Can Small and Synthetic Benchmarks Drive Modeling Innovation? A Retrospective Study of Question Answering Modeling Approaches

Nelson F. Liu **Tony Lee** **Robin Jia** **Percy Liang**
Computer Science Department, Stanford University
{nfliu, tonyhlee, robinjia, pliang}@cs.stanford.edu

Abstract

Datasets are not only resources for training accurate, deployable systems, but are also *benchmarks* for developing new *modeling approaches*. While large, natural datasets are necessary for training accurate systems, are they necessary for driving modeling innovation? For example, while the popular SQuAD question answering benchmark has driven the development of new modeling approaches, could synthetic or smaller benchmarks have led to similar innovations?

This counterfactual question is impossible to answer, but we can study a necessary condition: the ability for a benchmark to recapitulate findings made on SQuAD. We conduct a retrospective study of 20 SQuAD modeling approaches, investigating how well 32 existing and synthesized benchmarks *concur* with SQuAD—i.e., do they rank the approaches similarly? We carefully construct small, targeted synthetic benchmarks that do not resemble natural language, yet have high concurrence with SQuAD, demonstrating that naturalness and size are not necessary for reflecting historical modeling improvements on SQuAD. Our results raise the intriguing possibility that small and carefully designed synthetic benchmarks may be useful for driving the development of new modeling approaches.

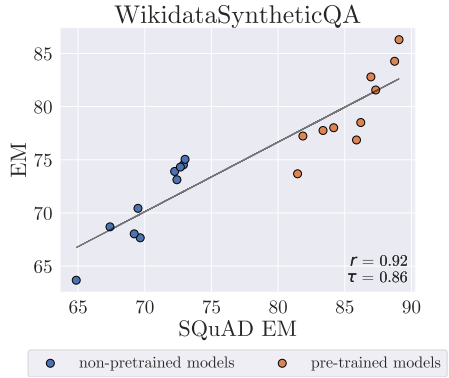

Passage Snippet: *LiveScript is the same as JScript . Asynchronous JavaScript and XML namesake Mocha . JavaScript belonged to Sun Microsystems . .xml user Office Open XML . ...*
Question: *XXXXX namesakes software architecture*
Answer: *Mocha*

Figure 1: We construct WikidataSyntheticQA, a small, synthetic question answering benchmark derived from Wikidata triples that does not resemble natural language, yet has high concurrence with SQuAD on 20 modeling approaches.

1 Introduction

NLP datasets such as SQuAD (Rajpurkar et al., 2016) and SNLI (Bowman et al., 2015) are simultaneously useful as *resources* for training deployable, accurate systems and *benchmarks* for comparing and developing modeling approaches. To support these goals, NLP datasets are most commonly human-constructed—they contain natural language examples that are written by humans (e.g., via crowdsourcing).

The success and popularity of large human-constructed datasets has led to two commonly-held intuitions about dataset construction:

1. *Datasets that more closely resemble natural language are superior to those that to do not* (Manning, 2006; Kwiatkowski et al., 2019; de Vries et al., 2020).

2. *Larger datasets are superior to smaller datasets* (Marcus et al., 1993; Bowman et al., 2015; Rajpurkar et al., 2016; Bajgar et al., 2017).

These intuitions about naturalness[1] and size hold when the goal is to use the dataset as a resource to produce accurate systems for real users.

Preprint version for arXiv; the authors welcome feedback. Contact nfliu@cs.stanford.edu for updated drafts.

[1] We broadly define *natural* text data as data that resembles natural language.

However, popular datasets like SQuAD or SNLI are not intended to fit the needs of any particular user population; systems solely trained on these datasets alone are not intended for deployment.[2] Instead, these datasets are more frequently used as *benchmarks* that facilitate the development of better *modeling approaches*. The hope is that the community gleans broadly-applicable insights from the process of improving benchmark performance and comparing various modeling approaches.

While large, natural datasets are necessary for building accurate systems, they may not be necessary for driving modeling innovation. For example, while SQuAD has driven the development of new modeling approaches, could synthetic or smaller benchmarks have led to similar innovations? Such benchmarks have the potential to help researchers iterate faster on ideas by reducing benchmark construction costs, speeding up experiments, and improving our ability to diagnose and understand model failures.

This counterfactual question is impossible to answer, since it requires predicting how alternative benchmarks would have influenced the ideas and intuitions of NLP researchers over many years. Furthermore, whether a benchmark leads to particular modeling innovations is heavily dependent on extrinsic factors that cannot be systematically controlled, such as the other benchmarks and ideas that existed at the time. However, we can focus on the intrinsic factors and study whether a necessary (but not sufficient) condition holds: can small, synthetic benchmarks recapitulate findings on SQuAD?

We examine this question by conducting a retrospective study of previously-proposed SQuAD modeling approaches. In particular, we ask whether modeling approaches originally developed for and evaluated on SQuAD are ranked similarly by 32 existing and synthesized benchmarks.

We say that two benchmarks have high *concurrence* if they rank a set of modeling approaches similarly, i.e., if approaches that yield performance improvements on one benchmark also produce performance improvements on the other. To measure concurrence between two benchmarks, we first evaluate 20 different modeling approaches on each benchmark—all evaluation is in-domain, using each benchmark's original *i.i.d.* train-test split. Then, we compare the performance trends of our modeling approaches on one benchmark against the performance trends on the other benchmark.

Although such a retrospective study only considers intrinsic factors, it is nonetheless a useful starting step towards raising and better understanding the counterfactual question (see §6.1 for further discussion of qualifications of this study).

We start by exploring the landscape of existing extractive question answering benchmarks (§3), observing first that many existing human-constructed benchmarks have high concurrence with SQuAD (§3.1). We also find that many cloze question-answering benchmarks that use human-written passages (e.g., the Children's Book Test and LAMBADA) have high concurrence with SQuAD, indicating that benchmarks without natural questions can still offer challenges with relevance to natural language data (§3.2). Lastly, we see that existing synthetic benchmarks (e.g., the bAbI tasks) have low concurrence with SQuAD, confirming popular intuition (§3.3).

Given that cloze benchmarks can have high concurrence with SQuAD, how artificial can a benchmark be while still maintaining high concurrence with SQuAD? To study this question (§4.1), we first construct FuzzySyntheticQA, a cloze-format synthetic benchmark that focuses on fuzzy pattern-matching between question and passage tokens. FuzzySyntheticQA has high concurrence with SQuAD on *non-pretrained* modeling approaches, despite its lack of resemblance to natural language (§4.1). This demonstrates that concurrence with non-pretrained models can be achieved with surprisingly simple data, since non-pretrained models primarily learn surface-level patterns from their training data. However, this benchmark has low concurrence with SQuAD on *pre-trained* models. We hypothesize that FuzzySyntheticQA has low concurrence with SQuAD on English-pretrained models because its relative lack of language structure makes the usual benefits of pre-training irrelevant.

Is it possible to construct a non-natural language benchmark where language pre-training helps? We construct WikidataSyntheticQA, a richer synthetic benchmark derived from Wikidata triples. This benchmark also does not resemble natural language, but has simple language-like sentential structure (§4.2). Nonetheless, WikidataSyntheticQA has high concurrence with SQuAD, establish-

[2] Though datasets unfit for producing systems on their own can still have utility for building accurate systems when used in conjunction with other more realistic resources (e.g., via data augmentation).

ing that synthetic benchmarks *can* reflect historical progress on human-constructed benchmarks.

Furthermore, we show that training on only 20K SQuAD examples is sufficient for high concurrence with the original SQuAD benchmark (§5).

In summary, our results demonstrate that naturalness and size are not necessary for building a benchmark that recapitulates findings made on SQuAD—benchmarks with varying naturalness and size can offer challenges with relevance to natural language, and large amounts of natural data do not necessarily make benchmarks worthwhile research targets. Although our results are inconclusive about whether synthetic benchmarks counterfactually could have driven modeling innovation, they raise the intriguing possibility—we encourage the community to revisit and carefully consider the merits and limitations of synthetic benchmarks with still a critical eye.

2 Measuring Concurrence Between Benchmarks

We say that two benchmarks have high *concurrence* when they rank a set of modeling approaches similarly. We measure the performance of a modeling approach when trained and tested on one benchmark with its performance when trained and tested on another benchmark—we use each benchmark's original *i.i.d.* train-test split, so all evaluation is in-domain. Repeating this process for many modeling approaches, we can assess whether performance gains *between* modeling approaches are generally preserved when moving from one benchmark to another.

Formally, define a benchmark B as a pair of datasets $(D_{\text{train}}, D_{\text{test}})$, where $D_{\text{train}} \subseteq \mathcal{X} \times \mathcal{Y}$ and $D_{\text{test}} \subseteq \mathcal{X} \times \mathcal{Y}$ for an input space $\mathcal{X}$ and an output space $\mathcal{Y}$. A *system* is a function $s : \mathcal{X} \to \mathcal{Y}$ (i.e., a function that takes a member of the input space and returns a member of the output space). A *modeling approach* is a function a that takes in a training dataset D_{train} and outputs a system s. Then, let EVAL denote an evaluation function, where $\text{EVAL}(a, B)$ returns the performance (under the given evaluation function EVAL) of a modeling approach a when trained and tested on benchmark B. Finally, let $\text{CONCUR}(B_1, B_2; M, \text{EVAL})$ be the *concurrence* between the benchmarks B_1 and B_2 with respect to a set of modeling approaches M and the evaluation function EVAL. Let $a \sim \text{uniform}(\mathcal{A})$, where $\text{uniform}(\mathcal{A})$ denotes the uniform distribution over the set of modeling approaches $\mathcal{A}$. Define random variables $P_1 = \text{EVAL}(a, B_1)$ and $P_2 = \text{EVAL}(a, B_2)$. Then, we define

$$\text{CONCUR}(B_1, B_2; \mathcal{A}, \text{EVAL}) = \text{CORR}(P_1, P_2)$$

where CORR is some correlation function.

We use the SQuAD exact match (EM) metric as our evaluation function EVAL of choice, and we consider the Pearson correlation coefficient (r) and the Kendall rank correlation coefficient (τ) as our correlation functions CORR. The former measures whether the relationship between model performance on the two benchmarks is approximately linear, whereas the latter measures whether pairwise rank comparisons between models are preserved between benchmarks. As a rough guideline, benchmarks with high concurrence often have $\tau > 0.80$, though interpreting concurrence often requires more than comparing the overall correlation.

To assess concurrence in this work, we use a representative set of 20 diverse modeling approaches for SQuAD introduced between 2016 to 2020 ($\mathcal{A}$). These modeling approaches include RaSoR (Lee et al., 2016), BiDAF (Seo et al., 2017), DocumentReader (Chen et al., 2017), QANet (Yu et al., 2018), BiDAF++ (Clark and Gardner, 2018), MnemonicReader (Hu et al., 2017), FusionNet (Huang et al., 2018), BERT (Devlin et al., 2019), ALBERT (Lan et al., 2020), RoBERTa (Liu et al., 2019), ELECTRA (Clark et al., 2020), and SpanBERT (Joshi et al., 2020). See Appendix A for a full list of the modeling approaches used to calculate concurrence, as well as implementation details.

Non-pretrained Modeling Approaches. 10 of our 20 modeling approaches are non-pretrained and introduced during the first two years after SQuAD's release (mid-2016 to late-2018).

These non-pretrained modeling approaches predominately use recurrent neural networks (most commonly LSTMs; Hochreiter and Schmidhuber, 1997) with static word representations (e.g., GloVe; Pennington et al., 2014) to produce continuous vector representations of passage and question tokens. Then, attention mechanisms combine the question and passage representations, and a softmax classifier predicts the start and end positions of the answer span within the passage.

During these years, improvements in the SQuAD state of the art came from the design of better end-

to-end neural network architectures. Researchers mostly focused their efforts on (1) designing better sequence encoders for passages and questions (Lee et al., 2016; Yang et al., 2017; Yu et al., 2018, *inter alia*) and (2) proposing improved attention mechanisms for question-passage interactions (Wang and Jiang, 2017; Seo et al., 2017; Wang et al., 2017; Huang et al., 2018, *inter alia*).

Pre-trained Modeling Approaches. 10 of our 20 modeling approaches are pre-trained and introduced after late-2018, when the introduction of BERT (Devlin et al., 2019) shifted the community focus from designing elaborate task-specific neural network architectures (embodied by our non-pretrained SQuAD approaches) to designing better pre-training procedures and objectives (often keeping the neural network architecture fixed). For example, BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and SpanBERT (Joshi et al., 2020) each use the same neural network architecture, but adopt different pre-training processes and objectives.

This shift also resulted in fewer researchers explicitly focusing on SQuAD itself—while non-pretrained SQuAD modeling approaches often exclusively evaluated on SQuAD, pre-trained modeling approaches are often evaluated on many benchmarks to highlight their generality.

All pre-trained modeling approaches largely follow the same procedure for fine-tuning on extractive question answering. The tokenized passage and question are concatenated into a single packed sequence and fed as input to the pre-trained model, which produces vector representations for each token. Then, a softmax classifier predicts the start and end positions of the answer span within the passage.

Evaluating Modeling Approaches. We evaluate each modeling approach on each benchmark with the same training hyperparameters used for SQuAD, as well as 5 additional randomly sampled hyperparameter settings.

3 How Well Do Existing Benchmarks Concur With SQuAD?

We begin by studying how well existing benchmarks concur with SQuAD. In particular, we consider three broad categories of extractive question answering benchmarks: human-constructed benchmarks, cloze-format benchmarks, and synthetic benchmarks, in order of decreasing naturalness. See Appendix B for details about benchmark preprocessing and Appendix D for full results for all models on all benchmarks.[3]

3.1 Many Existing Human-Constructed Benchmarks Have High Concurrence With SQuAD

We start by investigating how well human-constructed benchmarks concur with SQuAD.

Setup. Human-constructed extractive question answering benchmarks contain natural language questions and passages that are written by humans. Most extractive question answering benchmarks are human-constructed, and such benchmarks have driven much progress. We study how 5 human-constructed benchmarks concur with SQuAD—NewsQA (Trischler et al., 2017), NaturalQuestions (Kwiatkowski et al., 2019), DROP (Dua et al., 2019), HotpotQA (Yang et al., 2018), and QAMR (Michael et al., 2018).

In particular, we use the versions of NewsQA, NaturalQuestions, DROP, and HotpotQA from the MRQA 2019 shared task (Fisch et al., 2019), which adapted several existing benchmarks to conform to a unified extractive format. QAMR questions were originally collected at the sentence level, but we concatenate these sentences to reconstruct the original passages they were sourced from. See Appendix C.1 for examples from the human-constructed benchmarks we study.

