

Designing a Minimal Retrieve-and-Read System for Open-Domain Question Answering

Sohee Yang^{1*} and Minjoon Seo^{1,2}

¹KAIST AI ²NAVER Corp.

{sohee.yang, minjoon}@kaist.ac.kr

Abstract

In open-domain question answering (QA), retrieve-and-read mechanism has the inherent benefit of interpretability and the easiness of adding, removing, or editing knowledge compared to the parametric approaches of closed-book QA models. However, it is also known to suffer from its large storage footprint due to its document corpus and index. Here, we discuss several orthogonal strategies to drastically reduce the footprint of a retrieve-and-read open-domain QA system by up to 160x. Our results indicate that retrieve-and-read can be a viable option even in a highly constrained serving environment such as edge devices, as we show that it can achieve better accuracy than a purely parametric model with comparable docker-level system size.¹

1 Introduction

Open-domain question answering (QA) is the task of finding answers to generic factoid questions. In recent literature, the task is largely approached in two ways, namely *retrieve & read* and *parametric*. The former solves the problem by first retrieving documents relevant to the question from a large knowledge source and then reading the retrieved documents to find out the answer (Lee et al., 2019; Guu et al., 2019; Karpukhin et al., 2020; Lewis et al., 2020; Izacard and Grave, 2021). The latter, also known as closed-book QA, generates the answer in a purely parametric end-to-end manner (Brown et al., 2020; Roberts et al., 2020).

While a parametric model enjoys the benefit in terms of system size that they do not require additional knowledge source like a retrieve & read system does, its fundamental limitations are that their predictions are not so interpretable and they are not suitable for dynamic knowledge source as

* Most of the work was done while the author was working at NAVER Corp.

¹Our code and model weights are available in <https://github.com/clovaai/minimal-rmr-qa>.




Figure 1: System footprint vs. Exact Match (EM) accuracy on EfficientQA dev set. System footprint is measured by the command `du -h /` inside the standalone docker container as stated in the EfficientQA competition guideline. The red plot from left to right shows a path of reducing the size of an open-domain QA system with DPR from 77.5GB to 484.68MB by successively applying each of the strategies in Section 2. The storage footprints of the baseline systems with T5 are calculated assuming the use of the lightweight docker image and post-training compression methods applied to our system.

it is difficult to add, remove, or edit knowledge in the parametric model. These limitations are well-addressed by the retrieve & read mechanism, which makes it often more suitable for real-world products. However, it is known to suffer from its large storage footprint due to its document corpus and index, especially compared to the parametric model that only needs to store the parameters (Izacard et al., 2020; Fajcik et al., 2021; Lewis et al., 2021).

Building an interpretable and flexible open-domain QA system and reducing its system size are both important in real-world scenarios; the system must be able to quickly adapt to the changes of the world and be deployed in a highly constrained serving environment such as edge devices. Hence, to get the best of both worlds, it is worthwhile to explore the trade-off between the storage budget and the accuracy of a retrieve & read system.

Well-known approaches for reducing the size of a neural network include pruning (Han et al., 2016),

quantization (Zafzir et al., 2019), and knowledge distillation (Hinton et al., 2014). In this paper, we utilize some of these generic approaches and combine them with problem-specific techniques to size down a conventional retrieve & read system. We first train a passage filter and use it to reduce the corpus size (Section 2.1). We further apply parameter sharing strategies and knowledge distillation to make a single-encoder lightweight model that can perform both retrieval and reading (Section 2.2). In addition, we adopt multiple engineering tricks to make the whole system even smaller (Section 2.3).

We verify the effectiveness of our methods on the dev set and test set of EfficientQA² (Min et al., 2021). By applying our strategies to a recent extractive retrieve & read system, DPR (Karpukhin et al., 2020), we reduce its size by 160x with little loss of accuracy, which is still higher than the performance of a purely parametric T5 (Roberts et al., 2020) baseline with a comparable docker-level storage footprint. In Appendix A.5, we also report the performance on two more open-domain QA datasets, Natural Questions (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017), to test the generalizability of our methods and suggest a future research direction.

2 Method

In this section, we discuss three techniques for reducing the storage footprint of a generic retrieve & read system, namely passage filtering (Section 2.1), unifying retriever and reader into a single model through parameter sharing (Section 2.2), and post-training compression (Section 2.3). We assume that the initial system takes the conventional composition of a trainable (neural) retriever with a question encoder and a passage encoder that create dense vectors used for search, a neural extractive reader (possibly with passage ranking), and a text corpus and the corresponding index that serve as the knowledge source. Figure 1 shows how we start from one such retrieve & read system and apply each of the methods, in the order they are introduced in this section, to successively reduce its system footprint without sacrificing much accuracy.

2.1 Passage Filtering

Index and corpus files can take up a significant portion of the storage footprint of a retrieve & read

²A recently-hosted open-domain QA challenge at NeurIPS 2020 in which 18 teams participated. <https://efficientqa.github.io>

system if a large text corpus is utilized as the knowledge source. Therefore, to drastically reduce the system size, we train a binary classifier and use it to exclude passages that are relatively unlikely to be useful for question answering.

