

# Augmenting Pre-trained Language Models with QA-Memory for Open-Domain Question Answering

Wenhu Chen, Pat Verga, Michiel de Jong<sup>†</sup>, John Wieting, William W. Cohen

Google Research, University of Southern California<sup>†</sup>

{wenhuchen, patverga, jwieting, wcohen}@google.com, msdejong@usc.edu

## Abstract

Existing state-of-the-art methods for open-domain question-answering (ODQA) use an open book approach in which information is first retrieved from a large text corpus or knowledge base (KB) and then reasoned over to produce an answer. A recent alternative is to retrieve from a collection of previously-generated question-answer pairs; this has several practical advantages including being more memory and compute-efficient. Question-answer pairs are also appealing in that they can be viewed as an intermediate between text and KB triples: like KB triples, they often concisely express a single relationship, but like text, have much higher coverage than traditional KBs. In this work, we describe a new QA system that augments a text-to-text model with a large memory of question-answer pairs, and a new pre-training task for the latent step of question retrieval. The pre-training task substantially simplifies training and greatly improves performance on smaller QA benchmarks. Unlike prior systems of this sort, our QA system can also answer multi-hop questions that do not explicitly appear in the collection of stored question-answer pairs.

## 1 Introduction

Open-domain question answering (ODQA) is a well-studied knowledge-intensive task. State-of-the-art methods require retrieving relevant knowledge from a large corpus or datastore before reasoning over this retrieved evidence. Most existing methods retrieve documents (Chen et al., 2017a; Lee et al., 2019; Karpukhin et al., 2020) or structured KB triples (Verga et al., 2021). Recently, a few works have proposed retrieving from a collection of question-answer (QA) pairs—an approach made feasible by advances in scalable automatic question generation. In this setting, a new question is answered by retrieving paraphrases from a question index, and returning the associated answer

(Xiao et al., 2021; Lewis et al., 2021). Notably, the RePAQ system from (Lewis et al., 2021) won the 2020 EfficientQA competition (Min et al., 2021), outperforming closed-book QA (CBQA) models by a significant margin and matching the prior SoTA performance on NQ (Kwiatkowski et al., 2019).

A collection of QA pairs is appealing for several reasons. As opposed to text passages and much like a KB triple, QA pairs are often concise, tending to express a single relationship. However, unlike KB triples, QA collections have good coverage of actually asked questions like those in standard open QA datasets. RePAQ demonstrated several advantageous properties such as memory and computational efficiency, strong selective QA performance (i.e. selectively abstaining from answering), and effective ensembling with text-retrieval QA systems.

However, question-retrieval QA systems have several limitations as well. First, there is *no large-scale supervised data for question-question retrieval*. This contrasts with the step of retrieving text given a question, where supervised data is used to build retrievers like DPR (Karpukhin et al., 2020). To address this, RePAQ uses a latent-retrieval training process (similar to REALM (Lee et al., 2019)), in which the retriever is trained using the downstream end loss from the QA task. This requires asynchronously updating the index as training proceeds, a process that is complex and computationally expensive. This is also a problem for domains with limited QA data: as we will show, RePAQ’s performance is disappointing on smaller datasets like WebQuestions (Berant et al., 2013), containing only 3K training instances. To address this problem, we introduce a novel *pre-training task for question retrieval*, which can be applied to any text-QA dataset, which greatly improves performance on smaller datasets.

A second problem is that RePAQ is limited to answering questions explicitly stored in the index, or paraphrases of such questions. This contrasts

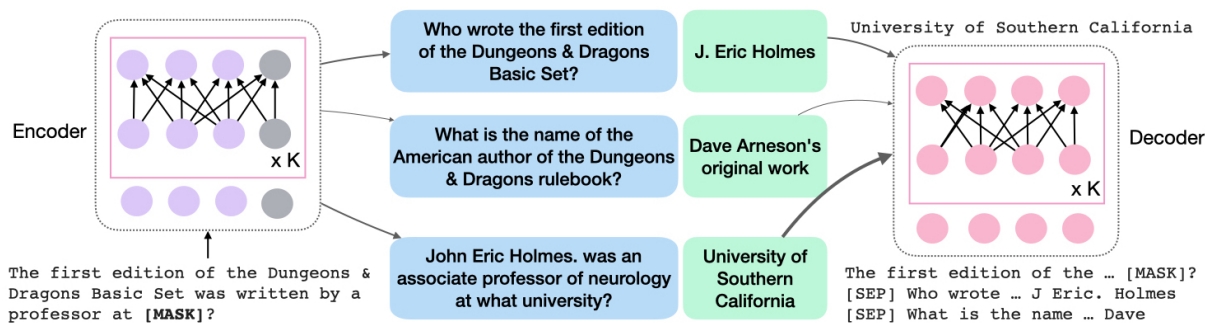


Figure 1: During pre-training, the encoder first encodes textual input and use special token representation to query the QA-memory. The retrieved QA-pairs are integrated to the decoder to generate outputs.

with QA systems that retrieve from KBs, which can typically generate complex queries that combine the atomic triples in the KB. To address this, we present an extended model that answers multi-hop questions by iteratively retrieving from a question-answer corpus, the first question-retrieval-based QA system that addresses this task.

In more detail, we propose a new QA-Memory-Augmented Transformer (QAMAT) with better compositionality paired with a lower complexity training strategy. QAMAT is based on a T5 encoder-decoder (Raffel et al., 2020) paired with an integrated key-value memory (Khandelwal et al., 2019; Borgeaud et al., 2021) populated with question-answer pairs (See Figure 1). Given an input, the encoder generates a query representation scored against the QA memory and retrieves the top-K relevant QA pairs. The encoder then reprocesses the input along with the retrievals forming a QA-injected representation which is passed to the decoder to attend to and generate.

To reduce the training (fine-tuning) sample complexity, we propose to first pre-train QAMAT on a large-scale corpus to teach the model to retrieve and interpret QA pairs. We construct the pre-training corpus by leveraging existing methods for question generation, producing a very large set of potentially interesting questions from text passages (Zhou et al., 2017; Alberti et al., 2019; Lewis et al., 2021). For each QA pair and the passage it was generated from, we mask the answer and train the model to fill the mask by retrieving and using an appropriate QA pair. We show that pre-training greatly boosts the model’s performance and helps the model generalize to different domains. For example, the pre-trained model can achieve a zero-shot performance of 40% EM on NQ and TriviaQA without *any* fine-tuning.

