APPROXIMATE NEAREST NEIGHBOR NEGATIVE CONTRASTIVE LEARNING FOR DENSE TEXT RETRIEVAL

Anonymous authors
Paper under double-blind review

ABSTRACT

Conducting text retrieval in a dense representation space has many intriguing advantages. Yet the end-to-end learned dense retrieval (DR) often underperforms word-based sparse retrieval. In this paper, we first theoretically show the learning bottleneck of dense retrieval is due to the domination of uninformative negatives sampled locally in batch, which yield diminishing gradient norms, large stochastic gradient variances, and slow learning convergence. We then propose Approximate nearest neighbor Negative Contrastive Learning (ANCE), a learning mechanism that selects hard training negatives globally from the entire corpus, using an asynchronously updated ANN index. Our experiments demonstrate the effectiveness of ANCE on web search, question answering, and in a commercial search environment, showing ANCE dot-product retrieval nearly matches the accuracy of BERT-based cascade IR pipeline, while being 100x more efficient. We also empirically validate our theory that negative sampling with ANCE better approximates the oracle gradient-norm based importance sampling, thus improves the convergence of stochastic training.

1 INTRODUCTION

Many language systems rely on text retrieval as their first step to find relevant information. For example, search ranking (Nogueira & Cho, 2019), open domain question answering (OpenQA) (Chen et al., 2017), and fact verification (Thorne et al., 2018) all first retrieve relevant documents for their later stage reranking, machine reading, and reasoning models. All these later-stage models enjoy the advancements of deep learning techniques (Rajpurkar et al., 2016; Wang et al., 2018), while, the first stage retrieval still mainly relies on matching discrete bag-of-words, e.g., BM25, which has become the bottleneck of many systems (Nogueira & Cho, 2019; Luan et al., 2020; Zhao et al., 2020).

Dense Retrieval (DR) aims to overcome the sparse retrieval bottleneck by matching texts in a continuous representation space learned via deep neural networks (Lee et al., 2019; Karpukhin et al., 2020; Luan et al., 2020). It has many desired properties: fully learnable representation, easy integration with pretraining, and efficiency support from approximate nearest neighbor (ANN) search (Johnson et al., 2019). These make dense retrieval an intriguing potential choice to fundamentally overcome some intrinsic limitations of sparse retrieval, for example, vocabulary mismatch (Croft et al., 2010).

A key challenge in dense retrieval is to construct proper negative instances when learning the representation space (Karpukhin et al., 2020). Unlike in reranking (Liu, 2009) where the training negatives are naturally the irrelevant documents from previous retrieval stages, in first stage retrieval, DR models need to distinguish relevant documents from *all irrelevant ones* in a corpus with millions or billions of documents. As illustrated in Fig. 1, these *global negatives* are quite different from negatives retrieved by sparse models.

Recent research explored various ways to construct negative training instances for dense retrieval (Huang et al., 2020; Karpukhin et al., 2020)., e.g., using contrastive learning (Faghri et al., 2017; Oord et al., 2018; He et al., 2019; Chen et al., 2020a) to select hard negatives in current or recent mini-batches. However, as observed in recent research (Karpukhin et al., 2020), the in-batch local negatives, though effective in learning word or visual representations, are not significantly better than spare-retrieved negatives in representation learning for dense retrieval. In addition, the accuracy of dense retrieval models often underperform BM25, especially on documents (Lee et al., 2019; Gao et al., 2020b; Luan et al., 2020).

In this paper, we first theoretically analyze the convergence of dense retrieval training with negative sampling. Using the variance reduction framework (Alain et al., 2015; Katharopoulos & Fleuret, 2018; Johnson & Guestrin, 2018), we show that, under conditions commonly met in dense retrieval, local in-batch negatives lead to diminishing gradient norms, resulted in high stochastic gradient variances and slow training convergence — the local negative sampling is the bottleneck of dense retrieval's effectiveness.

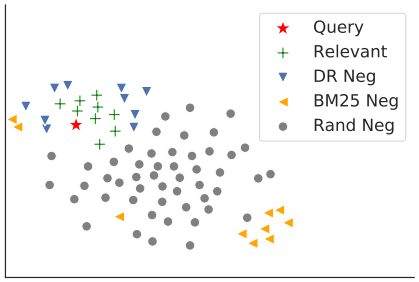

Figure 1: T-SNE (Maaten & Hinton, 2008) representations of query, relevant documents, negative training instances from BM25 (BM25 Neg) or randomly sampled (Rand Neg), and testing negatives (DR Neg) in dense retrieval.

Based on our analysis, we propose Approximate nearest neighbor Negative Contrastive Estimation (ANCE), a new contrastive representation learning mechanism for dense retrieval. Instead of random or in-batch local negatives, ANCE constructs global negatives using the being-optimized DR model to retrieve from the entire corpus. This fundamentally aligns the distribution of negative samples in training and of irrelevant documents to separate in testing. From the variance reduction point of view, these ANCE negatives lift the upper bound of per instance gradient norm, reduce the variance of the stochastic gradient estimation, and lead to faster learning convergence.

We implement ANCE using an asynchronously updated ANN index of the corpus representation. Similar to Guu et al. (2020), we maintain an Inferencer that parallelly computes the document encodings with a recent checkpoint from the being optimized DR model, and refresh the ANN index used for negative sampling once it finishes, to keep up with the model training. Our experiments demonstrate the advantage of ANCE in three text retrieval scenarios: standard web search (Craswell et al., 2020), OpenQA (Rajpurkar et al., 2016; Kwiatkowski et al., 2019), and in a commercial search engine's retrieval system. We also empirically validate our theory that the gradient norms on ANCE sampled negatives are much bigger than local negatives and thus improve the convergence of dense retrieval models.[1]

2 Preliminaries

In this section, we discuss the preliminaries of dense retrieval and its representation learning.

Task Definition: Given a query q and a corpus C, the first stage retrieval is to find a set of documents relevant to the query $D^+ = \{d_1, ..., d_i, ..., d_n\}$ from C ($|D^+| \ll |C|$), which then serve as input to later more complex models (Croft et al., 2010). Instead of using sparse term matches and inverted index, *Dense Retrieval* calculates the retrieval score $f()$ using similarities in a learned embedding space (Lee et al., 2019; Luan et al., 2020; Karpukhin et al., 2020):

$$f(q, d) = \text{sim}(g(q; \theta), g(d; \theta)), \tag{1}$$

where $g()$ is the representation model that encodes the query or document to dense embeddings. The encoder parameter θ provides the main capacity, often fine-tuned from pretrained transformers, e.g., BERT (Lee et al., 2019). The similarity function (sim()) is often simply cosine or dot product, to leverage efficient ANN retrieval (Johnson et al., 2019; Guo et al., 2020).

