



PROD: Progressive Distillation for Dense Retrieval

Zhenghao Lin
School of Informatics, Xiamen
University
zhenghaolin@stu.xmu.edu.cn

Yeyun Gong*
Microsoft Research Asia
yegong@microsoft.com

Xiao Liu
Microsoft Research Asia
xiaoliu2@microsoft.com

Hang Zhang
Microsoft Research Asia
v-zhang@microsoft.com

Chen Lin[†]
School of Informatics, Xiamen
University
chenlin@xmu.edu.cn

Anlei Dong
Jian Jiao
Jingwen Lu
Microsoft

Daxin Jiang
Rangan Majumder
Microsoft

Nan Duan
Microsoft Research Asia
nanduan@microsoft.com

ABSTRACT

Knowledge distillation is an effective way to transfer knowledge from a strong teacher to an efficient student model. Ideally, we expect the better the teacher is, the better the student performs. However, this expectation does not always come true. It is common that a strong teacher model results in a bad student via distillation due to the nonnegligible gap between teacher and student. To bridge the gap, we propose **PROD**, a **PRO**gressive Distillation method, for dense retrieval. PROD consists of a *teacher progressive distillation* and a *data progressive distillation* to gradually improve the student. To alleviate catastrophic forgetting, we introduce a regularization term in each distillation process. We conduct extensive experiments on seven datasets including five widely-used publicly available benchmarks: MS MARCO Passage, TREC Passage 19, TREC Document 19, MS MARCO Document, and Natural Questions, as well as two industry datasets: BING-REL and BING-ADS. PROD achieves the state-of-the-art in the distillation methods for dense retrieval. Our 6-layer student model even surpasses most of the existing 12-layer models on all five public benchmarks. The code and models are released in <https://github.com/microsoft/SimXNS>.

CCS CONCEPTS

• **Information systems** → **Document representation; Retrieval models and ranking.**

KEYWORDS

Neural Ranking Models, Dense Retrieval, Knowledge Distillation

*Equal contribution.

[†]Chen Lin is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW '23, April 30–May 04, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9416-1/23/04...\$15.00

<https://doi.org/10.1145/3543507.3583421>

ACM Reference Format:

Zhenghao Lin, Yeyun Gong, Xiao Liu, Hang Zhang, Chen Lin, Anlei Dong, Jian Jiao, Jingwen Lu, Daxin Jiang, Rangan Majumder, and Nan Duan. 2023. PROD: Progressive Distillation for Dense Retrieval. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, April 30–May 04, 2023, Austin, TX, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3543507.3583421>

1 INTRODUCTION

In recent years, pre-trained models have made breakthroughs in various NLP tasks, including question answering [7], summarization [21, 32], and dense retrieval [17, 52]. To further improve the performance of end tasks, large models are proposed [2, 42]. Despite their successful applications on small or medium scale benchmarks, the efficiency issue of model inference becomes a problem. In the practical scenarios of dense retrieval, online systems need to retrieve the relevant documents from a large number of candidates, and answer the user queries in time. Therefore, an efficient small model is particularly critical in dense retrieval applications.

To take into account both performance and efficiency, knowledge distillation techniques have been widely used [15, 26, 40]. In previous work, Zeng et al. [50] proposed a curriculum learning method for dense retrieval distillation. Ren et al. [35] and Zhang et al. [52] proposed the re-ranker as teacher method. Lin et al. [24] proposed an in-batch negative distillation method with ColBERT [18]. Lu et al. [26] proposed interaction distillation, cascade distillation, and dual regularization to bring the re-ranker with retriever. These methods demonstrate the importance of knowledge distillation in dense retrieval tasks from different perspectives. However, when the gap between student and teacher is very large, how to better close the gap is a big challenge. Both the experiments in the previous NLG task [55] and our experiments in the dense retrieval task provide ample support for the proposition that *the performance of the teacher and the student is not positively correlated*.

To solve this problem, there is a branch of work utilizing progressive distillation, in which teachers are dynamically adjusted to distill students. Apart from refining the pretraining workflow [43], progressive distillation methods have achieved great success in various down-stream tasks in NLP and CV [1, 14, 36, 39]. The main idea of the existing progressive distillation work is to use stronger

teachers while educating students. However, there are two shortcomings. First, this idea has not been proven to be useful in dense retrieval yet. Second, the existing studies [1, 14, 36, 39, 43] mainly focus on teacher models, neglecting the importance of training data in the process of progressive distillation.

In this paper, we propose a progressive distillation method, PROD, to minimize the gap between the teacher and the student. PROD consists of two progressive mechanisms: *teacher progressive distillation* (TPD) and *data progressive distillation* (DPD). In TPD, we gradually improve the capability of teachers by using different architectures, enabling student model to learn knowledge progressively. In DPD, we start to let students learn from all the data, and then gradually select samples that the student is confused about for strengthening. In each progressive step, we introduce regularization loss to avoid catastrophic forgetting of the knowledge memorized in the previous step.

The motivation of PROD is from two aspects: 1) When the gap between teacher and student is huge, a stronger teacher is not necessarily better than a relatively weak teacher. Such as, the university professor may not be more suitable than a kindergarten teacher to teach a kindergarten student. Therefore, we design TPD to enhance teachers gradually. 2) There are different knowledge suitable for the student model to learn at different stages, such as middle school textbooks are suitable for middle school students to learn. Thus, we design DPD to select the appropriate (not too easy or hard) knowledge for the student to learn.

We conduct extensive experiments on five widely-used benchmarks (MS MARCO Passage, TREC Passage 19, TREC Document 19, MS MARCO Document, and Natural Questions) and two industry datasets (BING-REL and BING-ADS). The results of extensive experiments on five popular benchmarks show the effectiveness of PROD, and performance on two industry datasets also demonstrates the commercial value of PROD.

2 RELATED WORK

This work is related to two lines of work.

2.1 Dense Retrieval

Compared with sparse retrieval methods [6, 30, 47], dense retrieval has the potential to find hidden semantic correlations between queries and passages. Several directions have been explored to improve the performances of the popular dual encoder structure, including finding hard negatives with higher qualities [17, 33, 46], multi-vector interactions [18], and the joint training of retrievers and re-rankers [35, 52].

Meanwhile, another line of the work studied knowledge distillation for dense retrieval by using a single teacher model [35, 53], multiple teacher models with joint training techniques [26] and curriculum learning [50]. Our work is in line with these work, while having the key difference that we choose one teacher model with progressive ability level and focusing on the confusing data with progressive difficulty in each training stage.

2.2 Knowledge Distillation

The knowledge distillation [11] has been widely studied for decades. There are a variety of off-the-shelf practical techniques, such as the

response-based method [19], the feature-based method [37] and distilling the attention scores [49] with Kullback–Leibler divergence (KLD) or mean squared error (MSE). Recently, a rising group of work focuses on the relationships between the teacher model and the student model. Some elaborate the knowledge transmission methods [4, 56], some introduce middle models [28], and others work on distillation with multiple teachers [45]. It can be noted that how to use a more reasonable teacher model to guide the learning of the student model has gradually become the current key direction.

