

Long Document Re-ranking with Modular Re-ranker

Luyu Gao

Language Technologies Institute
Carnegie Mellon University
Pittsburgh, PA, USA
luyug@cs.cmu.edu

Jamie Callan

Language Technologies Institute
Carnegie Mellon University
Pittsburgh, PA, USA
callan@cs.cmu.edu

ABSTRACT

Long document re-ranking has been a challenging problem for neural re-rankers based on deep language models like BERT. Early work breaks the documents into short passage-like chunks. These chunks are independently mapped to scalar scores or latent vectors, which are then pooled into a final relevance score. These encode-and-pool methods however inevitably introduce an information bottleneck: the low dimension representations. In this paper, we propose instead to model full query-to-document interaction, leveraging the attention operation and modular Transformer re-ranker framework. First, document chunks are encoded independently with an encoder module. An interaction module then encodes the query and performs joint attention from the query to all document chunk representations. We demonstrate that the model can use this new degree of freedom to aggregate important information from the entire document. Our experiments show that this design produces effective re-ranking on two classical IR collections Robust04 and ClueWeb09, and a large-scale supervised collection MS-MARCO document ranking.¹

CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking**; • **Computing methodologies** → **Neural networks**.

KEYWORDS

Neural IR, Document Re-ranking, Deep Learning

ACM Reference Format:

Luyu Gao and Jamie Callan. 2022. Long Document Re-ranking with Modular Re-ranker. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22), July 11–15, 2022, Madrid, Spain*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3477495.3531860>

1 INTRODUCTION

Deep pre-trained Transformer [19] language models (LM) like BERT [4] have been widely adopted in search and re-ranking and achieve state-of-the-art performance [3, 16, 23]. These models are composed of Transformer layers, which use self-attention operations [19] to contextualize and interact query and document [3, 16].

¹Our code is available at https://github.com/luyug/mores_plus.



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR '22, July 11–15, 2022, Madrid, Spain

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8732-3/22/07.

<https://doi.org/10.1145/3477495.3531860>

Despite their success, Transformer LMs have pre-set input length limits, typically 512, due to complexity considerations [4, 15, 23]. As a result, re-ranking long documents beyond this length has been a challenging problem. Prior research designed techniques to accommodate the neural LM. For example, some attempts to strategically limit the input length, which however prohibits full interaction among all query and document tokens. Others have turned to more efficient Transformers that use lower complexity but also lower capacity sparse self-attention. In this paper, we present an alternative method, re-ranking documents with a modular re-ranker and a new form of attention operation, query to document cross attention.

Modular re-rankers such as MORES [7] provide an alternative framework to common re-rankers. With MORES, query and document are independently encoded. A Transformer Interaction Module then makes relevance estimates using light-weight query-to-document attention. In this paper, we introduce a modification to this framework, MORES+, for long document re-ranking. We break a candidate document into passage-like chunks and encode them separately using identical Transformer encoders. A joint *query-to-all-chunk* attention then models interaction between query and the *full* document. The model can flexibly control the attention weights to pick up important information. This design roughly resembles the human reading comprehension process: after browsing the article and the question, the reader will refer back to the relevant pieces in the original article to come up with answers.

We present experiments with classical few-shot document IR datasets, ClueWeb09 and Robust04, as well as a recent large-scale dataset, MS-MARCO document ranking. We found that the MORES+ re-ranker produces better ranking quality on detailed natural language queries where interaction between query and document is more sophisticated.

The rest of the paper is organized as follows. Section 2 provides a background on document re-ranking methods with Transformers. Section 3 describes the MORES modular re-ranker and how we generalize it to MORES+ for document re-ranking. Section 4 presents experimental results with MORES+.

