

# Injecting Domain Adaptation with Learning-to-hash for Effective and Efficient Zero-shot Dense Retrieval

Nandan Thakur<sup>1</sup>, Nils Reimers<sup>2</sup> and Jimmy Lin<sup>1</sup>

<sup>1</sup>David R. Cheriton School of Computer Science, University of Waterloo

<sup>2</sup>Cohere.ai

## Abstract

Dense retrieval overcome the lexical gap and has shown great success in ad-hoc information retrieval (IR). Despite their success, dense retrievers are expensive to serve across practical use cases. For use cases requiring to search from millions of documents, the dense index becomes bulky and requires high memory usage for storing the index. More recently, learning-to-hash (LTH) techniques, for e.g., BPR and JPQ, produce binary document vectors, thereby reducing the memory requirement to efficiently store the dense index. LTH techniques are supervised and finetune the retriever using a ranking loss. They outperform their counterparts, i.e., traditional out-of-the-box vector compression techniques such as PCA or PQ. A missing piece from prior work is that existing techniques have been evaluated only in-domain, i.e., on a single dataset such as MS MARCO. In our work, we evaluate LTH and vector compression techniques for improving the downstream zero-shot retrieval accuracy of the TAS-B dense retriever while maintaining efficiency at inference. Our results demonstrate that, unlike prior work, LTH strategies when applied naively can underperform the zero-shot TAS-B dense retriever on average by up to 14% nDCG@10 on the BEIR benchmark. To solve this limitation, in our work, we propose an easy yet effective solution of injecting domain adaptation with existing supervised LTH techniques. We experiment with two well-known unsupervised domain adaptation techniques: GenQ and GPL. Our domain adaptation injection technique can improve the downstream zero-shot retrieval effectiveness for both BPR and JPQ variants of the TAS-B model by on average 11.5% and 8.2% nDCG@10 while both maintaining 32× memory efficiency and 14× and 2× speedup respectively in CPU retrieval latency on BEIR. All our code, models, and data are publicly available at <https://github.com/thakur-nandan/income>.

## Keywords

Compression Algorithms, Zero-shot Information Retrieval, Domain Adaptation, Learning to Hash

## 1. Introduction

Dense retrieval has become a core component within several downstream NLP and web-search tasks such as question-answering [1, 2], semantic similarity [3], conversational search [4], entity retrieval [5], fact-checking [6] and passage retrieval [7, 8, 9, 10, 11]. In dense retrieval, dual-encoder models encode semantically correlated queries and documents and represent them as spatially close embeddings, i.e., dense vector representations. Dense retrieval can be efficiently conducted via approximate nearest neighbor (ANN) search at inference [12].

Dual-encoder models encode and store document embeddings within an index, which starts to become bulky, once the documents start to increase within the corpus. It becomes expensive to serve dense retrievers practically. For example, storing 21M (million) passages within a dense index requires about 65 GB of memory [13]. As nearest-neighbor (exact) search in large vector

---

*ReNeuIR'23: Workshop on Reaching Efficiency in Neural Information Retrieval, July 23–27, 2023, Taipei, Taiwan*



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

spaces is rather slow, ANN techniques such as HNSW [14] are popularly used, which results in even higher memory requirements such as 150 GB [13]. As a result, for searching on a very large corpus, we require a high-memory and expensive machine to host the bulky index, which hinders the practical application of dense retrieval.

Aware of the issue, there have been recent efforts to improve the efficiency of dense retrieval models by compressing the vector representation to lower memory, i.e., space requirements. Traditionally, out-of-the-box unsupervised strategies such as dimension reduction with principal component analysis (PCA) or product quantization (PQ) [15] have been popular for retrieval [16, 17]. More recently, supervised learning-to-hash (LTH) techniques have been introduced such as BPR [13] and JPQ [18]. LTH techniques map the original dense embeddings into a Hamming space (binary space) and obtain the low-dimensional binary codes. Unfortunately, all existing vector compression methods are proposed under the single-domain retrieval assumption. However, in practice, many retrieval applications span across various domains with diverse distributions; often with scarce training data [19].

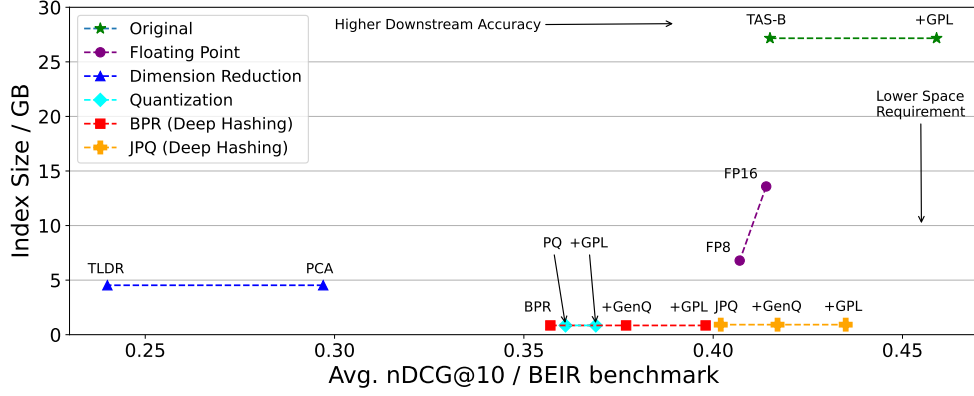
In our work, we focus on the Pareto frontier between downstream zero-shot retrieval accuracy while maintaining the efficiency of compression methods. We evaluate TAS-B [9], a robust generalizable dense retriever on 18 diverse retrieval datasets on the BEIR benchmark [19]. We evaluate out-of-the-box vector compression techniques to efficiently compress the embeddings, with the least drop in zero-shot retrieval effectiveness. We show that prior work on hashing strategies focused on single-domain retrieval fails severely under domain shifts in zero-shot retrieval, and can underperform the zero-shot TAS-B dense retriever on average by up to 14% nDCG@10 on average on BEIR.

