

DPTDR: Deep Prompt Tuning for Dense Passage Retrieval

Zhengyang Tang^{1,2}, Benyou Wang², and Ting Yao¹

¹Tencent Inc.

²The Chinese University of Hong Kong, Shenzhen

{zhytang, tessieyao}@tencent.com, wangbenyou@cuhk.edu.cn

Abstract

Deep prompt tuning (DPT) has gained great success in most natural language processing (NLP) tasks. However, it is not well-investigated in dense retrieval where fine-tuning (FT) still dominates. When deploying multiple retrieval tasks using the same backbone model (e.g., RoBERTa), FT-based methods are unfriendly in terms of deployment cost: each new retrieval model needs to repeatedly deploy the backbone model without reuse. To reduce the deployment cost in such a scenario, this work investigates applying DPT in dense retrieval. The challenge is that *directly applying DPT in dense retrieval largely underperforms FT methods*. To compensate for the performance drop, we propose two model-agnostic and task-agnostic strategies for DPT-based retrievers, namely *retrieval-oriented intermediate pretraining* and *unified negative mining*, as a general approach that could be compatible with any pre-trained language model and retrieval task. The experimental results¹ show that the proposed method (called DPTDR) outperforms previous state-of-the-art models on both MS-MARCO and Natural Questions. We also conduct ablation studies to examine the effectiveness of each strategy in DPTDR. We believe this work facilitates the industry, as it saves enormous efforts and costs of deployment and increases the utility of computing resources.

1 Introduction

Fine-tuning (FT) has been a de facto approach for effective dense passage retrieval (Karpukhin et al., 2020; Xiong et al., 2020) based on pre-trained language models (PLM). However, FT is unfriendly for industrial deployment in multi-task scenarios. Imaging for cloud service providers or infrastructure teams of search companies, each retrieval model (w.r.t., an individual task) necessarily redeploys a backbone model since the weights of

the backbone model in each task are fine-tuned and therefore slightly different. That dramatically increases hardware costs and inefficiency.

Recently, prompt tuning (PT) (Liu et al., 2021a) is a lightweight alternative to FT, which does not need storing a full copy of the backbone model for each task. One variant of PT, namely Deep Prompt Tuning (DPT; Li and Liang, 2021; Liu et al., 2021b), exhibits comparable performances with FT in various NLP tasks. DPT enjoys parameter-efficient (Houlsby et al., 2019) characteristics, of which the resulting prompts are light-weighted and can be easily passed to an online PLM service, thus overcoming the above challenge of FT. This paper asks: *whether can we replace FT by DPT with comparable performance to SOTA FT methods in dense passage retrieval?* With comparable performance, DPT is much more friendly in deployment than FT.

DPT usually freezes weights in the backbone models and alternatively trains deep prompts inserted; the latter has much fewer parameters than the former. However, freezing most weights in DPT hinders its adaptability and therefore possibly harms performance. Experimental results in Sec. 4.2.2 also demonstrate *directly applying DPT in dense retrieval largely underperforms FT methods*.

To make DPT comparable to FT in dense retrieval, a natural solution is *retrieval-oriented intermediate pretraining (RIP)*, which warms up the text representation via contrastive learning. Though it is not a novel idea (Lee et al., 2019; Gao and Callan, 2021b; Izacard et al., 2021), there exist two different pretraining ways tailored for DPT-based retrievers. One is to pre-train deep prompts while freezing the PLM backbone and use the pre-trained prompts to initialize a DPT retriever. The other is to pre-train a PLM directly and initialize a DPT retriever using the pre-trained PLM; in contrast to prior works (Gao and Callan, 2021b), we intend to allow any PLM easily pre-trained for DPT so that

¹Our code is available at <https://github.com/tangzhy/DPTDR>

users may employ their own PLMs, and thus we deliberately remove the workload to modify any model structures. Surprisingly, empirical findings in Sec. 4.4 show that this choice yields better performance than carefully modified PLMs (Gao and Callan, 2021b). Furthermore, we propose a *unified negative mining (UNM)* to merge retrieved negatives from many retrievers including BM25 and dense retrievers, in order to provide diverse and hard negatives for DPT training.

