

A Self-supervised Joint Training Framework for Document Reranking

Xiaozhi Zhu¹, Tianyong Hao¹, Sijie Cheng³, Fu Lee Wang⁴, and Hai Liu^{1,2*}

¹School of Computer Science, South China Normal University, China

²Guangzhou Key Laboratory of Big Data and Intelligent Education, Guangzhou 510000, China

³Fudan University, China, ⁴Hong Kong Metropolitan University, China

{2020022975, haoty}@m.scnu.edu.cn

815106263@163.com, pwang@hkmu.edu.hk, namelh@gmail.com

Abstract

Pretrained language models such as BERT have been successfully applied to a wide range of natural language processing tasks and also achieved impressive performance in document reranking tasks. Recent works indicate that further pretraining the language models on the task-specific datasets before fine-tuning helps improve reranking performance. However, the pre-training tasks like masked language model and next sentence prediction were based on the context of documents instead of encouraging the model to understand the content of queries in document reranking task. In this paper, we propose a new self-supervised joint training framework (SJTF) with a self-supervised method called Masked Query Prediction (MQP) to establish semantic relations between given queries and positive documents. The framework randomly masks a token of query and encode the masked query paired with positive documents, and use a linear layer as a decoder to predict the masked token. In addition, the MQP is used to jointly optimize the models with supervised ranking objective during fine-tuning stage without an extra further pre-training stage. Extensive experiments on the MS MARCO passage ranking and TREC Robust datasets show that models trained with our framework obtain significant improvements compared to original models.

1 Introduction

The document ranking task is to generate a ranked list of candidate documents based on their relevance scores to a given query posed in natural language, which has been a longstanding problem that has been widely studied over natural language processing (NLP) and question answering. Pre-trained language models (PLMs) such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and ELECTRA (Clark et al., 2020),

have achieved impressive results on various NLP tasks and have outperformed conventional document ranking methods (Hui et al., 2018; Mitra and Craswell, 2019) for powerful contextual representation capability. In recent years, several studies (Karpukhin et al., 2020; Qu et al., 2021) have used pre-trained language models as dual-encoder to separately encode queries and documents for dense document retrieval. One of the most common approaches (Nogueira and Cho, 2019) uses a PLMs as an interaction-based reranker for passage ranking, which fine-tunes BERT simply with an extra linear layer on the top of BERT and using a special vector [CLS] to produce relevance score for each query document pair. Inspired by the fact that contextualized embeddings produced by PLMs are essential for the success of pre-trained models, CEDR (MacAvaney et al., 2019) employed the classification vector into existing neural ranking models and PARADE (Li et al., 2020) used a transformer module for passage-level representation aggregation to obtain performance improvements.

In contrast to the approaches of using contextual representation for reranking, prior works (Sun et al., 2019; Gururangan et al., 2020) suggest that further pre-train the PLMs on within-task training unsupervised data is able to learn domain-specific and task-specific language patterns effectively. To better understand the complex sentence relations, UED (Yan et al., 2021) transformed original next sentence prediction (NSP) task in BERT to a new sentence relation prediction (NSR) task. Gu et al. (2020) proposed a novel selective masking strategy to focus on masking the important tokens and then train a model to reconstruct input for further pre-training the PLMs to learn task-specific patterns. However, these approaches typically perform the pre-training task on task-specific corpus to understand the context of passages, while fail to consider the passages as the context of the given query to capture semantic consistency. In addition, further

* Corresponding author: Hai Liu.

pre-training on task-specific domain datasets entails additional time cost and computational cost.

Therefore, this paper proposes a self-supervised joint training framework for encouraging the model to understand the content of queries based on the context of passages, where the auxiliary self-supervised method is combined with the ranking task to fine-tune pre-trained model. Specifically, the self-supervised joint training framework (SJTF) extends the typical reranking pipeline with the auxiliary self-supervised method MQP and a decoder for predicting the masked token after the pre-trained model. On the one hand, the self-supervised approach enables the model to establish semantic relations between queries and positive passages to better identify relevant passages from a large number of candidate passages. On the other hand, the proposed training framework reduces the training time by simultaneously performing the self-supervised method and the ranking task in the fine-tuning stage, while the time of the ranking model is not increased in the inference stage. We evaluate the proposed training framework on two widely used document ranking datasets MS MARCO and Robust04. The experimental results indicate that the models trained with the proposed SJTF framework obtain a performance improvement against original models.

