

Progressively Pretrained Dense Corpus Index for Open-Domain Question Answering

Wenhan Xiong*, Hong Wang*, William Yang Wang

University of California, Santa Barbara

{xwhan, hongwang600, william}@cs.ucsb.edu

Abstract

Commonly used information retrieval methods such as TF-IDF in open-domain question answering (QA) systems are insufficient to capture deep semantic matching that goes beyond lexical overlaps. Some recent studies consider the retrieval process as maximum inner product search (MIPS) using dense question and paragraph representations, achieving promising results on several information-seeking QA datasets. However, the pretraining of the dense vector representations is highly resource-demanding, *e.g.*, requires a very large batch size and lots of training steps. In this work, we propose a sample-efficient method to pretrain the paragraph encoder. First, instead of using heuristically created pseudo question-paragraph pairs for pretraining, we use an existing pretrained sequence-to-sequence model to build a strong question generator that creates high-quality pretraining data. Second, we propose a simple progressive pretraining algorithm to ensure the existence of effective negative samples in each batch. Across three open-domain QA datasets, our method consistently outperforms a strong dense retrieval baseline that uses 6 times more computation for training. On two of the datasets, our method achieves more than 4-point absolute improvement in terms of answer exact match.

1 Introduction

With the promise of making the vast amount of information buried in text easily accessible via user-friendly natural language queries, the area of open-domain QA has attracted lots of attention in recent years. Existing open-domain QA systems are typically made of two essential components (Chen et al., 2017). A *retrieval module* first retrieves a compact set of paragraphs from the whole corpus

(such as Wikipedia) that includes millions of documents. Then a *reading module* is deployed to extract an answer span from the retrieved paragraphs.

Over the past few years, much of the progress in open-domain QA has been focusing on improving the reading module of the system, which only needs to process a small number of retrieved paragraphs. Specifically, improvements include stronger reading comprehension models (Wang et al., 2018b; Yang et al., 2019; Xiong et al., 2020; Min et al., 2019a) and paragraph reranking models (Wang et al., 2018a; Lin et al., 2018) that assign more accurate relevance scores to the retrieved paragraphs. However, the performance is still bounded by the retrieval modules, which simply rely on traditional IR methods such as TF-IDF or BM25 (Robertson and Zaragoza, 2009). These methods retrieve text solely based on n-gram lexical overlap and can fail on cases when deep semantic matching is required and when there are no common lexicons between the question and the target paragraph.

While neural models have proven effective at learning deep semantic matching between text pairs (Bowman et al., 2015; Parikh et al., 2016; Chen et al., 2018; Devlin et al., 2019), they usually require computing question-dependent paragraph encodings (*i.e.*, the same paragraph will have different representations when considering different questions), which is formidable considering space constraints and retrieval efficiency in practice. More recent studies (Lee et al., 2019; Chang et al., 2020; Guu et al., 2020) show that such a dilemma can be resolved with large-scale matching-oriented pretraining. These approaches use separate encoders for questions and paragraphs and simply model the matching between the question and paragraph using inner products of the output vectors. Thus, these systems only need to encode all paragraphs in a question-agnostic fashion, and the resulted dense corpus index could be fixed and reused

* Equal Contribution.

- Our code is available at <https://github.com/xwhan/ProQA.git>.

for all possible questions. While achieving significant improvements over the BM25 baseline across a set of information-seeking QA datasets, existing pretraining strategies are highly sample-inefficient and typically require a large batch size (up to thousands), such that diverse and effective negative question-paragraph pairs could be included in each batch. When using a small batch size in our experiments, the model ceases to improve after certain updates. Given that a 12G GPU can only store around 10 samples with the BERT-base architecture at training time, the wider usage of these methods to corpora with different domains (*e.g.*, non-encyclopedic web documents or scientific publications) is hindered given modest GPU hardware.

In this work, we propose a simple and sample-efficient method for pretraining dense corpus representations. We achieve stronger open-domain QA performance compared to an existing method (Lee et al., 2019) that requires 6 times more computation at training time. Besides, our method uses a much smaller batch size and can be implemented with only a small number of GPUs, *i.e.*, we use at most 4 TITAN RTX GPUs for all our experiments. In a nutshell, the proposed method first uses a pretrained sequence-to-sequence model to generate high-quality pretraining data instead of relying on heuristics to create pseudo question-paragraph pairs; for the training algorithm, we use clustering techniques to get effective negative samples for each pair and progressively update the clusters. Our method’s efficacy is further validated through ablation studies, where we replicate existing methods that use the same amount of resources. For the downstream QA experiments, we carefully investigate different finetuning objectives and show the different configurations of the retrieval and span prediction losses have nontrivial effects on the final performance. We hope this analysis could save the efforts on trying out various finetuning strategies of future research that focus on improving the retrieval component of open-domain QA systems.

The main contributions of this work include:

- We show the possibility of pretraining an effective dense corpus index for open-domain QA with modest computation resources.
- Our data generation strategy demonstrates that pretrained language models are not only useful as plug-and-play contextual feature extractors: they could also be used as high-quality data generators for other pretraining tasks.

- We propose a clustering-based progressive training paradigm that improves the sample-efficiency of dense retrieval pretraining and can be easily incorporated into existing methods.

2 Framework

We begin by introducing the network architectures used in our retrieval and reading comprehension models. Next, we present how to generate high-quality question-paragraph pairs for pretraining and how we progressively train the retrieval model with effective negative instances. Finally, we show how to finetune the whole system for QA.

