

Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets

Patrick Lewis^{†‡}, Pontus Stenetorp[‡], Sebastian Riedel^{†‡}

[†]Facebook AI Research; [‡]University College London
{plewis, sriedel}@fb.com; p.stenetorp@cs.ucl.ac.uk

Abstract

Ideally Open-Domain Question Answering models should exhibit a number of competencies, ranging from simply memorizing questions seen at training time, to answering novel question formulations with answers seen during training, to generalizing to completely novel questions with novel answers. However, single aggregated test set scores do not show the full picture of what capabilities models truly have. In this work, we perform a detailed study of the test sets of three popular open-domain benchmark datasets with respect to these competencies. We find that 30% of test-set questions have a near-duplicate paraphrase in their corresponding train sets. In addition, we find that 60–70% of answers in the test sets are also present in the train sets. Using these findings, we evaluate a variety of popular open-domain models to obtain greater insight into what extent they can generalize, and what drives their overall performance. We find that all models perform substantially worse on questions that cannot be memorized from train sets, with a mean absolute performance difference of 61% between repeated and non-repeated data. Finally we show that simple nearest-neighbor models outperform a BART closed-book QA model, further highlighting the role that train set memorization plays in these benchmarks.

1 Introduction

Open-domain Question Answering (ODQA) is a task that examines the ability of models to produce answers to natural language factoid questions drawn from an open set of domains. ODQA has received significant attention for its potential practical applications, and more recently as a popular method to analyse how well NLP systems can capture and recall factual knowledge. This interest in ODQA as a challenging “knowledge-intensive” task has led to a flurry of recent works that have

driven test-set performance on standard ODQA datasets to new heights (Lee et al., 2019; Guu et al., 2020; Karpukhin et al., 2020; Lewis et al., 2020; Izacard and Grave, 2020, inter alia). Whilst there have been several works examining other kinds of QA datasets (Manjunatha et al., 2018; Kaushik and Lipton, 2018; Sugawara et al., 2018, 2020), however, we know comparatively little about how the questions and answers are distributed in ODQA benchmarks, making it hard to understand and contextualize the results we are observing.

In this work, we address these issues via an analysis of the test sets of three popular ODQA datasets, namely WebQuestions (Berant et al., 2013), TriviaQA (Joshi et al., 2017) and Open NaturalQuestions (Kwiatkowski et al., 2019; Lee et al., 2019). We identify three types of desired behaviour of a trained ODQA system, in increasing order of difficulty: 1) to recall the answer to a question that the model has seen at training time. 2) to answer novel questions at test time and choose an answer from the set of answers it has seen during training. 3) to answer novel questions which have answers which are not contained in the training data. It is not clear to what extent our current ODQA datasets measure each of these three behaviours.

To address this, we stratify the test sets of these datasets. Firstly, we split the test data by whether answers in the test set also appear somewhere in the train set. We find that 58–71% of test answers also occur somewhere in the training data, demonstrating that the majority of the test data does not probe for answer generalization. Secondly, we annotate 1,000 question-answer pairs from each test set for near-duplicate questions in their respective train sets. We find that a surprisingly high 28–34% have paraphrased questions in the training data, the vast majority of which are near-duplicates differing by one or two words. This result implies that 30% of the test set of these datasets only probe for how

Dataset	% Answer overlap	% Question overlap
NaturalQuestions	63.6	32.5
TriviaQA	71.7	33.6
WebQuestions	57.9	27.5

Table 1: Test-train set overlap for the datasets.

well models can memorize question-answer pairs seen at training.

Equipped with these observations, we compute the performance of several recently proposed ODQA models on our test subsets. We test both *Open-book* approaches, which leverage retrieval from a large corpus of documents and *Closed-book* approaches, which focus on training large parametric seq2seq models with no external knowledge source (Roberts et al., 2020). We find that test data with train-overlapping data contribute the bulk of the overall performance of all the models studied.