2 BACKGROUND

The revolution in pre-trained Transformer language models [4, 17] had a large impact on IR [3]. Encoder-only LMs like BERT [4] are the current state-of-the-art for text re-ranking. They can a) generate contextualized representations [17] to help enhance the language understanding capability of the re-rankers and b) perform deep interactions between query and document [3]. Other forms of pre-trained LMs also exist. Encoder-decoder LMs are pre-trained on sequence-to-sequence tasks and have previously been used in IR for tasks like document expansion [12, 18].

Despite their advantages, Transformer LM re-rankers cannot be naively used for long document ranking since most pre-trained LMs have a maximum input length constraint of 512. Prior research circumvented this by splitting documents into passage-like short text chunks. The passages are first separately judged and passage-level results are further pooled to generate the final prediction. Score-pooling systems simply take max/sum/first of the chunk scores [3]. These systems are typically named after their underlying LM and pooling method. For example, BERT-maxp [3] refers to a BERT-based system with max pooling. A generalization is representation aggregation systems like PARADE [13]. They take the low dimension CLS vector representations from each passage and learn an additional aggregation model to combine these representations. The aggregation model can any neural model, e.g. a Transformer in PARADE-Transformer [13]. These pooling systems make it possible to apply deep LMs to document re-ranking. However, they prohibit direct query and full document interaction. The interactions always happen at the passage level and the passage scores or low dimension representations are bottlenecks in the model.

Alternatively, some research has explored the use of long-text Transformers [2] that use sparse attention patterns [10]. Sparse attention comes with the cost of lower capacity than vanilla attention [2].

A third class of approaches tries to first select the important short chunk of document, the “key block”, and feed it into a vanilla Transformer [11, 14]. These models’ success depends on a) the quality of the selected block, and more importantly b) the assumption that a smaller key block exists to represent the full long document.

3 METHODOLOGIES

In this section, we first give some preliminaries on deep re-rankers based on the Transformer LM. We then briefly discuss the MORES modular re-ranker and how we generalize it for document re-ranking.

3.1 Preliminaries

Recent state-of-the-art systems in text re-ranking are based on encoder-only pre-trained Transformer LMs. They cast re-ranking as a sequence pair relation prediction task. Query, document, and a special CLS token are concatenated and fed into the Transformer LM. Query and document go through a series of Transformer layers together, interact, and produce contextualized representations H .

$$H = \text{LM}(\text{concatenate}(\text{CLS}; \text{query}; \text{document})) \quad (1)$$

A final score is computed as a dot product between the representation of the CLS vector H_{cls} and a projection vector \mathbf{v}_{proj} .

$$\text{score} = \mathbf{v}_{\text{proj}}^T H_{cls} \quad (2)$$

Effectively, the Transformer model here serves as a deep non-convex similarity function between query and document. These models, however, have two disadvantages: 1) the aforementioned LM input length constraint, which is typically 512, and, 2) query and document are entangled in a black-box fashion. The former means these models do not natively support long documents. The latter makes it hard to design an easy fix. In this paper, we instead turn to modular re-rankers [7].

3.2 Modular Re-ranker

The Modular Transformer-based re-ranker, or in short MORES, was introduced by Gao et al. [7]. In MORES, query (prepended with a CLS) and candidate documents are first separately encoded by two Transformer encoders.

$$H^d = \text{Encoder}^{\text{doc}}(\text{document}) \quad (3)$$

$$H^q = \text{Encoder}^{\text{qy}}(\text{query}) \quad (4)$$

These representations H^q and H^d remain independent until they are fed into an Interaction Module (IM).

$$H^{in} = \text{IM}(H^q, H^d) \quad (5)$$

The IM is a third Transformer that performs self-attention over the query representation H^q . In addition, query-to-document cross attention is performed in each layer. Given an intermediate query representation q , the cross-attention generates q' .

$$q' = \text{Attend}(q, H^d) \quad (6)$$

This operation interacts query and document and is of complexity linear to query/document length. The final CLS vector is projected to generate a relevance prediction score.