To resolve this, in our work, we propose an easy yet effective solution of injecting domain adaptation with existing supervised LTH strategies to improve zero-shot generalization without sacrificing efficiency at inference (both memory and retrieval latency). We focus on the unsupervised setting and experiment with two popular synthetic training data generation-based techniques: GenQ [19] and GPL [20]. For both techniques, we use MS MARCO [21] as the source and keep the rest of the BEIR datasets as the target, as previously experimented in [19, 20, 22]. Figure 1 summarizes the overall results presented in our work. Our injection technique can improve the naive implementations for both the BPR and JPQ variants of TAS-B by 11.5% and 8.2% nDCG@10 on average while maintaining its  $32\times$  memory efficiency and  $14\times$  and  $2\times$  speedup in retrieval latency in comparison to TAS-B dense retriever on BEIR.

## 2. Background Knowledge

### 2.1. TAS-B Dense Retriever

TAS-B [9] is a DistilBERT-based [23] siamese dense retrieval model proposed by Hofstätter et al. [9]. A Balanced Topic Aware Sampling (TAS-B) strategy was implemented to form training batches for optimizing retrieval effectiveness in a data-efficient manner. It was one of the earliest dense retrieval models to successfully exploit knowledge distillation, using dual supervision from two teacher models: a cross-encoder model and ColBERT [24]. The TAS-B model has previously been shown to achieve strong generalization performance on BEIR in Thakur et al. [19], outperforming ANCE [8] and DPR [2]. Due to its strong generalization capabilities, we chose TAS-B for our work.



**Figure 1:** Summary of index size (estimated for MS MARCO [21] in GB) versus effectiveness measured as downstream zero-shot retrieval accuracy, i.e. average nDCG@10 on BEIR [19] for all vector compression models evaluated in our work. A higher nDCG@10 and lower index size is desired. Injecting GPL improves BPR and JPO (zero-shot) by 11.5% and 8.2% respectively without any increase in index size.

## 2.2. Binary Passage Retriever (BPR)

BPR, as introduced by Yamada et al. [13], is a binary retriever trained under supervision in a learning-to-hash setup. BPR builds upon the existing dual-encoder architecture ( $f_\theta$ ; parameterized by  $\theta$ ) and adds a hash layer, which implements a  $\text{sign}(\cdot)$  function to convert the dense embedding  $f_\theta(p)$  into a binary hash  $h(p) = \text{sign}(f_\theta(p))$  where  $h(p) \in \{-1, 1\}^{dim}$ . This architecture helps us save memory as it reduces the size of the dense index by storing lightweight binary hash values in a corpus. During inference, similarity scores are computed in two stages during inference using a retrieve and rerank design. First, BPR retrieves top- $k$  passages for an input query hash  $h(q)$  and across all hashed passages  $h(p)$  using hamming loss. Next, BPR reranks the top- $k$  passage hashes via dot product with the query embedding  $\{\langle e(q), h(p_i) \rangle | i \leq k\}$ .

BPR is trained to represent the binary passage hashes  $h(p)$  by optimizing both ranking and contrastive losses. The ranking loss function with the approximated query hash  $\tilde{h}(q)$  and passage hash  $\tilde{h}(p)$  is defined as:

$$\mathcal{L}_{\text{rank}} = \sum_{j=1}^{n-1} \max(0, -(\tilde{h}(q_i) \cdot \tilde{h}(p_i^+) - \tilde{h}(q_i) \cdot \tilde{h}(p_{i,j}^-)) + \alpha) \quad (1)$$

where  $\alpha$  is a ranking constant. Similarly, the infoNCE [25] loss function  $\mathcal{L}_{\text{infoNCE}}$  teaches the model to improve reranking with only passage hashes (second-stage inference). The loss function with the original query embedding  $e(q)$  is defined as:

$$\mathcal{L}_{\text{infoNCE}} = -\log \frac{e^{e(q_i) \cdot h(p_i^+)}}{e^{e(q_i) \cdot h(p_i^+)} + \sum_{j=1}^{n-1} e^{e(q_i) \cdot h(p_{i,j}^-)}} \quad (2)$$

Finally, the model is trained by optimizing for both losses:  $\mathcal{L}_{\text{BPR}} = \mathcal{L}_{\text{rank}} + \mathcal{L}_{\text{infoNCE}}$ .

### 2.3. Joint Optimization of Query Encoder and Product Quantization (JPQ)

JPQ, as introduced by Zhan et al. [18], is a binary retriever trained to jointly optimize the query encoder and PQ index together. In traditional PQ, due to the separation between encoding and compression, the PQ training cannot benefit from the first-stage supervised ranking. JPQ overcomes this challenge and learns to optimize the position of the PQ centroid embeddings, by changing the centroid embeddings to push the relevant document closer to the query embedding and the irrelevant documents further away. At inference, the similarity is computed between the reconstructed quantized passage hash from trained PQ (quantized) centroid embeddings  $h(p)$  and the original query embedding  $e(q)$ . The approximate similarity score is calculated by computing the dot product:  $s(q, p) = \langle f_\theta(q), h(p) \rangle$ .

JPQ model is trained to minimize the pairwise loss function  $l$  between the query embedding  $e(q)$  and reconstructed quantized positive  $h(p^+)$  and negative passage  $h(p^-)$  using InfoNCE [25] loss. The final loss function is given as:

$$\mathcal{L}_{\text{JPQ}} = - \sum_{i=0}^{n-1} \sum_{i=0}^{n-1} l(e(q) \cdot h(p^+), e(q) \cdot h(p^-))$$

## 3. Domain Adaptation Injection

As our experimental results will show in Section 5, existing LTH strategies such as BPR are not robust against domain shifts within BEIR. To overcome this, we present an easy and effective approach to inject domain adaptation within existing LTH strategies: BPR [13] and JPQ [18].

### 3.1. Injecting GenQ Technique

The GenQ approach, presented by Thakur et al. [19] uses a T5-base model trained on MSMARCO is used to generate top- $k$  synthetic queries  $\hat{q}_i$  for a given passage  $p_i$  from each target corpora in BEIR [19] individually. For injection within existing learning-to-hash (LTH) systems such as BPR and JPQ, we do not change the underlying training setup. We feed the top- $k$  synthetic query-passage training pairs  $(\hat{q}_i, p_i)$  and use them for finetuning BPR and JPQ to better generalize across out-of-domain retrieval tasks in BEIR.