These issues seem to be more acute for closed-book models. Strikingly, we find that a closed-book BART-based model (Lewis et al., 2019) is incapable of producing answers not observed at training time, and achieves very low scores on non-overlapping questions, suggesting this model is only capable of memorizing question-answer pairs from training time. With this in mind, we build simple nearest-neighbor models which outperform this BART model, despite having virtually no capacity to generalize beyond training data.

In summary, we make the following contributions: 1) We provide insights into how answer entities are distributed between dataset splits for ODQA datasets 2) We provide annotated subsets of ODQA test sets indicating whether test-time questions are duplicates of training time questions.¹ 3) We evaluate a variety of models on our dataset splits, and derive insights into what kinds of question answering behaviour different models achieve.

2 Datasets

We analyse three widely used ODQA datasets, WebQuestions (Berant et al., 2013), TriviaQA (Joshi et al., 2017), and Open-NaturalQuestions (Lee et al., 2019). All three datasets consist of factual natural language questions and short multi-token answers, differing slightly in the style of questions and format of answers.

¹Our data and evaluation code is available at <https://github.com/facebookresearch/QA-Overlap>

WebQuestions consists of 3,778 train and 2,032 test instances. Questions were obtained by mining a search engine, and answers are Freebase entities (Bollacker et al., 2008) annotated by crowdworkers. The ODQA task consists of predicting the name of the Freebase entity. We use the standard train/test splits from Berant et al. (2013) and the development split from Karpukhin et al. (2020), which was randomly split from the train set.

TriviaQA consists of 78,785 train, 8,837 development and 11,313 test instances obtained by scraping trivia websites. Answers are Wikipedia entities, and any alias for the answer entity is considered a correct answer. We use the ODQA splits, which correspond to the unfiltered-train and unfiltered-dev reading comprehension splits (Lee et al., 2019; Min et al., 2019, 2020b; Karpukhin et al., 2020).

Open-NaturalQuestions consists of search engine questions with crowdsourced answer spans in Wikipedia articles. The ODQA version consists of question-answer pairs from NaturalQuestions which have short answer spans less than 6 tokens in length. We use the standard open-domain splits in our experiments, consisting of 79,168 train, 8,757 development and 3,610 question answer pairs.

For all three datasets, the canonical train, development and test splits were obtained by randomly splitting the question-answer pairs, and there are no exact duplicate questions in any dataset. We exclude development data from our overlap analyses, and focus purely on train-test overlap to explicitly assess the effects of training memorization.

3 Test-Train Overlaps

We explore two ways of examining the test sets based on overlaps between training and test data. Consider a question-answer pair (q, a) from the test set D_{test} where the answer consists of at least one answer reference $a = \{s_1..s_n\}$. We can consider *answer overlap* where there exists at least one $(q', a') \in D_{\text{train}}$ which shares at least one answer reference with (q, a) . We can also consider *question overlap*, where there exists some $(q'', a'') \in D_{\text{train}}$ where q'' is a duplicate of q , such that q and q'' are paraphrases and have the same answer.

Answer Overlap Following Rajpurkar et al. (2016), we apply answer normalization (lower-casing, stripping punctuation, removing articles and normalizing whitespace) on answer references

Open NaturalQuestions		TriviaQA		WebQuestions	
Overlapping	Non-overlapping	Overlapping	Non-overlapping	Overlapping	Non-overlapping
Phil Simms	Cloves	David Bowie	Death in the afternoon	Harvard	Queen Victoria
Brian Johnson	Matt Monro	Battle of camlann	Clash of the Titans	Alderaan	Brasilia
8	1,020 – 1,080 kg	Heligoland	ice-cream sundae	India	Paddington
the Indians	Hermann Ebbinghaus	Henry VII	Camshaft	2011	Tom Corbett
the 1830s	Matt Flinders	Niagra Falls	Cumberland	Zeus	Gary

Table 2: Randomly sampled overlapping and non-overlapping answers from all three test sets.