$$\text{score} = \mathbf{v}_{\text{proj}}^T H_{cls}^{in} \quad (7)$$

3.3 Modular Long Document Re-ranker

The flexible design of MORES allows us to further modify the modules to support long documents. Given a long document that does not fit into a typical LM re-ranker, we first break it into chunks, document $\rightarrow \{c_1, \dots, c_n\}$, and encode each chunk independently.

$$H_j^d = \text{Encoder}^{\text{doc}}(c_j) \quad j = 1, \dots, n \quad (8)$$

We modify the Interaction Module to take all or the chunks’ representations at once,

$$H^{in} = \text{IM}(H^q, [H_1^d; H_2^d; \dots; H_n^d]) \quad (9)$$

where a *single joint query-to-all-chunk cross attention* operation is performed over concatenated *all* chunks.

$$q' = \text{Attend}(q, \text{concatenate}(H_1^d; H_2^d; \dots; H_n^d)) \quad (10)$$

This joint attention allows flexible query-document token-level interaction: the query can freely attend to segments in the entire document without chunk boundary limitation. This operation is a generalization of the cross-attention operation in Equation 6. Since the full document is only being *attended to* here, this operation remains a complexity linear to document length. We call this new design MORES+, indicating a modification over MORES to support long document re-ranking.

For familiar readers, we note this design shares a similar spirit with Fusion-in-Decoder [9] which uses sequence-to-sequence model to combine multiple retrieved contexts for question answering.

3.4 Complexity and Efficiency Considerations

Here we compare the time complexities of score-pooling systems, representation aggregation systems, and MORES+. We use BERT-maxp [3] and PARADE-Transformer [13] as concrete examples. We consider a long document of length D , a query of length Q , and a chunk size C for all systems. This gives us $O(D/C)$ chunks. We

	Query Dependent	Query Independent	Total
BERT-maxp	$O(CD + 2QD + Q^2D/C)$	-	$O(CD + 2QD + Q^2D/C)$
PARADE-Transformer	$O(CD + 2QD + Q^2D/C + D^2/C^2)$	-	$O(CD + 2QD + Q^2D/C + D^2/C^2)$
MORES+	$O(QD + Q^2)$	$O(CD)$	$O(CD + QD + Q^2)$

Table 1: Complexity analysis of attention operations in score-pooling system BERT-maxp, representation aggregation system PARADE-Transformer and our proposed MORES+. This considers a length D document, a length Q query, and length C chunks.

focus our analysis here on the (self and cross) attention operations’ cost of the Transformers. Note that the other major cost factor is the feed-forward network, which is point-wisest and remains a cost linear to input length across all systems.

With both score pooling systems (BERT-maxp) and representation aggregation systems (PARADE-Transformer), each document chunk is concatenated with the query and separately fed into a Transformer. Self attention costs $O((C + Q)^2)$ each. A total of $O(CD + 2QD + Q^2D/C)$ for $O(D/C)$ chunks.

The aggregator in score aggregation systems has an extra cost. In PARADE-Transformer, self-attention costs $O(D^2/C^2)$ for latent vectors from $O(D/C)$ chunks.

In MORES+, each of the documents is encoded independently to the query. Self attention costs $O(C^2)$ for each chunk, a total of $O(CD)$. Query encoding costs $O(Q^2)$. Finally, the interaction module performs cross attention from a length Q query to $O(D/C)$ of the length C chunks, or equivalently D document tokens. This costs $O(QD)$ time.

Note that similar to MORES, MORES+ encodes documents independent to the query. These query independent representations can therefore be pre-computed offline. In comparison, both BERT-maxp and PARADE-Transformer start from the concatenated query and document chunks. They, therefore, do not permit any pre-computation. In Table 1, we record the complexities of all three systems. Note that a common situation is that a document is a few times longer than a chunk and a chunk is much longer than a query, i.e. $D > C \gg Q$. With pre-computation, it is possible to move a costly $O(CD)$ term offline. On the other hand, when everything is run online, MORES+ has complexity similar to compared systems.