Among the studies on knowledge distillation, there is a branch of work on progressive distillation, in which teachers are dynamically adjusted to distill students. In spite of the success on the workflow of pretraining [43], progressive distillation methods are widely proven to be effective in down-stream tasks, such as image classification [36], image generation [39], GLUE [14, 36], question answering [36], and cross-modal representation learning [1]. In each task, the main idea of the existing progressive distillation work is to use stronger teachers while educating students. Different from the above studies, PROD is adopt to the research field of dense retrieval, using a progressive method from two perspectives, teacher and data, in which teachers with different architectures are applied in a progressive order and more and more confusing data is mined to fill the performance gap between the teacher and the student.

3 METHODOLOGY

3.1 Task Description

We follow the basic definition of dense text retrieval [17, 18, 46, 52]. Assume we have a query set $\mathcal{Q} = \{q_1, q_2, \dots, q_n\}$ containing n queries and a passage set $\mathcal{P} = \{p_1, p_2, \dots, p_m\}$, our target is to find the most relevant passages in corpus \mathcal{P} for each query q_i .

3.2 Dual Encoder

The currently popular architecture for dense retrieval is dual encoder (DE), which can encode queries and passages into dense vectors \mathbf{p} and \mathbf{q} , respectively, and calculate the similarity scores through the inner product as:

$$s_{de}(\mathbf{q}, \mathbf{p}) = E_Q(\mathbf{q})^T \cdot E_P(\mathbf{p}) \quad (1)$$

where $E_Q(\cdot)$ and $E_P(\cdot)$ are the query encoder and the passage encoder, respectively. Based on the embedding vectors, existing solutions generally employ approximate nearest neighbor (ANN) search algorithms like faiss [16].

3.3 Cross Encoder

In order to capture the fine-grained relationships between a pair of query and passage, the cross encoder (CE) is often used as the re-ranking model, rearranging the retriever outputs to improve the retrieval performances. In particular, the concatenation of a query \mathbf{q} and a passage \mathbf{p} are with the special token [SEP] is the input of CE. The similarity score is calculated by the [CLS] embedding vector of CE as:

$$s_{ce}(\mathbf{q}, \mathbf{p}) = \mathbf{w}^T \cdot E_{ce}([\mathbf{q}; \mathbf{p}]) \quad (2)$$

where $E_{ce}(\cdot)$ is an encoder that can be initialized with any pre-trained language models, $[\cdot]$ is concatenation operation, and \mathbf{w} is the linear projector.

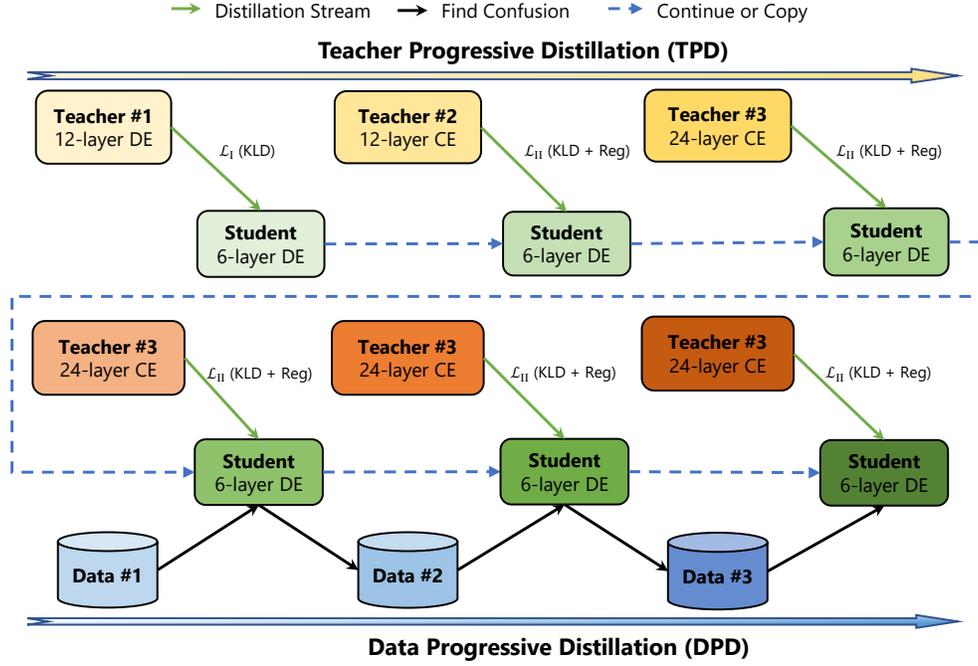


Figure 1: The framework of PROD. In the TPD, we used three types of teachers, namely, 12-layer DE, 12-layer CE, and 24-layer CE. In the DPD, we used 24-layer CE to mine confusing data and iterative distillation training. Each step of distillation is continued with the 6-layer DE student in the previous step. Furthermore, to alleviate catastrophic forgetting, we employ a regularization loss item to maintain the stability of distillation.

3.4 Knowledge Distillation

Knowledge distillation is a simple yet effective approach to boost a small model (student model) with the additional knowledge from a big model (teacher model) [11]. In the task of dense retrieval, the student model is usually trained with **hard labels** such as the relevance annotations for each query-passage pair using a contrastive loss. Additionally, it can also benefit from comparing with the prediction distribution of the teacher model, *aka soft labels*, by decreasing a measurement like KLD and MSE, indicating the differences between the prediction distributions [26, 35, 41].

4 METHOD

We show the framework of the progressive distillation method, namely PROD, in Figure 1. The main idea of PROD is to gradually minimize the gap between a well-trained teacher model (24-layer CE) and the target student model (6-layer DE) by two sequential procedures, which are called *teacher progressive distillation* (TPD) and *data progressive distillation* (DPD).

4.1 Teacher Progressive Distillation

Previous work found that different data instances would have various difficulties for a particular family of models to learn [8, 44]. Similarly, we believe that the probability distributions of different teacher models when selecting the relevant passages from a candidate pool are also different. Therefore, when facing the various difficulties of data instances in training, we use three different teacher models with different ability levels to civilize the student model gradually. More specifically, we employ three progressive teacher models, *i.e.*, a 12-layer DE, a 12-layer CE and a 24-layer

CE, to boost a 6-layer DE student model¹ in order. Before learning from a teacher model, we retrieve the top- k negatives [46] with the current student checkpoint, randomly sample a subset and train a teacher model using the mined hard negatives.

Distilling with a DE Teacher. We first use the original data to warm up the 12-layer teacher DE and the 6-layer student DE, and then use the isomorphic distillation, *i.e.*, 12-layer DE \rightarrow 6-layer DE. The loss function can be divided into two parts.

