

NeurIPS 2020 EfficientQA Competition: Systems, Analyses and Lessons Learned

Sewon Min², Jordan Boyd-Graber³, Chris Alberti¹, Danqi Chen⁴, Eunsol Choi⁵,
Michael Collins¹, Kelvin Guu¹, Hannaneh Hajishirzi², Kenton Lee¹,
Jennimaria Palomaki¹, Colin Raffel^{1,6}, Adam Roberts¹, Tom Kwiatkowski¹

¹Google Research ²University of Washington ³University of Maryland
⁴Princeton University ⁵University of Texas at Austin ⁶University of North Carolina

{sewon,hannaneh}@cs.washington.edu jbg@umiacs.umd.edu
danqi@cs.princeton.edu eunsol@cs.utexas.edu
{craffel,adarob,tomkwiat}@google.com

Patrick Lewis^{T1}, Yuxiang Wu^{T1}, Heinrich Küttler^{T1}, Linqing Liu^{T1},
Pasquale Minervini^{T1}, Pontus Stenetorp^{T1}, Sebastian Riedel^{T1,3},
Sohee Yang^{T2}, Minjoon Seo^{T2}, Gautier Izacard^{T3,7}, Fabio Petroni^{T3,7},
Lucas Hosseini^{T3,7}, Nicola De Cao^{T3}, Edouard Grave^{T3,7}, Ikuya Yamada^{T4},
Sonse Shimaoka^{T4}, Masatoshi Suzuki^{T4}, Shumpei Miyawaki^{T4}, Shun Sato^{T4},
Ryo Takahashi^{T4}, Jun Suzuki^{T4}, Martin Fajcik^{T5}, Martin Docekal^{T5},
Karel Ondrej^{T5}, Pavel Smrz^{T5}, Hao Cheng^{T6}, Yelong Shen^{T6}, Xiaodong Liu^{T6},
Pengcheng He^{T6}, Weizhu Chen^{T6}, Jianfeng Gao^{T6}, Barlas Oguz^{T7}, Xilun Chen^{T7},
Vladimir Karpukhin^{T7}, Stan Peshterliev^{T7}, Dmytro Okhonko^{T7},
Michael Schlichtkrull^{T7}, Sonal Gupta^{T7}, Yashar Mehdad^{T7}, Wen-tau Yih^{T7}

^{T1}UCLNLP & Facebook AI ^{T2}NAVER Clova ^{T3}Facebook AI Paris & London
^{T4}Studio Ousia, Tohoku University & RIKEN ^{T5}Brno University of Technology
^{T6}Microsoft Research & Dynamic 365 AI ^{T7}Facebook AI

Abstract

We review the EfficientQA competition¹ from NeurIPS 2020². The competition focused on open-domain question answering (QA), where systems take natural language questions as input and return natural language answers. The aim of the competition was to build systems that can predict correct answers while also satisfying strict on-disk memory budgets. These memory budgets were designed to encourage contestants to explore the trade-off between storing large, redundant, retrieval corpora or the parameters of large learned models. In this report, we describe the motivation and organization of the competition, review the best submissions, and analyze system predictions to inform a discussion of evaluation for open-domain QA.

¹<http://efficientqa.github.io/>

²<https://neurips.cc/Conferences/2020/CompetitionTrack>

1 Introduction

Open-domain question answering (QA) is emerging as a benchmark method of measuring computational systems’ abilities to read, represent, and retrieve knowledge [1, 2, 3, 4]. Recently, this task has been addressed by a diverse set of approaches that navigate multiple documents [5, 6], index large corpora of text [4, 7, 8], represent world knowledge in the parameters of a neural network [9], or consolidate knowledge from multiple passages [10, 11].

The EfficientQA competition, held at NeurIPS 2020, required contestants to build self-contained systems that contain all of the knowledge required to answer open-domain questions. There were no constraints on how the knowledge is stored—it could be in documents, databases, the parameters of a neural network, or any other form. However, three competition tracks encouraged systems that store and access this knowledge using the smallest number of bytes, including code, corpora, and model parameters. Specifically, EfficientQA had four tracks: 1) best accuracy overall (unconstrained); 2) best accuracy, system size under 6GiB; 3) best accuracy, system size under 500MiB; 4) smallest system to get 25% accuracy. These memory budgets were designed to encourage contestants to explore the trade-off between storing and accessing large, redundant, retrieval corpora, structured data stores, or the parameters of large learned models.

This paper summarizes the findings from the competition. Section 2 describes the competition in detail; Section 3 presents a description of the best performing models in each track; Section 4 introduces a new human evaluation of system accuracy, and Section 5.2 provides an analysis of this evaluation and a comparison to the automatic evaluation used during development. Finally, Section 6 pits the top systems against human trivia experts.

1.1 Key takeaways

The top submissions in each of EfficientQA’s four tracks significantly outperformed the provided baselines. All top submissions use a retrieval corpus and a neural network answering module. However, the nature of the retrieval corpus and answering module differs drastically across the tracks (Table 1).

Unrestricted and 6GiB tracks The top submissions to the unrestricted track and the 6GiB track (Section 3.1; 3.2) outperformed the state-of-the-art baselines from April 2020 by nearly 20%. They achieved this improvement by combining the state-of-the-art retrieval systems [8, 12] with answer generation [11]; leveraging the state-of-the-art in text generation [13] and text encoding [14]; modeling not only text but also tables and lists from Wikipedia; and combining the extractive and generative answer prediction. The top submissions to the 6GiB track additionally massively reduced the size of their indexed corpus and made use of the state-of-the-art in compression, with minimal impact on answer accuracy.

500MiB and smallest tracks To get under 500MB (Section 3.3), the systems made more drastic changes. The submission from NAVER Clova drastically reduced the size of their indexed corpus and reused a single Transformer model for the retriever and the reader, winning the 500MiB track according to the human evaluation. The even smaller UCLNLP-FB system took a novel approach in generating a large corpus of question-answer pairs, indexing it, and retrieving the most similar question to the input question. This approach, with two systems with different sizes in question-answer corpus, won both the 500MiB track and the smallest 25% track according to the automatic

evaluation.

Automatic vs. human evaluation The human evaluation supports the observation that automatic metrics often incorrectly penalize correct predictions in the open-domain setting [1, 9]. We also investigate the effect of question ambiguity on evaluation—the questions are from Natural Questions [15] and are often ambiguous without the associated evidence document [16]. In Section 4 we define an annotation scheme that supports multiple estimations of accuracy, corresponding to different definitions of correctness for answers to ambiguous questions. Almost all systems’ accuracy increased by 20%-25% under the strictest definition of correctness. The increase doubled when we relaxed the definition of correctness to permit any semantically valid interpretation of the question. However, we found that this relaxation led to the acceptance of many answers that are pragmatically unlikely to be correct and, in Section 5.2, we present a discussion of this trade-off as well as suggestions for future evaluation. We will use the newly annotated answers to provide extra estimations of system accuracy on the continued EfficientQA leaderboard³.

2 Competition Overview

2.1 Data

The competition uses English questions and answers from the Natural Questions dataset [15, NQ]: real user questions, issued to the Google search engine, along with multiple reference answers from Wikipedia⁴ We choose NQ over related tasks such as SQuAD [17] and TriviaQA [18] because the real user questions in NQ are interpretable outside of a document context (unlike SQuAD), and less amenable to traditional IR approaches than over-complete trivia questions in TriviaQA [4].

While the original NQ task was posed as a reading comprehension task, in which evidence documents are provided, recent work has adapted the NQ data to the open-domain setting [4, 19, 5, 6, 7, 9, 8] by taking examples with short answers (up to five tokens) and discarding evidence documents. In the open-domain setting, NQ contains 88k one-way annotated training examples and 4k five-way annotated development examples. For EfficientQA, we introduce a new test and development set constructed in the same way as the original NQ, but labeled slightly after (early 2019 rather than through 2018).⁵ Our test set was kept hidden from contestants, and submissions were made by uploading solutions to an automatic leaderboard.

2.2 Measuring System Size

All submissions to the three restricted tracks were submitted as self-contained Docker⁶ images. We defined system size to be the on-disk, at-rest size of this image. We chose this approach to avoid confusion about distinctions between data, model parameters, and code. However, this choice did lead to a significant engineering efforts for the very smallest systems, which were not able to build on the standard Docker templates for the predominant deep-learning libraries. While the organizers did provide a small reference system based on T5 [9] that used TensorFlow serving⁷, this did not

³<https://ai.google.com/research/NaturalQuestions/efficientqa>

⁴NQ was collected over the course of 2017 and 2018. Most solutions to the open-domain variant use a Wikipedia dump from December 2018. See [4] for details.

