

# RetroMAE: Pre-Training Retrieval-oriented Language Models Via Masked Auto-Encoder

Shitao Xiao<sup>1†</sup>, Zheng Liu<sup>2†</sup>, Yingxia Shao<sup>1\*</sup>, Zhao Cao<sup>2</sup>

1: Beijing University of Posts and Telecommunications, Beijing, China

2: Huawei Technologies Ltd. Co., Shenzhen, China

{stxiao, shaoyx}@bupt.edu.cn, {liuzheng107, caozhao1}@huawei.com

## Abstract

Despite pre-training’s progress in many important NLP tasks, it remains to explore effective pre-training strategies for dense retrieval. In this paper, we propose **RetroMAE**, a new retrieval oriented pre-training paradigm based on Masked Auto-Encoder (MAE). RetroMAE is highlighted by three critical designs. 1) **A novel MAE workflow**, where the input sentence is polluted for encoder and decoder with different masks. The sentence embedding is generated from the encoder’s masked input; then, the original sentence is recovered based on the sentence embedding and the decoder’s masked input via masked language modeling. 2) **Asymmetric model structure**, with a full-scale BERT like transformer as encoder, and a one-layer transformer as decoder. 3) **Asymmetric masking ratios**, with a moderate ratio for encoder: 15~30%, and an aggressive ratio for decoder: 50~70%. Our framework is simple to realize and empirically competitive: the pre-trained models dramatically improve the SOTA performances on a wide range of dense retrieval benchmarks, like **BEIR** and **MS MARCO**. The source code and pre-trained models are made publicly available at <https://github.com/staoxiao/RetroMAE> so as to inspire more interesting research.

## 1 Introduction

Dense retrieval is important to many web applications. By letting semantically correlated query and document represented as spatially close embeddings, dense retrieval can be efficiently conducted via approximate nearest neighbour search, such as PQ (Jegou et al., 2010; Xiao et al., 2021, 2022a) and HNSW (Malkov and Yashunin, 2018). Recently, large-scale language models have been widely used as the encoding networks for dense retrieval (Karpukhin et al., 2020; Xiong et al.,

2020; Luan et al., 2021). The mainstream models, e.g., BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), T5 (Raffel et al., 2019), are usually pre-trained by token-level tasks, like MLM and Seq2Seq. However, the sentence-level representation capability is not fully developed in these tasks, which restricts their potential for dense retrieval.

Given the above defect, there have been increasing interests to develop retrieval oriented pre-trained models. One popular strategy is to leverage self-contrastive learning (Chang et al., 2020; Guu et al., 2020), where the model is trained to discriminate positive samples from data augmentation. However, the self-contrastive learning can be severely limited by the data augmentation’s quality; besides, it usually calls for massive amounts of negative samples (He et al., 2020a; Chen et al., 2020). Another strategy relies on auto-encoding (Gao and Callan, 2021; Lu et al., 2021; Wang et al., 2021), which is free from the restrictions on data augmentation and negative sampling. The current works are differentiated in how the encoding-decoding workflow is designed, and it remains an open problem to explore more effective auto-encoding framework for retrieval oriented pre-training.

We argue that two factors are critical for the auto-encoding based pre-training: 1) the reconstruction task must be demanding enough on encoding quality, 2) the pre-training data needs to be fully utilized. We propose RetroMAE (Figure 1), which optimizes both aspects with the following designs.

- **A novel MAE workflow.** The pre-training follows a novel masked auto-encoding workflow. The input sentence is polluted twice with two different masks. One masked input is used by encoder, where the sentence embedding is generated. The other one is used by decoder: joined with the sentence embedding, the original sentence is recovered via masked language modeling (MLM).

- **Asymmetric structure.** RetroMAE adopts an asymmetric model structure. The encoder is a

†. The two researchers make equal contributions to this work and are designated as co-first authors.

\*. Corresponding author.

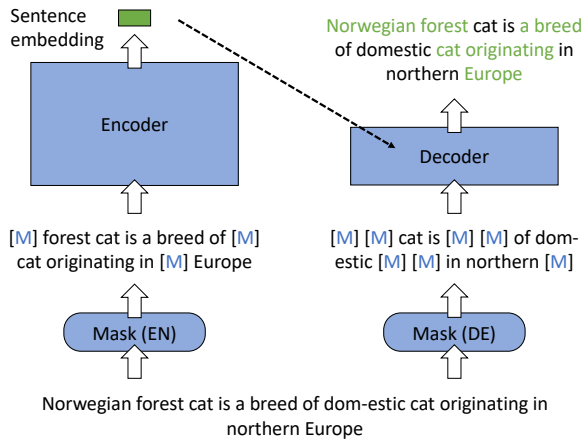


Figure 1: RetroMAE. The encoder utilizes a full-scale BERT, whose input is moderately masked. The decoder is a one-layer transformer, whose input is aggressively masked. The original input is recovered based on the sentence embedding and the decoder’s input via MLM.

full-scale BERT, which is able to generate discriminative embedding for the input sentence. In contrast, the decoder follows an extremely simplified structure, i.e., a single-layer transformer, which is learned to reconstruct the input sentence.

- **Asymmetric masking ratios.** The encoder’s input is masked at a moderate ratio: 15~30%, which is slightly above its traditional value in MLM. However, the decoder’s input is masked at a much more aggressive ratio: 50~70%.