By incorporating RIP and UNM, we implement a **Deep Prompt Tuning** method in **Dense Retrieval** tasks, called *DPTDR*. The experimental results show that DPTDR outperforms previous state-of-the-art models on both MS-MARCO and Natural Questions. We also conduct extensive experiments and find that: i) when combined with RIP and UNM, DPT is able to obtain comparable performance with FT in dense retrieval and exhibits insensitivity to prompt length, and ii) both RIP and UNM are effective in improving the performance. The contributions of this paper can be summarized as follows:

- To our best knowledge, this is the first work to apply DPT in dense retrieval. We bring forward two essential strategies, namely retrieval-oriented intermediate pretraining and unified negative mining, allowing DPT to match FT’s performance and be compatible with any PLM.
- Experiments show that DPTDR outperforms previous state-of-the-art models on MS-MARCO and Natural Questions and examine the effectiveness of the above strategies.
- We believe this work facilitates the industry, as it saves enormous efforts and costs of deployment and increases the utility of computing resources.

2 Related Work

2.1 Deep Prompt Tuning

DPT originates from prompting and prompt tuning (Liu et al., 2021a). Given some discrete or continuous prompts, PLMs like GPT-3 (Brown et al., 2020) can achieve impressive zero-shot and few-shot performances for knowledge-intensive tasks. However, studies find that prompt tuning fails to perform well for moderate-size models (Liu et al., 2021b). Thus, DPT (Li and Liang, 2021; Liu et al.,

2021b) is proposed by inserting prompts at deep layers to steer PLMs towards desired directions more capably. It obtains comparable performance to FT across a range of NLP tasks. DPTDR is mainly related to DPT, focusing on dense passage retrieval instead of NLP. There also exist works of pretraining prompts for prompt tuning (Gu et al., 2021), which shows effectiveness in few-shot learning using billion-size models, as we will explore as well in the context of DPT.

2.2 Dense Retrieval

Pretraining We have witnessed a series of unsupervised pretraining works proposed for dense retrieval, such as ICT, BFS, WLP, and independent cropping (Lee et al., 2019; Chang et al., 2020; Izacard et al., 2021). Following works also try to pre-train retriever and reader jointly for question answering (Guu et al., 2020). coCondenser (Gao and Callan, 2021b) follows a contrastive learning framework using Condenser structure (Gao and Callan, 2021a) by adding an explicit decoder to learn representations better. There are also semi-supervised and weakly-supervised works. DPR-PAQ (Oğuz et al., 2021) pre-trains a PLM using 65-million-size synthetic QA pairs on the target corpus. GTR (Ni et al., 2021) pre-trains T5 (Raffel et al., 2019) on 2-billion-size community QA pairs from T5-base to T5-xxlarge. We follow unsupervised contrastive learning as our pretraining strategy for DPTDR. However, we aim to ensure compatibility with any PLM, thus resulting in different sample building processes and model structure choices.

Negative mining DPR (Karpukhin et al., 2020) proposes to train retrievers using BM25 negatives. ANCE (Xiong et al., 2020) extends that by mining negatives periodically from previously-trained dense retrievers. RocketQA and RocketQAv2 (Qu et al., 2021; Ren et al., 2021) introduce the idea of denoised negative sampling by selecting negatives with high confidence scored by a re-ranker. DPTDR unifies the above into a general negative mining strategy.

3 Methodology

In this section, we first formalize the application of DPT in dense retrieval. We then describe the two strategies of RIP and UNM for DPT-based retrievers.

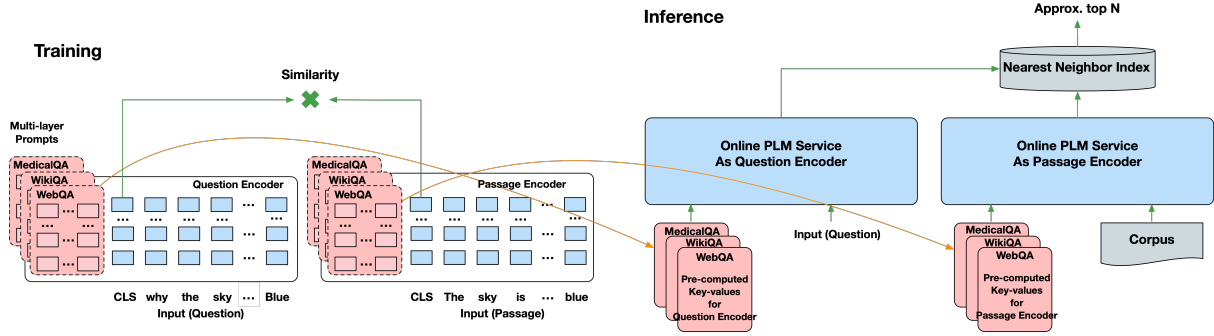


Figure 1: The framework of DPTDR. We first perform RIP which results in a PLM (the blue blocks) that can be used as the backbone for DPT training and deployed once as online PLM services. Then we train deep prompts (i.e., DPT) for different retrieval tasks such as WebQA, WikiQA, and MedicalQA (the pink blocks), during which we may employ UNM to improve performances. For inference, we can send tokenized input, together with trained prompts of their corresponding task, to online PLM services to get dense vectors.