⁵This temporal change caused slight degradation in performance for all systems.

⁶<https://www.docker.com>

⁷<https://www.tensorflow.org/tfx/serving/docker>

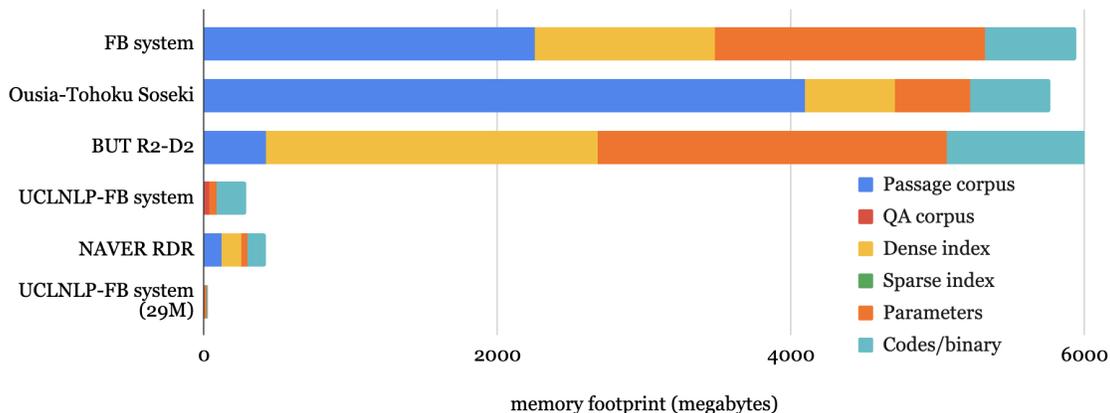


Figure 1: Memory footprint of each system component.

support most submissions, and the very smallest systems required clever compilation strategies on top of their modeling enhancements.

2.3 QA Metrics

Automatic evaluation The accuracy of each systems’ predicted answers is judged against reference annotations, annotated by five human workers. We follow the literature in using exact match between predicted and reference answers after minor normalization [4].

Human evaluation Due to the ambiguities inherent in language and question-answering in general, five reference answers are often not exhaustive, and systems predict correct answers that are judged incorrect according to the automatic metrics. To rectify this, we sent predictions from each of the systems for further human rating by three separate raters, to get a better estimation of accuracy. Details of the annotation scheme are given in Section 4 and an analysis of the ratings is presented in Section 5.2.

2.4 Competition Schedule

The competition was announced in June 2020, along with baselines and tutorials. The official leaderboard for restricted settings was launched on September 14th, 2020, and participants had two months to submit their systems (November 14th, 2020). Test set results in restricted settings are published to a public leaderboard. On November 15th, 2020, after the leaderboard was frozen, questions on the test set were released and participants submitted the predictions from their systems for an unrestricted track. Finally, predictions from the top systems from each leaderboard were sent for human evaluation, which was completed in the end of November. The same set of top systems was invited to present their systems at the NeurIPS event on December 12th and in this paper. The human vs. computer competition was held on December 6th (details in Section 6).

In total, we had 39 submissions from 18 unique teams, seven of which were affiliated or co-affiliated with universities.

Model	Affiliation	retr	pruned W	answer	others
Unrestricted track					
REALM	Organizers	p	X	ext	
DPR	Organizers	p	X	ext	
MS UnitedQA	Microsoft & Dynamics 365	p	X	ext+gen	
FB Hybrid	Facebook AI	p	X	gen	lists/tables
6GiB track					
DPR-subset	Organizers	p	train data	ext	
T5-XL+SSM	Organizers	X	-	gen	
FB system	FAIR-Paris&London	p	learned	gen	lists, lrzip compression
Ousia-Tohoku Soseki	Studio Ousia, Tohoku U & RIKEN	p	page view	ext	ZPAQ compression
BUT R2-D2	Brno U of Technology	p	learned	ext+gen	
500MiB track					
T5-Small+SSM	Organizers	X	-	gen	
UCLNLP-FB system	UCL & FAIR	(q,a)	-	-	data augmentation
NAVER RDR	NAVER Clova	p	learned	ext	single Transformer, KD
25% smallest track					
T5-XL+SSM	Organizers	X	-	gen	
UCLNLP-FB system (29M)	UCL & FAIR	(q,a)	-	-	data augmentation

Table 1: A list of the baselines and systems from participants, along with the team affiliation and key distinction between systems. *retr* means what the systems retrieves, e.g., passages (p), question-answer pairs ((q, a)), or none (X). *pruned W* indicates how Wikipedia is pruned; X indicates there is no pruning and full Wikipedia is used. *answer* indicates whether the answer is extracted (*ext*) or generated (*gen*). Other key distinctions are shown in the last column; *KD* indicates knowledge distillation.

3 Systems

We describe a set of provided baselines and systems from participants. Systems from participants include the top 1–3 submissions per track based on automatic measure, considering margins between the submissions. Table 1 summarizes the team affiliation and key distinctions between the systems, and Figure 1 shows the memory footprint of each system component.

3.1 Unconstrained track

3.1.1 Baselines

REALM/DPR Both REALM [7] and DPR [8] use the retrieval-reader framework. They retrieve the top K passages from Wikipedia by mapping the question and each passage to dense vectors and employing a maximum inner product search. The top K passages are then fed into an extractive reader which predicts the start and the end position of the answer. A key distinction between REALM and DPR is that REALM is trained with a self-supervised objective and a joint objective, while DPR is trained in a pipeline manner with gold positives and distantly supervised negatives.

3.1.2 Top systems

Both of the top two systems make use of the previous state-of-the-art system, such as DPR [8] and Fusion-in-Decoder [11]. They are greatly enhanced by a better training objective, aggregating the

answers from the extractive and generative models, or incorporating lists and tables from Wikipedia.

MS UnitedQA The UnitedQA system consists of three components: **Retrieval**, **Reading**, and **Re-ranking**. First, it uses a BERT-based [20] dense retriever [8] to fetch top 100 relevant passages from the English Wikipedia dump for a given question. Second, a hybrid approach, combining both generative (T5 [13, 11]) and extractive (ELECTRA [14]) readers, is used to produce answer candidates from the noisy collection of retrieved passages. In addition, several techniques are applied to facilitate the training of both readers. Specifically, posterior differential regularization (PDR) [21], improved loss-term [22] are explored for enhancing the extractive model, and adversarial training [23] approach, attention bias [24] are used to improve the generative model respectively. At the final re-ranking stage, it combines the generative and extractive model predictions with linear interpolation, and produce the final answer.

On the development set, a relative improvements of 13% is achieved by using PDR & improved loss-term to the extractive reader, 20% by using ELECTRA-large to the BERT-base extractive reader, 4% by using adversarial training & attention bias to the generative reader, and 5% by using the hybrid approach to the best single model.

FB Hybrid This system uses a retriever-reader architecture. The *retriever* is a combination of DPR [8] and Generation Augmented Retrieval (GAR) [12], where each of them retrieves 50 passages. The DPR encoders are trained iteratively, where better hard negatives are mined at each step using the model at the previous step—two iterations are found to be sufficient. In addition, the dense index includes lists and tables, as well as regular text passages. Specifically, 455907 tables and infoboxes are processed from the Natural Questions training corpus using a trivial linearization, which concatenates the text representation of each row with a newline character.⁸ Tables are chunked at 100 tokens with the header included in each chunk. The *reader* is the same as Fusion-in-Decoder (FiD) [11] with T5-large [13], which is given 100 passages from the retriever and generates an answer.

On the development set, relative improvements by an iterative training of DPR, incorporating lists and tables, and using GAR in addition to DPR are 3%, 5% and 4%, respectively. More details of the system can be found in [25].

3.2 6GiB track

3.2.1 Baselines

DPR-subset We create a variant of DPR with pruned Wikipedia. Specifically, we only keep passages from the pages that are paired with the questions on the training set of Natural Questions. This trick reduces the number of passages from 21M to 1.6M.