The above designs of RetroMAE are favorable to the pre-training effectiveness thanks to the following reasons. Firstly, the auto-encoding is made **more demanding on encoding quality**. The conventional auto-regression may attend to the prefix during the decoding process; and the conventional MLM only masks a small portion (15%) of the input tokens. By comparison, RetroMAE aggressively masks most of the input for decoding. As such, the reconstruction will be not enough to leverage the decoder’s input alone, but heavily depend on the sentence embedding. Thus, it will force the encoder to capture in-depth semantics of the input. Secondly, it ensures **training signals to be fully generated** from the input sentence. For conventional MLM-style methods, the training signals may only be generated from 15% of the input tokens. Whereas for RetroMAE, the training signals can be derived from the majority of the input. Besides, knowing that the decoder only contains one-single layer, we further propose the **enhanced decoding** on top of two-stream attention (Yang et al., 2019) and position-specific attention mask (Dong et al., 2019). As such, 100% of the tokens can

be used for reconstruction, and each token may sample a unique context for its reconstruction.

RetroMAE is simple to realize and empirically competitive. We merely use a moderate-amount of data (Wikipedia, BookCorpus, MS MARCO corpus) for pre-training, where a BERT base scale encoder is learned. For the zero-shot setting, it produces an average score of **45.2 on BEIR** (Thakur et al., 2021); and for the supervised setting, it may easily reach an MRR@10 of **41.6 on MS MARCO** passage retrieval (Nguyen et al., 2016) following standard knowledge distillation procedures. Both values are unprecedented for dense retrievers with the same model size and pre-training conditions. We also carefully evaluate the impact introduced from each of the components, whose results may bring interesting insights to the future research.

## 2 Related works

Dense retrieval is widely applied to web applications, like search engines (Karpukhin et al., 2020), advertising (Lu et al., 2020; Zhang et al., 2022) and recommender systems (Xiao et al., 2022b). It encodes query and document within the same latent space, where relevant documents to the query can be efficiently retrieved via ANN search. The encoding model is critical for the retrieval quality. Thanks to the recent development of large-scale language models, e.g., BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and T5 (Raffel et al., 2019), there has been a major leap-forward for dense retrieval’s performance (Karpukhin et al., 2020; Luan et al., 2021; Lin et al., 2021).

The large-scale language models are highly differentiated in terms of pre-training tasks. One common task is the masked language modeling (MLM), as adopted by BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), in which the masked tokens are predicted based on their context. The basic MLM is extended in many ways. For example, tasks like entity masking, phrase masking and span masking (Sun et al., 2019; Joshi et al., 2020) may help the pre-trained models to better support the sequence labeling applications, such as entity resolution and question answering. Besides, tasks like auto-regression (Radford et al., 2018; Yang et al., 2019) and Seq2Seq (Raffel et al., 2019; Lewis et al., 2019) are also utilized, where the pre-trained models are enabled to serve NLG related scenarios. However, most of the generic pre-trained models are based on token-level tasks,

where the sentence representation capability is not effectively developed (Chang et al., 2020). Thus, it may call for a great deal of labeled data (Nguyen et al., 2016; Kwiatkowski et al., 2019) and sophisticated fine-tuning methods (Xiong et al., 2020; Qu et al., 2020) to ensure the pre-trained models’ performance for dense retrieval.

To mitigate the above problem, recent works propose retrieval oriented pre-trained models. The existing methods can be divided as the ones based on self-contrastive learning (SCL) and the ones based on auto-encoding (AE). The SCL based methods (Chang et al., 2020; Guu et al., 2020; Xu et al., 2022) rely on data augmentation, e.g., inverse cloze task (ICT), where positive samples are generated for each anchor sentence. Then, the language model is learned to discriminate the positive samples from the negative ones via contrastive learning. However, the self-contrastive learning usually calls for huge amounts of negative samples, which is computationally expensive. Besides, the pre-training effect can be severely restricted by the quality of data augmentation. The AE based methods are free from these restrictions, where the language models are learned to reconstruct the input sentence based on the sentence embedding. The existing methods utilize various reconstruction tasks, such as MLM (Gao and Callan, 2021) and auto-regression (Lu et al., 2021; Wang et al., 2021; Li et al., 2020), which are highly differentiated in terms of how the original sentence is recovered and how the training loss is formulated. For example, the auto-regression relies on the sentence embedding and prefix for reconstruction; while MLM utilizes the sentence embedding and masked context. The auto-regression derives its training loss from the entire input tokens; however, the conventional MLM only learns from the masked positions, which accounts for 15% of the input tokens. Ideally, we expect the decoding operation to be demanding enough, as it will force the encoder to fully capture the semantics about the input so as to ensure the reconstruction quality. Besides, we also look forward to high data efficiency, which means the input data can be fully utilized for the pre-training task.

### 3 Methodology

We develop a novel masked auto-encoder for retrieval oriented pre-training. The model contains two modules: a BERT-like encoder  $\Phi_{enc}(\cdot)$  to generate sentence embedding, and a one-layer trans-

former based decoder  $\Phi_{dec}(\cdot)$  for sentence reconstruction. The original sentence  $X$  is masked as  $\tilde{X}_{enc}$  and encoded as the sentence embedding  $\mathbf{h}_{\tilde{X}}$ . The sentence is masked again (with a different mask) as  $\tilde{X}_{dec}$ ; together with  $\mathbf{h}_{\tilde{X}}$ , the original sentence  $X$  is reconstructed. Detailed elaborations about RetroMAE are made as follows.

#### 3.1 Encoding

The input sentence  $X$  is polluted as  $\tilde{X}_{enc}$  for the encoding stage, where a small fraction of its tokens are randomly replaced by the special token [M] (Figure 2. A). We apply a moderate masking ratio (15~30%), which means the majority of information about the input will be preserved. Then, the encoder  $\Phi^{enc}(\cdot)$  is used to transform the polluted input as the sentence embedding  $\mathbf{h}_{\tilde{X}}$ :

$$\mathbf{h}_{\tilde{X}} \leftarrow \Phi_{enc}(\tilde{X}_{enc}). \quad (1)$$

We apply a BERT like encoder with 12 layers and 768 hidden-dimensions, which helps to capture the in-depth semantics of the sentence. Following the common practice, we select the [CLS] token’s final hidden state as the sentence embedding.