Let the set of indices of all passages in the corpus be J_{total} . To create the training data, we split J_{total} into two disjoint sets, J_{train}^+ and J_{train}^- , such that the former contains the indices of the passages we would like to include in the minimal retrieve & read system.³ Denoting $E(\cdot)$ as a trainable dense encoder which maps a passage to a d -dimensional embedding, such that $v_j = E(p_j) \in \mathbb{R}^d$, the score $s_j = v_j^\top w$, where $w \in \mathbb{R}^d$ as a learnable vector, represents how likely a passage p_j would hold the answer to an input question. The classifier is trained with binary cross entropy on the minibatches of half-positive and half-negative passages, $J_{\text{train}}^{+'}$ and $J_{\text{train}}^{-'}$ drawn from J_{train}^+ and J_{train}^- , respectively.

During training, we sample several checkpoints and evaluate them using the hit ratio on a validation set: $\text{hit}_{\text{val}} = |J_{\text{val}}^{[1:|J_{\text{val}}^+|]} \cap J_{\text{val}}^+| / |J_{\text{val}}^+|$, where J_{val}^+ is a set of indices of the ground truth passages that hold the answer for the questions in the validation set and $J_{\text{val}}^{[1:|J_{\text{val}}^+|]}$ is the set of indices j of the passages whose inferred score s_j is in the top- $|J_{\text{val}}^+|$ scores sorted in descending order, among all s_j such that $j \in J_{\text{val}}^+ \cup J_{\text{val}}^-$. J_{val}^- is a disjoint set randomly sampled from J_{train}^- .

We select the checkpoint with the highest hit_{val} and calculate s_j for all p_j , where $j \in J_{\text{total}}$, using the selected checkpoint. Then, we retrieve $J_{\text{subset}} = J_{\text{total}}^{[1:n]}$, the set of indices of the n top-scoring passages, to indicate the passages to include in our minimal retrieve & read system.

2.2 Retriever-Reader with Single Encoder

In this subsection, we introduce how to obtain a unified retriever-reader with a single encoder (which results in a smaller system footprint) that can perform both retrieval and reading without much drop in accuracy. The unified retriever-reader is trained by successively applying (1) retriever encoder sharing, (2) distilling a reader into the retriever-reader network, and (3) iterative finetuning.

2.2.1 Lightweight Encoder and Embedding Dimension Reduction To make the system small, we utilize a lightweight pretrained encoder. Specifi-

³All the details, including how we split the data, are in Appendix A.1.

cally, MobileBERT (Sun et al., 2020) (4.3x smaller than BERT-base (Devlin et al., 2019)) is employed as the encoder of our retriever-reader model.

We use the dense embedding vectors of the passages in the knowledge source as the index. Therefore, reducing the embedding dimension results in a linear decrease in the index size. We use only the first 128 dimensions (out of 512) to encode the questions and passages.

2.2.2 Retriever Encoder Sharing Let $E_\psi(\cdot)$ and $E_\phi(\cdot)$ be the the question encoder and passage encoder of a retriever, where each of the encoders produces a vector for question q and passage p .

We share the parameters of the encoders, so that $\psi = \phi = \theta$, and differentiate the question inputs from passages inputs using an additional input signal: different token type ids of 0 for questions and 1 for passages. The retrieval score for a pair of question q and passage p is calculated as $\text{sim}_\theta(q, p) = E_\theta(q, 0) \cdot E_\theta(p, 1)$.

We minimize the negative log-likelihood of selecting the passage which holds the answer, namely the positive passage, while training on mini-batches that consist of questions that are each paired with one positive passage and several negative passages. This procedure creates a retriever with a single encoder of parameters θ that can encode both questions and passages.⁴

2.2.3 Unified Retriever-Reader Through Knowledge Distillation The previous subsection describes how to make a retriever that holds only one encoder. Here, we further train the parameters of the retriever θ so that it can also acquire the ability of a reader; we make a unified retriever-reader model that shares all the encoder parameters and eliminate the need for a separate reader. Specifically, using a fully trained reader of parameters ξ as the teacher, we adopt knowledge distillation to transfer its reading ability to the unified retriever-reader network. The training starts after initializing the parameters of the retriever-reader as θ , which is obtained from the retriever encoder sharing procedure described in the previous subsection.

Let $J_{\text{read}} \subset J_{\text{subset}}$ be the set of indices of the passages whose retrieval score $\text{sim}_\omega(q, p_j)$, calculated

⁴In a setting where the index is frozen (addition or editing of index items does not occur), the system does not need a passage encoder. However, we assume a self-contained system with the full ability to update the index, so the passage encoder is considered in the system composition.

for question q using a retriever with parameters ω ⁵, is among the top- k' scores for all $j \in J_{\text{subset}} \setminus J_{\text{read}}$ serves as the candidate pool of the indices of the training set passages.

During training, for question q , a set of passages $P_q = \{p_i | 1 \leq i \leq m\}$ where $m \geq 2$ is sampled from $\{p_j | j \in J_{\text{read}}\}$ to construct a part of the training batch, such that only p_1 contains the answer to question q among $p_i \in P_q$.