Learning with Negative Sampling: The effectiveness of DR resides in learning a good representation space that maps query and relevant documents together, while separating irrelevant ones. The learning of this representation often follows standard learning to rank (Liu, 2009): Given a query q, a set of relevant document D^+ and irrelevant ones D^-, to find the best θ^* that:

$$\theta^* = \text{argmin}_\theta \sum_q \sum_{d^+ \in D^+} \sum_{d^- \in D^-} l(f(q, d^+), f(q, d^-)). \tag{2}$$

The loss $l()$ can be binary cross entropy (BCE), hinge loss, or negative log likelihood (NLL).

A unique challenge in dense retrieval, targeting first stage retrieval, is that the irrelevant documents to separate are from the entire corpus ($D^- = C \setminus D^+$). This often leads to millions of negative

[1]Code and trained models are in the supplementary material and will be open-sourced.

instances, which have to be sampled in training:

$$\theta^* = \text{argmin}_\theta \sum_q \sum_{d^+ \in D^+} \sum_{d^- \in \hat{D}^-} l(f(q, d^+), f(q, d^-)). \tag{3}$$

A natural choice is to sample negatives $\hat{D}^-$ from top documents retrieved by BM25. However, they may bias the DR model to merely learn sparse retrieval and do not elevate DR models much beyond BM25 (Luan et al., 2020). Another way is to sample negatives in local mini-batches, e.g., as in contrastive learning (Oord et al., 2018; Chen et al., 2020a), however, these local negatives do not significantly outperform BM25 negatives (Karpukhin et al., 2020; Luan et al., 2020).

3 Analyses on The Convergence of Dense Retrieval Training

In this section, we provide theoretical analyses on the convergence of representation training in dense retrieval. We first show the connections between learning convergence and gradient norms, then the bounded gradient norms by uninformative negatives, and finally, how in-batch local negatives are ineffective under common conditions in dense retrieval.

Convergence Rate and Gradient Norms: Let $l(d^+, d^-) = l(f(q, d^+), f(q, d^-)$ be the loss function on the training triple (q, d^+, d^-), P_{D^-} the negative sampling distribution for the given (q, d^+), and p_{d^-} the sampling probability of negative instance d^-, a stochastic gradient decent (SGD) step with importance sampling (Alain et al., 2015) is:

$$\theta_{t+1} = \theta_t - \eta \frac{1}{N p_{d^-}} \nabla_{\theta_t} l(d^+, d^-), \tag{4}$$

with θ_t the parameter at t-th step, θ_{t+1} the one after, and N the total number of negatives. The scaling factor $\frac{1}{N p_{d^-}}$ is to make sure Eqn. 4 is an unbiased estimator of the full gradient.

Then we can characterize the converge rate of this SGD step as the movement to optimal θ^*. Following derivations in variance reduction (Katharopoulos & Fleuret, 2018; Johnson & Guestrin, 2018), let $g_{d^-} = \frac{1}{N p_{d^-}} \nabla_{\theta_t} l(d^+, d^-)$ the weighted gradient, the convergence rate is:

$$\begin{aligned}
\mathbb{E}\Delta^t &= ||\theta_t - \theta^*||^2 - \mathbb{E}_{P_{D^-}}(||\theta_{t+1} - \theta^*||^2) && (5)\\
&= ||\theta_t||^2 - 2\theta_t^T \theta^* - \mathbb{E}_{P_{D^-}}(||\theta_t - \eta g_{d^-}||^2) + 2\theta^{*T} \mathbb{E}_{P_{D^-}}(\theta_t - \eta g_{d^-}) && (6)\\
&= -\eta^2 \mathbb{E}_{P_{D^-}}(||g_{d^-}||^2) + 2\eta\theta_t^T \mathbb{E}_{P_{D^-}}(g_{d^-}) - 2\eta\theta^{*T} \mathbb{E}_{P_{D^-}}(g_{d^-}) && (7)\\
&= 2\eta \mathbb{E}_{P_{D^-}}(g_{d^-})^T (\theta_t - \theta^*) - \eta^2 \mathbb{E}_{P_{D^-}}(||g_{d^-}||^2) && (8)\\
&= 2\eta \mathbb{E}_{P_{D^-}}(g_{d^-})^T (\theta_t - \theta^*) - \eta^2 \mathbb{E}_{P_{D^-}}(g_{d^-})^T \mathbb{E}_{P_{D^-}}(g_{d^-}) - \eta^2 \text{Tr}(\mathcal{V}_{P_{D^-}}(g_{d^-})). && (9)
\end{aligned}$$

This shows we can obtain better convergence rate by sampling from a distribution P_{D^-} that minimizes the variance of the gradient estimator, $\mathbb{E}_{P_{D^-}}(||g_{d^-}||^2)$, or $\text{Tr}(\mathcal{V}_{P_{D^-}}(g_{d^-}))$ as the estimator is unbiased. There exists an optimal distribution that:

$$p^*_{d^-} = \text{argmin}_{p_{d^-}} \text{Tr}(\mathcal{V}_{P_{D^-}}(g_{d^-})) \propto ||\nabla_{\theta_t} l(d^+, d^-)||_2, \tag{10}$$

which is to sample proportionally to per instance gradient norm. This is a well known result in importance sampling (Alain et al., 2015; Johnson & Guestrin, 2018). It can be proved by applying Jensen's inequality on the gradient variance and then verifying that Eqn. 10 achieves the minimum. We do not repeat this proof and refer to Johnson & Guestrin (2018) for exact derivations.

Intuitively, an negative instance with larger gradient norm is more likely to reduce the training loss more, while those with diminishing gradients are not informative. Empirically, the correlation of gradient norm and training convergence is also observed in BERT fine-tuning (Mosbach et al., 2020).

Diminishing Gradients of Uninformative Negatives: The oracle distribution in Eqn. 10 is too expensive to compute and the closed form of gradient norms can be complicated in deep neural networks. Nevertheless, for MLP networks, Katharopoulos & Fleuret (2018) derives an upper bound of the per sample gradient norm:

$$||\nabla_{\theta_t} l(d^+, d^-)||_2 \leq L\rho ||\nabla_{\phi_L} l(d^+, d^-)||_2, \tag{11}$$

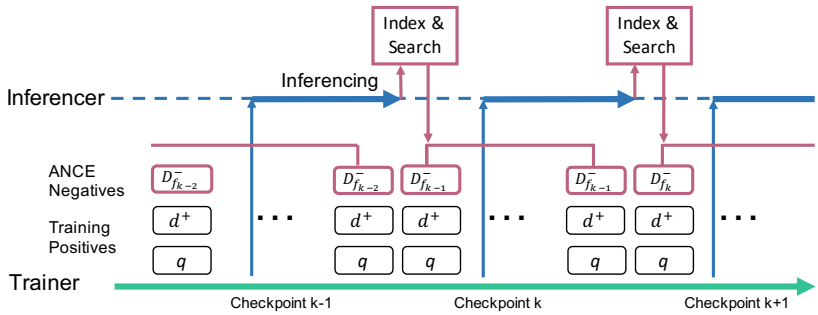

Figure 2: ANCE Asynchronous Training. The Trainer learns the representation using negatives from the ANN index. The Inferencer uses a recent checkpoint to update the representation of documents in the corpus and once finished, refreshes the ANN index with most up-to-date encodings.

where L is the number of layers, ρ is composed by pre-activation weights and gradients in intermediate layers, and $||\nabla_{\phi_L} l(d^+, d^-)||_2$ is the gradient w.r.t. the last layer. Intuitively, the intermediate layers are more regulated by various normalization techniques; the main moving piece is $||\nabla_{\phi_L} l(d^+, d^-)||_2$ (Katharopoulos & Fleuret, 2018).