T5-XL+SSM T5 [13] is a text-to-text Transformer language model utilizing an encoder-decoder architecture that was pre-trained using a “span corruption” objective on a cleaned and de-duplicated subset of Common Crawl. Roberts et al. [9] demonstrated surprisingly good results on Natural Questions Open (among other open-domain benchmarks) in a “closed-book” setting, where the model was only able to access the knowledge stored in its parameters after fine-tuning on the task. Furthermore, the authors achieved significantly improved results with additional pre-training using

⁸This linearization is found to be more effective than converting each row to a statement.

a “salient span masking” objective [7] before fine-tuning. The XL model has approximately 3 billion parameters.

3.2.2 Top systems

All of the top three systems are retrieval-based, and reduce the memory footprint by drastically pruning the Wikipedia corpus. The pruning is either done by a learned neural network or based on a page view. They also made use of model quantization and the state-of-the-art in compression to limit their system’s footprint. They additionally enhanced the previous state-of-the-art [8, 11], as done in an unrestricted track.

FB system This system is based on a retriever-reader approach. The retriever is an ensemble of a dense retriever and GAR. The dense retriever is initialized from BERT-base and is trained by distilling the cross-attention scores of the reader [26]. The reader is Fusion-in-Decoder (FiD) [11], initialized from T5 large [13]. GAR [12] is a retriever with generative query augmentation, based on BM25. The Lucene⁹ index is built on the fly.

Our text corpus initially included 26M passages including the plain text and the lists, but went through an article-level filtering to include 18.8M passages. Specifically, we train a linear classifier, where each Wikipedia article is represented by its title and list of categories. The positive examples to train the classifier are articles retrieved by DPR on the training data. The negative examples are iteratively obtained: articles that are not the positive examples but obtain high classification scores by the classifier from the previous iteration are chosen.

In addition to the base docker image and PyTorch [27], the three main sources of memory usage are: the model parameters, the dense index, and the external source of knowledge (Wikipedia). The model weights are stored using float16, taking 1.6GB of memory. The dense index is compressed by relying on three strategies, described in [28]. First, the vector representations go through dimension reduction from 768 to 256, and further quantized through product quantization, using 2 bits per dimension. The text of Wikipedia is compressed using lrzip¹⁰, leading to a memory usage of 2.2GB.

Ousia-Tohoku Soseki Soseki is an open-domain QA system that adopts a two-stage approach consisting of *Retriever* for passage retrieval, and *Reader* for reading comprehension. Given a question, Retriever obtains top-k candidate passages from Wikipedia. Then, Reader selects the most relevant passage from them and extracts an answer from the passage.

Retriever is based on a re-implementation of the DPR model [8], where two BERT-base-uncased models are used to embed questions and passages. We also quantize passage embeddings to reduce the system size. Embeddings of Wikipedia passages are precomputed and stored using Faiss [29]. Reader is a reading comprehension model based on ELECTRA-large [14]. We added one dense layer for selecting the most relevant passage, and two dense layers for detecting the start and end positions of the answer. Both Retriever and Reader are trained on the NQ-open dataset.

To reduce the system size under 6GB, we compressed the models, passages, and other data files using ZPAQ¹¹ and excluded Wikipedia passages with less than 40 monthly page views, resulting in 18M passages. These tricks do not drop the accuracy.

⁹<https://lucene.apache.org/>

¹⁰<https://github.com/ckolivas/lrzip>

¹¹<http://mattmahoney.net/dc/zpaq.html>

BUT R2-D2 R2-D2 is composed of a dense retriever, a re-ranker and two readers. The dense retriever is based on RoBERTa [30] and is trained via an objective known from DPR. It retrieves $K = 400$ passages from a pruned version of Wikipedia, reduced from 21M to 1.6M passages. Pruning is done via a simple binary classifier based on RoBERTa trained on the data created from the golden passages and negative passages randomly sampled from the index. This classifier obtains 90.22% accuracy. The re-ranker (based on Longformer [31]) concatenates retrieved passages, assigns a score for each passage, and selects $V = 20$ passages. The extractive reader (based on ELECTRA [14]) reads V passages in a similar way as in [32]. It is trained via the marginal likelihood objective combined with the compound objective. The generative reader follows the Fusion-in-Decoder schema and generates the answers. R2-D2 aggregates the output from these two readers using two fusioning methods. First, it reranks the top spans from the extractive reader by feeding them to a generative reader and combines its log likelihood with the log likelihood from the extractive reader through a linear combination. They are then further aggregated with abstractive answers from the generative reader that are generated independently, through another linear combination. The parameters are stored in float16 and compressed using ZIP.

Based on the result on the dev set, pruning Wikipedia corpus drops the accuracy by 6% relative. Aggregating the answers from two readers outperforms extractive reader only and generative reader only by 10% and 9% relative, respectively.

3.3 500MiB track

3.3.1 Baselines

T5-Small+SSM This is a smaller version of the T5-XL+SSM baseline above, with 512 hidden dimension instead of 4096, 6 layers instead of 24, and 8 attention heads instead of 64. It contains 60 million parameters in total.

3.3.2 Top systems

The top performing submissions in this track made more drastic changes to get under 500MiB. They have completely different approaches from each other.

UCLNLP-FB system This system is based on the approach from [33], consisting of a database of question-answer pairs and a retriever which returns the answer of the most similar stored question for an input question. This approach is attractive as it performs well using low parameter-count models, and less space is needed to store question-answer pairs than full Wikipedia passages. An improvement over [33] is made by increasing the coverage of the database by generating question-answer pairs, improving the retrieval accuracy, and including a re-ranker.

The *Question-Answer Pair Generator* is similar to [34, 35]. First, a passage selection model $P(c)$ is trained to identify appropriate passages from Wikipedia. For high-probability passages (w.r.t. $P(c)$), the system performs Named Entity Recognition to extract likely answers a . Next, the system generates questions q from (c, a) using a BART [36] model $P(q | a, c)$ fine-tuned on Natural Questions. Finally, the system filters question-answer pairs with a “global consistency filter”—if an open-domain QA system (Fusion-in-Decoder [11]) is given q and generates an answer that is consistent with a , the (q, a) pair is added to the database. This is because generated (q, a) pairs are sometimes incorrect or unanswerable without the passage they were generated from. The

final database consists of Natural Questions, the EfficientQA development set and 2.3M generated question-answer pairs.

The *Retriever* consists of TF-IDF and a bi-encoder retriever based on ALBERT-base [37], trained using a set-up similar to RAG [38]. A hundred question-answer pairs are retrieved from TF-IDF and are fed into the bi-encoder retriever to get top-10 question-answer pairs.¹² The *Reranker* is a cross-encoder module based on ALBERT-large, trained discriminatively using hard negatives produced by the retriever. It reranks the top-10 question-answer pairs from the retriever, and the answer of the top-1 question-answer pair is chosen as the final answer.

The system is further compressed using TFLite, quantization, and Alpine Linux.

NAVER RDR RDR is a lightweight retrieve-and-read system, consisting of a single Transformer (MobileBERT [39]) that performs both retrieval and reading, an index of dense passage vectors and the filtered Wikipedia corpus. The Transformer serves as a dense retriever as DPR does with two differences: (1) a single Transformer is used to encode both the question and the passage, and (2) the embeddings are defined as the first 128 out of 512 dimensions of the output vectors from the Transformer that correspond to the [CLS] token. The same Transformer serves as an extractive reader by producing the scores of the start and the end of the answer span as well as the passage reranking score. A distillation technique is used for training, by minimizing the KL divergence between its start, end, and reranking scores and those of a fully trained DPR reader. The Transformer is further finetuned in an iterative manner as a retriever and as a reader, which significantly boosts the performance.

The index consists of 1.2M 128-dimensional INT8 vectors (scalar-quantized and unsigned), which are the dense embeddings of the subset of the 21M passages in Wikipedia, filtered by a RoBERTa [30]-based binary classifier trained with logistic regression to exclude uninformative passages. Positives to train this classifier are the top 200 passages for each question on Natural Questions dataset and EfficientQA development set, retrieved by [40], a DPR retriever further finetuned on hard negatives using knowledge distillation from a DPR reader. Negatives are uniformly drawn from the set of 21M passages, excluding positives.

On the development set, pruning Wikipedia and dimension reduction of the dense embeddings reduce the memory footprint by 11.5x and 1.8x, and decreases the accuracy by 2.7% and 1.4%, respectively. Using a single MobileBERT for retrieval and reading with knowledge distillation and iterative fine-tuning reduces the memory footprint by 30% and improves the accuracy by 1.8%.

3.4 25% smallest track

3.4.1 Baselines

T5-XL+SSM Among the provided baselines, the smallest system with an accuracy of over 25% is T5-XL+SSM, the same system described in Section 3.2.1. This system is 5.65GB and achieves an accuracy of 28%.