### 3.2. Injecting GPL Technique

GPL (Generative Pseudo Labeling) presented by Wang et al. [20], is an extension of the GenQ technique. GPL combines query generation with a pseudo-labeling approach to filter out low-quality generated queries. More formally, for a generated query  $\hat{q}$  from the passage  $p^+$ , the method first retrieves top- $k$  matching passages  $\{p_1^-, p_2^-, \dots, p_k^-\}$  from our corpus using an existing retriever. GPL next randomly samples passages to create the synthetic triplets  $(\hat{q}, p^+, p^-)$ . Finally, a cross-encoder (CE) model is used to pseudo-label the similarity score  $s^+ = (\hat{q}, p^+)$  as well as  $s^- = (\hat{q}, p^-)$  and compute the score-margin  $\delta = s^+ - s^-$ .

Both BPR and JPQ models train the retriever model using infoNCE loss [25] which is found sensitive to synthetic low-quality generated queries generated using GenQ, as shown previously

in [20]. Hence, in our work, we improve the existing infoNCE loss and replace it effectively with a MarginMSE loss [26]. The first-stage ranking loss function  $\mathcal{L}_{\text{rank}}$  for BPR is unchanged. In our work, instead of  $\mathcal{L}_{\text{InfoNCE}}$ , we will be replacing it in both BPR and JPQ with:

$$\mathcal{L}_{\text{GPL}} = - \sum_{i=0}^{n-1} |(e(\hat{q}_i) \cdot h(p_i^+) - e(\hat{q}_i) \cdot h(p_i^-)) - (\text{CE}(\hat{q}_i, p_i^+) - \text{CE}(\hat{q}_i, p_i^-))|^2,$$

where  $\text{CE}(\hat{q}_i, p)$  is the pseudo-label or similarity score from a cross-encoder teacher model. Note, we can use both the T5 query generation and cross-encoder model from pre-existing models that have been trained on the source domain like MSMARCO.

## 4. Experimental Setup

### 4.1. Datasets and Evaluation

For all our experiments, we use the MS MARCO passage ranking dataset [21] as training data for our BPR and JPQ base implementations. It has 8.8M passages and 532.8K query-passage pairs labeled as relevant in the training set. We use the BEIR benchmark [19] to measure the zero-shot generalization accuracy of our models. The BEIR benchmark contains 18 different zero-shot retrieval evaluation datasets with diverse domains and tasks. Following prior work in [19, 20], we evaluate the zero-shot retrieval accuracy at nDCG@10.

### 4.2. Baseline Methods

In our work, we evaluate four different categories of vector compression baselines: (1) original (without compression), (2) floating point, (3) dimension reduction, and (4) quantization. For all baselines, we use the already finetuned TAS-B [9] model as the backbone and evaluate them in a zero-shot retrieval setup across 18 diverse datasets from the BEIR benchmark [19].

- **Original:** We evaluate two dense retriever baselines: (1) “zero-shot” baseline with TAS-B from Thakur et al. [19] and (2) “upper-baseline” of TAS-B + GPL from Wang et al. [20].
- **Floating point:** We reduce the floating-point precision of the zero-shot TAS-B model from the default 32 bits [19] to 16 bits (fp16) and 8 bits (fp8) respectively.
- **Dimensional reduction:** We fit (1) principle component analysis (PCA) transformation across all embeddings produced by TAS-B for each test dataset in BEIR. We reduce the dimension of the passage embedding from 768 to 128 dimensions. (2) We apply the identical approach for evaluating TLDR by Kalantidis et al. [27].
- **Quantization:** We decompose in (1) product quantization (PQ) [15] the original  $d$  dimensional vector into  $k$  centroids, quantized and stored with  $n$  bits. We apply product quantization with the zero-shot TAS-B model. (2) We evaluate the TAS-B model with GPL across each dataset in BEIR and apply product quantization (PQ+GPL).

### 4.3. Experimental Setup

In this section, we describe the implementation details which are necessary for baselines, training learn-to-hash strategies, and domain-adaptation techniques. All our code can be found publicly for reproduction purposes. All our experiments only require a single NVIDIA RTX A6000 GPU resource containing a maximum of 48 GB VRAM. We implement all our models using the HuggingFace’s Transformers package [28]. For all our experiments, we use the TAS-B model<sup>1</sup> as backbone containing a maximum sequence length of 350 tokens.

**Baselines.** We use faiss [12] for implementing all our vector compression baselines, except TLDR. For the floating-point baselines, we use the faiss ScalarQuantizer<sup>2</sup> in flat mode for both fp8 and fp16. For PCA, we use faiss PCAMatrix<sup>3</sup> in the flat mode and downward project the original dimension of TAS-B model from 768 into 128 dimensions. For PCA, we normalized both the query and document embeddings and computed cosine similarity with additional whitening. We did not find any significant improvement by fitting the MSMARCO dataset with each BEIR test dataset, therefore we trained PCA on each BEIR test corpus individually. For TLDR,<sup>4</sup> we train one additional multi-layer perceptron (MLP) layer with 2048 dimensions using 5 nearest neighbors to sample positive pairs and downward project the TAS-B model to 128 dimensions. For the PQ implementation, we use faiss ProductQuantizer<sup>5</sup> in flat mode. To achieve a compression of 32 $\times$ , we configure PQ with  $k = 96$  centroids and  $n = 8$  bits respectively. We also experimented with OPQ [29], however, we found it to underperform naive PQ. Similar to PCA, we did not find any significant improvement by fitting the MSMARCO dataset with each BEIR test dataset, therefore we train PQ on each BEIR test corpus individually.

**Learning-to-Hash techniques.** For training BPR and JPQ models, we first replicate the original training setup of BPR<sup>6</sup> and JPQ<sup>7</sup> using the TAS-B model as the backbone. Next, we choose TAS-B model (already finetuned on MS MARCO following [9]) as the backbone and further finetune on the MS MARCO triplets<sup>8</sup> publicly available in Sentence Transformers repository to convert the existing TAS-B dense retriever into binary retrievers. As we show in Appendix (Table 3), we find the TAS-B model to outperform original NQ-trained BPR [13] and STAR-based JPQ [18] models.<sup>9</sup> In our experiments, this additional fine-tuning stage is inexpensive and quick and only requires only a maximum of 50K additional steps (4-5 hours in a single NVIDIA RTX A6000 GPU) to convert TAS-B into its binary retriever equivalents, i.e., BPR [13] and JPQ [18] as shown in the Appendix (Figure 2).