Results and Discussion. We find that many human-constructed benchmarks have high concurrence with SQuAD (Figure 2). This suggests that the specific nuances that go into building human-constructed benchmarks may not drastically affect concurrence—natural language collected from humans may be largely sufficient for high concurrence with SQuAD.

It is worth noting that MRQA DROP has significantly lower r and τ compared to the other benchmarks—the results of non-pretrained and pre-trained approaches are best fit by two separate linear models (creating a "knee"). We hypothesize that this occurs because the benchmark was constructed with a model in the loop (SQuAD-trained BiDAF); crowdworkers were required to

[3] See nelsonliu.me/papers/can-small-and-synthetic-benchmarks-drive-innovation/vis for interactive visualizations of full results.

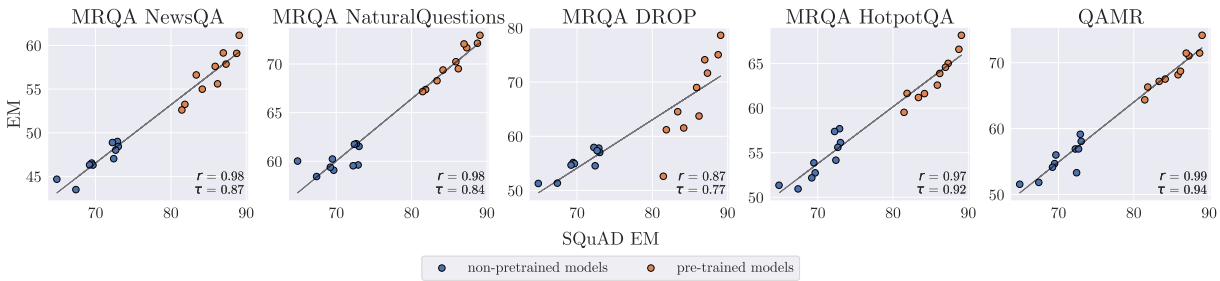

Figure 2: Many human-constructed benchmarks have high concurrence with SQuAD, suggesting that the specific nuances that go into building human-constructed benchmarks may not drastically affect concurrence—natural language collected from humans may be largely sufficient for high concurrence with SQuAD. MRQA DROP has lower overall concurrence, but we see higher concurrence within the non-pretrained and pre-trained subgroups.

write questions that this adversarial baseline could not answer.[4] Although τ and r are lower overall, we see high concurrence within the non-pretrained and pre-trained subgroups.

3.2 Many Existing Cloze Benchmarks Have High Concurrence With SQuAD

Cloze extractive question answering benchmarks contain cloze questions, which are "fill-in-the-blank" statements where the answer is masked. The cloze format enables collecting massive amounts of data at very low cost, since examples can be constructed by eliding spans from naturally-occurring text. Although their passages are often natural language, their questions are not natural (e.g., they do not begin with wh question words).

At the time of its release, SQuAD was touted for its naturalness, especially when compared to the largest previously-introduced question answering benchmarks, which were automatically-generated and used the semi-synthetic cloze format (Hermann et al., 2015; Hill et al., 2016; Onishi et al., 2016). SQuAD questions were crowdsourced—a core tenet behind its construction was importance of naturalness, which was lacking from the cloze-format benchmarks of the time. Despite their lack of naturalness, collecting cloze-format benchmarks is much cheaper than crowdsourcing—do cloze benchmarks recapitulate findings made on SQuAD?

Setup. We study the Children's Book Test (CBT; Hill et al., 2016), LAMBADA (Paperno et al., 2016), CNN (Hermann et al., 2015), and ReCoRD (Zhang et al., 2018) benchmarks.

In particular, we follow prior work (Kadlec et al., 2016; Dhingra et al., 2017; Sordoni et al., 2016, *inter alia*) and focus on subsets of CBT where the elided answer token is either a common noun (CBT-CN) or named entity (CBT-NE). In addition, we use a subsampled version of the CNN benchmark with 100K training examples to save computation.

Unlike extractive question answering benchmarks, where the answer can be an arbitrary span from the passage, cloze benchmarks are often multiple-choice (e.g., answers in CNN and ReCoRD must be single-token entities). To evaluate our extractive question answering modeling approaches on these cloze benchmarks, we simply discard the multiple-choice answers and otherwise treat the benchmark as extractive question answering. Given that converting from cloze to extractive question answering results in information loss (since systems no longer have access to answer choices), the resulting extractive benchmarks are more difficult than their original formulations. See Appendix C.2 for examples from the cloze benchmarks we study.

Results and Discussion. We find that CBT and LAMBADA have high concurrence with SQuAD (Figure 3)—despite their lack of natural questions, cloze benchmarks can recapitulate historical findings on SQuAD.

The CNN benchmark only has moderate concurrence with SQuAD due to a pair of outlier modeling approaches—DocumentReader and DocumentReader without external linguistic features, which both attain ∼69.5 EM on SQuAD and ∼72.5 EM on CNN (Figure 3). We hypothesize that these two modeling approaches are outliers

[4] Taori et al. (2020) observe a similar "knee" when evaluating zero-shot generalization of image classifiers to filtered benchmarks.

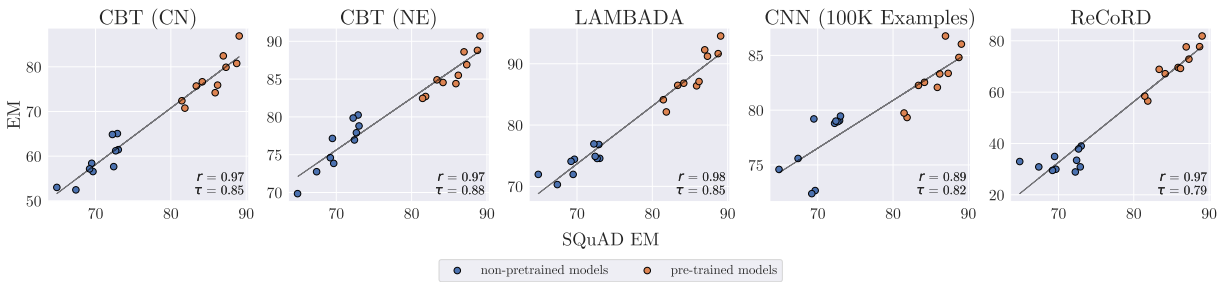

Figure 3: Despite their lack of natural questions, cloze benchmarks like CBT-CN, CBT-NE, and LAMBADA have high concurrence with SQuAD.

because they adopted SQuAD-specific preprocessing strategies that may have also influenced architecture design (see Appendix A.3 for details and discussion).

ReCoRD exhibits low overall concurrence due to the poor performance of older non-pretrained modeling approaches (Figure 3); they all achieve less than 40 EM (the most-common entity baseline is approximately 32.5 EM; Wang et al., 2019). We hypothesize that this stems from ReCoRD's construction procedure, wherein a filtering step removed all instances that were correctly answered by a strong pre-BERT modeling approach (SAN, which attains a reported SQuAD 1.1 development set EM of 76.24; Liu et al., 2018). Many of the non-pretrained models are weaker than SAN and thus do poorly on the SAN-filtered ReCoRD benchmark, leading to low concurrence between ReCoRD and SAN on non-pretrained models. Furthermore, for the (mostly pre-trained) modeling approaches that outperform SAN on SQuAD, ReCoRD and SQuAD have high concurrence.

These results raise the question: is it possible that we may not have needed to construct SQuAD? Since cloze benchmarks like CBT already existed and can recapitulate historical findings on SQuAD, it is possible that CBT may have led to similar modeling innovations as SQuAD. It is worth reiterating that the ability to recapitulate historical findings on SQuAD is a necessary (but not sufficient) condition for driving similar modeling innovations as SQuAD; our retrospective study only considers intrinsic factors, and not extrinsic factors such as the particular benchmarks and ideas that existed at the time (see § 6.1 for further discussion of qualifications of this study).

Given that cloze benchmarks, which are not completely natural, can have high concurrence with SQuAD, how far can we go? How artificial can a benchmark be, while still maintaining high concurrence with SQuAD?

3.3 Existing Synthetic Benchmarks Have Low Concurrence With SQuAD

Synthetic extractive question answering benchmarks contain questions and passages that are programmatically generated (and possibly not even natural language). Synthetic benchmarks offer precise control of benchmark contents, which enables targeted evaluation of specific phenomena (e.g., compositional reasoning; Lake and Baroni, 2018). However, this control often comes at the cost of reduced data diversity—natural language is complex and long-tailed, and synthetic benchmarks are unlikely to cover these nuances.

The bAbI task suite (Weston et al., 2016) is a notable prior attempt at building synthetic question-answering benchmarks. The bAbI benchmark has 20 tasks that each attempt to focus on a particular skill required of a competent dialogue system (e.g., fact retrieval, subject-object relations, counting). The textual data is generated from the state of a simulated toy environment. The creators of the bAbI benchmark hoped that focusing on individual phenomena would enable the community to effectively identify and improve upon targeted shortcomings of models.

Setup. To assess how existing synthetic benchmarks concur with SQuAD, we study the bAbI reading comprehension task suite. In particular, we consider the 11 tasks that can be losslessly converted to an extractive format (Tasks 1, 2, 3, 4, 5, 11, 12, 13, 14, 15, 16). For each task, we use the two officially-released data settings: one setting has 900 training examples and 100 development examples, and the other has 9,000 training examples and 1,000 development examples. In this section, we focus on the setting with 900 training examples, since all modeling approaches do nearly perfectly on al-

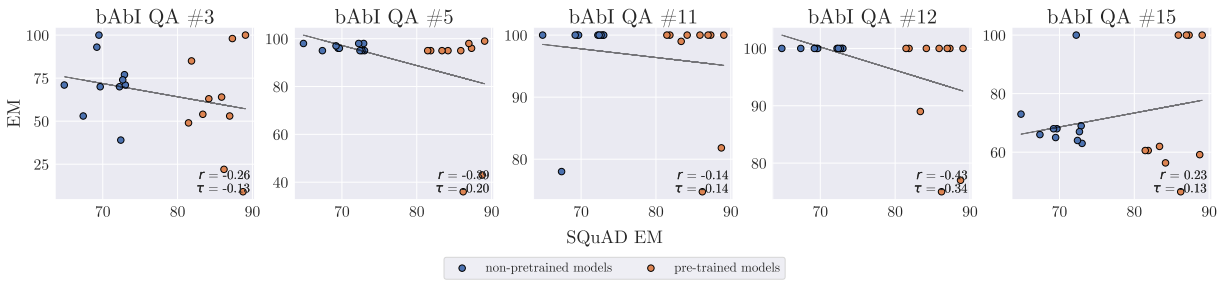

Figure 4: Many modeling approaches perform perfectly or at near-chance performance on the bAbI tasks, limiting their ability to recapitulate historical modeling progress on SQuAD. We show results on a representative subset of tasks here, see Appendix D.3 for the full results on all tasks.

most all tasks with 9,000 examples (Appendix D.3). See Appendix C.3 for examples from the existing synthetic benchmarks we study.

Results and Discussion. We find that the bAbI tasks have low concurrence with SQuAD, confirming popular intuition that existing synthetic benchmarks do not reflect historical improvements on natural benchmarks (Figure 4). Systems often achieve perfect performance on the bAbI tasks, and inability to solve the bAbI tasks does not meaningfully correlate with SQuAD performance.

As a result, focusing on bAbI as benchmark for driving modeling development is unlikely to have led to the same innovations as focusing on SQuAD. For instance, while pre-training consistently yields large gains on SQuAD, the improvements are inconsistent on bAbI. It is worth reiterating that lack of concurrence with SQuAD is not necessarily bad. The bAbI task suite certainly tests *some* types of reasoning, although it also does not address other phenomena found in natural language (or even SQuAD).

4 Constructing Synthetic Benchmarks With High SQuAD Concurrence

Our observation that existing cloze benchmarks can have high concurrence with SQuAD indicates that semi-synthetic benchmarks can nonetheless concur with SQuAD. However, we also find that existing synthetic datasets have low concurrence, so achieving high concurrence not trivial. In this section, we are able to ultimately strike a middle ground by using Wikidata triples to construct WikidataSyntheticQA, a synthetic benchmark that has high concurrence with SQuAD.

Given that cloze-style benchmarks can concur with SQuAD, we focus on generating synthetic cloze examples. Constructing such an example entails generating a passage, a cloze query, and an answer.

Passage Snippet: ... *chests Melchior divorced might whereof 37th Kadima milling raved Salib melanocephala Pilgrims chop Prosser draftsmanship 203 Caesarius madam Deconstruction Guevara Amalia ...*
Question: *Pigs corncrake XXXXX 286 airmanship Kition gracious Modernism Raul*
Answer: *chop*

Figure 5: Example passage and question-answer pair from FuzzySyntheticQA. Despite its lack of sequential structure (let alone natural language structure), this benchmark concurs with SQuAD on non-pretrained modeling approaches.

4.1 FuzzySyntheticQA Has High Concurrence With SQuAD on Non-pretrained Models

Setup. Many SQuAD questions can be answered by exploiting lexical overlap between question and passage tokens (Weissenborn et al., 2017). We construct FuzzySyntheticQA, a synthetic benchmark that does not resemble natural language but specifically targets this fuzzy pattern-matching, and show that is has high concurrence with SQuAD on non-pretrained models.

Figure 5 shows a sample passage and question-answering pairs. FuzzySyntheticQA has 10,000 question-answer pairs over 2,000 passages (5 questions per passage) for training and 10,000 question-answer pairs over 2,000 passages (5 questions per passage) for evaluation.

Passage Generation. We generate the passage by randomly sampling 150 tokens from the uniform distribution over a token vocabulary. The token vocabulary is taken from the WikiText-2 training set (Merity et al., 2017) and has 68,429 types.

Answer Generation. We generate the answer token by randomly selecting a token from the generated passage.

Cloze Question Generation. We generate the cloze question as follows. First, we extract the answer token's local context (up to 10 tokens) and mask out the answer token. Then, we corrupt the cloze question by randomly replacing its tokens with related tokens,[5] locally permuting its tokens (within 3 positions), and applying word dropout (with rate 0.2). This procedure is inspired by the noisy cloze generation method of Lewis et al. (2019).

