

QAMPARI: A Benchmark for Open-domain Questions with Many Answers

Samuel Joseph Amouyal Tomer Wolfson Ohad Rubin Ori Yoran
Jonathan Herzig Jonathan Berant

Blavatnik School of Computer Science, Tel Aviv University, Israel
{samuel.amouyal, ohad.rubin, joberant}@cs.tau.ac.il

Abstract

Existing benchmarks for open-domain question answering (ODQA) typically focus on questions whose answers are all in a single paragraph. By contrast, many natural questions, such as “*What players were drafted by the Brooklyn Nets?*” have a long list of *answers* extracted from multiple paragraphs. Answering such questions requires retrieving and reading many passages from a large corpus. We introduce QAMPARI, an ODQA benchmark, where answers are lists of entities, spread across many paragraphs. We created QAMPARI by (a) generating questions with multiple answers from Wikipedia’s knowledge graph and tables, (b) automatically pairing answers with supporting evidence in Wikipedia paragraphs, and (c) manually paraphrasing questions and validating each answer. Across a wide range of ODQA models, we find that QAMPARI is challenging in terms of both passage retrieval and answer generation, with models reaching an F_1 score of 32.8 at best. We view QAMPARI as a valuable resource for ODQA research, which will aid to develop models that handle a broad range of question types, including single and multi-answer questions.

1 Introduction

Open-domain question answering (ODQA) is a core language understanding task concerned with answering factoid questions over large document collections (Voorhees and Tice, 2000; Brill et al., 2002). Due to its wide applicability, ODQA has received substantial attention in recent years (Chen et al., 2017; Lee et al., 2019; Karpukhin et al., 2020). Typically, systems tackling ODQA tasks follow the “retrieve-and-read” paradigm, where a *retriever* first retrieves a set of candidate passages, followed by a *reader* which receives the retrieved passages and produces the final answer.

The retrieve-and-read paradigm has been effective for benchmarks such as Natural Questions

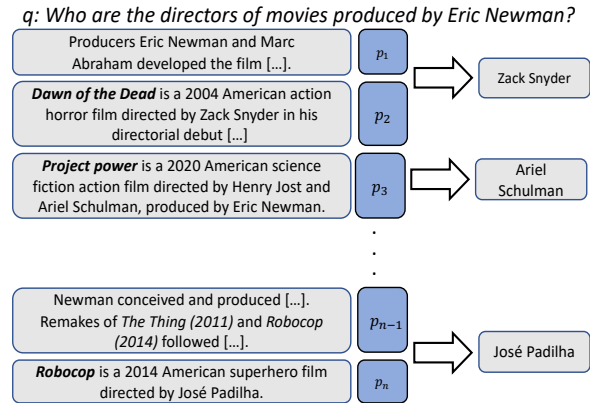


Figure 1: An example from QAMPARI with a generated question q , a subset of its evidence Wikipedia passages (left, p_i) and their corresponding answer.

(NQ) (Kwiatkowski et al., 2019) and TriviaQA (Joshi et al., 2017), where the answer is a single phrase from a single passage. However, in many cases, a question might have *many* answers, spread across multiple passages. Consider the example in Fig. 1. Eric Newman produced multiple movies, so finding them, along with their directors, requires incorporating information from many passages. Such questions pose two main challenges to retrieve-and-read systems. First, as there are multiple answers that can be far apart, the reader must reason over a long text sequence to generate all correct answers. Second, since the reader is computationally constrained to process at most K passages, the retriever must score all necessary passages at its top- K results, which is challenging and even impossible when the number of relevant passages exceeds K .

Nevertheless, research on multi-answer questions has largely been underexplored. While previous works proposed questions that involve reading multiple passages, the number of passages was quite small. AMBIGQA (Min et al., 2020) studied ambiguous questions from NQ with several answers. However, as 70% of its questions have at most two answers, retrieve-and-read models

can be adapted to AMBIGQA. HOTPOTQA (Yang et al., 2018) focused on multi-hop reasoning, but its questions require no more than two passages to answer. WIKINLDB (Thorne et al., 2021) is a benchmark for testing reasoning over multiple facts. However, WIKINLDB restricted its text corpus to databases of 1,000 facts at most, making it significantly smaller than standard ODQA corpora. Moreover, these facts are model-generated utterances rather than natural language passages. Multi-answer questions are also rare in real-world user questions (?Kwiatkowski et al., 2019), which can be attributed to the performance bias of existing systems. Namely, people mostly pose questions that they can successfully get answers to with current technology. This does not diminish the importance of multi-answer questions (‘Which drugs are effective against skin cancer?’; ‘Which plants can be grown in an apartment?’), which constitute an important research challenge.

In this work, we present QAMPARI, a benchmark for **Questions with many Answers over Multiple Paragraphs, Indeed**. All questions in QAMPARI have at least 5 answers, with an average of 13 answers. Examples are semi-automatically generated using two data sources, Wikidata (Vrandečić and Krötzsch, 2014) and Wikipedia tables. We automatically generate multi-answer questions of the form “What/Who has [relation] with [entity]?” and convert these into pseudo-language using manually defined templates. Then, we verify that our questions are answerable from Wikipedia by automatically extracting evidence passages for all their answers. Finally, we use crowdsourcing to validate example correctness, and paraphrase questions into natural language (Wang et al., 2015). To further enrich our data we also generate *composition* questions, that compose two relations (as in Fig. 1), and *intersection* questions, such as “What movies were produced and directed by Clint Eastwood?”. Overall, QAMPARI contains 2K development and test questions and more than 60K training examples – see Tab. 1 for some examples.

We evaluate a large suite of baselines, including models from the retrieve-and-read family as well as a closed-book question answering model (Roberts et al., 2020), and find that they struggle on QAMPARI. In the retrieve-and-read setup, we experiment with both BM25 (Robertson and Zaragoza, 2009) and DPR (Karpukhin et al., 2020) retrievers, followed by either (a) a RAG-like reader (Lewis

et al., 2020) that given each retrieved passage either decodes an answer or abstains, or (b) an FiD reader (Izacard and Grave, 2021) that takes the encoded representations of multiple passages and decodes the list of answers directly.