Then, we train the unified retriever-reader network with parameters θ using a multitask loss $\mathcal{L}_{\text{read}} + \mathcal{L}_{\text{ret}}$, such that the former is used to train the reader part of the network, and the latter is used to keep training the retriever part. The resulting retriever-reader model has the ability to perform both retrieval and reading.

$\mathcal{L}_{\text{read}}$ is designed to distill the knowledge of a reader teacher into the reader part of the retriever-reader student; the KL divergence between the sharpened and softmaxed answer span scores of the teacher and the student, $D_{KL}(\mathcal{P}_{\xi, q}^{\text{span}} || \mathcal{P}_{\theta, q}^{\text{span}})$. If the teacher reader additionally contains a passage ranker, distillation is also jointly done on the passage ranking scores (m -dim vector outputs).

Retrieval loss \mathcal{L}_{ret} is jointly optimized in a multitask-learning manner to prevent the retriever part of the unified network from forgetting the retrieval ability while training the reader part. The loss can either be the negative log-likelihood described in the previous subsection or another knowledge distillation objective function with a fully trained retriever teacher. If the reader teacher used for $\mathcal{L}_{\text{read}}$ has a passage ranker, the passage ranking score of the teacher can serve as the distillation target (Yang and Seo, 2020).

2.2.4 Iterative Finetuning of Unified Retriever-Reader We have observed that finetuning the unified retriever-reader for a few more epochs leads to better retrieval and reading performance. While the most simple method is to jointly train the model on the standard reader loss and retriever loss⁶, we additionally try iterative finetuning of each of the retriever and reader part as described in Algorithm 1. The motivation here is to apply a loose reconstruction constraint $\mathcal{L}_{\text{recon}}$ to keep the retrieval score as it is before and after the model is optimized for reading, with an assumption that this would be helpful

⁵The retriever with parameters ω is the retriever used with the teacher reader of parameters ξ .

⁶The marginal negative log-likelihood of all the correct answer spans in the positive passage and the negative log-likelihood of positive passage p_1 being selected, respectively.

to alleviate the train-inference discrepancy in the input distribution of the reader, created because the unified retriever-reader is not trained using a pipelined manner (training the reader on top of the retrieval result of a fixed retriever).

Algorithm 1 A single iterative finetuning step on the unified retriever-reader with parameters θ at time t

Input $\theta^{(t)}$ (parameters of the model at time t), knowledge distillation temperature τ , and training batch of question q and passages $P_q = \{p_i | 1 \leq i \leq m\}$ drawn from J_{read} such that $m \geq 2$, $Y(q, p_1) = 1$, and $Y(q, p_i) = 0, \forall 2 \leq i \leq m$. (batch size of 1 is assumed here for a simple presentation)

Output Updated parameters $\theta^{(t+1)}$

- 1: $\ell^{(t)} \leftarrow [E_{\theta}^{(t)}(p_1, 1), \dots, E_{\theta}^{(t)}(p_m, 1)]^T E_{\theta}^{(t)}(q, 0)$
 - 2: $\hat{\theta}^{(t)} \leftarrow \text{GradientUpdate}(\mathcal{L}_{\text{read}}(q, p_1, \dots, p_m); \theta^{(t)})$
 - 3: $\hat{\ell}^{(t)} \leftarrow [E_{\hat{\theta}}^{(t)}(p_1, 1), \dots, E_{\hat{\theta}}^{(t)}(p_m, 1)]^T E_{\hat{\theta}}^{(t)}(q, 0)$
 - 4: $\mathcal{L}_{\text{recon}} \leftarrow D_{\text{KL}}(\text{softmax}(\ell^{(t)}/\tau) || \text{softmax}(\hat{\ell}^{(t)}/\tau))$
 - 5: $\mathcal{L}_{\text{nl}} \leftarrow \text{CrossEntropy}(\text{softmax}(\hat{\ell}^{(t)}), Y)$
 - 6: $\theta^{(t+1)} \leftarrow \text{GradientUpdate}(\mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{nl}}; \hat{\theta}^{(t)})$
-

2.3 Post-Training Compression Techniques

In addition to the training methods to decrease the corpus, index, and model size, several post-training engineering tricks are applied to compress the system footprint further: (1) INT8 quantization of index items, (2) saving model weights as FP16, (3) resource compression, and (4) utilizing token IDs as the corpus instead of raw texts.

INT8 Quantization of Index Items The dense embeddings that serve as the items in the search index are of type FP32 in the default state. INT8 quantization can be applied to reduce the index size by four times with a little bit of drop in the accuracy. We make use of the quantization algorithm implemented in FAISS (Johnson et al., 2019) `IndexScalarQuantizer`⁷. During inference, the embeddings are de-quantized, and the search is performed on the restored FP32 vectors.

Saving Model Weights as FP16 Half precision can be used to size down the model weights of originally FP32 tensors with almost no drop in accuracy. In PyTorch, this can be done by calling `.half()` on each FP32 tensor in the model checkpoint.

In TensorFlow, model graphs saved as the data type of FP16 may result in unacceptably slow inference according to the used hardware. We have found out that keeping the tensor types of the graph as FP32 but making the actual assigned values as

⁷<https://github.com/facebookresearch/faiss/blob/v1.5.2/IndexScalarQuantizer.cpp>

FP16 enables a higher compression ratio when the model weights are compressed as described below.