Results and Discussion. Surprisingly, FuzzySyntheticQA has high concurrence with SQuAD on non-pretrained modeling approaches (r = 0.95, τ = 0.78; Figure 6). However, English-pretrained modeling approaches perform worse than non-pretrained approaches, resulting in low concurrence overall. These results demonstrate that benchmarks that lack much language structure can track historical progress in non-pretrained models on SQuAD. Furthermore, non-pretrained modeling approaches that achieve better SQuAD performance are also better token-level pattern-matchers, and token-level pattern-matching alone is sufficient to recapitulate historical results on SQuAD.

To examine whether the ineffectiveness of pre-training (and the resulting low overall concurrence) can be trivially attributed to the relative lack of linguistic structure in the passages, we experiment with generating FuzzySyntheticQA questions from passages taken from the English Wikipedia. Such a benchmark exhibits similar levels of low concurrence to simply sampling from the uniform distribution over tokens, indicating that low concurrence comes from more than just the lack of natural language passages (see Appendix F); simply making our passages more closely resemble natural language will not yield high concurrence.

Given that we observed that existing cloze benchmarks *can* have high concurrence with SQuAD, the cloze format is not necessarily at fault for the synthetic fuzzy pattern-matching benchmark's low overall concurrence with SQuAD. Furthermore, since taking passages from English Wikipedia does not increase overall concurrence, the low concurrence does not stem from lack of natural passages. These observations jointly imply that the low concurrence may come from our cloze query generation procedure, which currently only involves surface-level noise—what matters is not necessarily the surface-level appearance and features of a benchmark, but rather the phenomena and reasoning abilities required to solve its examples.

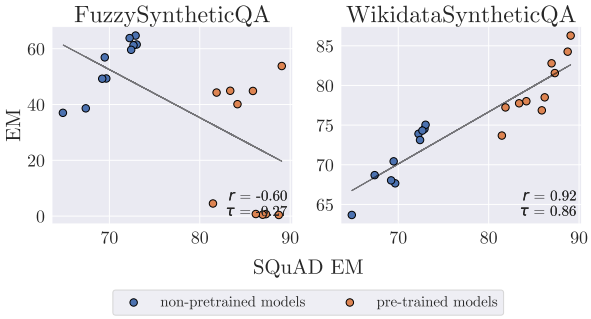

Figure 6: **Left:** While FuzzySyntheticQA has high concurrence with SQuAD on non-pretrained modeling approaches, English pre-training does not increase performance, leading to low overall concurrence. **Right:** Despite lacking natural language structure, WikidataSyntheticQA has high concurrence with SQuAD.

4.2 WikidataSyntheticQA Has High Concurrence With SQuAD

Setup. We examine whether it is possible to build synthetic cloze benchmarks that do not resemble natural language, yet require some of the reasoning capabilities necessary to handle complex natural language phenomena in SQuAD and other human-constructed benchmarks. Knowledge graphs like Wikidata are rich sources of complex relations between entities, which enables us to increase the complexity of question-passage token relations in our generated examples. We constructed WikidataSyntheticQA, a benchmark derived from Wikidata triples; Figure 7 shows a sample passage and question-answering pairs. Our synthetic benchmark has 10,000 question-answer pairs over 8,376 passages (between 1 and 5 questions per passage) for training and 9,835 question-answer pairs over 1,967 passages (5 questions per passage) for evaluation.

Wikidata Preliminaries. Wikidata is a knowledge graph connecting entities via relations. Wiki-

[5] A token can be replaced with one of its 100 approximate nearest neighbor tokens in the vocabulary, measured by vector distance in the pre-trained English FastText embeddings (300-dimensional, trained with subword information on Common Crawl).

Passage Snippet: *Mae Jemison profession astronaut . STS-47 orbits completed 126.0 . STS-47 crew member Mae Carol Jemison . Mae Jemison worked for NASA . Mae C. Jemison award received Rachel Carson Award . Mae Jemison birthplace Decatur . ...*
Question: *human spaceflight orbits completed XXXXX*
Answer: *126.0*
Question: *Rachel Carson Award honor received by XXXXX*
Answer: *Mae C. Jemison*
Question: *Human XXXXX The River City*
Answer: *Decatur*

Figure 7: Example passage and question-answer pairs from WikidataSyntheticQA. Despite its lack of natural language passages or questions, this benchmark concurs with SQuAD.

data entities and relations include a *label*, the most common name that an entity is known by, and *aliases*, alternative names for entities. For example, the entity Mae_C._Jemison has the label *"Mae C. Jemison"*, with aliases *"Mae Jemison"* and *"Mae Carol Jemison"*. Similarly, the relation place_of_birth has label *"place of birth"*, with aliases such as *"birthplace"*, *"born in"*, *"location of birth"*, etc. We treat labels and aliases as potential surface realizations of entities and relations.

Generation Preliminaries. A prerequisite of our generation process is a set of Wikidata triples. To select these triples, we first randomly choose a seed entity from the 10,000 Wikidata entities with the highest PageRank score (Page et al., 1999; Thalhammer and Rettinger, 2016).[6] We then extract the triples of the seed entity, as well as the triples of all entities connected to the seed entity. Finally, we randomly subsample 50 triples for use in generation.

Passage Generation. We use the following procedure to generate passages. Given the set of 50 Wikidata triples, we realize triples into textual surface forms by selecting a random Wikidata label or alias for each triple element. For example, the triple (Mae_C._Jemison, place_of_birth, Decatur) may be realized as the string *"Mae Jemison birthplace The River City"*. The final passage is formed by concatenating the realizations of all triples, adding a delimiter token between realized triples to mimic sentence structure.

[6]We used pre-computed PageRank scores from github.com/athalhammer/danker. The scores were computed on the Wikipedia dump from 2020-04-26.

Answer Generation. We generate an answer span by selecting a random triple used in the passage generation process, and then choosing a random element of that triple. The passage realization of this random element is the answer span.

Cloze Question Generation. We generate the cloze question by first taking the triple used to generate the answer span and masking out the answer's triple element.

Then, we optionally and randomly replace the predicate with its inverse (if one exists), reversing the subject and the object to maintain consistency. For example, given the triple (Mae_C._Jemison, employer, NASA), with NASA as the answer element, the resultant cloze query would be (Mae_C._Jemison, employer, [MASK]). This cloze query can be transformed to the query ([MASK], employee, Mae_C._Jemison), since employer and employee are inverse relations of each other.

We also optionally and randomly replace the remaining unmasked entity (i.e., the triple subject or object that was not masked) with one of its hypernyms. For example, the cloze query ([MASK], employer, NASA) could become ([MASK], employer, space_agency). A reasonable natural question corresponding to this cloze query would be *"Who was employed by a space agency?"*, challenging models to know that NASA is a space agency (more formally, that the relation (NASA, instance_of, space_agency) holds).

Results and Discussion. We find that WikidataSyntheticQA has high concurrence with SQuAD, demonstrating that a synthetic benchmarks without natural language passages or questions *can* recapitulate findings made on SQuAD (Figure 6). This result indicates that naturalness is not a necessary quality of benchmarks with high concurrence with SQuAD.

We hypothesize that WikidataSyntheticQA has higher SQuAD concurrence than our synthetic fuzzy pattern-matching benchmark because correctly answering its examples often requires reasoning about hypernymy relations between entities and inverse relations between predicates—it is conceivable that pre-trained modeling approaches are better-equipped to handle and use these lexical relations. In addition, the Wikidata aliases provide sufficient lexical variation such that the benchmark is not trivially solvable through string

pattern-matching (removing aliases from the generation procedure results in near-perfect performance from all modeling approaches).

In contrast, high performance on FuzzySyntheticQA simply requires reversing a surface-level noise function and matching similar tokens in the passage and question—we hypothesize that the majority of the synthetic pattern-matching benchmark's difficulty comes from implicitly learning the FastText similarity matrix used to replace words with nearest neighbors.

5 Benchmark Size Minimally Affects Concurrence With SQuAD

Larger training datasets yield better models with higher end-task accuracy, but are larger training datasets necessary for comparing modeling approaches? Smaller benchmarks are cheaper to construct, and their reduced size often means that running experiments requires less compute and time, enabling researchers to iterate on ideas faster.

SQuAD was notable for its size—with 87,599 question-answer pairs for training, it was almost two orders of magnitude larger than any previous manually-labeled reading comprehension benchmark (e.g., MCTest; Richardson et al., 2013). Did the SQuAD have to be large to be an effective benchmark? To study this question, we measure how smaller benchmarks recapitulate historical findings on SQuAD.

Setup. We only modify the number of training examples and use the original SQuAD development set for evaluation in all experiments; larger evaluation sets are critical for well-powered model comparisons (Card et al., 2020). In addition, the standard 80%/20% train-test split (or 80%/10%/10% train-development-test split) means that evaluation sets are generally much smaller than training sets, so reducing the amount of evaluation data has a marginal effect on overall benchmark cost and computation requirements compared to reducing the amount of training data.

Results and Discussion. We find that modulating SQuAD training set size minimally affects concurrence with the full SQuAD benchmark, beyond a baseline threshold of 20K examples (Figure 8). Our results suggest that a reduced SQuAD training dataset with only 20K examples may have been sufficient to spur the same modeling innovation as the original training split with 85K+ examples.

6 Discussion

6.1 Qualifications of This Study

This work is motivated by an unanswerable causal question: *Can small and synthetic benchmarks drive modeling innovation?* Our results are inconclusive regarding this unanswerable question, but they do not rule out the possibility. In this section, we discuss the qualifications of this study that prevent us from drawing stronger conclusions.

Gaps Between the Causal and the Counterfactual Question. While our work is motivated by the causal question, much of this study discusses an equally unanswerable counterfactual question: *If SQuAD was never constructed, could a small and synthetic benchmark have led to the same eventual modeling innovation?* While the causal question considers *new* benchmarks, the counterfactual question considers whether an alternative benchmark could have replaced an existing benchmark. Since the counterfactual question focuses on a particular benchmark and task formulation, while the causal question extends beyond individual benchmarks and tasks, an answer to the counterfactual question for a particular benchmark and task formulation may not hold for the causal question (e.g., due to possible SQuAD-specificity).

SQuAD-Specificity. We focus on SQuAD in this work, since the majority of extractive question answering modeling approaches are originally developed for the benchmark. Although extractive question answering is a flexible format for evaluating modeling approaches on a wide range of phenomena (Gardner et al., 2019; Fisch et al., 2019), focusing on a different benchmark or task (e.g., machine translation on WMT 2014 EN-DE) may significantly change the set of modeling approaches we use to calculate concurrence, potentially affecting our conclusions.

Our focus on SQuAD does not imply that it is a gold standard for progress in the field—for example, its crowdsourcing process results in passages and questions that contain high lexical overlap. Despite these flaws, it has proven to be a useful surrogate testbed for some of the core building blocks of language understanding (e.g., predicate-argument structure). While real-world question-answering benchmarks may more-closely resemble particular target user populations, it is unclear whether they target dramatically different phenomena than those found in SQuAD. Given that we empirically ob-

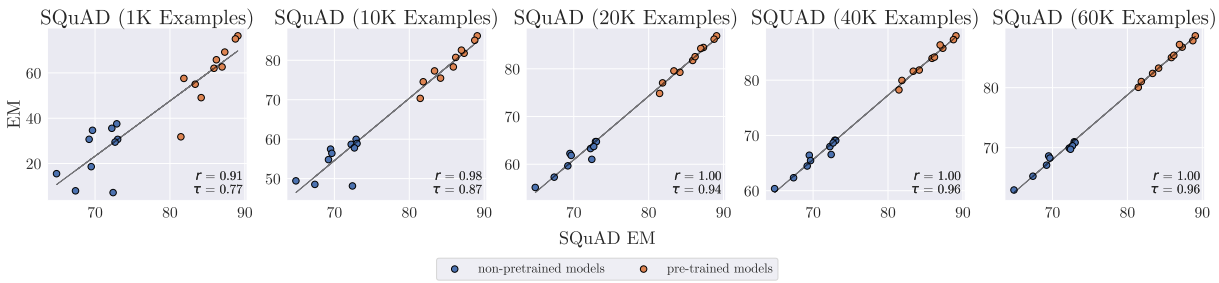

Figure 8: Modulating SQuAD training set size minimally affects concurrence with the full SQuAD benchmark, beyond a baseline threshold of 20K examples. From these results, it is conceivable that the creators of SQuAD may not have needed to collect 85K+ training examples, and that 20K could have sufficed.

serve that many human-constructed benchmarks have high concurrence with SQuAD (§3.1), we believe that our results and conclusions generalize beyond SQuAD.

It is also worth noting that our results are also inherently biased by the publication process and researcher motivations—all of the modeling approaches we used to measure concurrence came from published papers, which were perhaps easier to publish because they evaluated on a popular benchmark like SQuAD. As a result, there is an implicit selection bias toward modeling approaches that improve over prior work on SQuAD (since authors often choose not to write about techniques that do not approach or exceed the state of the art).

Extrinsic Factors. The counterfactual question is impossible to answer because it heavily depends on extrinsic factors. For example, the NLP community's reception of a benchmark is crucial for its impact and ability to influence modeling innovation.

At the most basic level, a benchmark will likely have little impact if researchers are not interested in working on it (perhaps because it is small or does not resemble natural language). To fairly answer whether a small and synthetic benchmark would have led to the same modeling innovation as SQuAD requires controlling for the impact of SQuAD's popularity on its ability to influence modeling innovation.

More fundamentally, benchmarks affect how researchers think. Modeling approaches are not developed on benchmarks in isolation; researchers talk to each other about common interests and are influenced by other ideas in the community. As a result, the unique state of the research community at the particular point in time that a benchmark is introduced is inextricably linked to how researchers develop new modeling approaches on the benchmark—this state cannot be replicated or controlled. In addition, the content of benchmarks themselves inspire researchers; whether by error analysis of system output or studying the training data itself, the data within benchmarks may be useful for sparking researcher creativity.

Finally, modeling approaches flow freely across benchmarks and tasks (in part due to the aforementioned extrinsic factors), and it is difficult to ascertain the influence of a particular benchmark in developing a new modeling approach. For example, while self-attention gained widespread popularity after its application to machine translation as a core component of the Transformer (Vaswani et al., 2017), it had previously been successfully applied to natural language inference on SNLI (Cheng et al., 2016; Parikh et al., 2016) and reading comprehension on SQuAD (Wang et al., 2017), although the basic idea of attention came earlier from translation (Bahdanau et al., 2015). Even more indirectly, while neural network-based approaches to NLP had shown sporadic promise before their widespread acceptance (Bengio et al., 2003; Collobert and Weston, 2008; Collobert et al., 2011, *inter alia*), their successful application to computer vision tasks (e.g., Krizhevsky et al., 2012, *inter alia*) was influential for their current popularity in NLP.