2.1 Model Architectures

Notations We introduce the following notations which will be used through our paper. The goal of open-domain QA is to find the answer derivation (p, s) from a large text corpus \mathcal{C} given a question q , where p is an evidence paragraph and s is a text span within p . The start and end token of s are denoted as $\text{START}(s)$ and $\text{END}(s)$ respectively. We refer the retrieval module as $P_\theta(p|q)$, with learnable parameters θ . Similarly, we refer the reading comprehension module as $P_\phi(s|p, q)$, which can be decomposed as $P_\phi(\text{START}(s)|p, q) \times P_\phi(\text{END}(s)|p, q)$. We use D^k to represent the top-k paragraphs from the retrieval module; a subset of $D^* \in D^k$ represents the paragraphs in D^k that cover the correct answer; for each paragraph $p \in D^*$, we define S_p^* as all the spans in p that match the answer string.

The Retrieval Module We use two isomorphic encoders to encode the questions and paragraphs, and the inner product of the output vectors is used as the matching score. The encoders are based on the BERT-base architecture. We add linear layers $\mathbf{W}_q \in \mathbb{R}^{768 \times 128}$ and $\mathbf{W}_p \in \mathbb{R}^{768 \times 128}$ above the final representations of the [CLS] token to derive the question and paragraph representations:

$$\begin{aligned} h_q &= \mathbf{W}_q \text{BERT}_Q(q)([\text{CLS}]) \\ h_p &= \mathbf{W}_p \text{BERT}_P(p)([\text{CLS}]), \end{aligned}$$

The matching score is modeled as $h_q^\top h_p$. Thus, the probability of selecting p given q is calculated as:

$$P_\theta(p|q) = \frac{e^{h_q^\top h_p}}{\sum_{p' \in \mathcal{C}} e^{h_q^\top h_{p'}}}.$$

In practice, we only consider the top-k retrieved paragraphs \mathcal{C} for normalization.

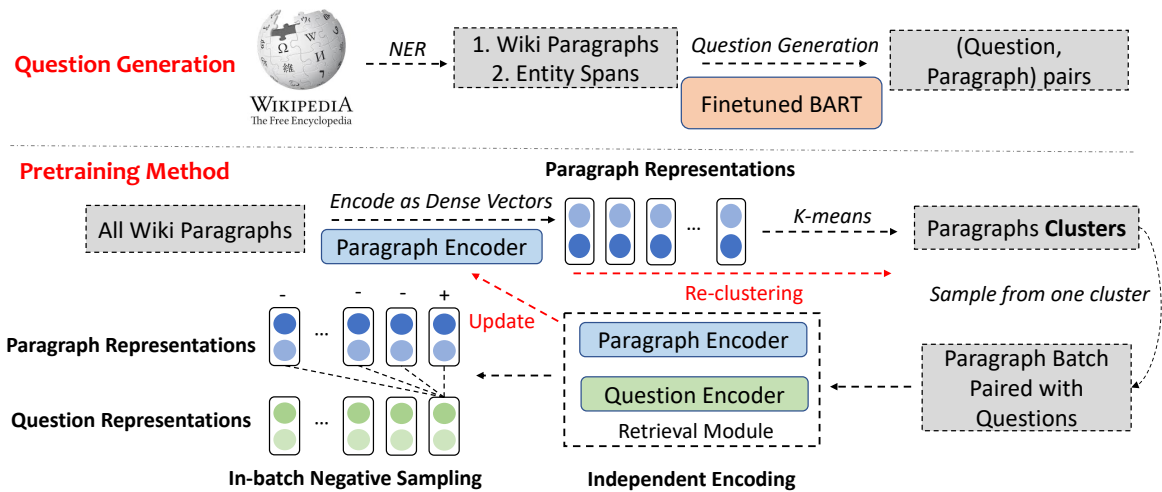


Figure 1: An overview of the progressive pretraining approach.

The Reading Module The architecture of our reading comprehension model is identical to the one in the original BERT paper (Devlin et al., 2019). We use two independent linear layers to predict the start and end position of the answer span. At training time, when calculating the span probabilities, we apply the shared-normalization technique proposed by Clark and Gardner (2018), which normalizes the probability across all the top-k retrieved paragraphs. This encourages the model to produce globally comparable answer scores. We denote this probability as $P_{\phi}^{sn}(s|p, q)$ in contrast to the original formulation $P_{\phi}(s|p, q)$ that normalizes the probability within each paragraph.

2.2 The Pretraining Method

We now describe how to pretrain the retrieval module using a better data generation strategy and a progressive training paradigm. Figure 1 depicts the whole pretraining process.

Pretraining Data Generation Previous dense retrieval approaches usually rely on simple heuristics to generate synthetic matching pairs for pretraining, which do not necessarily reflect the underlying matching pattern between questions and paragraphs. To minimize the gap between pretraining and the end task, we learn to generate high-quality questions from the paragraphs using a state-of-the-art pretrained seq2seq model, *i.e.*, BART (Lewis et al., 2019). More specifically, we finetune BART on the original NaturalQuestions dataset (Kwiatkowski et al., 2019) such that it learns to generate questions given the groundtruth answer string and the groundtruth paragraph (la-

beled as *long answer* in NaturalQuestions). We concatenate the paragraph and the answer string with a separating token as the input to the BART model. We find this simple input scheme is effective enough to generate high-quality questions, achieving a 55.6 ROUGE-L score on the dev set. Samples of the generated questions could be found in the appendix. Afterward, we use spaCy¹ to recognize potential answer spans (named entities or dates) in all paragraphs in the corpus and use the finetuned BART model to generate the questions conditioned on the paragraph and each of the potential answers.