In summary, the contributions of our paper are as follows:

- A self-supervised joint training framework (SJTF) is proposed to improve the representation learning without additional pre-training.
- A strategy to integrate the SJTF to existing passage reranking methods is proposed without architecture modification and inference time increasing.
- Experiments on standard datasets show that reranking models with SJTF integration achieve significant performance improvement.

2 Related Work

In recent years, pretrained language models such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and ELECTRA (Clark et al., 2020) had substantially outperformed the traditional neural ranking models like DRMM (Guo et al., 2016), Co-PACRR (Hui et al., 2018) and Conv-KNRM (Dai et al., 2018). Nogueira and Cho (2019) first

employed the pretrained language model BERT to passage reranking tasks using the classification vector [CLS]. CEDR (MacAvaney et al., 2019) incorporated the contextualized embeddings of BERT into existing IR models for document ranking. PARADE (Li et al., 2020) utilized Transformer (Vaswani et al., 2017) blocks to aggregate relevance passage-level representations to predict a document ranking. These approaches provided different perspectives on score prediction using the relevance representation vectors produced by pre-trained language models.

To better improve the representation learning of pre-trained language models in the target domain, Sun et al. (2019) further pre-trained BERT with masked language model (MLM) and next sentence prediction (NSP). Gururangan et al. (2020) observed that the less relevant the pre-trained corpus was to the target corpus, the more the pre-trained language model would benefit from further pre-training. In this paper, the proposed joint training framework could be used in conjunction with those reranking models in the fine-tuning phase, which reduced the time spent in further pre-training.

Representation learning has been shown to be critical on natural language tasks and has a significant influence on downstream tasks (Devlin et al., 2019; Peters et al., 2018; Yan et al., 2021). Devlin et al. (2019) adopted the self-supervised MLM task to encourage the model to learn contextual representations by predicting the masked token from the context. From the perspective of optimizing the masking strategy, RoBERTa (Liu et al., 2019) modified the static masking strategy by dynamically masking the input examples during the training stage, while (Gu et al., 2020) proposed a selective masking strategy that masked important words rather than any word in the sentence.

For understanding the content of documents, Cross-Thought (Wang et al., 2020) proposed a method for recovering masked words from documents that contain the most important information in a nearby sequence. However, understanding the content of both the query and the document is crucial in the question answering tasks (Zhang et al., 2019; Mudrakarta et al., 2018; Nogueira et al., 2019b), while these methods focused on representation learning for understanding the content of the document. The aim of our self-supervised task MQP is to predict the masked query word based on the semantic consistency of the query and the