These extrinsic factors are impossible to study and capture, and our retrospective study thus focuses solely on intrinsic factors of whether benchmarks can recapitulate historical modeling improvements.

6.2 Direct vs. Indirect Improvements

An overarching goal of NLP research is to develop and improve systems to better serve real users.

Direct improvements to systems are one way of achieving these goals. For example, by collecting data that better represents target user populations and tuning hyperparameters or the model architecture, researchers can dramatically improve the performance of deployed systems on real users. Furthermore, evaluating proposed system changes with real-world inputs from users interacting with a deployed system is invaluable for inspiring confidence that such changes genuinely improve the user experience (and that researchers are not "climbing up the wrong hill"; de Vries et al., 2020). Furthermore, real-world data is uniquely representative of domain- and user population- specific challenges required for improving the system of interest. If the goal is to directly improve deployed systems, real-world data is often indispensable.

However, real-world data is not always practically available—existing real-world benchmarks (Nguyen et al., 2016; He et al., 2018; Kwiatkowski et al., 2019) cover only a fraction of possible user populations and needs. Furthermore, constructing new real-world benchmarks is difficult and inaccessible for the many groups without access to real users interacting with a real system, and privacy concerns can also limit the ethical collection and distribution of such benchmarks (Richards and King, 2014; Metcalf and Crawford, 2016; Paullada et al., 2020, *inter alia*).

Direct improvements are not the only way to improve the quality of real systems—*indirect* improvements, which are not developed to improve a particular real-world application, can yield general improvements across a range of tasks and systems—the vast majority of academic research is making indirect improvements.[7] For example, Hochreiter and Schmidhuber (1997) developed the LSTM to better handle long-range dependencies in sequential data, evaluating them on a range of synthetic tasks. However, LSTMs also improve the performance of many real-world systems (e.g., Wu et al., 2016, *inter alia*).

While natural data is essential for developing direct improvements, it is not necessary for developing indirect improvements. In this work, we question what benchmarks enable us to effectively develop indirect improvements and use concurrence as a link between indirect and direct improvements.

[7]Even work on real-world benchmarks are often not direct improvements, since their examples may be filtered (in the case of NaturalQuestions) and user inputs change over time.

6.3 Can Synthetic Benchmarks Drive Modeling Innovation?

Although WikidataSyntheticQA has high concurrence with SQuAD and its construction does not require the existence of SQuAD, it is insufficient for determining whether synthetic benchmarks can drive modeling innovation. To further explore this possibility, one could construct and release a synthetic benchmark and see whether it spurs useful progress in the NLP community.[8]

Benefits of Synthetic Benchmarks. Synthetic benchmarks have many practical advantages, especially when compared to crowdsourced benchmarks. Synthetic benchmarks are unique because they enable fine-grained control over and understanding of benchmark contents; this makes them particularly useful for isolating the challenges required for handling specific tail phenomena. Although the experiments in this work consider the *i.i.d.* setting, synthetic benchmarks also enable researchers to flexibly and programmatically define subgroups for training and evaluation, making it easy to evaluate whether systems extrapolate across subgroups. Benchmarks creators should have a strong understanding of what challenges their new benchmarks pose, why their new benchmarks are difficult, and what high benchmark performance means about model capabilities—synthetic benchmarks enable us to precisely characterize these concerns.

However, synthetic benchmarks are certainly not a panacea, and it is important to note that synthetic benchmarks cannot completely replace natural language benchmarks for model development. For example, natural language data serves a unique role in inspiring model developers and can help build crucial intuition. Furthermore, although synthetic benchmarks may show promise for facilitating targeted progress, there is a real concern that synthetic benchmark creators may oversimplify and fail to capture the challenges associated with handling the natural language phenomena of interest—the

[8]Note that this still does not answer the causal question, since we cannot know what would have happened if we had not released the synthetic benchmark—it is possible that the community would have made the modeling innovations regardless.

risk of "climbing the wrong hill" is higher when using synthetic benchmarks, since their simplicity can make it easier for researchers to develop synthetic benchmark-specific approaches and ignore the complexity of the problems they are addressing (Lighthill, 1973).

In addition, although human-constructed benchmarks frequently contain spurious correlations that enable systems to achieve high absolute performance with simple heuristics (Gururangan et al., 2018; Poliak et al., 2018; Tsuchiya, 2018), the use of synthetic benchmarks does not necessarily make these spurious correlations more or less prevalent—it is entirely possible to build synthetic benchmarks with unintended spurious correlations.

Ultimately, synthetic benchmarks have distinct benefits and drawbacks when compared to conventional human-constructed benchmarks. Our results provide initial evidence that they may be useful for driving modeling innovation, and we encourage the community to further explore their potential uses and limitations.

7 Related Work

To the best of our knowledge, there are no comparable meta-analyses that study how modeling *approaches* generalize *across* benchmarks or what qualities of benchmarks are (un)necessary for such generalization.

7.1 Transferability and Out-of-Domain Generalization

Motivated by the danger of adaptive overfitting due to test set reuse (Dwork et al., 2015), a recent line of work examines whether *systems* have overfit to particular test sets by taking existing systems and testing them on newly-constructed test sets (Recht et al., 2019; Yadav and Bottou, 2019; Miller et al., 2020). Recent work has also studied whether higher-performing systems are more robust by measuring the extent to which system improvements transfer from in-domain to out-of-domain test sets (Taori et al., 2020; Djolonga et al., 2020).

In contrast, our work examines whether *modeling approaches* generalize across benchmarks (a training dataset, test dataset, and evaluation metric). Our experiments *train and test* modeling approaches on a variety of existing and newly-constructed benchmarks, rather than testing systems on a new test set built for particular benchmark of interest. In this regard, our work is similar to the study of Kornblith et al. (2019), who find that performance improvements on ImageNet are well-correlated with performance improvements on other benchmarks.

7.2 Synthetic Data in NLP

Synthetic Data During Training and Testing: Synthetic Benchmarks. Synthetic benchmarks only use synthetic data during training *and* testing for the purpose of comparing modeling approaches (Weston et al., 2016; Lake and Baroni, 2018; Chevalier-Boisvert et al., 2018; Saxton et al., 2019; Kim and Linzen, 2020; Ruis et al., 2020; Hui et al., 2020; Keysers et al., 2020). Rather than proposing a new synthetic benchmark for the community to work on, we seek to better understand and connect their potential for driving modeling innovation to natural benchmarks via concurrence.

Synthetic Data During Training: Data Augmentation. There is much existing work on synthetic data augmentation, which generates additional synthetic training data to improve system performance on non-synthetic test data—these results indicate that synthetic data can clearly be useful during training. For example, Andreas (2020) and Jia and Liang (2016) produce synthetic training examples for sequence modeling by recombining existing training examples, doing so with the aim of providing systems with a compositional inductive bias. Geva et al. (2020) improve the numerical reasoning ability of pre-trained language models by further pre-training on automatically-generated numerical and textual synthetic data.

Synthetic Data During Testing: Challenge Sets. Motivated in part by observations that systems can perform well on benchmarks by exploiting benchmark-specific predictive features (Gururangan et al., 2018; Poliak et al., 2018; Tsuchiya, 2018), a recent line of work has used synthetic test data to probe their weaknesses (Jia and Liang, 2017; Naik et al., 2018; Richardson et al., 2020; Si et al., 2020, *inter alia*). For example, Jia and Liang (2017) demonstrate that appending automatically-generated suffixes to passages drastically reduces the performance of state-of-the-art extractive question answering systems.

Synthetic test data has also been useful for better understanding what high-performing systems learn from benchmarks (Marvin and Linzen, 2018; McCoy et al., 2019; Richardson and Sabharwal, 2020, *inter alia*). For example, McCoy et al. (2019) use

templates to construct natural language inference examples to show that models trained on popular natural language inference benchmarks adopt fallible syntactic heuristics.

8 Conclusion

Although large natural datasets are crucial ingredients for training accurate, deployable NLP systems, we find that naturalness and size are not necessary qualities of benchmarks that recapitulate progress on SQuAD. Benchmarks with varying naturalness and size can offer challenges with relevance to natural language.

These results underscore that large amounts of natural data do not necessarily make benchmarks worthwhile research targets, and raise the possibility that synthetic benchmarks have the potential to play a greater role in driving the development of new modeling approaches. We hope the community further studies synthetic benchmarks to better understand their potential, utility, and limitations.

Acknowledgments

We thank Jesse Dodge, Matt Gardner, John Hewitt, Omer Levy, Kevin Lin, Julian Michael, Alison Ng, Roy Schwartz, Megha Srivastava, Eric Wallace, and Tianyi Zhang for their thoughtful feedback on and discussion about this work. NL is supported by an NSF Graduate Research Fellowship under grant number DGE-1656518.

We would also thank Pengxiang Cheng for sharing the the LAMBADA benchmark converted to extractive question answering format (enabling us to replicate the LAMBADA setup of Cheng and Erk, 2020). We also thank Xin Liu for publicly releasing an implementation of the MnemonicReader model, and Felix Wu for publicly releasing implementations of the FusionNet and DocumentReader models. We are grateful to Max Bartolo for providing valuable advice regarding fine-tuning pretrained language models on small benchmarks and sharing details about the hyperparameter optimization performed in Bartolo et al. (2020).

We acknowledge the Python community (Van Rossum and Drake Jr., 1995) for developing the core set of tools that enabled this work, including PyTorch (Paszke et al., 2019), NLTK (Bird and Loper, 2004), Jupyter (Kluyver et al., 2016), Matplotlib (Hunter, 2007), seaborn (Waskom and the seaborn development team, 2020), numpy (Oliphant, 2006; Walt et al., 2011; Harris et al., 2020), pandas (McKinney, 2010; pandas development team, 2020), and SciPy (Virtanen et al., 2020).

References

Jacob Andreas. 2020. Good-enough compositional data augmentation. In *Proc. of ACL*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. of ICLR*.

Ondrej Bajgar, Rudolf Kadlec, and Jan Kleindienst. 2017. Embracing data abundance. In *Proc. of ICLR (Workshop)*.

Max Bartolo, Alastair Roberts, Johannes Welbl, Sebastian Riedel, and Pontus Stenetorp. 2020. Beat the AI: Investigating adversarial human annotation for reading comprehension. *Transactions of the Association for Computational Linguistics*, 8:662–678.

Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *The Journal of Machine Learning Research*, 3:1137–1155.

Erik Bernhardsson and the Annoy development team. 2020. github.com/spotify/annoy.

Steven Bird and Edward Loper. 2004. NLTK: The natural language toolkit. In *Proc. of ACL (System Demonstrations)*.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proc. of EMNLP*.

Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. With little power comes great responsibility. In *Proc. of EMNLP*.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Proc. of ACL*.

Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *Proc. of EMNLP*.

Pengxiang Cheng and Katrin Erk. 2020. Attending to entities for better text understanding. In *Proc. of AAAI*.

Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. 2018. BabyAI: A platform to study the sample efficiency of grounded language learning. In *Proc. of ICLR*.

Christopher Clark and Matt Gardner. 2018. Simple and effective multi-paragraph reading comprehension. In *Proc. of ACL*.

Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *Proc. of ICLR*.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proc. of ICML*.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL*.

Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William W. Cohen, and Ruslan Salakhutdinov. 2017. Gated-attention readers for text comprehension. In *Proc. of ACL*.

Josip Djolonga, Jessica Yung, Michael Tschannen, Rob Romijnders, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Matthias Minderer, Alexander D'Amour, Dan Moldovan, Sylvan Gelly, Neil Houlsby, Xiaohua Zhai, and Mario Lucic. 2020. On robustness and transferability of convolutional neural networks. ArXiv:2007.08558.

Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Proc. of NAACL*.

Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon Roth. 2015. Preserving statistical validity in adaptive data analysis. In *Proc. of STOC*.

Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. MRQA 2019 shared task: Evaluating generalization in reading comprehension. In *Proc. of MRQA*.

Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, and Sewon Min. 2019. Question answering is a format; when is it useful? ArXiv:1909.11291.

Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proc. of NLP-OSS*.

Mor Geva, Ankit Gupta, and Jonathan Berant. 2020. Injecting numerical reasoning skills into language models. *Proc. of ACL 2020*.

Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. 2018. Annotation artifacts in natural language inference data. In *Proc. of NAACL*.

Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Rıo, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. Array programming with NumPy. *Nature*, 585(7825):357–362.

Wei He, Kai Liu, Jing Liu, Yajuan Lyu, Shiqi Zhao, Xinyan Xiao, Yuan Liu, Yizhong Wang, Hua Wu, Qiaoqiao She, Xuan Liu, Tian Wu, and Haifeng Wang. 2018. DuReader: a chinese machine reading comprehension dataset from real-world applications. In *Proc. of MRQA*.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa

Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Proc. of NeurIPS*.

Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The Goldilocks Principle: Reading children's books with explicit memory representations. In *Proc. of ICLR*.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Minghao Hu, Yuxing Peng, and Xipeng Qiu. 2017. Reinforced Mnemonic Reader for machine reading comprehension. ArXiv:1705.02798v3.

Hsin-Yuan Huang, Chenguang Zhu, Yelong Shen, and Weizhu Chen. 2018. FusionNet: Fusing via fully-aware attention with application to machine comprehension. In *Proc. of ICLR*.

David Yu-Tung Hui, Maxime Chevalier-Boisvert, Dzmitry Bahdanau, and Yoshua Bengio. 2020. BabyAI 1.1. ArXiv:2007.12770.

John D. Hunter. 2007. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95.

Robin Jia and Percy Liang. 2016. Data recombination for neural semantic parsing. In *Proc. of ACL*.

Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proc. of EMNLP*.

Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.

Mandar Joshi, Eunsol Choi, Omer Levy, Daniel Weld, and Luke Zettlemoyer. 2019. pair2vec: Compositional word-pair embeddings for cross-sentence inference. In *Proc. of NAACL*.

Rudolf Kadlec, Martin Schmid, Ondrej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. In *Proc. of ACL*.

Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. Measuring compositional generalization: A comprehensive method on realistic data. In *Proc. of ICLR*.

Najoung Kim and Tal Linzen. 2020. COGS: A compositional generalization challenge based on semantic interpretation. In *Proc. of EMNLP*.

Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica Hamrick, Jason Grout, Sylvain Corlay, Paul Ivanov, Damián Avila, Safia Abdalla, Carol Willing, and Jupyter development team. 2016. Jupyter Notebooks - a publishing format for reproducible computational workflows. In *Proc. of EPUB*.

