these documents be at the start of a model’s context window. Motivated by this finding and the shortcomings of a sliding window approach, we propose a novel algorithm that partitions a ranking to depth  $k$  and processes documents top-down. Unlike sliding window approaches, our algorithm is inherently parallelizable due to the use of a pivot element, which can be compared to documents down to an arbitrary depth concurrently. In doing so, we reduce the number of expected inference calls by around 33% when ranking at depth 100 while matching the performance of prior approaches across multiple strong re-rankers.

## Keywords

Neural Retrieval, Large Language Models, Efficiency, Selection Algorithm

## 1. Introduction

Until recently, neural approaches to ranking entailed the direct comparison of a document<sup>1</sup> to only one or two other distinct texts at a time: either a “point-wise” comparison to the query [1, 2] or a “pair-wise” comparison to a query and one other reference document [3]. Recent developments in Large Language Models (LLMs) improve the natural language understanding of transformer-based models and vastly increase the number of tokens a model can effectively process compared to previous models [4, 5, 6]. These two improvements allow for a “list-wise” ranking approach at inference time, where multiple documents can all be evaluated simultaneously. The key benefit of such a list-wise approach is that the ranker is directly

---

Woodstock’22: Symposium on the irreproducible science, June 07–11, 2022, Woodstock, NY

\*Corresponding author.

✉ a.parry.1@research.gla.ac.uk (A. Parry); Sean.MacAvaney@glasgow.ac.uk (S. MacAvaney);

Debasis.Ganguly@glasgow.ac.uk (D. Ganguly)

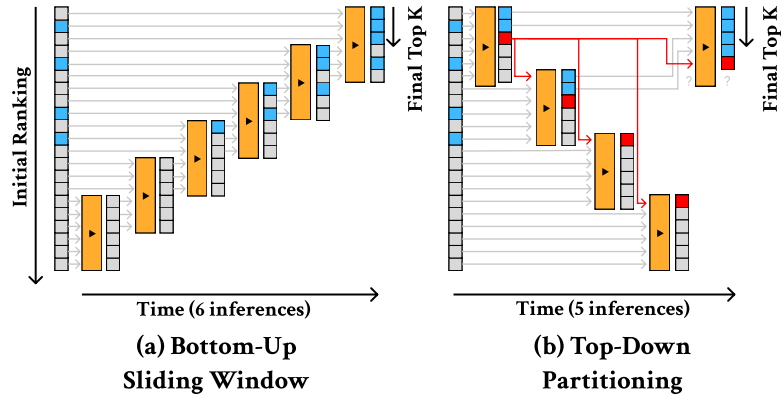
🌐 <https://parry-parry.github.io/> (A. Parry); <https://macavaney.us/> (S. MacAvaney); <https://gdebasis.github.io/>

(D. Ganguly)

🆔 0000-0001-5446-8328 (A. Parry); 0000-0002-8914-2659 (S. MacAvaney); 0000-0003-0050-7138 (D. Ganguly)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>1</sup>In practice, these models are conducted over shorter passages of text. But for ease of reading, we simply refer to them as “documents”.



**Figure 1:** Bottom-up sliding window vs our top-down partitioning approach. Relevant documents are colored blue. In the case of our approach, the pivot document marked red is identified in the first pass and can be compared to other partitions concurrently with red arrows denoting the use of the pivot element.

given context to judge each document by considering other potentially relevant documents. State-of-the-art list-wise ranking approaches output a permutation of the original ranking instead of explicit scores [7, 8]. The need to encode multiple documents and output multiple tokens to determine document order makes these models computationally expensive due to the quadratic complexity of attention, meaning that encoding multiple documents inherently takes more time than current point-wise architectures.

In existing work, there is still a limit to the number of documents that can be ranked in a single pass (typically around 20 documents [8]). Therefore, a larger ranking must be ‘batched’ before inference. A strategy is needed to select which passages to ‘batch’ together and how to combine them, given that a batch permutation gives no indicator of absolute relevance across the full ranking—only the relative ranking of documents within a batch. Due to the absence of scores, this problem becomes closer to a classical selection or partitioning algorithm, such as quick select [9].

The prevailing approach in list-wise ranking uses a sliding window over a chosen ranking depth from bottom to top [7, 10, 8] as illustrated in Figure 1(a). By using an overlapping ‘stride’ over windows, documents that are initially lower in the ranking can bubble to the top (in principle [10]). However, in this approach, the ranking of each window is dependent on the processing of the documents below, leading to the repeated scoring of multiple documents to enter the top- $k$  and, therefore, an inefficient partitioning scheme. We alleviate this problem with a rank-biased partitioning algorithm, which allows for parallel comparison of multiple batches since each batch only depends on a ‘pivot’ from the first batch (Figure 1(b)). Additionally, noted positional bias may be amplified by the repeated re-scoring of a potentially relevant document to reach a high rank. Though prior work has investigated this effect under the effect of a sliding window [10, 11], we investigate the in-window precision of list-wise rankers to better understand their behaviour agnostic of chosen partitions.

Our key intuition relates to the purpose of such expensive rankers; such an approach is inherently precision-orientated and, therefore, could find use in a broad set of domains that

rely on either an exact match or a small set of highly relevant documents [12, 13]. With such applications in mind, we identify a pivot element at the  $k^{\text{th}}$  rank, which is then compared with all documents down to an arbitrary depth. Any document considered more relevant than the pivot element is necessarily a candidate for the top  $k$  results. A final scoring is performed over all candidates to yield a final top- $k$ . We additionally constrain the number of candidates that can be selected in a given iteration to allow a budget for latency. In improving both the computational expense and latency of list-wise rankers we not only aim to improve their application at scale but make their use as annotators of training data more affordable as list-wise distillation becomes an increasingly common component of state-of-the-art ranking models [8, 14, 15]. We outline the following research questions that form the core contributions of this work.

**RQ-1:** How does the distribution of relevance within a single context window change effectiveness?

We find that due to the list-wise nature of LLM-based rankers, they present a bias to list order even within a single context window. Retrieval effectiveness measured by nDCG@10 can vary greatly within a single window when comparing a reduced number of documents. We generally observe a bias for a correctly ordered ranking, particularly in the case of a small proportion of relevant documents. Notably, we find that a standard cross-encoder achieves nearly identical performance to list-wise rankers over synthetic rankings.

**RQ-2:** How does efficiency trade-off affect performance in list-wise ranking?

Given our findings investigating **RQ-1**, we motivate our novel approach to list-wise ranking that matches the performance of sliding window ranking and reduces the number of LLM inferences while being parallelizable. We find that inferences can be reduced by up to 33% with no significant loss in performance whilst additionally reducing the sequential dependence between windows found in prior approaches, allowing parallelism to improve efficiency from both a compute and latency perspective.

**RQ-3:** How sensitive is top-down partitioning to the relevance of the contents of the first window?

We find that our new approach can be sensitive when the initial window of documents is imprecise. This is due to our use of a pivot element, which can affect overall performance if poorly chosen. However, we can alleviate this effect due to the design of our algorithm, which allows for a budget over the number of candidate documents to collect, which can be increased in the case of a weak first-stage retriever.

**RQ-4:** How does a larger candidate pool affect ranking performance?

In alleviating the effects of pivot element choice, we find that our algorithm successfully recovers from a sub-optimal pivot element when the allocated ranking budget is improved, increasing nDCG@10 by over two points. In doing so, we show that our algorithm can trade off efficiency for effectiveness given a weaker initial retrieved list.

With this work, we aim to bring attention to the inefficiencies of current approaches and how a task-specific approach can largely reduce the required inferences over computationally expensive language models. We release artefacts and experiment code to ensure reproducibility.<sup>2</sup>

---

<sup>2</sup>Github Repository

## 2. Background and Related Work

**The Ranking Problem.** A ranking is a permutation of a set ordered with respect to some scoring function. Within retrieval, for a corpus  $C = \{d_1, \dots, d_{|C|}\}$  and a user query text  $q$ , a ranking model returns a top- $k$  set of documents where  $k \ll |C|$ , ordered by the probability of relevance to  $q$  [16].