The first part is calculated by the output distribution of the student model and hard labels, which is called **hard loss**:

$$\mathcal{L}_h(q, \mathbb{P}^+, \mathbb{P}^-) = -\log \frac{\exp(s_{de}^S(q, \mathbb{P}^+))}{\sum_{\mathbb{P}' \in \{\mathbb{P}^+, \mathbb{P}^-\}} \exp(s_{de}^S(q, \mathbb{P}'))} \quad (3)$$

where \mathbb{P}^+ and \mathbb{P}^- is the relevant passage and negative passage pool of q , respectively; s_{de}^S is the similarity scores of the student model.

The other part measures the differences of the probability distributions between the teacher model and the student model given the same batch of data, which is called **soft loss**:

$$d_{de}^S = \frac{\exp(s_{de}^S(q, \mathbb{P}))}{\sum_{\mathbb{P}' \in \{\mathbb{P}^+, \mathbb{P}^-\}} \exp(s_{de}^S(q, \mathbb{P}'))} \quad (4)$$

$$d_{de}^T = \frac{\exp(s_{de}^T(q, \mathbb{P}))}{\sum_{\mathbb{P}' \in \{\mathbb{P}^+, \mathbb{P}^-\}} \exp(s_{de}^T(q, \mathbb{P}'))} \quad (5)$$

$$\mathcal{L}_{de \rightarrow de}(q, \mathbb{P}^+, \mathbb{P}^-) = D_{KL}(d_{de}^T / \tau, d_{de}^S / \tau) \quad (6)$$

where s_{de}^T is the similarity scores of the teacher model; τ is the temperature of distillation. At last, the final distillation loss with a DE teacher is a weighted sum of the before-mentioned two parts:

$$\mathcal{L}_I = \alpha_1 \mathcal{L}_h + \beta_1 \mathcal{L}_{de \rightarrow de} \quad (7)$$

¹We use the notation “ $X \rightarrow Y$ ” to denote the distillation stage of a teacher model X and a student model Y .

where α_1 and β_1 control the importance of hard loss and soft loss.

There are two important reasons why we use isomorphic distillation. (1) First, according to previous research on DE [17], adding more hard negatives when training DE can not lead to further improvements. Therefore, simply training a DE may be not suitable for difficult data instances. Meanwhile, the special in-batch negative techniques [17] for DE allows it to assign more appropriate similarity scores for easy negatives, which is more suitable for the first stage training of the student model. (2) Second, the performances of alone trained DE are limited. More specifically, when the student model distills to the later stage, it is difficult to quickly train a DE teacher model with better performance than the student model.

Distilling with CE Teachers. After learning with a 12-layer DE, we use two CEs as the teacher models. Since CE can capture the fine-grained correlations between a pair of query and passage, it is more suitable for training with difficult data. However, CE can not make an effective distinction between hard and easy negatives. Such probability distributions of CE teachers are somehow confusing and inconducive for the student model. Therefore, when deploying CE as the teacher model, we only use hard negatives in distillation.

In our preliminary study, we also find that it is critical to select a proper CE. Generally speaking, the CE with more parameters can give more accurate predictions, but the more accurate predictions may not benefit distillation [28, 55]. When the difficulties of training instances are not very conflating, the premature use of CE will lead to the performance loss of student model. We think the main reason is that CE with high performances will output so confident results that the predicted probability distribution is often unfavorable to the student model. Therefore, we perform 12-layer CE \rightarrow 6-layer DE before 24-layer CE \rightarrow 6-layer DE.

To overcome the catastrophic forgetting in learning with multiple teachers, we additionally follow Li and Hoiem [22] and Cao et al. [3], using regularization to maintain the stability of training. We first save a frozen copy of the student model as S' and involve a **regularization loss** item \mathcal{L}_r in distilling the student model S :

$$d_{de}^{S'} = \frac{\exp(s_{de}^{S'}(q, p))}{\sum_{p' \in \{p^+, p^-\}} \exp(s_{de}^{S'}(q, p'))} \quad (8)$$

$$\mathcal{L}_r(q, p^+, p^-) = D_{KL}(d_{de}^{S'} / \tau, d_{de}^S / \tau) \quad (9)$$

where $s_{de}^{S'}$ is calculated by the copied student S' .

The overall loss function consists of three parts: the hard loss in equation 3, the soft loss that adopts the relevant scores s_{ce}^T calculated by the CE teacher model, and the regularization loss \mathcal{L}_r . Finally, the loss function with a CE teacher is like:

$$d_{ce}^T = \frac{\exp(s_{ce}^T(q, p))}{\sum_{p' \in \{p^+, p^-\}} \exp(s_{ce}^T(q, p'))} \quad (10)$$

$$\mathcal{L}_{ce \rightarrow de}(q, p^+, p^-) = D_{KL}(d_{ce}^T / \tau, d_{de}^S / \tau) \quad (11)$$

$$\mathcal{L}_{II} = \alpha_2 \mathcal{L}_h + \beta_2 \mathcal{L}_{ce \rightarrow de} + \gamma \mathcal{L}_r \quad (12)$$

where α_2 , β_2 and γ control the relative importance.

4.2 Data Progressive Distillation

After learning from progressively stronger teachers, the student model will be improved with knowledge distillation. But there are still many confusing negative passages that lead to the disagreement between the teacher model and the student model, limiting the performance of the student model. Therefore, we aim to fill the

performance gap between the 24-layer CE teacher and the student model from the perspective of training data.

Intuitively, the data difficulty needs to be progressively raised to fit the capacity of the teacher, enhancing the student model in a clearer direction. Therefore, sequentially after the last step of TPD, our solution is to adopt an iterative data selection procedure (let us say there are N iterations), where each iteration consists of the following four steps:

(1) Retrieve the top- k negatives with the current student.

(2) Collect the queries that the teacher model can predict the positive as top-1 but the student model can only predict as top- k' . We construct a dataset \mathbb{D}^i for the i -th iteration as the collection of those queries, whose positive passages are the labeled ones and the negative passages are mined in step (1).

(3) Continually train the 24-layer CE teacher model with \mathbb{D}^i .

(4) Use \mathbb{D}^i to distill with the teacher and the loss function in equation 12.

5 EXPERIMENTS

5.1 Experimental Setting

We conduct experiments on several text retrieval datasets: MS MARCO Passage Ranking (MS-PAS) [29], TREC 2019 Deep Learning Track (TREC-PAS-19, TREC-DOC-19) [5], MS MARCO Document Ranking (MS-DOC) [29], Natural Questions (NQ) [20], and two industry datasets (BING-REL and BING-ADS). BING-REL contains multilingual positive and negative query-document pairs from Bing, where clicked documents are selected from search log data as positives and top- k retrieval documents are mined against a fixed doc corpus as the hard negatives following ANCE [46]. We collect high quality click data from Bing ads for BING-ADS (here we only select ads with match type *Phrase Match*²), the clicks resulting in user dwell time on landing pages for greater than 20 ms are considered as high quality clicks to exclude randomly accidental clicks. The corpus is formulated as “query, ad keyword”, the “ad keyword” is the keyword corresponding to the clicked ad. We use the data from November 2021 to July 2022 as training set, the data from August 1, 2022 to August 10, 2022 and from August 10, 2022 to August 20, 2022 as validation set and test set, respectively. The statistics are shown in Appendix A.