3.5 Initialization

Encoder-only LMs like BERT [4] and RoBERTa [15] do not perform cross attention and therefore do not contain weights that can initialize the interaction module. In the original MORES paper, the authors propose to borrow self-attention weights but also found the approach to be non-optimal [7]. Observing the emergence of pre-trained sequence-to-sequence encoder-decoder LMs [12, 18], we propose to instead initialize using their weights. In particular, we initialize with BART [12], using its encoder for our document encoder module and its decoder for our interaction module. The decoder-to-encoder weights are used as MORES+’s query-to-document attention weights. To accommodate this new initialization scheme, we reduce the query encoder module down to a single embedding layer and have queries encoded jointly in the interaction process. In other words, the flows of computation in BART are preserved.²

²Details can be found in our open open-source code.

Title	air traffic controller
Description	What are working conditions and pay for U.S. air traffic controllers?

Table 2: Example of Robust04 search topic (Topic 697).

	# Queries	# Docs	# Tokens / Document
Robust04	249	0.5 M	0.7K
ClueWeb09	200	50M	3.3K
MS-MARCO	0.37M	3.2M	1.3K

Table 3: Collection statistics.

4 EXPERIMENT SETUP

4.1 Datasets

We first verify the effectiveness of MORES+ systems on classical few-shot IR datasets. We use Robust04 and Clueweb09 and report nDCG@20, following Dai and Callan [3]. We consider two versions of queries: short keyword queries (title) and natural language query (description). An example provided by Dai and Callan [3] is shown in Table 2.

We also adopt the MS-MARCO document ranking dataset [1]. MS-MARCO document ranking is a recent popular large-scale supervised dataset constructed from Bing with substantially more training queries. Its leaderboard has received submissions from a variety of production-ready systems. Participants are given Train/Dev sets judgments and the submissions are evaluated based on the hidden Eval set performance. MS-MARCO helps demonstrate MORES+’s ability to handle contemporary real-world web search tasks. Queries in MS-MARCO document ranking are of natural language form. We report the official metric MRR@100 on Dev and Eval (leaderboard) query sets.

We show the three datasets’ statistics in Table 3. Note that all three have documents on average longer than the typical LM input length of 512. They help us test the effectiveness of long document re-rankers.

4.2 Implementations

Our MORES+ models are implemented based on BART code in the Huggingface transformer library [20]. The document encoder and interaction module both have 12 layers. We set each input chunk size to 512, similar to the maximum input sequence length of BART during pre-training [12]. We use a maximum of 3 chunks on Robust04 and MS-MARCO and a maximum of 6 chunks on ClueWeb09, preserving up to $\sim 1.5k$ tokens and $\sim 3k$ tokens in each input respectively. The models are trained with localized contrastive

estimation (LCE) loss [8]. Following previous work [10, 13], we first pre-train MORES+ on MS-MARCO passage ranking data to warm up the model.³

We use the commonly adopted Indri initial rankings released by Dai and Callan [3] on Robust04 and Clueweb09. Since one of our major baselines, the previous state-of-the-art system PARADE uses a different set of BM25+RM3 initial rankings on Robust04, we also use MORES+ to re-rank BM25+RM3 initial rankings on Robust04 released by Yang et al. [22]. We perform 5-fold cross-validation on Robust04 and ClueWeb09 as described in Dai and Callan [3].

On MS-MARCO, we follow top leaderboard systems and use a dense initial stage retriever. Our dense retriever is fine-tuned from a pre-trained coCondenser [6]. It combines a special dense pre-training architecture Condenser [5] and contrastive learning to pre-train a tailored model for dense retrieval. It allows us to perform a simple low-cost fine-tuning. For the leaderboard submission, we ensemble 4 models that are trained with different initialization.