Simon Kornblith, Jonathon Shlens, and Quoc V. Le. 2019. Do better imagenet models transfer better? In *Proc. of CVPR*.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Proc. of NeurIPS*.

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*, 7:453–466.

Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *Proc. of ICML*.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *Proc. of ICLR*.

Kenton Lee, Shimi Salant, Tom Kwiatkowski, Ankur Parikh, Dipanjan Das, and Jonathan Berant. 2016. Learning recurrent span representations for extractive question answering. ArXiv:1611.01436.

Patrick Lewis, Ludovic Denoyer, and Sebastian Riedel. 2019. Unsupervised question answering by cloze translation. In *Proc. of ACL*.

James Lighthill. 1973. Artificial intelligence: A general survey. In *Artificial Intelligence: a paper symposium*. Science Research Council London.

Xiaodong Liu, Yelong Shen, Kevin Duh, and Jianfeng Gao. 2018. Stochastic answer networks for machine reading comprehension. In *Proc. of ACL*.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach.

Christopher D. Manning. 2006. Local textual inference: It's hard to circumscribe, but you know it when you see it—and NLP needs it.

Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330.

Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. In *Proc. of EMNLP*.

R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *Proc. of ACL*.

Wes McKinney. 2010. Data structures for statistical computing in Python. In *Proc. of SciPy*.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer sentinel mixture models. In *Proc. of ICLR*.

Jacob Metcalf and Kate Crawford. 2016. Where are human subjects in big data research? the emerging ethics divide. *Big Data & Society*, 3(1).

Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. 2018. Crowdsourcing question-answer meaning representations. In *Proc. of NAACL*.

John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. 2020. The effect of natural distribution shift on question answering models. *Proc. of ICML*.

Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *Proc. of COLING*.

Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. In *Proc. of CoCo*.

Travis E. Oliphant. 2006. *A guide to NumPy*. Trelgol Publishing USA.

Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. Who did What: A large-scale person-centered cloze dataset. In *Proc. of EMNLP*.

Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.

Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proc. of ACL*.

Ankur P. Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proc. of EMNLP*.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit

Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Proc. of NeurIPS*.

Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. 2020. Data and its (dis) contents: A survey of dataset development and use in machine learning research. In *Proc. of ML-RSA*.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proc. of EMNLP*.

Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis only baselines in natural language inference. In *Proc. of *SEM*.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. Technical report, OpenAI.

Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for SQuAD. In *Proc. of ACL*.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proc. of EMNLP*.

Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. Do ImageNet classifiers generalize to ImageNet? *Proc. of ICML*.

Neil M. Richards and Jonathan H. King. 2014. Big data ethics. *Wake Forest Law Review*, 49:393.

Kyle Richardson, Hai Hu, Lawrence S. Moss, and Ashish Sabharwal. 2020. Probing natural language inference models through semantic fragments. In *Proc. of AAAI*.

Kyle Richardson and Ashish Sabharwal. 2020. What does my QA model know? devising controlled probes using expert knowledge. *Transactions of the Association for Computational Linguistics*, 8:572–588.

Matthew Richardson, Christopher J.C. Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proc. of EMNLP*.

Laura Ruis, Jacob Andreas, Marco Baroni, Diane Bouchacourt, and Brenden M. Lake. 2020. A benchmark for systematic generalization in grounded language understanding. In *Proc. of NeurIPS*.

David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. Analysing mathematical reasoning abilities of neural models. In *Proc. of ICLR*.

Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *Proc. of ICLR*.

Chenglei Si, Ziqing Yang, Yiming Cui, Wentao Ma, Ting Liu, and Shijin Wang. 2020. Benchmarking robustness of machine reading comprehension models. ArXiv:2004.14004.

Alessandro Sordoni, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2016. Iterative alternating neural attention for machine reading. ArXiv:1606.02245.

Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. 2020. Measuring robustness to natural distribution shifts in image classification. In *Proc. of NeurIPS*.

The pandas development team. 2020. github.com/pandas-dev/pandas.

Andreas Thalhammer and Achim Rettinger. 2016. PageRank on Wikipedia: Towards general importance scores for entities. In *Proc. of ESWC*.

Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordoni, Philip Bachman, and Kaheer Suleman. 2017. NewsQA: A machine comprehension dataset. In *Proc. of RepL4NLP*.

Masatoshi Tsuchiya. 2018. Performance impact caused by hidden bias of training data for recognizing textual entailment. In *Proc. of LREC*.

Guido Van Rossum and Fred L. Drake Jr. 1995. *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. of NeurIPS*.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*.

Harm de Vries, Dzmitry Bahdanau, and Christopher D. Manning. 2020. Towards ecologically valid research on language user interfaces. ArXiv:2007.14435.

Stéfan van der Walt, S. Chris Colbert, and Gaël Varoquaux. 2011. The NumPy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. In *Proc. of NeurIPS*.

Shuohang Wang and Jing Jiang. 2017. Machine comprehension using match-LSTM and answer pointer. In *Proc. of ICLR*.

Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proc. of ACL*.

Michael Waskom and the seaborn development team. 2020. github.com/mwaskom/seaborn.

Dirk Weissenborn, Georg Wiese, and Laura Seiffe. 2017. Making neural QA as simple as possible but not simpler. In *Proc. of CoNLL*.

Jason Weston, Antoine Bordes, Sumit Chopra, Sasha Rush, Bart van Merrienboer, Armand Joulin, and Tomas Mikolov. 2016. Towards AI-complete question answering: A set of prerequisite toy tasks. In *Proc. of ICLR*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proc. of EMNLP (System Demonstrations)*.

Felix Wu, Boyi Li, Lequn Wang, Ni Lao, John Blitzer, , and Kilian Q. Weinberger. 2019. FastFusionNet: New state-of-the-art for DAWNBench SQuAD. ArXiv:1902.11291.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. ArXiv:1609.08144.

Chhavi Yadav and Léon Bottou. 2019. Cold case: The lost MNIST digits. In *Proc. of NeurIPS*.

Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W. Cohen, and Ruslan Salakhutdinov. 2017. Words or characters? fine-grained gating for reading comprehension. In *Proc. of ICLR*.

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 *Proc. of EMNLP*.

Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. 2018. QANet: Combining local convolution with global self-attention for reading comprehension. In *Proc. of ICLR*.

Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. 2018. ReCoRD: Bridging the gap between human and machine commonsense reading comprehension. ArXiv:1810.12885.

Appendices

A Implementation Details of Modeling Approaches Evaluated

We evaluated a representative subset of 20 extractive question answering modeling approaches, published between 2016 to 2020 (Table 1). Below, we describe implementation details for all the modeling approaches evaluated.

Modeling Approach	SQuAD 1.1 Dev. EM	
	Our Reproduction	Published
RaSoR	64.9	66.4
BiDAF	67.4	67.7
DocumentReader	69.7	69.5
DocumentReader (no external features)	69.2	-
BiDAF++	69.5	71.6
MnemonicReader	73.0	71.8
MnemonicReader (no external features)	72.7	-
QANet	72.4	73.6
FusionNet	72.9	75.0
FusionNet (no external features)	72.2	-
BERT (base, uncased)	81.5	80.8
BERT (large, uncased)	84.2	84.1
BERT (large, uncased, whole-word masking)	87.3	86.7
ALBERT (base, V1)	81.9	82.3
ALBERT (xxlarge, V1)	89.1	89.3
RoBERTa (base)	83.4	-
RoBERTa (large)	87.0	88.9
ELECTRA (base)	85.9	84.5
SpanBERT (base)	86.2	-
SpanBERT (large)	88.7	88.1

Table 1: Published and reproduced SQuAD 1.1 EM of all 20 modeling approaches used for assessing concurrence. "-" indicates that the modeling approach has no published SQuAD 1.1 EM result.

A.1 RaSoR

We reimplement the RaSoR model of (Lee et al., 2016) with PyTorch in the AllenNLP (Gardner et al., 2018) framework, following the original paper as closely as possible. While the authors released an implementation of their method (github.com/shimisalant/rasor), the codebase is in Theano and inexplicably fails on passages that are significantly longer than those found in SQuAD (e.g., those found in the CNN benchmark).

A.2 BiDAF

We use the reimplementation of BiDAF (Seo et al., 2017) found in AllenNLP (Gardner et al., 2018).

A.3 DocumentReader (with and without external features)

We use an reimplementation of DocumentReader (Chen et al., 2017) released at github.com/felixgwu/FastFusionNet. The original DocumentReader approach uses external features from a part-of-speech tagger and named entity recognition system. To fairly compare to systems that do not use such external resources, we also run the models without these features. We keep the hand-crafted term-frequency and token exact match features defined in the DocumentReader paper.

We also make some changes to the DocumentReader preprocessing code. In particular, the original implementation (github.com/facebookresearch/DrQA) of these two modeling approaches (intended for training and evaluation on SQuAD) replaces all tokens without a pre-trained GloVe embedding (trained on 840B tokens from the Common Crawl) with a special unknown token—the reimplementation we use adopts the same practice. This preprocessing assumption works well for SQuAD, since the vast majority of SQuAD tokens also appear in the GloVe vocabulary. However, this preprocessing assumption does not apply to CNN—many of the special @entity*N* and @placeholder markers, which anonymize entities to prevent models from deriving answers from world knowledge, are not in the GloVe vocabulary. As a result, the original DocumentReader implementation maps them all to a single unknown token, effectively preventing the model from telling valid answer choices apart and yielding a model that performs no better than the majority baseline. Keeping these special tokens in the model's vocabulary enables differentiating between different entities in a passage, which naturally improves performance (and are the reported numbers)—however, the modeling approaches' improvements on SQuAD still do not transfer to CNN.

A.4 BiDAF++

We modify an AllenNLP (Gardner et al., 2018) reimplementation of the BiDAF++ Clark and Gardner (2018) model originally used in pair2vec (Joshi et al., 2019) for evaluation on SQuAD 2.0 (Rajpurkar et al., 2018).

A.5 MnemonicReader

We use an reimplementation of MnemonicReader (Hu et al., 2017; note the specific arXiv revision) released at github.com/HKUST-KnowComp/MnemonicReader. In particular, the reimplementation is of the vanilla MnemonicReader without reinforcement learning.

A.6 QANet

We use the reimplementation of QANet (Yu et al., 2018) found in AllenNLP (Gardner et al., 2018). This reimplementation was used as a baseline method for DROP (Dua et al., 2019).

A.7 FusionNet

We use an reimplementation of FusionNet (Chen et al., 2017) released at github.com/felixgwu/FastFusionNet. This reimplementation was used as a baseline in Wu et al. (2019). Drawing inspiration from DocumentReader, the FusionNet approach also uses external features from a part-of-speech tagger and named entity recognition system. As a result, we also run the models without these features to fairly compare to systems that do not use such external resources. We keep the hand-crafted term-frequency and token exact match features originally used in the FusionNet paper.

A.8 BERT (base, large, and wwm)

We use the HuggingFace Transformers (Wolf et al., 2020) library to fine-tune BERT (Devlin et al., 2019) on extractive question answering benchmarks. In particular, we use the base, uncased, BERT pre-trained model, the large, uncased, BERT pre-trained model, and the large, uncased, BERT model pre-trained with whole-word masking.

A.9 ALBERT (base and xxlarge)

We use the HuggingFace Transformers (Wolf et al., 2020) library to fine-tune ALBERT (Lan et al., 2020) on extractive question answering benchmarks. In particular, we use the base and xxlarge V1 ALBERT pre-trained models.

A.10 RoBERTa (base and large)

We use the HuggingFace Transformers (Wolf et al., 2020) library to fine-tune RoBERTa (Liu et al., 2019) on extractive question answering benchmarks. In particular, we use the base and large RoBERTa pre-trained models.

A.11 ELECTRA (base)

We use the HuggingFace Transformers (Wolf et al., 2020) library to fine-tune the ELECTRA base discriminator (Clark et al., 2020) on extractive question answering benchmarks.

A.12 SpanBERT (base and large)

We use the author-released codebase (github.com/facebookresearch/SpanBERT) to fine-tune SpanBERT (Joshi et al., 2020) on extractive question answering benchmarks. In particular, we use the base and large SpanBERT pre-trained models.

B Preprocessing Existing Benchmarks

B.1 Existing Human-Constructed Benchmarks

We use the MRQA NewsQA, MRQA DROP, and MRQA HotpotQA benchmarks exactly as released by the MRQA 2019 shared task (Fisch et al., 2019).

The passages in MRQA NaturalQuestions contain HTML entities (e.g., <P> and </P>). The tokenizers used in non-pretrained models frequently split these entities into separate tokens. For example, <P> may become <, P, and >. This is problematic because the entities are quite common in passages, and expanding them during tokenization drastically increases the passage lengths, which some non-pretrained modeling approaches cannot handle due to GPU memory efficiency.

HTML entities are tokenized like this because they contain non-alphanumeric characters. As a result, we normalize HTML entities by replacing the non-alphanumeric characters. For example, <P> becomes BPB, and </P> becomes EEPE. These tokens are correctly kept intact. It's possible that modeling approaches that use subword information will perform worse with these normalized HTML entities, but we empirically observe that this normalization does not have a measurable impact on model performance.

QAMR questions were originally collected at the sentence level, but we concatenate these sentences to reconstruct the original passages they were sourced from. We then pair these reconstructed passages with the original QAMR questions.

It's possible for questions to become unanswerable at the passage-level. One case of his happens when two sentences have the same question—we filter out such questions that are asked for multiple sentences in a reconstructed passage.

Questions can also become unanswerable if relations between entities change between sentences. For example, given the passage "Bill lived in California in 1920. Bill lived in Washington in 1921.", the question "Where did Bill live" is answerable within the context of a particular sentence, but not in the context of the entire passage. Manual examination of generated QAMR passages and questions suggests that this case is rather uncommon, but it may still introduce a small amount of noise into the benchmark.

B.2 Existing Cloze Benchmarks

To convert the CBT and CNN benchmarks to extractive format, we take the passages and question as-is. The answer span is designated as the first occurrence of the answer token in the passage.

To convert LAMBADA into extractive format, we follow the setup of Cheng and Erk (2020).

The ReCoRD benchmark is used as-is, since it includes span-level annotations of answer tokens in passages.

B.3 Existing Synthetic Benchmarks

We consider tasks 1, 2, 3, 4, 5, 11, 12, 13, 14, 15, 16. The other tasks cannot be converted to extractive format (e.g., they require "yes"/"no" answers that do not appear in passages).