Answer	Test Question	Train Question
Jason Marsden	who plays max voice in a goofy movie	who does max voice in a goofy movie
January 23 2018	when will the 2018 oscar nominations be announced	when are the oscar nominations for 2018 announced
Alan Shearer	who has scored more goals in the premier league	most goals scored by a premier league player
retina	where are the cones in the eye located	where are cone cells located in the eye
francisco pizarro	who led the conquest of the incas in south america	conquistador who defeated the incan empire in peru

Table 3: Randomly sampled test-train overlapping questions in Open NaturalQuestions. See Appendix A.2 for more examples, including examples from TriviaQA and WebQuestions

before searching for overlapping answer references for all (q, a) pairs in the test set – see Table 1. We find that 57.9% of test (q, a) pairs in WebQuestions have answer overlaps, with 63.6% and 71.7% for NaturalQuestions and TriviaQA respectively. We would naturally expect TriviaQA to have higher answer overlap as it has more answer references per question on average (13.7 references on average compared to 1.2 for NaturalQuestions and 2.4 for WebQuestions). Examples of answer overlaps are shown in Table 2.

Question-Overlap Unlike answer overlap, question overlap cannot be easily computed automatically, as searching for duplicates via rules or paraphrase classifiers may lead to both false positives and negatives. Thus, we turn to manual annotation to investigate question overlap. To obtain a representative sample for each dataset, we annotate a random subset of 1,000 (q, a) pairs for each test set. Annotators are shown a list of up to 50 training questions which have a similar answer reference.² This answer similarity function is designed for high recall to obtain a tight lower bound on question overlap. If there were no questions with similar answers in the train set, the question was automatically annotated as not overlapping. Three expert annotators annotated the remaining questions and indicated if any were paraphrases of the

²Training questions are selected for annotation if one of the following is true: they share an answer reference with a test question, a test answer reference is a sub-sequence of a training answer reference, or the other way around (a training reference answer is a sub-sequence of a test answer reference). If there are more than 50 such questions, the top 50 are chosen by the highest degree of word overlap with the test question.

test question and had the same answer.

The results from the annotation can be seen in Table 1 and examples of overlapping questions in Table 3. A sample of 100 2-way annotated examples indicated 93% agreement, corresponding to a Cohen’s Kappa of 0.85 (Cohen, 1960). What we observe is a high degree of question overlap, with between 27.5% and 33.6% of the 1,000 annotated test questions having a duplicate in the train set. It is also common to see several duplicates per test question, with an average of 2.8 duplicate questions per overlapping test question in NaturalQuestions.

4 Implications for Modelling

Earlier we identified three classes of answering behaviours: 1) questions that can be memorized at training time, 2) novel questions that can be answered with answers memorized at training time, 3) novel questions with novel answers. We refer to these behaviours as *Question memorization*, *Answer classification* and *QA generalization*.

Question memorization To perform well on the question-overlap subset, a model only needs to memorize (q, a) pairs at training time, then recognize which training question matches a test-time question. The reasoning required ranges from trivial duplicate detection for very similar questions such as “who played pink in pink floyd the wall” and “who played pink in the movie the wall”, to more challenging inference problems for subtler duplicates like “On which island in the North Sea did both St Aidan and St Cuthbert live?” and “irish born missionary st aidan founded a monastery in

Model		Open NaturalQuestions				TriviaQA				WebQuestions			
		Total	Question Overlap	Answer Overlap Only	No Overlap	Total	Question Overlap	Answer Overlap Only	No Overlap	Total	Question Overlap	Answer Overlap Only	No Overlap
Open book	RAG	44.5	70.7	34.9	24.8	56.8	83.0	54.7	29.2	45.5	82.6	45.8	21.1
	DPR	41.3	69.4	34.6	19.3	57.9	78.0	59.6	31.6	42.4	76.1	39.8	22.2
	FID	51.4	71.3	48.3	34.5	67.6	86.2	66.9	42.8	-	-	-	-
Closed book	T5-11B+SSM	35.2	76.5	21.0	8.4	-	-	-	-	42.8	82.1	44.5	22.0
	BART	26.5	67.6	10.2	0.8	26.7	67.3	16.3	0.8	27.4	76.1	20.7	1.6
Nearest Neighbor	Dense	26.7	69.4	7.0	0.0	28.9	81.8	11.2	0.0	26.4	81.9	17.1	0.0
	TF-IDF	22.2	56.8	4.1	0.0	23.5	69.8	5.1	0.0	19.4	68.1	8.7	0.0