For MS-PAS, We follow the existing work [35, 51], reporting **MRR@10**, **Recall@50** and **Recall@1k** on the dev set. For TREC-PAS-19, we select **nDCG@10** and **MAP@1k** as the evaluation metrics. For MS-DOC, we report **MRR@10** and **Recall@100** on the dev set. For TREC-DOC-19, we select **nDCG@10** and **Recall@100** as the evaluation metrics. For NQ, we choose **Recall@5**, **Recall@20** and **Recall@100** as the evaluation metrics. For BING-REL and BING-ADS, we also treat **MRR@10**, **Recall@5**, **Recall@20** and **Recall@100** as the evaluation metrics. We conduct significant tests based on the paired t-test with $p \leq 0.01$.

5.2 Baselines

We compare PROD with two groups of baselines. The first group contains sparse retrieval methods and dense retrieval methods without knowledge distillation or multiple vectors, including BM25 [47], DeepCT [6], docT5query [30], SPARTA [54], GAR [27], DPR

²<https://he1p.ads.microsoft.com/apex/index/3/en-us/50822>

[17], ANCE [46], RDR [48], Joint Top- k [38], DPR-PAQ [31], Ind Top- k [38], STAR [51], and ADORE [51].

The second group is about dense retrieval methods producing a single vector for each document and query enhanced by knowledge distillation, including Margin-MSE [12], TCT-CoBERT [23], TAS-B [13], SPLADE v2 [9], RocketQA v1 [33], RocketQA v2 [35], PAIR [34], and CL-DRD [50].

5.3 Implementation Details

Model Initialization. Following the settings of RocketQA v2 [35], both the 12-layer DE and the 12-layer CE use ERNIE-2.0-BASE as the encoders for the public datasets and BING-ADS. Besides, the 24-layer CE uses ERNIE-2.0-LARGE as the encoder. Please note that a well-pretrained checkpoint is not required for the student. Therefore, we adopt the first six layers of ERNIE-2.0-BASE as the 6-layer DE student for simplicity. For BING-REL, we use BERT-MULTILINGUAL-BASE as the encoders for 12-layer DE and 12-layer CE, simply taking the first six layers of BERT-MULTILINGUAL-BASE as the student.

Distillation Warming Up. Before TPD, we need to warm up the first teacher. We use random or BM25 negatives to train a 12-layer DE, retrieving the top- k negatives, retraining a 12-layer DE with the mined negatives as the first teacher for MS-PAS and MS-Doc. Especially, while retrieving the top- k negatives for NQ, we additionally filter positive passages by matching the passages and answers.

Hyper-parameter Setting. For mining the hard negatives, we choose $k = 1000$ for the public datasets and $k = 100$ for the industry datasets. In DPD, we set $k' = 15$ and the iteration number $N = 1$. We use AdamW [25] as the optimizer. Other detailed hyper-parameters for reproducing our experiments are shown in Appendix B.

5.4 Main Results

The results comparing PROD with the baselines on MS-PAS, TREC-PAS-19, TREC-Doc-19, NQ and MS-Doc are shown in Table 1, Table 2, Table 3, respectively. It can be easily observed that PROD achieves state-of-the-art results in all five datasets, which can be concluded from two perspectives.

With or Without Distillation. As we can see from Table 1 and Table 2, the baselines are divided into two groups, representing without and with distillation. Comparing against ADORE (the best on MS-PAS, TREC-PAS-19, MS-Doc and TREC-Doc-19 without distillation), Ind Top- k (the best on NQ without distillation), RocketQA v2 (the best on MS-PAS and NQ with distillation) and CL-DRD (the best on TREC-PAS-19 with distillation), PROD can achieve better performances in all the groups.

6-layer or 12-layer. In Table 1, the performance of the best 12-layer method RocketQA v2 is better than the best 6-layer method CL-DRD on MS-PAS. However, PROD further exceeds RocketQA v2 with a 6-layer architecture. Besides, among all the 12-layer baselines on NQ and MS-Doc, PROD achieve the best results, showing the effectiveness of the progressive distillation method even with a 6-layer student.

5.5 Comparison with Multi-Teacher Methods

To further prove the effectiveness of our method, we compare PROD with three methods that also incorporated with multiple teachers.

Random Batch [10] randomly selects a teacher in each batch of training. Merge Score [10] averages the soft labels of multiple teachers in training. And Merge Loss [10] adds up all the distillation loss of each teacher, before calculating gradients. Please note that all the methods use the same set of teachers as PROD, *i.e.*, 12-layer DE, 12-layer CE and 24-layer CE. We also report the results of training a 6-layer DE without distillation, which is denoted as “Pure Student”.

As we can see from the results of MS-PAS in Table 4, all the multi-teacher methods can lead to performance gain comparing with Pure Student. Besides, PROD can achieve the best performances among these multi-teacher methods even just uses TPD, which also reflects the effectiveness of our method.

5.6 Ablation Study

To investigate the effectiveness of TPD and DPD, we conduct a careful ablation study on the two parts. Please note that we use “12DE->12CE->24CE->DTD” to denote our method PROD, showing the specific steps inside.

Effect of TPD. The performances of each distillation step on MS-PAS and NQ are shown in Table 5. Compared with the original 6-layer DE without distillation, the performance of “12DE->12CE->24CE->DTD” on MS-PAS has been improved by about 7.7% in MRR@10, 7.1% in Recall@50 and 2.2% in Recall@1k, proving the overall effectiveness of PROD method. In addition, by comparing with the actual teacher variants used in PROD, we can easily find that each step of PROD has achieved stable improvement. We focus on the selection of the teacher model in each step and conduct further experiments to prove the correctness and necessity of the teacher’s order in each step of PROD.

- **Influence of Teacher #1.** How to choose Teacher #1 is a problem worth exploring. Therefore, in the first step of teacher progressive, we use three different teachers, “12DE”, “12CE” and “24CE”, to distill 6-layer DE. Although the experimental results show that “12DE” on MS-PAS is worse than “12CE” and “24CE” for MRR@10 and Recall@100, there are several important reasons why we still choose 12-layer DE as Teacher #1.

- (1) Comparing “12DE” with “12CE” and “24CE”, Recall@50 on MS-PAS are higher than CEs. We believe that this is because DE uses in-batch negatives, which increases the overall understanding of retrieval task, and is helpful for the student model to further learn more difficult knowledge.

- (2) Compared with more later distillation steps, it is surprising that we can not get better results by adjusting the order of DE distillation backward or completely abandoning DE. It shows that DE distillation mainly increases the model’s cognition of easy negatives. With the increase of data difficulty and student model performance in the training, it is too late to use DE distillation.

