

# Learning Diverse Document Representations with Deep Query Interactions for Dense Retrieval

Anonymous ACL submission

## Abstract

In this paper, we propose a new dense retrieval model which learns diverse document representations with deep query interactions. Our model encodes each document with a set of generated pseudo-queries to get query-informed, multi-view document representations. It not only enjoys high inference efficiency like the vanilla dual-encoder models, but also enables deep query-document interactions in document encoding and provides multi-faceted representations to better match different queries. Experiments on several benchmarks demonstrate the effectiveness of the proposed method, outperforming strong dual encoder baselines.

## 1 Introduction

Document retrieval plays an important role in information retrieval (IR) tasks such as web search and open domain question answering (Chen et al., 2017). Early works such as BM25-based retriever (Robertson and Zaragoza, 2009) rely on lexical term matching to calculate the relevance of a pair of texts. Recently, neural network based dense retrieval (Karpukhin et al., 2020) has gained traction in research community. Dense retrieval learns a neural encoder to map queries and documents into a dense, low-dimensional vector space, and is less vulnerable to term mismatch problem compared to lexical match-based approaches.

There are two architectures to model the relevance between queries and documents. Dual encoder architecture encodes query and document separately into fixed-dimensional vectors (Karpukhin et al., 2020), where the similarity between query and document is usually instantiated as a dot product or cosine similarity between their vectors. As there are no interactions between query and document, dual encoder approach permits efficient inference with vector space search on pre-computed document vectors. Cross encoder architecture feeds the concatenation of a query and

document pair into one encoder to calculate its relevance score (Nogueira and Cho, 2019). Compared to dual encoder, cross encoder is more accurate due to the deep interaction between query and document, but comes with computation costs infeasible for first-stage retrieval. It is highly desirable to design a retrieval model which can match the performance of the cross encoder approach while maintaining the inference efficiency of the dual encoder approach.

To this end, previous works mainly focus on two directions: late-interaction and distillation. The first solution is to design a hybrid architecture, where the early layers act as a dual encoder while the late layers work like a cross encoder (MacAvaney et al., 2020; Khattab and Zaharia, 2020; Humeau et al., 2020). Its effectiveness comes with the cost of retrieval latency due to the computation involved with late layers. Another solution is knowledge distillation (Hinton et al., 2015), using the cross encoder to augment the training data (Qu et al., 2021; Ren et al., 2021), or distilling the ranking scores or attention matrices of a more powerful cross encoder reranker to a dual encoder retriever (Hofstätter et al., 2021; Ren et al., 2021; Lu et al., 2022).

In this paper, we propose to achieve this goal by pre-computing the interaction-based representations. As depicted in Figure 1c, the document representations are obtained by feeding the concatenation of query and document through a cross encoder while the query representation is obtained in the same way as in the vanilla dual encoder. For every document, we use a query generation model to generate several queries which will each concatenate with the document to obtain a separate document representation.

Our model has the following advantages. Firstly, we can obtain document representation with deep query interactions without much additional inference cost. Additionally, we can naturally get multi-

041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053  
054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081

view document representations (Luan et al., 2021; Tang et al., 2021; Zhang et al., 2022) by treating the query as explicit view extractor.

We follow the popular contrastive learning paradigm for learning such representations. Experiments on several retrieval benchmarks demonstrate the effectiveness of the proposed approach.

To summarize, our main contributions are as follows:

- We propose a new model architecture for dense retrieval, which can benefit from deep query-document interaction with low inference latency and learn multi-view document representations to better match different queries.
- We show the effectiveness of this model over various baselines by experiments on several large-scale retrieval benchmarks.

## 2 Related Work

### 2.1 Dense Retrieval

Dense passage retrieval (DPR) (Karpukhin et al., 2020) learns a two-tower BERT encoder to represent question and passage as vectors and takes their dot product as relevance score. The training of such dense retrievers can be optimized with more sophisticated negative sampling strategy (Xiong et al., 2021; Qu et al., 2021; Hofstätter et al., 2021; Zhan et al., 2021; Yang et al., 2021), or knowledge distillation from a more powerful cross-encoder teacher (Qu et al., 2021; Ren et al., 2021; Hofstätter et al., 2021; Lu et al., 2022).