---

<sup>1</sup><https://huggingface.co/sentence-transformers/msmarco-distilbert-base-tas-b>

<sup>2</sup>[https://faiss.ai/cpp\\_api/struct/structfaiss\\_1\\_1ScalarQuantizer.html](https://faiss.ai/cpp_api/struct/structfaiss_1_1ScalarQuantizer.html)

<sup>3</sup>[https://faiss.ai/cpp\\_api/struct/structfaiss\\_1\\_1PCAMatrix.html](https://faiss.ai/cpp_api/struct/structfaiss_1_1PCAMatrix.html)

<sup>4</sup>TLDR repository: <https://github.com/naver/tldr>

<sup>5</sup>[https://faiss.ai/cpp\\_api/struct/structfaiss\\_1\\_1ProductQuantizer.html](https://faiss.ai/cpp_api/struct/structfaiss_1_1ProductQuantizer.html)

<sup>6</sup>BPR repository: <https://github.com/studio-ousia/bpr>

<sup>7</sup>JPQ repository: <https://github.com/jingtaozhan/JPQ>

<sup>8</sup>MS MARCO triplets: <https://sbert.net/datasets/msmarco-hard-negatives.jsonl.gz>

<sup>9</sup>Our method described is model agnostic, i.e., we can choose any model as the backbone for our method. We intuitively observe that a robust dense retriever model as a backbone in our method provides a higher downstream retrieval accuracy with binary retrievers.



**Domain adaptation injection.** To generate queries for both GenQ and GPL, we use docT5query [30] model finetuned on MS MARCO and generate at most 3 queries per passage using nucleus sampling with temperature 1.0,  $k = 25$  and  $p = 0.95$ . For BPR+GenQ injection, we follow [19] and finetune the TAS-B model as the backbone for 1 epoch and with a batch size of 75 for each individual BEIR synthetic query-passage dataset. We do not mine hard negatives for training with GenQ, i.e. negatives are randomly sampled within the batch. For BPR+GPL injection, we follow [20] and finetune the TAS-B model as the backbone for either 1 epoch or a maximum of 45K training steps with a batch size of 32. To retrieve hard negatives, we use three popular models finetuned on MS MARCO available in the SentenceTransformers repository: (1) *msmarco-distilbert-base-tas-b* (backbone model itself), (2) *msmarco-distilbert-base-v3*<sup>10</sup> and (3) *msmarco-MiniLM-L6-v3*<sup>11</sup>. Following [20], for each model, we retrieve a maximum of 50 negatives and uniformly sample one negative passage and one positive passage for each training query to form one training example. For the cross-encoder teacher, we use the *ms-marco-MiniLM-L6-v2*<sup>12</sup> model, following [20]. For JPQ+GenQ injection, we follow [18] and finetune the TAS-B model as the backbone for 2 epochs with a learning rate of  $1e^{-4}$ , training batch size of 256 and top-200 PQ-retrieved hard negatives. For JPQ+GPL injection, we follow the same setting as JPQ+GenQ injection and replace the infoNCE loss function with the MarginMSE loss [26]. Since labeling all top-200 hard negatives is computationally expensive with a cross-encoder at each training step with JPQ, we reduce and only label the top-25 hard negatives retrieved using the trainable PQ index at each training step.

## 5. Experimental Results

### 5.1. Zero-Shot Baseline Results on BEIR

From Table 1, we observe GPL improves zero-shot TAS-B by a huge margin in nDCG@10 on BEIR (0.415→0.459) which is consistent with prior work [20]. fp16 and fp8 to be strong baselines for upto 2-4× light compression within a single point difference nDCG@10 on BEIR. This shows a majority of the relevant information in dense retriever is contained within initial 16 or 8-bit floating decimals. Only ArguAna has a visible drop in nDCG@10 (0.429→0.321), where we likely suspect useful information is captured in between 8-16 bits of the query vector (which is a long passage in ArguAna). Dimension reduction finds a huge drop in zero-shot retrieval accuracy, and performs worst in terms of Pareto optimality. This indicates that useful information captured within the original 768 dimensions of the TAS-B dense retriever is lost during the downward projection into lower dimensions. For PCA, we observed whitening and cosine similarity have a visible improvement in nDCG@10 on average in BEIR (0.235→0.297). Contrary to TLDR, we did not observe any improvements with whitening or using cosine similarity. Quantization performs best in terms of Pareto optimality. Contrary to prior work in [29], OPQ first conducts a dimensionality reduction and is found to perform worse than PQ in zero-shot retrieval on BEIR. PQ even outperforms zero-shot BPR by a small margin of 0.4 nDCG@10 points. GPL only improves PQ marginally by 0.8 points (0.361→0.369) nDCG@10 as synthetic queries are added, whereas PQ is trained on domain-specific corpora.

<sup>10</sup><https://huggingface.co/sentence-transformers/msmarco-distilbert-base-v3>

<sup>11</sup><https://huggingface.co/sentence-transformers/msmarco-MiniLM-L-6-v3>