- (3) In Table 6, we can see the performance of different teacher models on MS-PAS and NQ, where CE performances are measured by reranking the results of DE. Although the performance of 12-layer DE is the worst, the distillation efficiency is the highest, that is, the student model is the closest to the teacher’s performance after distillation. This phenomenon is more obvious on NQ shown in Table 5, where the student model can perform best by using 12-layer DE as Teacher #1, while the performance of the 12-layer DE is the worst compared with CE. These results further support

Method	#Params	MS-Pas			TREC-Pas-19	
		MRR@10	Recall@50	Recall@1k	nDCG@10	MAP@1k
BM25 [47]	-	18.7	59.2	85.7	49.7	29.0
DeepCT [6]	-	24.3	69.0	91	55.0	34.1
docT5query [30]	-	27.2	75.6	94.7	64.2	40.3
ANCE [46]	12-layer (110M)	33.0	-	95.9	64.8	37.1
ADORE [51]	12-layer (110M)	34.7	-	-	68.3	41.9
TCT-ColBERT [23]	12-layer (110M)	33.5	-	96.4	67.0	39.1
RocketQA v1 [33]	12-layer (110M)	37.0	85.5	97.9	-	-
PAIR [34]	12-layer (110M)	37.9	86.4	98.2	-	-
RocketQA v2 [35]	12-layer (110M)	38.8	86.2	98.1	-	-
Margin-MSE [12]	6-layer (66M)	32.3	-	95.7	69.9	40.5
TAS-B [13]	6-layer (66M)	34.4	-	97.6	71.7	44.7
SPLADE v2 [9]	6-layer (66M)	36.8	-	97.9	72.9	-
CL-DRD [50]	6-layer (66M)	38.2	-	-	72.5	45.3
PROD	6-layer (66M)	39.3^{*†‡§}	87.0^{*†}	98.4^{*†‡}	73.3^{‡§}	48.4[§]

Table 1: The main results on MS-Pas and TREC-Pas-19. “#Params” represents the number of trainable parameters. We use the paired t-test with $p \leq 0.01$. The superscripts refer to significant improvements compared to PAIR(*), RocketQA v2(†), SPLADE v2(‡), CL-DRD(§).

Method	Recall@5	Recall@20	Recall@100
BM25 [47]	-	59.1	73.7
GAR [27]	60.9	74.4	85.3
DPR [17]	-	78.4	85.4
ANCE [46]	-	81.9	87.5
RDR [48]	-	82.8	88.2
Joint Top-k [38]	72.1	81.8	87.8
DPR-PAQ [31]	74.5	83.7	88.6
Ind Top-k [38]	75.0	84.0	89.2
RocketQA v1 [33]	74.0	82.7	88.5
PAIR [34]	74.9	83.5	89.1
RocketQA v2 [35]	75.1	83.7	89.0
PROD	75.6[*]	84.7^{*†}	89.6^{*†}

Table 2: The main results on NQ. All the baselines are 12-layer, while the student in PROD is 6-layer. We use the paired t-test with $p \leq 0.01$. The superscripts refer to significant improvements compared to Ind Top-k(*), RocketQA v2(†).

Method	MRR@10	Recall@50	Recall@1k
Pure Student	31.66	79.99	96.19
Random Batch	37.94	85.82	98.09
Merge Score	37.64	84.93	97.68
Merge Loss	38.09	86.10	98.09
TPD	38.75	86.56	98.41
TPD+DPD	39.34	87.06	98.44

Table 4: The results of different methods with multiple teachers on MS-Pas.

between teacher and student. To verify this opinion, we increase the gap between teacher and student, repeated the experiment of the second step on a 2-layer DE student model. The results are shown in Table 5, which makes the comparison of taking 24-layer CE or 12-layer CE as the teacher in the second step more obvious.

Moreover, in order to explore whether 12-layer CE distillation can also perform well in the third step, we replaced the 24-layer CE in the third step with 12-layer CE, which is denoted as “12DE->12CE->12CE”. By comparing with “12DE->12CE->24CE”, we can see that 24-layer CE, a more powerful teacher, must be used in the third step to achieve the best results, and this process must be transited from the 12-layer CE distillation.

Specifically, by comparing “12DE->12CE->24CE” with “12DE->24CE->24CE”, we can see that replacing the 12-layer CE with the 24-layer CE in the second step for distillation cannot achieve satisfactory performance. Therefore, it is necessary and reasonable for us to adopt two different architectures of CE, which are a 12-layer CE teacher in the second distillation step and a 24-layer CE teacher in the third distillation step in TPD.

Effect of DPD. We also care about how much DPD contributes to our framework PROD. By comparing “12DE->12CE->24CE->DTD” with “12DE->12CE->24CE” on MS-Pas and NQ, we can see that appending DTD after TPD can lead to consistent improvements on all the evaluation metrics used in both MS-Pas and NQ, showing the necessity of continual distillation using the confusing negative passages. Further more, to verify whether the performance gain comes from the confusing data, we design another experiment by continual distillation on all the queries with the refreshed hard

that the selection of distillation teachers at different steps should not take the performance of teacher models as the only standard.

• **Influence of Teacher #2.** In the second and third steps of teacher progressive, we adopt 12-layer CE and 24-layer CE as teacher models to distill student model. As shown in Table 5, using 24-layer CE as Teacher #2 is almost the same or even worse than 12-layer CE. We believe that this is caused by the excessive gap

Dataset	Teacher Variant	Student	MRR@10	Recall@1	Recall@5	Recall@20	Recall@50	Recall@100	Recall@1k
MS-PAS	-	6DE	31.66	20.14	47.45	68.95	79.99	85.59	96.19
	12DE*	6DE	35.69	23.65	52.54	73.94	84.09	89.31	97.59
	12CE	6DE	37.44	25.11	54.11	75.03	84.44	88.81	97.15
	24CE	6DE	36.18	23.80	53.08	74.57	84.15	89.13	97.46
	12DE->12CE*	6DE	38.09	25.34	55.54	76.75	86.12	90.83	97.72
	12CE->12DE	6DE	34.90	22.98	51.07	73.01	82.95	88.44	97.22
	12DE->24CE	6DE	38.08	25.46	55.46	76.76	85.72	90.53	97.82
	12DE->12CE->24CE*	6DE	38.75	25.89	55.86	77.68	86.56	91.33	98.41
	12CE->12CE->24CE	6DE	38.29	25.87	55.14	75.64	84.58	88.62	96.40
	12DE->12CE->12CE	6DE	37.88	24.94	55.53	76.62	86.12	90.76	98.01
	12DE->24CE->24CE	6DE	37.63	24.87	54.99	76.39	85.54	90.26	98.07
	12DE->12CE->24CE->DTD* (PROD)	6DE	39.34	26.66	56.35	78.09	87.06	91.52	98.44
	12DE->12CE->24CE->24CE	6DE	37.98	25.06	55.32	77.06	85.93	90.50	97.54
	-	2DE	27.31	17.18	40.83	61.65	72.79	79.79	93.57
	12DE	2DE	30.78	19.27	46.48	67.87	78.88	85.32	96.29
	12DE->12CE	2DE	34.10	22.41	50.59	71.15	80.62	86.32	95.56
12DE->24CE	2DE	32.92	21.68	48.47	69.84	80.24	85.95	95.49	
NQ	-	6DE	-	45.96	67.76	78.88	-	86.24	-
	12DE*	6DE	-	54.82	73.73	82.70	-	88.29	-
	12CE	6DE	-	52.16	71.54	81.78	-	88.24	-
	24CE	6DE	-	51.19	70.57	80.62	-	87.15	-
	12DE->12CE*	6DE	-	57.12	75.12	84.30	-	88.82	-
	12DE->24CE	6DE	-	56.87	74.75	84.22	-	88.81	-
	12DE->12CE->24CE*	6DE	-	57.20	75.54	84.66	-	89.48	-
	12DE->12CE->24CE->DTD* (PROD)	6DE	-	57.63	75.61	84.72	-	89.56	-