Recently, some work have been devoted to trading off the efficiency and effectiveness with a late-interaction architecture. Humeau et al. (2020) compress the query context into multiple dense vectors with a Poly-Encoder architecture. The relevance score is modeled by a attention-weighted sum of individual matching scores. Tang et al. (2021) further improve the multi-encoding scheme through  $k$ -means clustering over all document tokens’ embeddings. ColBERT (Khattab and Zaharia, 2020) learns word level representations for both query and document and calculates the relevance score with a MaxSim operation followed by a sum pooling aggregator. Although powerful, they cannot fully utilize the maximum inner product search (MIPS). In contrast, we employ a pre-interaction mechanism combined with a max pooler which is compatible with MIPS.

Multi-vector encoding is essential in these late-interaction models, but is also gradually borrowed to learn effective dense retrieval models. Luan et al. (2021) represent each document with its first  $k$  token embeddings. To learn multi-view document representations, Zhang et al. (2022) substitute the [CLS] token with  $k$  special [VIE] tokens as view extractors and propose a local contrastive loss with annealing temperature between different views. In comparison, our model learns diverse document representations through interactions with different queries.

### 2.2 Query Generation

Query generation (QG) is originally introduced to the IR community as a document expansion technique (Nogueira et al., 2019). Nogueira and Lin (2019) show that appending the T5-generated queries to the document before building the inverted index can bring substantial improvements over BM25. More recently, Mallia et al. (2021) use generated queries as term expansion to learn better sparse representations for documents.

In the context of dense retrieval, query generation is usually used for domain adaptation in data scarcity scenarios. For example, Ma et al. (2020) use QG model trained on general domain to generate synthetic queries on target domain for model training. To reduce noise in generated data, Wang et al. (2022) further introduce a cross encoder for pseudo labeling. Different from the previous work, we mainly use the generated queries to learn query-informed document representations.

## 3 Method

### 3.1 Task Definition

Given a query  $q$  and a collection of  $N$  documents  $\mathcal{D} = \{d_1, d_2, \dots, d_i, \dots, d_N\}$ , a retriever aims to find a set of  $K$  relevant documents  $\mathcal{D}_+ = \{d_{i_1}, d_{i_2}, \dots, d_{i_j}, \dots, d_{i_K}\}$ ,<sup>1</sup> by ranking the document in the corpus  $\mathcal{D}$  according to its relevance score with respect to the query  $q$ , for next stage re-ranking or downstream applications.

### 3.2 Dual Encoder

We first introduce the dual encoder (DE) architecture for dense retrieval. In this framework, a query encoder  $DE_q$  and a document encoder  $DE_d$  are used to encode the query and document into low-dimensional vectors, respectively. To measure their

<sup>1</sup>Usually  $K \ll N$ .

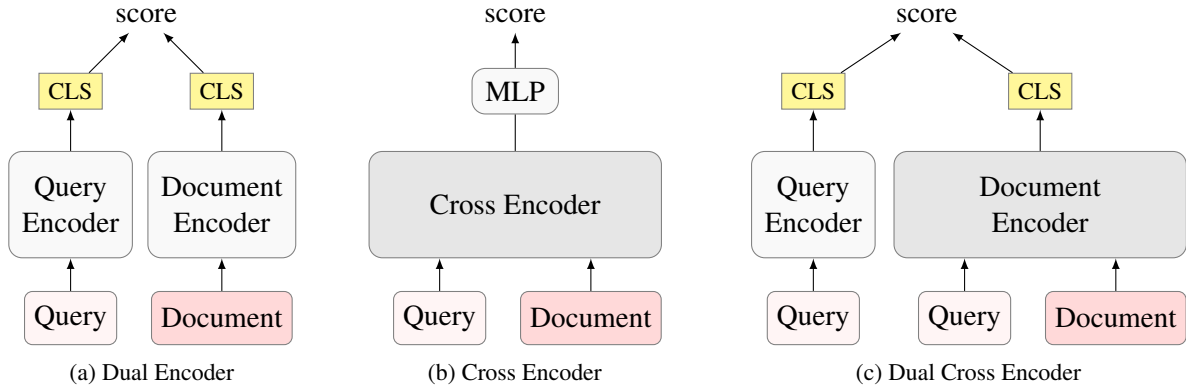


Figure 1: Illustration of different matching paradigms with different architectures.

relevance, a lightweight dot product between the two vectors is usually adopted to enable fast search,

$$s(q, d) = DE_q(q) \cdot DE_d(d). \quad (1)$$

The common design choice for the encoders is using multi-layer Transformers (Vaswani et al., 2017) initialized from pre-trained language models (PLMs), such as BERT (Devlin et al., 2019). How to get the representation from BERT is also an interesting question but beyond the scope of this paper. For simplicity, we directly take the [CLS] vector at the final layer as the text representation. The two encoders can share or use separate parameters. We tie the encoder parameters in main experiments but also provide results of untied parameters in ablation study.