Resource Compression Data compressors with a high compression ratio are effective at reducing the initial system footprint. Our observation is that bzip2 is better for binary files such as model weights or index of embedding vectors, whereas lzma is better for human-readable text files. System resources can also be compressed if necessary. We use `-9` option for both compressors.

Utilizing Token IDs as the Corpus A corpus file must be included in the system to get the actual text of the item retrieved by search (an embedding vector in our case). We have found out that using the file of the encoded token ids of the tokenized texts as the corpus, instead of the raw texts, is beneficial not only because it reduces the inference latency by preprocessing the texts, but also the compressed output size is often slightly smaller.

3 Experiments

Experimental Setup We apply our storage reduction methods to a recent extractive retrieve & read system, DPR (Karpukhin et al., 2020), which consists of three different BERT-base encoders: question encoder of the retriever, passage encoder of the retriever, and encoder of the reader with a ranker. All experiments are done on Naver Smart Machine Learning (NSML) Platform (Sung et al., 2017; Kim et al., 2018). The training and evaluation details are in Appendix A.1, A.2, and A.3.

Experimental Results Figure 1 shows how each of the discussed strategies changes DPR’s system size and Exact Match (EM) score on the EfficientQA dev set (see Table 3 and Table 4 in Appendix for details). Our starting point is a standalone open-domain QA system with DPR whose estimated size is 77.5 GB: 1.4 (system) + 0.8 (retriever) + 0.4 (reader) + 61 (index) + 13 (text) GB. The red plot shows from left to right one path to successively apply each strategy to reduce the system footprint to 484.69MB, which is 160 times smaller. Although the methods are described as sequential for easier presentation, the methods with filled markers and dotted lines are orthogonal to each other and thus can be applied in any other order. The methods with unfilled markers and solid lines are built on top of the previous method for each.

Sizing down the corpus from 21,015,325 to 1,224,000 (5.8%) passages (§2.1) decreases the sys-

tem footprint by a large margin of about 70.5GB with only 2.72% of drop in EM. Using a smaller passage embedding dimension of 128D (§2.2.1), changing the encoder to MobileBERT (§2.2.1), and sharing the encoders of the retriever (§2.2.2) save further 4.1GB of storage with little drop in accuracy of 1.28%. The process of unifying the retriever and reader into a single model (§2.2.3) drops EM by 1.11, but the accuracy increases by 2.77% (to 34.44%) with iterative finetuning (§2.2.4). In ablation studies on the three-step training procedure, omitting the knowledge distillation step drops EM by 1.5%, and omitting $\mathcal{L}_{\text{recon}}$ drops EM by 0.38%.

Applying post-training compression techniques further reduces the system footprint by a large margin while sacrificing little accuracy. EM changes to 34.39% with INT8 quantization, and the rest of the tricks do not affect the accuracy. Converting the PyTorch checkpoint to a binary for TensorFlow Serving to reduce system library dependency and applying bzip2 compression on some of the system resources creates the final system of 484.69MB with an accuracy of 34.33%. Figure 1 shows that this accuracy is higher than the performance of the parametric T5 (Roberts et al., 2020) baseline with a comparable docker-level system footprint.⁸

In Table 1, we show the test set accuracy of our final system and other baselines. In summary, the performance of our system is higher than all of the parametric baselines, and the accuracy drop from DPR is only 2.45% on the EfficientQA dev set and about 4% on the test set while reducing the system footprint to about 0.6% of the original size.

Our final system achieves the first place in the human (manual) evaluation and the second place in the automatic evaluation on “Systems Under 500MB Track” of the EfficientQA competition. While the accuracy of our system is 32.06% on the EfficientQA test set in the automatic evaluation, which is 1.38% behind the top-performing system (Lewis et al., 2021), its accuracy is 42.23% in the human evaluation which is 2.83% higher than the other system. Interestingly, when possibly correct answers are also counted as correct, the accuracy rises to 54.95% (7.58% higher than the other system). Please refer to Table 2 of Min et al. (2021) for more details.

In addition to the EfficientQA dataset, we also

⁸The accuracy of the T5 baselines are calculated using the SSM models finetuned on Natural Questions: https://github.com/google-research/google-research/tree/master/t5_closed_book_qa#released-model-checkpoints.

Table 1: System size and Exact Match (EM) score of several standalone open-domain QA systems, reported on the test set of EfficientQA. The value of the baseline systems except for DPR are reported in the EfficientQA leaderboard⁹. The system size of DPR is estimated as described in Section 3.

Model	EM	System Size	Mechanism
T5-1.1-small+SSM	18	486.61 MB	parametric
T5-1.1-XL+SSM	28	5.65 GB	parametric
REALM	35	27.19 GB	retrieve & read
DPR	36	77.5 GB	retrieve & read
+ Our Methods	32	484.69 MB	retrieve & read

perform experiments on open-domain Natural Questions (NQ) (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017) to test the generalizability of the proposed methods. The results and detailed analysis are presented in Appendix A.5.