For common learning to rank loss functions, for example, BCE loss and pairwise hinge loss, we can verified that (Katharopoulos & Fleuret, 2018):

$$l(d^+, d^-) \rightarrow 0 \Rightarrow ||\nabla_{\phi_L} l(d^+, d^-)||_2 \rightarrow 0 \Rightarrow ||\nabla_{\theta_t} l(d^+, d^-)||_2 \rightarrow 0. \tag{12}$$

Intuitively, negative samples with near zero loss have near zero gradients and contribute little to model convergence. The convergence of dense retrieval model training relies on the informativeness of constructed negatives.

Inefficacy of Local In-Batch Negatives: We argue that the in-batch local negatives are unlikely to provide informative samples due to two common properties of text retrieval.

Let D^{-*} be the set of informative negatives that are hard to distinguish from D^+, and b be the batch size, we have (1) $b \ll |C|$, the batch size is far smaller than the corpus size; (2) $|D^{-*}| \ll |C|$, that only a few negatives are informative and the majority of corpus is trivially unrelated.

Both conditions are easy to verify empirically in dense retrieval benchmarks. The two together make the probability that a random mini-batch includes meaningful negatives $p = \frac{b|D^{-*}|}{|C|^2}$ close to zero. Selecting negatives from local training batches is unlikely to provide optimal training signals for dense retrieval.

4 Approximate Nearest Neighbor Noise Contrastive Estimation

Our analyses show the importance, if not necessity, to construct negatives *globally* from the corpus. In this section, we propose *A*pproximate nearest neighbor *N*egative *C*ontrastive *E*stimation, (ANCE), which selects negatives from the entire corpus using an asynchronously updated ANN index.

ANCE samples negatives from the top retrieved documents via the DR model from the ANN index:

$$\theta^* = \text{argmin}_\theta \sum_q \sum_{d^+ \in D^+} \sum_{d^- \in D^-_{\text{ANCE}}} l(f(q, d^+), f(q, d^-)), \tag{13}$$

with $D^-_{\text{ANCE}} = \text{ANN}_{f(q,d)} \setminus D^+$ and $\text{ANN}_{f(q,d)}$ the top retrieved documents by $f()$ from the ANN index. By definition, D^-_{ANCE} are the hardest negatives for the current DR model: $D^-_{\text{ANCE}} \approx D^{-*}$. In theory, these more informative negatives have higher training loss, higher upper bound on the gradient norms, and will improve training convergence.

ANCE can be used to train any dense retrieval model. For simplicity, we use a simple set up in recent research (Luan et al., 2020) with BERT Siamese/Dual Encoder (shared between q and d), dot product similarity, and negative log likelihood (NLL) loss.

Asynchronous Index Refresh: During stochastic training, the DR model $f()$ is updated each mini-batch. Maintaining an update-to-date ANN index to select fresh ANCE negatives is challenging, as the index update requires two operations: 1) *Inference*: refresh the representations of all documents in the corpus with an updated DR model; 2) *Index*: rebuild the ANN index using updated representations. Although *Index* is efficient (Johnson et al., 2019), *Inference* is too expensive to compute per batch as it requires a forward pass on the entire corpus which is much bigger than the training batch.

Thus we implement an asynchronous index refresh similar to Guu et al. (2020), and update the ANN index once every m batches, i.e., with checkpoint f_k. As illustrated in Fig. 2, besides the Trainer, we run an Inferencer that takes the latest checkpoint (e.g., f_k) and recomputes the encodings of the entire corpus. In parallel, the Trainer continues its stochastic learning using $D^-_{f_{k-1}}$ from $\text{ANN}_{f_{k-1}}$. Once the corpus is re-encoded, the Inferencer rebuilds an updated ANN index (ANN_{f_k}) and feed it to the Trainer.

In this process, the ANCE negatives (D^-_{ANCE}) are asynchronously updated to "catch up" with the stochastic training. The gap between the ANN index and the DR model optimization depends on the allocation of computing resources between Trainer and Inferencer. Appendix A.3 shows an 1:1 GPU split is sufficient to minimize the influence of this gap.

5 Experimental Methodologies

This section describes our experimental setups. More details can be found in Appendix A.1 and A.2.

Benchmarks: The web search experiments are conducted on the TREC 2019 Deep Learning (DL) Track benchmark (Craswell et al., 2020), a standard large scale ad hoc retrieval dataset. We follow the official guideline and evaluate mainly in the retrieval setting, but also show the results when reranking the top 100 candidates from BM25.

The OpenQA experiments use the Natural Questions (NQ) (Kwiatkowski et al., 2019) and TriviaQA (TQA) (Joshi et al., 2017), following the settings from Karpukhin et al. (2020)[2]. The metrics are Coverage@20/100, which evaluate whether the Top-20/100 retrieved passages include the answer.

We also study the effectiveness of ANCE in the first stage retrieval of a commercial search engine's production system. We change the training of a production-quality DR model to ANCE, and evaluate the offline gains in various corpus sizes, encoding dimensions, and exact/approximate search.

Baselines: In TREC DL, we include best runs in relevant categories and refer to Craswell et al. (2020) for more baseline scores. We implement recent DR baselines that use the same BERT-Siamese, but vary in negative construction: random sampling in batch (Rand Neg), random sampling from BM25 top 100 (BM25 Neg) (Lee et al., 2019; Gao et al., 2020b) and the 1:1 combination of BM25 and Random negatives (BM25 + Rand Neg) (Karpukhin et al., 2020; Luan et al., 2020). We also compare with contrastive learning/Noise Contrastive Estimation, which uses hardest negatives in batch (NCE Neg) (Gutmann & Hyvärinen, 2010; Oord et al., 2018; Chen et al., 2020a). In OpenQA, we compare with DPR, BM25, and their combinations (Karpukhin et al., 2020).

Implementation Details: In TREC DL, recent research found MARCO passage training labels cleaner (Yan et al., 2019) and BM25 negatives can help train dense retrieval (Karpukhin et al., 2020; Luan et al., 2020). Thus, we include a "BM25 Warm Up" setting (BM25 $\rightarrow *$), where the DR models are first trained using MARCO official BM25 Negatives. ANCE is also warmed up by BM25 negatives. All DR models in TREC DL are fine-tuned from RoBERTa base (Liu et al., 2019). In OpenQA, we warm up ANCE using the released DPR checkpoints (Karpukhin et al., 2020).