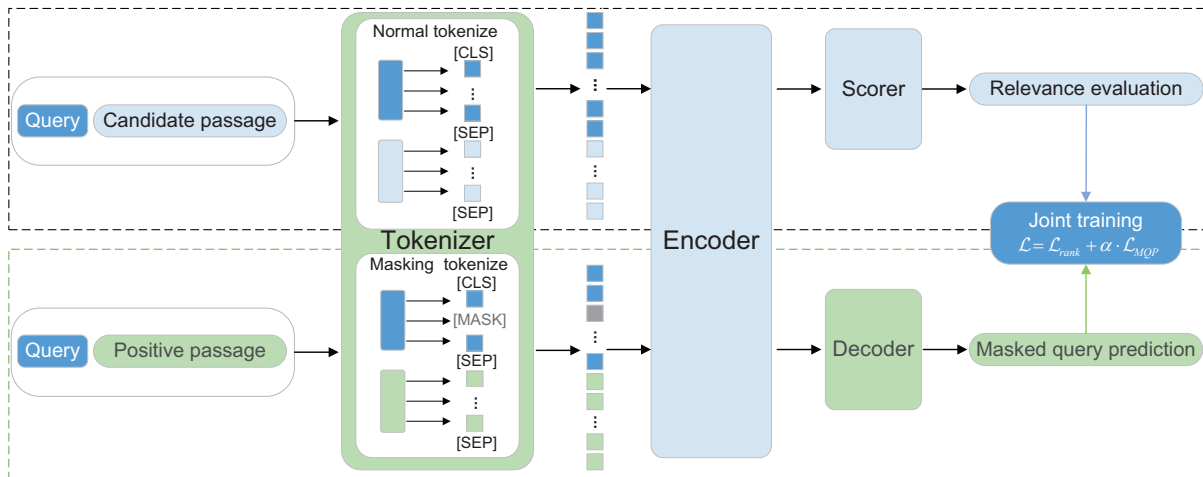


Figure 1: The architecture of joint training with Masked Query Prediction (MQP) task during fine-tuning stage.

positive passages, which forces the model to consider the positive passages as the context of the given query. The MQP differs from the standard self-supervised approaches of using unsupervised data (He et al., 2021; Hendrycks et al., 2019; You et al., 2021) in that the supervised signal is used to select relevant passage as the context of the query.

Multi-task learning (MTL) is an effective training setting for allowing model to obtain shared knowledge from several related supervised tasks in document ranking. Fun et al. (2021) enhanced the common representation learning using a retrieval optimized multi-task framework (ROM) for jointly training the retriever, reader and self-supervised tasks with a single encoder. UED (Yan et al., 2021) jointly trained both the ranking and query generation tasks to exploit the task relationships for enhancing the neural re-ranker. Liu et al. (2019) and Maillard et al. (2021) leveraged supervised data from related tasks to enhance the robustness of the model and generic knowledge representation learning. Although the datasets of related tasks are available, there may be differences in language patterns and data distribution between datasets. Our proposed joint training framework uses relevance labels to construct data on the target domain of ranking tasks without the requirement of external data, which is achievable on any of the question answering datasets.

3 The Approach

This section describes the proposed self-supervised joint training framework (SJTF), which employs a self-supervised method (MQP) to jointly fine-tune reranking models with the ranking task. In

general, the whole model consists of three parts: a pretrained encoder for producing interaction-based representation of the given query and passages, a scorer for calculating a precise relevance score for each query-passage pair and an extra decoder to predict the masked query token based on the interaction-based representation. The overall framework of our approach is shown in Figure 1.

In the task of passage reranking, a natural language question and a list of candidate passages retrieved by traditional methods or dense retriever are provided. The question is denoted as a k -length sequence of tokens $Q = \langle q_1, q_2, \dots, q_k \rangle$, while each candidate passage can be denoted by an m -length sequence of tokens $P = \langle p_1, p_2, \dots, p_m \rangle$. The passage reranking task requires the model to learn informative representation and produce a precise relevance score for each query-passage pairs to return the best permutation of candidate passages.

3.1 Radom Masking

In contrast to the MLM task that random masks the tokens of the passage, we assume that understanding the content of the given query is necessary and the semantic information of the positive passage can be used to infer the masked token in the query, which finally allows the model to learn the semantic relation between the query and the positive passage. The tokens to be masked in the query are selected randomly following a uniform distribution and replaced with special token [MASK], and for simplicity, each query token is considered here to have the same importance. Then the masked query and positive passage are concatenated

and preprocessed into a sequence $I_{mask} = \langle [CLS], q_1, \dots, [MASK], \dots, q_k, [SEP], p_1, \dots, p_m, [SEP] \rangle$, where special token $[CLS]$ indicates the start of a sentence, token $[MASK]$ indicates the masked token and token $[SEP]$ is a separator symbol.

The reason for masking the query is that the original query can be reconstructed by understanding the content of the processed query and the positive passage. However, if the goal is to predict the masked tokens in the passage, the semantic information of the query is redundant as it can be inferred from the context of the passages alone. Without requiring extra data augmentation, the framework SJTF only utilizes the relevance label and a simple masking strategy to construct the masked input sequence I_{mask} , which can be easily implemented in question-answering tasks.

