



HAL
open science

A Memory Efficient Baseline for Open Domain Question Answering

Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola de Cao, Sebastian Riedel,
Edouard Grave

► To cite this version:

Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola de Cao, Sebastian Riedel, et al.. A Memory Efficient Baseline for Open Domain Question Answering. 2021. <hal-03463488>

HAL Id: hal-03463488

<https://hal.science/hal-03463488v1>

Preprint submitted on 2 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

A Memory Efficient Baseline for Open Domain Question Answering

Gautier Izacard^{1,2,3} Fabio Petroni¹ Lucas Hosseini¹

Nicola De Cao^{1,4} Sebastian Riedel^{1,5} Edouard Grave¹

¹ Facebook AI Research ² ENS, PSL University ³ Inria

⁴ University of Amsterdam ⁵ University College London

gizacard|fabiopretroni|hoss|ndecao|sriedel|egrave@fb.com

Abstract

Recently, retrieval systems based on dense representations have led to important improvements in open-domain question answering, and related tasks. While very effective, this approach is also memory intensive, as the dense vectors for the whole knowledge source need to be kept in memory. In this paper, we study how the memory footprint of dense retriever-reader systems can be reduced. We consider three strategies to reduce the index size: dimension reduction, vector quantization and passage filtering. We evaluate our approach on two question answering benchmarks: TriviaQA and NaturalQuestions, showing that it is possible to get competitive systems using less than 6Gb of memory.

1 Introduction

The goal of open-domain question answering is to answer factoid questions about general topics (Voorhees et al., 1999). Since the evidence to answer the question is not provided to the system, a standard approach is to use an external source of knowledge such as Wikipedia. These systems, known as *retriever-reader*, start by retrieving relevant documents, or passages, from the knowledge source, and then process them with a reader model to produce the answer (Chen et al., 2017).

The downstream performance of these pipeline systems greatly depends on the quality of the retrieval module. Traditionally, queries and documents were represented as sparse vectors based on hand crafted term weighting such as BM25 (Robertson et al., 1995). More recently, dense representations obtained with neural networks have proven very effective for question answering, and related tasks. A limitation of these methods is the size of the index, which can be tens of gigabytes, as all the vector representations of the knowledge source

have to be kept in memory. This lead to the following question: *what is the more efficient way to store information for question answering?*

In this paper, we explore how non-parametric approaches to store information can be used for memory efficient question answering systems. More precisely, starting from a state-of-the-art pipeline using a dense retriever, we study how the size of the index can be reduced. We consider three main strategies to do so: dimension reduction of the dense representations, vector quantization and document filtering. We study the impact of these three techniques on the downstream performance, using two open-domain question answering benchmarks: TriviaQA and NaturalQuestions. We show that, with minimal loss of performance, it is possible to obtain QA systems using less than 6Gb of memory.

2 Related Work

Retriever-Reader Approach. Given a question, retriever-reader systems first retrieve relevant documents from an external knowledge source (Chen et al., 2017). These documents, along with the question, are then processed by a second model called the reader to produce the answer. In case where no gold spans are available, Clark and Gardner (2018) proposed a loss function based on global normalization over all the mentions of the answer. Wang et al. (2019) applied this approach to BERT pre-trained models, while Min et al. (2019) introduced an alternative method based on expectation-maximization. Roberts et al. (2020) used large pre-trained generative models, generating the answer without retrieving external documents, a setting referred to as *closed book*. Several works showed that generative QA models could be improved by retrieval, using the obtained documents as input of the sequence-to-sequence model (Min et al., 2020; Lewis et al., 2020; Izacard and Grave, 2020b).

Dense Embedding Retrieval. Traditional information retrieval is based on sparse vector representations, with hand crafted term weighting (Jones, 1972; Robertson et al., 1995). Given a query, a similarity score such as the dot product is used to retrieve the most relevant documents. Recently, it has been shown that dense representations, obtained with neural networks, were a good alternative to sparse representations. Different training schemes of the retriever have been considered, based on supervised learning (Karpukhin et al., 2020), latent modeling (Lee et al., 2019; Guu et al., 2020) or distillation (Izacard and Grave, 2020a). Finally, Luan et al. (2020) studied the impact of design choices, such as document length or vector dimension, on the retrieval performance. We differ from this work by evaluating on the downstream QA task, and exploring other techniques like quantization and document filtering.