When training models on QAMPARI alone, or in a multi-task setup with NQ, we observe that QAMPARI is challenging in terms of both passage retrieval and answer generation. Namely, the best model reaches an F_1 score of 32.8. Moreover, models return more than 80% of the correct answers in only 31.2% of the test examples, well below performance on single-answer datasets like NQ.

To summarize, QAMPARI is a challenging benchmark for evaluating the ability of ODQA models to handle questions with many answers over multiple passages. We advocate to evaluate ODQA models not on QAMPARI alone, but alongside benchmarks such as NQ and TriviaQA. Such joint evaluation will test models’ ability to handle both single- and multi-answer questions, an evaluation that the community is currently lacking. The QAMPARI benchmark, models and relevant code-base are available at: <https://samsam3232.github.io/qampari/>.

2 Dataset Construction

Each example in QAMPARI is a triple $(q, \mathcal{A}, \mathcal{P})$, where q is a question, \mathcal{A} is a set of answers and \mathcal{P} is a set of passages from our target corpus. An answer $a \in \mathcal{A}$ has 1-2 evidence passages from \mathcal{P} (see Fig. 1).

We define passages as consecutive sentences from our corpus (Wikipedia), that span on average 100 words. As our focus is multi-answer questions, examples in QAMPARI have $|\mathcal{A}| \geq 5$.

Overview We generate examples in two steps. First, we generate *simple questions* that involve a single entity and relation, e.g., “Who was drafted by the Brooklyn Nets?” (§2.1). Then, we expand such questions to generate *complex questions* with *intersection* and *composition* operations (§2.2).

To increase diversity, questions are generated from two data sources, Wikidata and Wikipedia tables. We first describe example generation over Wikidata, then briefly present the generation process from Wikipedia tables in §2.3. In both cases, we ensure answers can be derived from evidence passages in Wikipedia.¹ Tab. 1 presents examples

¹Wikipedia dump: 2021-08-01

from each data source and question type.

Notation We introduce notation for formal queries over Wikidata to explain example generation. Wikidata is a knowledge graph, \mathcal{K} , that can be viewed as a set of labeled edges (e_1, r, e_2) . Graph nodes $e_1, e_2 \in \mathcal{E}$ are entities connected by an edge labeled by the relation $r \in \mathcal{R}$. For example, a possible labeled edge is (BarackObama, ReceivedAward, NobelPeacePrize).

One can query \mathcal{K} by applying a relation r over an entity e , resulting in a *simple query* $r(e)$ whose *denotation* (answer set) is $\llbracket r(e) \rrbracket = \{e_i \mid (e_i, r, e) \in \mathcal{K}\}$. *Composition queries* are formed by applying a relation over the result of a simple query. We denote a composition query by $r_2(r_1(e))$, and its denotation is $\llbracket r_2(r_1(e)) \rrbracket = \{e_i \mid \exists e_j \text{ s.t. } (e_i, r_2, e_j) \in \mathcal{K} \wedge (e_j, r_1, e) \in \mathcal{K}\}$. Last, an *intersection query* $r_1(e_1) \sqcap r_2(e_2)$ corresponds to the intersection of two simple queries, $\llbracket r_1(e_1) \sqcap r_2(e_2) \rrbracket = \{e_i \mid (e_i, r_1, e_1) \in \mathcal{K} \wedge (e_i, r_2, e_2) \in \mathcal{K}\}$.

2.1 Simple Questions

Fig. 2 provides an overview of our procedure for creating *simple question* examples: (i) We manually define query templates, (ii) populate query templates using \mathcal{K} to create queries with a sufficiently large number of answers in \mathcal{K} , (iii) automatically identify evidence passages for the answers and filter out noisy examples, (iv) map query templates to question templates to obtain pseudo-language questions, and (v) validate answers and paraphrase pseudo-language questions through crowdsourcing. Next, we describe each of these steps in detail.

Generating query templates We manually select a set of 135 relations $\bar{\mathcal{R}} \subset \mathcal{R}$, which will be used in our query templates. We select frequent relations from Wikidata for which denotations contain many entities (e.g., ReceivedAward). The list of relations is in App. A. For each relation, we manually write a template to map queries to pseudo-language questions. For example, the template for ReceivedAward is “*Who received the award X?*”

Some relations are underspecified – for example, LocatedIn can describe the location of buildings, geographical features, and cities. When generating synthetic questions, this leads to vague questions such as “*What is located in Paris?*”. To address this, we manually split these to *typed relations* that specify the semantic type of their an-

swers/denotations. This is done using the type hierarchy given in Wikidata and given the type t of answer entities. We denote typed relations by r_t , and the denotation of $r_t(e)$ comprises all entities of type t returned by $r(e)$. For example, the entity The Louvre has type cultural organization, and we can map the relevant query template to the pseudo-language question “*Which cultural organization is located in Paris?*”.

Simple query generation We instantiate all possible simple queries using all $r \in \bar{\mathcal{R}}$ and entities e in Wikidata. For a relation r (or r_t), we keep the query $r(e)$ iff $|r(e)| \geq 5$. We denote this set of instantiated simple queries by \mathcal{S} , which contains 1,431,268 simple queries.

Finding evidence sentences For an ODQA benchmark, we must verify that every answer is found in our target corpus. We do this by identifying candidate evidence sentences from Wikipedia, and verifying they entail the answer, using a Natural Language Inference (NLI) model.

Specifically, every simple query-answer pair can be viewed as a triple (e_1, r, e_2) . We use a “distant supervision” approach (Mintz et al., 2009), similar to KELM (Agarwal et al., 2021), and define any sentence in the Wikipedia page of entity e_1 that contains the entity e_2 , or one of its Wikidata aliases, as a candidate evidence sentence (and vice versa in the page of e_2). E.g., in Fig. 2, the evidence for (BarackObama, ReceivedAward, NobelPeacePrize) appears on the page Barack Obama, where ‘*Nobel Peace Prize*’ appears.