To convert the tasks in the bAbI benchmark to extractive format, we take the passages and question as-is. While the bAbI benchmark does not provide character-level span annotations for answers, questions come with "supporting facts"—sentences in the passage that contain the answer. Thus, choose the first occurrence of the answer token in the supporting fact sentence as our answer span.

Some of the bAbI tasks, while useable in an extractive format in theory, cannot be trivially converted to the extractive format via the procedure above because the released benchmark's annotations do not

appear in the passage. For instance, consider Figure 9, which shows an example drawn from the training set of Task 15. The answer provided in the benchmark is "cat", although this token never appears in the passage—instead, "cats" does. In cases where the originally-labeled answer cannot be found in the supporting fact, but its pluralization is present, we use the pluralized answer as our answer span.

Passage: *Mice are afraid of cats. Gertrude is a mouse. Emily is a mouse. Wolves are afraid of sheep. Winona is a wolf. Jessica is a mouse. Cats are afraid of sheep. Sheep are afraid of cats.*

Question: *What is jessica afraid of?*
Answer: *cat*

Figure 9

C Examples From Existing Benchmarks

C.1 Examples From Existing Human-Constructed Benchmarks

Table 2 shows examples from the existing human-constructed benchmarks we study.

C.2 Examples From Existing Cloze Benchmarks

Table 3 shows examples from the existing cloze benchmarks we study.

C.3 Examples From Existing Synthetic Benchmarks

Table 4 shows examples from the existing synthetic benchmarks we study. The contents of this table are reproduced from Weston et al. (2016).

Benchmark	**Passage** (some parts shortened with ...)	Question	Answer
MRQA NewsQA	(CNET) – When Facebook Chief Executive Mark Zuckerberg recently announced a "Like" button that publishers could place on their Web pages, he predicted it would make the Web smarter and "more social". What Zuckerberg didn't point out is that widespread use of the Like button allows Facebook to track people as they switch from CNN.com to Yelp.com to ESPN.com, all of which are sites that have said they will implement the feature...	What does the like button allow?	Facebook to track people
MRQA NaturalQuestions	BPB A shooting schedule is a project plan of each day 's shooting for a film production . It is normally created and managed by the assistant director , who reports to the production manager managing the production schedule . Both schedules represent a timeline stating where and when production resources are used . EEPE	who 's job is it to schedule each day 's shooting	assistant director
MRQA DROP	Coming off their win over the Chargers, the Bills flew to Dolphin Stadium for a Week 8 AFC East duel with the Miami Dolphins. In the first quarter, Buffalo trailed early as Dolphins QB Chad Pennington completed a 2-yard TD pass to TE Anthony Fasano. The Bills responded with kicker Rian Lindell getting a 19-yard field goal. In the second quarter, Buffalo took the lead as Lindell got a 43-yard and a 47-yard field goal...	Which team allowed the most first half points?	Dolphins
MRQA HotpotQA	[PAR] [TLE] John M. Brown [SEP] John Mifflin Brown (September 8, 1817 – March 16, 1893) was a bishop in the African Methodist Episcopal (AME) church. He was a leader in the underground railroad. He helped open a number of churches and schools, including the Payne Institute which became Allen University in Columbia, South Carolina and Paul Quinn College in Waco, Texas. He was also an early principal of Union Seminary which became Wilberforce University [PAR] [TLE] Waco, Texas [SEP] Waco ( ) is a city which is the county seat of McLennan County, Texas, United States. It is situated along the Brazos River and I-35, halfway between Dallas and Austin. The city had a 2010 population of 124,805, making it the 22nd-most populous city in the state. The US Census 2016 population estimate is 134,432 The Waco Metropolitan Statistical Area consists of McLennan and Falls Counties, which had a 2010 population of 234,906. Falls County was added to the Waco MSA in 2013. The US Census 2016 population estimate for the Waco MSA is 265,207.	What city is the home to Paul Quinn College and sets on the Brazos River between Dallas and Austin?	Waco, Texas
QAMR	An additional problem to face the empire came as a result of the involvement of Emperor Maurice -LRB- r. 582 – 602 -RRB- in Persian politics when he intervened in a succession dispute . This led to a period of peace , but when Maurice was overthrown , the Persians invaded and during the reign of Emperor Heraclius -LRB- r. 610 – 641 -RRB- controlled large chunks of the empire , including Egypt , Syria , and Anatolia until Heraclius ' successful counterattack . In 628 the empire secured a peace treaty and recovered all of its lost territories .	Whose politics did the empire get involved with?	Persian

Table 2: Example passages, questions, and answers from the existing human-constructed benchmarks we study.

Benchmark	**Passage** (some parts shortened with ...)	**Question**	**Answer**
Children's Book Test (Common Nouns)	... Lady Latifa argued and urged her wishes , but in vain ; the prince was not to be moved . Then she called to the cupbearers for new wine , for she thought that when his head was hot with it he might consent to stay . The pure , clear wine was brought ; she filled a cup and gave to him . He said : ' O most enchanting sweetheart ! it is the rule for the host to drink first and then the guest . '	So to make him lose his head , she drained the XXXXX ; then filled it again and gave him .	cup
Children's Book Test (Named Entities)	... At last , however , the Sunball became aware how sad Letiko was . ... Then he sent them away , and called two hares to him , and said : ' Will you take Letiko home to her mother ? ' ' Yes , why not ? ' ' What will you eat and drink if you should become hungry and thirsty by the way ? ' ' We will eat grass and drink from streamlets . ' ' Then take her , and bring her home . '	Then the hares set out , taking XXXXX with them , and because it was a long way to her home they became hungry by the way .	Letiko
LAMBADA	sorry 's not going to win me my game tomorrow . my racket is . i ca n't believe i let you take it out of here in the first place ! " " but , dad , i 'm sure you made mistakes when you were a hippie teenager ! " " and i paid for them !	like you 're going to pay for my	racket
CNN	(@entity0) you 'll see some familiar faces in the @entity1 . @entity2 beat @entity3 66 - 52 on sunday , giving @entity4 ' coach @entity5 his 12th trip to the semifinals of the @entity6 men 's basketball tournament . @entity7 and @entity8 each scored 16 to help @entity2 win the @entity9 . @entity3 , led by 16 points from @entity10 , was hoping to earn its first trip to the @entity1 . here 's how the @entity1 , to be played in @entity11 , has shaped up : next saturday , @entity2 will face @entity12 in the first semifinal . in the next game , top seed @entity13 will battle @entity14	the @entity1 matchups : @placeholder vs. @entity12 and @entity13 vs. @entity14	@entity2
ReCoRD	Secretary of State Hillary Clinton on Monday tried to douse a political firestorm over the deadly assault on a U.S. diplomatic mission in Libya, saying she's responsible for the security of American diplomatic outposts. "I take responsibility," Clinton told CNN in an interview while on a visit to Peru. "I'm in charge of the State Department's 60,000-plus people all over the world, 275 posts. The president and the vice president wouldn't be knowledgeable about specific decisions that are made by security professionals. They're the ones who weigh all of the threats and the risks and the needs and make a considered decision." @highlight "What I want to avoid is some kind of political gotcha or blame game," Clinton says @highlight "I take this very personally," she says @highlight Diplomats need security but "can't hang out behind walls," she adds	Clinton also described a desperate scene in the @placeholder during the hours of the attack, as staff tried to find out what had happened.	State Department

Table 3: Example passages, questions, and answers from the existing cloze benchmarks we study.

Benchmark	Passage	Question	Answer
bAbI Task 1 (Single Supporting Fact)	Mary went to the bathroom. John moved to the hallway. Mary travelled to the office.	Where is Mary?	office
bAbI Task 2 (Two Supporting Facts)	John is in the playground. John picked up the football. Bob went to the kitchen.	Where is the football?	playground
bAbI Task 3 (Three Supporting Facts)	John picked up the apple. John went to the office. John went to the kitchen. John dropped the apple.	Where was the apple before the kitchen?	office
bAbI Task 4 (Two Argument Relations)	The office is north of the bedroom. The bedroom is north of the bathroom. The kitchen is west of the garden.	What is north of the bedroom?	office
bAbI Task 5 (Three Argument Relations)	Mary gave the cake to Fred. Fred gave the cake to Bill. Jeff was given the milk by Bill.	Who did Fred give the cake to?	Bill
bAbI Task 11 (Basic Coreference)	Daniel was in the kitchen. Then he went to the studio. Sandra was in the office.	Where is Daniel?	studio
bAbI Task 12 (Conjunction)	Mary and Jeff went to the kitchen. Then Jeff went to the park.	Where is Jeff?	park
bAbI Task 13 (Compound Coreference)	Daniel and Sandra journeyed to the office. Then they went to the garden. Sandra and John travelled to the kitchen. After that they moved to the hallway.	Where is Daniel?	garden
bAbI Task 14 (Time Reasoning)	In the afternoon Julie went to the park. Yesterday Julie was at school. Julie went to the cinema this evening.	Where did Julie go after the park?	cinema
bAbI Task 15 (Basic Deduction)	Sheep are afraid of wolves. Cats are afraid of dogs. Mice are afraid of cats. Gertrude is a sheep.	What is Gertrude afraid of?	wolves
bAbI Task 16 (Basic Induction)	Lily is a swan. Lily is white. Bernhard is green. Greg is a swan.	What color is Greg?	white

Table 4: Example passages, questions, and answers from the existing synthetic benchmarks we study.

D Full Results on Existing Benchmarks

D.1 Full Results on Existing Human-Constructed Benchmarks

Table 5 and Table 6 show the performance of each modeling approach on each existing human-constructed benchmark.

	MRQA NewsQA	MRQA NaturalQuestions	MRQA DROP
RaSoR	44.68	60.02	51.30
BiDAF	43.49	58.43	51.36
DocumentReader	46.30	59.08	54.96
DocumentReader (no external features)	46.32	59.39	54.69
BiDAF++	46.53	60.23	55.16
MnemonicReader	48.43	61.53	57.02
MnemonicReader (no external features)	48.01	61.80	57.35
QANet	47.03	61.74	54.56
FusionNet	49.00	59.62	57.82
FusionNet (no external features)	48.88	59.54	57.95
BERT (base, uncased)	52.61	67.16	52.63
BERT (large, uncased)	54.99	69.38	61.54
BERT (large, uncased, whole-word masking)	57.86	71.67	71.66
ALBERT (base, V1)	53.25	67.37	61.21
ALBERT (xxlarge, V1)	61.16	72.95	78.64
RoBERTa (base)	56.62	68.28	64.54
RoBERTa (large)	59.14	72.06	74.12
ELECTRA (base)	57.60	70.23	69.00
SpanBERT (base)	55.60	69.51	63.74
SpanBERT (large)	59.09	72.13	75.05

Table 5: Performance of modeling approaches when evaluated on MRQA NewsQA, MRQA NaturalQuestions and MRQA DROP.

	MRQA HotpotQA	QAMR
RaSoR	51.35	51.56
BiDAF	50.94	51.84
DocumentReader	52.74	56.00
DocumentReader (no external features)	52.18	54.14
BiDAF++	53.86	54.69
MnemonicReader	56.13	58.07
MnemonicReader (no external features)	55.60	56.92
QANet	54.16	53.31
FusionNet	57.69	59.14
FusionNet (no external features)	57.38	56.91
BERT (base, uncased)	59.53	64.36
BERT (large, uncased)	61.63	67.51
BERT (large, uncased, whole-word masking)	65.02	71.03
ALBERT (base, V1)	61.65	66.30
ALBERT (xxlarge, V1)	68.17	74.15
RoBERTa (base)	61.19	67.16
RoBERTa (large)	64.58	71.44
ELECTRA (base)	62.58	68.16
SpanBERT (base)	63.89	68.70
SpanBERT (large)	66.60	71.46

Table 6: Performance of modeling approaches when evaluated on MRQA HotpotQA and QAMR.

D.2 Full Results on Existing Cloze Benchmarks

Table 7 and Table 8 show the performance of each modeling approach on each existing cloze benchmark.

	CBT (CN)	CBT (NE)	LAMBADA
RaSoR	53.00	69.85	71.95
BiDAF	52.45	72.75	70.29
DocumentReader	56.55	73.85	74.42
DocumentReader (no external features)	57.15	74.60	74.08
BiDAF++	58.40	77.15	71.95
MnemonicReader	61.45	78.80	74.57
MnemonicReader (no external features)	61.20	77.90	74.55
QANet	57.65	76.95	74.89
FusionNet	65.05	80.25	76.83
FusionNet (no external features)	64.85	79.85	76.92
BERT (base, uncased)	72.40	82.45	84.13
BERT (large, uncased)	76.65	84.55	86.83
BERT (large, uncased, whole-word masking)	79.90	86.90	91.23
ALBERT (base, V1)	70.75	82.70	82.14
ALBERT (xxlarge, V1)	86.90	90.70	94.53
RoBERTa (base)	75.70	84.90	86.48
RoBERTa (large)	82.45	88.60	92.27
ELECTRA (base)	74.20	84.40	86.40
SpanBERT (base)	75.90	85.50	87.10
SpanBERT (large)	80.75	88.80	91.65

Table 7: Performance of modeling approaches when evaluated on CBT (CN), CBT (NE) and LAMBADA.

	CNN (100K Examples)	ReCoRD
RaSoR	74.59	32.97
BiDAF	75.59	30.88
DocumentReader	72.66	29.97
DocumentReader (no external features)	72.38	29.52
BiDAF++	79.20	34.93
MnemonicReader	79.46	39.01
MnemonicReader (no external features)	78.95	37.87
QANet	79.00	33.46
FusionNet	79.05	30.89
FusionNet (no external features)	78.80	28.91
BERT (base, uncased)	79.74	58.45
BERT (large, uncased)	82.54	67.18
BERT (large, uncased, whole-word masking)	82.72	72.85
ALBERT (base, V1)	79.33	56.54
ALBERT (xxlarge, V1)	86.03	81.87
RoBERTa (base)	82.26	68.88
RoBERTa (large)	86.77	77.63
ELECTRA (base)	82.08	69.61
SpanBERT (base)	83.31	69.23
SpanBERT (large)	84.81	77.72

Table 8: Performance of modeling approaches when evaluated on CNN (100K Examples) and ReCoRD.

