

# R2-D2: A Modular Baseline for Open-Domain Question Answering

Martin Fajcik, Martin Docekal, Karel Ondrej, Pavel Smrz

Brno University of Technology

{ifajcik, idocekal, ondrej, smrz}@fit.vutbr.cz

## Abstract

This work presents a novel four-stage open-domain QA pipeline R2-D2 (RANK TWICE, READ TWICE). The pipeline is composed of a retriever, passage reranker, extractive reader, generative reader and a mechanism that aggregates the final prediction from all system's components. We demonstrate its strength across three open-domain QA datasets: NaturalQuestions, TriviaQA and EfficientQA, surpassing state-of-the-art on the first two. Our analysis demonstrates that: (i) combining extractive and generative reader yields absolute improvements up to 5 exact match and it is at least twice as effective as the posterior averaging ensemble of the same models with different parameters, (ii) the extractive reader with fewer parameters can match the performance of the generative reader on extractive QA datasets<sup>1</sup>.

## 1 Introduction

Last year showed rapid progress in neural factoid open-domain question answering based on *retriever-reader* architecture (Open-QA). Such Open-QA systems (Chen et al., 2017) seek evidence for answering the questions inside the knowledge source using the *retriever* and then extract the answer from the retrieved knowledge using the *reader*. The knowledge source is often a large corpus of short snippets of natural language, so-called passages (e.g., taken from an encyclopedia).

The progress can be attributed to advances in neural retrieval methods (Karpukhin et al., 2020; Izacard and Grave, 2020; Khattab et al., 2020; Luan et al., 2021; Xiong et al., 2020, *inter alia*) that benefit from smarter negative sampling strategies or a better trade-off between complex question-passage interaction and its efficiency. It also can be attributed to reading methods that enable process-

ing large quantities of retrieved passages Izacard and Grave (2021). They compensate for a certain amount of the retrieval error and enable early aggregation of answer's evidence between passages.

This work demonstrates the relative improvement of 23-32% compared to last year's state-of-the-art DPR system (Karpukhin et al., 2020), while using the same knowledge source and the retriever. We propose a state-of-the-art Open-QA baseline composed of retriever, passage reranker, extractive reader, generative reader, and a novel component fusion approach. We follow the practice from information retrieval and show that our moderately sized reranker allows to reduce the passage count needed at the input of large reader models about four times. Our readers then take the best from both worlds. The extractive reader proposes a list of salient answer spans. The generative reader reranks these spans, seeing all the passages at once, or generates its own answer. The proposed pipeline is heterogeneous and modular, making it an ideal benchmark.

To sum up, our contributions are three-fold:

1. We present a simple novel approach to aggregate scores from all system components and show that combining extractive and generative approaches is superior to a posterior averaging ensemble of homogeneous models.
2. We show that the extractive reader can sometimes match the performance of the generative approaches without taking the advantage of the fusion between retrieved passages. This indicates that the evidence aggregation from multiple passages in the generative approaches is either not learned or not necessary to perform well on these datasets.
3. We push the state-of-the-art for two large and popular datasets, demonstrating what is achievable with the proposed approach, having the same knowledge source and the re-

<sup>1</sup>Our demo is available at <http://r2d2.fit.vutbr.cz/>. Code and preprocessed data are available at <https://github.com/KNOT-FIT-BUT/R2-D2>.

triever as in the previous works (Karpukhin et al., 2020; Izacard and Grave, 2021).

## 2 Open-QA Pipeline

We propose the R2-D2 (RANK TWICE, READ TWICE), 4-stage pipelined system that can choose whether to generate or to extract an answer. The parameters of each component in pipeline are estimated separately. It is composed of DPR passage retriever (Karpukhin et al., 2020), passage reranker (see subsection 2.1), and two readers. Figure 1 shows the diagram of our system. The first reader performs an extractive span-selection similar to Fajcik et al. (2020). The second reader is based on Fusion-In-Decoder (FiD) (Izacard and Grave, 2021).

Formally, given a question  $q \in \mathcal{Q}$  from the set of all possible questions  $\mathcal{Q}$  and the corpus  $\mathcal{C} = \{p_1, p_2, \dots, p_n\}$  composed of passages  $p_i$ , the retriever learns a ranking function  $\text{rank} : \mathcal{Q} \times \mathcal{C} \rightarrow \mathbb{R}$  that assigns a score to each passage. We assume each passage contains its passage title (e.g., title from the Wikipedia article).

Taking a top- $K$  scoring passages  $\mathcal{C}_r \subset \mathcal{C}$ , reranker again rescores  $\mathcal{C}_r$  scoring passages by learning a reranking function  $\text{rerank} : \mathcal{Q} \times \mathcal{C}_r \rightarrow \mathbb{R}$ . Note that while rank and rerank have similar signatures, the computational cost of rerank over the same amount of passages is drastically higher, as it computes fine-grained interaction between tokens of question and passage.

Next, the rescored passages are passed to two readers: the extractive reader reads top- $V$  passages  $\mathcal{C}_{rr} \subset \mathcal{C}_r$  independently of each other and assigns the probability  $P_e(a_e|q, \mathcal{C}_{rr})$  to each span  $a_e$  in the passages (see subsection 2.2). The FiD generative reader reads top- $V_2$  passages  $\mathcal{C}'_{rr} \subset \mathcal{C}_r$  jointly and generates an answer from probability space  $P_g(a_g|q, \mathcal{C}'_{rr})$  via greedy search.

Finally, R2-D2 aggregates the outputs from all components using two fusions (described in subsection 2.3).

### 2.1 Passage Reranker

The proposed passage reranker is based on transformer cross-encoder similar to Nogueira and Cho (2019); Luan et al. (2021). The input is the concatenation of question  $q \in \mathcal{Q}$  and passage  $p \in \mathcal{C}_r$  with a special SEP token between them. The passage consists of a title and context that are prepended with special start tokens and concate-

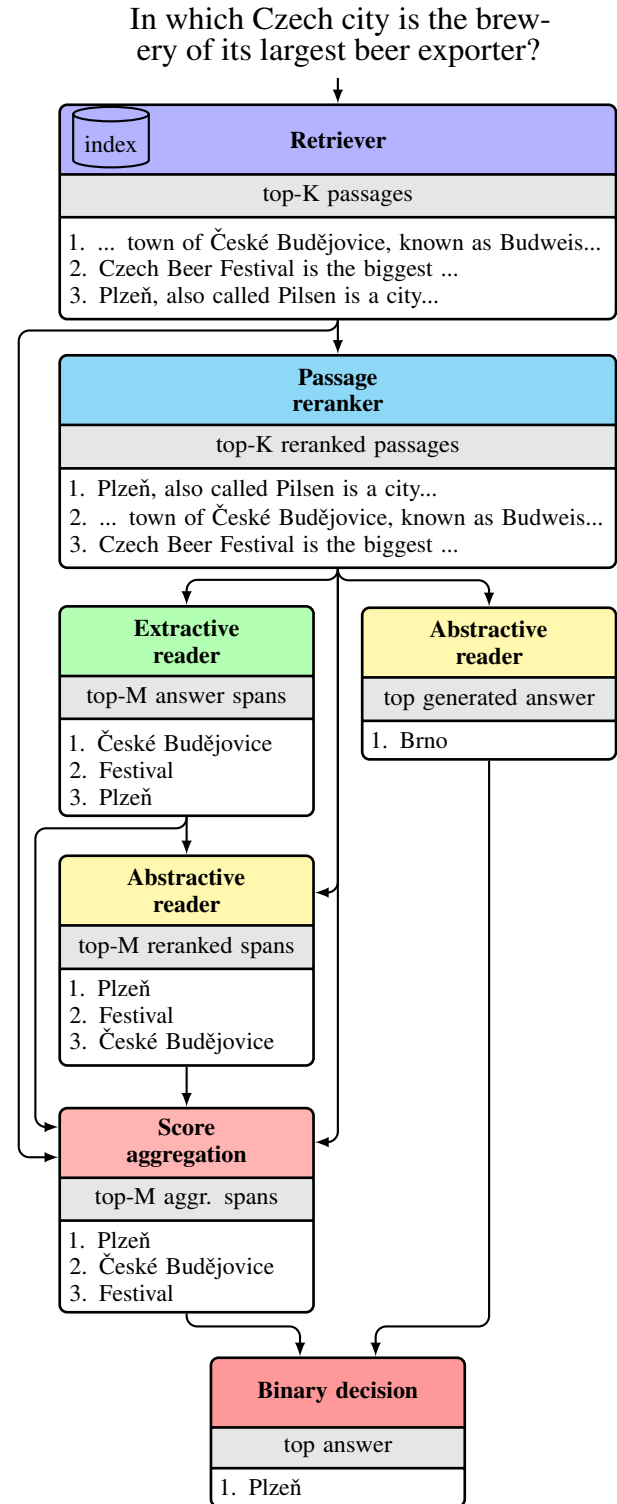


Figure 1: R2-D2 pipeline.

nated together. We denote the contextual representation of input token  $w$  obtained by the cross-encoder as  $\text{En}(p, q)[w] \in \mathbb{R}^d$ .