4 Related Works

There has recently been a line of work that targets to create storage-efficient open-domain QA systems, especially following the EfficientQA competition. Here, we introduce several approaches concurrent to ours that interested readers may refer to. Izacard et al. (2020) and Fajcik et al. (2021) explore the trade-off between storage budget and accuracy, and their retrieve & read systems take up only about 6GB with state-of-the-art performance. Lewis et al. (2021) propose a QA-pair retrieval system for open-domain QA, which enjoys the benefits of high flexibility and low latency. Their retriever answers 1100 questions per second with 41.2% accuracy on NQ, which rises to 47.7% when equipped with a reranker. The variants optimized for small system footprint are the winning systems of two storage-constrained tracks at EfficientQA. Min et al. (2021) review the EfficientQA competition with detailed analysis and summarize all of the top-performing systems.

5 Conclusion

We discuss several orthogonal approaches to reduce the system footprint of a retrieve-and-read-based open-domain QA system. The methods together reduce the size of a reference system (DPR) by 160 times with an accuracy drop of 2.45% and 4% on EfficientQA dev and test, respectively. We hope that the presented strategies and results can be helpful for designing future retrieve-and-read systems under a storage-constrained serving environment.

⁹<https://ai.google.com/research/NaturalQuestions/efficientqa>

Acknowledgements

The authors would like to thank the members of NAVER Clova for proofreading this paper. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

References

- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *NeurIPS*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Martin Fajcik, Martin Docekal, Karel Ondrej, and Pavel Smrz. 2021. Pruning the index contents for memory efficient open-domain qa. *arXiv preprint arXiv:2102.10697*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2019. Realm: Retrieval-augmented language model pre-training. In *ICML*.
- Song Han, Huizi Mao, and William J Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2014. Distilling the knowledge in a neural network. In *NIPS Deep Learning Workshop*.
- Gautier Izacard and Edouard Grave. 2021. Leveraging passage retrieval with generative models for open domain question answering. *EACL*.
- Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola De Cao, Sebastian Riedel, and Edouard Grave. 2020. A memory efficient baseline for open domain question answering. *arXiv preprint arXiv:2012.15156*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. In *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 *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. In *EMNLP*.
- Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, Nako Sung, and Jung-Woo Ha. 2018. Nsm1: Meet the mlaas platform with a real-world case study. *arXiv preprint arXiv:1712.05902*.
- 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. In *TACL*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. In *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. In *EMNLP*.
- 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. *arXiv preprint arXiv:2102.07033*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Sewon Min, Jordan Boyd-Graber, Chris Alberti, Danqi Chen, Eunsol Choi, Michael Collins, Kelvin Guu, Hannaneh Hajishirzi, Kenton Lee, Jennimaria Palomaki, Colin Raffel, Adam Roberts, Tom Kwiatkowski, Patrick Lewis, Yuxiang Wu, Heinrich Küttler, Linqing Liu, Pasquale Minervini, Pontus Stenetorp, Sebastian Riedel, Sohee Yang, Minjoon Seo, Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola De Cao, Edouard Grave, Ikuya Yamada, Sonse Shimaoka, Masatoshi Suzuki, Shumpei Miyawaki, Shun Sato, Ryo Takahashi, Jun Suzuki, Martin Fajcik, Martin Docekal, Karel Ondrej, Pavel Smrz, Hao Cheng, Yelong Shen, Xiaodong Liu, Pengcheng He, Weizhu Chen, Jianfeng Gao, Barlas Oğuz, Xilun Chen, Vladimir Karpukhin, Stan Peshterliev, Dmytro Okhonko, Michael Schlichtkrull, Sonal Gupta, Yashar Mehdad, and Wen tau Yih. 2021. Neurips 2020 efficientqa competition: Systems, analyses and lessons learned. *arXiv preprint arXiv:2101.00133*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *EMNLP*.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert:

a compact task-agnostic bert for resource-limited devices. In *ACL*.

Nako Sung, Minkyu Kim, Hyunwoo Jo, Youngil Yang, Jingwoong Kim, Leonard Lausen, Youngkwan Kim, Gayoung Lee, Donghyun Kwak, Jung-Woo Ha, and Sunghun Kim. 2017. Nsml: A machine learning platform that enables you to focus on your models. In *NeurIPS Systems for ML Workshop*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv preprint arXiv:1910.03771*.

Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *ICLR*.

Sohee Yang and Minjoon Seo. 2020. Is retriever merely an approximator of reader? *arXiv preprint arXiv:2010.10999*.

Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8bert: Quantized 8bit bert. In *NeurIPS EMC2 Workshop*.

A Appendix

A.1 Training Details of the Passage Filter

J_{total} is the set of all 21M passages that serve as the knowledge source of DPR. J_{train}^+ consists of the top-200 passages retrieved for each of the questions in Natural Questions (Kwiatkowski et al., 2019) train, dev, test set, and EfficientQA dev set. To retrieve the passages, we use the retriever of Yang and Seo (2020) trained on Natural Questions. Smoothed frequency $\lfloor \log_{10}(\text{freq}_{p_j}) + 1 \rfloor$ is considered to create a candidate pool of the positive passages, and oversampling from the pool is done to make $J_{\text{train}}^{+'}$, whereas $J_{\text{train}}^{-'}$ is randomly and uniformly sampled from J_{train}^- . The objective function is defined as follows:

$$\frac{\sum_{j \in J_{\text{train}}^{+'}} \log \sigma(s_j) + \sum_{j \in J_{\text{train}}^{-'}} \log [1 - \sigma(s_j)]}{|J_{\text{train}}^{+'}| + |J_{\text{train}}^{-'}|}.$$