Product Quantization. The process of quantization (Gray and Neuhoff, 1998) consists in mapping values from a continuous or large space (eg. `float32`) to a smaller discrete set (eg. `int8`). This process is the basis of most lossy data compression techniques. In vector, or block, quantization, all the elements of a vector are quantized simultaneously instead of independently. One of the simplest vector quantization method is to use k -means, and map each vector to the closest centroid. A better approach in high-dimension, known as product quantization (Jegou et al., 2010), is to subdivide each vector of dimension d into n sub-vectors of dimension d/n and to quantize them independently, using k -means.

3 System description

We briefly describe our topline QA system, following the retriever-reader approach.

For the reader we use a Fusion-in-Decoder model (Izacard and Grave, 2020b). This type of model leverages the sequence-to-sequence architecture to efficiently combine multiple passages in order to generate an answer. In the Fusion-in-Decoder model, the encoder of the sequence-to-sequence model is first applied independently on each passage concatenated with the question. Then, the representations are concatenated and the decoder is applied on the resulting representations to predict an answer. We initialize the Fusion-in-Decoder model with T5 (Raffel et al., 2019).

For the retriever, we use an approach based on

dense embeddings (Karpukhin et al., 2020). First, each passage p of the knowledge source is mapped to a d -dimensional vector using an embedding function E . Then, given an input question q , it is represented with the same embedding function. A similarity score between passage p and question q is obtained by the dot product between the two representations $S(q, p) = E(q)^T E(p)$. Passages with the highest similarity scores are then retrieved and processed by the reader. This operation is done efficiently by using a maximum inner product search library such as Faiss (Johnson et al., 2019) after pre-indexing all Wikipedia passages.

The retriever is trained with knowledge distillation, where the synthetic labels are obtained by aggregating the attention scores of the reader model (Izacard and Grave, 2020a). We start by training a reader with the passages obtained with DPR (Karpukhin et al., 2020). Then, a new retriever is obtained by distilling the aggregated cross-attention scores of the reader to the retriever. Finally, new passages are obtained with this retriever, and used to train a new reader model.

4 Compressing QA systems

In this section, we discuss the size of our topline system, and techniques to reduce it.

4.1 Initial system size

Models. Our system uses two neural models, the retriever and the reader. The retriever is based on the BERT base architecture, containing 110M parameters and thus weighting 0.22Gb in `float32`. The reader is based on T5 base and large, containing 220M and 770M parameters respectively, thus weighting 0.88Gb and 3.1Gb in `float32`.

Knowledge source. Our system follow standard open-domain QA practice and uses Wikipedia as knowledge source. We use the English Wikipedia¹ dump from Dec. 20, 2018, including lists, which weighs approximately 21Gb uncompressed. This can be reduced to 3.8Gb by using the compression softward `xz`, or even 3.1Gb with `lrzip`.

Dense index. Both dense and sparse representations lead to a large index, using a significant memory footprint. Following Karpukhin et al. (2020), we split Wikipedia into non-overlapping chunks of 100 words. Then, each chunk is represented by a vector of dimension 768 obtained with the