The effectiveness of this pre-training task means

that we can avoid the expensive latent training procedure used by RePAQ, and instead use an efficient two-stage training pipeline. In the first stage, we use a small local in-batch memory of QA pairs to optimize the QA pair encoder. We then freeze the encoder and construct the index for the global memory. In the second stage, we retrieve from this fixed global memory and continue to optimize the remaining parameters—including the parameters used to construct queries to the global memory—for better performance.

Lastly, we extend QAMAT to build QAMAT+, which iteratively retrieves from the memory to generate outputs. We demonstrate that QAMAT+ effectively chains multiple QA-pairs together to answer multi-hop questions in HotpotQA (Yang et al., 2018) and Musique (Trivedi et al., 2021). Such compositional reasoning capability is nonexistent in RePAQ (Lewis et al., 2021).

In summary, we develop a new **QA augmented architecture** which extends the lines of research considering QA pairs as a representation of knowledge as well as those on memory-augmented language models. When paired with our proposed **pre-training strategy** (section 4), we address many of the shortcomings of previous QA-indexing-based approaches leading to lower sample complexity training and the ability to perform **compositional reasoning** (subsection 3.5).

## 2 Related Work

### 2.1 Retriever-Reader Models

Retrieve-and-read models have been widely studied to address knowledge-intensive tasks and achieve state-of-the-art performance on most QA tasks. These methods use two models, one to retrieve from a passage index based on BM25 (Robertson and Zaragoza, 2009), and one to perform reading comprehension on the returned passages (Chen

et al., 2017b). More recently, deep retrieval models have gained more popularity to replace traditional string-similarity retriever.

DPR (Karpukhin et al., 2020) is a widely used supervised approach to achieve better results than BM25 on a large collection of text retrieval tasks (Thakur et al., 2021). Contrastive learning is used to train the deep retriever model to distinguish between annotated positive and mined negative candidates. More recently, ColBERT (Khattab and Zaharia, 2020) has been proposed to integrate more fine-grained late fusion between query and context to improve DPR.

Retrieval Augmented Generation (RAG) (Lewis et al., 2020), Fusion-in-Decoder (FiD) (Izard and Grave, 2021) and End-to-end training of Multi-Document Reader and Retriever (EmDR) (Singh et al., 2021) are proposed to read retrievals to extract or generate answers. These models require a trained retriever/reranker to obtain top-K results, which are fed to the reader to generate the answer. As discussed in section 1, our model provides better interpretability due to atomic knowledge representation. In subsection 5.4, we also demonstrate that our model’s inference speed is 5x faster.

## 2.2 Question Generation

The problem of question generation (Zhou et al., 2017) has attracted attention from the community in recent years. It has been used for data augmentation (Alberti et al., 2019) to improve current QA systems or to improve retrieval systems (Nogueira et al., 2019). Pan et al. (2021) also demonstrated that by connecting generated single-hop questions, we can train zero-shot multi-hop question answering systems. Besides QA, it has also been widely used in other domains like evaluating factual consistency of summarization (Eyal et al., 2019; Wang et al., 2020) or enhancing contextualized representation (Jia et al., 2021). Most related to our work is PAQ (Lewis et al., 2021), which aims to generate and use QA pairs as retrieval units for question answering. The efficacy of this data was further verified when it was used to train DPR, yielding better domain generalization (Oğuz et al., 2021).

## 2.3 Memory-Augmented Language Models

End-to-end memory-augmented language models aim to train a model to explicitly access external memory. The current work is focused on storing entities (Févy et al., 2020), entity mentions (Dhingra et al., 2019; Sun et al., 2021; de Jong et al.,

2022) or knowledge triples (Verga et al., 2021). Memory attention layers are then used to influence the computation of transformer layers. These entities and fact-centric memories are naturally atomic and interpretable, and models employing them have shown competitive performance on entity-focused QA datasets like Web-Question-SP (Yih et al., 2016) and ComplexWebQuestions (Talmor and Berant, 2018). However, these models are limited to integrating entity-centric knowledge and classifying the answer w.r.t a pre-defined entity list. For example, these models cannot handle questions with non-entity answers, e.g. number, date, noun phrases, etc, which are ubiquitous in various QA datasets like NQ (Kwiatkowski et al., 2019), SQuAD (Rajpurkar et al., 2016), or HotpotQA (Yang et al., 2018).

## 3 Our Model: QAMAT

### 3.1 Problem Definition

The input to our model is a piece of text  $X = x_1, \dots, x_n$ , where  $X$  is either a question during fine-tuning or a paragraph in pre-training. Pre-training is formulated as a span corruption task (Raffel et al., 2019): given an example in the pre-training corpus as  $(X, \{Q^k, A^k\}_{k=1}^m)$ , where  $A^1, \dots, A^m$  correspond to spans in the input  $X$ . We sample  $k$  spans from  $X$  as a cloze answer and replace all tokens within a span with a [MASK] token, and the model needs to recover all the answers. During fine-tuning, we add an artificial [MASK] in the question front, and let the model recover this as the answer. The pre-training/fine-tuning objective function is to maximize the masked language model objectives  $p(Y|X) = \sum_{m_i \in M} p(Y|X, m_i)p(m_i|X)$ , which marginalizes over the entire memory  $M$ . However, due to its intractability in a large-scale memory, we adopt an approximation to only sum over the top-K memory entries  $Top_K(M)$ .

We define the encoder function as  $f_\theta$ , which takes an input sequence  $X$  as input to generate a sequence of vector  $\mathcal{F}_\theta(X) \in \mathbb{R}^{n \times d}$ , where  $n$  is the input length and  $d$  is the hidden size. The designated position of  $\mathcal{F}_\theta(X)$  will be used as the query and memory representation, which are denoted as  $f_\theta(X; [\text{MASK}]) \in \mathbb{R}^d$  (at [MASK] position) and memory key/value as  $f_\theta(m_i^k; [\text{CLS}]) \in \mathbb{R}^d$  (at [CLS] position). For brevity, we leave out [MASK] and [CLS] and simply use  $f_\theta(\cdot)$ .