We finetune a RoBERTa-base (Liu et al., 2019) classifier with a batch size of 18 ($|J_{\text{train}}^{+'}| = |J_{\text{train}}^{-'}| = 9$), learning rate of 1e-5, dropout rate of 0.1, max norm gradient clipping of 2.0, and warmup steps of 1000, using one V100 GPU. We use the code of HuggingFace Transformers (Wolf et al., 2019), and no additional preprocessing is used on the data other than the tokenization for RoBERTa-base. We train the model for one epoch until all the positive passages oversampled according to the smoothed frequency are seen by the model at least once.

We evaluate the model on the validation set after every 2000 steps of gradient update. We compose J_{val}^+ as a set of indices of the passages that hold the answer for one of the questions in the EfficientQA dev set and are retrieved by an existing retriever as the most relevant passage to the question. We use $|J_{\text{val}}^+| = 893$ positive passages and $|J_{\text{val}}^-| = 10000 - 893 = 9107$ randomly selected negative passages. $n = 1,224,000$ passages are selected for use in our minimal retrieve & read system to fit in the storage budget of 500MB.

A.2 Training Details of the Retriever-Reader with Single Encoder

We have not searched for hyperparameters in almost all experiments on parameter sharing and mainly followed the training setup of DPR.¹⁰

For the experiments on EfficientQA, the training set of Natural Questions (Kwiatkowski et al., 2019) is used to train the models. The checkpoints

¹⁰<https://github.com/facebookresearch/DPR>

Table 2: Statistics of the number of questions in each dataset. The values in parenthesis denote the number of questions filtered by DPR preprocessing and used for actual training. There is no training data for EfficientQA, so we use the training set of Natural Questions to train the model for EfficientQA.

Dataset	Train	Dev	Test
EfficientQA	-	1,800	1,800
Natural Questions	79,168 (58,880)	8,757	3,610
TriviaQA	78,785 (60,413)	8,837	11,313

that report the best result on the EfficientQA dev set¹¹ are selected. Our code is built on top of the official implementation of DPR, so the datasets are preprocessed as done in the code of DPR. Table 2 shows the statistics of the datasets including Natural Questions and TriviaQA used for the experiments in Appendix A.5. We train the models using four to eight P40 or V100 GPUs.

Retriever Encoder Sharing We use MobileBERT (Sun et al., 2020) as the pretrained encoder. The encoder output vector is what corresponds to the [CLS] token in the input, and only the first 128 out of 512 dimensions are utilized to calculate the retrieval score. Following the setup of Karpukhin et al. (2020), we use a learning rate of $2e-5$, max norm gradient clipping of 2.0, warmup of 1237, sequence length of 256, in-batch negative training of 1 positive and 127 negatives, and training epochs of 40 to 70, applying early stopping if there is a lack of resource. The models are evaluated on the dev set after every epoch.

Unified Retriever-Reader Through Knowledge Distillation $k' = 200$ and $m = 24$ is used to create the training dataset. To train the unified retriever-reader, we use a learning rate of $1e-5$, max norm gradient clipping of 2.0, no warmup steps, sequence length of 350, batch size of 16, knowledge distillation temperature τ of 3, and training epochs of 16 to 30, applying early stopping when the score seems to be converged.

Since the reader teacher (DPR reader) has a ranker, $\mathcal{L}_{\text{read}}$ is defined as the sum of the KL divergence between the span scores and the KL divergence between the ranking scores of the teacher and the student. \mathcal{L}_{ret} also takes the passage ranking score from the ranker as the distillation target.

The model is evaluated at every 2000 steps, and we select the checkpoint with the highest aver-

¹¹https://github.com/google-research-datasets/natural-questions/blob/master/nq_open/NQ-open.efficientqa.dev.1.1.jsonl

age EM on m' retrieved passages, where $m' \in \{1, 10, 20, 30, 40, 50\}$, along with an acceptable reranking accuracy (how many times the positive passage is ranked at the top among 50 candidates).

Iterative Finetuning of Unified Retriever-Reader We finetune the model for at most six epochs. The rest of the hyperparameters are the same as described in the previous paragraphs.

A.3 Evaluation Details

The reported EM is the highest EM on m' retrieved passages where $m' \in \{1, 10, 20, \dots, 100\}$. The original code of DPR searches the answer only in the passage scored the highest by the passage ranker, and thus the answer span with the highest span score in the single passage is selected as the final answer. All of the EM scores presented in this work are also calculated this way.

On the other hand, we have found out that the end-to-end QA accuracy can be slightly increased by using the weighted sum of the passage ranking score $\mathcal{P}_{\text{rank}}$ and the answer span scores, $\mathcal{P}_{\text{start}}$ and \mathcal{P}_{end} for the start and end positions, respectively, to compare the answer candidates at inference time. Therefore, we have used this scoring method for the model submitted to the EfficientQA leaderboard. Specifically, we use $(1 - \lambda)(\log \mathcal{P}_{\text{start}} + \log \mathcal{P}_{\text{end}}) + 2\lambda \log \mathcal{P}_{\text{rank}}$ as the score. The answer spans with the top five weighted sum scores in each retrieved passage are selected as the candidate answers, and the one with the highest score is chosen as the final answer. We select $\lambda = 0.8$ based on the performance on the dev set. This method increases the dev set accuracy after the iterative finetuning stage (§2.2.3) from 34.44 to 34.61.