It is worth noting that the groundtruth answer paragraph supervision at this step could be eventually dropped and we could just use weakly supervised paragraphs to train the question generator; thus, our system becomes fully weak-supervised. As the pretraining process takes a long training time and lots of resources, we are unable to repeat the whole pretraining process with the weakly-supervised question generator. However, additional question generation experiments suggest that while using weakly-supervised paragraphs, the question generator still generates high-quality questions, achieving a ROUGE-L score of 49.6.²

In-batch Negative Sampling To save computation and improve the sample efficiency of pretraining, we choose to use in-batch negative sampling (Logeswaran and Lee, 2018) instead of gathering negative paragraphs for each question to pre-

¹<https://spacy.io>

²For reference, a state-of-the-art QG model (Ma et al., 2020) trained with strong supervision achieves 49.9 ROUGE-L on a similar QA dataset also collected from real-user queries.

train the retrieval module. Specifically, for each pair (q, p) within a batch B , the paragraphs paired with other questions are considered negative paragraphs for q . Thus, the pretraining objective for each generated question is to minimize the negative log-likelihood of selecting the correct p among all paragraphs in the batch:

$$\mathcal{L}_{pre} = -\log P_{\theta}(p|q). \quad (1)$$

A graphic illustration of this strategy is shown in Figure 1. As the batch size is usually very small compared to the number of all the paragraphs in the corpus, the pretraining task is much easier compared to the final retrieval task at inference time. In the whole corpus, there are usually lots of similar paragraphs and these paragraphs could act as strong distractors for each other in terms of both paragraph ranking and answer extraction. The desired retrieval model should be able to learn fine-grained matching instead of just learning to distinguish obviously different paragraphs. However, since existing dense retrieval methods typically use uniform batch sampling, there could be many easy negative samples in each batch, and they can only provide weak learning signals. Thus, a large batch size is usually adopted to include sufficient effective negative samples. Unfortunately, this is generally not applicable without hundreds of GPUs.

The Progressive Training Paradigm To provide effective negative samples, we propose a progressive training algorithm, as shown in the lower part of Figure 1. The key idea is to leverage the retrieval model itself to find groups of similar paragraphs. At a certain training step, we use the paragraph encoder at that moment to encode the whole corpus and cluster all (q, p) pairs into many groups based on the similarity of their paragraph encodings. These groups are supposed to include similar paragraphs and potentially related questions. Then, we continue our pretraining by sampling each batch from one of the clusters. By doing this, we can provide challenging and effective negative paragraphs for each question, even with small batch size. Every time we recluster the whole corpus, the model will be encouraged to learn finer-grained matching between questions and paragraphs. Algorithm 1 provides a formal description of the entire process. Note that our training algorithm shares spirits with Curriculum Learning (Bengio et al., 2009) and Self-Paced Learning (Kumar et al., 2010; Jiang et al., 2015), in which the models are trained with harder

Algorithm 1 The Clustering-based Progressive Pretraining

```

1: Input:
2: a) all  $(q, p)$  pairs from the question generation model;
3: b) the retrieval module  $BERT_Q$  and  $BERT_P$ ;
4: while not finished do
5:   Encode the whole corpus with  $BERT_P$ ;
6:   Clustering all paragraphs into  $C$  clusters using the
   dense encodings;
7:   for updates = 1:K do
8:     Random sample a paragraph cluster;
9:     Sample  $B$  paragraphs from the cluster;
10:    Fetch the corresponding questions;
11:    Calculate gradients wrt  $\mathcal{L}_{pre}$ ;
12:    if updates %  $U == 0$  then
13:      Update  $BERT_Q$  and  $BERT_P$ ;
14:    end if
15:  end for
16: end while

```

instances as the training progresses. Instead of utilizing a predefined or dynamically generated order of all the instances according to their easiness, our algorithm makes use of a dynamic grouping of all the training instances and is specifically designed for the efficient in-batch negative sampling paradigm.

2.3 QA Finetuning

Once pretrained, we use the paragraph encoder to encode the corpus into a large set of dense vectors. Following previous practice, we only finetune the question encoder and the reading module so that we can reuse the same dense index for different datasets. For every training question, we obtain the question representation h_q from the question encoder and retrieve the top-k paragraphs D^k on the fly using an existing maximum inner product search package. To train the reading module, we apply the shared-normalization trick and optimize the marginal probability of all matched answer spans in the top-k paragraphs:

$$\mathcal{L}_{reader} = -\log \sum_{p \in D^*} \sum_{s \in S_p^*} P_{\phi}^{sn}(s|p, q). \quad (2)$$

In addition to the reader loss, we also incorporate the “early” loss used by Lee et al. (2019), which updates the question encoder using the top-5000 dense paragraph vectors. If we define D_{5000}^* as those paragraphs in the top-5000 that contain the correct answer, then the “early” loss is defined as:

$$\mathcal{L}_{early} = -\log \sum_{p \in D_{5000}^*} P_{\theta}(p|q). \quad (3)$$

Thus our total finetuning loss is $\mathcal{L}_{early} + \mathcal{L}_{reader}$. Note this is different from the joint formulation