4.3 Baseline Systems

On the classical few-shot IR datasets, we compare with score pooling systems BERT-firstP/maxP/sumP [3], and latent representation aggregation system PARADE-CNN/Transformer [13]. We also consider long-text Transformers. QDS-Transformer [10] uses a special attention pattern for ranking. A few universal long-text Transformer baselines, Sparse-Transformer, Longformer, and Transformer-XH are also taken from [10]. A recent method BeST [11] that uses query-based key-block selection is also considered. We also borrow some earlier classical baseline systems from Dai and Callan [3] including SDM, RankSVM and Coord-Ascent as well as pre-BERT neural models DRMM and Conv-KNRM. Due to the high training cost of these long-text models, we directly copy previously reported numbers.

On MS-MARCO, we report the performance of the current top leaderboard systems, as shown in Table 5. Many of them are or are based on real-world production-ready systems [21, 24].

5 EXPERIMENT RESULTS

5.1 Classical Few-shot IR

Table 4 shows the performance of the baseline systems and MORES+ on Robust04 and ClueWeb09. We see that MORES+ achieves the new state-of-the-art in ranking effectiveness.

On Title queries, many neural re-rankers, despite being of higher capacity, do not show a clear advantage over lighter classical systems. Recall that these Title queries are typically of one or a few words, which are simple but ambiguous. As argued by Dai and Callan [3], they typically require less of the language understanding that neural models excel at. Similarly, these keyword-like queries do not require sophisticated query document interaction. In our experiments, we see on Robust04, the best PARADE systems, which pool passage latent vectors into global relevance scores, have similar performance to MORES+⁴.

On the other hand, when moving over to Description queries that are written in natural language, neural models start to show

³Training details are available in our open-source code.

⁴PARADE results should be compared with MORES+^{RM3} entries, the BM25+RM3 results which use the same initial retriever.

Model	nDCG@20			
	Robust04		ClueWeb09	
	Title	Description	Title	Description
Classical Systems				
SDM	0.427	0.427	0.279	0.235
RankSVM	0.420	0.435	0.289	0.245
Coord-Ascent	0.427	0.441	0.295	0.251
Neural Systems				
DRMM	0.422	0.412	0.275	0.245
Conv-KNRM	0.416	0.406	0.270	0.242
Sparse-Transformer	0.449	-	0.274	-
Longformer-QA	0.448	-	0.276	-
Transformer-XH	0.450	-	0.283	-
BERT-firstP [3]	0.444	0.491	0.286	0.272
BERT-maxP [3]	0.469	0.529	0.293	0.262
BERT-sumP [3]	0.467	0.524	0.289	0.261
QDS-Transformer [10]	0.457	-	0.308	-
BeST [11]	0.487	0.537	-	-
PARADE [13]				
- CNN ^{RM3}	0.563	0.610	-	-
- Transformer ^{RM3}	0.566	0.613	-	-
MORES+	0.556	0.612	0.353	0.329
MORES+ ^{RM3}	0.569	0.641	-	-

Table 4: Results on Robust04 and ClueWeb09. Results not available are denoted with ‘-’. RM3: These systems use BM25+RM3 initial rankings [22] instead of Indri initial rankings [3].

their advantages on language understanding. In particular, models that are based on pre-trained LMs have a clear performance margin over classical systems. Meanwhile, the direct interaction between the query and the full long document also proves to be of decent utility. We see that on both Robust04 and ClueWeb09 Description queries, MORES+ has a clear advantage over all previous methods, including score-pooling systems and latent representation pooling systems. As many search engines/portals are taking up more user-friendly natural language form queries, we believe the improvements brought by MORES+ will have further future impacts.

We also see that MORES+ can outperform sparse Transformer based models. Both sparse Transformer re-rankers and MORES+ attempt to preserve granular query-document interaction. Sparse Transformers re-rankers do so by making the self-attention pattern sparse, effectively lowering its cost at the cost of capacity. As discussed by Li et al. [13], these unique patterns of sparse attention are not optimal for re-ranking. In comparison, MORES+ uses another form of query-to-document attention which shows better effectiveness here. On the other hand, sparse Transformers are more efficient than other models. Note that, as discussed in subsection 3.4, modular re-rankers support document representation pre-computations, which can substantially speed up re-ranking by tens of times. We focus on effectiveness impacts in this paper and leave the exploration of efficiency to future work.