Table 5: The results of ablation study on MS-PAS and NQ. “*” means the actual steps in PROD. For simplicity, we use “nDE” and “nCE” to denote the n-layer DE and n-layer CE, respectively. We use “A->B” to denote continual distillation using A and B as the teachers in turn with refreshed top-k negatives before learning from each teacher.

Dataset	Teacher	MRR@10	Recall@1	Recall@5	Recall@20	Recall@50	Recall@100	Recall@1k
MS-PAS	12DE	35.77	23.50	52.72	74.18	84.10	89.14	97.61
	12CE	40.81	27.88	58.58	79.04	86.52	90.81	97.61
	24CE	41.96	28.84	59.81	80.00	87.89	92.23	97.61
NQ	12DE	-	52.94	73.15	82.70	-	88.68	-
	12CE	-	60.38	79.46	86.02	-	88.68	-
	24CE	-	64.78	80.57	86.24	-	88.68	-

Table 6: The performance of the teachers in the first step of TPD on MS-PAS and NQ. For simplicity, we use “nDE” and “nCE” to denote the n-layer DE and n-layer CE, respectively. The performances of CEs are the reranking results based on the retrieval output of the 12-layer DE.

negatives using 24-layer CE teacher, which is denoted as “12DE->12CE->24CE->24CE”. By comparing “12DE->12CE->24CE->DTD” with “12DE->12CE->24CE->24CE”, we can observe performance improvements on MS-PAS. We think the reason may be that the noises in the entire training set obstructs the student’s learning, which also means that mining a small amount of confusing data in DPD is necessary and improves the efficiency of distillation.

5.7 Results on Industry Datasets

We conduct experiments on the two industry datasets by comparing each steps in PROD. Please note that the scales of the industry datasets are larger than those of the public datasets. Therefore, considering the cost of computational resources to reach the model’s convergence, we only use 12-layer DE and 12-layer CE in PROD. For BING-REL, we simply evaluate the last checkpoint after training

and report the results on the dev set. The results on BING-REL and BING-Ads are illustrated in Table 7.

By comparing the results of each step in PROD in the two industry datasets, it is obvious that the performances increase consistently when more teachers or DPD are applied. It indicates that PROD can lead to significant performance gain in practical industrial scenarios. Additionally, after distilling using 12-layer CE, the performances of the 6-layer DE student can outperform the 12-layer DE teacher. Besides, we can see that the performances of the 12-layer CE teacher are much better than the student even after distilling with that teacher. But if applying DPD afterward, the performance of the student is much closer to the 12-layer CE teacher, which supports the fact that PROD has the potential to better alleviate the gap between the teacher and the student.

Dataset	Teacher	Student	MRR@10	R@5	R@20	R@100
BING-REL	-	6DE	39.51	54.27	70.78	83.02
	12DE	6DE	42.46	56.82	73.81	85.92
	12DE->12CE	6DE	43.81	58.80	74.18	85.72
	PROD	6DE	44.37	59.71	74.72	85.79
	12DE	-	42.51	57.71	73.89	85.81
	12CE	-	49.92	66.21	80.45	85.81
BING-ADS	-	6DE	22.30	33.54	57.39	82.44
	12DE	6DE	22.98	34.32	57.90	82.80
	12DE->12CE	6DE	24.07	35.54	59.24	83.71
	PROD	6DE	24.82	36.93	60.47	84.30
	12DE	-	23.46	35.25	59.59	84.32
	12CE	-	25.37	37.18	60.02	84.32

Table 7: The results on two industry datasets BING-REL and BING-ADS. R@n is short for Recall@n. We use “nDE” and “nCE” to denote the n-layer DE and n-layer CE, respectively. We use “A->B” to denote continual distillation using A and B as the teachers in turn with refreshed top-k negatives before learning from each teacher. The performances of CEs are based on the output of the 12-layer DE.

Method	MRR@10	Recall@50
TPD	38.75	86.56
ST-(1,2] & TT-(0,1]	39.13	86.89
ST-(1,5] & TT-(0,1]	39.26	87.02
ST-(1,15] & TT-(0,1]	39.34	87.06
ST-(5,20] & TT-(0,5]	39.29	86.82
ST-(0,15] < TT-(0,15]	39.30	86.98
ST-(0,31] < TT-(0,31]	39.13	87.12

Table 8: The impact of different strategies of selecting the confusing negative passages on MS-PAS. “ST” and “TT” means student’s and teacher’s top ranking passages, respectively. “A & B” denotes the intersection of A and B. “A < B” represents the passages that have higher relevance scores in B than A.

6 INTRINSIC EVALUATION

We explore the intrinsic properties of PROD from several aspects.

6.1 Confusing Negative Selection in DPD

We first explore the impact of different confusing negative passages selection strategies in DPD. All experiments are based on the same student model after TPD, the results are shown in Table 8. The experiment results show that different selection strategies can improve the student model performance to a certain extent. Besides, all the strategies outperform TPD. Among them, “ST-(1,15] & TT-(0,1]” performs best in MRR@10, which is the strategy we finally use to select the confusing data.

6.2 Iteration Number of DPD

In this section, we discuss the impact of different iteration number N of DPD by gradually increasing it. As results shown in Table 9, when the iteration is at early stage, the improvement of DPD is obvious. Specifically, MRR@10 reaches the best 37.44 when $N = 2$ and Recall@50 reaches the best 87.09 when $N = 3$. However, when the iteration number is bigger than 3, the performance of the student

Setting	MRR@10	Recall@5	Recall@20	Recall@50
$N = 0$	38.75	55.86	77.68	86.56
$N = 1$	39.34	56.35	78.09	87.06
$N = 2$	39.43	56.39	77.95	86.89
$N = 3$	39.28	56.36	78.14	87.09
$N = 4$	39.17	56.38	78.11	86.93
$N = 5$	39.08	56.16	78.11	86.95

Table 9: The results of different iteration number N of DPD on MS-PAS.