<sup>12</sup><https://huggingface.co/cross-encoder/ms-marco-MiniLM-L-6-v2>

Model (→)	Original		Floating Point		Dim. Reduction		Quantization		Deep Hashing (BPR)			Deep Hashing (JPQ)		
Dataset (↓)	TAS-B	+GPL	fp16	fp8	PCA	TLDR	PQ	+GPL	BPR	+GenQ	+GPL	JPQ	+GenQ	+GPL
Memory	1×	1×	2×	4×	6×	6×	32×	32×	32×	32×	32×	32×	32×	32×
MS MARCO	0.408	0.408	0.407	0.408	0.315	0.266	0.358	0.358	0.397	0.397	0.397	0.400	0.400	0.400
TREC-COVID	0.481	0.700	0.481	0.484	0.324	0.287	0.426	0.552	0.400	0.535	0.593	0.601	0.702	0.717
BioASQ	0.383	0.442	0.359	0.357	0.153	0.097	0.253	0.311	0.285	0.353	0.388	0.326	0.362	0.384
NFCorpus	0.319	0.345	0.319	0.317	0.267	0.195	0.295	0.303	0.286	0.296	0.299	0.312	0.320	0.324
NQ	0.463	0.483	0.463	0.461	0.343	0.265	0.402	0.389	0.447	0.333	0.457	0.446	0.404	0.456
HotpotQA	0.584	0.582	0.584	0.583	0.388	0.170	0.497	0.430	0.482	0.474	0.536	0.552	0.577	0.593
FiQA-2018	0.300	0.344	0.300	0.300	0.207	0.156	0.251	0.281	0.255	0.285	0.309	0.289	0.288	0.295
Signal-1M (RT)	0.289	0.276	0.290	0.287	0.235	0.237	0.263	0.241	0.241	0.270	0.278	0.269	0.271	0.280
TREC-NEWS	0.377	0.421	0.377	0.377	0.241	0.134	0.306	0.331	0.338	0.367	0.268	0.366	0.363	0.381
Robust04	0.427	0.437	0.428	0.428	0.280	0.172	0.348	0.334	0.309	0.337	0.374	0.398	0.392	0.416
ArguAna	0.429	0.557	0.429	0.321	0.333	0.374	0.326	0.348	0.316	0.348	0.354	0.389	0.453	0.443
Touche-2020	0.162	0.255	0.163	0.164	0.098	0.135	0.186	0.169	0.167	0.255	0.259	0.176	0.179	0.220
CQADupStack	0.314	0.357	0.314	0.314	0.270	0.202	0.272	0.309	0.287	0.329	0.336	0.296	0.329	0.337
Quora	0.835	0.836	0.835	0.835	0.773	0.829	0.825	0.813	0.852	0.840	0.853	0.830	0.849	0.847
DBPedia	0.384	0.384	0.384	0.384	0.304	0.235	0.351	0.322	0.335	0.304	0.348	0.370	0.352	0.385
SCIDOCS	0.149	0.169	0.149	0.147	0.116	0.104	0.124	0.132	0.130	0.137	0.152	0.134	0.148	0.155
FEVER	0.700	0.759	0.700	0.700	0.386	0.213	0.629	0.607	0.571	0.559	0.539	0.664	0.681	0.719
Climate-FEVER	0.228	0.235	0.228	0.227	0.105	0.103	0.192	0.167	0.177	0.201	0.204	0.194	0.192	0.226
SciFact	0.643	0.674	0.643	0.640	0.514	0.407	0.567	0.596	0.548	0.565	0.623	0.622	0.647	0.655
AVERAGE	0.415	0.459	0.414	0.407	0.297	0.240	0.361	0.369	0.357	0.377	0.398	0.402	0.417	0.435

**Table 1**

Retrieval accuracy (nDCG@10) by different compression models on BEIR [19] using TAS-B [9] as backbone. models starting with (+) denotes addition, for e.g. BPR+GPL or JPQ+GenQ. Any compression model outperforming both the original baselines, i.e., TAS-B (zero-shot) and TAS-B+GPL (upper-bound), is highlighted with **green**, whereas if a model outperforms any one of them is highlighted with **yellow**.

## 5.2. Learning-to-Hash Results on BEIR

From Table 1, BPR performs competitively in the in-domain retrieval setting on MS MARCO. BPR drops marginally by only 1.1 points nDCG@10 (0.408→0.397) in comparison with TAS-B dense retriever while gaining 32× memory compression in the tradeoff. Likewise, the JPQ model only shows a drop of 0.8 points nDCG@10 on MS MARCO. This follows that LTH techniques are effective in learning in-domain distributions as shown in prior work [18, 13]. However, when evaluated zero-shot on BEIR, the performance gap widens. BPR underperforms the TAS-B zero-shot retriever by a substantial margin of 5.8 points nDCG@10. BPR even underperforms the naive quantization PQ baseline by a small margin of 0.4 points nDCG@10. In contrast, we find JPQ as a better robust learning-to-hash (LTH) algorithm. JPQ underperforms the TAS-B dense retriever by 1.3 points and is 4.1 points at nDCG@10 better than BPR. We likely suspect this is due to quantization being more robust to domain shifts, however, an in-depth analysis is required for a better explanation which we keep as future work.

## 5.3. Domain Adaptation Injection Results on BEIR

We summarize the domain adaptation injection results with GPL and GenQ in Table 1. First, GenQ injection with BPR and JPQ leads to an improvement in specific BEIR datasets which span a single (focused) domain such as FiQA, SciFact, and BioASQ. Datasets in BEIR which cover a broad set of topics (e.g. NQ, DBPedia), GenQ performs worse than the BPR, due to the reasons previously outlined in [20, 19]. Overall, GenQ is a successful domain adaptation technique and improves both the zero-shot retrieval accuracy of BPR and JPQ by 2.0 and 1.5 points nDCG@10 on BEIR, while maintaining 32× memory efficiency.



**Table 2**

Efficiency results reported for all vector compression methodologies measured using query latency (in milliseconds; evaluated on CPU and GPU) and index size (in MB) for two datasets from BEIR [19]: DBPedia [31] and MS MARCO [21]. We randomly sample 1M passages from DBPedia for evaluation. We report the mean and standard deviation of retrieval latency. Lower latency and memory are desired.

Model (↓)	DBPedia [31] (1M)			MS MARCO [21] (8.8M)		
	Latency (CPU)	Latency (GPU)	Index Sizes	Latency (CPU)	Latency (GPU)	Index Sizes
Original	277.5 ± 51.81 ms	7.3 ± 0.42 ms	3072.00 MB	2915.5 ± 310.94 ms	63.1 ± 0.36 ms	27162.08 MB
PCA	55.1 ± 5.61 ms	4.3 ± 2.27 ms	514.76 MB	447.7 ± 38.47 ms	31.3 ± 1.49 ms	4529.77 MB
PQ	67.6 ± 5.63 ms	65.1 ± 5.07 ms	96.79 MB	516.7 ± 11.71 ms	524.7 ± 28.98 ms	849.60 MB
BPR	19.7 ± 2.31 ms	n.a. <sup>a</sup>	96.00 MB	150.1 ± 18.78 ms	n.a. <sup>a</sup>	848.82 MB
JPQ	170.4 ± 12.14 ms	8.0 ± 1.66 ms	104.79 MB	1325.2 ± 257.64 ms	75.2 ± 0.75 ms	920.34 MB

<sup>a</sup> We were unable to evaluate the retrieval latency of the BPR model using the faiss GPU index, as the original authors [13] optimized the inference code based on numpy. This is also documented here in this issue.