To fit long documents in BERT-Siamese, ANCE uses the two settings from Dai & Callan (2019b), FirstP which uses the first 512 tokens of the document, and MaxP, where the document is split to 512-token passages (maximum 4) and the passage level scores are max-pooled. The max-pooling operation is natively supported in ANN. The ANN search uses the Faiss IndexFlatIP Index (Johnson et al., 2019). We use 1:1 Trainer:Inference GPU allocation, index refreshing per 10k training batches, batch size 8, and gradient accumulation step 2 on 4 GPUs. For each positive, we uniformly sample one negative from ANN top 200. We measured ANCE efficiency using a single 32GB V100 GPU, on a cloud VM with Intel(R) Xeon(R) Platinum 8168 CPU and 650GB of RAM memory.

[2]https://github.com/facebookresearch/DPR.

Table 1: Results in TREC 2019 Deep Learning Track. Results not available are marked as "n.a.", not applicable are marked as "–". Best results in each category are marked bold.

	MARCO Dev Passage Retrieval		TREC DL Passage NDCG@10		TREC DL Document NDCG@10	
	MRR@10	**Recall@1k**	**Rerank**	**Retrieval**	**Rerank**	**Retrieval**
Sparse & Cascade IR						
BM25	0.240	0.814	–	0.506	–	0.519
Best DeepCT	0.243	n.a.	–	n.a.	–	0.554
Best TREC Trad Retrieval	0.240	n.a.	–	0.554	–	0.549
BERT Reranker	–	–	**0.742**	–	0.646	–
Dense Retrieval						
Rand Neg	0.261	0.949	0.605	0.552	0.615	0.543
NCE Neg	0.256	0.943	0.602	0.539	0.618	0.542
BM25 Neg	0.299	0.928	0.664	0.591	0.626	0.529
DPR (BM25 + Rand Neg)	0.311	0.952	0.653	0.600	0.629	0.557
BM25 $\rightarrow$ Rand	0.280	0.948	0.609	0.576	0.637	0.566
BM25 $\rightarrow$ NCE Neg	0.279	0.942	0.608	0.571	0.638	0.564
BM25 $\rightarrow$ BM25 + Rand	0.306	0.939	0.648	0.591	0.626	0.540
ANCE (FirstP)	**0.330**	**0.959**	0.677	**0.648**	0.641	0.615
ANCE (MaxP)	–	–	–	–	**0.671**	**0.628**

Table 2: Retrieval results (Answer Coverage at Top-20/100) on Natural Questions (NQ) and Trivial QA (TQA) in the setting from Karpukhin et al. (2020).

	Single Task		Multi Task	
	NQ	**TQA**	**NQ**	**TQA**
Retriever	Top-20/100	Top-20/100	Top-20/100	Top-20/100
BM25	59.1/73.7	66.9/76.7	–/–	–/–
DPR	78.4/85.4	79.4/85.0	79.4/86.0	78.8/84.7
BM25+DPR	76.6/83.8	79.8/84.5	78.0/83.9	79.9/84.4
ANCE	**81.9/87.5**	**80.3/85.3**	**82.1/87.9**	**80.3/85.2**

Table 3: Relative gains in the first stage retrieval of a commercial search engine. The gains are from changing the training of a production DR model to ANCE.

Corpus Size	Dim	Search	Gain
250 Million	768	KNN	+18.4%
8 Billion	64	KNN	+14.2%
8 Billion	64	ANN	+15.5%

6 Evaluation Results

In this section, we first evaluate the effectiveness and efficiency of ANCE. Then we empirically study the convergence of ANCE training.

6.1 Effectiveness and Efficiency

The results on TREC 2019 DL benchmark are in Table 1. ANCE significantly outperforms all sparse retrieval, including DeepCT, which uses BERT to learn term weights (Dai et al., 2019). Among all different negative construction mechanisms, ANCE is the only one that elevates BERT-Siamese to robustly exceed the sparse methods in document retrieval. It also outperforms DPR in passage retrieval in OpenQA (Table 2). ANCE's effectiveness is even more observed in real production (Table 3), where switching to ANCE training improves a production model stably by around 15%.

Among all DR models, ANCE has the smallest gap between its retrieval and reranking accuracy, showing the importance of global negatives in training retrieval models. ANCE retrieval nearly matches the accuracy of the cascade IR with interaction-based BERT Reranker. This overthrows a previously-held belief that modeling term-level interactions is necessary in search (Xiong et al., 2017; Qiao et al., 2019). *With ANCE, we can learn a representation space that effectively captures the finesse of search relevance.*

Table 4 measures the efficiency ANCE (FirstP) in TREC DL document retrieval. The online latency

Table 4: Efficiency of Sparse and Dense IR.

Operation	Offline	Online
BM25 Index Build	3h	–
BM25 Retrieval	–	37ms
BERT Rerank	–	1.15s
Sparse IR Total (BM25 + BERT)	–	**1.42s**
ANCE Inference		
Encoding of Corpus/Per doc	10h/4.5ms	–
Query Encoding	–	2.6ms
ANN Retrieval (batched q)	–	9ms
Dense Retrieval Total	–	**11.6ms**
ANCE Training		
Encoding of Corpus/Per doc	10h/4.5ms	–
ANN Index Build	10s	–
Neg Construction Per Batch	72ms	–
Back Propagation Per Batch	19ms	–

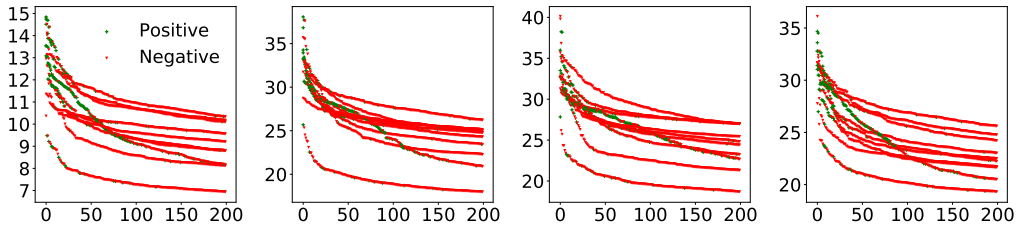

(a) ANCE FirstP (100%) (b) NCE Neg (0%) (c) Rand Neg (0%) (d) BM25+Rand (7%)

Figure 3: The top DR scores for 10 random TREC DL testing queries. The x-axes list their ranking order. The y-axes are their retrieval scores minus corpus average. All models are warmed up by BM25 Neg. The percentages are the overlaps between the testing and training negatives near convergence.