Method	Dataset		
	NaturalQuestions-Open	WebQuestions	CuratedTREC
DrQA (Chen et al., 2017)	-	20.7	25.7
R ³ (Wang et al., 2018a)	-	17.1	28.4
DSQA (Lin et al., 2018)	-	25.6	29.1
HardEM (Min et al., 2019a)	28.1	-	-
PathRetriever (Asai et al., 2020)	32.6	-	-
WKLM (Xiong et al., 2020)	-	34.6	-
GraphRetriever (Min et al., 2019b)	34.5	36.4	-
ORQA (Lee et al., 2019)	33.3	36.4	30.1
ProQA(Ours)	37.4	37.1	34.6

Table 1: Open-domain QA results in terms of exact answer match (EM). The first part of the table shows results from methods that use the traditional IR component. Note that these methods retrieve more paragraphs (typically dozens) than dense retrieval methods listed in the second part of the table, which only finds answers from the top-5.

used by Lee et al. (2019) and Guu et al. (2020), which consider the paragraphs as latent variables when calculating $P(s|q)$. We find the joint objective does not bring additional improvements, especially after we use shared normalization. More variants of the finetuning objectives will be discussed in §3.5. At inference time, we use a linear combination of the retrieval score and the answer span score to rank the answer candidates from the top-5 retrieved paragraphs. The linear combination weight is selected based on the validation performance on each tested dataset.

3 Experiments

3.1 Datasets

We center our studies on QA datasets that reflect real-world information-seeking scenarios. We consider **1) NaturalQuestions-Open** (Kwiatkowski et al., 2019; Lee et al., 2019), which includes around 10K real-user queries (79,168/8,757/3,610 for train/dev/test) from Google Search; **2) WebQuestions** (Berant et al., 2013), which is originally designed for knowledge base QA and includes 5,810 questions (3,417/361/2,032 for train/dev/test) generated by Google Suggest API; **3) CuratedTREC** (Baudis and Sedivý, 2015), which includes 2,180 real-user queries (1,353/133/694 for train/dev/test) from MSNSearch and AskJeeves logs. Compared to other datasets such as SQuAD (Rajpurkar et al., 2016) and TriviaQA (Joshi et al., 2017), questions in these datasets are created without the presence of ground-truth answers and the answer paragraphs, thus are less likely to have lexical overlap with the paragraph.

3.2 Essential Implementation Details

For pretraining, we use a batch size of 80 and accumulate the gradients every 8 batches. We use the Adam optimizer (Kingma and Ba, 2015) with learning rate $1e-5$ and conduct 90K parameter updates. Following previous work (Lee et al., 2019), we use the 12-20-2018 snapshot of English Wikipedia as our open-domain QA corpus. When splitting the documents into chunks, we try to reuse the original paragraph boundaries and create a new chunk every time the length of the current one exceeds 256 tokens. Overall, we created 12,494,770 text chunks, which is on-par with the number (13M) reported in previous work. These chunks are also referred to as paragraphs in our work. For progressive training, we recluster all the chunks with k-means around every 20k updates using the paragraph encodings.

While finetuning the modules for QA, we fix the paragraph encoder in the retrieval module. For each question, we use the top-5 retrieved paragraphs for training and skip the question if the top-5 paragraphs fail to cover the answer. The MIPS-based retrieval is implemented with FAISS (Johnson et al., 2019). On NaturalQuestions-Open, we finetune for 4 epochs. To save the finetuning time on this large dataset, we only use a subset (2,000 out of 8,757) of the original development set for model selection. For WebQuestions and CuratedTREC (both of them are much smaller), we finetune for 10 epochs. The optimizer settings are consistent with the pretraining. Hyperparameters and further details can be found in the appendix.

3.3 QA Performance

Following existing studies, we use the exact match (EM) as the evaluation metric, which indicates the percentage of the evaluation samples for which the

Method	EM	model size	batch size	# updates
ORQA	33.3	330M	4096	100K
T5	36.6	11318M	-	-
REALM	40.4	330M	512	200K
ProQA	37.4	330M	80*8	90K

Table 2: Resource comparison with SOTA models. EM scores are measured on NaturalQuestions-Open. *batch size* and *updates* all refer to the dense index pretraining. Note that REALM uses ORQA to initialize its parameters and we only report the numbers after ORQA initialization. “80*8” indicates that we use a batch size of 80 and accumulate the gradients every 8 batches.

predicted span matches the groundtruth answers. In Table 1, we first show that our progressive method (denoted as **ProQA**) is superior to all of the open-domain QA systems (the upper part of the table) that use conventional IR methods, even though we only use the top-5 paragraphs to predict the answer while these methods use dozens of retrieved paragraphs. For the dense retrieval methods, we compare with ORQA (Lee et al., 2019), which is most relevant to our study but simply uses pseudo question-paragraph pairs for pretraining and also requires a larger batch size (4,096). We achieve much stronger performance than ORQA with much fewer updates and a limited number of GPUs. To the best of our knowledge, this is the first work showing that an effective dense corpus index can be obtained without using highly expensive computational resources. The reduced requirement of computation also makes our method easier to replicate for corpora in different domains.