3.4.2 Top systems

UCLNLP-FB system (29M) This system is the same system as UCLNLP-FB system in the 500MiB track with following minor modifications, to further decrease the memory: (1) the retriever

¹²TF-IDF is used to reduce the inference time rather than memory footprint, as there was a time limit in the competition.

is just TF-IDF instead of a combination of TF-IDF and a bi-encoder, (2) the reranker uses ALBERT-base instead of ALBERT-large, and (3) there are 40k generated question-answer pairs, instead of 2.3M.

4 Human Annotations of Correctness

The EfficientQA development and test sets have up to five reference answers per question. Due to the variability of language, these five answer strings are often not exhaustive, and systems predict correct answers that are judged incorrect according to the automatic metrics. To rectify this, we sent predictions from each of the systems in Section 3 for human rating to get a better estimation of accuracy.

Each system prediction was sent for rating by three separate annotators: 1) the annotator first works on understanding the meaning and intent of the question (with a web search if necessary). 2) The annotator then determines whether the question is ambiguous, i.e., whether the question can lead to multiple different answers depending on factors such as: when the query was asked; where the query was asked; some unspoken intent of the questioner; or the opinion of the person giving the answer. 3) Finally, the annotator determines whether each answer is “definitely correct” (correct given a usual interpretation of the question), “possibly correct” (could be correct, given some interpretation of the question), or “definitely incorrect”.

Since the original Natural Questions data was collected more than a year before the start of the EfficientQA competition, the denotation of some questions may have changed over time (e.g. “who won the last season of bake-off”). Rather than determine a single correct point in time for these questions, we asked our annotators to assume that the query could have been asked at any time since the web has existed and choose the “possibly correct” label for answers that may or may not have been correct when the question was asked.

The final rating is an aggregation of ratings from three annotators: if at least 2/3 raters determined it to be “definitely correct”, the label is “definitely correct”. If at least 2/3 raters determined it to be either “definitely correct” or “possibly correct”, the label is “possibly correct”. The pairwise agreement of the human ratings is as follows: 69.2% (Cohen’s $\kappa = 53.8$) for 3-way ratings, 85.7% (Cohen’s $\kappa = 71.4$) for whether the prediction is definitely correct or not, and 76.7% (Cohen’s $\kappa = 53.4$) for whether the prediction is possibly correct or not.¹³

Section 5.1 ranks systems according to both the “definitely correct” and “possibly correct” ratings, and Section 5.2 details an analysis of the human annotations. We recommend to use “definitely correct” ratings as the main human evaluation performance because “possibly correct” ratings are often overly permissive (Section 5.2).

5 Results & Analyses

5.1 Results

All of the five systems in the unrestricted track and the 6GiB track significantly outperform the state-of-the-art (Table 2) at the beginning of the competition—DPR (36.6%) and REALM (35.9%).

¹³An agreement in whether the question is ambiguous or not is 61.3% (Cohen’s $\kappa = 22.6$). We conjecture the low agreement is due to an intrinsic difficulty in finding ambiguity [16], and consider the question to be ambiguous if at least one annotator rated it as ambiguous.

Model	Automatic eval	Human eval	
		Definitely	Possibly
<i>Unrestricted track</i>			
MS UnitedQA	54.00	65.80 (+21.9%)	78.12 (+44.7%)
FB Hybrid	53.89	67.38 (+25.0%)	79.88 (+48.2%)
<i>6GiB track</i>			
FB system	53.33	65.18 (+22.2%)	76.09 (+42.7%)
Ousia-Tohoku Soseki	50.17	62.01 (+23.6%)	73.83 (+47.2%)
BUT R2-D2	47.28	58.96 (+24.7%)	70.33 (+49.2%)
<i>500MiB track</i>			
UCLNLP-FB system	33.44	39.40 (+17.8%)	47.37 (+41.7%)
NAVER RDR	32.06	42.23 (+31.7%)	54.95 (+71.4%)
<i>25% smallest track</i>			
UCLNLP-FB system (29M)	26.78	32.45 (+21.2%)	41.21 (+53.9%)

Table 2: Summary of the result. For human evaluation result, relative improvements over the automatic evaluation are indicated in parenthesis. Following our analysis of the annotations, we determine that the ‘Definitely correct’ human ratings provide the best estimation of answer quality.

Systems in the 6GiB track approach the unrestricted track’s accuracy; for instance, the accuracy FB system is comparable to the accuracy of the top systems in the unrestricted track. The improvements in the 500MiB track are also impressive; both the top two systems significantly beat T5-small (17.6%).

Discrepancy between automatic eval and human eval Human raters find 13% and 17% of the predictions that do not match the reference answers to be definitely correct or possibly correct, respectively, overall increasing the accuracy of the systems. Most systems showed 17–25% and 41–54% improvement in accuracy when using definitely correct and possibly correct human evaluation respectively, compared to automatic evaluation metric which only consider exact string match to existing reference answers. An exception is NAVER RDR, which achieves significantly larger improvements (32% and 71%, respectively). We also found that when the gap in automatic measure between systems is marginal (around or smaller than 1%), human evaluation may change the rankings between the models.

Agreement between system predictions Figure 2 (left) shows the agreement between system predictions, based on exact match in automatic evaluation. The largest agreement is made between FB Hybrid and FB system, likely because they are both based on DPR and Fusion-in-Decoder. Agreements between systems in the unrestricted and the 6GiB tracks are generally higher, likely because they are all based on retrieval-reader framework and pruning Wikipedia does not hurt too much. Two systems in the 500MiB track have smaller agreement with the other systems, and agree with each other even less.

Ensemble oracle accuracy of the systems Figure 2 (right) reports the ensemble oracle accuracy for each system pair, which considers a prediction to be correct if either system prediction

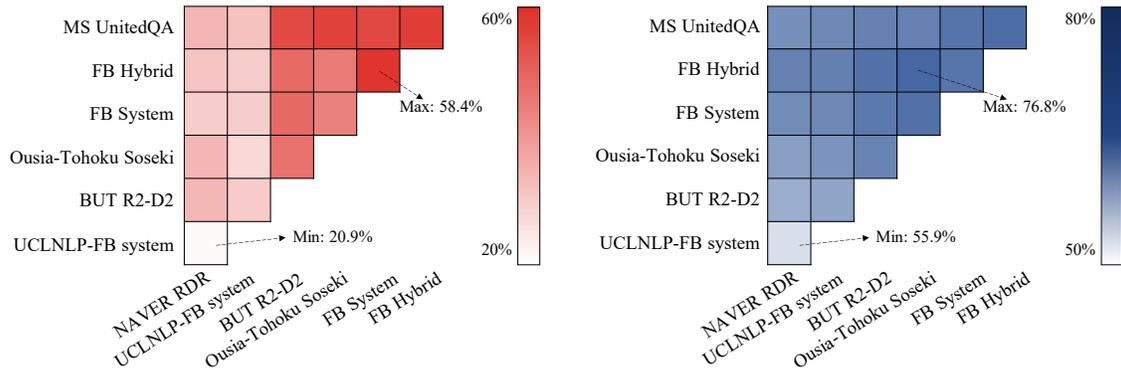


Figure 2: (Left) Agreement between system predictions. (Right) Ensemble oracle accuracy, which considers a prediction correct if at least one of the system predictions is correct (based on “definitely correct” human evaluation).

is correct. The FB Hybrid & Ousia-Tohoku Soseki pair achieves the highest ensemble oracle accuracy, indicating that their system predictions are substantially different from each other compared to other pairs with top performing systems.

5.2 Analyses

We present an analysis of the 50 sample questions where at least one prediction does not match with the gold answer, but is judged as correct by human raters, with a “definitely correct” label or a “possibly correct” label, respectively. The samples are largely divided into three classes: **as valid as gold** (judged as either definitely correct or possibly correct), **valid but not the best** (judged as possibly correct), and **closer to invalid** (judged as possibly correct). We describe fine-grained categories and their percentage¹⁴ (**definitely correct answers** and **possibly correct answers**) here, with statistics summarized in Table 3 and the examples shown in Table 5.

The following describes categories on predictions that are as valid as the gold answers.