3.2 Masked Query Prediction

The pre-trained language model BERT (Devlin et al., 2019) uses the MLM task to learn contextual language representations of individual texts in a large corpus. However, establishing contextual semantic relations between queries and candidate passages is critical in the question-answering domain. To achieve this, we propose a self-supervised auxiliary approach called Masked Query Prediction (MQP) so that the model uses positive passages as the context of the query and predicts masked token in conjunction with visible query tokens, which allows the model to extract semantic relation between the query and positive passage, thus pick relevant passages out of a large number of candidate passages.

After passing the masked input sequence I_{mask} constructed in section 3.1 through the BERT-like encoder which is shared with a passage reranking task, the representation vector $T_{mask} \in \mathbb{R}^d$ of the masked query token in the last layer is obtained as:

$$T_{masked} = Encoder(q_{masked}, p_{pos}) \quad (1)$$

Finally, the masked token representation vector T_{masked} is fed into a decoder implemented by a neural linear layer to predict the original token. The parameters of MQP module are optimized by the cross-entropy loss function \mathcal{L}_{MQP} which is defined as :

$$\mathcal{L}_{MQP} = - \sum_{m=1}^M \log P(t_m), \quad (2)$$

where M is the number of masked token and $P(t_m)$ denotes the probability that the token t_m is predicted over the whole vocabulary. For each query and positive passage pair, a single query token is replaced with the special token $[MASK]$ for the self-supervised method. Note that the MQP decoder is only used for predicting the masked token during the fine-tuning stage, while in the inference phrase only the encoder and scorer are used for reranking passages. The decoder for the MQP task is implemented with one linear layer, meaning that this requires only a small number of additional neural network parameters to be trained, which makes the MQP task easy to be extended to the existing passage reranking methods without significant modifications to the model architecture.

3.3 Passage Ranking

Following the settings of Nogueira and Cho (2019), a pair of query q and candidate passage p_i is packed as an input sequence $I = \langle [CLS], q_1, \dots, q_k, [SEP], p_1, \dots, p_m, [SEP] \rangle$, and the BERT-like pretrained language model is employed as a passage encoder E that produces a relevance representation vector for each QA pair. During fine-tuning the model, the $[CLS]$ vector $T_{cls} \in \mathbb{R}^d$ from the last layer of encoder is regarded as the final interaction-based representation:

$$T_{cls}^i = Encoder(q, p_i) \quad (3)$$

For passage reranking task, the representation vector T_{cls} of query-passage pairs is calculate by a scorer S_{ranker} which generates a relevance score to quantify their relevance. The relevance score of i -th pairs of query q and candidate passage P_i is denoted as:

$$Score_{cls}^i = S_{ranker}(T_{cls}^i; \theta), \quad (4)$$

where the scorer S_{ranker} can be implemented by a linear layer at top of the BERT or by an elaborate scoring module such as KNRM (Dai et al., 2018), SAN (Kingma and Ba, 2015) or PARADE (Li et al., 2020), and θ contains the set of parameters of the scorer module.

Compared with the point-wise ranking loss used in (Nogueira and Cho, 2019), the pair-wise margin ranking loss discriminates the positive and negative examples by relative distance, allowing the model to learn the margins between the positive and negative examples to give an appropriate relevance score. Therefore, the reranking module is

optimized by the ranking loss \mathcal{L}_{rank} as Equation (5):

$$\mathcal{L}_{rank} = \max(0, y \cdot (Score_{cls}^k - Score_{cls}^{k+1}) + \gamma), \quad (5)$$

where $y = -1$ if the $Score_{cls}^k$ is higher than $Score_{cls}^{k+1}$, and vice-versa for $y = 1$. γ is a hyperparameter that controls the margin of positive and negative examples.