#### 3.2 Decoding

The input sentence  $X$  is polluted as  $\tilde{X}_{dec}$  for the decoding stage (Figure 2. B). The masking ratio is more aggressive than the one used by the encoder, where 50~70% of the input tokens will be masked. The masked input is joined with the sentence embedding, based on which the original sentence is reconstructed by the decoder. Particularly, the sentence embedding and the masked input are combined into the following sequence:

$$\mathbf{H}_{\tilde{X}_{dec}} \leftarrow [\mathbf{h}_{\tilde{X}}, \mathbf{e}_{x_1} + \mathbf{p}_1, \dots, \mathbf{e}_{x_N} + \mathbf{p}_N]. \quad (2)$$

In the above equation,  $\mathbf{e}_{x_i}$  denotes the embedding of  $x_i$ , to which an extra position embedding  $\mathbf{p}_i$  is added. Finally, the decoder  $\Phi_{dec}$  is learned to reconstruct the original sentence  $X$  by optimizing the following objective:

$$\mathcal{L}_{dec} = \sum_{x_i \in \text{masked}} \text{CE}(x_i | \Phi_{dec}(\mathbf{H}_{\tilde{X}_{dec}})), \quad (3)$$

where CE is the cross-entropy loss. As mentioned, we use a one-layer transformer based decoder. Given the aggressively masked input and the extremely simplified network, the decoding becomes challenging, which forces the generation of high-quality sentence embedding so that the original input can be recovered with good fidelity.

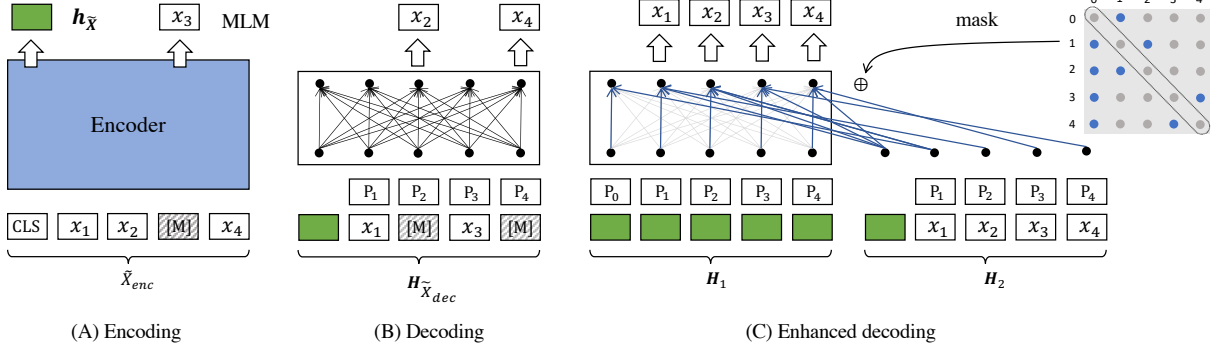


Figure 2: RetroMAE pre-training workflow. (A) Encoding: the input is moderately masked and encoded as the sentence embedding (the green rectangle). (B) Decoding: the input is aggressively masked, and joined with the sentence embedding to reconstruct the masked tokens (the shadowed tokens). (C) Enhanced encoding: all input tokens are reconstructed based on the sentence embedding and the visible context in each row (defined in Eq. 7); the main diagonal positions are filled with  $-\infty$  (grey), and positions for the visible context are filled with 0 (blue).

### 3.3 Enhanced Decoding

One limitation about the decoding process is that the training signals, i.e., the cross-entropy loss, can only be derived from the masked tokens. Besides, every masked token is always reconstructed based on the same context, i.e.,  $\mathbf{H}_{\tilde{X}_{dec}}$ . We argue that the pre-training effect can be further enhanced providing that 1) **more training signals** can be derived from the input sentence, and 2) the reconstruction task can be performed based on **diversified contexts**. To this end, we propose the **enhanced decoding** inspired by two-stream self-attention (Yang et al., 2019) and position-specific attention mask (Dong et al., 2019). Particularly, we generate two input streams:  $\mathbf{H}_1$  (query) and  $\mathbf{H}_2$  (context), for the decoding operation (Figure 2. C):

$$\begin{aligned} \mathbf{H}_1 &\leftarrow [\mathbf{h}_{\tilde{X}} + \mathbf{p}_0, \dots, \mathbf{h}_{\tilde{X}} + \mathbf{p}_N], \\ \mathbf{H}_2 &\leftarrow [\mathbf{h}_{\tilde{X}}, \mathbf{e}_{x_1} + \mathbf{p}_1, \dots, \mathbf{e}_{x_N} + \mathbf{p}_N]. \end{aligned} \quad (4)$$

where  $\mathbf{h}_{\tilde{X}}$  is the sentence embedding,  $\mathbf{e}_{x_i}$  is the token embedding (no token is masked in this place), and  $\mathbf{p}_i$  is the position embedding. We introduce the position-specific attention mask  $\mathbf{M} \in \mathbb{R}^{L \times L}$ , where the self-attention is performed as:

$$\begin{aligned} \mathbf{Q} &= \mathbf{H}_1 \mathbf{W}^Q, \mathbf{K} = \mathbf{H}_2 \mathbf{W}^K, \mathbf{V} = \mathbf{H}_2 \mathbf{W}^V; \\ \mathbf{M}_{ij} &= \begin{cases} 0, & \text{can be attended,} \\ -\infty, & \text{masked;} \end{cases} \\ \mathbf{A} &= \text{softmax}\left(\frac{\mathbf{Q}^T \mathbf{K}}{\sqrt{d}} + \mathbf{M}\right) \mathbf{V}. \end{aligned} \quad (5)$$