Aligning Wikipedia sentences to Wikidata can lead to false positives. E.g., for the triple (TheGoonies, HasScreenwriter, StevenSpielberg), most mentions of Spielberg in the page TheGoonies are not as a screenwriter. To account for this, we use an off-the-shelf NLI model.² For every answer, we consider each candidate evidence sentence along with its two preceding sentences, and check whether they entail the hypothesis phrase describing the triple (e_1, r, e_2) . We use templates to phrase triples as short declarative sentences (“*The Goonies has Steven Spielberg as screenwriter*”). An answer is *validated* if there is an evidence sentence that entails the triple. Manual analysis shows this

²huggingface.co/ynie/roberta-large-snli_mnli_fever_anli_R1_R2_R3-nli

Data source	Question type	Question	Answer example
Wikidata	Simple	Who is or was a member of the Australian Army?	George Macarthur-Onslow
	Intersection	What movie produced by Jerry Ward was also directed by Vincent Sherman?	Hard Way
	Composition	From which country did Seattle Storm make draft selections?	Australia
Wiki. tables	Simple	What magazine is a satirical magazine?	The Clinic
	Composition	What are the museums found in Concord, Massachusetts?	The Wayside

Table 1: Example questions and one representative answer for all data sources and question types.




Figure 2: An overview of example generation for simple questions.

process eliminates 70% of false positives, while removing only 7.5% of the correct alignments.

Query filtering After finding evidence sentences, we only keep queries that at least 80% of their answers were validated and their number of validated answers is between 5 and 200. The resulting set contains 60,792 simple queries, where each query has a set of validated answers, \mathcal{A} , and of passages \mathcal{P} that contain the identified evidence sentences.³

2.2 Complex Questions

To increase diversity, we expand simple queries to composition and intersection queries, for which answers require reading two passages.

Intersection Intersection queries are generated by finding two simple queries such that the size of the intersection of their denotations is at least 5. To avoid improbable questions such as “*Which competition was won by Manchester City and had Manchester City as a participant?*”, we add a constraint that the denotation of one of the simple queries cannot be a subset of the other. Formally, the set of intersection queries are all queries $r_1(e_1) \sqcap r_2(e_2)$ such that $|\llbracket r_2(e_2) \sqcap r_1(e_1) \rrbracket| \geq 5$, $\llbracket r_1(e_1) \rrbracket \not\subseteq \llbracket r_2(e_2) \rrbracket$, and $\llbracket r_2(e_2) \rrbracket \not\subseteq \llbracket r_1(e_1) \rrbracket$.

Pseudo-language questions are generated by heuristically combining the two simple questions, for example “*Which television program had Chris Carter as screenwriter and had Frank Spotnitz as*

screenwriter?”. There is no need to perform answer validation since all of the underlying intersecting answers were already validated.

Composition To create composition queries, we manually handpick a set of 423 relations $\mathcal{R}_{\text{comp}} \subset \mathcal{R}$ (list in our codebase), in a process similar to simple queries. Then, we generate all the possible composition queries $r_2(r_1(e))$ such that $r_1(e) \in \mathcal{S}$, $r_2 \in \mathcal{R}_{\text{comp}}$, and $|\llbracket r_2(r_1(e)) \rrbracket| \geq 5$. An example composition query is “*What is the height of buildings located in Dubai?*”.

Unlike intersection queries, in composition queries we need to validate that our new triples (e_i, r_2, e_j) , where $e_j \in \llbracket r_1(e) \rrbracket$, are indeed supported by Wikipedia sentences. We use the same procedure to find evidence sentences for triples (e_i, r_2, e_j) , and consider an answer e_i as *validated* if both (e_i, r_2, e_j) and (e_j, r_1, e) can be aligned to Wikipedia. We keep all complex queries where 80% of the answers are validated. Finally, we manually define templates for relations in $\mathcal{R}_{\text{comp}}$ to generate pseudo-language questions.

2.3 Questions from Wikipedia Tables

To further diversify QAMPARI, we create an analogous pipeline for generating simple and composition questions from Wikipedia tables, with more open-ended relations compared to Wikidata. We briefly describe this pipeline.

We look at all Wikipedia tables with title “*List of X*” that have at least 5 rows, in total, 1,897 tables. We find the “key” column, c_{key} in each table using the table classifier from Talmor et al. (2021),

³We keep a single evidence passage for every triple.

which outputs the column of entities that the table describes. For example, in the table *List of nuclear whistle blowers*, c_{key} is ‘name’ and specifies the whistle-blower names. This naturally creates simple questions of the form “Who or what is X?”.

Simple questions are expanded to composition questions by looking at non-key columns, $c_{\text{non-key}}$ and asking what rows in the table have the value v in column $c_{\text{non-key}}$. For example, what is the value in the column ‘Year’ for nuclear whistle-blowers.

Questions from Wikipedia are validated using a procedure similar to Wikidata. For each answer entity e , we validate that the Wikipedia page for e contains the relevant words that are part of the name of the table as well as the value (for composition questions), and only keep questions where 80% of the table rows are validated and the number of validated answers is at least 5. Overall, we generate 170 simple questions and 6,036 composition questions using this process.

2.4 Data Split

QAMPARI contains a training set, whose goal is to teach the model to handle multi-answer questions. However, we do not want the model to memorize how particular Wikidata relations map to text patterns. Consequently, we perform a *relation split*, randomly splitting the set $\bar{\mathcal{R}}$ into two equally-sized sets $\bar{\mathcal{R}}_{\text{train}}$ and $\bar{\mathcal{R}}_{\text{test}}$. Simple queries are assigned to the train/test set based on their relation, composition queries $r_2(r_1(e))$ are assigned to the test set iff either r_1 or r_2 are in $\bar{\mathcal{R}}_{\text{test}}$, and intersection queries $r_1(e_1) \sqcap r_2(e_2)$ are placed in the test set iff both r_1 and r_2 are in $\bar{\mathcal{R}}_{\text{test}}$.

We now create the train/development/test split (Tab. 2). The main bottleneck in our example generation pipeline is validation of the test set through crowdsourcing (§2.5), since each question requires validating all of the answers. Thus, we pre-determine the test set to contain 1,000 simple questions (830 from Wikidata, 170 from Wikipedia tables) and 1,000 complex questions (400 Wikidata composition questions, 400 Wikidata intersection questions, 200 Wikipedia tables composition questions). For simple Wikidata questions, we sample 830 questions such that the distribution over relations from $\bar{\mathcal{R}}_{\text{test}}$ is roughly uniform. All Wikipedia tables simple questions are placed in the test set, and for complex questions we randomly sample the pre-determined number from the set of generated questions. Last, the test set is randomly split in half to a development set and test set. We also

sub-sample training set examples, such that each relation appears in at most 1,000 examples.