Now we can define the reranking function for

passage rescoring as

$$\text{rerank}(q, p) = \text{En}(p, q)[\text{CLS}]^\top w \quad (1)$$

where  $w \in \mathbb{R}^d$  is a trainable vector and CLS is the special token added at the start of an input sequence. Finally, we define the following formula<sup>2</sup>

$$\mathbf{P}_{rr}(p|q, \mathcal{C}_r) = \text{softmax}_{p \in \mathcal{C}_r}(\text{rerank}(q, p))_p \quad (2)$$

to assign a probability to the case that passage  $p$  contains answer to the question  $q$ .

**Training.** The model input for each question is exactly one positive sample supplemented with hard negatives from the retriever. The ground truth passage, annotated the same way as in Karpukhin et al. (2020), is primarily used as a positive sample. If the ground truth is unknown, the positive sample is the best retriever passage containing the answer. The hard negatives are uniformly sampled from retriever’s top- $K$  results that do not contain the answer. The used loss function is the cross-entropy.

## 2.2 Extractive Reader

Extractive reader estimates the probability  $\mathbf{P}_e(a_e|q, \mathcal{C}_{rr})$ . It is the probability of a span  $a_e$  from top- $V$  passage  $p \in \mathcal{C}_{rr}$  being an answer to a question  $q$ . We decompose the  $\mathbf{P}_e(a_e|q, \mathcal{C}_{rr})$  into four probabilities of:

- token  $s$  being starting token of an answer span,
- token  $e$  being ending token of an answer span,
- tokens  $s$  and  $e$  being boundary tokens of an answer span (Fajcik et al., 2020),
- passage  $p$  containing an answer for the question  $q$  (inner reranker) as in Karpukhin et al. (2020).

To obtain the final probability used in test-time, we compute their product<sup>3</sup>. These probabilities are defined as:

$$\mathbf{P}_*(\cdot|q, \mathcal{C}_{rr}) = \text{softmax}(s_*)_i, \quad (3)$$

where  $*$  may stand for a *start*, *end*, *joint*, and a *passage*. The  $i$  is an index of a given element, and the  $s_*$  is a vector of scores for each element among all passages in  $\mathcal{C}_{rr}$ . So the *softmax* normalization

<sup>2</sup>Formal definition of softmax over a set is described in the Appendix D.

<sup>3</sup>We tried decoding from the subsets of these probabilities in Appendix E not observing significant difference.

sum goes through all the passages. On the other hand, the  $s_*$  scores are estimated by the model with just a single passage at its input (Clark and Gardner, 2018). The scores are as follows:

$$s_{start}^i = \text{En}(p, q)[s]^\top w_{start} \quad (4)$$

$$s_{end}^i = \text{En}(p, q)[e]^\top w_{end} \quad (5)$$

$$s_{joint}^i = (W_j \text{En}(p, q)[s] + b_j)^\top \text{En}(p, q)[e] \quad (6)$$

$$s_{passage}^i = \text{En}(p, q)[\text{CLS}]^\top w_p. \quad (7)$$

Where  $w_*, b_j \in \mathbb{R}^h$ ,  $\text{En}(p, q)[\cdot] \in \mathbb{R}^h$ , and  $W_j \in \mathbb{R}^{h \times h}$  are all trainable parameters.

We omit the spans of a title and question for answer span selection. Therefore the final answer can be selected only from the context.

The following training objective with independently marginalized components is used:

$$\begin{aligned} & -\log \sum_{s \in \text{starts}(\mathcal{C}_{rr})} \mathbf{P}_{start}(s|q, \mathcal{C}_{rr}) \\ & -\log \sum_{e \in \text{ends}(\mathcal{C}_{rr})} \mathbf{P}_{end}(e|q, \mathcal{C}_{rr}) \\ & -\log \sum_{j \in \text{boundaries}(\mathcal{C}_{rr})} \mathbf{P}_{joint}(j|q, \mathcal{C}_{rr}) \\ & -\log \sum_{p \in \mathcal{C}_{rr}} \mathbf{P}_{passage}(p|q, \mathcal{C}_{rr}). \end{aligned} \quad (8)$$

The sums are going through target annotations (starts, ends, etc.) obtained by the distant supervision approach.

## 2.3 Component Fusion

To produce the final answer, R2-D2 aggregates the log-probabilities of all system components via linear combinations tuned on validation data.

Firstly, the log-probabilities of all system components for top- $M$  answer spans proposed by the extractive reader are aggregated. Formally, assume the  $\mathcal{A}_q$  is the set of top- $M$  answer spans from  $\mathbf{P}_e(a|q, \mathcal{C}_{rr})$  for question  $q$ . The generative model performs the **answer reranking** evaluating the log-probability of the answer spans

$$\{\log \mathbf{P}_g(a|q, \mathcal{C}'_{rr}) : a \in \mathcal{A}_q\}. \quad (9)$$

Next a logistic regression loss (11) is minimized to perform **score aggregation**. It combines the scores across the R2-D2 components to maximize the correct answer span probability over dataset  $\mathcal{D}$ .

This dataset is composed of the top- $M$  outputs of the extractive reader with the correct answer.

$$x(a) = [\mathbf{P}_e(a) \mathbf{P}_g(a) \mathbf{P}_r(p_a) \mathbf{P}_{rr}(p_a)] \quad (10)$$

$$- \sum_{(\mathcal{A}_q, gt) \in \mathcal{D}} \log \operatorname{softmax}_{a \in \mathcal{A}_q} (w^\top \log x(a) + b)_{gt} \quad (11)$$

Here  $p_a$  denotes the passage containing the answer span  $a$ ,  $\mathcal{A}_q$  is a set of proposed answer spans,  $gt$  is the correct answer span, distribution dependencies are dropped for clarity and only the logistic regression parameters  $w, b$  are tuned in this step.

Finally, we theorized the correct answer span might not always be available in the passage set  $\mathcal{C}_{rr}$ , but the generative reader might be able to generate the answer from its parameters and the evidence given in passages. We introduce the binary classifier, which decides whether to select the best span answer from answer aggregation step or a free-form answer generated via FiD. Given that  $s_{agg}(q) = \max_{a \in \mathcal{A}_q} w^\top x(a) + b$  is the best span score and  $s_g^*(q) = \log \mathbf{P}_g(a_q^* | q, \mathcal{C}'_{rr})$  is the log-probability of the answer  $a_q^*$  obtained via greedy decoding for question  $q$ , a classifier is trained via binary cross-entropy  $BCE(l, t)$  with log-odds ratio  $l$  and target  $t$  to do the **binary decision**

$$\sum_{(e, t) \in \mathcal{D}} BCE(w^\top [s_{agg}(e); s_g^*(e)] + b, t). \quad (12)$$

Here, the training dataset  $\mathcal{D}$  contains only cases where either the extractive or the abstractive prediction is correct (but not both).

### 3 Experimental Setup

Our models are implemented in PyTorch (Paszke et al., 2019) using Transformers (Wolf et al., 2020). We use 12GB GPU to train the passage reranker, 48GB GPU for the generative reader, and 16x 32GB GPUs to train the extractive reader with  $V = 128$  passages at its input. The inference runs on 12GB GPU. In all experiments, we used Adam optimizer with a decoupled weight decay (Loshchilov and Hutter, 2019). Our models are evaluated by two metrics:

**Exact match (EM)** measures the proportion of examples, for which the system prediction matched at least one annotated ground-truth answer. We use the script from Lee et al. (2019)<sup>4</sup>.