We also define a broadcast operator  $B_k^n(x)$  to




Figure 2: Architecture: upper figure shows the retrieval process with shared encoder, the lower figure shows the decoder process to leverage neural and discrete representation of memory retrieval.

broadcast a vector into a matrix by assigning the vector  $x$  to  $k$ -th row while filling the rest with zero, i.e.  $B_k^n(x) = [\mathbf{0}, \dots, x^T, \dots, \mathbf{0}]$ .

### 3.2 Dense Retriever

The memory  $M$  contains separate key and value components, where the key  $m_i^k$  contains a question, and the corresponding value  $m_i^v$  contains the question-answer concatenation. To retrieve the top- $k$  QA-pairs from the memory, we use our encoder  $f_\theta$  to encode  $X$  and  $m_i$  separately and select the top- $K$  entries  $Top_K(M)$  based on their inner product, i.e.  $Top_K(M) = \text{argmax}_{m_i \in M} f_\theta(X) \cdot f_\theta(m_i^k)$ .

### 3.3 Neural Memory Integration

After the model retrieves the Top- $K$  candidates, their corresponding memory values  $m_i^v$  needs to be leveraged into the encoder to influence the decoder outputs in a differentiable fashion. We write our objective  $p(Y|X)$  as:

$$\begin{aligned}
& \sum_{m_i \in Top_K(M)} p(Y|X, m_i) p(m_i|X) \\
&= \sum_{m_i \in Top_K(M)} p(m_i|X) g_\theta(Y | \mathcal{F}_\theta(X) + B_k^n[f_\theta(m_i^v)]) \\
&\approx g_\theta(Y | \sum_{m_i \in Top_K(M)} p(m_i|X) (\mathcal{F}_\theta(X) + B_k^n[f_\theta(m_i^v)])) \\
&= g_\theta(Y | \mathcal{F}_\theta(X) + B_k^n[\sum_{m_i \in Top_K(M)} p(m_i|X) f_\theta(m_i^v)]) \\
& p(m_i|X) = \frac{e^{f_\theta(X) \cdot f_\theta(m_i^k)}}{\sum_{m_i \in Top_K(M)} e^{f_\theta(X) \cdot f_\theta(m_i^k)}}
\end{aligned}$$

The probability  $p(Y|X, m_i)$  is parameterized by a decoder function  $g_\theta$ , which takes a memory-infused encoder representation  $\mathcal{F}_\theta(X) + B_k^n[f_\theta(m_i^v)]$  as input. We approximate this marginal probability by pulling weighted summation inside the decoder function  $g_\theta$  to derive an aggregated memory-infused encoder representations  $\mathcal{F}_\theta(X) + B_k^n[\dots]$ . The retrieval weight  $p(m|X)$  is calculated as the softmax over the retrieval score over top- $K$  items. For simplicity,  $H(X, Top_K(M), p(m|X))$  is used to denote this encoder representation, thus the ob-

jective can be written as follows:

$$p(Y|X) = g_\theta(Y | H(X, Top_K(M), p(m|X))) \quad (1)$$

As shown in the upper part of Figure 2, we first use weighted-sum over the neural representation of retrieved memory entries  $f_\theta(m_i^v)$  and then simply add it to the encoder representation to infuse the retrieved QA-pair information. These two operations are both differentiable, which makes it possible to train retriever latently. In essence, the retriever will increase weights  $p(m_i|X)$  on more relevant memory items instead of irrelevant ones.

### 3.4 Neural + Discrete Memory Integration

A disadvantage of adopting weighted-sum  $\sum_i p(m_i|X) f_\theta(m_i^v) \in \mathbb{R}^d$  is that all the information from all of the top- $K$  documents are overly compressed into a  $d$ -dimension vector, whereas the token retrieval representation contains more information. Therefore, we propose to add a fine-grained token-wise representation  $\hat{H}(X, Top_K(M))$  to help the model access the retrieved discrete values  $m_i$  directly. The representation is obtained by encoding the concatenation of the input  $X$  and retrieved discrete tokens  $\hat{X} = \text{Concat}[m_k; \dots; m_1; X] \in \mathcal{R}^{(n+k|m)| \times d}$ .

Such discrete memory integration greatly enriches the representation for  $m_i$  and enables cross-attention between the query and retrieval, addressing the bottleneck problem. However, such discrete representation cannot propagate gradients back to the retriever. Finally, we propose to combine the neural memory  $H(X, \cdot)$  and discrete memory  $\hat{H}(X, \cdot)$  integration to combine their merits.

$$\begin{aligned}
& p(Y|X) \\
&= g_\theta(Y | H'(X, Top_K(M), p(m|X)) + \lambda \hat{H}(X, Top_K(M)))
\end{aligned} \quad (2)$$

where the  $\lambda$  is the balancing factor to weight the two representations. We use  $H'(\dots) = [\mathbf{0}; H(\dots)]$  to represent the concatenation of zero-matrix  $\mathbf{0} \in \mathbb{R}^{k|m| \times d}$ , which has consistent dimension with  $\hat{H}$ .



After leveraging  $\hat{H}$ , our model demonstrates significant improvements on the downstream tasks with 14% on TriviaQA and 10% on HotpotQA.

$H(X, Top_K(M), p(m|X))$  is only used to lazily train the retriever, after training, we can drop it and only use the concatenated representation  $\hat{H}(X, Top_K(M))$  as the encoder representation. The decoder  $g_\theta$  will attend to  $\hat{H}$  and perform a greedy search over vocabulary to generate output.

### 3.5 Multi-hop Extension

To further extend QAMAT’s capability to perform compositional reasoning, we propose a cascaded architecture (depicted in Figure 3) known as QAMAT+, where the model learns to perform multiple rounds of retrieval before feeding the augmented inputs to the decoder. Specifically for two-hop reasoning, we use  $X$  as the query to retrieve a first-round of top-K memory values  $Top_K(M; 1)$  with our learned retriever  $f_\theta$  described in subsection 3.2. Next, we augment the query by concatenating the retrieved values as  $X^1 = [Top_K(M; 1); X]$ . This new query  $X^1$  is used to perform a second round of retrieval to obtain additional top-K memory values,  $Top_K(M; 2)$ . Based on  $Top_K(M; 2)$ , we compute the hybrid encoder representation  $H(X^1, Top_K(M; 2))$  and  $\hat{H}(X^1, Top_K(M; 2))$  to compute  $p(Y|X^1; \theta)$ .