2.5 Crowdsourcing

Correctness validation For every question and answer, we present a crowdsourcing worker with the question, the answer, and links to the Wikipedia page (or pages for complex questions) with the evidence passage. We ask the worker to check if the question can be answered from the given pages, using the text only (no infoboxes or tables).

Since the vast majority of examples are correct, we test worker performance by injecting wrong answers in 10% of the cases and reject workers that fail to identify wrong answers. Moreover, we manually verify 5% of examples marked as *correct* and all examples marked as *incorrect*, and again reject low-performing workers. Overall, 24 annotators validated 30,259 answers for an average pay of 12.5\$ per hour. We find that our process for generating examples is accurate, with 96.6% of the answers validated. Non-validated questions were replaced until 2,000 questions were validated. A question is defined non-validated if its number of distinct answers goes below 5. Snapshots from the presented tasks are in App. C.

Paraphrasing Since our questions are in pseudo-language, we follow past work (Wang et al., 2015) and ask workers to re-phrase 3,000 questions in the training set and the entire development/test set. We restrict this task to US or UK workers who pass a qualification test. We randomly verified half of the paraphrases for each worker for quality assurance.

3 Dataset Analysis

QAMPARI contains 61,911 training examples, 1,000 development examples and 1,000 test examples. Tab. 1 provides example questions of each question type and data sources. We describe key statistics in Tab. 2. Test examples in QAMPARI have 13.23 answers on average and a median of 7 answers. For comparison, the number of answers per question is substantially higher than in AmbigQA (Min et al., 2020), where the median is 2. On average, simple questions have more answers than complex ones while being shorter in length. We note that since test and development questions were manually re-phrased by annotators they are generally shorter than the training questions.

Figure 3a presents a binned distribution of the number of answers per question in the development

		Total	Simp. WD	Simp. WP	Inter. WD	Comp. WD	Comp. WP
# Questions	train	61,911	28,574	-	2,301	25,200	5,836
	dev + test	2,000	830	170	400	400	200
Mean # Answers	train	13.25	16.65	-	9.19	9.74	13.35
	dev + test	13.23	15.69	23.84	8.94	8.77	11.51
Median # Answers	train	8.0	9.0	-	7.0	7.0	8.0
	dev + test	7.0	7.5	17.0	7.0	6.0	7.0
Mean Question len.	train	12.69	8.78	-	16.69	15.18	19.47
	dev + test	9.51	7.91	8.61	11.65	10.35	10.99

Table 2: QAMPARI questions breakdown by their type (**Simple**, **Intersection** or **Composition** questions) and underlying data source (**WD** for Wikidata, **WP** for Wikipedia tables).




Figure 3: Left: Distribution of the number of answers per example. Right: Proportion of questions per number of added answers in *ExtendedSet*.

and test sets. Roughly half of the questions have 8 or more answers, with 20% having more than 15 answers and 3.5% with over 50 answers.

Extended set As discussed in §2.5, we manually validate each answer in QAMPARI is supported by sentences from Wikipedia. However, Wikipedia might contain additional correct answers. To alleviate this issue, we manually annotate additional gold answers for a subset of test questions, and name it the *ExtendedSet*. We randomly sampled 200 questions from the test set and had an author manually annotate as many additional answers as possible in 12 minutes, per question. This process is not guaranteed to be complete, as it would require manually reviewing all of Wikipedia. Moreover, questions with hundreds of gold answers (“*Who worked for Burton F. C?*”) would incur hours of annotation, which is too expensive. This is similar to work in open information extraction (?), where creating the full gold set of triples is not feasible. Fig. 3 plots the number of added answers per question on the extended set. In 30% of the questions, we did not add any answer, and the median/average/maximum number of added answers are 2/3.13/16 respectively. Evaluation on the test set and the extended set in §4.3 shows that model precision on the extended set is somewhat higher, but does not alter model ranking, illustrating the reliability our test set.

4 Experimental Evaluation

4.1 Models

Retriever For retrieval, we experiment with both sparse and dense retrieval models on Wikipedia. As discussed in §2, we chunk Wikipedia into passages of consecutive sentences, using NLTK’s sentence tokenizer, where each passage is 100 words on average. For all retrievers, we evaluate retrieval accuracy of the top-200 passages returned per question.

We use BM25 (Robertson and Zaragoza, 2009) as a strong sparse retrieval model. BM25 scores question-passage pairs based on their lexical similarity. It has been shown that BM25 is notoriously hard to beat using unsupervised retrieval methods (Izacard et al., 2021; Ram et al., 2022), and achieves comparable performance to that of supervised methods (?). As our dense retriever we finetune on QAMPARI a DPR model (Karpukhin et al., 2020) trained on NQ. We finetune DPR in the typical contrastive manner (in-batch training), with one positive and one negative passage per question. Positives are sampled from the evidence passages, and negatives are sampled from the top-10 highest scoring passages, according to BM25, which do not contain the answer.

Reader We experiment with two readers – a Passage-Independent Generator (PIG), which reads each passage independently (a-là RAG (Lewis et al., 2020)), and a Fusion-in-Decoder (FiD) model (Izacard and Grave, 2021), which reads multiple passages simultaneously.

PIG is an encoder-decoder model that takes each of the retrieved passages as input and decodes a single answer or outputs “*Not Relevant*” to indicate there is no answer. The final output is the union of all decoded answers across retrieved passages. We initialize PIG with T5-large (Raffel et al., 2019) and train with standard maximum likelihood. We use evidence passages as positive examples and the top

scoring retrieved passage that is not an evidence passage and does not contain an answer (or its aliases) as a negative example.

FiD encodes each of the retrieved passages along with the input question. Its decoder then attends to the encoded representation and outputs a list of answers. We initialize FiD using a pretrained T5-Large model (Raffel et al., 2019) and train with standard maximum likelihood.