3.1 DPT in Dense Retrieval

Let C be a corpus consisting N passages, denoted by p_1, p_2, \dots, p_N . Given a question q , the task of dense retrieval is to find a passage p_i that is considered relevant to the question.

The dual-encoder Normally, a dual-encoder is applied. First its passage encoder $E_p(\cdot)$ embeds a passage p to a d -dimensional dense vector. Then a vector search index (Johnson et al., 2019) of passages is built for retrieval. At inference time, the question encoder $E_q(\cdot)$ embeds the question q to a d -dimensional dense vector, and k passages closet to the question based on the vector similarity will be retrieved. In practice, the similarity score is computed as the inner product:

$$s(q, p) = E_q(q) \cdot E_p(p). \quad (1)$$

For PLM-based dual-encoder, we usually take the representation at the first token (e.g., [CLS] symbol in BERT (Devlin et al., 2018)) as the output dense vector.

Deep prompt tuning We then apply DPT in the PLM-based dual-encoder, as illustrated in the left part of Figure 1. To prepend multi-layer prompts for the dual-encoder, we initialize a trainable prefix matrix M of dimension $l \times d$ for each layer of the PLM, where l is the length of the prompt and d is the hidden size of the PLM. Since the prompt resides at the deep layers of PLM, it has a full capacity to steer the PLM towards the desired direction and output meaningful dense vector for questions and passages. Note that a verbalizer (Schick and Schütze, 2020) plays a vital role in mapping words to labels in canonical prompt tuning. However, we remove it in dense retrieval since the output dense

vector is what we need. Let E'_p as the prompted passage encoder and E'_q as the prompted question encoder, and the similarity score is computed:

$$s'(q, p) = E'_q(q) \cdot E'_p(p). \quad (2)$$

Training The objective of the training is to learn dense vectors so that the similarity between relevant pairs of questions and passages ranks higher than irrelevant ones. Given a pair of question q and positive passage p_i , along with n negative passages, we optimize the loss function as the negative log-likelihood of the positive passage:

$$L(q_i, p_i^+, \{p_{i,j}^-\}_{j=1}^n) = -\log \frac{e^{s'(q_i, p_i^+)}}{e^{s'(q_i, p_i^+)} + \sum_{j=1}^n e^{s'(q_i, p_{i,j}^-)}}. \quad (3)$$

Generating negative passages is critical for the performance, and we will explain it in Sec. 3.3. During training, we freeze parameters of the backbone PLM and only update the deep prompts, where approximately 0.1%-0.4% parameters of a PLM get trained.

Inference As illustrated in the right part of Figure 1, since the backbone PLM is frozen, it is possible to deploy it ahead as online PLM services and then pass the trained prompts as pre-computed key values together with tokenized inputs to get dense vectors. It is at the core of how we save efforts and costs of deployment and increase the utility of computing resources. In practice, the cloud service providers or infrastructure teams of search companies are able to focus on the PLM as a central service, while users can quickly train deep prompts

for different retrieval tasks and obtain efficient and compelling retrieval performances without any deployment.

Although DPT brings in many advantages, it is worth noting that it does not accelerate the inference speed because the forward computation is not reduced but increased slightly.

3.2 Retrieval-oriented Intermediate Pretraining (RIP) for DPT

The goal of RIP is to either pre-train deep prompts or PLMs using contrastive learning. We first describe the task as follows. Let C denote a corpus consisting N passages. For a passage p_i , we split it into l sentences, denoted by s_i^1, \dots, s_i^l . Given a sentence s_i^j , the task of pretraining is to distinguish its context sentence s_i^j from sentences of other passages s_k^l , where $k \neq i$. Formally, we randomly select a pair of sentences from each passage as context sentences to form a batch of training data $B = \{s_i^1, s_i^2\}_{i=1}^m$, where m is the batch size. Then we define the contrastive loss for s_i^j over the batch as:

$$L_c(s_i^j) = -\log \frac{e^{s(s_i^1, s_i^2)}}{\sum_{k=1}^m \sum_{l=1}^2 \sum_{ij \neq kl} e^{s(s_i^j, s_k^l)}}. \quad (4)$$

In contrast to prior works(Gao and Callan, 2021b; Izacard et al., 2021), we directly sample sentences as opposed to text spans. Since sampling text spans is a non-trivial technique where factors such as the probability of short sentences and how to keep the spans linguistically meaningful can have a complicated effect on the pretraining, we remove this complexity in our approach. We also conduct an experiment observing sentences work even better than text spans on MS-MARCO corpus (Sec. 4.4).