Until the resurgence of neural networks and, more specifically, the transformer architecture [17], ad-hoc search primarily involved term weights from exact lexical matches [18]; neural architectures overcome the problem of exact term matching via soft-matching of terms using pre-trained language models [19, 20]. Multiple variants of the transformer architecture are applied in text retrieval, with most approaches falling into one of the following two paradigms. The first is a bi-encoder, which separately encodes a query and document into a dense latent space and executes a vector similarity search to determine relevance [1, 21], a variation of this architecture, learned sparse retrieval uses sparse vectors frequently projected over the vocabulary [22, 23]. The second, a cross-encoder, treats retrieval as a classification problem [2, 24, 25], jointly encoding a query and document before outputting a relevance score for the pair. List-wise ranking can be seen as a cross-encoder variant due to its jointly encoding queries and multiple documents [7]. Though a list-wise loss criterion has been applied in Learn-To-Rank [26] before neural retrieval, due to context window constraints, neural rankers were trained with point-wise or pair-wise loss functions. Recently, approaches using both T5 [14] and ELECTRA [11] architectures have applied list-wise loss criteria as larger batch sizes can be accommodated in modern compute instances.

**List-wise Ranking.** List-wise ranking is the ranking of multiple documents in a single inference step. We follow the notation convention of Tang et al. [27] in which an  $n$ -permutation of an  $n$ -ranking  $R = \{R_i\}_{i=1}^n$  is defined as  $\sigma : \{d_1, \dots, d_n\} \rightarrow \{d_1, \dots, d_n\}$ . Given  $R$ , a list-wise model  $\theta$  outputs a permutation  $\sigma$  ordered by relevance to  $q$ . A neural approach generally performs permutation in discrete space via the decoding of tokens to determine order. For an initial ranking  $R$  and query text  $q$ ,  $\sigma(R) = \text{PERMUTE}(R, q; \theta)$ , where the permutation  $\sigma$  is generated auto-regressively through greedy decoding; however, our approach is not constrained only to methods applying autoregressive decoders. Sun et al. [7] first showed the capability of GPT variants [5] as zero-shot list-wise rankers, with Ma et al. [28] achieving strong performance with a different prompt. Multiple works then distilled GPT-based rankers into smaller architectures in both a decoder-only [29, 8] and encoder-decoder setting [30], finding that the distilled model could outperform its teacher with suitable data augmentation [8]. Zhang et al. [10] then showed that comparable performance could be achieved without using the GPT API and additionally, investigated position bias across multiple windows; we instead investigate in-window precision to understand better the behaviour of list-wise rankers agnostic of a chosen partitioning algorithm.

Current approaches use sliding windows over a ranking to allow for inference in the case that the inclusion of all candidate documents would exceed the maximum context window of many models without excessive truncation. These methods chose a partition size  $w$  where  $w \ll |R|$  and progressively re-rank results from bottom to top by taking strides over windows. These methods effectively cannot be applied top to bottom as doing so provides an opportunity for highly ranked documents to be penalised as opposed to finding relevant documents that are ranked lower by

a first-stage retriever, defeating the point of a cascading ranking pipeline. This naturally leads to inefficient search due to the need to sequentially search over each partition from bottom to top, particularly in the case of repeated re-ranking [8]. Recent work has attempted to partially alleviate these inefficiencies by applying a “set-encoder” architecture [11], which can process all documents in a single window; however, such an architecture requires annotated data from larger list-wise rankers which still employ a sliding window as it outperforms a full list-wise ranking, as such our work can not only improve efficiency in a search setting but allow for the efficient annotation of training data for more efficient rankers. Meng et al. [31] instead apply auxiliary models to dynamically adjust re-ranking depth conditioned on estimates of relevance to improve efficiency by re-ranking a reduced number of documents, where this approach is applied pre-inference, we instead employ a dynamic algorithm which allows early stopping in the case of a suitable number of relevant documents being found, however, our work is not independent of pre-inference approaches, and they could be combined to improve efficiency within this paradigm. Beyond list-wise approaches, Qin et al. [32] proposed a zero-shot pair-wise approach in which, much like a classic sorting algorithm, comparisons are made between elements to produce a ranked list. This parallel with sorting algorithms led to an expansion of this approach by Zhuang et al. [33], finding that efficiency and performance could be further improved by taking inspiration from common sorting algorithms. Due to the pair-wise nature of these approaches, under current architectures, such approaches are less efficient empirically; however, they remain compelling due to the absence of an expensive training stage.

**Dynamic Pruning.** Our approach takes inspiration from classic dynamic pruning algorithms in retrieval. The key notion is that one does not need to consider documents that have no chance of entering the top- $k$  during scoring. Such an idea has been applied to approximate a top- $k$  ranking over sparse indices in a document-at-a-time fashion [34]. For a document  $d$  and query  $q$ , the query-dependent score of  $d$  can be expressed as:

$$\mathcal{S}(d, q) = \sum_{t \in q \cap d} \alpha_t w(t, d) \quad (1)$$

where in Equation 1,  $\alpha_t$  is the importance of a term  $t$  within  $q$  and  $w(d, t)$  is the weight assigned to  $t$  within  $d$ . For each query term the upper bound  $UB_t$  of its contribution to  $d$  is bounded by  $UB_t \geq \alpha_t w(d, t)$ , an upper bound on the query-dependent score of a document  $d$  is then expressed as  $\sum_{t \in q \cap d} UB_t \geq \mathcal{S}(d, q)$ . If this bound is lower than the current  $k^{\text{th}}$  document, the document cannot enter the top- $k$ ; as such, the scoring of many documents can be cut short or eliminated based on this estimation [35]. In recent years, dynamic pruning has been successfully applied to neural approaches, including re-ranking [36] and learned sparse retrieval [37]. What differentiates our algorithm for list-wise ranking is that we cannot make in-document or term-at-a-time comparisons and, therefore, use an entire document text as a threshold.

### 3. Methodology

We now outline our methodology for investigating in-window precision for list-wise rankers and our proposed algorithm.

### 3.1. Position and Relevance

Depending on an initial ranking, an order-sensitive re-ranker must be able to find relevant documents throughout a ranked list regardless of initial position. When a window-based approach is employed, an imprecise ranking may lead to the scattered distribution of relevant documents contrary to the distillation data commonly used to train these list-wise rankers as illustrated in Figure 1. As such, we consider how precision can be affected by both order and relevance distribution. Due to this sensitivity, prior work has proposed the use of a cascading process in which a list is repeatedly re-ranked to compensate for the initial imprecision of a first-stage retriever [8]. We look to investigate this behaviour agnostic of any given algorithm to diagnose the in-window effectiveness of list-wise rankers.

More concretely, we consider relevance binary for the purposes of this investigation, much like in a Cranfield setting. A ranked list  $R$  of size  $k$  can be composed with a ratio  $\frac{1}{r}$  denoting the proportion of relevant documents within  $R$ . Given a set of relevant documents  $D^+$  and non-relevant documents  $D^-$ , we choose  $\frac{k}{r}$  elements from  $D^+$  and  $\frac{k}{1-r}$  elements from  $D^-$ . Under an order-insensitive ranker with oracle knowledge, it is clear that the measured performance of  $R$  would be monotonically increasing with  $r$  agnostic of the ordering of  $R$ . However, we posit that current list-wise rankers could show reduced robustness to ordering and that we can better understand the mechanism of current approaches over an imprecise initial ranking in investigating effectiveness under these conditions.

### 3.2. Proposed Algorithm

Given the order sensitivity of existing approaches, we consider the need for a relevant document which initially lies at a low rank to be repeatedly re-ranked to reach the top of a ranking can have negative effects when re-ranking an imprecise list and furthermore clearly shows the inefficiency of a sliding window approach.

Our core idea is to remove the dependence of a sequence of partitions of a ranking on each other. Currently, under a sliding window  $w$  with a stride  $s$ , dependencies prevent parallelism and other efficiency improvements due to the bottom-up nature of such an approach. To improve efficiency, we partition a ranking in a top-down fashion compared to a highly ranked pivot element  $p$  to select candidates under a budget  $b$  for each iteration. At a high level, consider the upper bounds described in Section 2; our approach instead uses a pivot *document*  $p$  at rank  $k$  as a threshold as opposed to term-based thresholds. Agnostic of list-wise approaches, we assume that implicitly, a document  $d$  is scored with respect to  $q$ . If  $\mathcal{S}(q, d) < \mathcal{S}(q, p)$  then  $d$  cannot enter the top- $k$  ranking. In practice, we directly use the permutation of a set of documents to determine the above inequality as opposed to exact scores, as if a document is ranked lower than  $p$ , it is considered less relevant by a model.