Figure 3: QAMAT+ architecture: Multi-Hop framework for question-answer memory integration.

## 4 Training

### 4.1 Pre-training Corpus

Our QA-pairs are constructed by combining 30M deduplicated QA-pairs from PAQ (Lewis et al., 2021)(originally 65M, we delete paraphrases to keep a subset) and 30M additional QA-pairs generated from our own pipeline. The additional QA-pairs are populated from non-overlapping passage blocks to increase the knowledge coverage over Wikipedia. Our QA generation pipeline is similar to (Lewis et al., 2021) but trained solely on SQuAD 2.0 (Rajpurkar et al., 2018) and filtered with a cheap reading comprehension model rather than FiD (Izcard and Grave, 2021), the details are




Figure 4: Two stage training procedure: in-batch training with a batch-specific memory and end-to-end gradient updates, global training with a fixed global memory and partial gradient updates.

described in the Appendix. The final statistics of our QA-memory is described in Table 1, where the total size is comparable to RePAQ.

Memory	Size	#Passages	Training Data
Dedup-PAQ	30M	10M	NaturalQuestions
Additional	30M	10M	SQuAD 2.0
Combined	60M	20M	-

Table 1: The breakdown statistics of our QA corpus.

We denote the entire memory as  $M$  and formulate the pre-training corpus as  $\{X, \{Q^k, A^k\}_{k=1}^m\}$ , where  $X$  is the passage aligned with multiple QA-pairs  $\{Q^k, A^k\}_{k=1}^m$  generated from it.

### 4.2 End-to-End Training

During training, the retrieval process is integrated into the model’s training loop. The most widely adopted approach to accomplish this is approximate nearest neighbor search (ANNS) efficiently implemented by several libraries like ScaNN (Guo et al., 2020), FAISS (Johnson et al., 2019), etc. These libraries require a fixed set of dense vectors to construct the index and perform a Nearest-Neighbor search using approximate algorithms. However, our memory encoder  $f_\theta$  is continuously updated, which poses great challenges for ANNS index building. REALM (Guu et al., 2020) and RePAQ (Lewis et al., 2021) use an asynchronous index building sub-process to refresh the index every K steps, which is known to be extremely computationally expensive, especially with a large memory. To avoid such expensive computation overhead, we are inspired by TOME (de Jong et al., 2022) to adopt a two-stage training as shown in Figure 4.

**In-Batch Pre-training** In the first stage, instead of using the whole memory, we propose a batch-specific memory that concatenates the positive, random negative, and hard negative en-

tries from each instance in the batch. Assuming we have a batch size of  $B$  containing examples  $\{X_i, \{Q_i^k, A_i^k\}_{k=1}^K\}_{i=1}^B$ . For each example there exist  $K$  positive QA-pairs generated from the given context  $X_i$ . Additionally, we mine  $K$  hard negative QA-pairs  $\{\bar{Q}_i^k, \bar{A}_i^k\}_{k=1}^K$  for each input  $X_i$  to increase retrieval difficulty. This hard negative mining is done with BM25 (Robertson and Zaragoza, 2009) similar to DPR (Karpukhin et al., 2020). We construct the in-batch memory by aggregating the  $K \times B$  positive QA-pairs and  $K \times B$  hard negative memory entries, so the in-batch memory  $M$  contains a total of  $2K \times B$  QA-pairs (roughly a few thousand). Due to the small size of the memory, we can construct the memory index very efficiently. Thus, it enables us to continuously update the memory encoder parameters  $f_\theta$  to achieve strong QA-pair retrieval performance.

**Global Pre-training and Fine-Tuning** In this stage, we first freeze the memory encoder  $f_\theta$  to generate memory-key embedding for the entire memory to build its index. We then incorporate the on-device approximate search algorithm<sup>1</sup> to perform the nearest-neighbor search over the memory index to retrieve the top-K QA-pairs. Formally, we propose to maximize the same objective as Equation 2 but with stop-gradient applied to  $p(m|X)$  term. In this step, the model will only update the query model  $f_\theta$  and the decoder model  $g_\theta$ . During fine-tuning, we follow the same recipe as the global pre-training. Instead of feeding masked passages as inputs, we use questions with pseudo [MASK] token in the front as the input.

### 4.3 Multihop Extension

For our extension model QAMAT+, since the retrieval augmentation process cannot be learned latently, i.e. the gradient propagation is blocked in the concatenation step, we add additional supervision to maximize the groundtruth retrieval probability  $p(m_1|X)$  for the first-round retrieval  $m_1$ . We add such retrieval supervision objective to the original objective  $p(Y|X^1)$ , where  $X^1$  is the retrieval-augmented inputs as described in subsection 3.5.

<sup>1</sup><https://github.com/google-research/language/tree/master/language/mentionmemory>

## 5 QA Experiments

### 5.1 Implementation Details

Our model is based on the T5-base or large architecture implemented in JAX<sup>2</sup> and pre-trained on 32 TPUs on Google Cloud<sup>3</sup>. During in-batch training, our query and index encoder  $f_\theta$  are shared and initialized from the T5 encoder (during global training the index encoder is fixed and the query encoder continues to be updated). Our decoder  $g_\theta$  is similarly initialized from the T5 decoder. In total, we construct  $\sim 60M$  question-answer pairs as the global memory. The memory key is the question tokenized by T5 sentencepiece model into 32 tokens, and the memory value is the answer concatenated with its question tokenized into 40 tokens. The memory is indexed by a pre-computed matrix  $M^k \in \mathbb{R}^{|M| \times d}$  computed based on its keys (questions). The corresponding top-K memory values (question+answer) will be fetched.