Under the contrastive learning task, there exist two pretraining ways tailored for DPT, depending on the pre-trained objects (i.e., the deep prompts or the PLM backbone).

Pre-train deep prompts One is to pre-train deep prompts with a vanilla PLM. Later we initialize a DPT-based retriever using the pre-trained deep prompts and the vanilla PLM. However, experiments in Sec. 4.4 show that it suffers from catastrophic forgetting and exhibits no superior performance to randomly initialized prompts.

Pre-train the PLM The other is to pre-train a PLM, and then we initialize a DPT-based retriever

using randomly-initialized deep prompts and the pre-trained PLM. Notice that we intend to allow any PLM to be easily pre-trained for DPT so that users may employ their own PLMs. Thus we contrast prior works such as coCondenser(Gao and Callan, 2021b), a state-of-the-art model structure in contrastive pretraining, by removing the workload to modify any model structures. Surprisingly, it yields better performance than coCondenser in Table 8. Therefore, we refer RIP strategy as pretraining of PLMs for the rest.

For any PLM, We also intend to remain its original self-supervised tasks, such as masked language modeling(MLM; Devlin et al., 2018; Sun et al., 2019), denoted as L_s . Therefore, the final loss of pretraining over the batch is:

$$L = \frac{1}{2m} \sum_{i=1}^m \sum_{j=1}^2 L_s(s_i^j) + L_c(s_i^j). \quad (5)$$

After pretraining, the resulting model can be deployed once as online services and taken as the backbone model for DPT training.

3.3 Unified Negative Mining (UNM)

We also develop unified negative mining for DPT, as interpreted as "Multiple Retrievers & Hybrid Sampling." "Multiple Retrievers" is to incorporate negatives from as many retrievers as we can. We use a BM25 retriever as the initial retriever and train a DPT-based retriever using BM25 negatives. Later we treated retrieved negatives from the BM25 retriever and the first DPT-based retriever as un-denoised hard negatives. Users are allowed to introduce any other retrievers if possible. "Hybrid Sampling" is to select denoised hard negatives from un-denoised hard negatives retrieved by the above multiple retrievers. We borrow an existing re-ranker released by RocketQA (Qu et al., 2021) and select those negatives with high confidence. For training the final DPT-based retriever, we mix the denoised hard negatives, un-denoised hard negatives, and easy negatives from in-batch or cross-batch training.

We believe unified negative mining is critical for the performance of DPT-based retrievers, as it provides negatives of high quality and diversity.

4 Experiments

4.1 Experimental Setting

Datasets and metrics We experiment with two popular dense retrieval datasets, including MS-

Table 1: The statistics of MS-MARCO and Natural Questions.

Dataset	#q in train	#q in dev	#q in test	#passages
MS-MARCO	502,939	6,980	6,837	8,841,823
Natural Questions	58,812	6,515	3,610	21,015,324

MARCO (Bajaj et al., 2016) and Natural Questions (NQ; Karpukhin et al., 2020). The statistics of the datasets are listed in Table 1. MS-MARCO is constructed from Bing’s search query logs and web documents retrieved by Bing. Natural Question contains questions from Google Search. For evaluation, we report official metrics MRR@10, RECALL@1000 for MS-MARCO, and RECALL at 5, 20, and 100 for NQ. All models are trained on a single server with 8 NVIDIA Tesla A100 GPUs.

Settings in DPT We use RoBERTa-large-size models as the backbone for DPT training. Hyper-parameters are explored as below.

- **Learning rate** We search for 1e-2, 5e-3, 7e-3, 5e-4, 5e-5, 5e-6 with prompts’ length of 32, where 7e-3 performs relatively better than others and is set for the main experiment.
- **Training epochs** For training epochs, we search for 3, 6, 10 with a learning rate 7e-3 on MS-MARCO, where 10 performs best and is set for the main experiment. We also set training epochs as 60 for NQ for acceptable time cost.
- **Prompt length** We search for 8, 16, 32, 64, 128, as is discussed in Sec. 4.3. We use 128 for the main experiment.
- **Reparametrization** We also conduct experiments for prompts with or without MLP reparametrization, as is discussed in Sec. 4.3. We use non-reparametrization for the main experiment.