Table 4: Exact Match scores on our dataset splits. “Total” is the overall performance on the dataset. “Question Overlap” is the subset with train-test question overlap, and probes for simple question memorization. “Answer Overlap Only” is subset without train-test question overlap, but with train-test answer overlap, which probes for answer classification. “No overlap” is the subset with no train-test answer overlap and probes for QA generalization

653 on which english island which is also the name of a 1970s uk folk-rock band?”. Manual annotation of 100 question-overlap pairs showed 81% were simple duplicates differing by one or two words, 14% needed some paraphrase recognition ability, and 5% needed more sophisticated language understanding. To measure performance on question memorization, we build a test subset of (q, a) pairs with question overlap to the train set.

Answer Classification For a model to handle answer-overlap questions, a classifier over train set answers would be sufficient, as answers never appear at test time that do not appear at training time. We build a test subset of (q, a) pairs which have answer overlap, but not question overlap. Question-overlap pairs are excluded as they are significantly easier and would inflate scores.

QA Generalization Here, models cannot rely on memorizing their training data. To measure performance on this most challenging split, we build a test subset of (q, a) pairs which do not have answer overlap with the train set. We also note that we expect higher frequency answers, such as countries, integers and public figures will naturally appear less often in this subset. As such, models that perform well on the head of the answer distribution may struggle to perform well here, despite being able to perform some generalization at test time.

Next we briefly describe the models included in our analysis. For published models, we obtain test set predictions directly from the authors.

4.1 Open-Book Models

Open-book Models first retrieve relevant documents from Wikipedia and then either ex-

tract or generate answers conditioned on those documents. We consider Dense Passage Retriever (DPR) (Karpukhin et al., 2020), a pipeline model which retrieves documents using dense embeddings, before feeding them into a conventional reader-reranker which extracts spans of text as answers. We also include Retrieval-Augmented Generation (Lewis et al., 2020), a model that jointly learns to retrieve and generate answers in a seq2seq framework. Finally we include Fusion-in-Decoder (FID) (Izacard and Grave, 2020), a pipeline model which retrieves 100 documents and fuses them so that the decoder can attend to all documents at once. We do not include FID results on WebQuestions as the authors did not use it in their original work.

4.2 Closed-Book Models

Closed-book models store the knowledge required to answer their questions entirely within the model’s parameters, rather than in an external corpus. Typically these models consist of seq2seq transformers directly fine-tuned on (q, a) pairs. In our analysis, we train a BART-large closed-book model, which is trained with questions as input and generates (q, a) pairs as output. Checkpoints are selected by Exact Match score on a development set. We also include a more powerful T5-11B model from (Roberts et al., 2020). We use the T5-11B model pretrained with a special “Salient Span Masking” objective (Guu et al., 2020), which improves downstream ODQA performance. We do not include TriviaQA results for T5 since their model used a different data splitting scheme.

4.3 Nearest-Neighbor Models

Given the high levels of train-test overlaps in these datasets, we also experiment with some simple nearest-neighbor models. Here, we simply retrieve a (q, a) pair from the train set based on question similarity to the test question, and return its answer. We experiment with two models, one using TF-IDF and the other using maximum inner product of question embeddings from the DPR retriever. These models cannot generalize to non-overlapping answers, and have limited capacity to answer non-overlapping questions. However, they are attractive from the perspective of model size and efficiency (Min et al., 2020a).

4.4 Results

Question Memorization Earlier, we found that $\sim 30\%$ of test set questions overlap with the train set. The “Question overlap” columns in Table 4 shows performance on Question Memorization. Comparing this column with the total performance column shows that all models perform significantly higher on memorizable questions. This effect is most pronounced for closed book models. The T5-11B performs especially well for question memorization on both NaturalQuestions and WebQuestions. This suggests that its very large capacity, coupled with more powerful question understanding may allow it to store and recall training questions more effectively than other models.