FiD is computationally expensive, as its decoder attends to a large number of encoded tokens and the generated output is long. Thus, we can only fit the top-50 passages returned by the retriever on a single A100 GPU.

Closed-book question answering We also experiment with a closed-book setting, where the QA model generates answers from knowledge encoded in its parameters without any evidence passages. We initialize our closed-book QA model with T5-SSM with 3B parameters (Roberts et al., 2020), and train it with standard maximum likelihood – the question is provided as input, and the model is trained to generate the gold set of answers.

Zero-shot We test the zero-shot ability of Open AI’s *text-davinci-003*, from the Instruct-GPT family (?). We use GPT-3 in: (a) closed-book QA setup; (b) as a multi-passages reader. In the closed-book setup, the model receives only the question and is asked to provide a list of answers. In the reader setup, the model gets the question and the 15 highest-ranking passages from BM25 (the maximal number that fits in the context) and is asked to output a list of answers.

4.2 Experimental Setup

We created QAMPARI as a benchmark to be evaluated alongside additional ODQA benchmarks, such as NQ. Since it is semi-automatically generated, one can develop models tailored for QAMPARI. However, our goal is to have a single model that performs well across a wide variety of question types. Thus, we train and test models in a multi-task setup, on both NQ and QAMPARI, in addition to a QAMPARI only setting. We also train our models on NQ only and evaluate them on QAMPARI, to verify QAMPARI’s training set indeed improves answering questions with many answers.

Our main metrics are recall, precision, and F_1 . Specifically, for test example $(q, \mathcal{P}, \mathcal{A})$, and a predicted set of answers $\mathcal{A}_{\text{pred}}$, recall, precision, and F_1 are standardly computed by comparing \mathcal{A} and

	ARECALL@K		ERECALL@K	
	BM25	DPR	BM25	DPR
K=10	24.6	21.9	11.1	11.1
K=25	37.4	31.5	28.4	16.2
K=50	46.6	39.6	38.7	20.8
K=100	54.6	47.1	47.6	25.5
K=200	61.0	55.2	55.6	30.2

Table 3: Retriever test results.

$\mathcal{A}_{\text{pred}}$, allowing for aliases (i.e., a gold answer is covered if it or one of its aliases are in $\mathcal{A}_{\text{pred}}$). The model scores are averaged across examples. To get a sense of the average accuracy across examples, we measure the fraction of examples with F_1 of at least 0.5 ($\%F_1 \geq 0.5$) and the fraction with recall of at least 0.8 ($\%\text{Recall} \geq 0.8$). For NQ, we report the standard exact match (EM) metric.

We evaluate the retriever with RECALL@K, that is, the fraction of answers that appear in the top-K retrieved passages, averaged across examples. This metric comes in two flavors: (a) Answer RECALL@K (ARECALL@K): for every gold answer whether it or one of its aliases appear in the top-K retrieved passages. It is a loose metric since an answer can appear in a passage that does not provide any evidence to support the answer; (b) Evidence RECALL@K (ERECALL@K): since we have evidence paragraphs for every answer, we consider for every gold answer the fraction of evidence passages in the top-K retrieved passages. This is a strict metric since an answer can sometimes be answered by passages other than the ones we identified.

4.3 Results

Tab. 3 presents passage retrieval results on QAMPARI test. Scores for ARECALL@200 for BM25 and DPR are 61.0% and 55.2%, respectively. As for ERECALL@K, results are unsurprisingly lower. BM25 retrieves 55.6% of the evidence passages with K=200, while DPR retrieves only 30.2% of evidence passages.⁴ Overall, DPR pretrained on NQ and finetuned on QAMPARI performs worse than BM25. This is in line with ? who showed that, when tested on questions with *rare entities*, DPR performs worse than BM25. We hypothesize that rare entities in QAMPARI questions may account for DPR’s lower performance.

Tab. 4 lists results on the test sets of QAMPARI and NQ. Overall, performance on QAMPARI is

⁴While ERECALL@K for DPR is substantially lower than BM25, observe that ARECALL@K is better correlated with QA metrics (Tab. 4), as DPR retrieves non-evidence passages that still lead to the correct answer.

		Rec.	Prec.	F ₁	%Rec _{≥.8}	%F ₁ ≥.5
FiD-BM25	QO	25.1	36.8	28.3	6.8	24.2
	MT	26.9	37.7	29.7	7.4	25.6
FiD-DPR	QO	7.8	39.1	12.5	0	3.6
	MT	7.8	41.3	12.5	0	2.6
PIG-BM25	NQO	34.6	19.3	20.8	18.5	11.9
	QO	43.1	30.7	31.0	26.7	26
	MT	47.9	28.2	30.5	31.2	22.3
PIG-DPR	NQO	9.0	13.7	8.4	0.5	2.6
	QO	36.2	41.1	32.8	15.7	30.7
	MT	34.1	44.8	32.4	15	31.3
Closed book Reader	ZS	12.9	17.4	13.8	1.9	9.5
	ZS	20.0	22.8	18.8	5.8	13.8
Closed book	QO	1.7	7.3	2.6	0	0.3

Table 4: QAMPARI test results. **QO**: models trained on QAMPARI only; **NQO**: models trained on NQ only; **MT**: Multi-task training with NQ; **ZS**: Zero-shot setup.

low. FiD-DPR and PIG-DPR are more precision-oriented with FiD-DPR achieving precision of 41.3 and PIG-DPR a precision of 44.8. PIG-BM25 is recall-oriented, achieving recall of 47.9. Overall, PIG variants perform best, with small differences between PIG-BM25 and PIG-DPR, and both are slightly higher than FiD-BM25.

When training on both NQ and QAMPARI (MT), performance on NQ (47.2 with BM25 and 53.1 with DPR) is similar to that reported by [Izacard and Grave \(2021\)](#) (44.1 with BM25 and 51.4 with DPR). When training on NQ only, results on QAMPARI are significantly lower than when training also on QAMPARI, showing that training on QAMPARI improves performance on multi-answer questions, as expected. The lower performance on QAMPARI compared to NQ, despite the fact that NQ’s EM evaluation metric is much more strict than the metrics used for QAMPARI, illustrates the challenge in answering multi-answer questions.