Dataset	Method	MRR@10	Recall@50	Recall@1k
MS-PAS	PROD	39.34	87.06	98.44
	PROD w/o \mathcal{L}_r	39.20	86.93	98.31
Dataset	Method	Recall@5	Recall@20	Recall@100
NQ	PROD	75.61	84.72	89.56
	PROD w/o \mathcal{L}_r	75.30	84.52	89.45

Table 10: The impact of the regularization loss item \mathcal{L}_r in PROD on MS-PAS and NQ.

model decreases. We think the reason may be the long training steps of a small amount of data, which leads to the inevitable knowledge forgetting and overfitting. Therefore, considering the training steps and the average performance, we choose $N = 1$.

6.3 Regularization Loss Item \mathcal{L}_r

Last but not least, we explore the influence of the regularization loss item \mathcal{L}_r in distillation. Since the essence of the regularization loss item is to use the model in the previous step as a teacher and distill the current students, it is particularly effective in the situation where training data is scarce or the training is unstable and easy to overfit. In order to observe the effect of the regularization loss item more intuitively, we show the results on MS-PAS and NQ in Table 10. In the experiment, we compare between the distillation results with and without the regularization loss item. The results show that the regularization loss item not only stabilizes the model performances but also improves distillation effect, effectively alleviating the trend of overfitting in distillation training.

7 CONCLUSION

In this paper, we propose a novel distillation method PROD for dense retrieval. Concretely, we design *teacher progressive distillation* and *data progressive distillation* to gradually improve the performance of the student model. Extensive experiments on five widely-used benchmarks show that PROD can effectively improve the performance of the student model, achieving new state-of-art within the the distillation methods for dense retrieval, even surpassing some existing 12-layer models.

ACKNOWLEDGMENTS

Zhenghao Lin and Chen Lin are supported by National Natural Science Foundation of China (No.61972328). We also thank the reviewers for their insightful comments.

REFERENCES

- [1] Alex Andonian, Shixing Chen, and Raffay Hamid. 2022. Robust Cross-Modal Representation Learning with Progressive Self-Distillation. In *CVPR*. IEEE, 16409–16420.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [3] Pengfei Cao, Yubo Chen, Jun Zhao, and Taifeng Wang. 2020. Incremental Event Detection via Knowledge Consolidation Networks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.), 707–717. <https://doi.org/10.18653/v1/2020.emnlp-main.52>
- [4] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. 2021. Distilling Knowledge via Knowledge Review. In *CVPR*. 5008–5017.
- [5] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2020. Overview of the TREC 2019 deep learning track. *CoRR* abs/2003.07820 (2020). [arXiv:2003.07820](https://arxiv.org/abs/2003.07820) <https://arxiv.org/abs/2003.07820>
- [6] Zhuyun Dai and Jamie Callan. 2019. Deeper Text Understanding for IR with Contextual Neural Language Modeling. In *SIGIR*. 985–988.
- [7] 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* (2018).
- [8] Kavin Ethayarajah, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding Dataset Difficulty with V-Usable Information. In *Proceedings of the 39th International Conference on Machine Learning*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.), Vol. 162. 5988–6008. <https://proceedings.mlr.press/v162/ethayarajah22a.html>
- [9] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE v2: Sparse Lexical and Expansion Model for Information Retrieval. *CoRR* abs/2109.10086 (2021). [arXiv:2109.10086](https://arxiv.org/abs/2109.10086) <https://arxiv.org/abs/2109.10086>
- [10] Takashi Fukuda, Masayuki Suzuki, Gakuto Kurata, Samuel Thomas, Jia Cui, and Bhuvana Ramabhadran. 2017. Efficient Knowledge Distillation from an Ensemble of Teachers. In *INTERSPEECH*. ISCA, 3697–3701.
- [11] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. *CoRR* abs/1503.02531 (2015).
- [12] Sebastian Hofstätter, Sophia Althammer, Michael Schröder, Mete Sertkan, and Allan Hanbury. 2020. Improving Efficient Neural Ranking Models with Cross-Architecture Knowledge Distillation. *CoRR* abs/2010.02666 (2020). [arXiv:2010.02666](https://arxiv.org/abs/2010.02666) <https://arxiv.org/abs/2010.02666>
- [13] 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 *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.), 113–122. <https://doi.org/10.1145/3404835.3462891>
- [14] Shaoyi Huang, Dongkuan Xu, Ian En-Hsu Yen, Yijue Wang, Sung-En Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Sanguthevar Rajasekaran, Hang Liu, and Caiwen Ding. 2022. Sparse Progressive Distillation: Resolving Overfitting under Pretrain-and-Finetune Paradigm. In *ACL (1)*. Association for Computational Linguistics, 190–200.
- [15] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351* (2019).
- [16] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Transaction's on Big Data* 7, 3 (2021), 535–547. <https://doi.org/10.1109/TBDATA.2019.2921572>
- [17] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. 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*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.), 6769–6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- [18] 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*, Jimmy X. Huang, Yi Chang, Xueqi Cheng, Jaap Kamps, Vanessa Murdock, Ji-Rong Wen, and Yiqun Liu (Eds.), 39–48. <https://doi.org/10.1145/3397271.3401075>
- [19] Taehyeon Kim, Jaehoon Oh, Nakyl Kim, Sangwook Cho, and Se-Young Yun. 2021. Comparing Kullback-Leibler Divergence and Mean Squared Error Loss in Knowledge Distillation. In *IJCAI* 2628–2635.
- [20] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: a Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics* 7 (2019), 452–466. https://doi.org/10.1162/tacl_a_00276
- [21] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461* (2019).
- [22] Zhizhong Li and Derek Hoiem. 2018. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 12 (2018), 2935–2947. <https://doi.org/10.1109/TPAMI.2017.2773081>
- [23] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2020. Distilling Dense Representations for Ranking using Tightly-Coupled Teachers. *CoRR* abs/2010.11386 (2020). [arXiv:2010.11386](https://arxiv.org/abs/2010.11386) <https://arxiv.org/abs/2010.11386>
- [24] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-batch negatives for knowledge distillation with tightly-coupled teachers for dense retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP (ReplANLP-2021)*, 163–173.
- [25] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Proceedings of the 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net. <https://openreview.net/forum?id=Bkg6RiCqY7>
- [26] 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 (2022). <https://doi.org/10.48550/arXiv.2205.09153>
- [27] Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. Generation-Augmented Retrieval for Open-Domain Question Answering. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1–6, 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, 4089–4100. <https://doi.org/10.18653/v1/2021.acl-long.316>
- [28] Seyed-Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. 2020. Improved Knowledge Distillation via Teacher Assistant. In *AAAI*. 5191–5198.
- [29] 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 *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 (CEUR Workshop Proceedings, Vol. 1773)*, Tarek Richard Besold, Antoine Bordes, Artur S. d'Avila Garcez, and Greg Wayne (Eds.). http://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf
- [30] Rodrigo Frassetto Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019. Document Expansion by Query Prediction. *CoRR* abs/1904.08375 (2019).
- [31] Barlas Oguz, Kushal Lakhotia, Anshit Gupta, Patrick Lewis, Vladimir Karpukhin, Aleksandra Piktus, Xilun Chen, Sebastian Riedel, Scott Yih, Sonal Gupta, and Yashar Mehdad. 2022. Domain-matched Pre-training Tasks for Dense Retrieval. In *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10–15, 2022*, Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, 1524–1534. <https://doi.org/10.18653/v1/2022.findings-naacl.114>
- [32] Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. *arXiv preprint arXiv:2001.04063* (2020).
- [33] 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 *NAACL-HLT*. 5835–5847.
- [34] Ruiyang Ren, Shangwen Lv, Yingqi Qu, Jing Liu, Wayne Xin Zhao, Qiaoqiao She, Hua Wu, Haifeng Wang, and Ji-Rong Wen. 2021. PAIR: Leveraging Passage-Centric Similarity Relation for Improving Dense Passage Retrieval. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021 (Findings of ACL, Vol. ACL/IJCNLP 2021)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.), 2173–2183. <https://doi.org/10.18653/v1/2021.findings-acl.191>
- [35] 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. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.), 2825–2835. <https://doi.org/10.18653/v1/2021.emnlp-main.224>
- [36] Mehdi Rezagholizadeh, Aref Jafari, Puneeth Salad, Pranav Sharma, Ali Saheb Pasand, and Ali Ghodsi. 2021. Pro-KD: Progressive Distillation by Following the Footsteps of the Teacher. *CoRR* abs/2110.08532 (2021).
- [37] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. FitNets: Hints for Thin Deep Nets. In *ICLR*