<sup>4</sup><https://cutt.ly/rkZNIer>

**Accuracy@K** measures the proportion of examples, for which the ground-truth answer string is present in top-K retrieved passages. We match the string exactly as Karpukhin et al. (2020)<sup>5</sup>.

#### 3.1 Datasets and Data Pre-processing

We evaluate our models on three datasets. Their statistics are available in Table 1. To train the reranker we filter out examples, which do not contain golden passage or exact match in top- $K$  retrieved passages. To train the extractive reader, only examples with exact match in a golden passage or top-1 retrieved passage are kept. Both filtering strategies are closely described in Appendix C.

**NQ-Open** (Kwiatkowski et al., 2019; Lee et al., 2019) or NaturalQuestions-Open consists of real user queries obtained from Google search engine. The maximum length of each answer is at most 5 tokens. Each training and development sample contains 1 annotated answer, while test data contain 5-way answer annotation.

**TQ-Open** (Joshi et al., 2017) or TriviaQA-Open consists of question-answer pairs from 14 different trivia quiz websites. Each question contains human annotated answer and a set of answer aliases gathered from Wikipedia. We use the unfiltered version.

**EfficientQA** (Min et al., 2021) is a dataset collected the same way as NQ-Open through 2019 and thus may contain more questions without evidence in our corpus than NQ-Open. We use the officially released dev set for testing<sup>6</sup> models trained on NQ-Open training data.

Additionally, we also report results according to train-test set overlaps discovered by Lewis et al. (2021) in Appendix H.

#### 3.2 Models and Pipeline

**Retriever.** We use BERT-based DPR from the official checkpoint<sup>7</sup>. Each passage is represented via 768-dimensional embedding. We use a multiset checkpoint for TQ-Open, as the checkpoint for TQ directly isn't officially released. We use the same knowledge corpus containing 21,015,320 passages based on 12-20-2018 Wikipedia snapshot as Karpukhin et al. (2020). In inference time, the retriever passes  $K = 200$  passages  $\mathcal{C}_r$  to reranker.

<sup>5</sup><https://cutt.ly/0luNhx4>

<sup>6</sup>The test set was not released during our experiments.

<sup>7</sup><https://github.com/facebookresearch/DPR>



Dataset	Train	Dev	Test
NQ-Open	79,168	8,757	3,610
- filt. reranker	71,238	-	-
- filt. ext. reader	61,755	-	-
- w/ golden passage	58,876	6,515	-
TQ-Open	78,785	8,837	11,313
- filt. reranker	69,346	-	-
- filt. ext. reader	62,332	-	-
- w/ golden passage	60,413	6,760	-
EfficientQA	-	-	1,800

Table 1: Dataset statistics. The filt. lines report how many examples are kept for training the reranker (filt. reranker) and extractive reader (filt. ext. reader). The lines w/ golden passage denote how many examples from the set contain golden passage annotation.

**Passage reranker.** We use the RoBERTa-base (Liu et al., 2019) and truncate the inputs to a maximum length of 256. The linear scheduler with 0.1 warmup proportion is used, the number of epochs is 5 and the model is validated every 40,000 optimization steps. We use learning rate  $1.6 \cdot 10^{-4}$  and batch size 8. In training, the model reranks 24 passages per question with negatives uniformly sampled from top-400 passages retrieved by DPR. During the inference, top- $K$  ( $K = 200$ ) retriever passages are rescored and passed to readers.

**Extractive reader.** The extractive reader encoder is based on pre-trained ELECTRA-large. Its inputs are truncated if they are longer than the allowed maximum size (512 tokens). During the training phase, all spans from all  $p \in \mathcal{C}_r$ <sup>8</sup> that match<sup>9</sup> with at least one of the known answers are selected as target annotations. Therefore the annotations might appear in the wrong context. The extractive reader reads the top- $V = 128$  passages during the training phase and when it is used without the reranker. To demonstrate the effect of reranker, the reader reads only the top- $V = 24$  passages if the reranker is used. We use a linear scheduler with a warmup for the first 20,000 steps for all models. The maximum number of training steps is 200,000. The model is validated every 20,000 steps, and the best checkpoint among validations is selected. The learning rate is  $2 \cdot 10^{-5}$  and the optimization step was done after each training example.

**Generative reader.** We utilize T5-large (Raffel et al., 2020) and use a concatenation of question,

passages and their respective titles at the Fusion-in-Decoder’s input the same way as Izacard and Grave (2020). We truncate each passage to the length of 250 tokens for NQ. For TQ, as questions are significantly longer, we truncate whole inputs to the same size. Following FiD for TQ, we use only human-generated answer. In training, the golden passage always comes first, if available, and we take the rest of passages as ranked by retriever up to  $V_2$  passages. Izacard and Grave (2021) trained FiD with  $V_2 = 100$  passages at its input. However, such approach requires tremendous amount of GPU memory, and thus requires employing speed-memory trade-offs such as gradient checkpointing (Chen et al., 2016). Unlike the original approach, we use only  $V_2 = 25$  passages in our FiD. We note that in practice combining reranker with shorter-context FiD yields results similar to original implementation with much lower memory consumption and better throughput in the R2-D2 setting<sup>10</sup>. We analyze the speed of our implementation in Appendix I. Other hyperparameters are similar to the original work—batch size 64, learning rate  $5 \cdot 10^{-5}$  but no learning rate schedule. In test time, we decode an answer via greedy decoding.

## 4 Results and Analysis

The effectiveness of our approach is compared with the state-of-the-art in Table 2. Our system, composed of just the retriever and FiD reader R1-D1 (Generative), shows inferior performance compared to FiD-large. This is most likely caused by 4 times fewer passages at its input, as in Izacard and Grave (2021). In contrast, our ELECTRA based extractive reader R1-D1 (Extractive) shows large gains compared to extractive state-of-the-art, while having the same retriever as DPR. We hypothesize this may be caused by ELECTRA pre-training method, which shows strong performance through variety of tasks and we further show that it is also due to training and inference with large input size of 128 passages and better objective (discussed in Section 4.2 and Appendix G). Only system that matches the performance of our extractive reader is the concurrent work on UnitedQA-E (Cheng et al., 2021), which uses advanced regularization and HardEM techniques. We note that these are orthogonal to our approach and could potentially lead to further improvements.

<sup>8</sup>Note that we train on data from retriever, not reranker.

<sup>9</sup>Matching strategies are described in Appendix C.

<sup>10</sup>Due to the numerous decoder computations in answer re-ranking.

	Method	NQ	TQ	# $\theta$
Extractive	BM25+BERT (Mao et al., 2020)	37.7	60.1	110M
	Hard EM (Min et al., 2019a)	28.1	50.9	110M
	Path Retriever (Asai et al., 2020)	32.6	-	447M
	Graph Retriever (Min et al., 2019b)	34.5	56.0	110M
	ORQA (Lee et al., 2019)	33.3	45.0	220M
	REALM (Guu et al., 2020)	40.4	-	660M
	ProQA (Xiong et al., 2021)	34.3	-	220M
	DPR (Karpukhin et al., 2020)	41.5	56.8	220M
	RDR (Yang and Seo, 2020)	42.1	57.0	110M
	GAR+DPR (Mao et al., 2020)	43.8	-	626M
	ColBERT (Khattab et al., 2020)	48.2	63.2	440M
	RIDER (GAR+DPR) (Mao et al., 2021)	48.3	-	626M
	UnitedQA-E (Cheng et al., 2021)	51.8	68.9	440M
	Generative	BM25+SSG (Mao et al., 2020)	35.3	58.6
T51.1+SSM (Roberts et al., 2020)		35.2	61.6	11B
RAG (Lewis et al., 2020)		44.5	56.8	516M
DPR+SSG (Min et al., 2020)		42.2	-	516M
FiD-base (Izcard and Grave, 2021)		48.2	65.0	333M
FiD-large (Izcard and Grave, 2021)		51.4	67.6	848M
FiD-large++ (Izcard et al., 2020)		54.7	<b>73.3</b>	848M
UnitedQA-G (Cheng et al., 2021)		52.3	68.6	880M
UnitedQA (Ens. E+G+G) (Cheng et al., 2021)		54.7	70.5	1.87B
Ours	R1-D1 (Generative)	49.9	65.4	848M
	R1-D1 (Extractive)	50.8	65.0	445M
	R2-D2 (21M)	<b>55.0</b>	69.9	1.29B
	R2-D2 (21M) w/ HN-DPR	<b>55.9</b>	-	1.29B

Table 2: Comparison with the state-of-the-art in EM. # $\theta$  denotes the estimated amount of model parameters. Symbol - reports the result only for smaller system with 220M parameters.