As outlined in Algorithm 1, given a set of documents  $A_{i-1} \subseteq R$  being the previous iteration or the full ranked list, we process the top- $w$  candidates  $L$  and take a pivot element<sup>3</sup>  $p$  at rank  $k$  with all elements above the pivot added to a candidate set  $A_i$  with maximum size  $b$ , which is modified over each iteration; all other elements are added to a backfill set  $B$ , which persists throughout execution to retain all documents which are estimated to be less relevant than the

---

<sup>3</sup>we choose  $k = \frac{w}{2}$  for comparison to the commonly used stride of 5.

pivot. The remainder of the ranked list  $R$  can be processed in parallel in partitions of size of  $w - 1$ .

The pivot  $p$  is prepended to the next set  $L$ , which is then ordered by model  $\theta$  to find documents ranked above  $p$ . Any elements of  $L$  ranked above  $p$  are added to  $A$ , and the remainder is added to  $B$ . If no element in  $R$  is greater than the current pivot, we know no element can enter the top- $k$ ; therefore, our top set is already sorted as  $L$ . If candidates are found, the algorithm is executed on  $A_i$  constrained to budget size  $b$  (Algorithm 1 Line 14). For our main evaluation, the candidate size equals the window size. Therefore, a single extra iteration is required after traversing a full ranked list. Our algorithm incorporates the desirable rank bias of favoring the top-ranked documents because it works top-down to select a candidate pool constrained by budget  $b$ . As a document at rank 20 is more likely to be relevant than a document at rank 100, when a set of  $b$  candidates is found, we posit that there may be no need to search the remaining partitions.

For  $|R| > w$ , a sliding window approach always has a worst-case number of inferences, inferences  $(R) = \frac{|R|}{s} - 1$  which in practise could not be improved via early stopping due to the the bottom to top application of a sliding window. As top-down partitioning does not use a stride over a sliding window, our expectation over the number of inferences is a function of the desired search depth, the window size, and the budget  $b$ . Window size beyond the top- $w$  is reduced due to the need for the pivot element to be included in each inference. As we do not know how many documents will be considered more relevant than the pivot, identifying  $p_i$  at the beginning of the  $i^{\text{th}}$  iteration, we define the set of possible ranking iterations  $A$  as follows. Given the  $0^{\text{th}}$  element  $A_0 = R$ :

$$A_i = \{A_i : |A_i| < b \wedge \forall d \in A_{i-1}, \text{rank}(d) < \text{rank}(p_{i-1})\} \quad (2)$$

We can then estimate a worst-case number of inferences by taking the sum of inferences over all iterations defined in Equation 2. Empirically, this does not occur as the algorithm either finds  $b$  candidates before processing all partitions or the ranker is sufficiently precise to meet the condition  $|A_i| = k - 1$  early.

$$\text{inferences}(R) = \sum_{i=0}^{|A|} 1 + \frac{|A_i| - w}{w - 1} \quad (3)$$

Equation 3 degenerates to inferences  $(R) = 2 + \frac{|R| - w}{w - 1}$  in the case of  $b = w$ . We need a single inference to find the initial pivot at each stage; we then have  $|A_i| - w$  documents remaining to process with a reduced window size  $w - 1$  as we must account for the pivot element. We then require a final inference to order the set of  $w$  candidates. Though both a sliding window and our algorithm are of  $O(n)$  complexity, empirically, both early stopping and the lack of a stride lead to reduced computation.

---

**Algorithm 1: Top-Down Partitioning**

---

**Input:** Candidate set:  $A_{i-1}$ , Backfill set:  $B$ , Query:  $q$ , Window size:  $w$ , Budget:  $b$ , Optimization cutoff:  $c$ , List-wise ranker:  $\theta$

**Output:**  $\sigma(A_{i-1})$ :  $A_{i-1}$  ordered by relevance to  $q$

```
1 begin
2    $L \leftarrow \{d_1, \dots, d_w, d \in A_{i-1}\}$  // First  $w$  documents
3    $R \leftarrow A_{i-1} - L$  // Remaining documents
4    $L \leftarrow \text{PERMUTE}(L, q; \theta)$  // Order current partition
5    $p \leftarrow L[k]$  // Take pivot element
6    $A_i \leftarrow L[1 : k]$ 
7    $B \leftarrow B \cup L[k + 1 : |L|]$ 
8   while  $|A_{i-1}| > 0$  and  $|A_i| < b$  do
9      $L \leftarrow \{d_1, \dots, d_{w-1}, d \in R\}$ 
10     $R \leftarrow R - L$ 
11     $L \leftarrow \text{PERMUTE}(p \cup L, q; \theta)$ 
12     $A_i \leftarrow A_i \cup L[1 : p]$  // Elements  $> p$ 
13     $B \leftarrow B \cup L[p + 1 : |L|]$  // Elements  $< p$ 
14  return  $(|A_i| = k - 1) ? A_i \cup p \cup B : \text{pivot}(A_i) \cup p \cup B$ 
```

---

## 4. Evaluation

We now describe the evaluation of top-down partitioning for list-wise re-ranking and our observations from empirical evidence.

### 4.1. Experiment Setup

We outline the following research questions to assess our novel algorithm:

**RQ-1:** How does the distribution of relevance within a window change effectiveness when order changes?

**RQ-2:** How does efficiency trade-off affect performance in list-wise ranking?

**RQ-3:** How sensitive is top-down partitioning to the relevance of the contents of the first window?

**RQ-4:** How does a larger candidate pool affect ranking performance?

### 4.2. List-Wise Ranking Approaches.

We compare approaches to ranking with sequence-to-sequence list-wise rankers. We do not consider variants of pairwise re-ranking [32, 33] as their efficiency and performance were found to be inferior to the following approaches [29].

**Single Window (Single):** Re-ranks the top- $w$  documents from a first stage ranking where  $w$  is window size. Assumes sufficient recall within the top- $w$  of a ranked list retrieved by a first stage model.

**Sliding Window (Sliding):** For a given window size  $w$  and stride  $s$  re-ranks documents to depth  $D$  using a sliding window from bottom to top.

**Top-Down Partitioning (TDPart):** Our approach outlined in Algorithm 1. Performs an initial search over a top- $w$  window to find a pivot at cutoff  $k$  before searching to depth  $D$  comparing each window to the pivot.



### 4.3. Datasets

We conduct our core experiments on the MSMARCO corpus [38], a collection of over 8.8 million documents commonly used to train neural ranking models. As we are primarily concerned with the precision of the top- $k$  re-ranked documents, we choose to evaluate our approach on the TREC Deep Learning 2019 and 2020 test queries [39, 40], which have densely annotated queries in contrast to the official MSMARCO DEV set [38], we refer the reader to the original model papers for an overview of each model’s performance on sparse benchmarks. To evaluate the robustness of our approach in an out-of-domain setting we take a subset of the BEIR benchmark sets [41], due to the greater computational expense of current list-wise rankers multiple sparsely annotated test sets would be infeasible under our compute budget, as such we choose test sets with a smaller set of densely annotated topics namely, TREC Covid [42] and Touche [43]. TREC COVID is a test benchmark with a corpus of 171k texts and 50 annotated queries with topics focusing on health information relating to the COVID-19 virus. Touche is an argument retrieval benchmark with a corpus of 382k texts and 49 annotated queries with topics relating to both societally important and everyday information needs which require arguments to provide evidence for a given perspective.

We evaluate nDCG@1, 5, 10, and P@10 using the `ir_measures` library; following standard MSMARCO conventions, we consider a relevance judgment of 2 or more as relevant and a relevance judgement of 1 for other datasets. Our algorithm aims to reduce the inferences required when applying list-wise rankers while maintaining effectiveness, so we conduct equivalence tests (paired TOST) with  $p < 0.05$  and a 5% bounds between our methods and baselines. Additionally we report the average number of inferences required per query for both in- and out-of-domain experiments and the number of possible parallel inferences.