The output  $\mathbf{A}$ , together with  $\mathbf{H}_1$  (because of the residual connection) are used to reconstruct the original input (other operations, like layer-norm

and FFN, are omitted from our discussion). Finally, the following objective will be optimized:

$$\mathcal{L}_{dec} = \sum_{x_i \in X} \text{CE}(x_i | \mathbf{A}, \mathbf{H}_1). \quad (6)$$

Knowing that the decoder only consists of one single transformer layer, each token  $x_i$  is reconstructed based on the context which are visible to the  $i$ -th row of matrix  $\mathbf{M}$ . In this place, the following rules are applied to generate the position specific attention mask  $\mathbf{M}$ :

$$\mathbf{M}_{ij} = \begin{cases} 0, & x_j \in s(X_{\neq i}), \text{ or } j_{i \neq 0} = 0 \\ -\infty, & \text{otherwise.} \end{cases} \quad (7)$$

The sampled tokens,  $s(X_{\neq i})$ , and the 1st position (except for the 1st row) will be visible when reconstructing  $x_i$ . The diagonal elements, i.e.,  $x_i$  for the  $i$ -th row, will always be excluded, which means they will always be masked; as a result, each token cannot attend to itself during the reconstruction.

We summarize the pre-training workflow with the enhanced decoding as Algorithm 1. Note that the original masked language modeling task in BERT is kept in encoder. The corresponding loss, denoted as  $\mathcal{L}_{enc}$ , is added with the decoder's loss, which formulates the final training loss. The following features need to be emphasized for our pre-training workflow. Firstly, the reconstruction task is **demanding** given the aggressive masking ratio and the extremely simplified network of decoder. Secondly, we may derive **abundant** pre-training signals from the unsupervised corpus since all tokens within each input sentence can be used for the reconstruction. Finally, the pre-training is **simple**



---

**Algorithm 1: RetroMAE**

---

```
1 begin
2    $\tilde{X}_{enc} \leftarrow \text{mask}(X)$ ;
3    $\mathbf{h}_{\tilde{X}} \leftarrow \Phi_{enc}(\tilde{X}_{enc})$ ;
4    $\mathbf{H}_1, \mathbf{H}_2 \leftarrow \text{Eq. 4}$ ;
5    $\mathbf{M} \leftarrow \text{Eq. 7}$ ;
6    $\mathbf{A} \leftarrow \text{based on } \mathbf{H}_1, \mathbf{H}_2, \mathbf{M} \text{ as Eq. 5}$ ;
7    $\mathcal{L}_{dec} \leftarrow \text{Eq. 6}$ ;
8   model update w.r.t.  $\mathcal{L}_{enc} + \mathcal{L}_{dec}$ ;
```

---

to realize: 1) there are no requirements on sophisticated data augmentation and negative sampling, and 2) the computation cost is maintained to be similar with the conventional BERT/RobERTa style pre-training given the simplicity of decoder.

## 4 Experimental Studies

We evaluate the retrieval performance of the sentence embedding generated by RetroMAE’s pre-trained encoder, where two major issues are explored. 1) RetroMAE’s impact on zero-shot and supervised dense retrieval, in comparison with the generic pre-trained models and the retrieval oriented pre-trained models. 2) The impact from the four technical factors in RetroMAE: the enhanced decoding, the decoder size, the decoder’s masking ratio, and the encoder’s masking ratio.

### 4.1 Experiment Settings

The following **datasets** are utilized for the pre-training and evaluation of RetroMAE.

- **Pre-training.** We reuse the same pre-training corpora as the ones utilized by BERT (Devlin et al., 2019): **English Wikipedia** and **BookCorpus**. Corresponding datasets are also frequently used by the previous works on retrieval oriented pre-training, such as SEED (Lu et al., 2021) and Condenser (Gao and Callan, 2021). Following coCondenser (Gao and Callan, 2022), we also use **MS MARCO** corpus to analyze the effect of in-domain pre-training for dense retrieval (which we find critical to the performance on MS MARCO but unnecessary to the performances on other datasets).

- **Evaluation.** We use two datasets to evaluate the retrieval performance after supervision. 1) **MS MARCO** (Nguyen et al., 2016), which contains queries from Bing Search. We use its passage retrieval task, which contains 502,939 training queries and 6,980 evaluation queries (Dev).

The answer needs to be retrieved from 8.8 million candidate passages. 2) **Natural Questions** (Kwiatkowski et al., 2019), which consists of queries from Google. There are 79,168 training queries, 8,757 dev queries, and 3,610 testing queries. The answer is retrieved from 21,015,324 Wikipedia passages. We evaluate the zero-shot retrieval performance on **BEIR benchmark** (Thakur et al., 2021). It fine-tunes the pre-trained model with MS MARCO queries, and evaluate the zero-shot transferability on the other 18 datasets. The evaluation data covers dense retrieval tasks across different domains, such as question answering, fact checking, bio-medical retrieval, news retrieval, and social media retrieval, etc.

We consider three types of **baseline methods** in our experimental studies<sup>1</sup>.

- **Generic models.** The following generic pre-trained language models are used in our experiments. 1) **BERT** (Devlin et al., 2019), which is the most popular pre-trained language model in practice. It is also frequently fine-tuned as the encoding network for dense retrievers (Karpukhin et al., 2020; Xiong et al., 2020). 2) **RobERTa** (Liu et al., 2019), which is an enhanced replication of BERT with substantially enlarged training data and improved training settings. 3) **DeBERTa** (He et al., 2020b), which further improves BERT and RobERTa with disentangled attention mechanism and an enhanced mask decoder; it is one of the strongest pre-trained models on NLU benchmarks, such as GLUE and SuperGLUE.