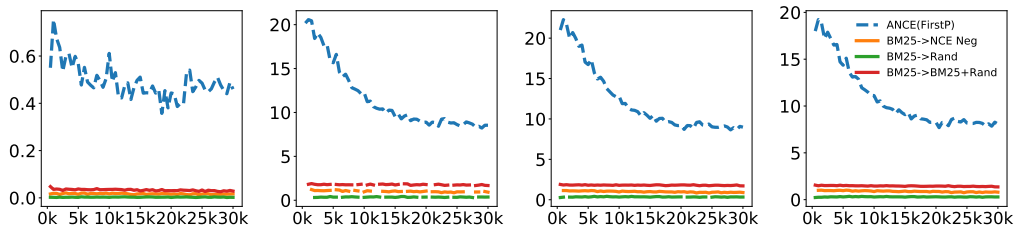

(a) Training Loss (b) Grad Norm (Bottom) (c) Grad Norm (Middle) (d) Grad Norm (Top)

Figure 4: The training loss and gradient norms with different negatives in DR training, after BM25 warm up. The gradient norms are the average of the bottom (1-4), middle (5-8), and top (9-12) BERT layers. The x-axes are training steps.

is on one query and 100 retrieved documents. DR with standard batching provides a *100x speed up compared to BERT Rerank*, a natural benefit from the Siamese network and pre-computable document encoding. In ANCE training, the bulk of computing is to update the encodings of the training corpus using new checkpoints. Assuming the model used to sample negatives and to be learned is the same, this is inevitable but can be mitigated by asynchronous index refresh.

6.2 Empirical Analyses on Training Convergence

In this experiment, we study the learning convergence of dense retrieval and empirically verify our theoretical analyses in Section 3.

We first show the long tail distribution of search relevance in dense retrieval. As plotted in Fig. 3, there are a few instances per query with significant higher retrieval scores, while the majority form a long tail. In retrieval/ranking, the key challenge is to distinguish the relevant ones among those highest scored ones; the rest is trivially irrelevant. We also empirically measure the probability of local in-batch negatives including informative negatives (D^{-*}), by their overlap with top 100 highest scored negatives. This probability, either using NCE Neg or Rand Neg, is *zero*, the same as our theory assumes. In comparison, the overlap between BM25 Neg with top DR retrieved negatives is 15%, while that of ANCE negatives starts at 63% and converges to 100% by design.

Then we empirically validate our theory that local negatives lead to lower loss, bounded gradient norm, and thus slow convergence. The training loss and pre-clip gradient norms during DR training are plotted in Fig. 4. As expected, the uninformative local negatives are trivial to separate, yielding near-zero training loss, while ANCE global negatives are much harder and maintain a high training loss. The same with our theoretical assumption, the gradient norms of local negatives are indeed restricted closely to zero. In comparison, the gradient norms on ANCE global negatives are bigger by orders of magnitude. This confirms ANCE better approximates the oracle importance sampling distribution $p^*_{d^-} \propto ||\nabla_{\theta_t} l(d^+, d^-)||_2$ and improves learning convergence.

6.3 DISCUSSIONS

We use BERT-Siamese and NLL loss to be consistent with recent research. We have experimented with cosine similarity and BCE/hinge loss. They produce more bounded retrieval scores and even smaller gradient norms on local negatives. But the retrieval accuracy is not much better. We also include some additional experiments in Appendix. Appendix A.2 discusses the surprisingly small overlap (<25%) between dense retrieval results and sparse retrieval results. DR is a fundamentally different approach and more studies are required to understand its behavior. Appendix A.3 and A.4 provide ablation studies on the asynchronous gaps and hyperparameters. Appendix A.5 includes case studies where we observe that the irrelevant documents from ANCE are often still "semantically related", unlike in sparse retrieval where errors are more because of matched query terms in unrelated documents.

7 RELATED WORK

In early research on neural information retrieval (Neu-IR) (Mitra et al., 2018), a common belief was that the interaction models, those that specifically handle term level matches, are more effective though more expensive (Guo et al., 2016; Xiong et al., 2017; Nogueira & Cho, 2019). Many techniques are developed to reduce their cost, for example, distillation (Gao et al., 2020a) and caching (Humeau et al., 2020; Khattab & Zaharia, 2020; MacAvaney et al., 2020). ANCE shows that a properly trained representation-based BERT-Siamese is in fact as effective as the interaction-based BERT ranker. This finding will motivate many new research explorations in Neu-IR.

Deep learning has been used to improve various components of sparse retrieval, for example, term weighting (Dai & Callan, 2019b), query expansion (Zheng et al., 2020), and document expansion (Nogueira et al., 2019). Dense Retrieval chooses a different path and conducts retrieval purely in the embedding space via ANN search (Lee et al., 2019; Chang et al., 2020; Karpukhin et al., 2020; Luan et al., 2020). This work demonstrates that a simple dense retrieval system can achieve SOTA accuracy, while also behaves dramatically different from classic retrieval. The recent advancement in dense retrieval may raise a new generation of search systems.

Recent research in contrastive representation learning also shows the benefits of sampling negatives from a larger candidate pool. In computer vision, He et al. (2019) decouple the negative sampling pool size with training batch size, by maintaining a negative candidate pool of recent batches and updating their representation with momentum. This enlarged negative pool significantly improves unsupervised visual representation learning (Chen et al., 2020b). A parallel work (Xiong et al., 2020) improves DPR by sampling negatives from a memory bank (Wu et al., 2018) — in which the representations of negative candidates are frozen so more candidates can be stored. Instead of a bigger local pool, ANCE goes all the way along this trajectory and constructs negatives globally from the entire corpus, using an asynchronously updated ANN index.

Besides being a real world application itself, dense retrieval is also a core component in many other language systems, for example, to retrieval relevant information for grounded language models (Khandelwal et al., 2019; Guu et al., 2020), extractive/generative QA (Karpukhin et al., 2020; Lewis et al., 2020b), and fact verification (Xiong et al., 2020), or to find paraphrase pairs for pretraining (Lewis et al., 2020a). There dense retrieval models are either frozen or optimized indirectly by signals from their end tasks. ANCE is orthogonal with those lines of research and focuses on the representation learning for dense retrieval. Its better retrieval accuracy can benefit many language systems.

8 CONCLUSION

In this paper, we first provide theoretical analyses on the convergence of representation learning in dense retrieval. We show that under common conditions in text retrieval, the local negatives used in DR training are uninformative, yield low gradient norms, and contribute little to the learning convergence. We then propose ANCE to eliminate this bottleneck by constructing training negatives globally from the entire corpus. Our experiments demonstrate the advantage of ANCE in web search, OpenQA, and the production system of a commercial search engine. Our studies empirically validate our theory that ANCE negatives have much bigger gradient norms, reduce the stochastic gradient variance, and improve training convergence.

REFERENCES

Guillaume Alain, Alex Lamb, Chinnadhurai Sankar, Aaron Courville, and Yoshua Bengio. Variance reduction in sgd by distributed importance sampling. *arXiv preprint arXiv:1511.06481*, 2015.

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

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

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading wikipedia to answer open-oomain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pp. 1870–1879, 2017.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020a.

Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.

Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. Overview of the trec 2019 deep learning track. In *Text REtrieval Conference (TREC)*. TREC, 2020.

W Bruce Croft, Donald Metzler, and Trevor Strohman. *Search engines: information retrieval in practice*, volume 520. Addison-Wesley Reading, 2010.

Zhuyun Dai and Jamie Callan. Context-aware sentence/passage term importance estimation for first stage retrieval. *arXiv preprint arXiv:1910.10687*, 2019a.