Answer Classification The “Answer overlap only” column in Table 4 shows performance on answer classification. Answer classification has a large drop in performance compared to question memorization, dropping by an average of 45% Exact Match. Open-book models handle this setting better than closed book models. The BART model in particular struggles here, scoring only 10.2%.

QA Generalization The “No overlap” column in Table 4 shows performance on QA generalization. All models significantly lose performance on QA generalization, highlighting the shortcomings of the total performance metric. For example, we may expect the FID state-of-the-art model to answer half of NaturalQuestions-style questions correctly, but once we account for repeated questions and answers, it can only answer about a third of questions correctly. This difference is even more pronounced for other models, with an average absolute drop of 25% with respect to total performance.

Nearest-Neighbor Models The bottom two rows of Table 4 show results for nearest-neighbor models. TF-IDF, despite being completely untrained, can answer about 20% of test questions correctly, purely by retrieving questions from the train sets. The dense retrieval model outperforms the BART closed-book model on NaturalQuestions and TriviaQA. Further, the dense nearest neighbor model also outperforms the significantly more complex DPR open-book model on TriviaQA and WebQuestions on the question overlap subset.

5 Related Work

Examining what behaviours are learnt by models has received attention in language understanding tasks, such as GLUE (Wang et al., 2018), which includes tools for probing for different reasoning types. There has also been critical and careful analysis of QA systems and datasets. Chen et al. (2016), Sugawara et al. (2020) and Kaushik and Lipton (2018) analyse the difficulty of various machine reading datasets, and Manjunatha et al. (2018) show that visual QA models memorize common question-answer relationships in training data. Févry et al. (2020) analyse various closed-book models’ TriviaQA predictions. Kwiatkowski et al. (2019) note that the machine reading NaturalQuestions dataset has train-test overlap of Wikipedia titles, and provide baselines for “long-answer” QA. Verga et al. (2020) observe answer overlap effects in a related modality (knowledgebase QA), but do not consider question overlap.

6 Conclusion

We performed an analysis of popular open-domain QA datasets. We found that 60% of test set answers overlap with the train set and 30% of test set questions have at least one duplicate in the train set. Given these observations, we contextualize performance of seven ODQA models, stratifying by the extent of overlap, exploring how well these models generalize versus simply memorizing training data. It is clear that performance on these datasets cannot be properly understood by overall QA accuracy. In the future, a greater emphasis should be placed on more behaviour-driven evaluation, rather than pursuing single-number overall accuracy.