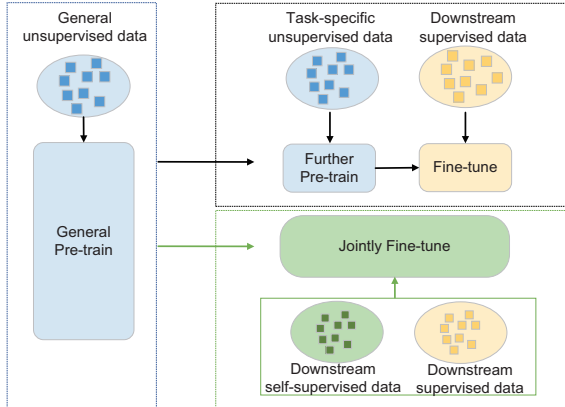


Figure 2: The overall self-supervised jointly training framework SJTF.

3.4 Joint Training

Different from the self-supervised tasks that are used for further pre-training on the downstream dataset (Fun et al., 2021; Liu et al., 2019), our proposed self-supervised method MQP is combined with the ranking task in the fine-tuning phase, since it utilizes the relevance label information to construct masked input sequences. As shown in Figure 2, the joint training strategy simplifies the training procedure without further pre-training phases and training resources required. In the fine-tuning stage, the loss is defined as a linear combination of passage reranking loss and masked query prediction loss as:

$$\mathcal{L} = \mathcal{L}_{rank} + \alpha \cdot \mathcal{L}_{MQP} \quad (6)$$

The hyperparameter α is assigned by different values to tradeoff between passage reranking and masked query prediction. As the MQP is a secondary task used to help the model understand the query content, where the loss weight α is usually set to a lower value, resulting in the parameters of the reranking model still being optimized primarily by the passage reranking task.

4 Experiments and Results

4.1 Dataset

The proposed method is extended to existing ranking models and evaluated on two widely used datasets: MS MARCO Passage Ranking (Nguyen et al., 2016) and TREC Robust 2004 (Voorhees). The statistics of these two datasets are shown in Table 1.

MS MARCO Passage Ranking dataset is a large-scale dataset consisting of real anonymous questions from the Bing search engine and 8.8 million candidate passages for passage reranking task. The training set contains about 500 thousand positive query-passage pairs and each query has one relevant passage on average. The development set and evaluation set contains 6980 queries and 6837 queries respectively, where the relevance labels are provided for the development set only.

Robust04 dataset is a newswire collection used by TREC 2004 Robust track, which comprises 250 queries and 0.5 million documents (TREC Disks 4 and 5). Following the setting of CEDR (MacAvaney et al., 2019), we use the same five folds cross-validation with three folds for training, one fold for validation and one fold for testing.

4.2 Evaluation Metrics

Following the previous works, three widely used evaluation metrics are adopted to measure the performance of the proposed approach, including MRR@10 for the MS MARCO Passage Ranking dataset, P@20 and NDCG@20 evaluated by trec_eval¹ for the Robust04 dataset. The result reported for model performance are averaged over all test folds on the Robust04 dataset.

MRR@10(Mean Reciprocal Rank) This metric considers the reciprocal rank of the first relevant passage in ranked list to a given query as the precision. For MS MARCO passage ranking task only provides binary label and does not specify relative ranking order between passages, thus the MRR metric is used for evaluation.

P@20 The top-20 precision is defined as the proportion of relevant documents which are ranked in the top 20 candidate documents.

NDCG@20 NDCG is used to measure the discrepancy between the ranked list and the correct ranking list, which evaluates the ranking performance of models.

¹https://trec.nist.gov/trec_eval

Table 1: The statistics of datasets MS MARCO Passage Ranking and TREC Robust04.

Datasets	Queries	Avg. of word length	Passages	Avg. of word length
MS MARCO	516,756	5.97	8,841,823	56.58
Robust04	250	2.65	528,155	577.82

4.3 Baselines

The proposed training framework is integrated with the following methods and compared the reranking results with the original methods:

BM25+BERT (Nogueira and Cho, 2019) utilizes the traditional unsupervised ranking method BM25 as a first-stage retriever to generate a ranked list of candidate passages, and the relevance scores between query and candidate passages are produced by the BERT-base model with a linear layer.

BM25+BERT+FP (Gururangan et al., 2020) further pretrains the language model BERT with MLM objective on the target datasets before fine-tuning to learn the specific domain language patterns for document reranking task.