¹Under <https://creativecommons.org/licenses/by-sa/3.0/>

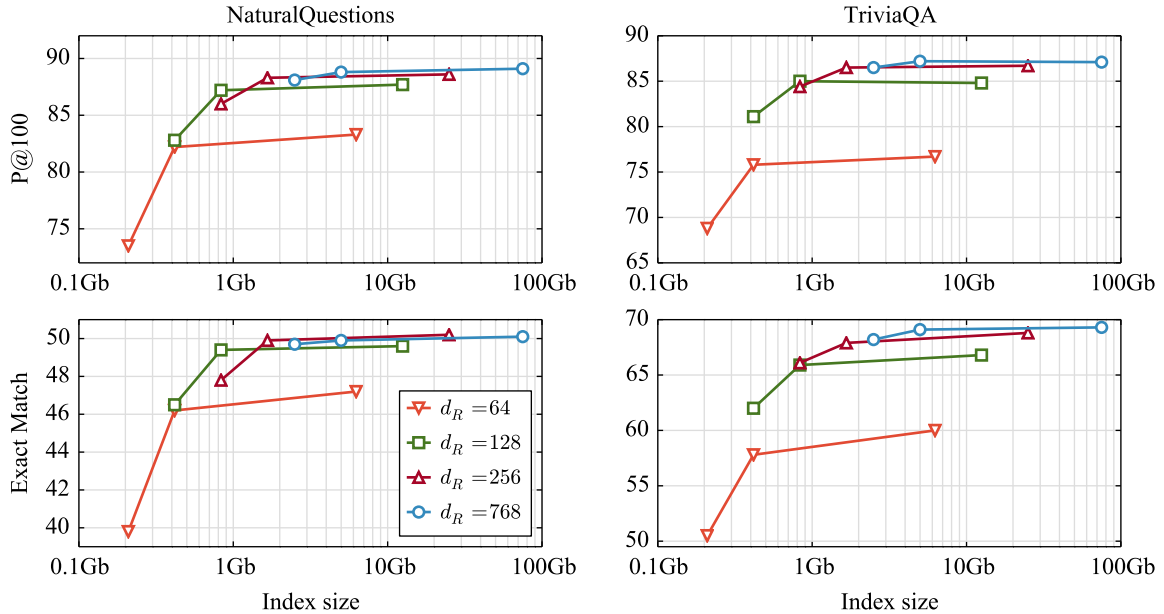


Figure 1: Performance as a function of index size, for different dimension and quantization. Each curve corresponds to a fixed dimension, the left point using 1 bit per dimension, the middle point 2 bit and the right point 32 bit.

retriever model. Assuming that vectors are stored in `float32`, this leads to an index size of 75Gb.

4.2 Reducing the index size

In the following, we discuss strategies to reduce the size of QA systems. As noted previously, most of the memory footprint is allocated to the index, and we will thus focus on reducing its size.

Dimension reduction. One of the simplest way to reduce the index size is to reduce the dimension d of the dense representations of passages. We consider to reduce the embedding dimension at train time, by adding a linear layer that maps the output of the network to a d_R -dimensional vector, followed by a LayerNorm. This follows the same solution adopted in (Luan et al., 2020), where the trade-off between the dimension reduction and the retrieval performance is studied. We empirically verified that adding the LayerNorm led to stronger performance than without it. An alternative strategy would be to use principal component analysis on the set of already trained embeddings.

Product quantization. A second way to reduce the index size is to use vector quantization, which is complementary to dimension reduction. Standard vectors are stored in memory by using 4 bytes per dimension when using `float32`. Without loss of performance, it is possible to use `float16`, lead-

ing to a compression factor of two. To compress representations more aggressively, we propose to use product quantization. In that case, embeddings are divided in n_v sub-vectors of dimension d/n_v , each being stored using n_b bits. If we use sub-vectors of dimension two and use one byte to store them, then, each dimension is using 0.5 bytes. This leads to a compression factor of eight compared to the original index.

Passage filtering. A last method to reduce the index size is to remove documents that are unlikely to be useful for question answering. To do so, we train a linear classifier, where each Wikipedia article is represented by its title and list of categories. As positive examples, we use articles retrieved with DPR on the training data. To obtain negative examples, we use self-training: we start by randomly sampling Wikipedia articles as negatives, and train a first classifier. Then, the articles classified as negative with the highest confidence are used as negatives for the next iteration of self-training. We perform a few iterations of this scheme, and apply the final classifier to filter the Wikipedia dump.

5 Experiments

We conduct experiments on two datasets: NaturalQuestions (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017) and follow the standard setting for open-domain QA (Karpukhin et al., 2020).




Figure 2: Performance as a function of the number of passages, in millions (full wikipedia contains 26M).

The final end-to-end performance of the system is evaluated using the Exact Match score. We also report the top-k retrieval accuracy (P@k): the percentage of questions for which at least one of the top-k retrieved passages contains the gold answer.

5.1 Technical details

Retriever. We initialize the retriever with the uncased BERT base model (Devlin et al., 2019), and use the same model to embed questions and passages. The representations of the final layer are averaged instead of using the representation corresponding to the [CLS] token. Furthermore, similarly to Khattab et al. (2020), we pad questions and passages to fixed-length sequences without using mask when applying the embedder function. We use sentence of length 40 tokens for the questions, and 200 tokens for the passages.