In Table 2, we compare our method other more recently published QA systems in terms of both performance and computation cost. It is worth noting that although we need to use a BART model to generate training questions for Wikipedia documents, the inference cost of the question generator is still much lower than the training cost of our system and is not significant for comparing the overall computation cost: with the same GPU hardware, generating all the questions takes less than 1/6 of the training time. When using the *batch size* \times *# of updates* to approximate the training FLOPs, we see that our system is at least 6 times more efficient than ORQA. Compared to the recent proposed T5 (Roberts et al., 2020) approach, which converts the QA problem into a sequence-to-sequence (decode answers after encoding questions) problem and relies on the large model capacity to answer questions without

Method	R@5	R@10	R@20
ProQA (90k)	52.0	61.0	68.8
ORQA* (90k)	20.4	29.0	37.2
ProQA (no clustering, 90k)	42.9	52.6	60.8
ProQA (no clustering; 70k)	43.8	53.5	61.3
ProQA (no clustering; 50k)	38.8	48.2	56.7

Table 3: Ablation studies on different pretraining strategies. The retrieval modules (Recall@k) are tested on WebQuestions. *Our reimplemention.

retrieving documents, our system achieves better performance and is also much faster at inference time, due to the much smaller model size. The state-of-the-art REALM model (Guu et al., 2020) uses a more complicated pretraining approach that requires asynchronously refreshing the corpus index at train time. As it relies on ORQA initialization and further pretraining updates, it is even more computational expensive at training time. Also, as our method directly improves the ORQA pretraining, our method could easily stack with the REALM pretraining approach.

Concurrent to our work, Karpukhin et al. (2020) show that it is possible to use the groundtruth answer paragraphs in the original NaturalQuestions dataset to train a stronger dense retriever. However, they use a larger index dimension (768) while encoding paragraphs and also retrieve more paragraphs (20~100) for answer extraction. As a larger index dimension naturally leads to better retrieval results (Luan et al., 2020) (despite sacrificing search efficiency) and using more paragraphs increases the recall of matched answer spans³, this concurrent result is not directly comparable to ours, and we leave the combination effect of efficient pretraining and strong supervision (*i.e.*, using human-labeled paragraphs) to future work.

3.4 Ablation Studies

To validate the sample efficiency of our method, we replicate the inverse-cloze pretraining approach from ORQA using the same amount of resource as we used while training our model, *i.e.*, the same batch size and updates (90k). We also study the effect of the progressive training paradigm by pretraining the model with the same generated data but without the clustering-based sampling. We test the retrieval performance on the WebQuestions test

³While using more paragraphs, we achieve 40.6 EM (compared to 41.5 EM in (Karpukhin et al., 2020)) even with a much smaller index dimension. Also, we did not use the gold paragraphs (strong supervision) to train our reading module.

set before any finetuning. We use Recall@k as the evaluation metric, which measures how often the answer paragraphs appear in the top-k retrieval. The results are shown in Table 3. We can see that for the non-clustering version of our method, the improvements are diminishing as we reach certain training steps, while the progressive training algorithm brings around 8% improvements on different retrieval metrics. This suggests the importance of introducing more challenging negative examples in the batch when the batch size is limited. Comparing the no-clustering version of our method against ORQA, we see that using our data generation strategy results in much better retrieval performance (more than 22% improvements on all metrics).

3.5 Analysis on Finetuning Objectives

Noting that different finetuning configurations have been used in existing studies, we conduct additional experiments to investigate the efficacy of different finetuning objectives and provide insights for future research that intends to focus on improving the model itself. Specifically, we study the effects of using a joint objective (Lee et al., 2019; Guu et al., 2020) and adding an additional reranking objective that is commonly used in sparse-retrieval QA systems (Min et al., 2019a; Xiong et al., 2020).

The joint objective treats the retrieved paragraphs as latent variables and optimizes the marginal probability of the matched answer spans in all paragraphs:

$$\mathcal{L}_{joint} = -\log \sum_{p \in D^*} P_{\theta}(p|q) \sum_{s \in S_p^*} P_{\phi}(s|p, q). \quad (4)$$

The reranking objective is usually implemented through a paragraph reranking module that uses a question-dependent paragraph encoder, *e.g.*, a BERT encoder that takes the concatenation of the question and paragraph as input. This kind of reranker has been shown to be beneficial to conventional IR methods since it can usually provide more accurate paragraph scores than the TF-IDF or BM25 based retriever while ranking the answer candidates. To implement this reranking module, we simply add another reranking scoring layer to our BERT-based span prediction module $P_{\phi}(s|p, q)$, which encodes the paragraphs in a question-dependent fashion. At inference time, we use the paragraph scores predicted by this reranking component instead of the pretrained retrieval model to guide our final answer selection.

id	Objective Settings			EM
	joint	rerank	shared-norm	
1	-	-	✓	38.5
2	✓	-	✓	38.3
3	-	✓	✓	38.2
4	✓	-	-	36.2
5	-	-	-	35.1

Table 4: Analysis on different finetuning objectives on NaturalQuestions-Open. EM scores are measured on the 2,000 validation samples we used for model selection.

Table 4 shows the results of different objective settings. Comparing the results of (4) and (5), we can see that the joint objective can bring some improvements when shared-normalization is not applied. However, it does not yield improvements when shared-normalization is applied, according to the results of (1) and (2). By comparing (1) and (3), we see that with the strong pretrained retrieval model, adding an extra reranking module that uses question-dependent paragraph encodings is no longer beneficial. This is partially because our pretrained retrieval model gets further improved during finetuning, in contrast to a fixed TF-IDF based retriever. Finally, from (1) and (5), we see that the shared normalization brings much larger improvements than the other factors. This aligns with the findings from an existing work (Wang et al., 2019) that only tested on SQuAD questions.