BERT+PACRR (MacAvaney et al., 2019) constructs query-document term similarity matrix in PACRR (Hui et al., 2018) using the interaction-based representation vectors generated by BERT.

PARADE (Li et al., 2020) aggregates passage-level relevance representations to predict a document relevance score, where the long document is split into several passages and each passage is encoded with a given query.

CEDR-KNRM (MacAvaney et al., 2019) incorporates the classification vector of a fine-tuned BERT into existing neural models KNRM (Dai et al., 2018) and leverage contextual information to improve ad-hoc document ranking.

4.4 Implementation Detail

For both datasets, following the setting of the first retrieval stage (Nogueira et al., 2019a), we employ BM25 (Robertson et al., 2009) as the retriever at the first stage to obtain a list of top-k candidate documents/passages for next reranking stage. The reranking models are optimized by Adam (Kingma and Ba, 2015) with a learning rate of $3e-6$ and a batch size of 8. The maximum sequence length is limited to 128 tokens for MS MARCO passage ranking dataset, while the documents are truncated to 800 tokens for Robust04 dataset. The encoder of CEDR and PARADE are the base-size pre-trained language model which consists of 12 transformer blocks, the hidden size as 768, and 12 self-attention

heads. For Robust04 dataset, the dropout function is used with the rate of 0.1 for improving model robustness. We set the value of margin hyperparameter γ of pairwise ranking loss to 1. A uniform distribution is used to random mask a token in a given query, where all query tokens are assumed to be of equal importance. Since the MQP loss serves as an auxiliary loss to enhance the interactive-based representation learning for encoder, the weight hyperparameter α is assigned to 0.2 for MS MARCO and 0.05 for Robust04 respectively. The experiments are conducted on a single GeForce RTX 3090 GPU.

4.5 Results

The experimental results of model+SJTF on the MS MARCO passage ranking and Robust04 datasets are presented in Table 2, where model+SJTF refers to the joint training of different reranking models applying to SJTF during the fine-tuning stage. From Table 2, the reranking models adopting SJTF have obtained significant improvements using paired t-test ($p < 0.05$) compared to original models.

Compared with BM25+BERT-base, the further pretraining method BM25+BERT+FP improves MRR@10 by 0.5% on the MS MARCO and NDCG@20 by 0.8% on the Robust04 datasets for better capturing the domain-specific language patterns. By jointly training the BERT-base model with SJTF framework during the fine-tuning stage, BERT+SJTF achieves a 1.3% improvement in terms of MRR@10 on the MS MARCO and 1.1% improvement of NDCG@20 on the Robust04, The experimental results suggest that learning the semantic dependencies of queries and relevant documents is more effective than further learning the contextual semantic information of documents. BERT+PACRR+SJTF outperforms the original approach by 0.5% on the MS MARCO dataset and 1.3% on the Robust04 dataset, which indicates that the performance of traditional reranking methods would benefit from an improved contextual representation. The results show that the models fine-tuned with the joint training framework yields a

Table 2: Performance comparison of all methods on the MS MARCO Passage Ranking and Robust04. '-' means that the PARADE method is not applicable because PARADE aggregates information from multiple overlapping passages. Significant paired t-test ($p < 0.05$) improvements over the original models are marked with †.

Method	MS MARCO Passage Ranking		TREC Robust04	
	MRR@10(Dev)	MRR@10(Eval)	P@20	NDCG@20
BM25+BERT	0.345	0.342	0.4042	0.4541
BM25+BERT+FP	0.355	0.347	0.4024	0.4624
BM25+BERT+SJTF	0.362†	0.355	0.4291†	0.4657†
BERT+PACRR	0.328	0.326	0.3969	0.4599
BERT+PACRR+SJTF	0.333†	0.331	0.4122†	0.4735†
PARADE	-	-	0.4604	0.5399
PARADE+SJTF	-	-	0.4701†	0.5455
CEDR-KNRM	0.337	0.331	0.4729	0.5388
CEDR-KNRM+SJTF	0.347†	0.337	0.4851†	0.5565†

better interaction-based representation based on the semantic relationships between the query and candidate passages, which leads to a significant performance improvement without the model structure change.