PIG-DPR has much higher recall than FiD-DPR, showing that going over 200 passages independently (PIG) leads to higher recall than jointly reasoning over 50 passages (FiD). Moreover, the solid performance of PIG-DPR indicates that QA performance is more correlated with ARecall@K than ERecall@K (Tab. 3).

Finetuned closed-book performance is low with an F₁ of 2.6 for QAMPARI, which we attribute to the relation-based train/test split (§2.4). This guarantees that there is no overlap between train and test questions. ? have shown that mitigating such train-test overlap causes a drop in QA performance, with a drastic drop being observed in closed-book models.

Zero-shot results The performance of zero-shot models is lower than finetuned retrieve-and-read models, as expected. However, *text-davinci-003*’s performance in the closed book setup is impressive and significantly better than finetuned T5-3B.

ExtendedSet results We report results for FiD and PIG on the *ExtendedSet* (see §3) in §F. As expected, considering additional correct answers improves the precision of all models. Since changes to recall are small, the overall F₁ is higher when considering manual annotations. Importantly, ranking across models does not change, and the absolute performance remains low, suggesting that our test set can be safely used for evaluation.

Oracle analysis To disentangle retrieval from answer extraction, we run PIG and FiD in an oracle setup, where we assume a perfect retriever and run our readers on the gold evidence passages only. Performance of both models greatly improves in this setup, with larger gains for PIG. This shows that developing better retrieval mechanisms for multi-answer questions can greatly benefit QAMPARI. FiD’s recall is still limited (47.5), illustrating the challenge of reading a large number of documents. Full oracle results are in §G (Tab. §9).

5 Related work

ODQA tasks have largely been dedicated to single-answer questions ([Berant et al., 2013](#); [Joshi et al., 2017](#); [Kwiatkowski et al., 2019](#)). The same applies for most multi-hop ODQA tasks ([Welbl et al., 2018](#); [Yang et al., 2018](#); ?). While they require 2-4 paragraphs, the answer is a single phrase. Multi-answer questions were introduced in the TREC QA tracks (??). However, evaluation was on 50 questions. ? introduced artificially generated multi-answer questions, but only for reading comprehension rather than ODQA. Concurrent to QAMPARI, ? introduced RoMQA, a benchmark containing multi-answer questions generated using Wikidata. While their setup is closest to ours, they evaluate on a subset of Wikipedia that is aligned to a subset of Wikidata.

Retrieve-and-read models are the prevailing approach in ODQA ([Chen et al., 2017](#); [Yang et al., 2019](#); [Lee et al., 2019](#); [Sachan et al., 2021](#)). Closed-book QA is an alternative approach, ([Roberts et al., 2020](#); [Tay et al., 2022](#)) but requires using high-capacity models.

6 Conclusions

We release QAMPARI, a dataset targeting ODQA models ability to answer multi-answer questions, and show that it is challenging for current state-of-the-art models. QAMPARI will aid develop models that answer a wide range of question types, including single- and multi-answer questions.

Limitations

A key limitation of QAMPARI is that the gold set of answers is incomplete. Thus, predicted answers might be correct but missing from the gold answer set. The *ExtendedSet* addresses this problem partially, allowing a more accurate model ranking, but even in this set all the correct answers are not part of the gold set. A second limitation is that our data generation process is mostly automatic and is thus amenable to reverse-engineering. Hence, we recommend evaluating models on QAMPARI along with additional benchmarks created with a different generation process. Last, our data generation process can only generate answers based on relations from Wikidata and relations that are in Wikipedia tables, and thus its scope does not generalize to arbitrary relations.

References

- Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2021. [Knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3554–3565, Online. Association for Computational Linguistics.
- Petr Baudiš and Jan Šedivý. 2015. [Modeling of the question answering task in the YodaQA system](#). In *Proceedings of the 6th International Conference on Experimental IR Meets Multilinguality, Multimodality, and Interaction - Volume 9283, CLEF'15*, page 222–228, Berlin, Heidelberg. Springer-Verlag.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Eric Brill, Susan Dumais, and Michele Banko. 2002. [An analysis of the AskMSR question-answering system](#). In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 257–264. Association for Computational Linguistics.
- 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.
- Jifan Chen and Greg Durrett. 2019. [Understanding dataset design choices for multi-hop reasoning](#). 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 4026–4032, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig, Ruslan Salakhutdinov, and William W. Cohen. 2020. [Differentiable reasoning over a virtual knowledge base](#). *CoRR*, abs/2002.10640.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. [Towards unsupervised dense information retrieval with contrastive learning](#). *arXiv preprint arXiv:2112.09118*.
- 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.
- 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.
- 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.

- 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, 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*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063.
- 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.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *CoRR*, abs/1910.10683.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Ori Ram, Gal Shachaf, Omer Levy, Jonathan Berant, and Amir Globerson. 2022. [Learning to retrieve passages without supervision](#). In *North American Association for Computational Linguistics (NAACL)*.
- 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.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Devendra Singh Sachan, Siva Reddy, William L. Hamilton, Chris Dyer, and Dani Yogatama. 2021. [End-to-end training of multi-document reader and retriever for open-domain question answering](#). In *Advances in Neural Information Processing Systems*.
- Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. [Simple entity-centric questions challenge dense retrievers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Haitian Sun, Pat Verga, Bhuwan Dhingra, Ruslan Salakhutdinov, and William W. Cohen. 2021. [Reasoning over virtual knowledge bases with open predicate relations](#). *CoRR*, abs/2102.07043.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hannaneh Hajishirzi, and Jonathan Berant. 2021. [Multimodal{qa}: complex question answering over text, tables and images](#). In *International Conference on Learning Representations*.
- Yi Tay, Vinh Q Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. 2022. [Transformer memory as a differentiable search index](#). *arXiv preprint arXiv:2202.06991*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- James Thorne, Majid Yazdani, Marzieh Saeidi, Fabrizio Silvestri, Sebastian Riedel, and Alon Halevy. 2021. [Database reasoning over text](#). 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 3091–3104, Online. Association for Computational Linguistics.
- Ellen M. Voorhees and Dawn M. Tice. 2000. [The TREC-8 question answering track](#). In *Proceedings of the Second International Conference on Language Resources and Evaluation (LREC'00)*, Athens, Greece. European Language Resources Association (ELRA).
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Communications of the ACM*, 57(10):78–85.