### 3.3 Cross Encoder

The cross encoder (CE) takes the concatenation of query and document as input and uses deep neural network to model their deep interactions. Given a pair of query and document consisting of multiple tokens, we feed their concatenation through a cross encoder to get the interaction-aware representation,

$$\mathbf{r} = CE(q + d). \quad (2)$$

Then a multi-layer perceptron (MLP) is applied on top of the interaction-aware representation to predict the relevance score,

$$s(q, d) = MLP(\mathbf{r}). \quad (3)$$

The cross encoder is also usually instanced as a multi-layer Transformer network initialized from BERT. It can model term-level interactions between query and document, providing more fine-grained relevance estimation.

### 3.4 Dual Cross Encoder

We present our dual cross encoder where the document encoder acts as a cross encoder whereas the query encoder works like a dual encoder. Specifically, the query representation and document representation with query interaction are calculated as

$$\mathbf{q} = DE_q(q), \quad (4)$$

$$\mathbf{d} = CE_d(q + d). \quad (5)$$

Their similarity is measured by a dot product like in the vanilla dual encoder,

$$s(q, d) = \mathbf{q} \cdot \mathbf{d}. \quad (6)$$

**Query Generation.** Note that the query from the query encoder side and document encoder side do not necessarily have to be the same since we only have access to the gold query for documents appearing in the training set. It is impractical to manually write potential queries for each document in the whole corpus. Hence, we use a T5 model (Raffel et al., 2020) fine-tuned on the doc-to-query task to generate queries for each document. We empirically adopt 10 queries decoded with top- $k$  sampling strategy (Fan et al., 2018) to encourage the query generation diversity.

The advantages of this architecture are two-fold. On the one hand, it can model the query-document interaction in the pre-computed document representations. On the other hand, it can enjoy the retrieval efficiency of the vanilla dual encoder by pre-computing the interaction-aware document representations.

### 3.5 Training

The conventional way to train a dense retriever requires a set of  $(q, d_+, d_-)$  pairs. The model is

trained by optimizing the contrastive loss,

$$L(q, d_+, \mathcal{D}_-) = -\log \frac{e^{s(q, d_+)}}{\sum_{d \in \{d_+\} \cup \mathcal{D}_-} e^{s(q, d)}}, \quad (7)$$

where  $\mathcal{D}_-$  contains a set of negative documents  $d_-$  for query  $q$ . Following Karpukhin et al. (2020), we include both BM25 hard negatives and in-batch negatives in  $\mathcal{D}_-$ .

**Constructing Positives and Negatives.** Fusing query information into document representation requires redefining the positive and negative pairs. For a given query  $q$ , our framework potentially permits four types of positive and negatives, namely,  $(q_+ + d_+)$ ,  $(q_+ + d_-)$ ,  $(q_- + d_+)$  and  $(q_- + d_-)$ . To train our model, we convert the traditional positive and negative pair from the training set into that in our framework with the mapping function

$$f : (q, d_+, d_-) \mapsto (q, q + d_+, q + d_-), \quad (8)$$

where the  $+$  is the concatenation operation used in cross encoder. This mapping leads to the positive of type  $(q_+ + d_+)$  and the following two types of negatives. We leave the exploration of other types of negatives to future work.

**Hard Negatives.** The negative documents  $d_-$  are usually randomly sampled from BM25 top-ranked documents. After the mapping function defined above, these negatives fall in the type of negative  $(q_+ + d_-)$ , which serve as hard negatives in our framework. This type of negatives can teach the model to learn more fine-grained information, as  $d_-$  is usually topically related to the gold query but cannot exactly answer the query. It also prevents our model from learning the shortcut, *i.e.*, only learning matching signals from the query side and ignoring the document side information.

**In-Batch Negatives.** To improve the training efficiency, we also adopt in-batch negatives to train our model. In our framework, the in-batch negatives belong to the negative type  $(q_- + d_-)$ . This type of negatives is simple and can enable the model to learn topic-level discrimination ability.

**Data Augmentation.** Regarding the generated queries as weakly annotated data, we can first pre-train our model on these noisy data as a warm-up stage and then fine-tune it on the human-annotated high-quality training set.

### 3.6 Inference

**Index** We encode the corpus following the same format as Equation 5, to get multi-view document representations with deep query interactions.