D.3 Full Results on Existing Synthetic Benchmarks

Table 9 and Table 10 and Table 11 show the performance of each modeling approach on each existing of the bAbI tasks (900 training examples).

	bAbI QA #1	bAbI QA #2	bAbI QA #3	bAbI QA #4
RaSoR	100.0	60.0	71.0	81.0
BiDAF	100.0	42.0	53.0	83.0
DocumentReader	100.0	63.0	70.0	100.0
DocumentReader (no external features)	100.0	76.0	93.0	100.0
BiDAF++	100.0	100.0	100.0	78.0
MnemonicReader	100.0	44.0	71.0	100.0
MnemonicReader (no external features)	100.0	100.0	74.0	100.0
QANet	100.0	42.0	39.0	85.0
FusionNet	100.0	84.0	77.0	100.0
FusionNet (no external features)	100.0	100.0	70.0	100.0
BERT (base, uncased)	100.0	80.0	49.0	81.0
BERT (large, uncased)	100.0	63.0	63.0	79.0
BERT (large, uncased, whole-word masking)	100.0	98.0	98.0	91.0
ALBERT (base, V1)	100.0	86.0	85.0	85.0
ALBERT (xxlarge, V1)	100.0	100.0	100.0	100.0
RoBERTa (base)	100.0	73.0	54.0	64.0
RoBERTa (large)	100.0	39.0	53.0	87.0
ELECTRA (base)	100.0	86.0	64.0	100.0
SpanBERT (base)	57.0	9.0	22.0	60.0
SpanBERT (large)	61.0	38.0	9.0	60.0

Table 9: Performance of modeling approaches when evaluated on bAbI QA #1, bAbI QA #2, bAbI QA #3 and bAbI QA #4.

	bAbI QA #5	bAbI QA #11	bAbI QA #12	bAbI QA #13
RaSoR	98.0	100.00	100.0	100.0
BiDAF	95.0	78.00	100.0	95.0
DocumentReader	96.0	100.00	100.0	100.0
DocumentReader (no external features)	97.0	100.00	100.0	100.0
BiDAF++	96.0	100.00	100.0	95.0
MnemonicReader	95.0	100.00	100.0	95.0
MnemonicReader (no external features)	95.0	100.00	100.0	100.0
QANet	95.0	100.00	100.0	95.0
FusionNet	98.0	100.00	100.0	100.0
FusionNet (no external features)	98.0	100.00	100.0	100.0
BERT (base, uncased)	95.0	100.00	100.0	97.0
BERT (large, uncased)	95.0	100.00	100.0	100.0
BERT (large, uncased, whole-word masking)	96.0	100.00	100.0	100.0
ALBERT (base, V1)	95.0	100.00	100.0	100.0
ALBERT (xxlarge, V1)	99.0	100.00	100.0	100.0
RoBERTa (base)	95.0	98.99	89.0	95.0
RoBERTa (large)	98.0	100.00	100.0	95.0
ELECTRA (base)	95.0	100.00	100.0	97.0
SpanBERT (base)	36.0	74.75	75.0	95.0
SpanBERT (large)	43.0	81.82	77.0	95.0

Table 10: Performance of modeling approaches when evaluated on bAbI QA #5, bAbI QA #11, bAbI QA #12 and bAbI QA #13.

Figure 10 shows how well the bAbI tasks (9000) training examples concur with SQuAD.

Table 12 and Table 13 and Table 14 show the performance of each modeling approach on each existing of the bAbI tasks (9000 training examples).

	bAbI QA #14	bAbI QA #15	bAbI QA #16
RaSoR	97.0	73.00	64.0
BiDAF	95.0	66.00	61.0
DocumentReader	96.0	68.00	63.0
DocumentReader (no external features)	99.0	68.00	64.0
BiDAF++	92.0	65.00	61.0
MnemonicReader	99.0	63.00	65.0
MnemonicReader (no external features)	99.0	67.00	65.0
QANet	62.0	64.00	58.0
FusionNet	100.0	69.00	64.0
FusionNet (no external features)	99.0	100.00	64.0
BERT (base, uncased)	84.0	60.56	50.0
BERT (large, uncased)	88.0	56.34	52.0
BERT (large, uncased, whole-word masking)	96.0	100.00	62.0
ALBERT (base, V1)	78.0	60.56	80.0
ALBERT (xxlarge, V1)	100.0	100.00	100.0
RoBERTa (base)	81.0	61.97	47.0
RoBERTa (large)	77.0	100.00	44.0
ELECTRA (base)	87.0	100.00	47.0
SpanBERT (base)	37.0	46.48	36.0
SpanBERT (large)	37.0	59.15	49.0

Table 11: Performance of modeling approaches when evaluated on bAbI QA #14, bAbI QA #15 and bAbI QA #16.

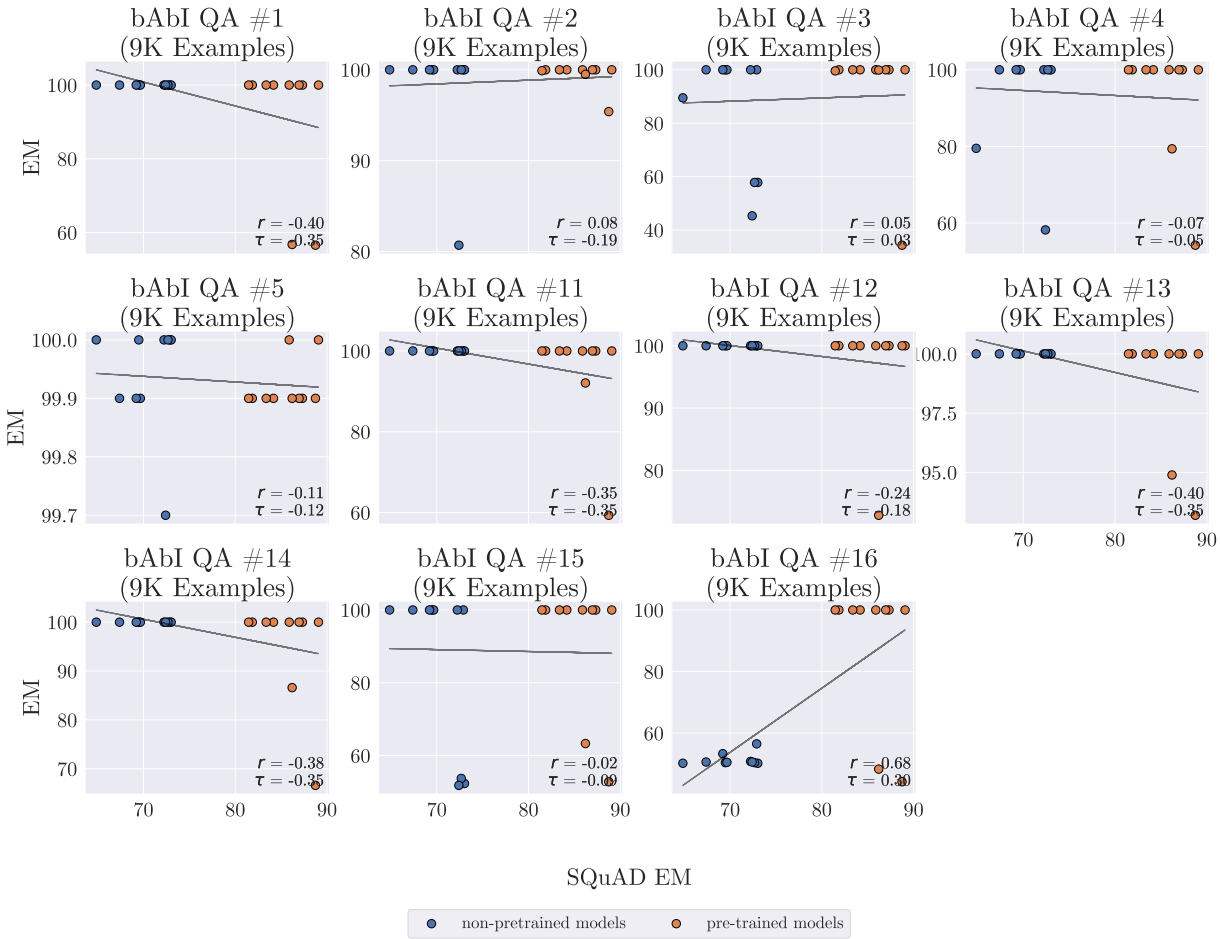

Figure 10: Many modeling approaches perform perfectly on bAbI tasks when training on 9,000 examples, limiting their ability to recapitulate historical modeling progress on SQuAD.

	bAbI QA #1 (9K)	bAbI QA #2 (9K)	bAbI QA #3 (9K)	bAbI QA #4 (9K)
RaSoR	100.00	100.0	89.5	79.50
BiDAF	100.00	100.0	100.0	100.00
DocumentReader	100.00	100.0	100.0	100.00
DocumentReader (no external features)	100.00	100.0	100.0	100.00
BiDAF++	100.00	100.0	100.0	100.00
MnemonicReader	100.00	100.0	57.8	100.00
MnemonicReader (no external features)	100.00	100.0	57.8	100.00
QANet	100.00	80.7	45.3	58.20
FusionNet	100.00	100.0	100.0	100.00
FusionNet (no external features)	100.00	100.0	100.0	100.00
BERT (base, uncased)	100.00	99.9	99.6	100.00
BERT (large, uncased)	100.00	100.0	100.0	100.00
BERT (large, uncased, whole-word masking)	100.00	100.0	100.0	100.00
ALBERT (base, V1)	100.00	100.0	100.0	100.00
ALBERT (xxlarge, V1)	100.00	100.0	100.0	100.00
RoBERTa (base)	100.00	100.0	100.0	100.00
RoBERTa (large)	100.00	100.0	100.0	100.00
ELECTRA (base)	100.00	100.0	100.0	100.00
SpanBERT (base)	56.77	99.5	99.9	79.37
SpanBERT (large)	56.57	95.4	34.3	54.21

Table 12: Performance of modeling approaches when evaluated on bAbI QA #1 (9K Examples), bAbI QA #2 (9K Examples), bAbI QA #3 (9K Examples) and bAbI QA #4 (9K Examples).

	bAbI QA #5 (9K)	bAbI QA #11 (9K)	bAbI QA #12 (9K)	bAbI QA #13 (9K)
RaSoR	100.0	100.00	100.0	100.00
BiDAF	99.9	100.00	100.0	100.00
DocumentReader	99.9	100.00	100.0	100.00
DocumentReader (no external features)	99.9	100.00	100.0	100.00
BiDAF++	100.0	100.00	100.0	100.00
MnemonicReader	100.0	100.00	100.0	100.00
MnemonicReader (no external features)	100.0	100.00	100.0	100.00
QANet	99.7	100.00	100.0	100.00
FusionNet	100.0	100.00	100.0	100.00
FusionNet (no external features)	100.0	100.00	100.0	100.00
BERT (base, uncased)	99.9	100.00	100.0	100.00
BERT (large, uncased)	99.9	100.00	100.0	100.00
BERT (large, uncased, whole-word masking)	99.9	100.00	100.0	100.00
ALBERT (base, V1)	99.9	100.00	100.0	100.00
ALBERT (xxlarge, V1)	100.0	100.00	100.0	100.00
RoBERTa (base)	99.9	100.00	100.0	100.00
RoBERTa (large)	99.9	100.00	100.0	100.00
ELECTRA (base)	100.0	100.00	100.0	100.00
SpanBERT (base)	99.9	92.08	72.8	94.89
SpanBERT (large)	99.9	59.32	100.0	93.19

Table 13: Performance of modeling approaches when evaluated on bAbI QA #5 (9K Examples), bAbI QA #11 (9K Examples), bAbI QA #12 (9K Examples) and bAbI QA #13 (9K Examples).

	bAbI QA #14 (9K)	bAbI QA #15 (9K)	bAbI QA #16 (9K)
RaSoR	100.0	100.00	50.2
BiDAF	100.0	100.00	50.6
DocumentReader	100.0	100.00	50.5
DocumentReader (no external features)	100.0	100.00	53.3
BiDAF++	100.0	100.00	50.4
MnemonicReader	100.0	52.30	50.2
MnemonicReader (no external features)	100.0	53.70	50.4
QANet	100.0	51.80	50.6
FusionNet	100.0	100.00	56.5
FusionNet (no external features)	100.0	100.00	50.8
BERT (base, uncased)	100.0	100.00	100.0
BERT (large, uncased)	100.0	100.00	100.0
BERT (large, uncased, whole-word masking)	100.0	100.00	100.0
ALBERT (base, V1)	100.0	100.00	100.0
ALBERT (xxlarge, V1)	100.0	100.00	100.0
RoBERTa (base)	100.0	100.00	100.0
RoBERTa (large)	100.0	100.00	100.0
ELECTRA (base)	100.0	100.00	100.0
SpanBERT (base)	86.6	63.30	48.3
SpanBERT (large)	66.6	52.78	44.2

Table 14: Performance of modeling approaches when evaluated on bAbI QA #14 (9K Examples), bAbI QA #15 (9K Examples) and bAbI QA #16 (9K Examples).

E FuzzySyntheticQA Construction Details

Figure 11 provides an overview of the construction of FuzzySyntheticQA.

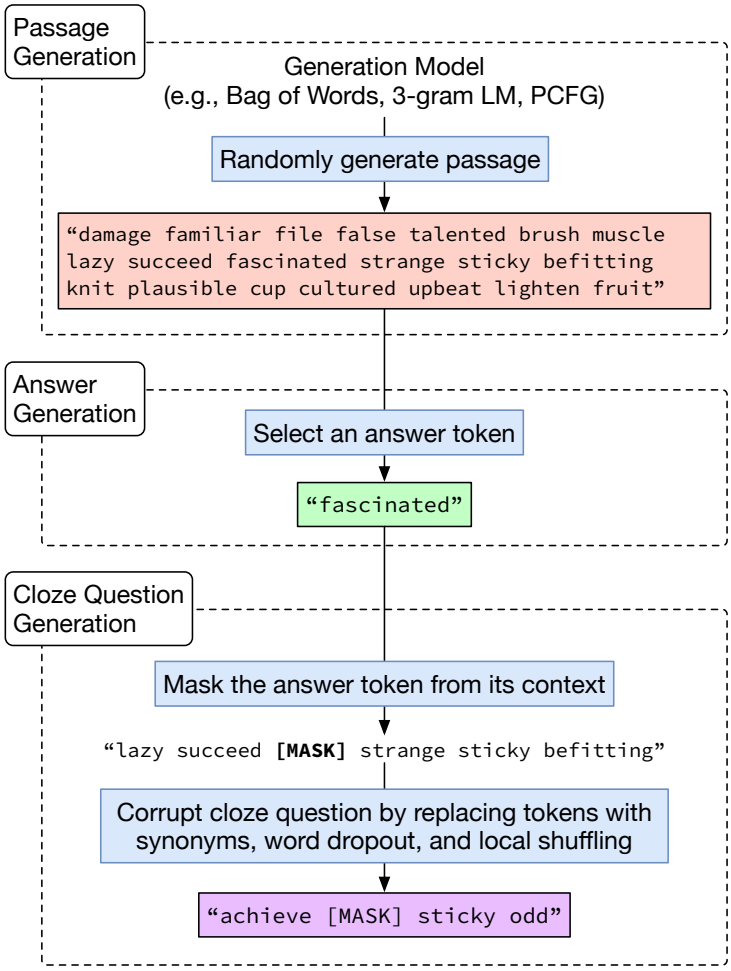

Figure 11: Constructing a FuzzySyntheticQA example by generating a passage, answer, and cloze question.