In addition, MORES+ also outperforms the key block method BeST. We can conceptually think of the cross attention operation in MORES+ as a soft information selection process. It can freely attend to useful pieces in the document, by assigning higher attention

MS-MARCO Document Ranking Leaderboard		
Model	MRR@100	
	Dev	Eval
UniRetriever	0.500	0.440
Group-HNS-Retrieval+Multi-Granularity-Rerank	0.496	0.436
ANCE MaxP + LongP / SEED-Encoder+LongP (ensemble)	0.487	0.427
hybrid retriever / misc. BERT-longp	0.481	0.424
PROP_step400K base + doc2query top1000(ensemble v0.2)	0.479	0.423
coCondenser + MORES+	0.501	0.436
- single (no ensemble)	0.493	-

Table 5: Results on MS-MARCO document ranking. Top 5 compared runs as of Jan 15th 2022 are recorded.

weights. This soft process does not put hard constraints nor require the pieces to be contiguous. In other words, MORES+ can more flexibly select “key information” beyond “key block”.

5.2 Large-scale Supervised IR


In Table 5, we show performance on the MS-MARCO document ranking dataset. We see that, at the time when this paper is written, MORES+ achieves 2nd best performance on the Eval set and the best on the Dev set. Early work [8] on the MS-MARCO document ranking has shown that a) it requires stronger initial retrievers than BM25, but b) stronger retriever rankings are harder to re-rank. This means, for our discussion, it is important to verify the re-ranker’s improvement when paired with a strong retriever. Here we see that MORES+ can successfully re-rank candidates produced by one of the state-of-the-art dense retrievers, achieving competitive results.

When making further comparisons with other runs, we would like to remind our readers that other leaderboard systems are typically heavily engineered systems that combine various techniques. In comparison, ours is a simple two-stage system with a pair of dense retrievers and a modular re-ranker. In the last row of Table 5, we show Dev set performance of the MORES+ system without ensemble. It can still achieve competitive results and outperform several top leaderboard systems.


5.3 Ablation: Number of Input Chunks

In this section, we consider an ablation study where we input smaller numbers of chunks to MORES+. Ideally, as the number of chunks increases, a good model should be able to pick up more information and generate better final relevance judgments. It should not be confused by extra noise in the input if the added chunks are not relevant. On the other hand, even with less sufficient information, the model should not fail catastrophically but retain a certain level of performance. In practice, reducing the number of input chunks can help lower the total computation amount. It helps lower search latency when search volume increases.

In Figure 1, we plot number of input chunks versus nDCG@20 on Robust04 and ClueWeb09 using description queries. Note here a special case is the single chunk situation, where MORES+ reduces to the original MORES system. This corresponds to the left-most data point in each plot. Here, MORES+ shows monotonic improvements with respect to the number of input chunks while remaining decently effective with a lower number of chunks.



(a) Robust04 performance measured with nDCG@20.



(b) ClueWeb09 performance measured with nDCG@20.

Figure 1: Effect of changing number of input chunks to MORES+. Robust04 has an average length of 0.7K and ClueWeb09 an average length of 3.3K. We used up to 3 chunks on Robust04 and up to 6 on ClueWeb09.

6 CONCLUSION

In this paper, we present MORES+, a new approach for re-ranking long documents. It deviates from previous long document systems that rely on an encoder-only Transformer and is based on a modular re-ranker framework, MORES. We introduce a granular token-level query document interaction method using query to document chunk cross attention. We use this method to augment the original MORES framework for long document re-ranking.

Our experiments demonstrate the effectiveness of MORES+ on both classical IR datasets with few training examples and contemporary large-scale supervised datasets. We found MORES+’s new form of token-level interaction is especially beneficial for natural language queries. It achieves a new state-of-the-art on Robust04 and ClueWeb09 using Description queries, with large improvements over baselines. On MS-MARCO document ranking, MORES+ also shows competitive performance, with second-best Eval and best Dev set performance.