Finally, we find that our R2-D2 system with 21M passages corpus is competitive even with FiD++, which uses DPR retriever improved via knowledge distillation, and 26M passage corpus, which also includes lists. Additionally, we evaluate our model with a better retrieval model (HN-DPR) based on the DPR checkpoint where hard negatives are mined using the retrieval model itself<sup>11</sup>. Note that we do not compare EfficientQA with state-of-the-art, as the previous works didn't reported results on dev set we use for testing.

#### 4.1 Reranker Performance

Next, we compare the performance of our retriever, reranker and reader with Accuracy@K in Figure 2. The passage reranker improves the accuracy consistently and we observe the same trend on other datasets (Appendix A). We also include analysis, where we rerank each passage  $p_i$  according to its  $s_{passage}^i$  score from extractive reader. We observe results similar or even better to reranker for  $K < 10$ , indicating the extractive reader reranks well on its own. However, in subsequent experiments we do not replace the reranker with reader because: (i) passage reranker has fewer parameters, (ii) extractive reading can run in parallel

<sup>11</sup><https://cutt.ly/Ux5Yt4h>

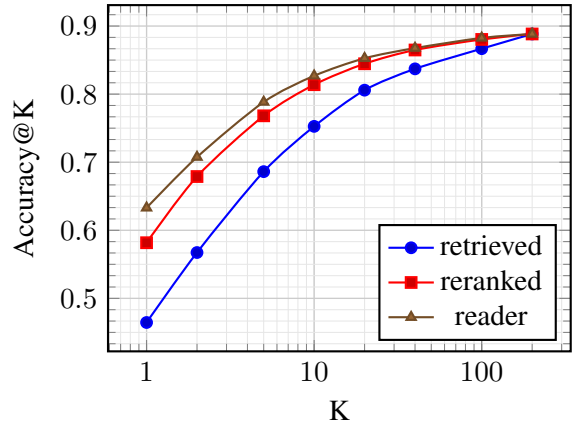


Figure 2: Accuracy@K on test-data of NQ-Open.

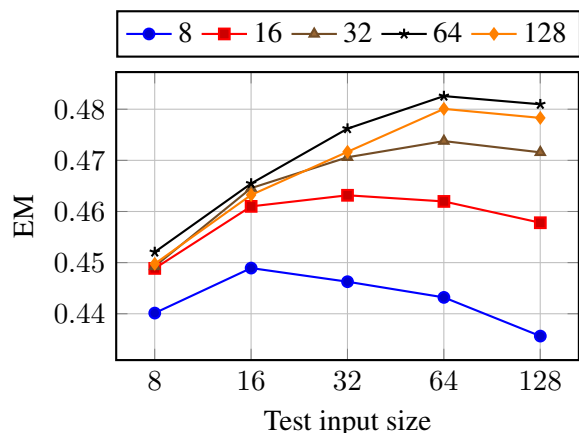


Figure 3: Influence of test input size on extractive reader's performance for various train input sizes (different curves) on NQ-Open test dataset.

with reranking and generative reading as extractive reader is not benefiting from reranking, and (iii) passage reranking scores often improve results during score aggregation (see Section 4.4).

#### 4.2 Extractive Reader Performance

In order to investigate the influence of the number of input passages on the extractive reader's performance, we trained multiple ELECTRA-base models, each with different input size. In test time, we evaluate each of them on various input sizes. Figure 3 shows that increasing train/test input size has a positive influence on extractive reader's performance. However, input size 128 doesn't seem to increase the performance anymore.

#### 4.3 Ablations

The ablations are listed in Table 3. We ablate results without using passage reranker, with separate readers and their combination and with different

Readers	Fusion	NQ-Open			TQ-Open			EfficientQA		
		ret.	+rr	$\Delta$	ret.	+rr	$\Delta$	ret.	+rr	$\Delta$
ext	-	50.78	50.72	-0.06	65.01	65.46	0.45	47.00	47.56	0.56
gen	-	49.92	50.69	0.77	65.38	<u>69.14</u>	3.76	44.83	47.33	2.50
ext+gen	naive	51.88	52.44	0.56	66.17	68.01	1.84	47.06	49.11	2.05
ext+gen	aggr	54.13	54.90	0.77	67.42	68.66	1.24	50.44	52.00	1.56
ext+gen	aggr+bd	54.07	<b>54.99</b>	0.92	67.37	<b>69.94</b>	2.57	49.72	<b>52.22</b>	2.50

Table 3: Ablation study. We report results for extractive (ext), generative (gen) and both readers (ext+gen) without (ret.) and with reranking (+rr). The  $\Delta$  column shows the exact match difference caused by passage reranking.

	$P_*$	$\emptyset$	$\{r\}$	$\{rr\}$	$\{r, rr\}$
NQ-Open	$\{e\}$	50.72	51.41	51.55	51.69
	$\{g\}$	52.44	52.88	53.35	53.19
	$\{e, g\}$	54.63	<b>55.10</b>	54.82	54.90
TQ-Open	$\{e\}$	65.54	65.64	65.60	65.61
	$\{g\}$	68.25	68.17	68.21	68.26
	$\{e, g\}$	68.45	68.57	<b>68.66</b>	<b>68.66</b>

Table 4: Results for different pipeline components used for score aggregation on NQ-Open a TQ-Open. See text for details.

stages of component fusion. Namely, performing a *naive* answer re-ranking by generative reader means the system chooses the most probable answer span among the top- $M$  spans provided by the extractive reader according to generative reader log-probabilities as shown in equation (9). Analogously, the *aggr* fusion denotes that the system chooses the most probable answer span according to aggregated scores, as in equation (11). Finally, the *aggr+bd* fusion denotes the binary decision, as shown in equation (12).

As expected, we observe that reranker improves the results consistently for generative model in all cases. The gains are especially large for TQ-Open (over 3.7 EM, underscored in Table 3). In fact, the results are comparable to Izacard and Grave (2021), suggesting that using the FiD reader with smaller context window and reranker is a reasonable alternative to memory inefficient FiD with large input size. Furthermore as expected, the extractive reader without reranker already has top-128 passages at the input, and improvements from the passage reranking are only negligible if any (less than 1 EM).

Finally, the results on NQ-Open and EfficientQA suggest applying the binary decision does not bring

Table 5: Results for binary decision on NQ-Open and TQ-Open for different aggregated pipeline components from Table 4.

large improvements over the score aggregation if any. However, notice that this is not the case for TQ-Open, where the generative reader performs significantly better compared to extractive reader, suggesting both component fusions play important role in the system.

#### 4.4 Component Fusion

Furthermore, we analyze the performance of each component combination in the score aggregation and its impact on the component fusion via binary decision. Both fusions are tuned on validation data and reported on test data of the NQ-Open and TQ-Open datasets. See Appendix B for analysis on additional datasets. Table 4 shows all relevant combinations of ranker  $r$ , reranker  $rr$ , extractive reader  $e$  and generative reader  $g$  probabilities used in score aggregation. In overall, we observe that combining retriever and reranker scores with the reader leads to better or equal performance. On NQ-Open, we observe minor improvements up to  $\sim 1$  EM. However, there is no difference on TQ-Open.