### 4.4. Models

We apply our approach to multiple state-of-the-art list-wise rankers that execute the PERMUTE function described in Section 3. RankGPT [7], which directly prompts GPT variants<sup>4</sup>, a decoder-only distilled ranker RankZephyr [8] based on the Zephyr-7B [44] architecture and LiT5 [30] based on a T5 [20] fusion-in-decoding model [13]. We allow a maximum of 4096 tokens per input in all cases. We also employ an oracle approach with access to test relevance judgments returning a permutation sorted by these judgments to ensure our algorithm can remain robust as underlying models improve. Due to tie breaks inherent when sorting by discrete relevance judgments, precision can vary. As our approach is fundamentally rank-biased, we investigate the effect of first-stage rankers on downstream precision. We re-rank BM25 [18], a common lexical model, and we use the Terrier implementation with default parameters [45]; SPLADE++ ED [23], a teacher-distilled learned sparse model and RetroMAE [46], a teacher-distilled bi-encoder using MAE pre-training over a brute force index. In experiments addressing **RQ-1**, we additionally contrast RankZephyr with monoELECTRA [25], a strong point-wise cross-encoder to control for order-sensitivity in our comparisons.

---

<sup>4</sup>We use GPT 3.5 (gpt-3.5-turbo-0125) due to budget constraints

## 4.5. In-Window Precision

As described in Section 3.1, we look to investigate how both ranking order and relevance distribution affect the in-window effectiveness of list-wise rankers. In investigating **RQ-1**, we define a ranking data generation process with oracle knowledge, which generates ranked lists given a ratio of relevant and non-relevant documents. To reduce the noise in this process, given the existence of false negatives in MSMARCO, we persist in a ranking between ratios within the generation process, only adding relevant documents as the ratio  $r$  increases. Using human relevance judgements for a test collection, we partition all judged documents  $D$  into relevant,  $D^+$ , and non-relevant,  $D^-$ , sets. We filter queries<sup>5</sup> such that all evaluated queries have at least  $k - 1$  judgements in both  $D^+$  and  $D^-$  for  $|R| = k$  to allow for the smooth interpolation between majority relevant and majority non-relevant ranked lists. For experiments with MSMARCO, we combine Deep Learning 2019 and 2020 queries, yielding 53 compatible queries for  $|R| = 20$ . For out-of-domain experiments, 50 TREC COVID queries and 16 Touche queries are valid. Naturally, we use the full set of queries when performing retrieval effectiveness evaluation outside of this synthetic setting.

To persist a ranked list, for an initial ratio  $r_0$ , we take documents from the sets described above. When generating for the next ratio  $r_1$ , where  $r_1 > r_0$ , we find the number of *new* relevant documents we need to include via  $n = r_1 - r_0$ . We remove  $n$  non-relevant documents from  $R$  and add  $n$  documents from  $D^+$ , ensuring that random noise is reduced and we can clearly measure the effect of adding new relevant documents across a given sample.

To investigate order-sensitivity, we again utilise oracle knowledge to sort by relevance judgement; on MSMARCO, this is a value in the range  $[0, 1, 2, 3]$  and generally, for BEIR datasets, a value in the range  $[0, 1, 2]$ . We sort each ranking in ascending (**ASC**), descending (**DESC**) and random order (**RANDOM**). In doing so, we investigate how the behaviour of these models varies with lists, which could be present at various points in the re-ranking process. For example, when re-ranking a strong first-stage, the top- $k$  is likely to be well ordered, approximating descending order with a majority of relevant documents present; however, when re-ranking a weaker first stage, when ranking lower ranked documents, a window may be composed of only a few relevant documents in either random or ascending order. As such, our synthetic conditions can have parallels with real ranked lists.

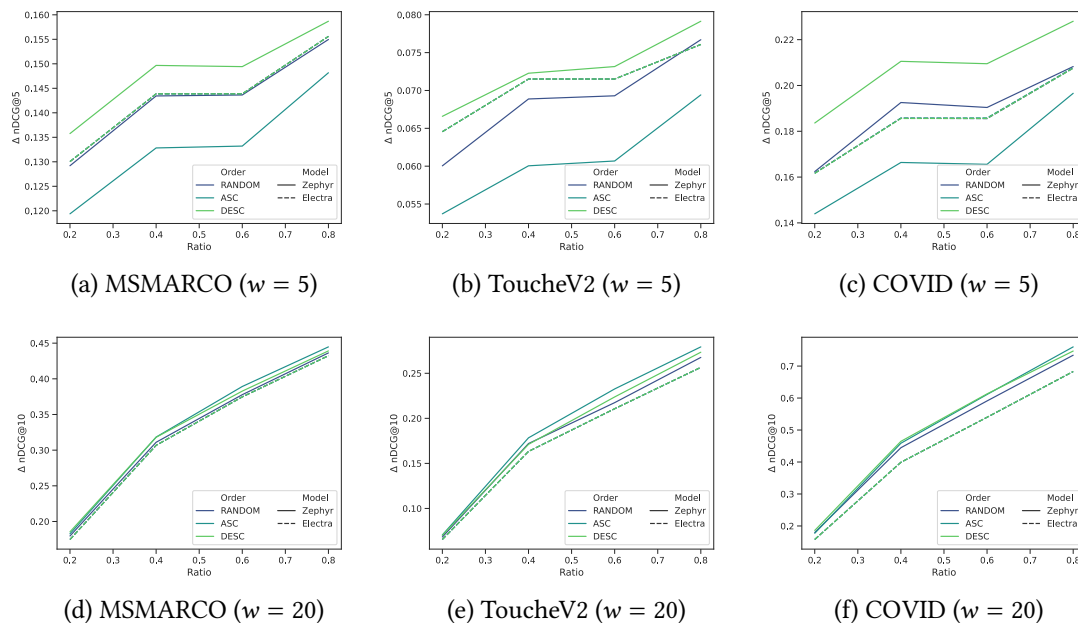
To prevent bias in the initialisation of each ranked list, we aggregate performance over 5 initial rankings for each query. We investigate ratios in the range  $[0.2, 0.4, 0.6, 0.8]$ , where a smaller ratio denotes a reduced proportion of relevant documents. We assess window sizes of 5 and 20 to observe how these conditions vary when a reduced number of documents are observed.

## 5. Results and Discussion

We now present and discuss our findings with respect to research questions outlined in Section 4.1.

---

<sup>5</sup>We use the same filtered set of queries for different window sizes to allow comparisons, i.e. both window sizes sample from the pool satisfying  $|R| = 20$ .



**Figure 2:** Effectiveness of a list-wise and point-wise ranker varying the proportion of relevant documents and the ranked order measuring nDCG@10 across MSMARCO Deep Learning, Touche and TREC COVID queries.  $w$  represents window size or, in this case, the number of documents ranked at once.

## 5.1. In-Window Effectiveness

In Figure 2, we present the effectiveness of RankZephyr<sup>6</sup> compared to monoELECTRA (order-invariant), controlling for the proportion of relevant documents (Ratio) and the ordering of the ranking. Addressing **RQ-1**, Figures 2 (a), (b), and (c) indicate that for a reduced list length, there is a clear distinction between orderings suggesting a greater sensitivity to relevant documents being placed at the beginning of a ranking. Additionally, observe that RankZephyr, when re-ranking from a ranked list in descending order, consistently outperforms monoELECTRA; however, random order is aligned with monoELECTRA in-domain, and performance varies out-of-domain.

In contrast, we observe that in Figures 2, RankZephyr presents minimal order sensitivity. Additionally, it is notable that monoELECTRA performs near-identically to RankZephyr under this synthetic setting. Given that a greater margin is observed between list-wise and point-wise methods out-of-domain, additionally observed by Schlatt et al. [11], we consider that “inter-document” interactions, may be unnecessary when priors are suitably modelled in-domain, with list-wise inference provided appropriate context when the target domain is not fully captured in a point-wise setting; however, clearly, list-wise distillation can improve performance, but list-wise inference may lead to a bias to rankings similar to that of a teacher model as discussed below in Section 5.2.