Yushi Wang, Jonathan Berant, and Percy Liang. 2015. Building a semantic parser overnight. In *Association for Computational Linguistics (ACL)*.

Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.

Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.

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. *arXiv preprint arXiv:1902.01718*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.

A Simple Relations

In Tab. 5. we gathered all the 135 relations we used to create our simple questions. The 423 relations used to create our composition questions can be found in our code base.

B Composition template

Composition questions overall template is: **What is the <comp_property> of <subtype> who/which <base_property>?**. All the templates are in our code base.

C Crowdsourcing Validation

Fig. 4 shows two screenshots of the task crowdsourcing workers performed.

D Experimental setup details

For both readers (FiD and PIG), we used T5-large which has 770 million parameters. We used an A100 to train both of them, FiD with a batch size of 8 and PIG with a batch size of 32 for a single GPU. We trained each of them for around 48 hours on two GPUs.

For FiD, we concatenated the answers using # as a separator. At evaluation time, there is no importance to the order of the answers.

For both PIG and FiD, all aliases of a given gold entity provided by Wikidata are used as additional correct answers. When verifying whether our model predicted an answer A, we verify whether it predicted A or any of its aliases. We performed an hyper parameter search around the learning rate, the number of training steps, the ratio of positive to negative (for PIG) and the number of times an NQ example will appear in each epoch (for multi task). Tab. 6 presents the parameters of the reported results.

We report the results of a single run with seed 0.

E Question type analysis

We break test performance of FiD-BM25 (MT) by question type (Tab. 7). Surprisingly, performance on simple questions is lower than complex questions, and intersection questions seem easiest. Possible explanations are: (a) simple questions have more answers (see Tab. 2), which makes them harder, and (b) models can predict the answer given just one evidence passage, due to “shortcuts” (Chen and Durrett, 2019), or parametric knowledge (Longpre et al., 2021).

F ExtendedSet Results

In Tab. 8 we present results analogous to those in Tab. 4 for the *ExtendedSet* with BM25. Precision improves by 5-6 points across models, while recall changes are smaller leading to an overall increase in F_1 . Nevertheless changes are not dramatic and model ranking remains constant, suggesting the full test set can be safely used.

G Development Set Results

In Tab. 9 we present results analogous to those in Tab. 4 for the development set.

is a	has author	located in	language	occupation
sex or gender	country of citizenship	part of	place of birth	located in
educated at	language spoken, written or signed	has part	played the sport	employer
genre	position held	cast member	country of origin	award received
place of death	made from material	creator	has participant	depicts
maintained by	operator	performer	member of political party	owned by
religion	headquarter location	participant	member of	position played
original language	competition class	publisher	role	record label
work location	director	doctoral advisor	residence	native language
place of publication	medical condition	winner	field of work	form or work
conflict	place of burial	instrument	composer	league
screenwriter	distribution format	producer	sponsor	ethnicity
voice actor	distributed by	participating team	academic degree	manufacturer
architectural style	fabrication method	present in work	production company	cause of death
military branch	manner of death	industry	director of photography	narrative location
original broadcaster	organizer	student of	location of creation	located in or next to body of water
architect	archives at	nominated for	country of registry	allegiance
movement	voice actor	noble title	based on	dedicated to
legislated by	location of formation	developer	contributor to creative work or subject	lyrics written by
located in protected area	tracklist	editor	presenter	religious order
from narrative universe	location of discovery	media franchise	commissioned by	political ideology
commemorates	port of registry	influenced by	indigenous to	operating area
translator	brand	interested in	designed by	illustrator
vessel class	costume designer	drafted by	coach of sports team	convicted of
scenographer	culture	significant place	executive producer	represented by
broadcast by	investor	cover art by	home port	collection creator
armament	inspired by	first appearance	choreographer	animator
source of energy	musical conductor	adapted by	sound designer	has written for
academic major	ratified by	business model	worshipped by	narrator
partnership with	colorist	art director	has work in the collection	military rank

Table 5: Simple relations

		Learning rate	# steps	pos. to neg.	# NQ examples
FiD-BM25	QO	0.00005	90k	-	-
	MT	0.00005	190k	-	2
FiD-DPR	QO	0.00005	85k	-	-
	MT	0.00005	190k	-	2
PIG-BM25	QO	0.000001	60k	1	-
	MT	0.000001	75k	1	1
PIG-DPR	QO	0.000001	60k	1	-
	MT	0.000001	75k	1	1
Closed book	QO	0.0001	95k	-	-

Table 6: Hyper parameters used for reported results.

	Recall	Precision	F ₁
Wikidata simple	21.3	30.7	23.1
Wikidata intersection	37.0	47.1	40.0
Wikidata composition	18.6	32.4	22.2
Wikipedia simple	9.1	20.6	11.5
Wikipedia composition	31.2	37.4	32.7

Table 7: Question type analysis of FiD-BM25, trained in MT setup on QAMPARI development set.

			QAMPARI				
			Recall	Precision	F ₁	% Recall \geq 0.8	% F ₁ \geq 0.5
FiD-BM25	QO	w.o. annotations	20.5	34.6	24.3	4.0	19.6
		w. annotations	23.3	40.6	27.8	4.5	25.1
FiD-BM25	MT	w.o. annotations	22.8	37.0	26.8	4.5	20.6
		w. annotations	25.7	42.9	30.6	5.0	24.6
PIG-BM25	QO	w.o. annotations	45.1	28.9	30.7	27.5	23
		w. annotations	42.7	33.6	32.8	24	29.5
PIG-BM25	MT	w.o. annotations	49.3	27.9	30.7	31.5	20.5
		w. annotations	47.1	33.1	33.2	27	26

Table 8: QAMPARI *ExtendedSet* results with (w.) and without (w.o.) the additional manual annotations. The best results with and without annotations are bolded. **QO**: models trained on QAMPARI only; **MT**: Multi-task training with NQ.