The impact of adding a binary decision after the score aggregation is shown in Table 5. Interestingly, the binary decision component significantly improves the performance only without reranked

Readers	Ensemble	EM	$\Delta_{ext}$	$\Delta_{gen}$
	-	46.79	-	-
ext	2 models	48.30	1.51	-
	3 models	48.59	1.80	-
	-	45.00	-	-
gen	2 models	46.30	-	1.30
	3 models	46.59	-	1.59
ext+gen	aggr	<b>49.92</b>	<b>3.13</b>	<b>4.92</b>

Table 6: Comparison between ensembling via posterior averaging and score aggregation on NQ-Open.

answer scores ( $\{e\}$  rows in both tables). However, fusing the generative and extractive reader via binary decision performs significantly worse on NQ-Open than fusing both readers together with score aggregation ( $\{e\}$  row in Table 5 vs.  $\{e, g\}$  row in Table 4). As already noted in ablations, we find this to be quite the opposite for TQ-Open. We hypothesize that the binary decision is strong in cases, where generative reader performs better to extractive reader (the case of TQ-Open). We argue that if the generative reader is better, the abstractive answer should be used far more often, than when it’s not. We support the hypothesis by analyzing the proportion of test samples, on which the binary decision component activated (i.e. an abstractive prediction was selected). On NQ-Open, the component almost never activated (only on 3.5% samples), but this proportion is much higher (26.6%) on TQ-Open.

#### 4.5 Comparison with Posterior Averaging

Finally, we compare our score aggregation with the ensemble computed via posterior probability averaging. In particular, we train three extractive and generative base-sized models initialized with different random seed. We do not use reranker in this experiment, and set train/test input size of extractive reader to 32. We assess the predictions using the averaged posterior probabilities and compare their average performance with score aggregation in Table 6. Concretely, we compare with average of all 2 model ensembles (2 models) and with an ensemble of all 3 checkpoints (3 models). We observe two to three times improvement of score aggregation over the posterior probability averaging on NQ-Open test data.

## 5 Related Work

**Passage reranking.** Previous work in QA based on neural nets used Bi-LSTM encoders (Wang et al., 2018; Lee et al., 2018) that score each document independently. Over time, Bi-LSTM were replaced by BERT-like transformer encoders (Qiao et al., 2019; Wang et al., 2019a). For document ranking, Nogueira et al. (2019) proposed a multi-stage architecture. The first stage scores each document independently, and the second estimates the more relevant document from all document pairs. Another document ranking approach uses the seq2seq model to generate a true or false answer to the document’s relevance to the query (Nogueira et al., 2020). Recent works have often focused on effective reranking. Xin et al. (2020) achieved inference speedup using early exiting, Jang and Kim (2020) proposed a smaller and faster model, and Mao et al. (2021) came up with a method which uses reader’s predictions to rerank the passages. Our reranker is most similar to Nogueira and Cho (2019); Luan et al. (2021), except that unlike in IR, we assume there is just one correct passage and thus train our model via categorical cross-entropy.

**Multipassage Reading Comprehension** Related work considers generative and extractive approaches towards modeling the reader. The generative reader generates an answer while conditioned on question alone (Roberts et al., 2020), or question with relevant passages (Lewis et al., 2020; Min et al., 2020). Izacard and Grave (2021) showed it suffices to concatenate the passages in the decoder of seq2seq model, increasing the amount of top-passages the model can depend on dramatically. The extractive reader used in Open-QA assumes that the answer is a continuous span string in located in retrieved paragraphs (Chen et al., 2017). Clark and Gardner (2018) proposed to aggregate the probabilities of distantly supervised answer matches via maximum marginal likelihood (MML). Lin et al. (2018) proposed to denoise distantly supervised answer string matches in MML via paragraph-ranker. Cheng et al. (2020a) experimented with different assumptions for MML, showing improvement when marginalizing over components of span probability independently. Fajcik et al. (2020) proposed to model joint span probability directly via compound objective, instead of modeling the probability of span’s start and end independently. Karpukhin et al. (2020) incorporated an independent passage classifier loss to his MML



objective. Our objective is similar to the last work, except that it uses joint component and also optimizes MML over relevant passages' probabilities.

**Component Fusion.** Yang et al. (2019) also combined BM25 ranker and reader scores via linear combination. Our work can be seen as an extension of this idea to combining the scores of all pipeline's components. Iyer et al. (2020) proposed a system which directly learns to rerank question-passage-answer triplets proposed via extractive model. However, reranking answers from their large extractive model via large reranker leads to ~1 EM improvement absolute, whereas R2-D2s score aggregation improves 4 to 5 EM w.r.t. the extractive reader. Concurrently with our work, Cheng et al. (2021) proposed hard voting ensembling scheme to combine the reader predictions. Firstly, each model from an ensemble produces its best prediction, then the votes for identical predictions are combined, omitting the scores produced by the individual models. The authors obtained best results using two FiD readers and single extractive reader, leading to 1.6 and 2.4 EM improvement on TQ-Open and NQ-Open, compared to their best single extractive or generative model.

## 6 Conclusion

This work proposed R2-D2, a novel state-of-the-art pipeline for open-domain QA based on 4 components: retriever, reranker, generative reader and extractive reader. We showed that employing a reranker is a reasonable alternative to using large passage counts at the input of both the extractive and the generative reader. Our results on NQ-Open and EfficientQA showed that the extractive and the generative reader could perform equally in Open-QA, although the generative reader is twice the size of the extractive reader. On the other hand, we observe the extractive reader underperforms on TQ-Open. We hypothesize, that the cause is (1) the complexity of trivia questions with many entities, which often require combining evidence from multiple passages — these are impossible to answer for the extractive reader by design — and (2) the expensive hyperparameter search, as we used NQ-Open hyperparameters also for TQ-Open. Contrary to belief based on the results on different datasets (Yang et al., 2019; Wang et al., 2019b; Izacard and Grave, 2021), we found the extractive reader can also benefit from larger input sizes, both in training and test time. Finally, we proposed a component

fusion, which allows merging the complementary behavior of generative and extractive approaches along with the ranking components and found it improves the results significantly. Due to its heterogeneous and modular nature, our pipeline forms an ideal base for future research of component integration in modern Open-QA.

## Acknowledgments

We would like to thank Jan Doležal for implementing an R2-D2 demo. This work was supported by the Czech Ministry of Education, Youth and Sports, subprogram INTERCOST, project code: LTC18054. The computation used the infrastructure supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90140).

## References

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2020. [Learning to retrieve reasoning paths over wikipedia graph for question answering](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *arXiv preprint arXiv:2004.05150*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *ArXiv*, abs/1604.06174.
- Hao Cheng, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2020a. [Probabilistic assumptions matter: Improved models for distantly-supervised document-level question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5657–5667, Online. Association for Computational Linguistics.
- Hao Cheng, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2020b. [Probabilistic assumptions matter: Improved models for distantly-supervised document-level question answering](#). In

- Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5657–5667, Online. Association for Computational Linguistics.
- Hao Cheng, Yelong Shen, Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2021. [UnitQA: A hybrid approach for open domain question answering](#). 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 3080–3090, Online. Association for Computational Linguistics.
- Christopher Clark and Matt Gardner. 2018. [Simple and effective multi-paragraph reading comprehension](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 845–855, Melbourne, Australia. Association for Computational Linguistics.
- Martin Fajcik, Josef Jon, Santosh Kesiraju, and Pavel Smrz. 2020. [Rethinking the objectives of extractive question answering](#). *arXiv preprint arXiv:2008.12804*.
- 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*.
- Srinivasan Iyer, Sewon Min, Yashar Mehdad, and Wen-tau Yih. 2020. [RECONSIDER: re-ranking using span-focused cross-attention for open domain question answering](#). *CoRR*, abs/2010.10757.
- Gautier Izacard and Edouard Grave. 2020. Distilling knowledge from reader to retriever for question answering. *arXiv preprint arXiv:2012.04584*.
- Gautier Izacard and Edouard 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, Online. Association for Computational Linguistics.
- 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*.
- Youngjin Jang and Harksoo Kim. 2020. [Document re-ranking model for machine-reading and comprehension](#). *Applied Sciences*, 10(21).
- Mandar Joshi, Eunsol Choi, Daniel 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, Vancouver, Canada. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Omar Khattab, Christopher Potts, and Matei Zaharia. 2020. Relevance-guided supervision for openQA with colBERT. *arXiv preprint arXiv:2007.00814*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Jinhyuk Lee, Seongjun Yun, Hyunjae Kim, Miyoung Ko, and Jaewoo Kang. 2018. [Ranking paragraphs for improving answer recall in open-domain question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 565–569, Brussels, Belgium. Association for Computational Linguistics.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. [Question and answer test-train overlap in open-domain question answering datasets](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008, Online. Association for Computational Linguistics.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Yankai Lin, Haozhe Ji, Zhiyuan Liu, and Maosong Sun. 2018. [Denoising distantly supervised open-domain question answering](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1736–1745, Melbourne, Australia. Association for Computational Linguistics.