The improvement of PARADE by splitting a long document into overlapping passages via sliding windows is less, increased by 0.55% on Robust04 datasets, suggesting that the semantic blocks in long documents are not all tightly semantically associated with queries and using these to predict masked query tokens only slightly enhances representation learning. In comparison with the results of PARADE, CEDR has been significantly improved by applying the SJTF framework, which indicates that the document-level semantic relations captured by the SJTF framework are beneficial for representation learning. The overall experimental results verify the effectiveness and generality of the training framework of encouraging the models to adopt the positive passage as the semantically consistent context for a given query.

4.6 Ablation Study

Inspired by the observations from previous works (Gu et al., 2020; Liu et al., 2019; Gururangan et al., 2020) that different masking strategies enable a model to learn various perspective knowledge during pre-training stage, we design three different self-supervised methods in terms of focusing on understanding the content of query or passages and using the query or passage as the context for masking prediction task.

Similar to the traditional MLM approach, we randomly mask a token of a positive document and

Table 3: The performance of BERT-base and CEDR-KNRM models by integrating different masking strategies on Robust04 dataset.

Method	NDCG@20	P@20
BM25+BERT	0.4541	0.4042
+ MPP	0.4539	0.4026
+ MQP _{ALL}	0.4567	0.3979
+ MQP	0.4657	0.4291
CEDR-KNRM	0.5388	0.4729
+ MPP	0.5433	0.4738
+ MQP _{ALL}	0.5464	0.4784
+ MQP	0.5565	0.4851

use a corresponding query and the remaining documents as contexts to predict the masked token, and this strategy is denoted as MPP. The aim of MPP is to encourage the model to learn task-specific language patterns and understand that the corresponding query and the positive documents are semantically consistent. However, it is possible to reconstruct the masked token based on the content of document itself without demanding the semantic information of the query. In addition, since MPP is a simple task for pre-trained models, it cannot effectively improve the ranking performance of models.

Considering that the documents retrieved by BM25 are relevant or partially relevant to the query at the word level, the second strategy is to treat each retrieved document (regardless of positive and negative documents) as the context of the masked query for the MQP task, which does not utilize the supervised signal and is denoted as MQP_{ALL}. In contrast to the first strategy, MQP_{ALL} allows the

models to focus on understanding the content of query and capture the semantic relationships between queries and documents. Table 3 illustrates the slight improvement achieved by MQP_{ALL} compared to the original model and the first strategy.

The self-supervised method MQP proposed in this paper is based on the assumption that although the documents retrieved by BM25 have word-level relevance to the query, the documents containing the query terms are not always relevant. Therefore, we use supervised signals to consider only positive documents as the context of masked queries, since the semantic information of negative documents is equal to noise for the queries. As shown in table 3, after filtering out negative documents by supervised signals, BERT-base and CEDR-KNRM achieve a significant improvement on NDCG@20 and P@20 metrics. Experimental results illustrate that the reasonable use of label informations to construct a self-supervised approach enables the models to obtain more accurate knowledges of the downstream task, which helps the models to identify correct documents for each question from a large collection.

4.7 Hyper-parameter Study

As a hyper-parameter described in Equation (6), the setting of α value controls the influence of the self-supervised approach on the representation learning of models. Figure 3 shows the result of different hyper-parameters setting on the Robust04 dataset. In general, a larger value α allows an excessive effect of the self-supervised method, which reduces the influence of the reranking task in optimizing models while achieving sub-optimal performance. Therefore, with α value greater than 0.05, the ranking performance of all three methods show a decreasing trend. For the CEDR method, a stable performance improvement is obtained using the proposed training framework, indicating that considering the positive documents as the contextual features of the query can significantly enhance the representation learning of the CEDR model. By gradually raising the α value from 0.01 to 0.05, the performance of both BERT-base and BERT+PACRR also gradually achieves the highest NDCG@20, which suggests that encouraging these models to focus on query understanding can lead to better reranking performance. From Figure 3, it can be seen that the α value of 0.05 is more robust for most of the models.