- **Semantically the same** (60%, 22%): The prediction is semantically equivalent to the gold, e.g., “about 930 BCE” and “around 930 BCE”.¹⁵
- **Open-ended question** (6%, 4%): There is a large set of distinct, plausible answers, mainly because the question is vague or open-ended.
- **Ambiguous entity/event references** (20%, 20%): There is a set of distinct answers because the question contains ambiguous references of entities or events. For instance, “Gold woman” in the example in Table 5 may refer to the fictional character (“Ayesh”) or the actress (“Elizabeth Debick”).
- **Different granularity** (14%, 6%): The prediction and the gold answer have different granularity, such as the year (1982) vs. the month (October 1982), or the city (Pawtucket) vs.

¹⁴The total is over 100% because one question may have multiple predictions that fall into different categories.

¹⁵A few cases in this category are due to inconsistent tokenization (e.g., “Mercedes-Benz” and “Mercedes - Benz”), which could potentially be automatically captured but is not handled by the automatic evaluation from the literature.

Category	definitely (%)	possibly (%)
<i>As valid as gold</i>		
Semantically the same	60	22
Open-ended question	6	4
Ambiguous entity/event references	20	20
Answer has different granularity	14	6
Gold is incorrect	2	6
<i>Valid but not the best</i>		
Ambiguous answer type	0	20
Answer is time-dependent	0	18
<i>Less plausible</i>		
Conflicting info in Wikipedia	0	4
Plausible only in certain condition	0	4
Mismatch with question intent	0	8
Incorrect prediction	0	2

Table 3: Analysis of predictions that are rated as *definitely correct* or *possibly correct* by human raters. Examples of each category are shown in Table 5. Note that the total is over 100% as one question may fall into multiple categories.

the state (Rhode Island).

- **Incorrect gold (2%, 6%)**: The gold annotation is incorrect, e.g., in Table 5, “1965” is incorrect as it is when the first attraction appeared in the theme park.

The following describes categories on possibly correct predictions that are as valid but not the best answers.

- **Ambiguous answer type (20%)**: There is a unique entity or event that can be the answer to the question, but there is ambiguity on which exact text should be presented as the answer. For instance, both the episode title (“Somber News”) and the air date (“February 23 2013”) are the valid answer to the question in Table 5.
- **Answer is time-dependent (18%)**: The answer depends on the time of the question being asked. Questions in this category usually involve recurring events such as sports games or elections.

Finally, we present categories for predictions that are closer to invalid answers.

- **Conflicting information in Wikipedia (4%)**: While the gold is the only valid answer, the English Wikipedia contains incorrect information that supports that prediction is the answer to the question. Consider the question in Table 5. While “Fort Hood” is an incorrect answer by a fact, the Wikipedia page¹⁶ states “Fort Hood is the most populous U.S. military installation in the world.”¹⁷

¹⁶https://en.wikipedia.org/wiki/Fort_Hood

¹⁷This question is used in the human vs. computer competition and the human team also predicted “Fort Hood”.

- **Plausible only in certain conditions (4%)**: Prediction may be valid in certain conditions, but is incorrect in general cases. For instance, for the question in the table, “president of India” may only be valid in India.
- **Mismatch with question intent (8%)**: The prediction is not definitely incorrect but is supposedly not the intended answer to the question. For instance, the example question in Table 5 is supposedly asking for the answer that is different from the great depression.
- **Incorrect prediction (2%)**: The prediction is definitely incorrect, e.g., “185 billion years ago” is not the correct answer to the last question in Table 5 as it is older than an estimated formation of the universe.

Our two main takeaways are as follows. First, the automatic evaluation confirms the observation of Voorhees and Tice [1] that it is insufficient in capturing semantically equivalent answers, which are responsible for 60% of the definitely correct predictions. Second, ambiguity arises frequently in the questions in different levels, allowing multiple, semantically different answers to be valid. This is consistent with Min et al. [16], which reported that around half of the questions in Natural Questions contain ambiguity, due to ambiguous references of entities, events or properties, or time-dependency of the answer. Based on our human evaluation, annotations on ambiguity have low agreement rate (61.3%, Cohen’s $\kappa = 22.6$), and predictions with the same level of plausibility are often marked as “definitely correct” or “possibly correct” by different human raters. We note the notions of “correctness” and “plausibility” are not binary, and are instead often dependent on pragmatic interpretation of the questioner’s intent. For example, the question “who has the most superbowl rings” could be read as “which person (including coaches) has the most superbowl rings”, “which player has the most superbowl rings”, “which team has the most superbowl rings”. All three annotators identified this question as being ambiguous but they disagreed about the validity of the different readings. The three raters were split three ways when rating the correct answer (“Pittsburgh Steelers”) for the last interpretation. Meanwhile there were no “incorrect” ratings, and 2/3 “definitely correct” ratings given to the correct answer (“Tom Brady”) for the second interpretation, despite the fact that two coaches have more superbowl rings. Clearly, the annotators are applying some personal interpretation of the questioner’s intent and answer plausibility.

While we believe that many real world questions do require some non-literal assumptions about the questioner’s intent, and we believe that the natural language processing community should not shy away from that task, we also acknowledge that there is work to be done in creating better, non-binary, definitions of correctness. Section 6 also contrasts these interpretations of correctness with the more rigid definition used by the Trivia community.

Performance on unambiguous questions To better understand the effect of ambiguity on the ranking of different solutions, we also evaluate system performance on a subset of the questions that are unambiguous. We define unambiguous questions to be those that (1) have at least three out of five reference answers contain valid short answers¹⁸, and (2) are not labeled as ambiguous by any of three human raters, resulting in 51.5% of the original set. Table 4 shows human evaluation on the original set and this subset of unambiguous questions. Most systems, except UCLNLP-FB system, achieve higher accuracy on unambiguous questions, with the first three systems achieving over or near 80%. Unsurprisingly, the gap between “definitely correct” accuracy and “possibly correct” accuracy is marginal on unambiguous questions.

¹⁸This follows the original Natural Questions approach of using annotator agreement to set a threshold for ‘high quality answers’.

Model	All		Unambiguous Qs	
	Definitely	Possibly	Definitely	Possibly
MS UnitedQA	65.80	78.12	78.24	81.18
FB Hybrid	67.38	79.88	82.65	85.59
FB system	65.18	76.09	79.12	81.47
Ousia-Tohoku Soseki	62.01	73.83	72.94	75.00
BUT R2-D2	58.96	70.55	69.71	72.06
UCLNLP-FB system	39.40	47.37	42.06	43.24
NAVER RDR	42.23	54.95	49.71	53.24
UCLNLP-FB system (29M)	32.45	41.21	28.53	30.29

Table 4: Human evaluation on the original set and a subset of unambiguous questions.

Importantly, the overall rankings are unchanged when we restrict our evaluation set to only unambiguous answers. This suggests that, while question ambiguity may lead to disagreement between annotators at a per-example level, it is not adversely impacting our ability to consistently rank solutions.

Characteristics of *easy* and *hard* questions Table 6 shows questions that are answered correctly by all seven systems (*easy*), or answered incorrectly by all seven systems (*hard*). A common feature of easy questions [41, 42] is that a sentence from Wikipedia provides an explicit support to the question. Such supporting sentences have high lexical overlap with the question and require little paraphrasing (e.g. “the highest population” vs. “the most populous”). Even when the question is not well-formed, all systems may find the correct answer if there is high lexical overlap between the question and the supporting sentence (e.g. the fifth example question in Table 6).

Many of the hard questions have their answers in the tables (e.g. the first question in Table 6). This likely makes them “hard” because all systems except FB Hybrid do not consider the tables, and even FB Hybrid may miss such cases. In the other cases, systems make mistakes even when there is the text that supports the answer (e.g. the next two questions in Table 6). Though the reason for the mistake would be different depending on the systems, we conjecture that such supporting sentences are harder to retrieve or do not support the answer as explicitly as other examples. Occasionally, the question is not well-formed, such as the last one in Table 6.

Error Analysis We further present an error analysis of the top two systems in the 500MiB track: UCLNLP-FB system and NAVER RDR. We choose these two systems because their approaches differ significantly from each other and the 500MiB track fits the main motivation of the EfficientQA competition. We randomly sample 100 questions from the development data. 40, 33 and 27 questions are answered correctly by both systems, one of the systems, and none of the systems, respectively.