Due to the discrepancy between the validation accuracy during and after training (described in detail in Appendix A.5), we select up to five checkpoints based on the dev set accuracy observed during training and evaluate them to obtain the one with the actual highest dev set accuracy after the iterative finetuning is done.

A.4 System Footprint

System footprint is measured by the command `du -h /` inside the standalone docker container right after its launching as stated in the EfficientQA competition guideline. The system footprint at runtime may be larger when the resources are initially compressed at the time of launching the container.

Table 3 shows from the top to bottom the detailed

Table 3: Detailed ablations on how the system size changes from 77.5 GB to 484.69 MB by applying each of the discussed methods from the top to bottom. The values in the table use MB as the unit. Decreased values are marked in red and increased values are marked in blue.

	Docker	Retriever	Reader	Index	Text File	Total
DPR	1,270	836	418	61,919	13,065	77,508
§2.1 Passage Filtering (21,015,325 → 1,224,000 passages)	1,270	836	418	3,681	756	6,961
§2.2.1 Embedding Dimension Reduction (768D → 128D)	1,270	836	418	614	756	3,894
§2.2.1 Lightweight Encoder (BERT → MobileBERT)	1,270	188	94	614	756	2,922
§2.2.2 Retriever Encoder Sharing	1,270	94	94	614	756	2,828
§2.2.3 Unified Retriever-Reader Through Knowledge Distillation	1,270	94	0	614	756	2,734
§2.2.4 Iterative Finetuning of Unified Retriever-Reader	1,270	94	0	614	756	2,734
§2.3 INT8 Quantization of Index Items	1,270	94	0	170	756	2,290
§2.3 Saving Model Weights as FP16	1,270	47	0	170	756	2,243
§2.3 Resource Compression	1,270	42	0	145	187	1,644
§2.3 Utilizing Token IDs as the Corpus	1,270	42	0	145	177	1,634
§3 TF Serving, Minimizing Library Dependencies, Fusing Index into Model Graph	312	177	0	0	177	666
§2.3 System Resource Compression	130	177	0	0	177	484

ablations on how the system size changes from 77.5 GB to 484.69 MB by applying each of the methods discussed in Section 2. The values in the table use MB as the unit. Decreased values are marked in red and increased values are marked in blue.

The docker image is initially assumed to be bitnami/pytorch:1.4.0¹², and it changes to python:3.6.11-slim-buster¹³ after adopting TensorFlow (TF) Serving that does not require heavy system libraries as PyTorch does. The most lightweight docker image with python uses Alpine, but TF Serving does not run on an Alpine docker container due to the lack of support of system library requirements.

A.5 Experiments: NQ and TriviaQA

Experimental Setup Most of the details to train the models on Natural Questions (NQ) and TriviaQA (Trivia) follow what is written in Appendix A.1 and Appendix A.2, and here we describe only the differences. To train the passage filter, we use log with base 2 instead of 10 for Trivia due to its higher validation set accuracy. The questions used to create the training data are from the train and dev set of the datasets which correspond to the targets of the filter models. To train the unified retriever-reader through knowledge distillation, a batch size of 8 with gradient accumulation steps of 2 is used to train the models using only four V100 GPUs. The maximum number of training epochs is set to 30, but training is stopped around the 16th epoch to shorten the training time even when the scores do not seem to be fully converged.

¹²<https://hub.docker.com/r/bitnami/pytorch>

¹³https://hub.docker.com/_/python

For iterative finetuning, a batch size of 8 with gradient accumulation steps of 2 is again used with four V100 GPUs. The maximum number of training epochs is also set to 30, but the training is stopped before the 10th epoch.

Experimental Results Figure 2 shows the EM and docker-level system footprint when each of the discussed strategies is applied to DPR. In the case of the EfficientQA dataset, the step-wise evaluation result on the test set cannot be reported because the answer set is not publicly available. On the other hand, for NQ and Trivia, we present the step-wise accuracy on the test set along with that on the dev set to show how the strategies affect the actual performance on the test set. The evaluation results of all the cases are presented in Table 4.

Let us define the relative performance drop at step t as the percentage of $\frac{EM_{t-1} - EM_t}{EM_{t-1}}$ where EM_t is the EM score at the t -th phase. As shown in the figures and the table, applying the methods to different datasets does not show consistent trends. Because the EfficientQA dataset is constructed in the same way as NQ (Min et al., 2021), the trends on these two datasets are similar except that changing the backbone from BERT to MobileBERT (§2.2.1) results in a significant relative performance drop of 8.21% on the dev set of NQ while the value is only 0.18% on EfficientQA. On the other hand, the same change results in about 4% of relative performance *gain* on Trivia. A different phenomenon also appears when the retriever encoders are shared (§2.2.2) that the accuracy rises on EfficientQA and NQ while it drops on Trivia.