References

- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic Parsing on Freebase from Question-Answer Pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, SIGMOD '08, pages 1247–1250, Vancouver, Canada. Association for Computing Machinery.
- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. [A Thorough Examination of the CNN/Daily Mail Reading Comprehension Task](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2358–2367, Berlin, Germany. Association for Computational Linguistics.
- Jacob Cohen. 1960. [A Coefficient of Agreement for Nominal Scales](#). *Educational and Psychological Measurement*, 20(1):37–46. Publisher: SAGE Publications Inc.
- Thibault Févry, Livio Baldini Soares, Nicholas FitzGerald, Eunsol Choi, and Tom Kwiatkowski. 2020. [Entities as Experts: Sparse Memory Access with Entity Supervision](#). *arXiv:2004.07202 [cs]*. ArXiv: 2004.07202.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: Retrieval-Augmented Language Model Pre-Training](#). *arXiv:2002.08909 [cs]*. ArXiv: 2002.08909.
- Gautier Izacard and Edouard Grave. 2020. [Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering](#). *arXiv:2007.01282 [cs]*. ArXiv: 2007.01282.
- 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 Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense Passage Retrieval for Open-Domain Question Answering](#). *arXiv:2004.04906 [cs]*. ArXiv: 2004.04906.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How Much Reading Does Reading Comprehension Require? A Critical Investigation of Popular Benchmarks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, Brussels, Belgium. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural Questions: a Benchmark for Question Answering Research](#). *Transactions of the Association of Computational Linguistics*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent Retrieval for Weakly Supervised Open Domain Question Answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension](#). *arXiv:1910.13461 [cs, stat]*. ArXiv: 1910.13461.
- Patrick Lewis, Ethan Perez, Aleksandara 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](#). *arXiv:2005.11401 [cs]*. ArXiv: 2005.11401.
- Varun Manjunatha, Nirat Saini, and Larry S. Davis. 2018. [Explicit Bias Discovery in Visual Question Answering Models](#). *arXiv:1811.07789 [cs]*. ArXiv: 1811.07789.
- Sewon Min, Jordan Boyd-Graber, Chris Alberti, Danqi Chen, Eunsol Choi, Michael Collins, Kelvin Guu, Hannaneh Hajishirzi, Kenton Lee, Jennimaria Palomaki, Colin Raffel, Adam Roberts, Tom Kwiatkowski, Patrick Lewis, Yuxiang Wu, Heinrich Küttler, Linqing Liu, Pasquale Minervini, Pontus Stenetorp, Sebastian Riedel, Sohee Yang, Minjoon Seo, Gautier Izacard, Fabio Petroni, Lucas Hosseini, Nicola De Cao, Edouard Grave, Ikuya Yamada, Sonse Shimaoka, Masatoshi Suzuki, Shumpei Miyawaki, Shun Sato, Ryo Takahashi, Jun Suzuki, Martin Fajcik, Martin Docekal, Karel Ondrej, Pavel Smrz, Hao Cheng, Yelong Shen, Xiaodong Liu, Pengcheng He, Weizhu Chen, Jianfeng Gao, Barlas Oğuz, Xilun Chen, Vladimir Karpukhin, Stan Peshterliev, Dmytro Okhonko, Michael Schlichtkrull, Sonal Gupta, Yashar Mehdad, and Wen-tau Yih. 2020a. [NeurIPS 2020 EfficientQA Competition: Systems, Analyses and Lessons Learned](#). *arXiv:2101.00133 [cs]*. ArXiv: 2101.00133.

- Sewon Min, Danqi Chen, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2019. [A Discrete Hard EM Approach for Weakly Supervised Question Answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2851–2864, Hong Kong, China. Association for Computational Linguistics.
- Sewon Min, Danqi Chen, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2020b. [Knowledge Guided Text Retrieval and Reading for Open Domain Question Answering](#). *arXiv:1911.03868 [cs]*. ArXiv: 1911.03868.
- 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.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How Much Knowledge Can You Pack Into the Parameters of a Language Model?](#) *arXiv:2002.08910 [cs, stat]*. ArXiv: 2002.08910.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What Makes Reading Comprehension Questions Easier?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219, Brussels, Belgium. Association for Computational Linguistics.
- Saku Sugawara, Pontus Stenetorp, Kentaro Inui, and Akiko Aizawa. 2020. [Assessing the Benchmarking Capacity of Machine Reading Comprehension Datasets](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8918–8927. AAAI Press.
- Pat Verga, Haitian Sun, Livio Baldini Soares, and William W. Cohen. 2020. [Facts as Experts: Adaptable and Interpretable Neural Memory over Symbolic Knowledge](#). *arXiv:2007.00849 [cs]*. ArXiv: 2007.00849.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

A Appendices for Question and Answer Test-Train Overlap in Open-Domain Question Answering Datasets

A.1 Test subset sizes

Table 5 gives details of how many instances are in each test set, as well as the number of instances in each test subset used in the main paper.

A.2 Additional Question Overlap Examples

Tables 6, 7 and 8 give more question overlap examples for the three datasets.