- **Retrieval oriented models.** We consider two types of retrieval oriented pre-trained models in our experiments. One is based on self-contrastive learning, where the following methods are included. 1) **SimCSE** (Gao et al., 2021), in which the language model is learned to discriminate different views of the anchor sentence augmented by drop-out. 2) **LaPraDoR** (Xu et al., 2022), an enhancement of conventional ICT pre-training (Guu et al., 2020; Chang et al., 2020); it proposes to train the query and document encoder iteratively so that the scale of negative samples can be expanded. 3) **DiffCSE** (Chuang et al., 2022), which enhances SimCSE with the jointly utilization of self-contrastive learning and conditional difference prediction. The other one is based on auto-encoding, in which the following methods are included. 4) **Condenser** (Gao and

1. For all baseline methods, we use their officially released pre-trained models for our experiments.

Datasets	BERT	RoBERTa	DeBERTa	LaPraDoR	SimCSE	DiffCSE	SEED	Condenser	RetroMAE
TREC-COVID	0.615	0.649	0.688	0.478	0.460	0.492	0.627	0.750	<b>0.772</b>
BioASQ	0.253	0.279	0.290	0.252	0.263	0.258	0.308	0.322	<b>0.421</b>
NFCorpus	0.260	0.243	0.238	0.310	0.260	0.259	0.278	0.277	<b>0.308</b>
NQ	0.467	0.413	0.452	0.454	0.435	0.412	0.446	0.486	<b>0.518</b>
HotpotQA	0.488	0.448	0.474	0.513	0.502	0.499	0.541	0.538	<b>0.635</b>
FiQA-2018	0.252	0.291	0.299	0.288	0.250	0.229	0.259	0.259	<b>0.316</b>
Signal-1M(RT)	0.204	0.229	0.243	0.241	0.262	0.260	0.256	0.261	<b>0.265</b>
TREC-NEWS	0.362	0.385	0.378	0.286	0.356	0.363	0.358	0.376	<b>0.428</b>
Robust04	0.351	0.384	0.378	0.299	0.330	0.343	0.365	0.349	<b>0.447</b>
ArguAna	0.265	0.395	0.297	<b>0.499</b>	0.413	0.468	0.389	0.298	0.433
Touche-2020	0.259	<b>0.299</b>	0.271	0.137	0.159	0.168	0.225	0.248	0.237
CQADupStack	0.282	0.278	0.279	0.309	0.290	0.305	0.290	<b>0.347</b>	0.317
Quora	0.787	0.509	0.846	0.837	0.844	0.850	<b>0.852</b>	0.853	0.847
DBPedia	0.314	0.275	0.271	0.334	0.314	0.303	0.330	0.339	<b>0.390</b>
SCIDOCS	0.113	0.111	0.106	0.150	0.124	0.125	0.124	0.133	<b>0.150</b>
FEVER	0.682	0.683	0.594	0.511	0.623	0.641	0.641	0.691	<b>0.774</b>
Climate-FEVER	0.187	0.222	0.160	0.173	0.211	0.200	0.176	0.211	<b>0.232</b>
SciFact	0.533	0.539	0.543	0.531	0.554	0.523	0.575	0.593	<b>0.653</b>
AVERAGE	0.371	0.368	0.378	0.367	0.369	0.372	0.391	0.407	<b>0.452</b>

Table 1: Zero-shot dense retrieval performances on BEIR benchmark (measured by NDCG@10).

Callan, 2021), where the sentence embedding is joined with the intermediate hidden-states from encoder for MLM. 5) **SEED** (Lu et al., 2021), where the sentence embedding is used to recover the original input via auto-regression.

• **Implementation details.** RetroMAE utilizes bi-directional transformers as its encoder, with 12 layers, 768 hidden-dim, and a 30522-token vocabulary (same as BERT base). The decoder is a one-layer transformer. The default masking ratios are 0.3 for encoder and 0.5 for decoder. The model is trained for 8 epochs, with AdamW optimizer, batch-size 32 (per device), learning rate 1e-4. The training is on a machine with 8× Nvidia A100 (40GB) GPUs. The models are implemented with PyTorch 1.8 and HuggingFace transformers 4.16. We adopt the official script<sup>2</sup> from BEIR to prepare the models for their zero-shot evaluation. For supervised evaluations, we use **DPR** (Karpukhin et al., 2020) and **ANCE** (Xiong et al., 2020) to fine-tune the pre-trained models. We also evaluate the models’ performance on MS MARCO with standard knowledge distillation. Particularly, we train one BERT base scale cross-encoder over hard negatives returned by the ANCE-finetuned bi-encoder; then, we further finetune the bi-encoder by minimizing the KL-divergence with the cross-encoder.

2. [https://github.com/beir-cellar/beir/blob/main/examples/retrieval/training/train\\_msmarco\\_v3.py](https://github.com/beir-cellar/beir/blob/main/examples/retrieval/training/train_msmarco_v3.py)

## 4.2 Main Results

We analyze the zero-shot performances in Table 1, where RetroMAE achieves remarkable advantages: it produces the best empirical results on most of the datasets, and it surpasses the strongest baseline by **+4.5%** in total average. The leap-forward of zero-shot performance is much larger than any of the improvements made by the baseline models. The improvement is indeed thrilling knowing that it is not from the increasing of model scale or the enrichment of pre-training data, but purely from the upgrade of pre-training algorithm. We further demonstrate the supervised evaluations in Table 2 and 3, where the pre-trained models are fine-tuned with DPR and ANCE. The baselines are partitioned into two groups according to whether they are generic pre-trained models or retrieval oriented ones. It can be observed that RetroMAE maintains notable advantages over the baselines. With DPR fine-tuning, it outperforms the strongest baselines by +1.96% (MRR@10) and +1.48% (Recall@10) on the two datasets; with ANCE fine-tuning, corresponding advantages become +1.42% (MRR@10) and +1.41% (Recall@10).