Denoting  $\mathbf{d}_j^i$  as the  $i$ -th view of the  $j$ -th document  $d_j \in \mathcal{D}$ , we have

$$q_j^i \sim P_{QG}(q|d_j), \quad (9)$$

$$\mathbf{d}_j^i = CE_d(q_j^i + d_j), \quad (10)$$

where  $P_{QG}(q|d)$  denotes the query generation model,  $i \in \{1, \dots, k\}$  and  $j \in \{1, \dots, N\}$ .

**Retrieval** When a query comes, we encode it with the query encoder to get its contextualized representation  $\mathbf{q}$  as in Equation 4. We adopted multi-vector encodings for a document  $d_j$ , the relevance score between the query  $q$  and the document  $d_j$  is taken as the max pooling of its different views' scores,

$$s(q, d_j) = \max_i \mathbf{q}^T \mathbf{d}_j^i. \quad (11)$$

This operation is compatible with MIPS for efficiency optimization,<sup>2</sup>

$$p = \arg \max_d \mathbf{q}^T \mathbf{d}. \quad (12)$$

## 4 Experiment

In this section, we evaluate our model on different retrieval benchmarks and compare it with various baselines.

### 4.1 Datasets

We conduct experiments on the following retrieval benchmarks.

**MS MARCO** is a retrieval benchmark that originates from a machine reading comprehension dataset containing real user queries collected from Bing search and passages from web collection (Bajaj et al., 2016). We evaluate our model on the passage retrieval task. The corpus contains about 8.8M passages. The training set consists of about 500k annotated query-document pairs. The dev set has 6980 annotated queries. Since the test set is not publicly available, we evaluate on the dev set following previous work.

**TREC Deep Learning (DL)** tracks provide test sets with more elaborate annotations to evaluate the real capacity of ranking models. We evaluate on the

<sup>2</sup>Note that to get top- $K$  documents, we first retrieve  $10K$  documents to ensure that we have at least  $K$  documents after pooling.

2019 and 2020 test set (Craswell et al., 2020b,a). The 2019 test set contains 43 annotated queries and the 2020 test set contains 54 annotated queries. Both of them share the same corpus with the MS MARCO passage retrieval benchmark.

## 4.2 Evaluation Metrics

Following previous work, we mainly evaluate the retrieval performance on MS MARCO passage retrieval benchmark with  $MRR@10$  but also report the score of  $Recall@1000$ . For datasets from TREC DL tracks, we evaluate with  $nDCG@10$ .

## 4.3 Baselines

We mainly compare our model against the DPR (Karpukhin et al., 2020) baseline with a dual encoder architecture, but also report results of the following models most related to ours.

- BM25 (Robertson and Zaragoza, 2009) is the traditional lexical retriever.
- DocT5Query (Nogueira and Lin, 2019) appends generated queries to the document before building the inverted index of BM25.
- DeepImpact (Mallia et al., 2021) learns sparse representation for documents using generated queries as expanded terms.
- ANCE (Xiong et al., 2021) trains the DPR model with iterative hard negative mining strategy. We include this baseline since this technique is used in ME-BERT and DRPQ.
- ME-BERT (Luan et al., 2021) utilizes the first  $k$  token embeddings as multi-vector encodings for documents and adopts max pooling for score aggregation.
- DRPQ (Tang et al., 2021) improves over ME-BERT by performing a  $k$ -means over all tokens’ embeddings and utilizing a attention-based score aggregator.
- ColBERT (Khattab and Zaharia, 2020) represents query and document at token-level and uses a MaxSim pooler followed by a sum aggregator to calculate the relevance score.

## 4.4 Implementation

We implement our model based on the `tevatron` toolkit (Gao et al., 2022). For a fair comparison

with our model, we re-implement the DPR baseline using the same set of hyperparameters.

We train all the models on 8 NVIDIA Telsa V100 GPUs with 32GB memory. We initialize all the encoders with `bert-base-uncased`. The max sequence length is 16 for query and 128 for passage. The number of positive and negative passages follows a ratio of 1:7 for each sample. We set the batch size to 32. We use both officially provided BM25 negatives and in-batch negatives to train the models. We use Adam optimizer with the learning rate of  $5 \times 10^{-6}$ , linear decay with 10% warmup steps.