We follow coCondenser (Gao and Callan, 2021b) for other hyper-parameters (e.g., parameter sharing, batch size, warm-up ratio, and mixed-precision training).

Settings in RIP We choose to pre-train vanilla RoBERTa-large for RIP, whose model size appears more common for DPT (Li and Liang, 2021; Liu et al., 2021b) and is consistent with the above DPT training. We remain RoBERTa’s original self-supervised task (MLM; Liu et al., 2019). To

compare our approach with coCondenser (Gao and Callan, 2021b), we also pre-train a coCondenser RoBERTa-large. Since coCondenser modifies the PLM by adding a carefully designed Condenser structure, we follow their structural setting using an equal split, 12 early layers, and 12 late layers. We split the passages into sentences on both MS-MARCO and NQ Wikipedia as the training corpus. The models are trained using AdamW optimizer with a learning rate 1e-4, weight decay of 0.01, linear learning rate decay, and a batch size of 2K. We train 8 epochs for MS-MARCO and 4 epochs for NQ Wikipedia.

Settings in UNM For un-denoised hard negatives, we randomly select 30 out of the top 200 retrieved negatives from multiple retrievers. For denoised hard negatives, we select negatives with a score less than 0.1 output by an existing re-ranker (Qu et al., 2021).

Baselines We use the following baselines. **coCondenser** (Gao and Callan, 2021b) designs a complicated pretraining model structure on top of a vanilla PLM. **DPR-PAQ** (Oğuz et al., 2021) pre-trains a RoBERTa-large using 65-million-size synthetic QA pairs. Since the data is created by a model trained on NQ (Kwiatkowski et al., 2019) and Trivia QA (Joshi et al., 2017), it can be considered a semi-supervised method. It is also comparable to us as both of us use RoBERTa-large. **GTR** (Ni et al., 2021) pre-trains T5 encoder (Rafael et al., 2019) using 2-billion size community QA pairs. It also provides results across all model size ranges from T5-base to T5-xxlarge. The massive training corpus and model size establish a SOTA performance.

We also include some standard baselines including sparse retrieval systems (BM25, DeepCT (Dai and Callan, 2019), DocT5Query (Nogueira et al., 2019), and GAR (Mao et al., 2020)) and dense retrieval systems (DPR (Karpukhin et al., 2020), ANCE (Xiong et al., 2020), ME-BERT (Luan et al., 2020), and RocketQA (Qu et al., 2021)). We also include RocketQAv2 (Ren et al., 2021) as it jointly trains the retriever and reranker using hybrid sam-

Table 2: Passage retrieval results on MS-MARCO Dev and Natural Questions Test. We copy the results from the original papers. The best and second-best results are in bold and underlined fonts respectively.

Methods	PLM	MS-MARCO Dev		Natural Questions Test		
		MRR@10	R@1000	R@5	R@20	R@100
BM25	-	18.7	85.7	-	59.1	73.7
DeepCT(Dai and Callan, 2019)	-	24.3	91.0	-	-	-
docT5query(Nogueira et al., 2019)	-	27.7	94.7	-	-	-
GAR(Mao et al., 2020)	-	-	-	-	74.4	85.3
DPR(Karpukhin et al., 2020)	BERT-base	-	-	-	78.4	85.4
ANCE(Xiong et al., 2020)	RoBERTa-base	33.0	95.9	-	81.9	87.5
ME-BERT(Luan et al., 2020)	BERT-large	34.3	-	-	-	-
RocketQA(Qu et al., 2021)	ERNIE-base	37.0	97.9	74.0	82.7	88.5
RocketQAv2(Ren et al., 2021)	ERNIE-base	<u>38.8</u>	98.1	75.1	83.7	89.0
coCondenser(Gao and Callan, 2021b)	Condenser	38.2	98.4	75.8	84.3	89.0
DPR-PAQ(Oğuz et al., 2021)	RoBERTa-large	34.0	-	<u>76.9</u>	<u>84.7</u>	<u>89.2</u>
GTR(Ni et al., 2021)	T5-base	36.6	98.3	-	-	-
	T5-large	37.9	99.1	-	-	-
	T5-xxlarge	38.5	98.9	-	-	-
	T5-xxlarge	<u>38.8</u>	<u>99.0</u>	-	-	-
DPTDR	RoBERTa-large	39.1	98.9	77.5	85.1	89.4

pled negatives.