During in-batch pre-training, we use a large batch size of 512 and a learning rate of 1e-3, where each example contains a positive Q-A pair and 7 hard negative QA-pairs mined through BM25 (Robertson and Zaragoza, 2009). The in-batch memory contains a total of 4096 entries, we set Top-k of 4 and update over all the modules. After 100K steps of in-batch pre-training, we switch to global pre-training with global memory retrieval. We decrease the batch size to 32 and enlarge Top-K to 16 for larger memory. We update only the query encoder and decoder for another 100K steps. Finally, we set K to 32 to fine-tune on downstream datasets with a decreased learning rate of 5e-4.

### 5.2 Datasets

We evaluate our framework on the three most widely used single-hop open-domain question-answering datasets and two multi-hop open-domain question-answering datasets

**NQ-Open** The NaturalQuestions (Kwiatkowski et al., 2019) dataset consists of naturally occurring Google queries and their answers. We follow Lee et al. (2019) to keep questions that have a "short answer type". It consists of 79168 training examples, 8757 dev examples, and 3610 test examples.

**TriviaQA** The TriviaQA dataset is a collection of trivia question-answer pairs that were scraped from the web (Joshi et al., 2017). We use their unfiltered

<sup>2</sup><https://github.com/google-research/t5x>

<sup>3</sup><https://cloud.google.com/tpu/>

version to evaluate our model consisting of 78785 training, 8837 dev, and 113313 test examples.

**WebQuestions** The WebQuestion dataset contains questions that were sampled from Google Suggest API (Berant et al., 2013). The answers are annotated from FreeBase, the training set contains 3417 examples, the dev set contains 361 examples, and the test set contains 2032 examples.

**HotpotQA** The HotpotQA dataset contains questions generated by human workers by reading two passages (Yang et al., 2018). The questions are designed to require multiple hops and include both bridge questions and comparison questions. The training set contains a total of 90564 examples, the dev-set contains 7405 examples for evaluation.

**Musique** The Musique dataset contains questions created by composing multiple questions from existing single-hop questions and was constructed to contain less bias and artifacts (Trivedi et al., 2021). In our experiments, we consider only the subset of 2-hop questions, resulting in a training set of 14376 examples and a dev set of 1252 examples for evaluation. While the dataset was originally designed as a distractor setting (given a question and a small number of passages, return the answer), we instead consider an open-domain setting.

### 5.3 Baselines

We compare our model with baselines from the following categories. 1) CBQA large language models (T5 XXL), which directly outputs an answer without retrieval. 2) Entity/KG memory-augmented models that use memory attention to incorporate entity-level features into language models (Entities-as-Experts (EaE) (Férvy et al., 2020), Fact-Injected Language Model (FiM) (Verga et al., 2021), MentionMemory (TOME) (de Jong et al., 2022)). 3) Retrieve-and-read model, which retrieves passages to pass to a reader model which predicts the answer. 4) QA-retrieval models, which train a retriever to collect QA-pairs from a large datastore, and then rerank these QA-pairs (top 50-100) with original query with cross-attention. The highest-ranked answer is returned as the final answer.

### 5.4 Single-Hop Results

Our results are summarized in Table 2 which reports exact-match (EM) score.

**Comparison with RePAQ** Our main comparison is with the previous best QA-retrieval-based ap-

proach "RePAQ w/ rerank (XXL ALBERT)". This model has a similar number of parameters to QAMAT (Large). Without using an explicit re-ranking procedure, our model performs slightly worse on NQ but obtains significant gains on TriviaQA and WebQuestion. Especially on WebQuestion, which only contains 3K training examples, RePAQ performs significantly worse than the other datasets because it requires a high volume of examples to update the retriever from scratch. With our proposed pre-training strategy, QAMAT can initialize from a much better checkpoint to decrease the sample complexity, yielding an absolute 6% EM improvement. Additionally, without any fine-tuning, we demonstrate that our model already achieves promising results across these datasets, nearly matching the performance of "RePAQ w/o rerank"<sup>4</sup>.

**Comparison with retrieve-and-read models** In comparison to this class of model, QAMAT roughly matches the performance of RAG, though it still lags behind the SoTA model FiD. However, FiD requires reading 100 passages, i.e. 20K tokens while our best model works more efficiently by only reading top-32 QA-pairs, i.e. 1.2K tokens. To investigate the speed difference between these approaches, we compared their inference speeds using the same hardware (32 Google Cloud v3 TPUs). We found that QAMAT can answer 240 Qs/sec, while FiD only answers 50 Qs/sec, a **5x inference time speedup** over FiD.

### 5.5 Multi-hop Results

Since the document corpora source of HotpotQA and Musique are different from single-hop QA datasets, we adopt question generation model trained on SQuAD 2.0 (Rajpurkar et al., 2018) to generate questions for these two datasets. To create the document corpora, we gather all of the provided positive and negative documents, obtaining 500K passages for HotpotQA and 85K passages for Musique. We then use the trained generation models to populate 3M QA pairs for HotpotQA and 500K QA pairs for Musique. These QA pairs are then used as the memory source for QAMAT+, simulating a (slightly smaller) open-domain setup. When training QAMAT+ on Musique, we initialize from HotpotQA’s in-Batch pre-trained checkpoint, which can bring 5-7% F1 improvement.

<sup>4</sup>It’s worth noting that the question generation models are trained using some of these datasets’ training data so this is not truly “zero-shot” performance.

Model (Test Set)	NQ	TQA	WQ
T5-3B (Roberts et al., 2020)	30.4	35.1	33.6
T5-11B (Roberts et al., 2020)	32.6	42.3	37.2
EaE (Férvy et al., 2020)	-	43.2	-
FILM (Verga et al., 2021)	-	29.1	-
TOME-2 (de Jong et al., 2022)	-	53.4	-
DensePhrases (Lee et al., 2021)	40.9	50.7	-
REALM (Guu et al., 2020)	40.4	55.8	40.7
DPR (Karpukhin et al., 2020)	41.5	57.9	42.4
RAG-Seq (Lewis et al., 2020)	44.5	56.8	45.2
FID (Izacard and Grave, 2021)	48.2	65.0	-
RePAQ (Lewis et al., 2021)	41.2	38.8	29.4†
RePAQ+Rerank (Lewis et al., 2021)	47.6	50.7	37.6†
QAMAT Zero-Shot (Base)	37.9	34.1	25.9
QAMAT Zero-Shot (Large)	39.8	40.0	25.1
QAMAT Fine-tuned (Base)	44.5	53.2	43.0
QAMAT Fine-tuned (Large)	45.5	54.8	43.6

Table 2: The main experimental results on single-hop question answering datasets (NQ=NaturalQuestions, TQA=TriviaQA, WQ=WebQuestions), † means Best-effort replication using our own implementation.