In the preliminary study without data augmentation, we train both models for 10 epochs. To make full use of generated queries, we first pre-train the models for 10 epochs on the corpus with a batch size of 256 and only in-batch negatives, and then fine-tune the models for 20 epochs till convergence on the training set. We haven’t tuned other hyperparameters. The pre-training stage takes about 15 hours and the fine-tuning stage takes about 8 hours.

During inference, we use `IndexFlatIP` of the `faiss` library (Johnson et al., 2021) to index the corpus and perform an exact search.

## 4.5 Results

Table 1 illustrates the evaluation results of our model and the baselines.

We first compare our model against the DPR dual encoder baseline. We can observe substantial improvements in terms of  $MRR@10$  and  $nDCG@10$  across all these datasets, which demonstrate the effectiveness of our approach. The  $Recall@1k$  also exhibits a slight improvement.

Our approach is also competitive with other baselines. On MS MARCO, it surpasses other baselines and is comparable to ColBERT, while being more efficient. On TREC DL 19, the results are comparable to ME-BERT, which used a more powerful large-size model as backbone and the hard negative mining technique of ANCE. On TREC DL 20, our model even outperforms the ColBERT model.

## 4.6 Ablation Study

We conduct ablation studies on our model design choice.

### 4.6.1 Effect of Data Augmentation

We used the generated queries as data augmentation for pre-training. We ablate on the effect of pre-training in this section. The results of different

Model	PLM	MS MARCO		TREC DL 19	TREC DL 20
		MRR@10	Recall@1k	nDCG@10	nDCG@10
<b>Sparse</b>					
BM25	-	18.4	85.3	50.6	48.0
DocT5Query	-	27.7	94.7	64.8	61.9
DeepImpact	BERT <sub>base</sub>	32.6	94.8	69.5	65.1
<b>Dense</b>					
DPR	BERT <sub>base</sub>	31.4	95.3	59.0	62.1
ANCE	RoBERTa <sub>base</sub>	33.0	95.9	64.8	-
ME-BERT	BERT <sub>large</sub>	33.4	-	68.7	-
DRPQ	BERT <sub>base</sub>	34.5	96.4	-	-
ColBERT	BERT <sub>base</sub>	36.0	96.8	69.4	67.6
Ours	BERT <sub>base</sub>	36.0	96.4	68.3	68.9

Table 1: Evaluation results on MS MARCO passage retrieval benchmark and TREC DL track. DocT5Query and DeepImpact can be seen as the sparse counterparts of our model. Both ME-BERT and DRPQ learn multi-vector encodings for documents, and have used the hard negative mining technique proposed in ANCE. ColBERT learns term-level representations of both query and document for late interaction. Results not available are marked as ‘-’.

training stages on MS MARCO dev set are shown in Table 2.

MRR@10	Pretrain	Finetune	Full
DPR	25.6	31.4	34.2
Ours	26.1	33.2	36.0

Table 2: Ablation of different training stages on MS MARCO dev set. Pretrain: only use generated data for training; Finetune: only use data from training set for training; Full: Pretrain + Finetune.

We can see that using generated data for pre-training gives a MRR@10 score comparable to DocT5Query but lower than directly fine-tuning using data from the training set. The top- $k$  sampling decoding strategy in query generation may introduce some noise, which explains why the pre-training underperforms directly fine-tuning with high-quality data. However, the pre-training stage is still beneficial for the fine-tuning stage.

The results on TREC DL track are shown in Table 3. Our model still consistently outperforms the dual encoder baseline under different settings. The improvements are more significant on this benchmark since the annotation is more complete. Notably, our model without data augmentation is comparable to the DPR baseline with data augmentation on this benchmark.

#### 4.6.2 Effect of Sharing Parameters

Sharing the encoder parameters can reduce the number of model parameters to half. We tie our

nDCG@10	DL 19	DL 20
w/o Data Augmentation		
DPR	59.0	62.1
Ours	63.0	67.6
w/ Data Augmentation		
DPR	63.1	66.5
Ours	68.3	68.9

Table 3: Results on TREC DL track under different settings.

encoder parameters in main experiments but also provide ablation of untied encoder parameters in Table 4 to study its effect.

MRR@10	tie	untie
DPR	31.4	31.7
Ours	33.2	33.8

Table 4: Results of tie / untie encoder parameters on MS MARCO dev set.

We can observe that using two sets of encoder parameters gives slightly better performance but not so significantly. Using separate encoders brings more improvements to our model, which is normal since the nature of two encoders in our model is more asymmetric than that in the vanilla dual encoder.