Coupled with this finding, we observe that descending order is preferred in both in- and out-

<sup>6</sup>Similar trends were consistently observed for other list-wise rankers

**Table 1**

Comparing ‘Single Window’, ‘Sliding Window’, and ‘TDPart’ (ours) approaches on TREC Deep Learning benchmarks. We report TOST ( $p < 0.05$ ), equivalence to our approach denoted with underline. We report mean inferences with the mean number of inferences run in parallel in parentheses.

		SPLADE ++ ED							RetroMAE							BM25												
		nDCG@1			nDCG@5			P@10	nDCG@1			nDCG@5			P@10	nDCG@1			nDCG@5			P@10						
Model	Mode	DL19	DL20	DL19	DL20	DL19	DL20	DL19	DL20	N. Inf.	DL19	DL20	DL19	DL20	DL19	DL20	DL19	DL20	N. Inf.	DL19	DL20	DL19	DL20	DL19	DL20	N. Inf.		
Oracle	Single	<u>0.977</u>	<u>1.000</u>	<u>0.936</u>	<u>0.977</u>	<u>0.890</u>	<u>0.916</u>	<u>0.765</u>	<u>0.748</u>	1.0 (1.0)	0.953	0.981	0.925	0.934	0.863	0.874	0.749	0.696	1.0 (1.0)	0.907	0.935	0.823	0.811	0.719	0.715	0.560	0.489	1.0 (1.0)
Oracle	Sliding	<u>0.984</u>	<u>1.000</u>	<u>0.972</u>	<u>0.995</u>	<u>0.957</u>	<u>0.978</u>	<u>0.872</u>	<u>0.822</u>	9.0 (1.0)	<u>0.992</u>	<u>1.000</u>	<u>0.968</u>	<u>0.993</u>	<u>0.948</u>	<u>0.958</u>	<u>0.863</u>	<u>0.800</u>	9.0 (1.0)	<u>0.953</u>	<u>0.981</u>	<u>0.924</u>	<u>0.928</u>	<u>0.879</u>	<u>0.880</u>	<u>0.788</u>	<u>0.702</u>	8.9 (1.0)
Oracle	TDPart	0.984	1.000	0.972	0.995	0.956	0.976	0.872	0.820	7.0 (5.0)	0.992	1.000	0.967	0.993	0.945	0.958	0.853	0.800	7.1 (5.1)	0.946	0.981	0.913	0.925	0.858	0.872	0.765	0.687	7.4 (5.4)
Zephyr	Single	<u>0.849</u>	<u>0.855</u>	<u>0.808</u>	<u>0.831</u>	<u>0.777</u>	<u>0.795</u>	<u>0.672</u>	<u>0.637</u>	1.0 (1.0)	<u>0.810</u>	<u>0.867</u>	<u>0.792</u>	<u>0.815</u>	<u>0.758</u>	<u>0.778</u>	<u>0.665</u>	<u>0.607</u>	1.0 (1.0)	0.713	0.784	0.681	0.689	0.625	0.628	0.500	0.435	1.0 (1.0)
Zephyr	Sliding	<u>0.849</u>	<u>0.849</u>	<u>0.791</u>	<u>0.826</u>	<u>0.777</u>	<u>0.802</u>	<u>0.691</u>	<u>0.631</u>	9.0 (1.0)	<u>0.810</u>	<u>0.867</u>	<u>0.796</u>	<u>0.818</u>	<u>0.769</u>	<u>0.785</u>	<u>0.686</u>	<u>0.607</u>	9.0 (1.0)	0.713	0.846	0.715	0.720	0.707	0.722	0.637	0.535	8.9 (1.0)
Zephyr	TDPart	0.833	0.852	0.794	0.828	0.780	0.805	0.679	0.639	7.4 (5.4)	0.826	0.864	0.797	0.814	0.771	0.788	0.686	0.619	7.2 (5.2)	0.736	0.858	0.731	0.772	0.681	0.723	0.577	0.543	7.4 (5.4)
LiT5	Single	<u>0.802</u>	<u>0.827</u>	<u>0.795</u>	<u>0.797</u>	<u>0.763</u>	<u>0.763</u>	<u>0.665</u>	<u>0.598</u>	1.0 (1.0)	0.787	<u>0.858</u>	<u>0.773</u>	<u>0.776</u>	<u>0.739</u>	<u>0.748</u>	0.653	<u>0.585</u>	1.0 (1.0)	0.748	0.799	0.685	0.678	0.626	0.604	0.507	0.413	1.0 (1.0)
LiT5	Sliding	<u>0.802</u>	<u>0.818</u>	<u>0.798</u>	<u>0.771</u>	<u>0.774</u>	<u>0.763</u>	<u>0.691</u>	<u>0.600</u>	9.0 (1.0)	0.756	<u>0.855</u>	<u>0.770</u>	<u>0.781</u>	<u>0.741</u>	<u>0.750</u>	0.665	<u>0.578</u>	9.0 (1.0)	0.787	<u>0.833</u>	0.751	<u>0.735</u>	0.723	0.700	0.640	0.531	8.9 (1.0)
LiT5	TDPart	0.787	0.827	0.795	0.791	0.766	0.768	0.684	0.604	6.3 (4.3)	0.740	0.867	0.763	0.798	0.743	0.761	0.684	0.591	6.6 (4.6)	0.787	0.836	0.724	0.746	0.687	0.679	0.593	0.489	5.9 (3.9)
GPT 3.5	Single	0.833	0.787	<u>0.783</u>	<u>0.775</u>	<u>0.760</u>	<u>0.752</u>	<u>0.644</u>	0.607	1.0 (1.0)	0.841	0.827	<u>0.802</u>	0.789	<u>0.741</u>	<u>0.744</u>	0.630	<u>0.576</u>	1.0 (1.0)	0.698	0.765	0.647	0.666	0.602	0.588	0.484	0.393	1.0 (1.0)
GPT 3.5	Sliding	0.837	0.781	<u>0.786</u>	<u>0.779</u>	<u>0.754</u>	0.766	<u>0.642</u>	0.617	9.0 (1.0)	0.802	<u>0.799</u>	0.779	<u>0.752</u>	<u>0.740</u>	<u>0.727</u>	0.649	0.561	9.0 (1.0)	0.709	0.809	0.727	<u>0.704</u>	0.686	0.654	0.588	0.480	8.9 (1.0)
GPT 3.5	TDPart	0.802	0.806	0.779	0.788	0.753	0.752	0.642	0.598	7.4 (5.4)	0.837	0.796	0.787	0.758	0.736	0.739	0.637	0.578	7.2 (5.2)	0.686	0.802	0.693	0.708	0.654	0.637	0.556	0.459	7.5 (5.5)

of-domain settings for lower proportions of relevant documents, suggesting that a well-ordered ranking can further improve performance. We draw a parallel with repeated re-ranking over BM25 proposed by Pradeep et al. [8], in which initially, a relevant document may be missed as it may happen to be in a low position in the ranking and its rank is not sufficiently increased in the current sliding window. However, when re-ranking is again performed, the position of that document may have been moved such that the document rank is sufficiently improved to consider it at an earlier position in the context window, aligning with descending order for a low ratio as discussed above<sup>7</sup>.

When applying top-down partitioning, we consider that this effect may be reduced by placing the pivot element used for candidate selection at the first position in the ranking. Assuming that a ranked list is suitably well-ordered to allow the chosen pivot to be partially relevant, this should align with our findings that descending order is preferred for a reduced number of relevant documents.

## 5.2. In-Domain Effectiveness

**Inferences are Empirically Reduced.** Table 1 presents our core results over TREC Deep Learning queries, contrasting our algorithm with existing approaches. Rows 1-3 show an oracle ranker over each first stage. We observe statistical equivalence with a sliding window approach while consistently reducing the number of inferences required to achieve this performance. In doing so, we show that the comparison of a pivot element as opposed to a strided window is sufficient for list-wise inference.