Model (Dev Set F1 Score)	HPQ	MusQ
T5-3B (Roberts et al., 2020)	27.8	7.5
T5-11B (Roberts et al., 2020)	30.2	9.0
MDR+T5-Decoder (Xiong et al., 2020)	<b>62.6</b>	26.8
RePAQ (Lewis et al., 2021)†	47.8	18.6
QAMAT	42.0	16.7
QAMAT+	57.6	<b>29.8</b>

Table 3: The main experimental results on MultiHop QA datasets with QAMAT and QAMAT+, † means Best-effort replication using our own implementation.

In Table 3, we show that QAMAT+ achieves promising results on both multi-hop datasets, outperforming T5-CBQA and RePAQ by a large margin. Additionally, QAMAT+ performs considerably better than the single-hop QAMAT, demonstrating the effectiveness of performing multi-round retrieval. Though QAMAT+ still lags behind the document-based model (MDR+T5 Decoder) on HotpotQA, it surpasses it on the more challenging Musique dataset. These encouraging results suggest the potential for QAMAT+ to perform compositional reasoning over multiple QA-pairs, which greatly increases the coverage of QA datastore to cover more composite factual information.

## 5.6 Ablation Studies

**Number of Retrievals** To understand the properties of our model better, we first investigate the impact of the number of retrievals,  $K$ , on the model’s performance. We gradually increase the  $K$  to col-

Top-K	1	10	20	30
NQ-Recall@K	0.41	0.58	0.62	0.64
TriviaQA-Recall@K	0.46	0.66	0.70	0.72
NQ-EM@K	0.39	0.42	0.44	0.44
TriviaQA-EM@K	0.45	0.51	0.53	0.53

Table 4: The retrieval recall and EM score of different retrieval numbers on test sets.

Pre-training Stages	NQ	TQA	WQ
Only In-Batch	42.1	48.2	39.7
Only Global	26.0	28.9	26.1
In-Batch → Global	44.5	53.2	43.0

Table 5: Downstream EM performance of models when pre-trained using in-batch, global, or both stages.

lect the recall and final QA performance. The results are shown in Table 4. We observe that even though retrieval recall continues to increase beyond  $K > 20$ , the EM score saturates much earlier. Future research could improve performance further by developing decoders to more accurately exploit these larger retrievals sets.

**Importance of Two-Stage Pre-training** We next analyze the importance of the two-stage pre-training from section 4 by removing either the in-batch or global stage. From our results shown in Table 5, we can see that using in-batch pre-training alone leads to a degradation in performance when compared to the two-stage approach. This is likely because the model is never exposed to the full set of hard negatives which will be encountered when performing retrieval over the global memory. On the other hand, if we directly pre-train the global-memory model without any in-batch initialization, the retriever performance is nearly random and the decoder consequently learns to ignore the retrieval and simply memorize question-answer pairs.

## 6 Conclusion

In this paper, we propose a more accurate and efficient architecture to utilize QA-pairs as representation units of knowledge. Our proposed model QAMAT outperforms RePAQ significantly, while leveraging our less expensive training procedure. Furthermore, we show how a QA-backed model can perform compositional reasoning and address more complex queries. In the future, we hope to further close the gap with state-of-the-art document-based retrieve-and-read models and extend this approach to a broader set of tasks.



## Limitations

Our approach has several limitations: 1) we use generated question-answer pairs as a knowledge base, which are extracted from web documents. In order to maintain high quality and faithfulness, the question generation pipeline needs to be well trained with a sufficient amount of clean data. Such conditions might not hold for other domains outside of Wikipedia like biomedical text, thus the general QA-as-Knowledge-Base concept could require additional innovations to extend to other areas. 2) Our latent retrieval learning requires quasi paired data to learn the alignment between the query and memory. This is hard to satisfy in some domains with noisier data or only a very weak alignment between a query and the memory. 3) Our model requires mined intermediate retrieval signals to train QA-MAT+, which currently relies on lexical-overlap-based heuristics. In other cases, this may not be sufficient and instead might require a more principled design to mine better intermediate supervision.

## Ethical Statement

Our work encourages the model to ground on the existing knowledge populated from large textual collections. We believe it is a reasonable towards building more trustworthy and more robust machine learning models. Having better attributions to knowledge source could help humans better understand the model’s rationale for decision making. However, we do admit that the question generation models used to populate the QA knowledge base could potentially exacerbate the biases already present in the original Wikipedia data. We will keep working on this direction to minimize its potential negative impacts.

## References

Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. 2019. Synthetic qa corpora generation with roundtrip consistency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6168–6173.

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*, pages 1533–1544.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican,

George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2021. Improving language models by retrieving from trillions of tokens. *arXiv preprint arXiv:2112.04426*.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017a. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017b. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.

Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGerald, Fei Sha, and William Cohen. 2022. Mention memory: incorporating textual knowledge into transformers through entity mention attention. *International Conference on Learning Representations*.

Bhuvan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig, Ruslan Salakhutdinov, and William W Cohen. 2019. Differentiable reasoning over a virtual knowledge base. In *International Conference on Learning Representations*.

Matan Eyal, Tal Baumel, and Michael Elhadad. 2019. Question answering as an automatic evaluation metric for news article summarization. In *Proceedings of NAACL-HLT*, pages 3938–3948.

Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. Entities as experts: Sparse memory access with entity supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4937–4951.

Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pages 3887–3896. PMLR.

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Papatat, and Mingwei Chang. 2020. [Retrieval augmented language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.

Gautier Izacard and Édouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880.