## 5 Analysis

Our experimental results in the previous section demonstrate that it is indeed beneficial to incorporate query interactions into the document representations. The generated queries are crucial to the success of our model. In this section, we analyse the influence of query quality and query diversity to the retrieval performance.

### 5.1 On the Query Quality

The number of queries is an important factor in our framework. Too few queries have low diversity while too many queries will sacrifice efficiency. Thus we provide an analysis here to study its effect. We evaluate the query generation performance on the dev set of MS MARCO and reveal its relation with the retrieval performance.

To measure the generation quality, we calculate the maximum ROUGE-L score between generated queries and the gold query on the dev set. For retrieval performance, we report the MRR@10. Figure 2 illustrates the evolution of the two metrics with different number of queries.<sup>3</sup>




Figure 2: The evolution of ROUGE-L and MRR@10 on MS MARCO dev set when varying number of queries from 1 to 10.

We can see that as the number of queries grows, the retrieval performance becomes better because of the improved generation quality. The correlation between the two metrics is shown in Figure 3. The Pearson coefficient is 0.9958, indicating a strong positive correlation. Keep increasing the number of queries will consistently improve the retrieval performance but more marginally.

<sup>3</sup>Please refer to Table 6 in Appendix A for exact numbers.




Figure 3: Correlation between generation and retrieval performance on MS MARCO dev set.

### 5.2 On the Query Diversity

Intuitively, more diverse queries can potentially hit more types of queries. We used top- $k$  sampling strategy to encourage the query generation diversity. However, whether and to what extent the generated queries are diverse remains unclear. To this end, we adopt the self-BLEU (Zhu et al., 2018) to measure the query generation diversity for a document.

We partition the documents of MS MARCO dev set to subsets of different query diversity according to their self-BLEU-4 score and measure the retrieval performance on these subsets. The statistics are shown in Figure 4. First, we observe that most of the documents have high query generation diversity thanks to the top- $k$  sampling strategy (see Figure 4a). Second, the retrieval performance drops with higher diversity (see Figure 4b). One possible reason is that the QG model will generate more diverse queries when it doesn't know the right one. As such, higher diversity indicates lower quality (see Figure 4c) and the retrieval performance drops with lower generation quality (see Figure 3). It would be desirable to design a diversity metric that takes into account the generation quality.

### 5.3 Case Study

We conduct a case study on the dev set to intuitively compare our model and the dual encoder baseline, as well as to illustrate the QG performance.

Table 5 shows an example drawn from the MS MARCO dev set. Our model can retrieve the correct passage by generating the right query. DPR retrieves a hard negative passage where the content is corresponding to the query keywords but can not correctly answer the query. By generating queries, our model can better distinguish the




Figure 4: Statistics of query diversity on MS MARCO dev set. We divided the diversity into 5 levels based on an average division of the self-BLEU-4 score.

Query	how old is canada
Ours Rank 1	Canada was finally established as a country in 1867. It is 148 years old as of July 1 2015. Canada has been a country for 147 years. The first attempt at colonization occurred in 1000 A.D. by the Norsemen. There was no further European exploration until 1497 A.D. when the Italian sailor John Cabot came along. It then started being inhabited by more Europeans.
Generated Queries	when was canada established when was canada discovered what year was canada founded how long has canada been a country how old is canada
DPR Rank 1	it depends where you live but in Canada you have to be at least 16 years old.
Generated Queries	what is the legal age to be in canada how old do you have to be to live in canada how old do you have to be to enter canada as a citizen at what age can i go to canada to study in canada what is the minimum age to join the military

Table 5: Case Study on MS MARCO dev set.

521 difference among document meanings.

## 522 6 Discussion

523 The ranking task is usually approached with a two-  
524 stage pipeline: retrieve-then-rerank. The two stages  
525 usually use different architectures due to the effec-  
526 tiveness and efficiency trade-off. Dual encoder is  
527 more efficient for retrieval, while cross encoder is  
528 more powerful for reranking. How to take advan-  
529 tage of each other’s strengths for mutual improve-  
530 ments is a hot topic of research. We propose a new  
531 dual cross encoder architecture to benefit from both  
532 with a pre-interaction mechanism.

533 One limitation of our framework is that there  
534 exists a discrepancy between training and inference.  
535 We used the gold query to train the model but do  
536 not have access to the gold query during inference.  
537 Generating more queries would bridge this gap, but  
538 at the cost of efficiency. We wish to close this gap

with improved training strategy or improved query  
generation quality in the future.