Table 7 shows the breakdown of the predictions from **UCLNLP-FB system**, where 50 questions are answered correctly, and the other 50 are not. A majority of the error cases (47 out of 50) is due to retrieving an incorrect question. Out of 47, 25 cases retrieve the question with a different topic from the input questions, many of which discuss different entities (e.g., “don’t it make you feel like dancing” vs. “you make me feel like dancing”). The other 19 cases retrieve the question that discusses the same topic or contains the same key entity, but is asking about different details. For instance, in the example in Table 7, the input question asks the amount *spent for* a film, while the

retrieved question ask the amount *made by* the film. It is also worth noting that the system sometimes gets the correct answer from the retrieved question with a different meaning from the input question, e.g., “What republican is running for mayor of phoenix” vs. “Who is the current mayor of phoenix”, which have the same answer because the current mayor of Phoenix is Republican.

Table 8 shows the breakdown of the predictions from **NAVER RDR**, where 63 questions are answered correctly, and 37 are not. Failure due to pruning the gold passage is rare, being responsible for only 3%. More of the failure cases are due to missing the gold passage, either when choosing 80 passages through dense retrieval (12%), or during cross-attention re-ranking 80 passages to decide on the top 1 passage (15%). Finally, in 7% of the cases, the top 1 passage contains the gold answer but the system fails to extract the correct answer. The gold passages in this category have valid but implicit support of the answer, e.g., the last example in Table 8.

6 Trivia Experts vs Computer Systems

The questions in NQ are posed by humans to computers, and the competition attracted some of the strongest and most efficient QA systems available today. However, humans also answer questions for fun and recreation [43] and the ultimate goal of artificial intelligence is to create machines that answer questions as well as humans [44, known as the Turing test]. Moreover, existing comparisons of human question answering ability often use unskilled humans [17], leading to claims of computers “putting millions of jobs at risk” [45]. Or, in competitions with trivia experts [46], arcane rules of competitions can tilt the playing field toward computers [47] or use unnatural questions [48, 49].

6.1 A Fair Comparison

We advertised our competition to trivia enthusiasts on social media. Teams of up to eight players applied to be part of the competition. We selected five teams to participate in the preliminary competition (results in Section 6.2).

To create a fair competition and to showcase all of the tiers of the efficient QA competition, we offered three ways to answer each question where either humans or computers have more resources to answer a question (Table 9).

To complement the 500MiB systems, humans had to instantly signal when they knew the answer to a question. This reflects instant recall of a fact by a single individual. In the next phase in competition with the 6GiB systems, both humans and computers had more resources: the human team could discuss the answer for thirty seconds, arguing why they believe their answer is correct and computers had over ten times the memory. Finally, to focus on reading comprehension, unlimited systems faced off against the human teams who also had access to snippets from search results using the question as a query. As with the previous phase, they have thirty seconds to discuss their answer.

We selected questions for the human eval based on the following criteria:

- Diverse over topic, ensuring there were questions about history, literature, philosophy, sports, and popular culture. This results in fewer questions about sports and popular culture than the standard NQ distribution.
- Not tied to 2018. Because the NQ dataset came from 2017–2019 and it is difficult for humans to forget the last three years, we excluded questions that depend on the current date.

- Interesting questions. While not strictly adversarial, we wanted to showcase both human and computer ability, so we excluded questions that many humans would not know (e.g., “how many us states are there”) or questions with answers that are difficult to evaluate in the NQ framework (“how many words are in *Les Misérables*?”).
- To avoid the issues described in Section 2.3, we avoided questions that were overly ambiguous, or where the gold answers had the issues mentioned above (answer changes based on time the question was asked, unclear answer type, mismatch with question intention, etc.).

Thus, for the human competition, we exclude the broader interpretation of correct adopted in Section 5.2 to a standard we called “actually correct”, consistent with traditional trivia interpretations [43]. A human judge researched all of the answers to the questions and evaluated whether an answer was correct or not. This allows for more permissive answer lines such as “to pass away” to be accepted for questions like “what does it mean to cross over the rainbow bridge” even though the answer line only lists “heaven”. Again, consistent with trivia best practice, we only required full names in the case of confusion (e.g., “Theodore Roosevelt” or “Franklin D. Roosevelt” instead of just “Roosevelt”).

6.2 Preliminary Competition

To select which of the human teams faced off against the top computers in the final competition and to let the human teams practice this unconventional format, we had the human teams face off against the baseline systems: T5 (500MiB), DPR (6GiB), and REALM (Unlimited).

We set aside an hour for each team. Because better teams can get through more questions, to normalize comparisons per team, we computed the average margin per question (Table 10). Only two teams had higher scores than the computer baseline teams. The top team (Team B), which had multiple professors and trivia champions was clearly the best human team on this set of questions.

6.3 Final Game

The winning team (Team B) went up against the winning computer systems in a final match with fifty questions. Most questions were answered in Phase 1, and those that were not were typically answered in Phase 2 (Figure 3). Only four questions reached Phase 3; these questions that stumped both humans and computers:

- kimi no na wa what does it mean
- who said there are old pilots and there are bold pilots
- what is the second largest city in ethiopia
- who said i’ll make mincemeat out of you

Of these questions, only the second largest city in Ethiopia was answered correctly (the humans were working through a list of cities in Ethiopia they knew). The other questions represent questions that cannot be answered by Wikipedia (“old pilots”), questions that require language capabilities beyond question answering (“kimi no na wa”), or that have a more nuanced answer than NQ can provide. For example, on the questions “who said i’ll make mincemeat out of you”, the human team reasonably thought that the question meant who *originated* the phrase, while NQ’s answer was the cartoon character who is most associated with the phrase ([Klondike Kat](#)).

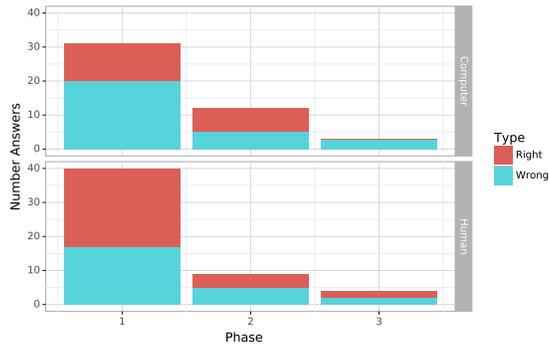


Figure 3: The human–computer competition is broken into three phases, where each phase allows the participants to use more resources to answer the questions. While most questions were answered in Phase 1 (where humans had a distinct advantage), in Phase 2 computers had a slight advantage (0.58 to 0.44 accuracy). However, there were diminishing returns in Phase 3; of the four questions that got to Phase 3, the computers did not answer any correctly.

6.4 Reflection

For question answering, there are multiple interpretations of the phrase “human evaluation”. Human evaluation of answers are important for exposing problems in the dataset, revealing ambiguity, and measuring whether the answers are useful. However, for the ultimate questions of where we have achieved artificial intelligence [44], we need fair comparisons with *skilled* humans. Moreover, to measure whether question answering systems are *useful* for people, we need to create socio-technical systems where humans and computers can answer questions together [50]. Moreover, trivia games are fun. They help illustrate the strengths and weaknesses of the underlying dataset and the QA methods with a spoonful of sugar to make otherwise dull evaluations more interesting.

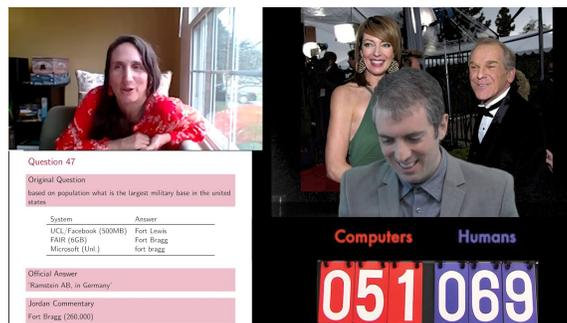


Figure 4: Screenshot from our final game between the top computer systems and trivia experts. Full videos of the human–computer competition are available at <https://go.umd.edu/2020eqa>.

7 Conclusions and Future Work

The EfficientQA competition was held to encourage research in open-domain question answering that focuses on both accuracy and memory footprint. Within two months—a relatively short period of time—39 systems from 18 unique teams participated. The top performing submissions in each of four tracks explore new approaches and significantly advance the state-of-the-art.

All top performing submissions in the unrestricted track and the 6GiB track have a retriever-reader framework, and are enhanced over the state-of-the-art through a better training objective, aggregation of the extractive and generative answer modules, or incorporation of the lists and tables. Systems in the 6GiB track additionally cleverly select a small subset of Wikipedia, only marginally sacrificing accuracy when combined with state-of-the-art compression.