		QAMPARI				
		Recall	Precision	F ₁	% Recall \geq 0.8	% F ₁ \geq 0.5
FiD-BM25	QO	23.3	35.6	26.3	5.9	22.7
	MT	23.9	34.2	26.3	6.0	22.4
FiD-DPR	QO	6.5	35.2	10.1	0	3.7
	MT	7.2	39.8	11.4	0.0	2.8
PIG-BM25	QO	41.4	26.4	28.0	25.3	21.0
	MT	43.7	26.9	28.9	26.6	22.0
PIG-DPR	QO	33.9	38.6	29.9	15.8	26.2
	MT	31.7	42.2	29.6	14.3	26.3
Closed book	QO	2.4	7.2	3.1	0.1	0.7
FiD-Oracle	MT	47.5	62.7	51.2	18.4	56.1
PIG-Oracle	MT	71.5	60.9	62.4	55.7	73.8

Table 9: QAMPARI development results. **QO**: models trained on QAMPARI only; **MT**: Multi-task training with NQ.

Question answer validation - Instructions (Click to collapse)

- **Mission explanation:**
- Our goal, with your help, is to teach the AI to answer questions.
- In this HIT, you will be presented with a machine-generated question, and a candidate answer. You will be provided with a Wikipedia link and we ask you to verify if indeed the answer is correct according to this source.
- An answer is correct if it can be found within the Wikipedia text. **Wikipedia text does not include the info box (the rectangle on the top right corner of the page, the crossed rectangle on the image below) or tables.**

George Washington

From Wikipedia, the free encyclopedia

"General Washington" redirects here. For other uses, see *General Washington* (*Washington*) and *George Washington* (*Washington*).

George Washington (February 22, 1732^[a] – December 14, 1799) was an American soldier, statesman, and Founding Father who served as the first president of the United States from 1789 to 1797. Appointed by the Continental Congress as commander of the Continental Army, Washington led the Patriot forces to victory in the American Revolutionary War, and presided at the Constitutional Convention of 1787, which established the Constitution of the United States and a federal government. Washington has been called the "Father of the Nation" for his manifold leadership in the formative days of the country.^[b]

Washington's first public office was serving as official Surveyor of Culpeper County, Virginia from 1748 to 1756. Subsequently, he received his initial military training (as well as a commission with the Virginia Regiment during the French and Indian War). He was later elected by the Virginia House of Burgesses and was named a delegate to the Continental Congress. Here he was appointed Commanding General of the Continental Army. With this title, he commanded American forces (allied with France) in the defeat and surrender of the British at the Siege of Yorktown during the American Revolutionary War. He resigned his commission after the Treaty of Paris was signed in 1763.

Washington played an indispensable role in adopting and ratifying the Constitution of the United States. He was then twice elected president by the Electoral College unanimously. As president, he implemented a strong, well-financed national government while remaining impartial in a fierce rivalry between cabinet members Thomas Jefferson and Alexander Hamilton. During the French Revolution, he proclaimed a policy of neutrality while endorsing the Jay Treaty. He set enduring precedents for the office of president, including the title "Mr. President", and his Farewell Address is widely regarded as a pre-eminent statement on republicanism.

Washington was a slaveowner who had a complicated relationship with slavery. Washington controlled a total of over 577 slaves at one time or another, who were forced to work on his farm and in his houses, including the White House. As president, he agreed laws passed by Congress that both prohibited and upheld slavery. He set aside that one of his slaves, William Lee, should be freed upon his death, and that the other 529 slaves must work for his wife and be freed on her death. She freed them during her lifetime to remove the incentive to hasten her death.^{[c][d]}

He endeavored to assimilate Native Americans into the Anglo-American culture but fought indigenous resistance during instances of tribal conflict. He was a member of the Anglican Church and the Freemasons, and he urged broad religious freedom in his roles as general and president. Upon his death, he was lauded by Henry "Light Horse Harry" Lee as "first in war, first in peace, and first in the hearts of his countrymen".^[e]

Washington has been commemorated by monuments, a federal holiday, national parks, geographical locations, including the national capital, the State of Washington, stamps, and currency, and many nicknames and polls rank him among the greatest U.S. presidents. In 1876 Washington was



- To verify the answer you do not need to read the entire Wikipedia page, you can simply search for all places where the answer or the entity in the question is mentioned.
 - If you cannot tell if the answer is correct, or if a fact contradicting the answer can be found, select "Cannot determine/Wrong answer".
 - The great majority of examples are correct, but it is important to catch the wrong ones.
 - If you cannot understand the question, select "I cannot understand the question".
- Note that questions are sometimes ungrammatical or phrased unnaturally, and this is fine. Only select this option if the question can not be understood.

- **Mission plan:**
- This is a long term task with thousands of future HITs.
- Please read carefully the examples since there is relevant information in the explanations.
- Qualification is given if a total of 4 HITs (40 questions) were completed in total.
- Score for qualification must be above 90% of correct answer during the qualification round.
- Answers will be randomly verified for quality assurance.

(a) Instructions

Examples (click to collapse / expand)

Correct Correct Cannot determine/Wrong Cannot determine/Wrong

Generated question:
Who was a student of George Crumb?

Generated answer:
John Bethune Melby

Wikipedia link:
John Bethune Melby

Expected annotation:
Answer is CORRECT

Explanation: The wikipedia page of the link is the Wikipedia page of the answer, therefore you need to look for the entity in the question (George Crumb). In the Wikipedia page, you can see: "He studied with Henry Weinberg, George Crumb, Peter Westergaard, J. K. Randall, and Milton Babbitt". That proves that the generated answer is correct.

Question 1:

Generated question	What films were produced by CBS Films?
Generated answer	Five Feet Apart
Wikipedia link	Five Feet Apart

Is answer correct?

Answer is CORRECT
 Cannot determine/Wrong answer
 I cannot understand the question

Optional. Tell us if you see a problem with this example, or with the instructions.

Question 2:

Generated question	Which landforms are located in the protected area Desolation Wilderness?
Generated answer	Lost Lake
Wikipedia link	Lost Lake (California)

(b) Task

Figure 4: Screenshots from crowdsourcing task.