## 541 7 Conclusion

542 We proposed a novel dense retrieval model to  
543 bridge the gap between dual encoder and cross  
544 encoder. In our framework, the document repre-  
545 sentations are obtained by pre-interacting with a  
546 set of generated pseudo-queries through a cross en-  
547 coder. Our approach enables multi-view document  
548 representation with deep query interaction while  
549 maintaining the inference efficiency of the dual en-  
550 coder approach. We demonstrated its effectiveness  
551 compared to dual encoder baseline via experiments  
552 on various retrieval benchmarks. In the future work,  
553 we would like to explore how to better incorporate  
554 generated queries for model training and how to  
555 improve the query generation quality for better re-  
556 trieval performance.



## References

Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2016. *Ms marco: A human generated machine reading comprehension dataset*.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. *Reading Wikipedia to answer open-domain questions*. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.

Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020a. *Overview of the TREC 2020 deep learning track*. In *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).

Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2020b. *Overview of the TREC 2019 deep learning track*. *CoRR*, abs/2003.07820.

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.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. *Hierarchical neural story generation*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.

Luyu Gao, Xueguang Ma, Jimmy J. Lin, and Jamie Callan. 2022. *Tevatron: An efficient and flexible toolkit for dense retrieval*. *ArXiv*, abs/2203.05765.

Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. *Distilling the knowledge in a neural network*. *CoRR*, abs/1503.02531.

Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. *Efficiently teaching an effective dense retriever with balanced topic aware sampling*. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 113–122. ACM.

Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2020. *Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring*. In *International Conference on Learning Representations*.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. *Billion-scale similarity search with gpus*. *IEEE Transactions on Big Data*, 7(3):535–547.

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.

Omar Khattab and Matei Zaharia. 2020. *Colbert: Efficient and effective passage search via contextualized late interaction over BERT*. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 39–48. ACM.

Yuxiang Lu, Yiding Liu, Jiaxiang Liu, Yunsheng Shi, Zhengjie Huang, Shikun Feng, Yu Sun, Hao Tian, Hua Wu, Shuaiqiang Wang, Dawei Yin, and Haifeng Wang. 2022. *Ernie-search: Bridging cross-encoder with dual-encoder via self on-the-fly distillation for dense passage retrieval*. *CoRR*, abs/2205.09153.

Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. *Sparse, dense, and attentional representations for text retrieval*. *Transactions of the Association for Computational Linguistics*, 9:329–345.

Ji Ma, Ivan Korotkov, Yinfei Yang, Keith Hall, and Ryan McDonald. 2020. *Zero-shot neural passage retrieval via domain-targeted synthetic question generation*. *arXiv preprint arXiv:2004.14503*.

Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonello, Nazli Goharian, and Ophir Frieder. 2020. *Efficient document re-ranking for transformers by precomputing term representations*. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 49–58. ACM.

Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonello. 2021. *Learning passage impacts for inverted indexes*. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1723–1727. ACM.

Rodrigo Nogueira and Jimmy Lin. 2019. *From doc2query to docttttquery*.

Rodrigo Frassetto Nogueira and Kyunghyun Cho. 2019. *Passage re-ranking with BERT*. *CoRR*, abs/1901.04085.

668 Rodrigo Frassetto Nogueira, Wei Yang, Jimmy Lin, and 725  
669 Kyunghyun Cho. 2019. [Document expansion by](#) 726  
670 [query prediction](#). *CoRR*, abs/1904.08375. 727

671 Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang 728  
672 Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and 729  
673 Haifeng Wang. 2021. [RocketQA: An optimized training](#) 730  
674 [approach to dense passage retrieval for open-](#) 731  
675 [domain question answering](#). In *Proceedings of the* 732  
676 *2021 Conference of the North American Chapter of* 733  
677 *the Association for Computational Linguistics: Hu-*  
678 *man Language Technologies*, pages 5835–5847, On-  
679 line. Association for Computational Linguistics.