Next, we find GPL to be highly effective in domain adaptation. GPL adopts the MarginMSE loss which learns effectively to rank using a cross-encoder teacher. BPR+GPL injection is able to improve zero-shot accuracy on 16/18 datasets, improving the average score of BPR by 11.5% (0.357→0.398) nDCG@10 on BEIR. We likely suspect this is due to BPR able to rerank better in the second stage with distilled knowledge from the cross-encoder. Likewise, we observe similar gains with the JPQ model. GPL injection improves the average score of JPQ by 8.2% (0.402→0.435) nDCG@10. GPL injection with JPQ surprisingly even outperforms the original zero-shot TAS-B dense retriever model by 2.0 points nDCG@10 on BEIR. Finally, JPQ+GPL underperforms the upper baseline of TAS-B + GPL by 2.4 points nDCG@10 on BEIR, while gaining a 32× memory efficiency as a Pareto tradeoff.

## 6. Efficiency: Memory and Retrieval Latency

In this section, we compare the efficiency of the vector compression models used in our study. We evaluate our models on two BEIR datasets: (1) 1M sampled passages in DBPedia and (2) MS MARCO. We naively measure the retrieval latency using brute force, i.e. exhaustive search and memory requirements using faiss [12] package. We measure the size of the faiss index (in MB) and evaluate the latency on average to retrieve a single query from each dataset (in milliseconds). We conduct experiments on a server with two Intel Xeon Silver 4210 CPUs and a single NVIDIA RTX A6000 GPU.

**Index sizes.** From Table 2, we observe that the TAS-B dense retriever performs well in terms of zero-shot retrieval accuracy, however, significantly increases the index size by several orders of magnitude. On the other hand, models with quantization such as PQ or supervised LTH techniques such as BPR and JPQ and their variants overall achieve a tiny index. All these techniques substantially compress the index and are able to encode MS MARCO [21] containing 8.8M passages under 1GB in contrast to TAS-B dense retriever which requires more than 27 GB.

**Retrieval latency.** From Table 2, we observe that for the CPU setting, overall all compression techniques are much faster compared to brute-force dense retrieval (original) models. PCA provides 5-6×, PQ provides 4-5×, BPR provides 14-19×, and JPQ provides 2× speedup on average in comparison to the TAS-B dense retriever. For the GPU setting, we find the brute-force

dense retrieval models are significantly faster. TAS-B dense retriever models gain a speedup of 40-45 $\times$  in contrast to CPU inference. PCA provides a 2 $\times$  speedup, whereas JPQ gets a comparable retrieval latency in comparison to the original dense retriever. Surprisingly, we find PQ does not improve its latency when evaluated in the GPU setting.

## 7. Conclusion

Supervised learning-to-hash (LTH) algorithms have been popular and effective across in-domain tasks in retrieval, however, fail to generalize to unseen domains during test time when applied naively. The algorithms are memory efficient but lack downstream accuracy in zero-shot retrieval containing specialized domains for e.g. BEIR benchmark and can severely underperform the zero-shot TAS-B dense retriever. In order to adapt these vector compression algorithms under severe domain shifts, we propose domain adaptation injection in our work. Domain adaptation with GPL improves zero-shot effectiveness of both BPR and JPQ with 11.5% and 8.2% nDCG@10, and maintains the 32 $\times$  memory efficiency and 14 $\times$  and 2 $\times$  CPU speedup on BEIR.

## 8. Limitations and Future Work

In our work, we find a huge improvement in both efficiency and downstream retrieval accuracy with domain-adaptation injection with BPR [13] and JPQ [18]. Our work has a few limitations which we briefly mention and motivate further future work.

- **Different compression algorithms.** In our work, we considered JPQ and BPR due to their popularity and effectiveness as shown in our preliminary results. In the future, we can extend our work to newer LTH algorithms, for e.g. RepCONC [32] or newer models, for e.g. boosted dense retriever [33].
- **Better backbone models.** We suspect the downstream retrieval accuracy on BEIR can be improved with stronger backbone models than TAS-B, for e.g. Contriever [34]. Due to the model-agnostic nature of our method, we can easily extend our work with more robust dense retrievers in the future.
- **Requires separate models.** Domain adaptation with BPR and JPQ both require training separate models for each domain with our technique. This can be quite cumbersome in case there are several domains for which the user would need to train multiple models.
- **Compute intensive.** Our domain-adaptation technique GPL leads to the highest downstream accuracy but is compute-intensive. For hard negative mining, each model is required to compute embeddings for the whole corpus. Next, the cross-encoder teacher increases model training complexity significantly. In the future, we can explore efficient teachers instead such as ColBERT [24] or TILDE [35].

## 9. Acknowledgments

This research was supported in part by the Canada First Research Excellence Fund and the Natural Sciences and Engineering Research Council (NSERC) of Canada. Computational resources were provided by Compute Canada. We would like to thank Kexin Wang for his helpful feedback in earlier drafts of the paper. Additionally, we would like to thank the anonymous reviewers from ReNeuIR 2023 for their helpful feedback and insightful comments.

## References

- [1] K. Lee, M.-W. Chang, K. Toutanova, Latent retrieval for weakly supervised open domain question answering, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 6086–6096. URL: <https://aclanthology.org/P19-1612>. doi:10.18653/v1/P19-1612.
- [2] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, W.-t. Yih, Dense passage retrieval for open-domain question answering, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics, Online, 2020, pp. 6769–6781. URL: <https://www.aclweb.org/anthology/2020.emnlp-main.550>. doi:10.18653/v1/2020.emnlp-main.550.
- [3] N. Reimers, I. Gurevych, Sentence-BERT: Sentence embeddings using Siamese BERT-networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 3982–3992. URL: <https://aclanthology.org/D19-1410>. doi:10.18653/v1/D19-1410.
- [4] S. Yu, Z. Liu, C. Xiong, T. Feng, Z. Liu, Few-Shot Conversational Dense Retrieval, Association for Computing Machinery, New York, NY, USA, 2021, p. 829–838. URL: <https://doi.org/10.1145/3404835.3462856>.
- [5] D. Gillick, S. Kulkarni, L. Lansing, A. Presta, J. Baldrige, E. Ie, D. Garcia-Olano, Learning dense representations for entity retrieval, in: Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 528–537. URL: <https://aclanthology.org/K19-1049>. doi:10.18653/v1/K19-1049.
- [6] C. Samarinas, W. Hsu, M. L. Lee, Improving evidence retrieval for automated explainable fact-checking, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations, Association for Computational Linguistics, Online, 2021, pp. 84–91. URL: <https://aclanthology.org/2021.naacl-demos.10>. doi:10.18653/v1/2021.naacl-demos.10.
- [7] L. Gao, Z. Dai, T. Chen, Z. Fan, B. Van Durme, J. Callan, Complement lexical retrieval model with semantic residual embeddings, in: D. Hiemstra, M.-F. Moens, J. Mothe, R. Perego, M. Potthast, F. Sebastiani (Eds.), Advances in Information Retrieval, Springer International Publishing, Cham, 2021, pp. 146–160. URL: [https://link.springer.com/chapter/10.1007/978-3-030-72113-8\\_10](https://link.springer.com/chapter/10.1007/978-3-030-72113-8_10).
- [8] L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. N. Bennett, J. Ahmed, A. Overwijk, Approximate nearest neighbor negative contrastive learning for dense text retrieval, in: International Conference on Learning Representations, 2021. URL: <https://openreview.net/forum?id=zeFrfgYZln>.
- [9] S. Hofstätter, S.-C. Lin, J.-H. Yang, J. Lin, A. Hanbury, Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling, Association for Computing Machinery, New York, NY, USA, 2021, p. 113–122. URL: <https://doi.org/10.1145/3404835.3462891>.
- [10] J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, S. Ma, Optimizing Dense Retrieval Model Training with Hard Negatives, Association for Computing Machinery, New York, NY, USA, 2021, p.