As for RetroMAE’s **pre-trained models over MS MARCO** corpus. With ANCE fine-tuning (Table 4), our method outperforms its peer approach coCondenser (Gao and Callan, 2022) by **+1.1%** on MRR@10 (which is of the same model size and same pre-training data). While under the knowl-

Methods	MS MARCO					Natural Question				
	MRR@10	MRR@100	R@10	R@100	R@1000	R@10	R@20	R@30	R@50	R@100
BERT	0.3170	0.3278	0.5801	0.8570	0.9598	0.7399	0.7925	0.8136	0.8396	0.8668
RoBERTa	0.3136	0.3258	0.5638	0.8478	0.9579	0.7150	0.7676	0.7939	0.8211	0.8476
DeBERTa	0.3186	0.3304	0.5824	0.8625	0.9654	0.7152	0.7778	0.8022	0.8269	0.8510
SimCSE	0.3193	0.3307	0.5907	0.8653	0.9699	0.7291	0.7864	0.8125	0.8391	0.8670
LaPraDoR	0.3191	0.3307	0.5833	0.8537	0.9602	0.7377	0.7920	0.8155	0.8399	0.8677
DiffCSE	0.3202	0.3311	0.5832	0.8561	0.9607	0.7393	0.7934	0.8155	0.8407	0.8673
SEED	0.3264	0.3374	0.5913	0.8535	0.9539	0.7454	0.7958	0.8208	0.8432	0.8701
Condenser	0.3357	0.3471	0.6082	0.8770	0.9683	0.7562	0.8053	0.8269	0.8501	0.8711
RetroMAE	<b>0.3553</b>	<b>0.3665</b>	<b>0.6356</b>	<b>0.8922</b>	<b>0.9763</b>	<b>0.7704</b>	<b>0.8172</b>	<b>0.8399</b>	<b>0.8604</b>	<b>0.8812</b>

Table 2: Supervised evaluation results based on DPR fine-tuning.

Methods	MS MARCO					Natural Question				
	MRR@10	MRR@100	R@10	R@100	R@1000	R@10	R@20	R@30	R@50	R@100
BERT	0.3460	0.3569	0.6220	0.8734	0.9642	0.7875	0.8227	0.8460	0.8601	0.8776
RoBERTa	0.3433	0.3543	0.6130	0.8705	0.9637	0.7629	0.8053	0.8277	0.8449	0.8698
DeBERTa	0.3396	0.3512	0.6016	0.8719	0.9670	0.7654	0.8097	0.8288	0.8479	0.8698
SimCSE	0.3520	0.3623	0.6276	0.8849	0.9738	0.7742	0.8194	0.8418	0.8626	0.8864
LaPraDoR	0.3456	0.3564	0.6129	0.8755	0.9640	0.7801	0.8247	0.8424	0.8590	0.8773
DiffCSE	0.3462	0.3571	0.6217	0.8748	0.9654	0.7853	0.8252	0.8410	0.8620	0.8784
SEED	0.3544	0.3653	0.6263	0.8812	0.9687	0.7803	0.8258	0.8449	0.8684	0.8870
Condenser	0.3635	0.3742	0.6388	0.8912	0.9722	0.7903	0.8325	0.8524	0.8668	0.8834
RetroMAE	<b>0.3822</b>	<b>0.3928</b>	<b>0.6677</b>	<b>0.9074</b>	<b>0.9807</b>	<b>0.8044</b>	<b>0.8443</b>	<b>0.8632</b>	<b>0.8776</b>	<b>0.8942</b>

Table 3: Supervised evaluation results based on ANCE fine-tuning.

Methods	MS MARCO			
	MRR@10	R@10	R@100	R@1000
coCondenser	0.382	-	-	0.984
RetroMAE	<b>0.393</b>	<b>0.675</b>	<b>0.918</b>	<b>0.985</b>

Table 4: RetroMAE vs. coCondenser on MS MARCO. Both models are fine-tuned by ANCE.

Methods	MS MARCO			
	MRR@10	R@10	R@100	R@1000
AR2	0.395	-	-	0.986
ColBERTv2	0.397	-	-	0.984
RocketQAv2	0.388	-	-	0.981
ERNIE-Search	0.401	-	-	0.982
RetroMAE	<b>0.416</b>	<b>0.709</b>	<b>0.927</b>	<b>0.988</b>

Table 5: RetroMAE vs. the recent dense retrievers; all models are fine-tuned by knowledge distillation.

edge distillation scenario (Table 5), RetroMAE surpasses a series of strong baselines proposed in the recent period, including the models with highly sophisticated distillation methods: AR2 (Zhang et al., 2021) by **+2.1%**, RocketQAv2 (Ren et al., 2021) by **+2.8%**, ERNIE-search (Lu et al., 2022) by **+1.5%**, and the late-interaction model ColBERTv2 (Santhanam et al., 2021) by **+1.9%**.

To summarize, the above results verify RetroMAE’s effectiveness from two aspects. 1) It substantially improves the pre-trained model’s **transferability**, which helps to result in superior zero-shot performances on out-of-domain datasets. 2) It provides a **strong initialization of dense retriever**; after fine-tuned with in-domain data, it gives rise to a high-quality supervised retrieval performance in the corresponding scenario. Besides the primary results, we may also have the following interesting observations.

Firstly, the advanced pre-training methods in generic areas do not necessarily contribute to the dense retrieval performances. Particularly, both RoBERTa and DeBERTa are major improvements of BERT; however, none of them is able to achieve better performances than BERT as they did on general NLU benchmarks. This observation further supports the motivation to develop retrieval oriented pre-trained models.