4 Error Analysis

To investigate the fundamental differences of the dense and sparse retrieval methods in open-domain QA, we conduct an error analysis using both the proposed method and a baseline system that uses TF-IDF and BM25 for retrieval. This baseline uses a similar retrieval pipeline as Min et al. (2019a) and is trained with the finetuning objective defined in Eq. 2. This sparse retrieval baseline achieves 29.7 EM on the official dev set of NaturalQuestions-Open while our method achieves 36.7 EM⁴. Figure 2 shows the Venn diagram of the error sets from both systems. Our key findings are summarized in the following paragraphs.

The difference retrieval paradigms could complement each other. First, according to the error set differences (shown by the white regions in Figure 2), a considerable portion of the error cases (9.1% of the devt set) of our dense-retrieval system

⁴Note that this number is different from the number in Table 4 as we did not use the whole dev set for model selection.




Figure 2: The error sets (from the official dev set of NaturalQuestions-Open) of our dense retrieval method and a baseline system using sparse retrieval methods like TF-IDF. We use circles to represent the error sets of both systems. The percentages show the relative size of each set in terms of the whole dev set.

does not occur in the sparse-retrieval system and vice versa. This suggests the necessity of incorporating different retrieval paradigm when building real-world applications. In fact, the hybrid approach has already been adopted by a phrase-level retrieval method (Seo et al., 2019) and concurrent studies (Luan et al., 2020; Karpukhin et al., 2020).

Both systems are underestimated. As shown in Figure 2, 54.2% of the questions cannot be correctly answered by either system. However, our manual inspection on 50 of the shared error cases suggests that around 30% of these errors are due to annotation issues (14%) or the ambiguous nature of real-user queries (16%). One obvious annotation issue is the incompleteness of the answer labels. Example questions include “*When did Brazil lose to in 2014 World Cup?*”, to which both “*Germany*” and “*Netherlands*” are correct answers. This issue occurs because the annotators of NaturalQuestions only have a local view of the knowledge source as they are only asked to label the answer span using one document. In terms of the ambiguous questions, many of them are due to constraints unspecified by the question words, such as “*What is the population of New York City?*” (the time constraint is implicit) or “*When did Justice League come out in Canada?*” (needs entity disambiguation). This kind of questions result from the information-seeking nature of the open-domain QA task where the users usually use the minimal number of words for searching and they are not aware of the potential ambiguous factors. To solve this kind of questions, an interactive QA system might be necessary. In the appendix, we show more ambiguous questions in which other kinds of constraints are missing.

5 Related Work

The task of answering questions without specifying specific domains has been intensively studied since the earlier TREC QA competitions (Voorhees, 1999). Studies in the early stage (Kwok et al., 2001; Brill et al., 2002; Ferrucci et al., 2010; Baudiš, 2015) mostly rely on highly sophisticated pipelines and heterogeneous resources. Built on the recent advances in machine reading comprehension, Chen et al. (2017) show that open-domain QA can be simply formulated as a reading comprehension problem with the help of a standard IR component that provides candidate paragraphs for answer extraction. This two-stage formulation is simple yet effective to achieve competitive performance while using Wikipedia as the only knowledge resource.

Following this formulation, a couple of recent studies have proposed to improve the system using stronger reading comprehension models (Yang et al., 2019; Wang et al., 2018b), more effective learning objectives (Clark and Gardner, 2018; Min et al., 2019a; Wang et al., 2019) or paragraph reranking models (Wang et al., 2018a; Lin et al., 2018; Lee et al., 2018). However, the retrieval components in these systems are still based on traditional inverted index methods, which are efficient but might fail when the target paragraph does not have enough lexicon overlap with the question.

In contrast to the sparse term-based features used in TF-IDF or BM25, dense paragraph vectors learned by deep neural networks (Zhang et al., 2017; Conneau et al., 2017) can capture much richer semantics beyond the n-gram term features. To build effective paragraph encoders tailed for the paragraph retrieval in open-domain QA, more recent studies (Lee et al., 2019; Chang et al., 2020; Guu et al., 2020) propose to pretrain Transformer encoders (Vaswani et al., 2017) with objectives that simulate the semantic matching between questions and paragraphs. For instance, Lee et al. (2019) uses the inverse cloze pretraining task to train a bi-encoder model to match a sentence and the paragraph in which the sentence belongs to. These approaches demonstrate promising performance but require a lot of resources for pretraining. The focus of this paper is to reduce the computational requirements of building an effective corpus index.

6 Conclusion

We propose an efficient method for pretraining the dense corpus index which can replace the tradi-

tional IR methods in open-domain QA systems. The proposed approach is powered by a better data generation strategy and a simple yet effective data sampling protocol for pretraining. With careful finetuning, we achieve stronger QA performance than ORQA that uses much more computational resources. We hope our method could encourage more energy-efficient pretraining methods in this direction such that the dense retrieval paradigm could be more widely used in different domains.

Acknowledgement

The research was partly supported by the J.P. Morgan AI Research Award. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agency.