We believe MORES+ opens new possibilities in document re-ranking using modular re-rankers. Broadly, it reminds people that instead of sticking with the original Transformer encoder architecture used in models like BERT, document ranking can also be solved with innovations in architecture. On the other hand, MORES+ shows competitive performance to real-world systems. Future work can explore the possibility of using MORES+ in real search tasks.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) grant IIS-1815528. Any opinions, findings, and conclusions in this paper are the authors’ and do not necessarily reflect those of the sponsor. The authors would like to thank Google’s TPU Research Cloud (TRC) for access to Cloud TPUs and the anonymous reviewers for the reviews.

REFERENCES

- [1] Daniel Fernando Campos, Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, Li Deng, and Bhaskar Mitra. 2016. MS MARCO: A Human Generated MACHine Reading COmprehension Dataset. *ArXiv abs/1611.09268* (2016).
- [2] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509* (2019).
- [3] Zhuyun Dai and Jamie Callan. 2019. Deeper Text Understanding for IR with Contextual Neural Language Modeling. In *The 42nd International ACM SIGIR Conference on Research & Development in Information Retrieval*.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.
- [5] Luyu Gao and Jamie Callan. 2021. Condenser: a Pre-training Architecture for Dense Retrieval. In *EMNLP*.
- [6] Luyu Gao and Jamie Callan. 2021. Unsupervised Corpus Aware Language Model Pre-training for Dense Passage Retrieval. *CoRR abs/2108.05540* (2021). <https://arxiv.org/abs/2108.05540>
- [7] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2020. Modularized Transformer-based Ranking Framework. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online. <https://doi.org/10.18653/v1/2020.emnlp-main.342>
- [8] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2021. Rethink Training of BERT Rerankers in Multi-Stage Retrieval Pipeline. In *ECIR*.
- [9] Gautier Izacard and Edouard Grave. 2021. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In *EACL*.
- [10] Jun-Yu Jiang, Chenyan Xiong, Chia-Jung Lee, and Wei Wang. 2020. Long Document Ranking with Query-Directed Sparse Transformer. *ArXiv abs/2010.12683* (2020).
- [11] Youngwoo Kim, Razieh Rahimi, Hamed Bonab, and James Allan. 2021. Query-driven Segment Selection for Ranking Long Documents. *Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (2021).
- [12] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. *ArXiv abs/1910.13461* (2020).
- [13] Canjia Li, Andrew Yates, Sean MacAvaney, Ben He, and Yingfei Sun. 2020. PA-RADE: Passage Representation Aggregation for Document Reranking. *ArXiv abs/2008.09093* (2020).
- [14] Minghan Li and Éric Gaussier. 2021. KeyBLD: Selecting Key Blocks with Local Pre-ranking for Long Document Information Retrieval. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2021).
- [15] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [16] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv:1901.04085* (2019).
- [17] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [18] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *ArXiv abs/1910.10683* (2019).
- [19] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*.
- [20] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv abs/1910.03771* (2019).
- [21] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. *ArXiv abs/2007.00808* (2021).
- [22] Wei Yang, Kuang Lu, Peilin Yang, and Jimmy J. Lin. 2019. Critically Examining the "Neural Hype": Weak Baselines and the Additivity of Effectiveness Gains from Neural Ranking Models. *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2019).
- [23] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, 5754–5764.
- [24] Jianjin Zhang, Zheng Liu, Weihao Han, Shitao Xiao, Rui Zheng, Yingxia Shao, Hao Sun, Hanqing Zhu, Premkumar Srinivasan, Denvy Deng, Qi Zhang, and Xing Xie. 2022. Uni-Retriever: Towards Learning The Unified Embedding Based Retriever in Bing Sponsored Search. *ArXiv abs/2202.06212* (2022).