Secondly, the auto-encoding style pre-training (adopted by SEED, Condenser, and RetroMAE) is empirically proved to be much more favorable for dense retrieval, given its dominance over the generic and self-contrastive learning based pre-

Factor	Setting	MS MARCO				Natural Questions			
		MRR@10	MRR@100	R@100	R@1000	MRR@10	MRR@100	R@100	R@1000
Decoding	w. enhance	<b>0.3553</b>	<b>0.6356</b>	<b>0.8922</b>	<b>0.9763</b>	<b>0.7704</b>	<b>0.8399</b>	<b>0.8604</b>	<b>0.8812</b>
	w.o. enhance	0.3462	0.6218	0.8813	0.9725	0.7562	0.8291	0.8540	0.8759
Decoder size (w.o.)	#layer = 1	0.3462	0.6218	0.8813	0.9725	0.7562	0.8291	0.8540	0.8759
	#layer = 2	0.3446	0.6217	0.8828	0.9729	0.7561	0.8289	0.8538	0.8759
	#layer = 3	0.3439	0.6223	0.8829	0.9730	0.7563	0.8290	0.8537	0.8760
Mask ratio (decoder)	0.15 (w.)	0.3496	0.6297	0.8905	0.9734	0.7608	0.8309	0.8554	0.8750
	0.5 (w.)	<b>0.3553</b>	<b>0.6356</b>	<b>0.8922</b>	<b>0.9763</b>	<b>0.7704</b>	<b>0.8399</b>	<b>0.8604</b>	<b>0.8812</b>
	0.9 (w.)	0.3514	0.6285	0.8905	0.9740	0.7609	0.8343	0.8562	0.8756
	0.15 (w.o.)	0.3440	0.6177	0.8802	0.9700	0.7519	0.8253	0.8523	0.8758
	0.7 (w.o.)	<b>0.3508</b>	<b>0.6262</b>	<b>0.8850</b>	<b>0.9738</b>	<b>0.7593</b>	<b>0.8327</b>	<b>0.8551</b>	<b>0.8760</b>
	0.9 (w.o.)	0.3441	0.6198	0.8803	0.9725	0.7576	0.8307	<b>0.8551</b>	0.8745
	Mask ratio (encoder)	0.15 (w.)	0.3501	0.6306	0.8890	0.9757	0.7703	<b>0.8404</b>	<b>0.8604</b>
0.3 (w.)	<b>0.3553</b>	<b>0.6356</b>	<b>0.8922</b>	<b>0.9763</b>	<b>0.7704</b>	0.8399	<b>0.8604</b>	<b>0.8812</b>	
0.9 (w.)	0.3365	0.6143	0.8750	0.9701	0.7599	0.8296	0.8508	0.8692	

Table 6: Ablation studies. (“w.”/“w.o.” indicates “with”/“without” using the enhanced decoding.)

trained models in both zero-shot evaluation and supervised evaluation.

Thirdly, the self-contrastive learning based pre-trained models bring very little improvements over the generic ones when fine-tuning is made available, as their performances are close to each other in both supervised and zero-shot evaluations. In fact, there is no supervise about this observation considering the similar results reported by recent studies on dense retrieval (Gao and Callan, 2021) (BERT against ICT) and image processing (El-Nouby et al., 2021) (BEiT (Bao et al., 2021) against MoCo/DINO). That is the pre-trained models from self-contrastive learning tend to have comparatively limited potential for fine-tuning. Therefore, we argue that the retrieval-oriented pre-training algorithm should be properly designed according to the condition of fine-tuning in specific scenarios.

### 4.3 Ablation Studies

We ablate RetroMAE with its supervised performance (DPR fine-tuning) in Table 6, where the following factors are analyzed: 1) decoding method, 2) decoder’s size, 3) decoder’s masking ratio, 4) encoder’s masking ratio. We may have the following observations from our results.

Firstly of all, we analyze the impact from the **decoding method**, i.e., whether the enhanced decoding is used or not. It can be observed that the enhanced decoding (w.) notably outperforms the basic decoding (w.o.). Such an empirical advantage can be attributed to the improved data efficiency of the enhanced decoding. Under the default masking ratio, the basic decoding merely samples 50% of

the input tokens for reconstruction, all of which are predicted based on the same context. With the enhanced decoding (Section 3.3), all of the input tokens can be utilized for reconstruction, and each of the tokens is reconstructed based on a unique context sampled as Eq. 7. As such, the enhanced decoding may obtain more sufficient and diversified training signals from the input data.

Secondly, we analyze the impact from **decoder size**, with the number of transformer layers increased from 1 to 3. Knowing that the enhanced decoding is only applicable for single-layer transformers, it is disabled for these experiments. Despite higher computation costs, the enlarged decoders do not bring any empirical gains. Besides, considering that the enhanced decoding has to be disabled for large decoders (#layer>1), which will severely harm the retrieval performances, the one-layer decoder is proved to be the best option.

Thirdly, we make an analysis for different **masking ratios of decoder**, whose value is increased from 0.15 to 0.9. We introduce two sets of experiments: one with the enhanced decoding (w.), and the other one without using enhanced decoding (w.o.). For both experiments, we observe substantial improvements of retrieval quality resulted from the aggressive masking ratios. For enhanced decoding (w.), the optimal performance is achieved at 0.5; without using enhanced decoding (w.o.), the optimal performance is reached at 0.7. This minor difference is probably because: unlike “w.” where all input tokens can be reconstructed, “w.o.” only reconstructs the masked tokens; thus, it may trade a larger ratio for the increasing of training signals.



Lastly, we study the **encoder’s masking ratio**. It is quite interesting that a slightly improved masking ratio of 0.3 also improves the empirical performances, compared with the commonly used value 0.15. For both encoder and decoder, the increased reconstruction difficulty can be the common reason why the increased masking ratios benefit the retrieval quality. However, different from decoder, a too aggressive ratio of encoder, e.g., 0.9, will severely harm the retrieval performance. This is because a too large masking ratio will prevent the generation of high-quality sentence embedding, considering that most of the useful information of the input sentence will be discarded.

The ablation studies are **concluded as follows**: 1) RetroMAE’s performance can be notably improved by the enhanced decoding; 2) the one-layer transformer is the best for decoder; 3) the retrieval quality can be improved from an aggressive masking ratio of decoder, and a moderately improved masking ratio of encoder.