Addressing **RQ-2**, as can be observed in the last column of Table 1, when applying state-of-the-art rankers, we again observe a reduced number of inferences across all cases when applying top-down partitioning, as expected by the analysis of our algorithm. Additionally, our approach exceeds or matches a sliding window’s performance measured by nDCG@10 in 79% of cases,

<sup>7</sup>This observation may be partially attributed to causal attention, which, in the case of descending order, the model can contextualise non-relevant documents with respect to already observed relevant documents however we leave verification of this conjecture to future work.

including RankGPT, which is zero-shot and known to have variable output [29]. Additionally, observe that under an oracle setting, our algorithm is more efficient depending on the precision of a first-stage retriever, showing that we can improve efficiency agnostic of the effectiveness of an underlying estimator. When utilising both decoder-only and encoder-decoder architectures, the number of inferences required does not show a linear trend, which could be attributed to the ability of each model to suitably identify relevant documents given their position within a current window; however, in each case, we see that our algorithm leads to at maximum .4 more inferences than an oracle setting on average per query. Crucially, the latency of our approach would be dramatically reduced due to the majority of inferences being parallelizable, as observed in the bracket of the final column of each Table. Particularly in the case of LiT5, which is based on the smaller T5 architecture, this could lead to greater potential for list-wise inference to scale under a greater number of concurrent queries. Due to the reduced window size required by our algorithm, we frequently observe a smaller final window as in our experiments, depth 100 does not cleanly divide by 19 ( $w - 1$ ); simply searching to an approximate depth which has the factor  $w - 1$  would further reduce inferences with minimal effect on performance.

**Sensitivity to Initial Ranking.** Addressing **RQ-3**, when applying our algorithm to strong first stages such as RetroMAE and SPLADE++ ED, it can be observed that in all cases, our approach is either statistically equivalent or outperforms the strongest baseline, being either a single window in the case that the top- $w$  documents are already sufficiently precise or a sliding window. Additionally, our approach generally outperforms a sliding window in terms of nDCG@5; however, we note that nDCG@1 is variable. We generally achieve statistical equivalence with a sliding window. This observation shows that our approach provides an effective trade-off between precision and latency whilst outperforming a single window where a sliding window fails. Of note is the statistical equivalence of all approaches when using a strong first stage and list-wise ranker; however, when considering a list-wise ranker, one would want to maximize performance such that a deeper search is essential. We propose that the multiple scoring of documents in a sliding window approach, which can ‘miss’ relevant documents due to the need for multiple scoring of the same document as observed by Zhang et al. [10] which may inversely pull non-relevant documents into the top- $k$ , explaining performance degradation over a single window as can be observed particularly under RankZephyr. Alternatively, when applying a weaker first stage, such as BM25, though our approach improves over a single window, we observe non-equivalent performance in P@10 when applying our approach versus a sliding window. A likely cause of this degradation is that due to the imprecision of the first  $w$  elements of a BM25 ranking, the chosen pivot element is insufficiently relevant to distinguish suitable candidates. Therefore, the candidate pool is polluted with partially or non-relevant items as the pivot is not sufficiently relevant to provide a challenging threshold. However, our approach more robustly ranks documents in contrast to a sliding window, as shown by improvements in nDCG@1 and nDCG@5, providing further evidence that non-relevant documents are accidentally ‘bubbled’ up to high ranks. Applying a list-wise ranker with our approach saves computation; therefore, allocating a small fraction of these savings to a stronger first stage, such as SPLADE++ ED, is worthwhile.

**Ablation of Budget.** Addressing **RQ-4**, we conduct an ablation of our approach to assess how a larger budget and, therefore, the opportunity to re-rank the top- $k$  multiple times can affect retrieval performance. We assess budgets in the range [20, 30, 40, 50]. In Figure 3, observe



**Figure 3:** Ablation of Budget across each first-stage ranker on the TREC Deep Learning 2019 test queries.

that for precise first stages, small improvements in effectiveness occur when the budget is increased apart from RankGPT, which exhibits high variance. Additionally, we find that larger values of the budget parameter can exceed reported performance on Deep Learning queries from prior works; however, we do not tune this value on a validation set, and any improvement is insignificant over previous state-of-the-art methods. Akin to the multiple re-ranking stages proposed by Pradeep et al. [8] in which the previous ranking by the list-wise ranker is again re-ranked; however, we find a budget increase somewhat unnecessary given a first-stage dense or sparse neural retriever as performance gains are insufficient compared to the increased number of inferences required for a larger budget. However, in the case of BM25, we see that increasing the budget improves performance across all models, particularly in the case of LiT5, where  $nDCG@10$  is improved from 0.687 to 0.715 when the budget is increased. Generally, we observe improvements of around 2 points of  $nDCG@10$ , allowing RankGPT to reach the performance of LiT5 when re-ranking SPLADE++ ED. We attribute this to the low precision of BM25, meaning that a greater budget is required to recover from a poor initial pivot value such that a progressive re-ranking is helpful. This explains small albeit insignificant degradation in  $nDCG@10$  performance compared to a sliding window in Table 1. Under our algorithm, progressive re-ranking does not require full re-ranking and is instead adaptive to the initial candidate set, improving efficiency.

### 5.3. Out-of-Domain Effectiveness

Table 2 presents the out-of-domain performance of list-wise rankers using different partitioning algorithms. We observe again that our approach is statistically equivalent to a sliding window approach under an oracle setting but note that unlike MSMARCO, a single window is generally insufficient in terms of recall. Our approach improves over a single window on TREC COVID in 2 of 3 models and, on Touche, outperforms a sliding window approach in terms of  $nDCG@10$  and  $P@10$  using both RankZephyr and LiT5 whilst reducing inferences by over 33%. We observe reduced performance assessed over TREC COVID using LiT5 compared to a sliding window approach, potentially showing that given the large training process of such rankers, without interventions (such as those applied in RankZephyr, e.g. order shuffling) to discriminate non-relevant documents from a pivot is more challenging. Over Touche both a sliding window and our approach generally degrade over a single window suggesting that these models cannot

**Table 2**

Comparing ‘Single Window’, ‘Sliding Window’, and ‘TDPart’ (ours) approaches on out-of-domain benchmarks. We report TOST ( $p < 0.05$ ), equivalence to our approach denoted with underline. We report mean inferences with the mean number of inferences run in parallel in parentheses.

Model	Mode	TREC COVID				Touche				N. Inf.
		nDCG@1	nDCG@5	nDCG@10	P@10	nDCG@1	nDCG@5	nDCG@10	P@10	
Oracle	Single	0.980	0.946	0.874	0.778	0.959	0.814	0.615	0.402	1.0 (1.0)
Oracle	Sliding	<u>1.000</u>	<u>0.989</u>	<u>0.983</u>	<u>0.974</u>	<u>1.000</u>	<u>0.977</u>	<u>0.877</u>	<u>0.673</u>	9.0 (1.0)
Oracle	TDPart	1.000	0.989	0.975	0.954	1.000	0.977	0.876	0.671	6.5 (4.5)
Zephyr	Single	0.890	0.813	0.754	0.632	0.480	0.383	0.356	0.265	1.0 (1.0)
Zephyr	Sliding	0.870	0.842	<u>0.805</u>	0.702	0.490	0.342	0.318	0.222	9.0 (1.0)
Zephyr	TDPart	0.910	0.854	<u>0.804</u>	0.696	0.480	0.369	0.340	0.241	5.8 (3.8)
LiT5	Single	0.820	0.754	<u>0.715</u>	0.608	0.429	0.366	0.340	0.255	1.0 (1.0)
LiT5	Sliding	0.850	0.763	0.744	0.644	0.439	0.330	0.323	0.239	9.0 (1.0)
LiT5	TDPart	0.820	0.769	0.725	0.618	0.418	0.337	0.331	0.251	5.2 (3.2)
GPT 3.5	Single	0.880	0.779	0.718	0.598	0.459	0.374	0.325	0.237	1.0 (1.0)
GPT 3.5	Sliding	0.910	0.833	0.774	0.656	0.357	0.322	0.308	0.235	9.0 (1.0)
GPT 3.5	TDPart	0.900	0.849	0.784	0.682	0.276	0.308	0.272	0.198	6.5 (4.5)

separate relevant and non-relevant documents when searching to a greater depth, this effect is most pronounced in RankGPT which presents large degradation in terms of nDCG@10. Though our approach is a step towards more robust ranking using list-wise rankers, performance in contrast to a single window is still reduced, and such cases should be studied further to compensate for domain shift. The lack of statistical equivalence between approaches over list-wise rankers would suggest that greater work is required to improve their robustness and, therefore, allow greater recall. We note a clear trend that stronger models are more effective when applying a pivot-based approach and generally outperform the exhaustive sliding window while providing a sufficient trade-off between performance and efficiency.