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

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. Transformer-XL: attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2978–2988, 2019.

Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. *arXiv preprint arXiv:1707.05612*, 2017.

Luyu Gao, Zhuyun Dai, and Jamie Callan. Understanding bert rankers under distillation. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval*, pp. 149–152, 2020a.

Luyu Gao, Zhuyun Dai, Zhen Fan, and Jamie Callan. Complementing lexical retrieval with semantic residual embedding. *arXiv preprint arXiv:2004.13969*, 2020b.

Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pp. 55–64, 2016.

Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. Accelerating large-scale inference with anisotropic vector quantization. *arXiv preprint arXiv:1908.10396*, 2020.

Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: a new estimation principle for unnormalized statistical models. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, pp. 297–304, 2010.

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

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.

Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2553–2561, 2020.

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

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

Tyler B Johnson and Carlos Guestrin. Training deep models faster with robust, approximate importance sampling. In *Advances in Neural Information Processing Systems*, pp. 7265–7275, 2018.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: a large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pp. 1601–1611, 2017.

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

Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. *arXiv preprint arXiv:1803.00942*, 2018.

Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*, 2019.

Omar Khattab and Matei Zaharia. Colbert: Efficient and effective passage search via contextualized late interaction over bert. *arXiv preprint arXiv:2004.12832*, 2020.

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

Victor Lavrenko and W Bruce Croft. Relevance-based language models. In *Association for Computing Machinery (ACM) Special Interest Group on Information Retrieval (SIGIR) Forum*, volume 51, pp. 260–267. ACM New York, NY, USA, 2017.

Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. Latent retrieval for weakly supervised open domain question answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 6086–6096, 2019.

Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. Pre-training via paraphrasing. *arXiv preprint arXiv:2006.15020*, 2020a.

Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*, 2020b.

Tie-Yan Liu. Learning to rank for information retrieval. *Foundations and trends in information retrieval*, 3(3): 225–331, 2009.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: a robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Yi Luan, Jacob Eisenstein, Kristina Toutanove, and Michael Collins. Sparse, dense, and attentional representations for text retrieval. *arXiv preprint arXiv:2005.00181*, 2020.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonellotto, Nazli Goharian, and Ophir Frieder. Efficient document re-ranking for transformers by precomputing term representations. *arXiv preprint arXiv:2004.14255*, 2020.

Bhaskar Mitra, Nick Craswell, et al. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 13(1):1–126, 2018.

Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. *arXiv preprint arXiv:2006.04884*, 2020.

Rodrigo Nogueira and Kyunghyun Cho. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085*, 2019.

Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. Document expansion by query prediction. *arXiv preprint arXiv:1904.08375*, 2019.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

Yifan Qiao, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. Understanding the behaviors of bert in ranking. *arXiv preprint arXiv:1904.07531*, 2019.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, 2016.

James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. The fact extraction and verification (FEVER) shared task. In *Proceedings of the 1st Workshop on Fact Extraction and VERification (FEVER)*, pp. 1–9, 2018.

Ellen M Voorhees. Variations in relevance judgments and the measurement of retrieval effectiveness. *Information Processing & Management*, 36(5):697–716, 2000.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: a multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.

Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. *arXiv preprint arXiv:1805.01978*, 2018.

Chenyan Xiong, Zhuyun Dai, Jamie Callan, Zhiyuan Liu, and Russell Power. End-to-end neural ad-hoc ranking with kernel pooling. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 55–64, 2017.

Wenhan Xiong, Xiang Lorraine Li, Srini Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Wen-tau Yih, Sebastian Riedel, Douwe Kiela, and Barlas Oğuz. Answering complex open-domain questions with multi-hop dense retrieval. *arXiv preprint arXiv:2009.12756*, 2020.

Ming Yan, Chenliang Li, Chen Wu, Bin Bi, Wei Wang, Jiangnan Xia, and Luo Si. Idst at trec 2019 deep learning track: Deep cascade ranking with generation-based document expansion and pre-trained language modeling. In *Text REtrieval Conference*. TREC, 2019.

Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul Bennett, and Saurabh Tiwary. Transformer-xh: multi-evidence reasoning with extra hop attention. In *International Conference on Learning Representations*, 2020.

Zhi Zheng, Kai Hui, Ben He, Xianpei Han, Le Sun, and Andrew Yates. Bert-qe: Contextualized query expansion for document re-ranking. *arXiv preprint arXiv:2009.07258*, 2020.

Table 5: Coverage of TREC 2019 DL Track labels on Dense Retrieval methods. Overlap with BM25 is calculated on top 100 retrieved documents.

	TREC DL Passage			TREC DL Document		
Method	**Recall@1K**	**Hole@10**	**Overlap w. BM25**	**Recall@100**	**Hole@10**	**Overlap w. BM25**
BM25	0.685	5.9%	100%	0.387	0.2%	100%
BM25 Neg	0.569	25.8%	11.9%	0.217	28.1%	17.9%
BM25 + Rand Neg	0.662	20.2%	16.4%	0.240	21.4%	21.0%
ANCE (FirstP)	0.661	14.8%	17.4%	0.266	13.3%	24.4%
ANCE (MaxP)	-	-	-	0.286	11.9%	24.9%

A APPENDIX

A.1 MORE EXPERIMENTAL DETAILS

More Details on TREC Deep Learning Benchmarks: There are two tasks in the Track: document retrieval and passage retrieval. The training and development sets are from MS MARCO, which includes passage level relevance labels for one million Bing queries (Bajaj et al., 2016). The document corpus was post-constructed by back-filling the body texts of the passage's URLs and their labels were inherited from its passages (Craswell et al., 2020). The testing sets are labeled by NIST accessors on the top 10 ranked results from past Track participants (Craswell et al., 2020).

TREC DL official metrics include NDCG@10 on test and MRR@10 on MARCO Passage Dev. MARCO Document Dev is noisy and the recall on the DL Track testing is less meaningful due to low label coverage on DR results. There is a two-year gap between the construction of the passage training data and the back-filling of their full document content. Some original documents were no longer available. There is also a decent amount of content changes in those documents during the two-year gap, and many no longer contain the passages. This back-filling perhaps is the reason why many Track participants found the passage training data is more effective than the inherited document labels. Note that the TREC testing labels are not influenced as the annotators were provided the same document contents when judging.

All the TREC DL runs are trained using these training data. Their inference results on the testing queries of the document and the passage retrieval tasks were evaluated by NIST assessors in the standard TREC-style pooling technique (Voorhees, 2000). The pooling depth is set to 10, that is, the top 10 ranked results from all participated runs are evaluated, and these evaluated labels are released as the official TREC DL benchmarks for passage and document retrieval tasks.