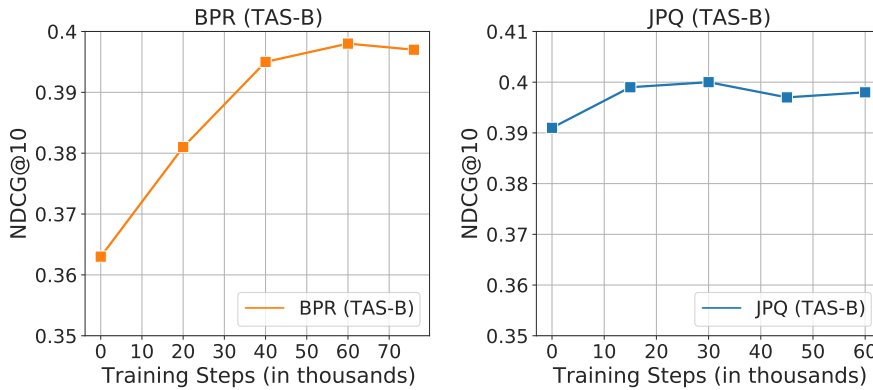
- 1503–1512. URL: <https://doi.org/10.1145/3404835.3462880>.
- [11] Y. Luan, J. Eisenstein, K. Toutanova, M. Collins, Sparse, Dense, and Attentional Representations for Text Retrieval, *Transactions of the Association for Computational Linguistics* 9 (2021) 329–345. URL: [https://doi.org/10.1162/tacl\\_a\\_00369](https://doi.org/10.1162/tacl_a_00369). doi:10.1162/tacl\_a\_00369.
- [12] J. Johnson, M. Douze, H. Jégou, Billion-scale similarity search with GPUs, *IEEE Transactions on Big Data* 7 (2019) 535–547. URL: <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8733051>.
- [13] I. Yamada, A. Asai, H. Hajishirzi, Efficient passage retrieval with hashing 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 (Volume 2: Short Papers)*, Association for Computational Linguistics, Online, 2021, pp. 979–986. URL: <https://aclanthology.org/2021.acl-short.123>. doi:10.18653/v1/2021.acl-short.123.
- [14] Y. Malkov, A. Ponomarenko, A. Logvinov, V. Krylov, Approximate nearest neighbor algorithm based on navigable small world graphs, *Information Systems* 45 (2014) 61–68. URL: <https://www.sciencedirect.com/science/article/pii/S0306437913001300>. doi:<https://doi.org/10.1016/j.is.2013.10.006>.
- [15] H. Jegou, M. Douze, C. Schmid, Product quantization for nearest neighbor search, *IEEE transactions on pattern analysis and machine intelligence* 33 (2010) 117–128. URL: <https://hal.inria.fr/inria-00514462v2/document>.
- [16] G. Izacard, F. Petroni, L. Hosseini, N. D. Cao, S. Riedel, E. Grave, A memory efficient baseline for open domain question answering, 2020. [arXiv:2012.15156](https://arxiv.org/abs/2012.15156).
- [17] X. Ma, M. Li, K. Sun, J. Xin, J. Lin, Simple and effective unsupervised redundancy elimination to compress dense vectors for passage retrieval, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 2854–2859. URL: <https://aclanthology.org/2021.emnlp-main.227>. doi:10.18653/v1/2021.emnlp-main.227.
- [18] J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, S. Ma, Jointly optimizing query encoder and product quantization to improve retrieval performance, *CIKM '21*, Association for Computing Machinery, New York, NY, USA, 2021, p. 2487–2496. URL: <https://doi.org/10.1145/3459637.3482358>. doi:10.1145/3459637.3482358.
- [19] N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, I. Gurevych, BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models, in: *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL: <https://openreview.net/forum?id=wCu6T5xFjeJ>.
- [20] K. Wang, N. Thakur, N. Reimers, I. Gurevych, GPL: Generative pseudo labeling for unsupervised domain adaptation of dense retrieval, in: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Seattle, United States, 2022, pp. 2345–2360. URL: <https://aclanthology.org/2022.naacl-main.168>. doi:10.18653/v1/2022.naacl-main.168.
- [21] T. Nguyen, M. Rosenberg, X. Song, J. Gao, S. Tiwary, R. Majumder, L. Deng, Ms marco: A human generated machine reading comprehension dataset, in: *CoCo@ NIPS*, 2016. URL: <https://www.microsoft.com/en-us/research/publication/>

ms-marco-human-generated-machine-reading-comprehension-dataset/.