In the other tracks, systems explore novel approaches and achieve impressive improvements over the baselines. UCLNLP-FB system generates a large corpus of question-answer pairs and retrieves the closest question to the input question, winning both of the 500MiB track and the smallest 25% track based on automatic evaluation. NAVER RDR, along with a clever method of dramatically filtering the Wikipedia passages, uses a single Transformer model to retrieve, rerank and extract answers from passages, and wins the 500MiB track based on human evaluation. Still, both systems are behind the unrestricted systems in performance by 20%, indicating significant room for improvements in memory-restricted settings.

A human analysis shows that automatic evaluations of QA systems are not sufficient for thorough evaluations. Human raters find 30% of the predictions that do not match reference answers but are nonetheless correct. This does not affect all systems equally: relative accuracy rises using definitely correct (between 18–32%) and possibly correct (between 42 and a whopping 71%). The rise is mainly due to automatic evaluation failing to capture semantically equivalent answers, time-dependence of the answers, or underlying ambiguity in the questions [16].

Future work in efficient open-domain question answering should continue to explore the tradeoff between system size, accuracy, and abstaining [51, 52]. Moreover, it is important to continue to refine the quality of QA evaluation: not all annotation and not all annotators are created equal. Using trivia enthusiasts for annotation and human benchmarks is a fun and effective evaluation of relative computer QA ability. We would encourage future leaderboards to use some element of human verification in important evaluations (e.g., an annual bake off), removing ambiguity in questions [16], crafting adversarial examples [53, 54, 55, 56], or evaluating whether a response is useful to a user [57, 50].

Acknowledgement

We thank all the participants for taking part and making this a successful competition. We thank Google for providing prizes for computer participants. Boyd-Graber is also supported by NSF Grant IIS-1822494. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the view of the sponsor.

References

- [1] Voorhees, E. M., D. M. Tice. Building a question answering test collection. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in*

- information retrieval*, pages 200–207. ACM, 2000.
- [2] Chen, D., A. Fisch, J. Weston, et al. Reading Wikipedia to answer open-domain questions. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2017.
 - [3] Seo, M., J. Lee, T. Kwiatkowski, et al. Real-time open-domain question answering with dense-sparse phrase index. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2019.
 - [4] Lee, K., M.-W. Chang, K. Toutanova. Latent retrieval for weakly supervised open domain question answering. In *Proceedings of the Conference of Association for Computational Linguistics*. 2019.
 - [5] Min, S., D. Chen, L. Zettlemoyer, et al. Knowledge guided text retrieval and reading for open domain question answering, 2019.
 - [6] Asai, A., K. Hashimoto, H. Hajishirzi, et al. Learning to retrieve reasoning paths over wikipedia graph for question answering. In *Proceedings of the International Conference on Learning Representations*. 2019.
 - [7] Guu, K., K. Lee, Z. Tung, et al. REALM: Retrieval-augmented language model pre-training. In *Proceedings of the Conference on Machine Learning*. 2020.
 - [8] Karpukhin, V., B. Oguz, S. Min, et al. Dense passage retrieval for open-domain question answering. In *Proceedings of Empirical Methods in Natural Language Processing*. 2020.
 - [9] Roberts, A., C. Raffel, N. Shazeer. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 2020.
 - [10] Lewis, P., E. Perez, A. Piktus, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the Conference on Neural Information Processing Systems*. 2020.
 - [11] Izacard, G., E. Grave. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*, 2020.
 - [12] Mao, Y., P. He, X. Liu, et al. Generation-augmented retrieval for open-domain question answering. *arXiv preprint arXiv:2009.08553*, 2020.
 - [13] Raffel, C., N. Shazeer, A. Roberts, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
 - [14] Clark, K., M.-T. Luong, Q. V. Le, et al. Electra: Pre-training text encoders as discriminators rather than generators. In *Proceedings of the International Conference on Learning Representations*. 2020.
 - [15] Kwiatkowski, T., J. Palomaki, O. Redfield, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*, 2019.
 - [16] Min, S., J. Michael, H. Hajishirzi, et al. AmbigQA: Answering ambiguous open-domain questions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 2020.

- [17] Rajpurkar, P., J. Zhang, K. Lopyrev, et al. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of Empirical Methods in Natural Language Processing*. 2016.
- [18] Joshi, M., E. Choi, D. Weld, et al. 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)*, vol. 1, pages 1601–1611. 2017.
- [19] Min, S., D. Chen, H. Hajishirzi, et al. A discrete hard EM approach for weakly supervised question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 2019.
- [20] Devlin, J., M.-W. Chang, K. Lee, et al. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota, 2019.
- [21] Cheng, H., X. Liu, L. Pereira, et al. Posterior differential regularization with f-divergence for improving model robustness, 2020.
- [22] Cheng, H., M.-W. Chang, K. Lee, et al. 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. Association for Computational Linguistics, Online, 2020.
- [23] Ju, Y., F. Zhao, S. Chen, et al. Technical report on conversational question answering, 2019.
- [24] Lewis, M., M. Ghazvininejad, G. Ghosh, et al. Pre-training via paraphrasing, 2020.
- [25] Oguz, B., X. Chen, V. Karpukhin, et al. Unified open-domain question answering with structured and unstructured knowledge. *arXiv preprint*, 2020.
- [26] Izacard, G., E. Grave. Distilling knowledge from reader to retriever for question answering, 2020.
- [27] Paszke, A., S. Gross, F. Massa, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035. Curran Associates, Inc., 2019.
- [28] Izacard, G., F. Petroni, L. Hosseini, et al. A memory efficient baseline for open domain question answering, 2020.
- [29] Johnson, J., M. Douze, H. Jégou. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734*, 2017.
- [30] Liu, Y., M. Ott, N. Goyal, et al. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [31] Beltagy, I., M. E. Peters, A. Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.

- [32] Fajcik, M., J. Jon, S. Kesiraju, et al. Rethinking the objectives of extractive question answering. *arXiv preprint arXiv:2008.12804*, 2020.
- [33] Lewis, P., P. Stenetorp, S. Riedel. Question and answer test-train overlap in open-domain question answering datasets, 2020.
- [34] Alberti, C., D. Andor, E. Pitler, et al. Synthetic QA corpora generation with roundtrip consistency. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2019.
- [35] Lewis, P., L. Denoyer, S. Riedel. Unsupervised question answering by cloze translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2019.
- [36] Lewis, M., Y. Liu, N. Goyal, et al. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2020.
- [37] Lan, Z., M. Chen, S. Goodman, et al. ALBERT: A lite BERT for self-supervised learning of language representations. In *Proceedings of the International Conference on Learning Representations*. 2020.
- [38] Lewis, P. S. H., E. Perez, A. Piktus, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*. 2020.
- [39] Sun, Z., H. Yu, X. Song, et al. MobileBERT: a compact task-agnostic bert for resource-limited devices. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2020.
- [40] Yang, S., M. Seo. Is retriever merely an approximator of reader? *arXiv preprint arXiv:2010.10999*, 2020.
- [41] Sugawara, S., K. Inui, S. Sekine, et al. What makes reading comprehension questions easier? In *Proceedings of Empirical Methods in Natural Language Processing*. 2018.
- [42] Kaushik, D., Z. C. Lipton. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proceedings of Empirical Methods in Natural Language Processing*. 2018.
- [43] Jennings, K. *Brainiac: adventures in the curious, competitive, compulsive world of trivia buffs*. Villard, 2006.
- [44] Turing, A. M. Computers & thought. *Mind*, pages 11–35, 1995.
- [45] Cuthbertson, A. Robots can now read better than humans, putting millions of jobs at risk, 2018.
- [46] Ferrucci, D., E. Brown, J. Chu-Carroll, et al. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3), 2010.
- [47] Boyd-Graber, J., B. Börschinger. What question answering can learn from trivia nerds. *ArXiv*, 10, 2020.