- Robin Jia, Mike Lewis, and Luke Zettlemoyer. 2021. Question answering infused pre-training of general-purpose contextualized representations. *arXiv preprint arXiv:2106.08190*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- 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 (Volume 1: Long Papers)*, pages 1601–1611.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Tom Kwiattkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Jinhyuk Lee, Mujeen Sung, Jaewoo Kang, and Danqi Chen. 2021. Learning dense representations of phrases at scale. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6634–6647.
- 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.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. Paq: 65 million probably-asked questions and what you can do with them. *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Sewon Min, Jordan Boyd-Graber, Chris Alberti, Danqi Chen, Eunsol Choi, Michael Collins, Kelvin Guu, Hannaneh Hajishirzi, Kenton Lee, Jennimaria Palomaki, et al. 2021. Neurips 2020 efficientqa competition: Systems, analyses and lessons learned. In *NeurIPS 2020 Competition and Demonstration Track*, pages 86–111. PMLR.
- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019. Document expansion by query prediction. *arXiv preprint arXiv:1904.08375*.
- Barlas Oğuz, Kushal Lakhota, Anchit Gupta, Patrick Lewis, Vladimir Karpukhin, Aleksandra Piktus, Xilun Chen, Sebastian Riedel, Wen-tau Yih, Sonal Gupta, et al. 2021. Domain-matched pre-training tasks for dense retrieval. *arXiv preprint arXiv:2107.13602*.
- Liangming Pan, Wenhua Chen, Wenhan Xiong, Min-Yen Kan, and William Yang Wang. 2021. Unsupervised multi-hop question answering by question generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5866–5880.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789.
- 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*, pages 2383–2392.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426.

- Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.
- Devendra Singh, Siva Reddy, Will Hamilton, Chris Dyer, and Dani Yogatama. 2021. End-to-end training of multi-document reader and retriever for open-domain question answering. *Advances in Neural Information Processing Systems*, 34.
- Haitian Sun, Pat Verga, Bhuwan Dhingra, Ruslan Salakhutdinov, and William W Cohen. 2021. Reasoning over virtual knowledge bases with open predicate relations. In *International Conference on Machine Learning*, pages 9966–9977. PMLR.
- Alon Talmor and Jonathan Berant. 2018. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 641–651.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2021. Musique: Multi-hop questions via single-hop question composition. *arXiv preprint arXiv:2108.00573*.
- Pat Verga, Haitian Sun, Livio Baldini Soares, and William Weston Cohen. 2021. Adaptable and interpretable neural memory over symbolic knowledge. In *Proceedings of NAACL-HLT*, pages 3678–3691.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020.
- Jinfeng Xiao, Lidan Wang, Franck Dernoncourt, Trung Bui, Tong Sun, and Jiawei Han. 2021. Open-domain question answering with pre-constructed question spaces. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 61–67.
- Wenhan Xiong, Xiang Li, Srinu Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Scott Yih, Sebastian Riedel, Douwe Kiela, et al. 2020. Answering complex open-domain questions with multi-hop dense retrieval. In *International Conference on Learning Representations*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380.
- Wen-tau Yih, Matthew Richardson, Christopher Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671. Springer.

## A Question Answer Pairs as Knowledge Base

We can see QA-pairs as a virtual knowledge graph, where the question template defines the relation, the topic entity in the question defines the head entity node, and the answer denotes the tail entity. A typical example is given in Figure 5, such compositionality makes the QA-pair more controllable and easy to reason over than documents.

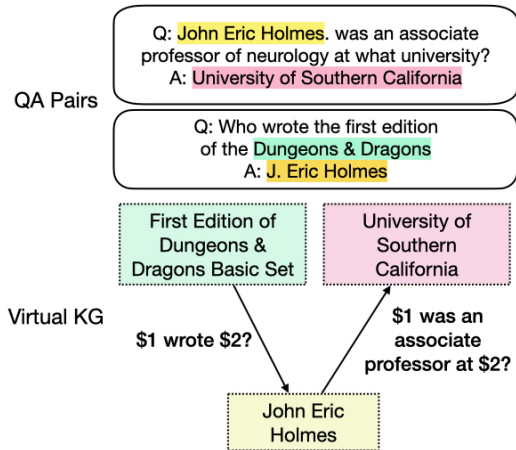


Figure 5: QA pairs can be seen as virtual knowledge base, where the question can represent complex relations connecting subject and answer.

## B Question Generation

Here, we use existing SQuAD datasets’  $\langle Q, A, \text{Document} \rangle$  triples (Rajpurkar et al., 2016) to train answer extraction, question generation model.

**Answer Extraction** Specifically, our answer extraction model takes a document as the input and trains an encoder-decoder model to generate a potential answer. We use beam search over the trained model to find the highest-likely answers in the given document. In our experiment, the answer extraction model is trained with the SQuAD dataset, where the document is given as the input, and the answer spans are the prediction targets.

**Question Generation** For the question generation model, we take the SQuAD dataset and use document + extracted answer as the input to generate questions as the outputs. This step is also accomplished by an encoder-decoder model. which is mainly purposed for reading comprehension problems, where the annotated questions are highly correlated with the document containing very few hallucinations. However, the questions in SQuAD (Ra-

ipurkar et al., 2016) could be contextualized or ambiguous, which could lead to ambiguity problems to hurt the retrieval performance. Therefore, we add question filtering to select the most accurate QA pairs.

**Question Filtering** For the question filtering model, we take the document + generated question to generate an answer. We compare the predicted answer vs. the original answer to see if they match each other. If not, the QA-pair will be filtered based on such inconsistency. We use a reading comprehension model trained with SQuAD to predict the answer. The predicted answer based on the document will match with the original QA-pair to decide its consistency. Such an option runs much faster, providing much higher recall but lower precision compared to the open-domain FiD filtering used in (Lewis et al., 2021).

We visualize our question generation pipeline mentioned above in Figure 6.