Dataset	Total	Answer Overlap	N Annotated	Question Overlap	Answer Overlap Only
NaturalQuestions	3610	2297	996	324	315
TriviaQA	11313	8112	1001	336	411
WebQuestions	2032	1176	998	274	299

Table 5: Number of instances in each test set subset

Answer	Test Question	Train Question
Bob Geldof	who played pink in pink floyd the wall	who played pink in the movie the wall
Daren Maxwell Kaga-soff	who played ricky in secret life of the american teenager	who played ricky on the secret life of the american teenager
Andy	who does april end up with on parks and rec	who does april marry in parks and rec
may 5 2017	when did gaurdians of the galaxy 2 come out	when is guardians of the galaxy vol 2 released
norman pritchard	who won the first medal in olympics for india	who won the first individual olympic medal for india
moira kelly	who does the voice of nala in the lion king	who played nala in the lion king movie
supreme court	who enforces the charter of rights and freedoms	who has final authority of interpretation of the canadian charter of rights and freedoms
554	most passing yards by nfl qb in a game	what is the nfl record for most passing yards in a single game
John Ross	who ran the fastest 40 yard dash in the nfl	who has the fastest 40 yard dash ever
international border ib	what is the name of india pakistan border	what is the border name between india and pakistan
Andrew Wright	who wrote when a man loves a woman	who wrote song when a man loves a woman
new england patriots	who has participated in the most super bowls	what nfl team has been to most super bowls

Table 6: Additional examples of test-train overlapping questions in Open NaturalQuestions

Answer	Test Question	Train Question
Picasso	Who painted "Boy With a Pipe" which, in May 2004, was sold for a record price of \$104 million?	Painted in 1905, the painting Garçon à la pipe was a famous painting by which famous artist who died in 1973?
Wensum	On what river is the city of Norwich	The English city of Norwich lies on which river?
Mantle	Comprising around two-thirds of the Earth's mass, what is found between the core of the Earth and its crust?	What do we call the layer of the Earth between its crust and its core?
Live and Let Die	In which James Bond film does actress Jane Seymour play Solitaire?	Jane Seymour played the character "Solitaire" in which Bond film?
Esau	Who, in the Bible, was the eldest son of Isaac?	In the Bible, who was the first born of Isaac?
Alanis Morissette	Who made the 1995 album 'Jagged Little Pill' which sold 33 million copies?	Who released the 1995 hit album "Jagged Little Pill"?
Excalibur	In British legend, what is the name of King Arthur's sword?	What was the name of King Arthur's sword?
Humidity	What is measured by a Hygrometer?	What does a hygrometer measure?
A Storm	On the Beaufort scale what is defined as force 11?	What is force 11 (eleven) on the Beaufort scale?
Jeremy Irons	Actress Sinead Cusack is married to which 'Oscar' winning actor?	Which actor is the husband of Sinead Cusack?
Sir Cloudesley Shovell	Who was the British Admiral who died in 1707 when four of his ships were wrecked in the Scilly Isles?	In 1707 a fleet of Navy ships was wrecked off the Scilly Islands. Who was the commander who lost his life in the disaster?
Tony Hart	Which famous individual created the 'Blue Peter' sailing ship logo?	Which artist designed the logo for UK television children's show 'Blue Peter'?

Table 7: Examples of test-train overlapping questions in TriviaQA

Answer	Test Question	Train Question
costa rica	where is isthmus of panama located on the map?	where is isthmus of panama located?
1986 world series	when's the last time the mets won the world series?	when did the mets win the pennant?
abbottabad	where was bin laden found and killed?	what country was osama bin laden killed in?
believer	what other movies has ryan gosling been in?	what movies does ryan gosling star in?
sculpture	what type of art did leonardo da vinci make?	what kind of art did leonardo da vinci produce?
origin of species	what book did charles darwin wrote in 1859?	what was the name of the book that charles darwin wrote?
morehouse college	what college did martin luther king jr go to?	where did dr. martin luther king jr. go to school?
communist state	what type of government did soviet union have?	what type of government does the former soviet union have?
turkish lira	what money to take to turkey?	what currency to take to side turkey?
spanish language	what is the most common language spoken in argentina?	what is language in argentina?
opera OR classical music	what music period did beethoven live in?	what music did beethoven composed?
harry s truman	who was president after franklin d. roosevelt?	who became president when roosevelt died in office?

Table 8: Examples of test-train overlapping questions in WebQuestions