To efficiently replace tokens with related tokens, we consider each token's 100 *approximate* nearest neighbors as replacement candidates. In particular, we use Annoy (Bernhardsson and the Annoy development team, 2020) to perform the approximate nearest neighboor look-ups. Similarities are derived from the Euclidean distance of normalized vectors between two tokens.

F Full Results on FuzzySyntheticQA

Figure 12 shows that changing the passage generation method in FuzzySyntheticQA has a minimal effect on concurrence. We experiment with generating passages from a 3-gram language model, a probabilistic context-free grammar, a large neural language model (GPT-2 1.5B; Radford et al., 2019), and by taking real Wikipedia paragraphs.

The 3-gram language model is trained with maximum likelihood estimation on WikiText-103 (Merity et al., 2017). The PCFG is trained with maximum likelihood estimation on the Penn Treebank (Marcus et al., 1993). Lastly, we take GPT-2 1.5B generations from the officially-released output samples (github.com/openai/gpt-2-output-dataset; generated with top-k truncated sampling with k = 40).

Table 15 and Table 16 show the performance of each modeling approach on each of our constructed synthetic fuzzy pattern-matching benchmarks.

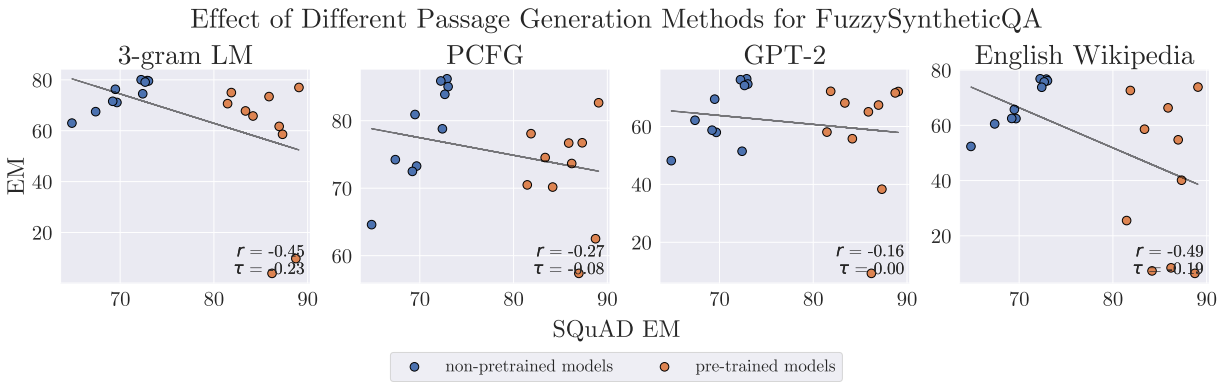

Figure 12: Even with progressively more natural passages, FuzzySyntheticQA continues to have low overall concurrence with SQuAD—this low concurrence is not trivially caused by the lack of natural passages, and simply making our passages more closely resemble natural language will not yield high concurrence.

	Synthetic Fuzzy Pattern-Matching	3-gram LM Synthetic Fuzzy Pattern-Matching	PCFG Synthetic Fuzzy Pattern-Matching
RaSoR	37.01	63.00	64.60
BiDAF	38.62	67.50	74.23
DocumentReader	49.32	71.11	73.28
DocumentReader (no external features)	49.24	71.57	72.49
BiDAF++	56.89	76.30	80.92
MnemonicReader	61.50	79.56	85.05
MnemonicReader (no external features)	61.24	79.13	83.91
QANet	59.60	74.53	78.80
FusionNet	64.71	79.72	86.21
FusionNet (no external features)	63.80	80.05	85.89
BERT (base, uncased)	4.51	70.65	70.49
BERT (large, uncased)	40.11	65.79	70.17
BERT (large, uncased, whole-word masking)	0.70	58.60	76.73
ALBERT (base, V1)	44.28	75.00	78.08
ALBERT (xxlarge, V1)	53.79	77.01	82.66
RoBERTa (base)	44.92	67.78	74.54
RoBERTa (large)	0.49	61.71	57.38
ELECTRA (base)	44.85	73.42	76.69
SpanBERT (base)	0.74	3.92	73.66
SpanBERT (large)	0.40	9.74	62.51

Table 15: Performance of modeling approaches when evaluated on Synthetic Fuzzy Pattern-Matching, 3-gram LM Synthetic Fuzzy Pattern-Matching and PCFG Synthetic Fuzzy Pattern-Matching.

	GPT-2 Synthetic Fuzzy Pattern-Matching	English Wikipedia Synthetic Fuzzy Pattern-Matching
RaSoR	48.20	52.37
BiDAF	62.16	60.52
DocumentReader	57.97	62.45
DocumentReader (no external features)	58.73	62.50
BiDAF++	69.45	65.74
MnemonicReader	74.67	76.15
MnemonicReader (no external features)	74.18	75.71
QANet	51.45	73.79
FusionNet	76.48	76.73
FusionNet (no external features)	76.17	76.85
BERT (base, uncased)	58.07	25.52
BERT (large, uncased)	55.78	7.29
BERT (large, uncased, whole-word masking)	38.34	40.13
ALBERT (base, V1)	72.16	72.62
ALBERT (xxlarge, V1)	72.09	73.86
RoBERTa (base)	68.14	58.60
RoBERTa (large)	67.41	54.76
ELECTRA (base)	65.07	66.33
SpanBERT (base)	9.26	8.40
SpanBERT (large)	71.61	6.40

Table 16: Performance of modeling approaches when evaluated on GPT-2 Synthetic Fuzzy Pattern-Matching and English Wikipedia Synthetic Fuzzy Pattern-Matching.

G WikidataSyntheticQA Construction Details

Figure 13 summarizes the data generation procedure for WikidataSyntheticQA.

Inverses of Properties. Some of our generated questions use the inverse relationships between two properties. To obtain the inverse relationship for a given property, we first retrieve its list of property constraints by using Wikidata property P2302 (property constraint). If Q21510855 (inverse constraint) is present, we then retrieve the corresponding property of this inverse relationship. If the inverse constraint is not present, we check the corresponding property of P7087 (inverse label item), which outputs the item with a label of the inverse relationship of the property.

Entity Hyponyms. Some of our generated questions replace entities with their hyponyms. To obtain the hyponyms for a given entity, we retrieve any object entities of the P31 (instance of) and P279 (subclass of) properties.

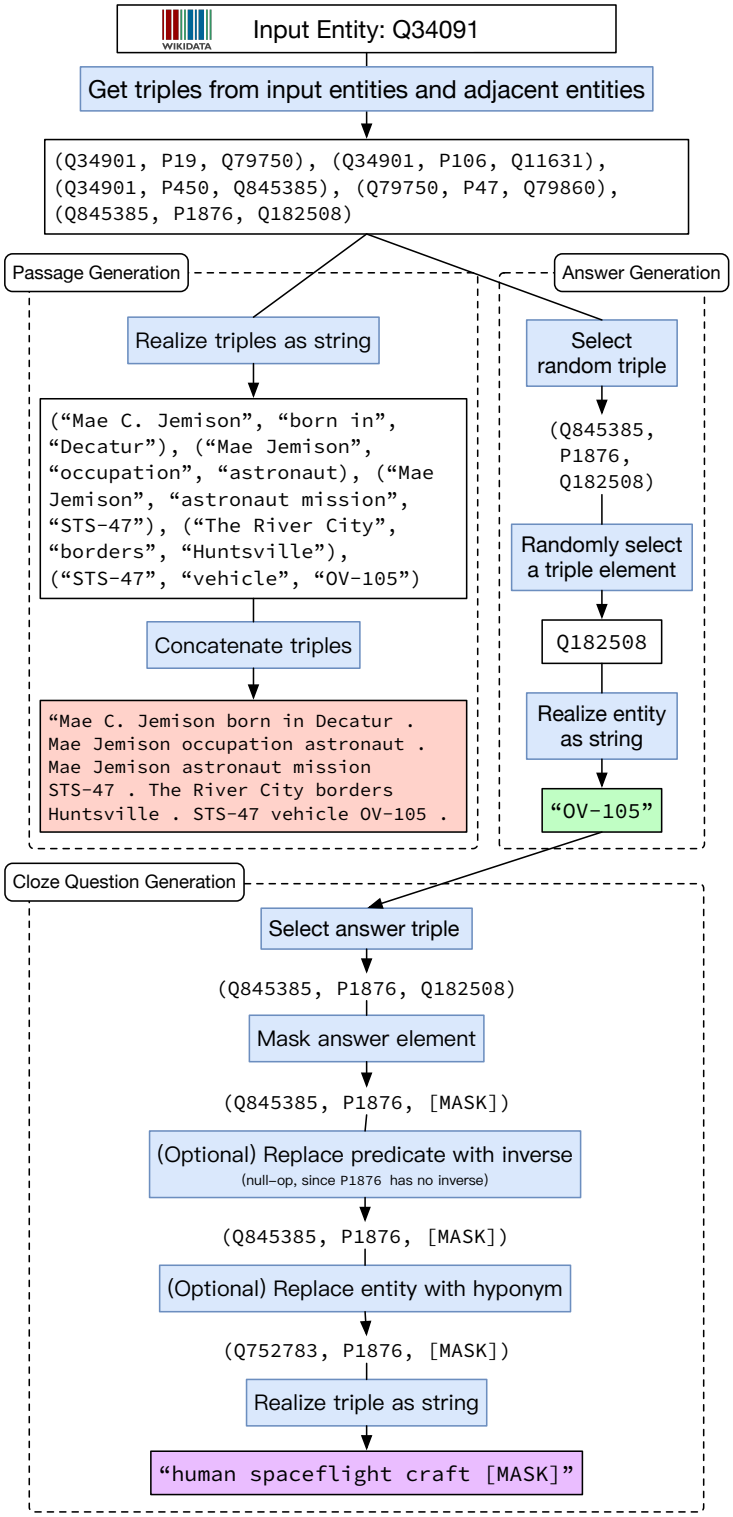

Figure 13: Constructing a WikidataSyntheticQA example by generating a passage, answer, and cloze question.

H Full Results on WikidataSyntheticQA

Table 17 shows the performance of each modeling approach on WikidataSyntheticQA.

	Synthetic Wikidata
RaSoR	63.67
BiDAF	68.69
DocumentReader	67.66
DocumentReader (no external features)	68.03
BiDAF++	70.43
MnemonicReader	75.04
MnemonicReader (no external features)	74.31
QANet	73.12
FusionNet	74.52
FusionNet (no external features)	73.90
BERT (base, uncased)	73.68
BERT (large, uncased)	78.01
BERT (large, uncased, whole-word masking)	81.56
ALBERT (base, V1)	77.23
ALBERT (xxlarge, V1)	86.29
RoBERTa (base)	77.75
RoBERTa (large)	82.79
ELECTRA (base)	76.86
SpanBERT (base)	78.50
SpanBERT (large)	84.26

Table 17: Performance of modeling approaches when evaluated on Synthetic Wikidata.

I Full Results on Subsampled SQuAD

Table 18 and Table 19 show the performance of each modeling approach on subsamples of the SQuAD benchmark.

	SQuAD 1.1		
	All	1K Examples	10K Examples
RaSoR	64.86	15.52	49.44
BiDAF	67.39	7.96	48.54
DocumentReader	69.66	34.66	56.42
DocumentReader (no external features)	69.21	30.69	54.82
BiDAF++	69.49	18.62	57.48
MnemonicReader	73.02	30.67	58.91
MnemonicReader (no external features)	72.67	29.46	57.79
QANet	72.41	7.18	48.15
FusionNet	72.90	37.52	59.97
FusionNet (no external features)	72.24	35.55	58.69
BERT (base, uncased)	81.46	31.80	70.34
BERT (large, uncased)	84.17	49.08	75.47
BERT (large, uncased, whole-word masking)	87.32	69.19	81.78
ALBERT (base, V1)	81.86	57.57	74.55
ALBERT (xxlarge, V1)	89.07	76.36	86.19
RoBERTa (base)	83.37	55.01	77.30
RoBERTa (large)	86.96	62.64	82.56
ELECTRA (base)	85.88	62.05	78.31
SpanBERT (base)	86.20	65.80	80.72
SpanBERT (large)	88.74	75.00	85.06

Table 18: Performance of modeling approaches when evaluated on SQuAD, SQuAD (1K Examples) and SQuAD (10K Examples).

	SQuAD 1.1		
	20K Examples	40K Examples	60K Examples
RaSoR	55.13	60.37	62.95
BiDAF	57.29	62.35	65.25
DocumentReader	61.84	65.45	68.27
DocumentReader (no external features)	59.66	64.47	67.09
BiDAF++	62.25	66.42	68.62
MnemonicReader	64.74	69.09	70.86
MnemonicReader (no external features)	63.71	68.65	70.32
QANet	61.02	66.55	69.74
FusionNet	64.74	69.14	70.98
FusionNet (no external features)	63.28	67.98	69.93
BERT (base, uncased)	74.84	78.24	80.05
BERT (large, uncased)	79.27	81.83	83.25
BERT (large, uncased, whole-word masking)	84.47	85.78	86.75
ALBERT (base, V1)	77.05	79.95	81.02
ALBERT (xxlarge, V1)	86.91	88.02	88.63
RoBERTa (base)	79.56	81.62	82.37
RoBERTa (large)	84.26	86.37	87.18
ELECTRA (base)	81.75	83.95	85.01
SpanBERT (base)	82.54	84.17	85.39
SpanBERT (large)	86.21	87.33	87.82

Table 19: Performance of modeling approaches when evaluated on SQuAD (20K Examples), SQUAD (40K Examples) and SQuAD (60K Examples).