The percentage of the final accuracy to the ac-

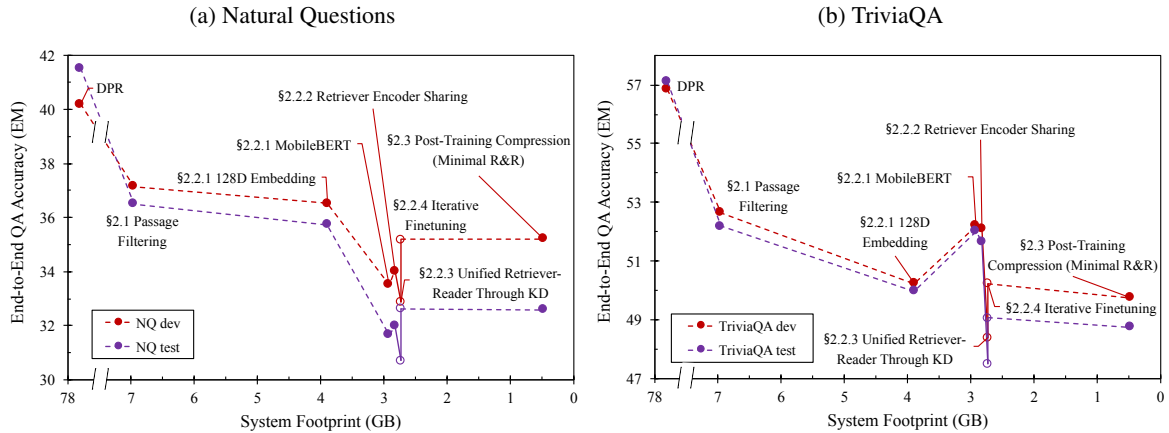


Figure 2: System footprint vs. Exact Match (EM) accuracy on Natural Questions and TriviaQA.

Table 4: Detailed ablations on how the Exact Match (EM) score on each dataset changes by applying each of the discussed methods from the top to bottom.

	EfficientQA Dev	NQ Dev	NQ Test	TriviaQA Dev	TriviaQA Test
DPR	36.78	40.20	41.52	56.84	57.10
§2.1 Passage Filtering	34.06	37.14	36.51	52.65	52.19
§2.2.1 128D Embedding	32.67	36.52	35.73	50.23	49.97
§2.2.1 MobileBERT	32.61	33.52	31.66	52.20	52.02
§2.2.2 Retriever Encoder Sharing	32.78	34.03	31.99	52.11	51.66
§2.2.3 Unified Retriever-Reader Through KD	31.67	32.88	30.72	48.40	47.51
§2.2.4 Iterative Finetuning	34.44	35.19	32.63	50.23	49.06
§2.3 Post-Training Compression	34.33	35.22	32.60	49.76	48.75

accuracy at the start also differs among the datasets: 93.3% and 89.0%¹⁴ on the EfficientQA dev and test set, 87.6%, 87.5%, and 85.4% on the NQ dev set, Trivia dev set, and Trivia test set, respectively, but 78.5% on the NQ test set. While the gap between the percentages on the dev and test set is small on Trivia, the value is considerably large on NQ. Also, the gap between the dev and test set accuracy divided by the latter is about 7% on EfficientQA and NQ, while it is only 2% on Trivia.

Meanwhile, a common observation is that passage filtering (§2.1), embedding dimension reduction (§2.2.1), and unifying the retriever and the reader through knowledge distillation (§2.2.3) consistently result in the drop of accuracy. The relative performance drop of each of the methods is 7.40%, 4.08%, and 3.39% on the EfficientQA dev set, 7.61%, 1.67%, and 3.38% on the NQ dev set, 12.07%, 2.14%, and 3.97% on the NQ test set, 7.37%, 4.60%, 7.12% on the Trivia dev set, and 8.60%, 4.25%, 8.03% on the Trivia test set.¹⁵

¹⁴36.0 is used as an approximation for the accuracy of DPR on the EfficientQA test set, which is reported as 36 in https://github.com/google-research-datasets/natural-questions/tree/master/nq_open.

¹⁵Figure 1 of Izacard et al. (2020) also shows the trade-off between the index size and system accuracy. Note that the im-

plementation details of their passage filtering and embedding dimension reduction are different from ours.

In the case of unifying the retriever and the reader into one model, one possible cause of the accuracy drop might have come from its currently suboptimal checkpoint selection method. From the moment the retriever and reader are unified into one model and jointly trained, the validation accuracy reported during training uses the outputs of the *initial* retriever parameters while the actual evaluation must use outputs of the *updated* retriever parameters at the time of validation. Due to this discrepancy, checkpoint selection based on the validation accuracy at training does not lead to the model with the actual highest dev set accuracy. The discrepancy may further necessitate measuring the true dev set accuracy at several different checkpoints (possibly with high validation accuracy during training) to choose the final model after iterative finetuning. To deal with this issue and fairly compare the best checkpoints, future research may be conducted to refresh the retrieval index during training as in the work of Guu et al. (2019); Xiong et al. (2021), so that the evaluation (and training) may not be done on the stale retrieval outputs.