- (Poster).
- [38] Devendra Singh Sachan, Mostofa Patwary, Mohammad Shoeybi, Neel Kant, Wei Ping, William L. Hamilton, and Bryan Catanzaro. 2021. End-to-End Training of Neural Retrievers for Open-Domain Question Answering. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, 6648–6662. <https://doi.org/10.18653/v1/2021.acl-long.519>
- [39] Tim Salimans and Jonathan Ho. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *ICLR*. OpenReview.net.
- [40] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [41] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz (Eds.). Association for Computational Linguistics, 3715–3734. <https://doi.org/10.18653/v1/2022.naacl-main.272>
- [42] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990* (2022).
- [43] Weiyue Su, Xuyi Chen, Shikun Feng, Jiaxiang Liu, Weixin Liu, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. ERNIE-Tiny : A Progressive Distillation Framework for Pretrained Transformer Compression. *CoRR* abs/2106.02241 (2021).
- [44] Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.), 9275–9293. <https://doi.org/10.18653/v1/2020.emnlp-main.746>
- [45] Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2022. Unified and Effective Ensemble Knowledge Distillation. *CoRR* abs/2204.00548 (2022).
- [46] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *Proceedings of the 9th International Conference on Learning Representations*. <https://openreview.net/forum?id=zeFrfgYzLn>
- [47] Peilin Yang, Hui Fang, and Jimmy Lin. 2017. Anserini: Enabling the Use of Lucene for Information Retrieval Research. In *SIGIR*. 1253–1256.
- [48] Sohee Yang and Minjoon Seo. 2020. Is Retriever Merely an Approximator of Reader? *CoRR* abs/2010.10999 (2020). [arXiv:2010.10999](https://arxiv.org/abs/2010.10999) <https://arxiv.org/abs/2010.10999>
- [49] Sergey Zagoruyko and Nikos Komodakis. 2017. Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR (Poster)*.
- [50] Hansi Zeng, Hamed Zamani, and Vishwa Vinay. 2022. Curriculum Learning for Dense Retrieval Distillation. In *SIGIR*. 1979–1983.
- [51] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. Optimizing Dense Retrieval Model Training with Hard Negatives. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.), 1503–1512. <https://doi.org/10.1145/3404835.3462880>
- [52] Hang Zhang, Yeyun Gong, Yelong Shen, Jiancheng Lv, Nan Duan, and Weizhu Chen. 2022. Adversarial Retriever-Ranker for Dense Text Retrieval. In *Proceedings of the 10th International Conference on Learning Representations*. <https://openreview.net/forum?id=MR7XubKUFb>
- [53] Kai Zhang, Chongyang Tao, Tao Shen, Can Xu, Xiubo Geng, Binxing Jiao, and Daxin Jiang. 2022. LED: Lexicon-Enlightened Dense Retriever for Large-Scale Retrieval. *CoRR* abs/2208.13661 (2022).
- [54] Tiancheng Zhao, Xiaopeng Lu, and Kyusong Lee. 2021. SPARTA: Efficient Open-Domain Question Answering via Sparse Transformer Matching Retrieval. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tammy Chakraborty, and Yichao Zhou (Eds.), 565–575. <https://doi.org/10.18653/v1/2021.naacl-main.47>
- [55] Chunting Zhou, Jiatao Gu, and Graham Neubig. 2020. Understanding Knowledge Distillation in Non-autoregressive Machine Translation. In *ICLR*.
- [56] Wangchunshu Zhou, Canwen Xu, and Julian J. McAuley. 2022. BERT Learns to Teach: Knowledge Distillation with Meta Learning. In *ACL (1)*. 7037–7049.

Dataset	Train	Dev	Test	#Doc
MS-PAS	502,939	6,980	-	8,841,823
TREC-PAS-19	-	-	200	8,841,823
MS-Doc	367,013	5,193	-	3,213,835
TREC-Doc-19	-	-	200	3,213,835
NQ	58,880	8,757	3,610	21,015,324
BING-REL	1,593,219	8,013	-	5,335,927
BING-ADS	8,306,968	53,219	52,590	2,866,527

Table 11: Statistics of the text retrieval datasets.

		Parameter	NQ	MS-PAS	MS-DOC	BING-REL	BING-ADS
Global		Max query len	32	32	32	64	32
		Max passage len	128	144	480	512	32
		Temperature	4.0	4.0	4.0	4.0	4.0
		Hard loss weight	0.1	0.1	0.1	0.1	0.1
		Soft loss weight	0.9	0.9	0.9	0.9	0.1
		Warmup	0.1	0.1	0.1	0.1	0.1
TPD	DE	Learning rate	5e-5	5e-5	5e-5	2e-5	5e-5
		Batch size	128	128	128	256	512
	CE	Distillation step	80000	40000	40000	16000	40000
		Negative num	1	1	1	2	1
DPD	DE	Learning rate	5e-5	5e-5	5e-5	2e-5	5e-5
		Batch size	64	64	64	64	256
	CE	Distillation step	20000	40000	40000	16000	16000
		Negative num	15	15	15	15	15
DPD	DE	Learning rate	1e-5	1e-5	1e-5	1e-5	1e-5
		Batch size	64	64	64	64	256
	CE	Distillation step	200	2000	2000	2000	2000
		Negative num	15	15	15	15	15

Table 12: Hyper-parameters for PROD.

A DATA STATISTICS

The statistics of used datasets are shown in Table 11.

B HYPER-PARAMETERS

The detailed hyper-parameters are shown in Table 12.