4.2 Experimental Results

4.2.1 Comparison with Existing Methods

Table 2 shows the dev set performance for MS-MARCO and test set performance for NQ. We can generally see that **DPTDR outperforms all the baselines in terms of MRR@10 on MS-MARCO and R@5 on NQ** and set a new SOTA in the two datasets.

We first compare DPTDR with DPR-PAQ. DPR-PAQ achieves competitive performance on NQ. It should be expected since it involves large semi-supervised pretraining on the NQ dataset. Nonetheless, DPTDR still outperforms DPR-PAQ by 0.6 points in R@5 although we use an unsupervised pretraining model. When we study the performance on MS-MARCO, DPR-PAQ fails to perform as consistently well as on NQ, which could result from domain mismatch of pretraining, and DPTDR outperforms it by a significant margin of 5.1 points in MRR@10.

Secondly, we compare DPTDR with GTR. GTR pre-trains T5 using 2-billion-size community QA pairs as a weakly-supervised pretraining. For such a scale of training corpus, we would expect that larger models consume the corpus more thoroughly and perform better on downstream tasks. As a result, GTR consistently boosts the performance on MS-MARCO with the model size increasing. How-

ever, DPTDR still outperforms GTR T5-xxlarge, a 10-billion-size model, and outperforms GTR T5-large by a noticeable margin of 1.2 points in MRR@10. It shows that model size is a positive contributor but not an absolute dominator for dense retrieval. Appropriate pretraining and negative mining can help improve performances using much more affordable computing resources. At the same time, note that DPT shall play a critical role in achieving comparable performance to FT with the help of RIP and UNM. We will validate this in Sec. 4.2.2.

Finally, we would like to compare DPTDR with coCondenser. Since coCondenser employs a pre-trained Condenser model(Gao and Callan, 2021a), we will conduct a more fair comparison in Sec. 4.4.

4.2.2 Comparing FT with and without RIP and UNM Strategies

To answer the raised question: *whether can we replace FT by DPT with comparable performance to SOTA FT methods in dense passage retrieval?* We conduct FT by following hyper-parameters of coCondenser (Gao and Callan, 2021b).

Comparison w/o RIP&UNM As a starter, we examine the effectiveness of directly replacing FT with DPT, which means we conduct training without RIP and UNM strategies. Thus we use the vanilla RoBERTa-large as the backbone model and BM25 negatives. As is shown in Table 3. We notice

Table 3: The comparison between FT and DPT with and without RIP and UNM strategies on MS-MARCO Dev and Natural Questions Test. **DPT with RIP&UNM** is the proposed method, a.k.a, ‘DPTDR’.

		MS-MARCO Dev		Natural Questions Test		
		MRR@10	R@1000	R@5	R@20	R@100
w/o RIP&UNM	FT	34.9	97.2	68.8	80.0	86.4
	DPT	32.7 (2.2 ↓)	96.3 (0.9 ↓)	66.5 (2.3 ↓)	78.5 (1.5 ↓)	85.5 (0.9 ↓)
w/ RIP&UNM	FT	39.4	99.0	77.0	85.4	89.2
	DPT	39.1 (0.3 ↓)	98.9 (0.1 ↓)	77.5 (0.5 ↑)	85.1 (0.3 ↓)	89.4 (0.2 ↑)

that DPT largely underperforms FT in this setting with a noticeable margin of 2.2 points in MRR@10 on MS-MARCO and 2.3 points in R@5 on NQ. It indicates that freezing most weights in DPT actually hinders its adaptability and therefore harms performance.

Comparison w/ RIP&UNM Next, we examine the performance of FT and DPT with RIP and UNM strategies. We use the RIP RoBERTa-large as the backbone model and UNM negatives. Table 3 shows that *i*) RIP and UNM improve the performances of both FT and DPT and *ii*) most importantly, DPT is comparable to FT under this setting, where the gap shrinks to only 0.3 points in MRR@10 on MS-MARCO, and DPT even slightly outperforms FT by 0.5 points in R@5 on NQ. As a result, we can see that when combined with RIP and UNM, DPT can obtain comparable performance with FT in dense retrieval.