680 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine 734  
681 Lee, Sharan Narang, Michael Matena, Yanqi 735  
682 Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the](#) 736  
683 [limits of transfer learning with a unified text-to-text](#) 737  
684 [transformer](#). *Journal of Machine Learning Research*, 738  
685 21(140):1–67. 739

686 Ruiyang Ren, Yingqi Qu, Jing Liu, Wayne Xin Zhao, 740  
687 QiaoQiao She, Hua Wu, Haifeng Wang, and Ji-Rong 741  
688 Wen. 2021. [RocketQAv2: A joint training method](#) 742  
689 [for dense passage retrieval and passage re-ranking](#). 743  
690 In *Proceedings of the 2021 Conference on Empirical* 744  
691 *Methods in Natural Language Processing*, pages 745  
692 2825–2835, Online and Punta Cana, Dominican Re- 746  
693 public. Association for Computational Linguistics. 747

694 Stephen E. Robertson and Hugo Zaragoza. 2009. [The](#) 748  
695 [probabilistic relevance framework: BM25 and be-](#) 749  
696 [yond](#). *Found. Trends Inf. Retr.*, 3(4):333–389. 750

697 Hongyin Tang, Xingwu Sun, Beihong Jin, Jingang 751  
698 Wang, Fuzheng Zhang, and Wei Wu. 2021. [Improving](#) 752  
699 [document representations by generating pseudo](#) 753  
700 [query embeddings for dense retrieval](#). In *Proceed-* 754  
701 *ings of the 59th Annual Meeting of the Association for* 755  
702 *Computational Linguistics and the 11th International*  
703 *Joint Conference on Natural Language Processing* 756  
704 *(Volume 1: Long Papers)*, pages 5054–5064, Online.  
705 Association for Computational Linguistics.

706 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob 757  
707 Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz 758  
708 Kaiser, and Illia Polosukhin. 2017. [Attention is all](#) 759  
709 [you need](#). In *Advances in Neural Information Pro-* 760  
710 *cessing Systems*, volume 30. Curran Associates, Inc.

711 Kexin Wang, Nandan Thakur, Nils Reimers, and Iryna 761  
712 Gurevych. 2022. [GPL: Generative pseudo labeling](#) 762  
713 [for unsupervised domain adaptation of dense retrieval](#). 763  
714 In *Proceedings of the 2022 Conference of the North* 764  
715 *American Chapter of the Association for Computa-* 765  
716 *tional Linguistics: Human Language Technologies*, 766  
717 pages 2345–2360, Seattle, United States. Association 767  
718 for Computational Linguistics.

719 Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, 768  
720 Jialin Liu, Paul N. Bennett, Junaid Ahmed, and 769  
721 Arnold Overwijk. 2021. [Approximate nearest neigh-](#) 770  
722 [bor negative contrastive learning for dense text re-](#) 771  
723 [trieval](#). In *International Conference on Learning* 772  
724 *Representations*.

Nan Yang, Furu Wei, Binxing Jiao, Daxing Jiang, and 725  
Linjun Yang. 2021. [xMoCo: Cross momentum con-](#) 726  
[trastive learning for open-domain question answering](#). 727  
In *Proceedings of the 59th Annual Meeting of the* 728  
*Association for Computational Linguistics and the* 729  
*11th International Joint Conference on Natural Lan-* 730  
*guage Processing (Volume 1: Long Papers)*, pages 731  
6120–6129, Online. Association for Computational 732  
Linguistics. 733

Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min 734  
Zhang, and Shaoping Ma. 2021. [Optimizing dense](#) 735  
[retrieval model training with hard negatives](#). In *SI-* 736  
*GIR '21: The 44th International ACM SIGIR Confer-* 737  
*ence on Research and Development in Information* 738  
*Retrieval, Virtual Event, Canada, July 11-15, 2021*, 739  
pages 1503–1512. ACM. 740

Shunyu Zhang, Yaobo Liang, Ming Gong, Daxin Jiang, 741  
and Nan Duan. 2022. [Multi-view document repre-](#) 742  
[sentation learning for open-domain dense retrieval](#). 743  
In *Proceedings of the 60th Annual Meeting of the* 744  
*Association for Computational Linguistics (Volume* 745  
*1: Long Papers)*, pages 5990–6000, Dublin, Ireland. 746  
Association for Computational Linguistics. 747

Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan 748  
Zhang, Jun Wang, and Yong Yu. 2018. [Txygen: A](#) 749  
[benchmarking platform for text generation models](#). 750  
In *The 41st International ACM SIGIR Conference on* 751  
*Research & Development in Information Retrieval,* 752  
*SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, 753  
pages 1097–1100. ACM. 754

## A Appendix

$k$	ROUGE-L	MRR@10
1	42.49	27.74
2	50.93	30.09
3	55.67	31.15
4	58.45	31.66
5	60.63	31.92
6	62.28	32.38
7	63.57	32.67
8	64.62	32.88
9	65.46	32.96
10	66.22	33.23

Table 6: Results of generation and retrieval performance on MS MARCO dev set when varying number of queries (correspond to Figure 2 and Figure 3).