Figure 3: Performance of various rerankers with different configurations of α on the Robust04 dataset using the metric nDCG@20.

5 Conclusion

This paper proposes a self-supervised joint training framework SJTF to improve the representation learning for document ranking. By integrating with this training framework, models are able to better understand the content of queries and capture the semantic relation between queries and positive documents, which guides the models to identify the relevant documents among a large number of candidate documents. Experimental results show that our proposed framework enhances the representation learning of reranking models and achieves better performance compared with baselines in the document reranking task.

Acknowledgements

We thank all anonymous reviewers for their insightful comments to import this paper. This work is supported by Natural Science Foundation of Guangdong Province (No. 2021A1515011339) and Guangzhou Key Laboratory of Big Data and Intelligent Education (No. 201905010009).

References

- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. [Convolutional neural networks for soft-matching n-grams in ad-hoc search](#). In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018*,

- Marina Del Rey, CA, USA, February 5-9, 2018, pages 126–134. ACM.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hengxin Fun, Sunil Gandhi, and Sujith Ravi. 2021. Efficient retrieval optimized multi-task learning. *arXiv preprint arXiv:2104.10129*.
- Yuxian Gu, Zhengyan Zhang, Xiaozhi Wang, Zhiyuan Liu, and Maosong Sun. 2020. [Train no evil: Selective masking for task-guided pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6966–6974, Online. Association for Computational Linguistics.
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. [A deep relevance matching model for ad-hoc retrieval](#). In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, pages 55–64. ACM.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2021. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*.
- Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. 2019. [Using self-supervised learning can improve model robustness and uncertainty](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 15637–15648.
- Kai Hui, Andrew Yates, Klaus Berberich, and Gerard de Melo. 2018. [Co-pacrr: A context-aware neural IR model for ad-hoc retrieval](#). In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, February 5-9, 2018*, pages 279–287. ACM.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Canjia Li, Andrew Yates, Sean MacAvaney, Ben He, and Yingfei Sun. 2020. [Parade: Passage representation aggregation for document reranking](#). *arXiv preprint arXiv:2008.09093*.
- 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*.
- Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. 2019. [CEDR: contextualized embeddings for document ranking](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 1101–1104. ACM.
- Jean Maillard, Vladimir Karpukhin, Fabio Petroni, Wen-tau Yih, Barlas Oguz, Veselin Stoyanov, and Gargi Ghosh. 2021. [Multi-task retrieval for knowledge-intensive tasks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1098–1111, Online. Association for Computational Linguistics.
- Bhaskar Mitra and Nick Craswell. 2019. [An updated duet model for passage re-ranking](#). *CoRR*, abs/1903.07666.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. [Did the model understand the question?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1896–1906, Melbourne, Australia. Association for Computational Linguistics.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human generated machine reading comprehension dataset. In *CoCo@ NIPS*.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. [Passage re-ranking with BERT](#). *CoRR*, abs/1901.04085.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019a. [Multi-stage document ranking with BERT](#). *CoRR*, abs/1910.14424.

- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019b. [Document expansion by query prediction](#). *CoRR*, abs/1904.08375.
- 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, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847, Online. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In *China National Conference on Chinese Computational Linguistics*, pages 194–206. Springer.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ellen M Voorhees. Overview of the trec 2004 robust retrieval track.
- Shuohang Wang, Yuwei Fang, Siqi Sun, Zhe Gan, Yu Cheng, Jingjing Liu, and Jing Jiang. 2020. [Cross-thought for sentence encoder pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 412–421, Online. Association for Computational Linguistics.
- Ming Yan, Chenliang Li, Bin Bi, Wei Wang, and Songfang Huang. 2021. A unified pretraining framework for passage ranking and expansion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4555–4563.
- Chenyu You, Nuo Chen, and Yuexian Zou. 2021. Self-supervised contrastive cross-modality representation learning for spoken question answering. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 28–39.
- Xiaoming Zhang, Mingming Meng, Xiaoling Sun, and Yu Bai. 2019. Factqa: question answering over domain knowledge graph based on two-level query expansion. *Data Technologies and Applications*.