References

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. [Learning to retrieve reasoning paths over wikipedia graph for question answering](#). In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia.
- Petr Baudiš. 2015. Yodaqa: a modular question answering system pipeline. In *19th International Student Conference on Electrical Engineering*, pages 1156–1165.
- Petr Baudis and Jan Sedivý. 2015. [Modeling of the question answering task in the yodaqa system](#). In *6th International Conference of the CLEF Association, CLEF 2015*, pages 222–228, Toulouse, France. Springer.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009*, pages 41–48, Montreal, Québec, Canada. ACM.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013*, pages 1533–1544, Seattle, Washington, USA. ACL.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Eric Brill, Susan T. Dumais, and Michele Banko. 2002. [An analysis of the askmsr question-answering system](#). In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing, EMNLP 2002*, Philadelphia, Pennsylvania, USA.
- Wei-Cheng Chang, Felix X. Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. [Pre-training tasks for embedding-based large-scale retrieval](#). In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Qian Chen, Zhen-Hua Ling, and Xiaodan Zhu. 2018. [Enhancing sentence embedding with generalized pooling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1815–1826, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Christopher Clark and Matt Gardner. 2018. [Simple and effective multi-paragraph reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018*, pages 845–855, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019*, pages 4171–4186, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- David A. Ferrucci, Eric W. Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John M. Prager, Nico Schlaefer, and Christopher A. Welty. 2010. [Building watson: An overview of the deepqa project](#). *AI Magazine*, 31(3):59–79.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *CoRR*, abs/2002.08909.

- Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G. Hauptmann. 2015. [Self-paced curriculum learning](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 2694–2700, Austin, Texas, USA. AAAI Press.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. [Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). *CoRR*, abs/2004.04906.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015*, San Diego, California, USA.
- M. Pawan Kumar, Benjamin Packer, and Daphne Koller. 2010. [Self-paced learning for latent variable models](#). In *24th Annual Conference on Neural Information Processing Systems, NIPS 2010*, pages 1189–1197, Vancouver, British Columbia, Canada. Curran Associates, Inc.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: a benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Cody C. T. Kwok, Oren Etzioni, and Daniel S. Weld. 2001. [Scaling question answering to the web](#). In *Proceedings of the Tenth International World Wide Web Conference, WWW 2001*, pages 150–161, Hong Kong, China. ACM.
- Jinhyuk Lee, Seongjun Yun, Hyunjae Kim, Miyoung Ko, and Jaewoo Kang. 2018. [Ranking paragraphs for improving answer recall in open-domain question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 565–569, Brussels, Belgium. Association for Computational Linguistics.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *CoRR*, abs/1910.13461.
- Yankai Lin, Haozhe Ji, Zhiyuan Liu, and Maosong Sun. 2018. [Denoising distantly supervised open-domain question answering](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 1736–1745, Melbourne, Australia. Association for Computational Linguistics.
- Lajanugen Logeswaran and Honglak Lee. 2018. [An efficient framework for learning sentence representations](#). In *6th International Conference on Learning Representations, ICLR 2018*, Vancouver, British Columbia, Canada.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2020. [Sparse, dense, and attentional representations for text retrieval](#). *CoRR*, abs/2005.00181.
- Xiyao Ma, Qile Zhu, Yanlin Zhou, Xiaolin Li, and Dapeng Wu. 2020. [Improving question generation with sentence-level semantic matching and answer position inferring](#). In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, New York, New York, USA. AAAI Press.
- Sewon Min, Danqi Chen, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019a. [A discrete hard EM approach for weakly supervised question answering](#). 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)*, pages 2851–2864, Hong Kong, China. Association for Computational Linguistics.
- Sewon Min, Danqi Chen, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019b. [Knowledge guided text retrieval and reading for open domain question answering](#). *CoRR*, abs/1911.03868.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. [A decomposable attention model for natural language inference](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, Austin, Texas. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016*, pages 2383–2392, Austin, Texas, USA. The Association for Computational Linguistics.

- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) *CoRR*, abs/2002.08910.
- Stephen E. Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Min Joon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur P. Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. [Real-time open-domain question answering with dense-sparse phrase index](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*, pages 4430–4441, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Thirty-first Annual Conference on Neural Information Processing Systems, NIPS 2017*, pages 5998–6008, Long Beach, California, USA.
- Ellen M. Voorhees. 1999. [The TREC-8 question answering track report](#). In *Proceedings of The Eighth Text REtrieval Conference, TREC 1999*, Gaithersburg, Maryland, USA. National Institute of Standards and Technology.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018a. [R³: Reinforced ranker-reader for open-domain question answering](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 5981–5988, New Orleans, Louisiana, USA. AAAI Press.
- Shuohang Wang, Mo Yu, Jing Jiang, Wei Zhang, Xiaoxiao Guo, Shiyu Chang, Zhiguo Wang, Tim Klinger, Gerald Tesauro, and Murray Campbell. 2018b. [Evidence aggregation for answer re-ranking in open-domain question answering](#). In *6th International Conference on Learning Representations, ICLR 2018*, Vancouver, British Columbia, Canada.
- Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nalapati, and Bing Xiang. 2019. [Multi-passage BERT: A globally normalized BERT model for open-domain question answering](#). 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 2019*, pages 5877–5881, Hong Kong, China. Association for Computational Linguistics.
- Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2020. [Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model](#). In *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. [End-to-end open-domain question answering with BERTserini](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yizhe Zhang, Dinghan Shen, Guoyin Wang, Zhe Gan, Ricardo Henao, and Lawrence Carin. 2017. [Deconvolutional paragraph representation learning](#). In *Thirty-first Annual Conference on Neural Information Processing Systems, NIPS 2017*, pages 4169–4179, Long Beach, California, USA.