- [48] Boyd-Graber, J., B. Satinoff, H. He, et al. Besting the quiz master: Crowdsourcing incremental classification games. In *Proceedings of Empirical Methods in Natural Language Processing*. 2012.
- [49] Rodriguez, P., S. Feng, M. Iyyer, et al. Quizbowl: The case for incremental question answering. *CoRR*, abs/1904.04792, 2019.
- [50] Feng, S., J. Boyd-Graber. What AI can do for me: Evaluating machine learning interpretations in cooperative play. In *International Conference on Intelligent User Interfaces*. 2019.
- [51] He, H., J. Boyd-Graber, K. Kwok, et al. Opponent modeling in deep reinforcement learning. In *Proceedings of the International Conference of Machine Learning*. 2016.
- [52] Rajpurkar, P., R. Jia, P. Liang. Know what you don't know: Unanswerable questions for SQuAD. In *Proceedings of the Association for Computational Linguistics*. 2018.
- [53] Jia, R., P. Liang. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031. Association for Computational Linguistics, Copenhagen, Denmark, 2017.
- [54] Wallace, E., P. Rodriguez, S. Feng, et al. Trick me if you can: Human-in-the-loop generation of adversarial question answering examples. *Transactions of the Association of Computational Linguistics*, 10, 2019.
- [55] Dua, D., Y. Wang, P. Dasigi, et al. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2368–2378. Minneapolis, Minnesota, 2019.
- [56] Bartolo, M., A. Roberts, J. Welbl, et al. Beat the AI: Investigating adversarial human annotation for reading comprehension. *Transactions of the Association for Computational Linguistics*, 8:662–678, 2020.
- [57] Fan, A., Y. Jernite, E. Perez, et al. Eli5: Long form question answering. *arXiv preprint arXiv:1907.09190*, 2019.

Category	Example
<i>As valid as gold</i>	
Semantically the same	Q: When did israel split into israel and judah Gold: about 930 BCE / Prediction: around 930 BCE Q: What movie is bring it on the musical based off of Gold: Bring It On / Prediction: 2000 film of the same name
Open-ended question	Q: City belonging to mid west of united states Gold: Des Moines / Prediction: kansas city
Ambiguous entity/event references	Q: Gold woman in guardians of the galaxy 2 Gold: Ayesha / Prediction: Elizabeth Debicki
Answer has different granularity	Q: When did the musical cats open on broadway Gold: 1982, 7 October 1982 / Prediction: october 1982 Q: Where did the industrial revolution in the united states begin Gold: in Pawtucket / Prediction: rhode island
Gold is incorrect	Q: When did universal studios become a theme park Gold: 1965 / Prediction: july 15 1964, 1964
<i>Valid but not the best</i>	
Ambiguous answer type	Q: When did naruto find out about jiraiya’s death Gold: Somber News / Prediction: february 23 2013
Answer is time-dependent	Q: Who goes to the big 12 championship game Gold: Oklahoma Sooners / Prediction: tcu horned frogs
<i>Less plausible</i>	
Conflicting info in Wikipedia	Q: Based on population what is the largest military base in the united states / Gold: Fort Bragg / Prediction: fort hood
Plausible only in certain condition	Q: Who appointed the chief justice of high court Gold: President / Prediction: president of india
Mismatch with question intent	Q: While the united states dealt with the depression it was also facing Gold: Racial tensions / Prediction: great depression, economic downturns
Incorrect	Q: When were banded iron formations formed on the sea floor Gold: Precambrian / Prediction: some 185 billion years ago

Table 5: Examples of predictions that are rated as *definitely correct* or *possibly correct* by human raters.

Easy questions

Q: What city in america has the highest population / Gold: New York City
New York City: New York City is the most populous city in the United States ...

Q: Who plays 11 in stranger things on netflix / Gold: Millie Bobby Brown
Millie Bobby Brown: She gained notability for her role as Eleven in the first season of Netflix science fiction series Stranger Things.

Q: Where did leave it to beaver take place / Gold: Mayfield
Leave It to Beaver: Leave It to Beaver is set in the fictitious community of Mayfield and its environs.

Q: The substance that is dissolved in the solution / Gold: solute
Solution: A solute is a substance dissolved in another substance, known as a solvent.

Q: This means that in dna cytosine is always paired with / Gold: guanine
Guanine: In DNA, guanine is paired with cytosine.

Hard questions

Q: What is the area code for colombo sri lanka / Gold: 036
 Telephone numbers in Sri Lanka: (Answer in the table)

Q: Which country is highest oil producer in the world / Gold: United States / Predictions: Russia
History of the petroleum industry in the United States: For much of the 19th and 20th centuries, the US was the largest oil producing country in the world.
List of countries by oil production: (Answer in the table)
Russia: Russia is the world's leading natural gas exporter and second largest natural gas producer, while also the second largest oil exporter and the third largest oil producer.

Q: Who did gibbs shoot in season 10 finale / Gold: Mendez
 Predictions: fbi agent tobias fornell (from all systems in unrestricted/6Gb track)
Past, Present and Future (NCIS) Gibbs realizes that Mendez is .. The show then returns to the scene depicted in a scene from season 10 finale's ... shoots and kills Mendez.

Q: The girl next door film based on a true story / Gold: loosely / Prediction: girl next door

Table 6: Samples of easy questions (all systems predict correct answers) and hard questions (all systems predict incorrect answers).

Correct	Captured by automatic eval (32%) Input: Who played apollo creed in the original rocky (A: carl weathers) Retrieved: Who played apollo creed in the first rocky movie (A: carl weathers)
	Captured by manual eval (18%) Input: When does the new season of snl air (A: September 29, 2018) Retrieved: When does the new season of snl 2017 start (A: September 30, 2017)
Incorrect	Retrieved Q is not related to input Q in topics (25%) Input: Who sings don't it make you feel like dancing Retrieved: Who sings the song you make me feel like dancing
	Retrieved Q is related but is asking different detail (19%) Input: How much did it cost to make bohemian rhapsody film Retrieved: How much money did the movie rhapsody make
	Retrieved Q is related but contains incorrect 'wh' word (3%) Input: When was the current season of top chef filmed Retrieved: What is the new season of top chef
	Retrieved Q is context-dependent (1%) Input: Who won the war of 1812 between russia and france Retrieved: Who did we fight in the war of 1812
	Retrieved Q has incorrect answer (2%) Input: When was the last time florida state had a losing record Retrieved: When was the last time florida state football had a losing season

Table 7: Breakdown of the predictions from UCLNLP-FB system on 100 random samples.

Correct	Captured by automatic eval (34%) Q: Who played apollo creed in the original rocky / Gold: Carl Weathers / Prediction: Carl Weathers Retrieved P: <i>Carl Weathers</i> : He is best known for portraying apollo creed in the "rocky" series of films.
	Captured by manual eval (29%) Q: what was the shelby in gone in 60 seconds / Gold: Shelby Mustang GT500 / Prediction: gt500 Retrieved P: <i>Eleanor</i> : The eleanor name is reused for a shelby mustang gt500 in the 2000 gone in 60 seconds remake.
Incorrect	No answer found in pruned Wiki (3%) Q: Who has been chosen for the 2017 saraswati samman Gold: Sitanshu Yashaschandra / Prediction Vijayendra saraswathi Retrieved P: <i>jayendra saraswathi</i> : The mutt' s pontiff sri vijayendra saraswathi performed the poojas for his guru and predecessor. Gold P: <i>Sitanshu Yashaschandra</i> : He received Saraswati Samman (2017) for his poetry collection..
	Fail to retrieve the gold passage (12%) Q: Who sings don't it make you feel like dancing / Gold: the Headpins / Prediction: leo sayer Retrieved P: <i>you make me feel like dancing</i> is a song by the british singer leo sayer ..
	Fail to rerank the gold passage (15%) Q: Who says all animals are equal but some are more equal than others Gold: The pigs / Prediction: aristotle Retrieved P: <i>moral status of animals in the ancient world</i> : Aristotle perceived some similarities between humans and other species and developed a sort of "psychological continuum", recognising that human and non-human animals differ only by degree in possessing certain temperaments.
	Fail to extract the correct answer (7%) Q: Which is the smallest continent in size in the world / Gold: Australia / Prediction: Asia Retrieved P: <i>Continent</i> ...Ordered from largest in area to smallest, they are: Asia, Africa, North America, South America, Antarctica, Europe, and Australia

Table 8: Breakdown of the predictions from NAVER RDR on 100 random samples.

Phase	Human Teams	Computer Systems	Points
1	Single player “buzz in”	500MiB	3
2	Team discussion	6GiB	2
3	Team discussion w/ search results	Unlimited	1

Table 9: Phases of the human–computer competition. On odd questions the humans go first, while on even questions the computers go first. If neither team has a correct answer in a phase, we move on to the next phase.

Team	Margin per question
B	0.4 ppq
C	0.1 ppq
D	-0.1 ppq
A	-0.2 ppq
E	-0.6 ppq

Table 10: In our preliminary round of the human–computer competition against the baseline systems, only two teams had positive margins against the computer. We invited the top team to compete in the final competition against the winning competition entrants.