- [22] J. Xin, C. Xiong, A. Srinivasan, A. Sharma, D. Jose, P. Bennett, Zero-shot dense retrieval with momentum adversarial domain invariant representations, in: Findings of the Association for Computational Linguistics: ACL 2022, Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 4008–4020. URL: <https://aclanthology.org/2022.findings-acl.316>.
- [23] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2020. [arXiv:1910.01108](https://arxiv.org/abs/1910.01108).
- [24] O. Khattab, M. Zaharia, ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT, Association for Computing Machinery, New York, NY, USA, 2020, p. 39–48. URL: <https://doi.org/10.1145/3397271.3401075>.
- [25] A. van den Oord, Y. Li, O. Vinyals, Representation learning with contrastive predictive coding, [arXiv preprint arXiv:1807.03748](https://arxiv.org/abs/1807.03748) (2019). URL: <https://arxiv.org/abs/1807.03748>.
- [26] S. Hofstätter, S. Althammer, M. Schröder, M. Sertkan, A. Hanbury, Improving efficient neural ranking models with cross-architecture knowledge distillation, 2021. [arXiv:2010.02666](https://arxiv.org/abs/2010.02666).
- [27] Y. Kalantidis, C. E. R. K. Lassance, J. Almazán, D. Larlus, TLDR: Twin learning for dimensionality reduction, Transactions on Machine Learning Research (2022). URL: <https://openreview.net/forum?id=86fhqdBUBx>, expert Certification.
- [28] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, A. Rush, Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online, 2020, pp. 38–45. URL: <https://aclanthology.org/2020.emnlp-demos.6>. doi:10.18653/v1/2020.emnlp-demos.6.
- [29] T. Ge, K. He, Q. Ke, J. Sun, Optimized product quantization, IEEE Transactions on Pattern Analysis and Machine Intelligence 36 (2014) 744–755. doi:10.1109/TPAMI.2013.240.
- [30] R. Nogueira, J. Lin, From doc2query to docTTTTTquery, 2019. URL: [https://cs.uwaterloo.ca/~jimmylin/publications/Nogueira\\_Lin\\_2019\\_docTTTTTquery-v2.pdf](https://cs.uwaterloo.ca/~jimmylin/publications/Nogueira_Lin_2019_docTTTTTquery-v2.pdf).
- [31] F. Hasibi, F. Nikolaev, C. Xiong, K. Balog, S. E. Bratsberg, A. Kotov, J. Callan, Dbpedia-entity v2: A test collection for entity search, in: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 1265–1268. URL: <https://doi.org/10.1145/3077136.3080751>. doi:10.1145/3077136.3080751.
- [32] J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, S. Ma, Learning Discrete Representations via Constrained Clustering for Effective and Efficient Dense Retrieval, Association for Computing Machinery, New York, NY, USA, 2022, p. 1328–1336. URL: <https://doi.org/10.1145/3488560.3498443>.
- [33] P. Lewis, B. Oguz, W. Xiong, F. Petroni, S. Yih, S. Riedel, Boosted dense retriever, in: Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Seattle, United States, 2022, pp. 3102–3117. URL: <https://aclanthology.org/2022.naacl-main.226>. doi:10.18653/v1/2022.naacl-main.226.
- [34] G. Izacard, M. Caron, L. Hosseini, S. Riedel, P. Bojanowski, A. Joulin, E. Grave, Unsupervised

- dense information retrieval with contrastive learning, *Trans. Mach. Learn. Res.* 2022 (2022). URL: <https://openreview.net/forum?id=jKN1pXi7b0>.
- [35] S. Zhuang, G. Zuccon, Fast passage re-ranking with contextualized exact term matching and efficient passage expansion, 2021. [arXiv:2108.08513](https://arxiv.org/abs/2108.08513).
- [36] J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, S. Ma, Optimizing dense retrieval model training with hard negatives, in: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, Association for Computing Machinery, New York, NY, USA, 2021, p. 1503–1512. URL: <https://doi.org/10.1145/3404835.3462880>. doi:10.1145/3404835.3462880.
- [37] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, 2019. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692).





**Figure 2:** The number of training steps (in thousands) required to finetune TAS-B [9] model for converting into BPR or JPQ. The y-axis denotes the nDCG@10 performance on the MSMARCO DEV [21] dataset. The training takes approximately 4-5 hours on a single NVIDIA RTX A6000 GPU. The in-domain MS MARCO performance starts to saturate after 30-40K steps.

Dataset ( $\downarrow$ ) / Model ( $\rightarrow$ )	JPQ (STAR) [18]	JPQ (TAS-B) (our work)	BPR (DPR) [13]	BPR (TAS-B) (our work)
Memory Efficiency	<b>32</b> $\times$	<b>32</b> $\times$	<b>32</b> $\times$	<b>32</b> $\times$
MSMARCO	0.402	0.400	0.130	0.397
TREC-COVID	<b>0.654</b>	0.601	0.201	<b>0.400</b>
BioASQ	0.306	<b>0.326</b>	0.040	<b>0.285</b>
NFCorpus	0.237	<b>0.312</b>	0.115	<b>0.286</b>
NQ	0.446	0.446	0.399	<b>0.447</b>
HotpotQA	0.456	<b>0.552</b>	0.240	<b>0.482</b>
FiQA-2018	<b>0.295</b>	0.289	0.081	<b>0.255</b>
Signal-1M (RT)	0.249	<b>0.269</b>	0.157	<b>0.241</b>
TREC-NEWS	<b>0.382</b>	0.366	0.161	<b>0.338</b>
Robust04	0.392	<b>0.398</b>	0.207	<b>0.309</b>
ArguAna	<b>0.415</b>	0.389	0.201	<b>0.316</b>
Touche-2020	<b>0.240</b>	0.176	0.073	<b>0.167</b>
CQADupStack	0.296	0.296	0.110	<b>0.287</b>
Quora	<b>0.852</b>	0.830	0.677	<b>0.852</b>
DBPedia	0.281	<b>0.370</b>	0.236	<b>0.335</b>
SCIDOCS	0.122	<b>0.134</b>	0.058	<b>0.130</b>
FEVER	<b>0.669</b>	0.664	0.325	<b>0.571</b>
Climate-FEVER	<b>0.198</b>	0.194	0.121	<b>0.177</b>
SciFact	0.507	<b>0.622</b>	0.221	<b>0.548</b>
AVERAGE	0.389	<b>0.402</b>	0.201	<b>0.357</b>

**Table 3**

Comparison of TAS-B as the model backbone for JPQ [18] and BPR [13] with original model backbones (STAR and DPR). The JPQ (STAR) model was trained on ( $M = 96$ ) centroids with STAR [36] (RoBERTa [37]) as the backbone. The BPR (DPR) was trained with DPR [2] as the backbone. All scores denote nDCG@10, **bold** denotes the best score. The table shows that for both JPQ and BPR, TAS-B as the model backbone outperforms previously used model backbones: STAR and DPR.