A Appendix

A.1 Further Implementation Details

The pretraining process takes 4 TITAN RTX GPUs, each with a 24G memory. We use the NVIDIA Apex package for mixed-precision training. 90K parameter updates take around 7 days to finish. For QA experiments, We use the **IndexIVFFlat** index for efficient search. We assign all the vectors to 100 Voronoi cells and only search from the closest 20 cells. The random seed is set as 3 for all QA datasets. We use a batch size of 8 (8 questions, each of them are paired with 5 paragraphs) for NaturalQuestions-Open and 1 for the other datasets. We limit the maximum answer length to 10 sub-word tokens. For NaturalQuestions-Open, we evaluate the model every 1000 updates and save the best checkpoint based on validation EM. For WebQuestions and CuratedTREC, we evaluate the model after every epoch. As neither of these two small datasets has an official dev set, we use a small split to find the best hyperparameters and then retrain the model with all the training questions. To accelerate training, especially for the early loss function which requires annotate the top5000 retrieved paragraphs, we pre-annotate the top10000 paragraphs retrieved by the untuned retrieval module and build an answer paragraph set for each question. At fine-tuning time, we direct check whether a particular paragraph is in the precomputed paragraph set, instead of doing string matching for each of the 5000 paragraphs. Our BERT implementations are based on huggingface Transformers⁵.

A.2 Qualitative Examples

Here we include more examples that complement the results and analysis of the paper. Table 5 shows the generated questions from the finetuned BART model and Table 6 complements the error analysis.

⁵<https://github.com/huggingface/transformers>

<p>Gold Paragraph: “Does He Love You” is a song written by Sandy Knox and Billy Stritch, and recorded as a duet by American country music artists Reba McEntire and Linda Davis. It was released in August 1993 as the first single from Reba’s album Greatest Hits Volume Two. It is one of country music ’s several songs about a love triangle.</p> <p>Original Question: Who sings does he love me with reba?</p> <p>Generated Question: Who sings with reba mcentire on does he love you?</p>
<p>Gold Paragraph: Invisible Man First edition Author Ralph Ellison Country United States Language English Genre Bildungsroman African-American literature social commentary Publisher Random House Publication date 1952 Media type Print (hardcover and paperback) Pages 581 (second edition) ...</p> <p>Original Question: How many pages is invisible man by ralph ellison?</p> <p>Generated Question: How many pages in the invisible man by ralph ellison?</p>
<p>Gold Paragraph: The Great Lakes (French: les Grands-Lacs), also called the Laurentian Great Lakes and the Great Lakes of North America, are a series of interconnected freshwater lakes located primarily in the upper mid-east region of North America, on the Canada–United States border, which connect to the Atlantic Ocean through the Saint Lawrence River. They consist of Lakes Superior, Michigan, Huron (or Michigan–Huron), Erie, and Ontario.</p> <p>Original Question: Where do the great lakes meet the ocean?</p> <p>Generated Question: Where do the great lakes of north america meet the atlantic?</p>
<p>Gold Paragraph: My Hero Academia: Two Heroes , Hepburn:Boku no Hiro Academia THE MOVIE: Futari no Hiro) is a 2018 Japanese anime superhero film based on the manga My Hero Academia by Kohei Horikoshi. Set between the second and third seasons of the anime series, the film was directed by Kenji Nagasaki and produced by Bones. Anime Expo hosted the film’s world premiere on July 5, 2018, and it was later released to theaters in Japan on August 3, 2018.</p> <p>Original Question: When does the new my hero academia movie come out?</p> <p>Generated Question: When does the my hero academia two heroes movie come out?</p>
<p>Gold Paragraph: Victoria’s Secret Store, 722 Lexington Ave, New York, NY Type Subsidiary Industry Apparel Founded June 12, 1977; 40 years ago (1977-06-12) Stanford Shopping Center, Palo Alto, California, U.S. Founder Roy Raymond Headquarters Three Limited Parkway, Columbus , Ohio , U.S. Number of locations 1,017 company - owned stores 18 independently owned stores Area served ...</p> <p>Original Question: Who was the creator of victoria’s secret?</p> <p>Generated Question: Who is the founder of victoria’s secret and when was it founded?</p>

Table 5: Samples of the generated questions. The answer spans are underlined. Here we show the generated questions for samples at the beginning of the official NaturalQuestions-Open dev data. We only skip the samples whose gold paragraphs are not natural paragraphs (e.g., incomplete sentences).

<p>Question: What is a ford mondeo in the usa?</p> <p>Annotated Answers: ford contour, mercury mystique, ford fusion <i>ambiguous;</i> could be asking about a particular car type (mid-sized car) instead of brand series</p>
<p>Question: air flow in the eye of a hurricane?</p> <p>Annotated Answers: no wind <i>ambiguous;</i> question itself is hard to understand</p>
<p>Question: Who wrote I’ll be there for you?</p> <p>Annotated Answers: Michael Skloff, Marta Kauffman, Allee Willis, David Crane, Phil Solem, Danny Wilde, The Rembrandts <i>ambiguous;</i> there are multiple songs having this name</p>
<p>Question: Where do you go for phase 1 training?</p> <p>Annotated Answers: army foundation college <i>ambiguous;</i> the meaning of phase 1 is vague, could have different meanings in different context</p>
<p>Question: When does the new spiderman series come out?</p> <p>Annotated Answers: August 19 , 2017 <i>ambiguous;</i> time constraint missing</p>
<p>Question: Where did the super bowl take place this year?</p> <p>Annotated Answers: minneapolis, minnesota <i>ambiguous;</i> the year cannot be inferred from the question words alone</p>

Table 6: Error cases that include ambiguous questions.