## 5 Conclusion

We propose RetroMAE, a novel masked auto-encoding framework to pre-train retrieval oriented language models: the input sentence is randomly masked for encoder and decoder, and the sentence embedding is joined with the decoder’s masked input to reconstruct the original input. We introduce an asymmetric model structure (full-scale encoder and single-layer decoder) and asymmetric masking ratios (a moderate ratio for encoder and an aggressive one for decoder), which makes the reconstruction sufficiently demanding. We also introduce the enhanced decoding, which makes the full utilization of the pre-training data. Our experiments on BEIR, MS MARCO, and Natural Question validate RetroMAE’s effectiveness, as significant improvements on both zero-shot and supervised evaluations can be achieved over the existing methods.

## 6 Limitations

So far, our empirical studies are performed based on BERT base scale transformers; meanwhile, merely a moderate amount of pre-training data is used (mainly due to the limitations on computation resources). Despite the demonstrated effectiveness, it’s still necessary to explore the impact from enlarged networks and increased pre-training data, as both factors were found to be important in recent works (Ni et al., 2021).

## Acknowledgements

Yingxia Shao’s work is supported by the National Natural Science Foundation of China (Nos. U1936104, 61902037, 62272054, 62192784), and CCF-Tencent Open Fund.

## References

- Hangbo Bao, Li Dong, and Furu Wei. 2021. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*.
- 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*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljagic, Shang-Wen Li, Wen-tau Yih, Yoon Kim, and James R. Glass. 2022. *Diffcse: Difference-based contrastive learning for sentence embeddings*. *CoRR*, abs/2204.10298.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in Neural Information Processing Systems*, 32.
- Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. 2021. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*.
- Luyu Gao and Jamie Callan. 2021. Condenser: a pre-training architecture for dense retrieval. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 981–993.
- Luyu Gao and Jamie Callan. 2022. Unsupervised corpus aware language model pre-training for dense passage retrieval. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2843–2853, Dublin, Ireland.

- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- 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*.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020a. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020b. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6769–6781.
- 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.
- 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*.
- Chunyu Li, Xiang Gao, Yuan Li, Baolin Peng, Xiuming Li, Yizhe Zhang, and Jianfeng Gao. 2020. Optimus: Organizing sentences via pre-trained modeling of a latent space. *arXiv preprint arXiv:2004.04092*.
- Jimmy Lin, Rodrigo Nogueira, and Andrew Yates. 2021. Pretrained transformers for text ranking: Bert and beyond. *Synthesis Lectures on Human Language Technologies*, 14(4):1–325.
- 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. *CoRR*, abs/1907.11692.
- Shuqi Lu, Di He, Chenyan Xiong, Guolin Ke, Waleed Malik, Zhicheng Dou, Paul Bennett, Tie-Yan Liu, and Arnold Overwijk. 2021. Less is more: Pretrain a strong Siamese encoder for dense text retrieval using a weak decoder. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2780–2791.
- Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. Twinbert: Distilling knowledge to twin-structured bert models for efficient retrieval. *arXiv preprint arXiv:2002.06275*.
- Yuxiang Lu, Yiding Liu, Jiayang Liu, Yunsheng Shi, Zhengjie Huang, Shikun Feng, Yu Sun, Hao Tian, Hua Wu, Shuaiqiang Wang, Dawei Yin, et al. 2022. Ernie-search: Bridging cross-encoder with dual-encoder via self on-the-fly distillation for dense passage retrieval. *arXiv preprint arXiv:2205.09153*.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. Sparse, dense, and attentional representations for text retrieval. *Transactions of the Association for Computational Linguistics*, 9:329–345.
- Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 42(4):824–836.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human generated machine reading comprehension dataset. In *CoCo@ NIPS*.
- 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*.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2020. Rocketqa: An optimized training approach to dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2010.08191*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- 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*.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2021. Colbertv2: Effective and efficient retrieval via lightweight late interaction. *arXiv preprint arXiv:2112.01488*.
- 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*.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*.
- Kexin Wang, Nils Reimers, and Iryna Gurevych. 2021. Tsdad: Using transformer-based sequential denoising auto-encoder for unsupervised sentence embedding learning. *arXiv preprint arXiv:2104.06979*.
- Shitao Xiao, Zheng Liu, Weihao Han, Jianjin Zhang, Defu Lian, Yeyun Gong, Qi Chen, Fan Yang, Hao Sun, Yingxia Shao, et al. 2022a. Distill-vq: Learning retrieval oriented vector quantization by distilling knowledge from dense embeddings. *arXiv preprint arXiv:2204.00185*.
- Shitao Xiao, Zheng Liu, Yingxia Shao, Tao Di, Bhuvan Middha, Fangzhao Wu, and Xing Xie. 2022b. Training large-scale news recommenders with pretrained language models in the loop. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4215–4225.
- Shitao Xiao, Zheng Liu, Yingxia Shao, Defu Lian, and Xing Xie. 2021. Matching-oriented product quantization for ad-hoc retrieval. *arXiv preprint arXiv:2104.07858*.
- 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*.
- Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. 2022. Laprador: Unsupervised pretrained dense retriever for zero-shot text retrieval. *arXiv preprint arXiv:2203.06169*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Hang Zhang, Yeyun Gong, Yelong Shen, Jiancheng Lv, Nan Duan, and Weizhu Chen. 2021. Adversarial retriever-ranker for dense text retrieval. *arXiv preprint arXiv:2110.03611*.
- Jianjin Zhang, Zheng Liu, Weihao Han, Shitao Xiao, Ruicheng Zheng, Yingxia Shao, Hao Sun, Hanqing Zhu, Premkumar Srinivasan, Weiwei Deng, et al. 2022. Uni-retriever: Towards learning the unified embedding based retriever in bing sponsored search. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4493–4501.