- 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*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. Sparse, dense, and attentional representations for text retrieval. *Transactions of the Association for Computational Linguistics*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2020. Generation-augmented retrieval for open-domain question answering. *arXiv preprint arXiv:2009.08553*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. Reader-guided passage reranking for open-domain question answering. *arXiv preprint arXiv:2101.00294*.
- 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. *arXiv preprint arXiv:2101.00133*.
- Sewon Min, Danqi Chen, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019a. [A discrete hard EM approach for weakly supervised question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2851–2864, Hong Kong, China. Association for Computational Linguistics.
- Sewon Min, Danqi Chen, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2019b. Knowledge guided text retrieval and reading for open domain question answering. *arXiv preprint arXiv:1911.03868*.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. [AmbigQA: Answering ambiguous open-domain questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5783–5797, Online. Association for Computational Linguistics.
- Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. [Document ranking with a pre-trained sequence-to-sequence model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, Online. Association for Computational Linguistics.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. [Multi-stage document ranking with BERT](#). *arXiv preprint arXiv:1910.14424*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Yifan Qiao, Chenyan Xiong, Zhenghao Liu, and Zhiyuan Liu. 2019. [Understanding the behaviors of BERT in ranking](#). *arXiv preprint arXiv:1904.07531*.
- 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.
- 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, Online. Association for Computational Linguistics.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018. [R3: Reinforced ranker-reader for open-domain question answering](#). In *AAAI*, pages 5981–5988.
- Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019a. [Multi-passage BERT: A globally normalized BERT model for open-domain question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5878–5882, Hong Kong, China. Association for Computational Linguistics.
- Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019b. [Multi-passage BERT: A globally normalized BERT model for open-domain question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5878–5882, Hong Kong, China. Association for Computational Linguistics.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ji Xin, Rodrigo Nogueira, Yaoliang Yu, and Jimmy Lin. 2020. [Early exiting BERT for efficient document ranking](#). In *Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing*, pages 83–88, Online. Association for Computational Linguistics.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate nearest neighbor negative contrastive learning for dense text retrieval. *arXiv preprint arXiv:2007.00808*.
- Wenhan Xiong, Hong Wang, and William Yang Wang. 2021. [Progressively pretrained dense corpus index for open-domain question answering](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2803–2815, Online. Association for Computational Linguistics.
- Sohee Yang and Minjoon Seo. 2020. Is retriever merely an approximator of reader? *arXiv preprint arXiv:2010.10999*.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. [End-to-end open-domain question answering with BERTserini](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 72–77, Minneapolis, Minnesota. Association for Computational Linguistics.



## A Additional Accuracy@K Analysis

Analysis of Accuracy@K on NQ-Open development data in Figure 4a, on EfficientQA data is shown in Figure 4b, and on TQ-Open development data in Figure 4c and test data in Figure 4d.

## B Additional Component Fusion Analysis

This section includes results analogical to Tables 4, 5 on EfficientQA and development data of NQ-Open and TQ-Open (Tables 7, 8).

	$P_*$	$\emptyset$	$\{r\}$	$\{rr\}$	$\{r, rr\}$
NQ-Open	$\{e\}$	48.38	48.94	48.85	49.14
	$\{g\}$	49.99	50.49	50.35	50.47
	$\{e, g\}$	51.79	<b>51.97</b>	51.82	51.80
TQ-Open	$\{e\}$	65.07	65.21	65.16	65.24
	$\{g\}$	67.68	67.72	67.73	67.76
	$\{e, g\}$	68.13	<b>68.19</b>	68.17	68.12
EfficientQA	$\{e\}$	47.56	48.33	48.89	48.72
	$\{g\}$	49.11	49.56	50.22	50.11
	$\{e, g\}$	50.78	51.67	50.89	<b>52.00</b>

Table 7: Score aggregation on validation data of NQ-Open, TQ-Open and EfficientQA.

	$P_*$	$\emptyset$	$\{r\}$	$\{rr\}$	$\{r, rr\}$
NQ-Open	$\{e\}$	50.65	51.24	51.01	51.17
	$\{g\}$	50.36	50.91	50.68	50.90
	$\{e, g\}$	52.24	<b>52.29</b>	52.27	52.07
TQ-Open	$\{e\}$	69.03	69.03	69.01	68.99
	$\{g\}$	69.54	69.46	69.62	69.70
	$\{e, g\}$	69.77	<b>69.79</b>	69.67	69.61
EfficientQA	$\{e\}$	48.33	50.06	49.39	49.67
	$\{g\}$	48.94	49.50	50.06	49.72
	$\{e, g\}$	50.78	51.83	50.94	<b>52.22</b>

Table 8: Binary decision on NQ-Open, TQ-Open and EfficientQA.

## C Data Pre-processing

This section describes how the training datasets for reranker and extractive reader are filtered, and how the distant supervision labeling is generated. Note not each example contains golden passage, as not each example can be mapped to the used dump

of Wikipedia. We use the same golden passage mapping as Karpukhin et al. (2020).

For passage reranking, the input must contain at least one positive example. We meet this condition either by adding a golden passage or searching for the passage with an answer in the top-400 results retrieved by DPR. In detail about the search, first the Simple tokenizer proposed in DrQA<sup>12</sup> tokenizes each passage and golden answer. The positive example is the best-scored tokenized passage that contains an exact match with one of the tokenized answers. Note the search proceeds in the same way as in DPR’s Accuracy@K implementation<sup>13</sup>.

The extractive reader is trained only on samples which contain exact match to at least one of the annotated answers in the top-1 passage, or golden passage if it is available. The exact match is performed on the subword token level (i.e. in ELECTRA’s tokenization).

Next, the span annotations are extracted from the passages at the reader’s input. Note each sample may contain multiple answers. The annotations for each answer in each sample are obtained differently in retrieved passages and in the golden passage. For retrieved passages, we search for the answer’s exact matches in passages, and use each match as target annotation. For golden passage, we also search for the answer’s exact matches in it. If there is none, the answer is soft matched with single sub-sequence of golden passage, which yields highest non-zero F1 score. The F1 soft match is also performed on the subword token level. Therefore answers with zero highest F1 soft match with golden passage and no exact match in any of the reader’s input passages are discarded.

### C.1 Upper Bound on F1 Matching

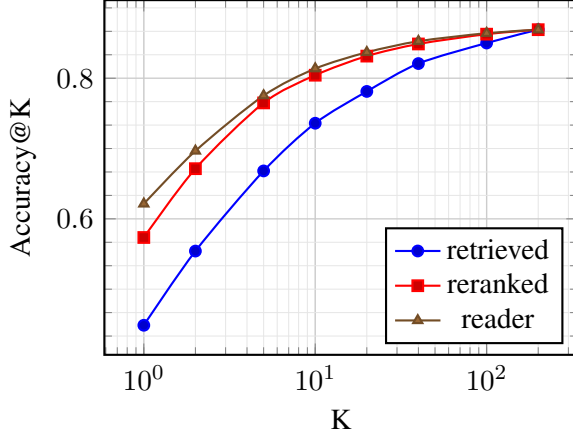
Because the brute-force computation of a span with the greatest nonzero F1 score is potentially very demanding, we found the length limit for spans that are worth searching (see Theorem C.2).

To compare brute-force with upper bound implementation, we run an experiment on 16,741 passages (retrieved for NQ-Open dev). The average time per passage for brute-force approach was 121 ms while it was only 9 ms for implementation that uses the upper bound optimization.