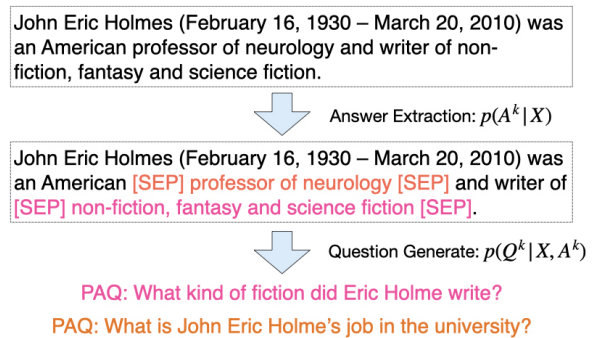


Figure 6: Question generation pipeline: Answers are extracted from passages and then questions are generated conditioned on that contextualized answer. This procedure is used to generate both our model’s QA memory and our pre-training data.

## C Ablation Study

We experiment with two variants of memory to see their performance difference.

### C.1 PAQ memory

The first version is the standard PAQ corpus (Lewis et al., 2021) containing 65M QA pairs, where these QA-pairs are generated by models trained on NQ (Kwiatkowski et al., 2019) and filtered through FiD model (Izacard and Grave, 2021) also trained on NQ (Kwiatkowski et al., 2019). This memory is highly precise due to ODQA-filtering process, however, it only covers information from 9M out of the 20M passage blocks used in DPR (Karpukhin et al., 2020).



Our memory contains 30M PAQ corpus being de-duplicated, i.e. only one question corresponds to an answer span. We generate 30M additional QA-pairs based on the left-out 10M documents from PAQ (Lewis et al., 2021) and add these complementary QA-pairs to form our 60M memory to increase the coverage. However, since our filtering procedure is based on reading comprehension, the precision of QA-pairs is lower than the original PAQ memory.

Memory	NQ	TriviaQA	WebQuestions
PAQ 65M	44.7	48.0	39.4
Ours 60M	44.5	53.2	43.0

Table 6: Impact of different memory over the downstream QA dataset performance.

As can be seen, from Table 6, using the most precise but low-coverage PAQ memory from PAQ (Lewis et al., 2021) yields the worse results on TriviaQA and WebQuestions. After adding an additional 30M PAQs to the memory generated by our pipeline, we are able to achieve 4-5% improvements on these two datasets while still maintaining NQ’s performance.

## C.2 Size of Pre-training Corpus

Next, we investigate the impact of the size of the pre-training corpus. As a baseline, we repurpose the aligned query-passage corpus used to train DPR (Karpukhin et al., 2020) which we adapt to our setting by simply reversing the pairs (120K passage -> question retrieval). Additionally, we vary the size of generated pre-training corpus (from 1M to 20M instances) to see its impact on the model’s final downstream performance. From Table 7, we can see that the smaller-sized pre-training corpus can drastically reduce the model’s performance, with up to a 5% drop seen on TriviaQA.

Pre-train Examples	NQ	TQA	WQ
120K	42.5	48.2	39.7
1M	42.8	48.8	40.2
5M	43.8	51.5	41.7
10M	44.3	52.1	42.5
20M	44.5	53.2	43.0

Table 7: Impact of pre-training corpus size on final downstream EM performance. The upper portion is pre-trained using the DPR-reverse corpus described in subsection 5.6 and the lower portion uses subsets of our generated pre-training corpus (subsection 4.1)




Figure 7: The impact of memory size on downstream QA EM performance.

**Size of Memory** Finally, we look at how big of memory we need to reach optimal downstream accuracy and how the model behaves with a smaller memory. As is shown in Figure 7, having a small memory of less than 5M entries does not improve over a model with no memory at all. Due to the lack of coverage, the model does not receive a useful signal from the retrieval and is subsequently not incentivized to utilize those retrievals when making a prediction. However, once the size of the memory increases beyond 15M we observe a steep increase in the final performance, indicating that the model is gradually learning to incorporate retrieved information retrievals to assist prediction.

## D MultiHop QA Training

In order to train the multi-hop QA model, we need to have intermediate supervision for the query augmentation process. Here we use a string-based match to derive what are the most possible intermediate questions from a collection of pre-generated QA pairs. We depict the mining process as Figure 8.

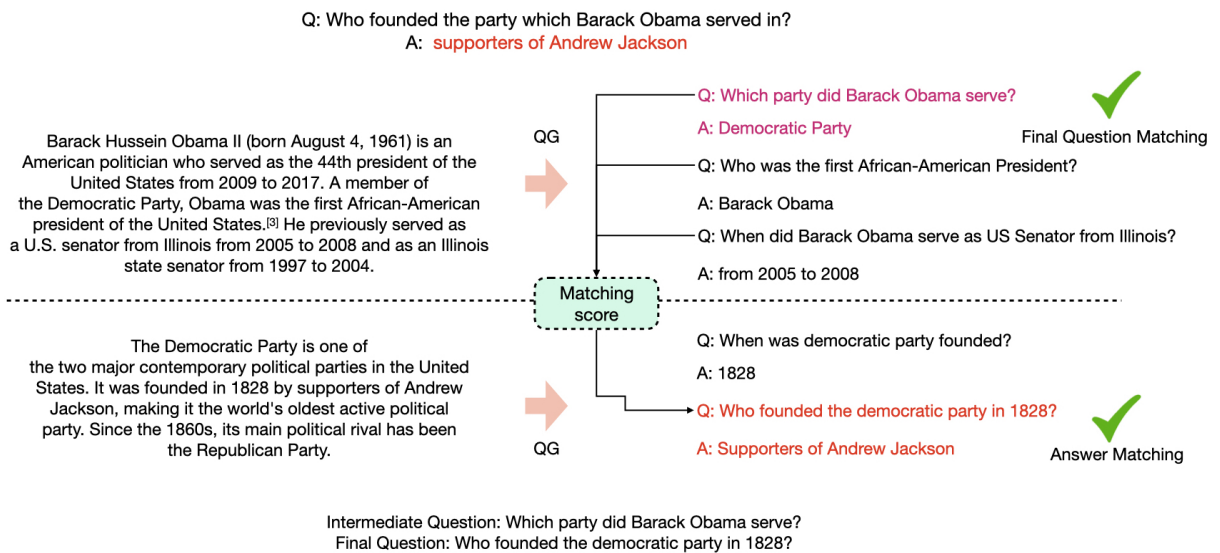


Figure 8: We first find the final question based on answer string matching with the pre-generated question, and then base on that to trace back the intermediate question.