## 6. Conclusion

This work introduces a novel algorithm for list-wise ranking, which, while improving efficiency, matches or exceeds prior work in and out-of-domain. In investigating the precision of such rankers agnostic of a particular partitioning algorithm, we further motivate our approach using a pivot element chosen from highly ranked documents. Using a pivot as opposed to a strided window, we can rank documents to an arbitrary depth with a reduced number of inferences and in parallel. This finding is essential in improving the efficiency of list-wise approaches as they find applications in both ranking and data annotation for more efficient rankers. Our proposed approach is statically equivalent to prior work in an oracle setting, suggesting that as underlying models improve, our algorithm can continue to improve efficiency with no performance degradation. Our approach is most effective with strong first-stage retrievers, where a single high-quality pivot element can be used for efficient candidate selection. In contrast, weak first stages might require a larger budget allowed by our algorithm for progressive re-ranking to overcome the initial imprecision. We identified limitations in out-of-

domain retrieval tasks, where discrimination between relevant and non-relevant documents becomes more challenging at greater re-ranking depths. This highlights the need for further research on improving the robustness of list-wise rankers, particularly for generalizability across different domains. In summary, we show that current approaches are sub-optimal, and further consideration must be given to the efficiency of list-wise ranking at an algorithmic level.

## References

- [1] V. Karpukhin, B. Oguz, S. Min, P. S. H. Lewis, L. Wu, S. Edunov, D. Chen, W. Yih, Dense passage retrieval for open-domain question answering, in: B. Webber, T. Cohn, Y. He, Y. Liu (Eds.), Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020, Association for Computational Linguistics, 2020, pp. 6769–6781. URL: <https://doi.org/10.18653/v1/2020.emnlp-main.550>. doi:10.18653/v1/2020.EMNLP-MAIN.550.
- [2] R. F. Nogueira, K. Cho, Passage re-ranking with BERT, CoRR abs/1901.04085 (2019). URL: <http://arxiv.org/abs/1901.04085>. arXiv:1901.04085.
- [3] R. Pradeep, R. F. Nogueira, J. Lin, The expando-mono-duo design pattern for text ranking with pretrained sequence-to-sequence models, CoRR abs/2101.05667 (2021). URL: <https://arxiv.org/abs/2101.05667>. arXiv:2101.05667.
- [4] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, S. Narang, G. Mishra, A. Yu, V. Y. Zhao, Y. Huang, A. M. Dai, H. Yu, S. Petrov, E. H. Chi, J. Dean, J. Devlin, A. Roberts, D. Zhou, Q. V. Le, J. Wei, Scaling instruction-finetuned language models, CoRR abs/2210.11416 (2022). URL: <https://doi.org/10.48550/arXiv.2210.11416>. doi:10.48550/ARXIV.2210.11416. arXiv:2210.11416.
- [5] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- [6] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, G. Lample, Llama: Open and efficient foundation language models, CoRR abs/2302.13971 (2023). URL: <https://doi.org/10.48550/arXiv.2302.13971>. doi:10.48550/ARXIV.2302.13971. arXiv:2302.13971.
- [7] W. Sun, L. Yan, X. Ma, S. Wang, P. Ren, Z. Chen, D. Yin, Z. Ren, Is chatgpt good at search? investigating large language models as re-ranking agents, in: H. Bouamor, J. Pino, K. Bali (Eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023, Association for Computational Linguistics, 2023, pp. 14918–14937. URL: <https://aclanthology.org/2023.emnlp-main.923>.



- [8] R. Pradeep, S. Sharifymoghaddam, J. Lin, Rankzephyr: Effective and robust zero-shot listwise reranking is a breeze!, *CoRR abs/2312.02724* (2023). URL: <https://doi.org/10.48550/arXiv.2312.02724>. doi:10.48550/ARXIV.2312.02724. arXiv:2312.02724.
- [9] C. A. R. Hoare, Algorithm 65: find, *Commun. ACM* 4 (1961) 321–322. URL: <https://doi.org/10.1145/366622.366647>. doi:10.1145/366622.366647.
- [10] X. Zhang, S. Hofstätter, P. Lewis, R. Tang, J. Lin, Rank-without-gpt: Building gpt-independent listwise rerankers on open-source large language models, *CoRR abs/2312.02969* (2023). URL: <https://doi.org/10.48550/arXiv.2312.02969>. doi:10.48550/ARXIV.2312.02969. arXiv:2312.02969.
- [11] F. Schlatt, M. Fröbe, H. Scells, S. Zhuang, B. Koopman, G. Zuccon, B. Stein, M. Potthast, M. Hagen, Set-encoder: Permutation-invariant inter-passage attention for listwise passage re-ranking with cross-encoders, 2024. arXiv:2404.06912.
- [12] P. S. H. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive NLP tasks, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020*. URL: <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>.
- [13] G. Izacard, P. S. H. Lewis, M. Lomeli, L. Hosseini, F. Petroni, T. Schick, J. Dwivedi-Yu, A. Joulin, S. Riedel, E. Grave, Few-shot learning with retrieval augmented language models, *CoRR abs/2208.03299* (2022). URL: <https://doi.org/10.48550/arXiv.2208.03299>. doi:10.48550/ARXIV.2208.03299. arXiv:2208.03299.
- [14] H. Zhuang, Z. Qin, R. Jagerman, K. Hui, J. Ma, J. Lu, J. Ni, X. Wang, M. Bendersky, 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 '23, Association for Computing Machinery, New York, NY, USA, 2023*, p. 2308–2313. URL: <https://doi.org/10.1145/3539618.3592047>. doi:10.1145/3539618.3592047.
- [15] C. Lassance, H. Déjean, T. Formal, S. Clinchant, Splade-v3: New baselines for splade, 2024. arXiv:2403.06789.
- [16] S. E. Robertson, The probability ranking principle in ir, *Journal of documentation* 33 (1977) 294–304.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, 2017*, pp. 5998–6008. URL: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- [18] S. E. Robertson, S. Walker, S. Jones, M. Hancock-Beaulieu, M. Gatford, Okapi at TREC-3, in: D. K. Harman (Ed.), *Proceedings of The Third Text REtrieval Conference, TREC 1994, Gaithersburg, Maryland, USA, November 2-4, 1994*, volume 500-225 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 1994, pp. 109–126. URL: <http://trec.nist.gov/pubs/trec3/papers/city.ps.gz>.
- [19] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional

- transformers for language understanding, in: J. Burstein, C. Doran, T. Solorio (Eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), Association for Computational Linguistics, 2019, pp. 4171–4186. URL: <https://doi.org/10.18653/v1/n19-1423>. doi:10.18653/V1/N19-1423.
- [20] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *J. Mach. Learn. Res.* 21 (2020) 140:1–140:67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [21] S. MacAvaney, A. Yates, A. Cohan, N. Goharian, CEDR: contextualized embeddings for document ranking, in: B. Piwowarski, M. Chevalier, É. Gaussier, Y. Maarek, J. Nie, F. Scholer (Eds.), Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019, ACM, 2019, pp. 1101–1104. URL: <https://doi.org/10.1145/3331184.3331317>. doi:10.1145/3331184.3331317.
- [22] S. MacAvaney, F. M. Nardini, R. Perego, N. Tonello, N. Goharian, O. Frieder, Expansion via prediction of importance with contextualization, in: J. X. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, Y. Liu (Eds.), Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020, ACM, 2020, pp. 1573–1576. URL: <https://doi.org/10.1145/3397271.3401262>. doi:10.1145/3397271.3401262.
- [23] T. Formal, B. Piwowarski, S. Clinchant, SPLADE: sparse lexical and expansion model for first stage ranking, *CoRR abs/2107.05720* (2021). URL: <https://arxiv.org/abs/2107.05720>. arXiv:2107.05720.
- [24] R. F. Nogueira, Z. Jiang, R. Pradeep, J. Lin, Document ranking with a pretrained sequence-to-sequence model, in: T. Cohn, Y. He, Y. Liu (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020, volume EMNLP 2020 of *Findings of ACL*, Association for Computational Linguistics, 2020, pp. 708–718. URL: <https://doi.org/10.18653/v1/2020.findings-emnlp.63>. doi:10.18653/V1/2020.FINDINGS-EMNLP.63.
- [25] R. Pradeep, Y. Liu, X. Zhang, Y. Li, A. Yates, J. Lin, Squeezing water from a stone: A bag of tricks for further improving cross-encoder effectiveness for reranking, in: M. Hagen, S. Verberne, C. Macdonald, C. Seifert, K. Balog, K. Nørkvåg, V. Setty (Eds.), Advances in Information Retrieval - 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10-14, 2022, Proceedings, Part I, volume 13185 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 655–670. URL: [https://doi.org/10.1007/978-3-030-99736-6\\_44](https://doi.org/10.1007/978-3-030-99736-6_44). doi:10.1007/978-3-030-99736-6\_44.
- [26] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, H. Li, Learning to rank: from pairwise approach to listwise approach, in: Proceedings of the 24th International Conference on Machine Learning, ICML '07, Association for Computing Machinery, New York, NY, USA, 2007, p. 129–136. URL: <https://doi.org/10.1145/1273496.1273513>. doi:10.1145/1273496.1273513.
- [27] R. Tang, X. Zhang, X. Ma, J. Lin, F. Ture, Found in the middle: Permutation self-consistency improves listwise ranking in large language models, *CoRR abs/2310.07712* (2023). URL: <https://doi.org/10.48550/arXiv.2310.07712>. doi:10.48550/ARXIV.2310.07712.