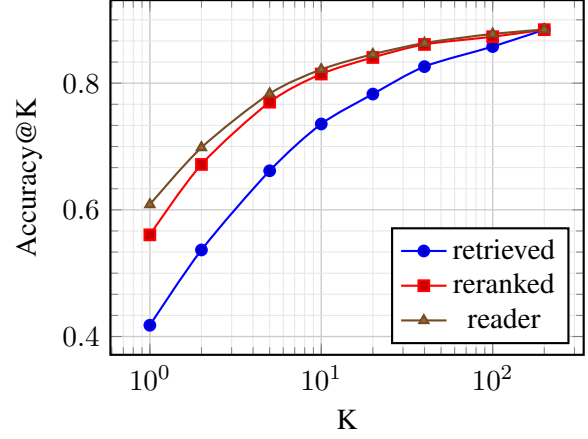
The soft match is described in Algorithm 1. It assumes that there is no exact match.

<sup>12</sup><https://github.com/facebookresearch/DrQA>

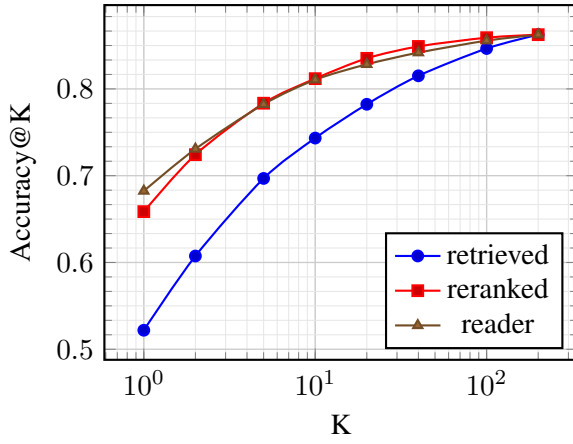
<sup>13</sup><https://github.com/facebookresearch/DPR>



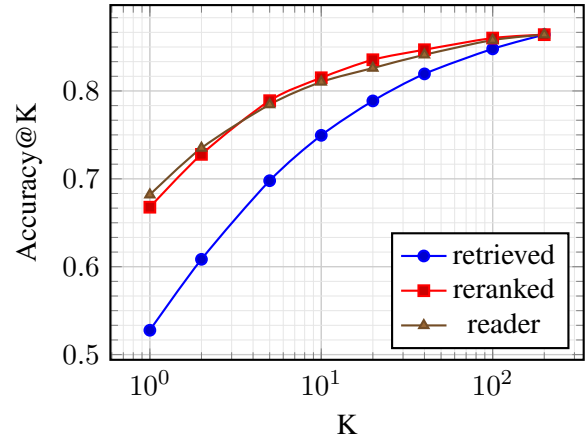
(a) NQ-Open (dev).



(b) EfficientQA.



(c) TQ-Open (dev).



(d) TQ-Open (test).

Figure 4: Analysis of Accuracy@K on different datasets.

**Algorithm 1** Soft match**Require:** set of spans  $S$  and answer span  $a$ 

```

1: function SOFTMATCH( $S, a$ )
2:    $actSize \leftarrow 1$ 
3:    $lenLimit \leftarrow 2$ 
4:    $bestSpan \leftarrow \text{None}$ 
5:    $bestScore \leftarrow 0$ 
6:   while  $actSize < lenLimit$  do
7:     for all  $t \in S$  of size  $actSize$  do
8:        $score \leftarrow F1(t, a)$ 
9:       if  $score > bestScore$  then
10:         $bestSpan \leftarrow t$ 
11:         $bestScore \leftarrow score$ 
12:         $lenLimit \leftarrow |a| \frac{|t| + |a| - s_{ta}}{s_{ta}}$ 
13:      $actSize \leftarrow actSize + 1$ 
14:   return  $bestSpan$ 

```

**Lemma C.1.** Let  $t$  and  $a$  be non-empty spans and  $0 < s_{ta} \leq |a|$  number of shared tokens for them. Then<sup>14</sup>

$$|t| \leq |a| \frac{|t| + |a| - s_{ta}}{s_{ta}}. \quad (13)$$

*Proof.* To prove it by contradiction assume that

$$|t| > |a| \frac{|t| + |a| - s_{ta}}{s_{ta}}, \quad (14)$$

then

$$s_{ta}|t| > |a||t| + |a||a| - |a|s_{ta}, \quad (15)$$

and also  $0 < s_{ta} \leq |a|$ , thus  $|a||a| - |a|s_{ta} \geq 0$ . Therefore even if we assume that  $|a||a| - |a|s_{ta} = 0$ . We get

$$\begin{aligned} s_{ta}|t| &> |a||t| \\ s_{ta} &> |a|, \end{aligned} \quad (16)$$

<sup>14</sup> $|x|$  symbolises number of tokens in span  $x$ .

which is in contradiction with  $0 < s_{ta} \leq |a|$ .  $\square$

**Theorem C.2.** *Let  $S$  be a set of non-empty spans,  $a$  an non-empty answer span,  $t$  non-empty trial span,  $0 < s_{ta} \leq |a|$  is number of shared tokens for  $t$  and  $a$ , and  $S_b = \{z | z \in S \wedge |z| \geq |a| \frac{|t|+|a|-s_{ta}}{s_{ta}}\}$ . Then the theorem states that*

$$\forall x \in S_b (\text{F1}(x, a) \leq \text{F1}(t, a)). \quad (17)$$

*Proof.* To prove it by contradiction assume that

$$\exists x \in S_b (\text{F1}(x, a) > \text{F1}(t, a)). \quad (18)$$

F1 score can be expressed as:

$$\text{F1}(b, c) = \frac{2s_{bc}}{|b| + |c|}, \quad (19)$$

thus

$$\frac{2s_{xa}}{|x| + |a|} > \frac{2s_{ta}}{|t| + |a|}. \quad (20)$$

From Lemma C.1  $|t| \leq |x|$ . Therefore  $s_{ta} < s_{xa}$ , to satisfy the inequality (in equation 20), and we know that  $0 < s_{xa} \leq |a|$ . So let the  $s_{xa} = |a|$  (the maximum) then

$$\begin{aligned} \frac{2|a|}{|x| + |a|} &> \frac{2s_{ta}}{|t| + |a|} \\ |a|(|t| + |a|) &> s_{ta}|x| + s_{ta}|a| \\ |x| &< |a| \frac{|t| + |a| - s_{ta}}{s_{ta}}, \end{aligned} \quad (21)$$

which is in contradiction with  $x \in S_b$ .  $\square$

## D Softmax Notation

Usually, softmax function  $\sigma : \mathbb{R}^K \rightarrow \mathbb{R}^K$  is defined as:

$$\sigma(v)_i = \frac{e^{v_i}}{\sum_{j=1}^K e^{v_j}}. \quad (22)$$

However, some parts of this work used variant of softmax that is defined as follows:

$$\text{softmax}_{x \in D}(f(x))_y = \frac{e^{f(y)}}{\sum_{x \in D} e^{f(x)}}, \quad (23)$$

where  $D$  is the input set,  $f : D \rightarrow \mathbb{R}$ ,  $y \in D$ .

## E Decoding the Distributions from the Extractive Reader

We analyzed the subsets of joint probability space over spans obtained via multiplication of distributions as explained in section 2.2 in Table 9. The factors of this space are the distribution given by the outer product of independent probability distributions  $\mathbf{P}_{start}(\cdot)\mathbf{P}_{end}(\cdot)^\top$  denoted as I, joint probability distribution  $\mathbf{P}_{joint}(\cdot)$  denoted as J, and passage distribution  $\mathbf{P}_{passage}(\cdot)$  denoted as C.

Factorization	NQ-dev	NQ-test	EfficientQA
I	48.32	50.58	47.33
J	48.53	<b>51.25</b>	<b>47.83</b>
I+J	<b>48.57</b>	50.83	<b>47.83</b>
I+C	48.22	50.55	47.22
J+C	48.49	51.11	47.56
I+J+C	48.50	50.86	47.67

Table 9: The results of the pipeline with different types of extractive reader’s distribution used for decoding. See text for details.

marginalizes independently	joint comp.	start&end comp.	EM
-	-	✓	45.42
-	✓	✓	45.41
✓	-	✓	45.71
✓	✓	-	<b>47.09</b>
✓	✓	✓	47.06

Table 10: Ablation of loss components on NQ-Open test dataset using ELECTRA-base model.