Reader. The reader model is initialized with a T5-base model unless otherwise specified. We use 100 retrieved passages per question to train the model, using AdamW (Loshchilov and Hutter, 2019) with a batch size of 64. We train the model with a peak learning rate of 10^{-4} . The learning rate increases linearly during 600 gradient steps, and decreases linearly during 14.4k steps. The best model is selected based on the exact match score evaluated on the dev set every 500 gradient steps.

5.2 Dimension reduction vs. quantization

We start by investigating if dimension reduction and product quantization are complementary, and what setting works best. For this, we train retriever-reader systems with different dimensions d_R of the index. Then, we use product quantization to compress the pre-computed index. The number of bits per subvector is fixed at 8, and we vary the number of subvectors. Results are reported in Fig. 1. It appears that dimension reduction and quantization are

Model	size	NQ	TriviaQA
Topline base	79Gb	50.4	69.8
Topline large	81Gb	54.7	73.3
Compressed base	2.1Gb	44.0	56.8
Compressed large	5.1Gb	53.6	71.3

Table 1: Test set accuracy of compressed systems and our topline models.

complementary, as pareto optimal systems often involve both. Combining them allows to obtain an index of reasonable size, 1.6Gb, with little accuracy loss: -0.2EM on NQ and -1.1EM on TriviaQA.

5.3 Passage filtering

Filtering passages allows to reduce the size of both the index and the knowledge source. In Figure 2 we progressively discard articles starting from the full Wikipedia containing 26M passages. We report results using two retriever-reader pipelines trained in the previous subsection with $d_R = 256$ and $d_R = 128$. For the first one we use a quantized index with 64 subvectors per passage, this leads to an index of size 1.67Gb for the complete Wikipedia. The second uses 16 subvectors per passage and its original size on the whole Wikipedia is 0.42Gb. As shown here it is possible to discard a significant fraction of Wikipedia articles.

5.4 Final models

Finally, we report in Table 1 the test performance of two compressed systems, and compare them to our state-of-the-art topline models. The first system, at 5.1Gb, is made of a reader large, index of dimension 256 with 2 bits per dimension and 18M of passages. The second system, at 2.1Gb, is made of a reader small, index of dimension 128 with with 2 bits per dimension and 10M passages. We observe that these systems, while significantly smaller than the topline, obtain competitive performance.

Discussion

In this paper, we explored how to compress retriever-reader pipelines based on dense representations. We studied different strategies, which are complementary, such as dimension reduction, vector quantization and passage filtering. When combining these different methods, we showed that it is possible to obtain systems smaller than 6Gb, which are competitive with the state-of-the-art.

References

- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Proc. ACL*.
- Christopher Clark and Matt Gardner. 2018. Simple and effective multi-paragraph reading comprehension. In *Proc. ACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL*.
- Robert M. Gray and David L. Neuhoff. 1998. Quantization. *IEEE transactions on information theory*, 44(6):2325–2383.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*.
- Gautier Izacard and Edouard Grave. 2020a. [Distilling knowledge from reader to retriever for question answering](#).
- Gautier Izacard and Edouard Grave. 2020b. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*.
- Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*.
- Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proc. ACL*.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.
- Omar Khattab, Christopher Potts, and Matei Zaharia. 2020. Relevance-guided supervision for openqa with colbert.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: a benchmark for question answering research. *TACL*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. In *Proc. ACL*.
- Patrick Lewis, Ethan Perez, Aleksandara 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. *arXiv preprint arXiv:2005.11401*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2020. Sparse, dense, and attentional representations for text retrieval. *arXiv preprint arXiv:2005.00181*.
- Sewon Min, Danqi Chen, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. A discrete hard EM approach for weakly supervised question answering. In *Proc. EMNLP-IJCNLP*.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. Ambigqa: Answering ambiguous open-domain questions. *arXiv preprint arXiv:2004.10645*.
- 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*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at TREC-3. *NIST Special Publication Sp*.
- Ellen M Voorhees et al. 1999. The TREC-8 question answering track report. In *TREC*.
- Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019. Multi-passage BERT: A globally normalized BERT model for open-domain question answering. In *Proc. EMNLP-IJCNLP*.