arXiv:2310.07712.

- [28] X. Ma, X. Zhang, R. Pradeep, J. Lin, Zero-shot listwise document reranking with a large language model, CoRR abs/2305.02156 (2023). URL: <https://doi.org/10.48550/arXiv.2305.02156>. doi:10.48550/ARXIV.2305.02156. arXiv:2305.02156.
- [29] R. Pradeep, S. Sharifmoghaddam, J. Lin, Rankvicuna: Zero-shot listwise document reranking with open-source large language models, CoRR abs/2309.15088 (2023). URL: <https://doi.org/10.48550/arXiv.2309.15088>. doi:10.48550/ARXIV.2309.15088. arXiv:2309.15088.
- [30] M. S. Tamber, R. Pradeep, J. Lin, Scaling down, litting up: Efficient zero-shot listwise reranking with seq2seq encoder-decoder models, CoRR abs/2312.16098 (2023). URL: <https://doi.org/10.48550/arXiv.2312.16098>. doi:10.48550/ARXIV.2312.16098. arXiv:2312.16098.
- [31] C. Meng, N. Arabzadeh, A. Askari, M. Aliannejadi, M. de Rijke, Ranked list truncation for large language model-based re-ranking, 2024. URL: <https://api.semanticscholar.org/CorpusID:269449617>.
- [32] Z. Qin, R. Jagerman, K. Hui, H. Zhuang, J. Wu, J. Shen, T. Liu, J. Liu, D. Metzler, X. Wang, M. Bendersky, Large language models are effective text rankers with pairwise ranking prompting, CoRR abs/2306.17563 (2023). URL: <https://doi.org/10.48550/arXiv.2306.17563>. doi:10.48550/ARXIV.2306.17563. arXiv:2306.17563.
- [33] S. Zhuang, H. Zhuang, B. Koopman, G. Zuccon, A setwise approach for effective and highly efficient zero-shot ranking with large language models, arXiv preprint arXiv:2310.09497 (2023).
- [34] A. Z. Broder, D. Carmel, M. Herscovici, A. Soffer, J. Y. Zien, Efficient query evaluation using a two-level retrieval process, in: Proceedings of the 2003 ACM CIKM International Conference on Information and Knowledge Management, New Orleans, Louisiana, USA, November 2-8, 2003, ACM, 2003, pp. 426–434. URL: <https://doi.org/10.1145/956863.956944>. doi:10.1145/956863.956944.
- [35] T. Strohman, W. B. Croft, Efficient document retrieval in main memory, in: W. Kraaij, A. P. de Vries, C. L. A. Clarke, N. Fuhr, N. Kando (Eds.), SIGIR 2007: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Amsterdam, The Netherlands, July 23-27, 2007, ACM, 2007, pp. 175–182. URL: <https://doi.org/10.1145/1277741.1277774>. doi:10.1145/1277741.1277774.
- [36] J. Leonhardt, H. Müller, K. Rudra, M. Khosla, A. Anand, A. Anand, Efficient neural ranking using forward indexes and lightweight encoders, ACM Transactions on Information Systems (2023). URL: <https://api.semanticscholar.org/CorpusID:264935522>.
- [37] Y. Qiao, Y. Yang, H. Lin, T. Yang, Optimizing guided traversal for fast learned sparse retrieval, Proceedings of the ACM Web Conference 2023 (2023). URL: <https://api.semanticscholar.org/CorpusID:258333719>.
- [38] T. Nguyen, M. Rosenberg, X. Song, J. Gao, S. Tiwary, R. Majumder, L. Deng, MS MARCO: A human generated machine reading comprehension dataset, CoRR abs/1611.09268 (2016). URL: <http://arxiv.org/abs/1611.09268>. arXiv:1611.09268.
- [39] N. Craswell, B. Mitra, E. Yilmaz, D. Campos, E. M. Voorhees, Overview of the TREC 2019 deep learning track, CoRR abs/2003.07820 (2020). URL: <https://arxiv.org/abs/2003.07820>. arXiv:2003.07820.
- [40] N. Craswell, B. Mitra, E. Yilmaz, D. Campos, Overview of the TREC 2020 deep learning track, in: E. M. Voorhees, A. Ellis (Eds.), Proceedings of the Twenty-Ninth Text REtrieval

Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16-20, 2020, volume 1266 of *NIST Special Publication*, National Institute of Standards and Technology (NIST), 2020. URL: <https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.DL.pdf>.

- [41] N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, I. Gurevych, BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models, *CoRR* abs/2104.08663 (2021). URL: <https://arxiv.org/abs/2104.08663>. arXiv: 2104.08663.
- [42] E. M. Voorhees, T. Alam, S. Bedrick, D. Demner-Fushman, W. R. Hersh, K. Lo, K. Roberts, I. Soboroff, L. L. Wang, Trec-covid, *ACM SIGIR Forum* 54 (2020) 1 – 12. URL: <https://api.semanticscholar.org/CorpusID:218581058>.
- [43] A. Bondarenko, M. Fröbe, M. Beloucif, L. Gienapp, Y. Ajjour, A. Panchenko, C. Biemann, B. Stein, H. Wachsmuth, M. Potthast, M. Hagen, Overview of touché 2020: Argument retrieval, in: A. Arampatzis, E. Kanoulas, T. Tsirikia, S. Vrochidis, H. Joho, C. Lioma, C. Eickhoff, A. Névél, L. Cappellato, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Springer International Publishing, Cham, 2020, pp. 384–395.
- [44] L. Tunstall, E. Beeching, N. Lambert, N. Rajani, K. Rasul, Y. Belkada, S. Huang, L. von Werra, C. Fourrier, N. Habib, N. Sarrazin, O. Sanseviero, A. M. Rush, T. Wolf, Zephyr: Direct distillation of lm alignment, 2023. arXiv: 2310.16944.
- [45] I. Ounis, G. Amati, V. Plachouras, B. He, C. Macdonald, D. Johnson, Terrier information retrieval platform, in: D. E. Losada, J. M. Fernández-Luna (Eds.), *Advances in Information Retrieval, 27th European Conference on IR Research, ECIR 2005, Santiago de Compostela, Spain, March 21-23, 2005, Proceedings*, volume 3408 of *Lecture Notes in Computer Science*, Springer, 2005, pp. 517–519. URL: [https://doi.org/10.1007/978-3-540-31865-1\\_37](https://doi.org/10.1007/978-3-540-31865-1_37). doi:10.1007/978-3-540-31865-1\_37.
- [46] S. Xiao, Z. Liu, Y. Shao, Z. Cao, Retromae: Pre-training retrieval-oriented language models via masked auto-encoder, in: Y. Goldberg, Z. Kozareva, Y. Zhang (Eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, Association for Computational Linguistics, 2022, pp. 538–548. URL: <https://doi.org/10.18653/v1/2022.emnlp-main.35>. doi:10.18653/v1/2022.EMNLP-MAIN.35.