More Details on Baselines: The most representative sparse retrieval baselines in TREC DL include the standard BM25 ("bm25base" or "bm25base_p"), Best TREC Sparse Retrieval ("bm25tuned_rm3" or "bm25tuned_prf_p") with tuned query expansion (Lavrenko & Croft, 2017), and Best DeepCT ("dct_tp_bm25e2", doc only), which uses BERT to estimate the term importance for BM25 (Dai & Callan, 2019a). These three runs represent the standard sparse retrieval, best classical sparse retrieval, and the recent progress of using BERT to improve sparse retrieval. We also include the standard cascade retrieval-and-reranking systems BERT Reranker ("bm25exp_marcomb" or "p_exp_rm3_bert"), which is the best run using standard BERT on top of query/doc expansion, from the groups with multiple top MARCO runs (Nogueira & Cho, 2019; Nogueira et al., 2019).

BERT-Siamese Configurations: We follow the network configurations in Luan et al. (2020) in all Dense Retrieval methods, which we found provides the most stable results. More specifically, we initialize the BERT-Siamese model with RoBERTa base (Liu et al., 2019) and add a 768×768 projection layer on top of the last layer's "[CLS]" token, followed by a layer norm.

Implementation Details: The training often takes about 1-2 hours per ANCE epoch, which is whenever new ANCE negative is ready, it immediately replaces existing negatives in training, without waiting. It converges in about 10 epochs, similar to other DR baselines. The optimization uses LAMB optimizer, learning rate 5e-6 for document and 1e-6 for passage retrieval, and linear warm-up and decay after 5000 steps. More detailed hyperparameter settings can be found in our code release.

A.2 OVERLAP WITH SPARSE RETRIEVAL IN TREC 2019 DL TRACK

As a nature of TREC-style pooling evaluation, only those ranked in the top 10 by the 2019 TREC participating systems were labeled. As a result, documents not in the pool and thus not labeled are all considered irrelevant, even though there may be relevant ones among them. When reusing TREC style relevance labels, it is very important to keep track of the "hole rate" on the evaluated systems, i.e., the fraction of the top K ranked results without TREC labels (not in the pool). A larger hole rate shows that the evaluated methods are very different

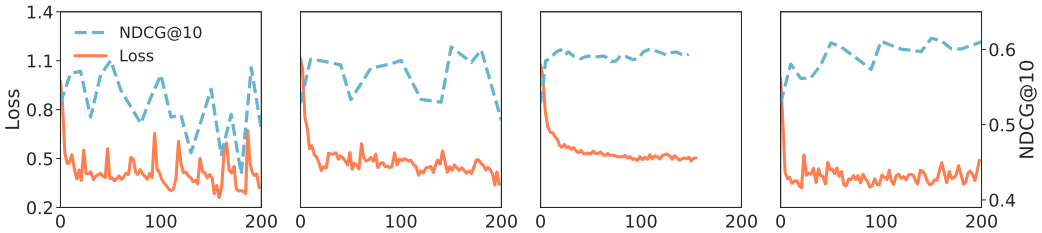

(a) 10k Batch; 4:4; 1e-5 (b) 20k Batch; 8:4; 1e-6 (c) 5k Batch; 4:8; 1e-6 (d) 10k Batch; 4:4; 5e-6

Figure 5: Training loss and testing NDCG of ANCE (FirstP) on documents, with different ANN index refreshing (e.g., per 10k Batch), Trainer:Inferencer GPU allocation, and learning rate (e.g., 1e-5). X-axes is the training steps in thousands.

Table 6: Results of several different hyperparameter configurations. "Top K Neg" lists the top k ANN retrieved candidates from which we sampled the ANCE negatives from.

	Hyperparameter			MARCO Dev Passage	TREC DL Document
	Learning rate	**Top K Neg**	**Refresh (step)**	**Retrieval MRR@10**	**Retrieval NDCG@10**
Passage ANCE	1e-6	200	10k	**0.33**	–
	1e-6	500	10k	0.31	–
	2e-6	200	10k	0.29	–
	2e-7	500	20k	0.303	–
	2e-7	1000	20k	0.302	–
Document ANCE	1e-5	100	10k	–	0.58
	1e-6	100	20k	–	0.59
	1e-6	100	5k	–	0.60
	5e-6	200	10k	–	**0.614**
	1e-6	200	10k	–	0.61

from those systems that participated in the Track and contributed to the pool, thus the evaluation results are not perfect. Note that the hole rate does not necessarily reflect the accuracy of the system, only the difference of it.

In TREC 2019 Deep Learning Track, all the participating systems are based on sparse retrieval. Dense retrieval methods often differ considerably from sparse retrievals and in general will retrieve many new documents. This is confirmed in Table 5. All DR methods have very low overlap with the official BM25 in their top 100 retrieved documents. At most, only 25% of documents retrieved by DR are also retrieved by BM25. This makes the hole rate quite high and the recall metric not very informative. It also suggests that DR methods might benefit more in this year's TREC 2020 Deep Learning Track if participants are contributing DR based systems.

The MS MARCO ranking labels were not constructed based on pooling the sparse retrieval results. They were from Bing (Bajaj et al., 2016), which uses many signals beyond term overlap. This makes the recall metric in MS MARCO more robust as it reflects how a single model can recover a complex online system.

A.3 Impact of Asynchronous Gap

Fig. 5 illustrates the behavior of asynchronous learning with different configurations. A large learning rate or a low refreshing rate (Figure 5(a) and 5(b)) leads to fluctuations as the async gap of the ANN index may drive the representation learning to undesired local optima. Refreshing as often as every 5k Batches yields a smooth convergence (Figure 5(c)), but requires twice as many GPU allocated to the Inferencer. A 1:1 GPUs allocation of Trainer and Inference with appropriate learning rates is adequate to minimize the impact of async gap.

A.4 Hyperparameter Studies

We show the results of some hyperparameter configurations in Table 6. The cost of training with BERT makes it difficult to conduct a lot hyperparameter exploration. Often a failed configuration leads to divergence early in training. We barely explore other configurations due to the time-consuming nature of working with pretrained language models. Our DR model architecture is kept consistent with recent parallel work and the learning configurations in Table 6 are about all the explorations we did. Most of the hyperparameter choices are decided solely using the training loss curve and otherwise by the loss in the MARCO Dev set. We found the training loss, validation NDCG, and testing performance align well in our (limited) hyperparameter explorations.

Table 7: Queries in the TREC 2019 DL Track Document Ranking Tasks where ANCE performs better than BM25. Snippets are manually extracted. The documents in the first disagreed ranking position are shown, where on all examples ANCE won. The NDCG@10 of ANCE and BM25 in the corresponding query is listed.