4.3 Analysis on DPT

Sensitivity on prompt length We also seek to understand how prompt length affects the performance of DPT-based retrievers. From Table 4, we observe that the performance of prompt length of 8 already achieves a strong MRR@10 at 38.6 on MS-MARCO. When we increase the length to 128, it makes the most robust performance of MRR@10 at 39.1. The longer prompt means more trainable parameters, which obtains more power to steer PLMs. However, we also want to point out that the DPT retriever exhibits insensitivity to prompt length since the performances are competitive overall across various lengths. Therefore, we choose 32 as the default prompt length along with other hyperparameters in the main experiment for the rest of the ablation studies on MS-MARCO to accelerate the training.

Impact of reparameterization Reparametrization plays an important role in DPT. Li and Liang,

Table 4: Sensitivity of prompt length on MS-MARCO Dev.

Prompt Length	MRR@10	R@1000
8	38.6	98.9
16	38.6	99.0
32	38.7	98.9
64	38.5	98.9
128	39.1	98.9

2021 point out that MLP reparametrization results in more stable and compelling performances, while Liu et al., 2021b find it still depends on different tasks. In dense retrieval, we aim to determine whether it has a positive effect. Table 5 presents the results on MS-MARCO. We observe that MLP reparametrization results in a performance drop in MRR@10 on MS-MARCO. Since MLP breaks the independence of inter-layer prompts, we conjecture this brings optimization difficulty for dense retrieval.

Table 5: Ablations of reparameterization on MS-MARCO Dev.

Reparameterization	MRR@10	R@1000
embedding	38.7	98.9
mlp	38.0	99.0

4.4 Analysis on RIP

Whether to pre-train deep prompts or not? We try to examine whether pre-trained deep prompts could improve the performance of DPT-based retrievers. We use BERT-base as our backbone model and pre-train deep prompts of length 32 without reparameterization. The pretraining tasks and corpus are exactly the same as Sec. 3.2. We initialize DPT-based retrievers using pre-trained and randomly-initialized prompts. As is shown in Table 6, the pre-trained prompts do not boost the performance over randomly initialized prompts on

MS-MARCO. It reveals that the deep prompts may easily suffer from catastrophic forgetting.

Table 6: Ablations of prompt initialization on MS-MARCO Dev.

Prompt Initialization	MRR@10	R@1000
Random	32.4	95.5
Pre-trained	32.4	95.5

RIP on text spans or sentences We also explore pretraining using randomly-sampled sentences versus randomly-sampled text spans. Since coCondenser(Gao and Callan, 2021b) releases their pre-trained model using randomly-sampled text spans, we directly use their model to examine the zero-shot performance. For sampling sentences, we use the same PLM and hyper-parameters based on coCondenser code² except changing the training corpus consisting of randomly-sampled sentences. Table 7 presents the zero-shot performance on MS-MARCO. The pretraining using sentences works better than the one using text spans. This is might be owing to that text-spans do not consider the (starting and ending) borders of natural sentences and therefore break their completeness in semantics.

Table 7: Zero-shot performance of coCondenser with different sampling granularity (i.e., sentences or spans) on MS-MARCO Dev.

Unit	MRR@10	R@1000
Spans Gao and Callan (2021b)	11.1	78.2
Sentences	15.4	87.2

RIP’s effectiveness and comparison with co-Condenser We also try to examine the effectiveness of RIP strategy and compare it with coCondenser (Gao and Callan, 2021b). Concretely, we take vanilla RoBERTa-large, coCondenser RoBERTa-large, and RIP RoBERTa-large as the backbone model for DPT training under the same setting. Table 8 presents their results in both zero-shot and full-shot settings on MS-MARCO. For vanilla RoBERTa-large, it performs extremely poorly in zero-shot experiments, and with no surprise, it performs worst in full-shot experiments among the three PLMs. For co-Condenser RoBERTa-large, it achieves a noticeable improvement over vanilla RoBERTa-large,

²<https://github.com/lu Yug/Condenser>

where MRR@10 of zero-shot performance becomes meaningful at 6.3, and MRR@10 of full-shot performance increases to 37.3. For RIP RoBERTa-large, we see it achieves the best performance in both zero-shot and full-shot experiments. We also borrow the analysis tool from Wang and Isola (2020), which takes l_{align} between semantically-related positive pairs and $l_{uniform}$ of representation space to measure the quality of PLM representations. For both the metrics, lower numbers are better. RIP is much better than the vanilla model in both alignment and uniformity, while co-Condenser works well in alignment but worse in uniformity.

Thus a question is raised: *does PLM need additional structures for contrastive pretraining?* Both zero-shot and full-shot experiments demonstrate that RIP works even better than a carefully modified model structure. Therefore, we conjecture that PLM’s multi-layer transformers could be already expressive enough for dense retrieval under an appropriate contrastive learning task. However, additional model structures may bring optimization difficulty, especially when the number of added parameters is large.