## F Passage Reranker Revision

In preliminary experiments of this work we used a Longformer encoder (Beltagy et al., 2020) with concatenated passages at it’s input to benefit from the early fusion between passages. In particular, the passages at the Longformer’s input were shuffled and concatenated, and we used presoftmax score computed from the first Longformer’s output representation of each passage as the *rerank* function. The passages were shuffled with a fixed seed in both, training and test time. Therefore each passage was scored not only according to the question but also according to other passages. However, we did not observe any significant benefits when we used the Longformer setup over a RoBERTa which scores each passage independently (see Table 11).

Readers	Fusion	NQ-Open (dev)			NQ-Open (test)			TQ-Open (test)			EfficientQA		
		Long.	RoB.	$\Delta$	Long.	RoB.	$\Delta$	Long.	RoB.	$\Delta$	Long.	RoB.	$\Delta$
ext	-	48.50	48.38	-0.12	50.86	50.72	-0.14	65.41	65.46	-0.05	47.67	47.56	-0.11
gen	-	49.34	49.40	0.06	51.50	50.69	-0.81	68.85	69.14	-0.29	47.33	47.33	0.00
ext+gen	naive	49.91	49.99	0.08	53.43	52.44	-0.99	67.82	68.01	-0.19	49.06	49.11	0.05
ext+gen	aggr	52.05	51.80	-0.25	54.96	54.90	-0.06	68.49	68.66	-0.17	51.56	52.00	0.44
ext+gen	aggr+bd	52.36	52.07	-0.29	55.01	54.99	-0.02	69.62	69.94	-0.32	51.06	52.22	1.16

Table 11: Exact match comparison of Longformer (Long.) and RoBERTa (RoB.) based passage reranker.

## G Ablating the Extractive Reader’s Objective

Firstly let us demonstrate that start and end components in the used loss (see equation 8) perform summation over both inter-passage and intra-passage combinations of starts and ends:

$$\begin{aligned}
& -\log \sum_{s \in S} P_{start}(s) - \log \sum_{e \in E} P_{end}(e) = \\
& = -\log \sum_{s \in S} P_{start}(s) \sum_{e \in E} P_{end}(e) = \quad (24) \\
& = -\log \sum_{s \in S, e \in E} P_{start}(s) P_{end}(e).
\end{aligned}$$

Where  $S = starts(C_{rr})$ ,  $E = ends(C_{rr})$  and distribution dependencies are dropped for clarity. Inter-passage combinations obviously do not correspond to a real answer. Even though that this loss does not reflect the task correctly, it achieves better results (see Table 10) than the following loss

$$-\log \sum_{c \in C_{rr}} \sum_{a_e \in answers(c)} P_e(a_e|q, C_{rr}) \quad (25)$$

that marginalizes components jointly, and thus the summation is done only through intra-passage start-end combinations. Such results agree with previous work (Cheng et al., 2020b).

Table 10 also shows that the joint component improves the independent loss variant, but not the other one that marginalizes jointly. We hypothesize that this is because the loss in equation 25 already considers only the intra-passage start-end pairs. Lastly, Table 10 shows that using just the joint and passage component is sufficient for NQ-Open, which agrees with Fajcik et al. (2020).

## H Results According to Question and Answer Test-Train Overlap

In addition to evaluation on the TQ-Open and NQ-Open shown in Table 2, we also report results

on subsets of these datasets in Table 12, as split by Lewis et al. (2021). We compare R2-D1 (retriever, reranker and extractive or generative reader, marked as gen and ext respectively) and R2-D2 (ext+gen) to official results on FiD (Izcard and Grave, 2021).

## I Inference Speed of Our Implementation

While optimizing the R2-D2’s inference speed was not the main focus of this paper, we show that even our unoptimized implementation can be used in practice in small scale. We analyze the speed of our implementation on NQ-Open test data in Table 13. The times were measured on a workstation with Intel Xeon Silver 4214 48-core CPU, 188GB RAM and Nvidia 2080Ti 12GB GPU. Table columns show settings with and without passage reranker. Table rows are split into two parts; *intermediate* rows show time spent by the pipeline’s single component (e.g., row ext. reader shows what time the pipeline spent by running just the ext. reader), and *total* rows show the total time taken by the whole pipeline. The retriever and reranker infer with batch sizes 32 and 100 respectively, the readers run with batch size 1.

We note that in retrieval, we do not use any approximate K-NN algorithm to facilitate retrieval of top-K nearest passages and instead do the dot product with the matrix of passages directly on the CPU. Secondly, we note that we do not parallelize the inference of generative reader and extractive reader. Thirdly, notice the difference in extractive reader’s speed with and without passage reranker is caused by its different input size (see details of extractive reader’s experiments setup in subsection 3.2). Finally, we compare the speed of our approach using FiD with 25 and 100 input passages, like in the original FiD implementation<sup>15</sup>. The ratios of our measurements are compared explicitly in Table 14.

<sup>15</sup>We simply pass 100 input passages to the model trained with 25 passages in the experiment.



Model	NQ-Open				TQ-Open			
	Total	Question Overlap	Answer Overlap Only	No Overlap	Total	Question Overlap	Answer Overlap Only	No Overlap
FiD	51.40	71.30	48.30	34.50	67.60	87.50	66.90	42.80
ext+gen	54.99	75.00	48.89	39.91	69.94	90.18	71.53	44.83
gen	50.69	70.06	46.98	34.04	69.14	87.50	70.32	44.83
ext	50.72	72.53	45.40	35.11	65.46	83.63	66.42	39.46
$\Delta_{\text{gen}}$	4.30	4.94	1.91	5.87	0.80	2.68	1.21	0.00
$\Delta_{\text{ext}}$	4.27	2.47	3.49	4.80	4.48	6.55	5.11	5.37

Table 12: Results on the overlapping and non-overlapping parts of test sets for NQ and TQ. *Total* column corresponds to overall result on the whole dataset, as reported before, *Question Overlap* corresponds to samples with train-test question overlap and answer overlap, *Answer Overlap Only* corresponds to samples with answer overlap, but no question overlap, and *No Overlap* corresponds to samples with no overlap between train and test sets.

	Modules	Rankers	
		retriever	+reranker
intermediate	retriever	0.21	0.21
	passage reranker	-	1.94
	ext. reader	2.21	0.35
	gen. reader (25)	0.55	0.55
	answer reranker (25)	3.11	3.11
	gen. reader (100)	1.85	-
	answer reranker (100)	11.67	-
total	ext	2.41	2.19
	gen (25)	0.76	2.70
	ext+gen (25)	6.08	6.16
	gen (100)	2.06	-
	ext+gen (100)	15.94	-

Table 13: Inference times on NQ-Open in seconds per question. See text for details.

Setup ratios	Modules			
	only gen.	ans. reranker	gen pipe.	ext+gen pipe.
gen(100) / gen(25)	3.36x	3.75x	2.71x	2.62x
rr+gen(25) / gen(25)	*1.00x	*1.00x	3.55x	1.01x
gen(100) / rr+gen(25)	*3.36x	*3.75x	0.76x	2.58x

Table 14: Ratios of inference times on NQ-Open. First two columns compare the speed in stage of generating abstractive answer (only gen.) and answer reranking (ans. reranker). The subsequent columns compare speed of whole pipeline just with generative reader and no component fusion (gen pipe.) and full R2-D2 pipeline (ext+gen pipe.). Row gen(100)/gen(25) compares the speedup of pipeline when using just 25 passages in FiD’s input (denoted as gen(25)) instead of 100 (denoted as gen(100)). Row rr+gen(25)/gen(25) shows speedup gained from not using passage reranker (denoted as rr). Row gen(100)/rr+gen(25) compares the speed of using rr and gen(25) instead of gen(100) (with no passage reranking). Results marked with \* are not affected by passage reranking component, as they only measure speed of pipeline’s individual component. For instance, table shows that doing answer reranking with generative reader with just 25 passages at its input runs 3.75x faster than doing answer reranking with generative reader that uses 100 passages.