	ANCE	**BM25**
Query:	qid (104861): Cost of interior concrete flooring	
Title:	Concrete network: Concrete Floor Cost	Pinterest: Types of Flooring
DocNo:	D293855	D2692315
Snippet:	For a concrete floor with a basic finish, you can expect to pay $2 to $12 per square foot...	Know About Hardwood Flooring And Its Types White Oak Floors Oak Flooring Laminate Flooring In Bathroom ...
Ranking Position:	1	1
TREC Label:	3 (Very Relevant)	0 (Irrelevant)
NDCG@10:	0.86	0.15
Query:	qid (833860): What is the most popular food in Switzerland	
Title:	Wikipedia: Swiss cuisine	Answers.com: Most popular traditional food dishes of Mexico
DocNo:	D1927155	D3192888
Snippet:	Swiss cuisine bears witness to many regional influences, ... Switzerland was historically a country of farmers, so traditional Swiss dishes tend not to be...	One of the most popular traditional Mexican deserts is a spongy cake ... (in the related questions section) What is the most popular food dish in Switzerland?...
Ranking Position:	1	1
TREC Label:	3 (Very Relevant)	0 (Irrelevant)
NDCG@10:	0.90	0.14
Query:	qid (1106007): Define visceral	
Title:	Vocabulary.com: Visceral	Quizlet.com: A&P EX3 autonomic 9-10
DocNo:	D542828	D830758
Snippet:	When something's visceral, you feel it in your guts. A visceral feeling is intuitive — there might not be a rational explanation, but you feel that you know what's best...	Acetylcholine A neurotransmitter liberated by many peripheral nervous system neurons and some central nervous system neurons...
Ranking Position:	1	1
TREC Label:	3 (Very Relevant)	0 (Irrelevant)
NDCG@10:	0.80	0.14

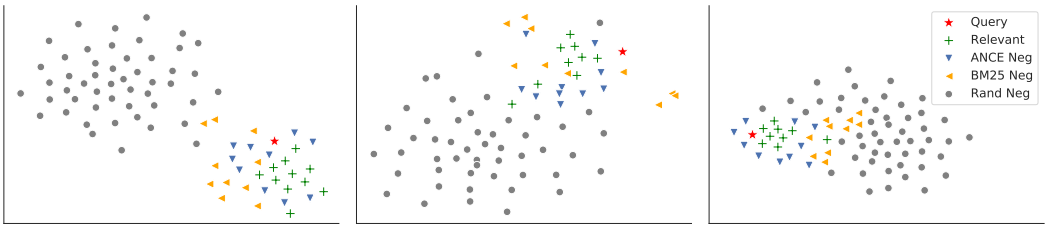

(a) 104861: interior flooring cost. (b) 833860: popular Swiss food (c) 1106007: define visceral

Figure 6: t-SNE Plots for Winning Cases in Table 7.

A.5 Case Studies

In this section, we show Win/Loss case studies between ANCE and BM25. Among the 43 TREC 2019 DL Track evaluation queries in the document task, ANCE outperforms BM25 on 29 queries, loses on 13 queries, and ties on the rest 1 query. The winning examples are shown in Table 7 and the losing ones are in Table 8. Their corresponding ANCE-learned (FirstP) representations are illustrated by t-SNE in Fig. 6 and Fig. 7.

In general, we found ANCE better captures the semantics in the documents and their relevance to the query. The winning cases show the intrinsic limitations of sparse retrieval. For example, BM25 exact matches the "most

Table 8: Queries in the TREC 2019 DL Track Document Ranking Tasks where ANCE performs worse than BM25. Snippets are manually extracted. The documents in the first position where BM25 wins are shown. The NDCG@10 of ANCE and BM25 in the corresponding query is listed. Typos in the query are from the real web search queries in TREC.

	ANCE	**BM25**
Query:	qid (182539): Example of monotonic function	
Title:	Wikipedia: Monotonic function	Explain Extended: Things SQL needs: sargability of monotonic functions
DocNo:	D510209	D175960
Snippet:	In mathematics, a monotonic function (or monotone function) is a function between ordered sets that preserves or reverses the given order... For example, if y=g(x) is strictly monotonic on the range [a,b] . . .	I'm going to write a series of articles about the things SQL needs to work faster and more efficienly. . .
Ranking Position:	1	1
TREC Label:	0 (Irrelevant)	2 (Relevant)
NDCG@10:	0.25	0.61
Query:	qid (1117099): What is a active margin	
Title:	Wikipedia: Margin (finance)	Yahoo Answer: What is the difference between passive and active continental margins
DocNo:	D166625	D2907204
Snippet:	In finance, margin is collateral that the holder of a financial instrument . . .	An active continental margin is found on the leading edge of the continent where . . .
Ranking Position:	2	2
TREC Label:	0 (Irrelevant)	3 (Very Relevant)
NDCG@10:	0.44	0.74
Query:	qid (1132213): How long to hold bow in yoga	
Title:	Yahoo Answer: How long should you hold a yoga pose for	yogaoutlet.com: How to do bow pose in yoga
DocNo:	D3043610	D3378723
Snippet:	so i've been doing yoga for a few weeks now and already notice that my flexiablity has increased drastically. . . . That depends on the posture itself . . .	Bow Pose is an intermediate yoga backbend that deeply opens the chest and the front of the body. . . Hold for up to 30 seconds . . .
Ranking Position:	3	3
TREC Label:	0 (Irrelevant)	3 (Very Relevant)
NDCG@10:	0.66	0.74

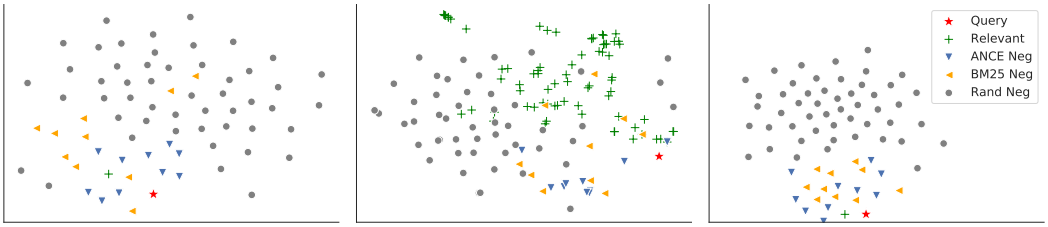

(a) 182539: monotonic function (b) 1117099: active margin (c) 1132213: yoga bow

Figure 7: t-SNE Plots for Losing Cases in Table 8.

popular food" in the query "what is the most popular food in Switzerland" but using the document is about Mexico. The term "Switzerland" only appears in the related question section of the web page.

The losing cases in Table 8 are also quite interesting. Many times we found that it is not that DR fails completely and retrieves documents not related to the query's information needs at all, which was a big concern when we

started research in DR. The errors ANCE made include retrieving documents that are related just not exactly relevant to the query, for example, "yoga pose" for "bow in yoga". In other cases, ANCE retrieved wrong documents due to the lack of the domain knowledge: the pretrained language model may not know "active margin" is a geographical terminology, not a financial one (which we did not know ourselves and took some time to figure out when conducting this case study). There are also some cases where the dense retrieved documents make sense to us but were labeled irrelevant.

The t-SNE plots in Fig. 6 and Fig. 7 show many interesting patterns of the learned representation space. The ANCE winning cases often correspond to clear separations of different document groups. For losing cases the representation space is more mixed, or there is too few relevant documents which may cause the variances in model performances. There are also many different interesting patterns in the ANCE-learned representation space. We include the t-SNE plots for all 43 TREC DL Track queries in the supplementary material. More future analyses of the learned patterns in the representation space may help provide more insights on dense retrieval.