4.5 Analysis on UNM

Ablation on UNM We try to understand how UNM affects performances. Table 9 presents the results on MS-MARCO. DPT using BM25 negatives achieves a baseline of MRR@10 at 36.8. When combining un-denoised hard negatives from multiple retrievers, we see that the performance achieves a noticeable improvement in MRR@10 by 1.5 points. When combining denoised hard negatives selected by a re-ranker, the performance further gets boosted of which MRR@10 increases by 0.4 points. The results demonstrate that both multiple retrievers and hybrid sampling positively contribute to dense retrieval.

5 Conclusion

In this paper, we investigate applying DPT in dense passage retrieval. To mitigate the performance drop of a vanilla DPT, We also propose two strategies, namely RIP and UNM, to enhance DPT and match the performance of FT. Experiments show that DPTDR outperforms previous state-of-the-art models on both MS-MARCO and Natural Questions and demonstrated the effectiveness of the above strategies. We believe this work facilitates the in-

Table 8: Ablations of different PLMs for DPT on MS-MARCO Dev.

Backbone PLM	Zero-shot				Full-shot	
	l_{align}	$l_{uniform}$	MRR@10	R@1000	MRR@10	R@1000
vanilla RoBERTa-large	161.4	-13.8	0.0	0.1	35.5	97.5
coCondenser RoBERTa-large	4.9	-12.9	6.4	63.3	37.3	98.0
RIP RoBERTa-large	21.9	-26.4	14.3	87.2	38.7	98.9

Table 9: Ablations of UNM on MS-MARCO Dev.

Neg Pool	MRR@10	R@1000
BM25 Neg	36.8	98.6
+ un-denoised Neg	38.3	98.9
+ denoised Neg	38.7	98.9

dustry, as it saves enormous efforts and costs of deployment and increases the utility of computing resources. In future work, we will explore scaling up the model size to further improve DPTDR.

Acknowledgment

Benyou Wang is funded by the CUHKSZ startup funding No. UDF01002678.

References

Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Wei-Cheng Chang, Felix X Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. Pre-training tasks for embedding-based large-scale retrieval. *arXiv preprint arXiv:2002.03932*.

Zhuyun Dai and Jamie Callan. 2019. Deeper text understanding for ir with contextual neural language modeling. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 985–988.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Luyu Gao and Jamie Callan. 2021a. Condenser: a pre-training architecture for dense retrieval. In *Proceed-*

ings of the 2021 Conference on Empirical Methods in Natural Language Processing, pages 981–993.

Luyu Gao and Jamie Callan. 2021b. Unsupervised corpus aware language model pre-training for dense passage retrieval. *arXiv preprint arXiv:2108.05540*.

Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.

Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Towards unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.

Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.

Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*.

- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2021b. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. *arXiv preprint arXiv:2110.07602*.
- 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*.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2020. Sparse, dense, and attentional representations for text retrieval. *arXiv preprint arXiv:2005.00181*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2020. Generation-augmented retrieval for open-domain question answering. *arXiv preprint arXiv:2009.08553*.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernández Ábrego, Ji Ma, Vincent Y Zhao, Yi Luan, Keith B Hall, Ming-Wei Chang, et al. 2021. Large dual encoders are generalizable retrievers. *arXiv preprint arXiv:2112.07899*.
- Rodrigo Nogueira, Jimmy Lin, and AI Epistemic. 2019. From doc2query to docttttquery. *Online preprint*.
- Barlas Oğuz, Kushal Lakhotia, Anchit Gupta, Patrick Lewis, Vladimir Karpukhin, Aleksandra Piktus, Xilun Chen, Sebastian Riedel, Wen-tau Yih, Sonal Gupta, et al. 2021. Domain-matched pre-training tasks for dense retrieval. *arXiv preprint arXiv:2107.13602*.
- 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.
- 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 preprint arXiv:1910.10683*.
- Ruiyang Ren, Yingqi Qu, Jing Liu, Wayne Xin Zhao, Qiaoqiao She, Hua Wu, Haifeng Wang, and Ji-Rong Wen. 2021. Rocketqav2: A joint training method for dense passage retrieval and passage re-ranking. *arXiv preprint arXiv:2110.07367*.
- Timo Schick and Hinrich Schütze. 2020. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate nearest neighbor negative contrastive learning for dense text retrieval. *arXiv preprint arXiv:2007.00808*.