Cheap and Good? Simple and Effective Data Augmentation for Low Resource Machine Reading

Hoang Van
University of Arizona
Tucson, AZ, USA
vnhh@arizona.edu

Vikas Yadav
IBM Research
Yorktown Heights, NY, USA
vikasy@ibm.com

Mihai Surdeanu
University of Arizona
Tucson, AZ, USA
msurdeanu@arizona.edu

ABSTRACT

We propose a simple and effective strategy for data augmentation for low-resource machine reading comprehension (MRC). Our approach first pretrains the answer extraction components of a MRC system on the augmented data that contains approximate context of the correct answers, before training it on the exact answer spans. The approximate context helps the QA method components in narrowing the location of the answers. We demonstrate that our simple strategy substantially improves both document retrieval and answer extraction performance by providing larger context of the answers and additional training data. In particular, our method significantly improves the performance of BERT based retriever (15.12%), and answer extractor (4.33% F1) on TechQA, a complex, low-resource MRC task. Further, our data augmentation strategy yields significant improvements of up to 3.9% exact match (EM) and 2.7% F1 for answer extraction on PolicyQA, another practical but moderate sized QA dataset that also contains long answer spans.

CCS CONCEPTS

• Information Systems → Question Answering.

KEYWORDS

Data augmentation, Question answering, Document retrieval

ACM Reference Format:

Hoang Van, Vikas Yadav, and Mihai Surdeanu. 2021. Cheap and Good? Simple and Effective Data Augmentation for Low Resource Machine Reading. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463099>

1 INTRODUCTION

Supervised neural question answering (QA) methods have achieved state-of-the-art performance on several datasets and benchmarks [3, 10, 13, 18]. However, their success is fueled by large annotated datasets, which are expensive to generate and not always available in low resource settings. Moreover, machine reading comprehension (MRC) QA with span selection from a large set of candidate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8037-9/21/07...\$15.00

<https://doi.org/10.1145/3404835.3463099>

Figure 1: Overview of our three key contributions: data augmentation (top), transfer learning (middle), and dynamic document retrieval (bottom). Top: we augment question-answer pairs with additional answer spans by moving the answer spans to left or right of the original spans by d characters. Middle: we train the MRC answer extraction (AE) component with original settings on the combined augmented + original data. During transfer learning, the models trained on the combined data continue tuning on the original (without augmentation) data. Bottom: at inference time, k candidate documents are retrieved by the data-augmented document retriever (trained using the scheme above) from the pool of N documents per question.

documents has been shown to require even more training data [4, 5, 7, 15].

In this work, we propose a simple and effective data augmentation method that tackles the shortage of training data for MRC. The intuition behind our idea is straightforward: rather than directly train a neural approach on the answer spans provided during training, we start by *pretraining* it to identify the rough context in which the answer appears. Then we use this pretrained neural model in two different ways. First, instead of starting from scratch, we initialize the answer extraction model with the weights from this model before training it. Second, at inference time, we use the pretrained model as an additional document retrieval component: rather than extracting answers from any document, we focus only on documents that contain contexts identified as likely to contain an answer.

The key contributions of our work are:

(1) We introduce a simple yet effective data augmentation method for MRC QA. Our approach generates additional training data by artificially moving the boundaries of answer spans in the training partition, and training a model on this data to identify the approximate context in which the answer is likely to appear.

(2) We implement two different strategies to take advantage of the above model. (a) First, we transfer its knowledge into the answer extraction model by initializing with its learned weights, and (b) secondly, at inference time, we use the above model (trained on the fuzzy answer boundaries) to rerank documents based on their likelihood to contain an answer context, and employ answer extraction just in the top documents with the highest scores.

(3) We demonstrate that utilizing approximate RC knowledge to both document retrieval and answer extraction results in considerable improvements in their performance on low resource RC tasks. In particular, after transferring the approximate RC knowledge, the BERT-based MRC model’s performance is improved by 3.46%, while the BERT-based document retriever’s ([2]) performance is improved by 4.33% (both in absolute F1 score) on TechQA, a real-world low-resource MRC dataset [2]. Further, our best model also outperforms the previous best document retrieval scores [2] by 15.12% in document retrieval accuracy (DRA) on TechQA. Similarly, our proposed approach also yields 3.9% EM and 2.7% F1 improvements for the BERT-based answer extractor on PolicyQA, another practical and challenging QA dataset that has long answer spans [1]. Our code and reproducibility checklist are publicly available.¹

2 APPROACH

Fig.1 highlights the three main components of our MRC approach: (a) data augmentation with fuzzy answer boundaries (top); (b) transfer learning into the answer extraction component (middle); and (c) improved document retrieval. We detail these modules next.

2.1 Data Augmentation

We pose the data augmentation problem [9, 11] as a positive example expansion problem. First, we sample which answers should be used for augmentation with probability p . Second, once a question-answer pair is selected for augmentation, we generate n additional fuzzy answer spans by moving the correct answer span to the left or right by d characters. Each generated answer span becomes a new positive training data point. For example, for the answer “*install gtk2 libraries*” in the sentence “*Libraries missing, install the gtk2 libraries (32 and 64 bit)*”, our data augmentation method generates the spans “*Libraries missing, install gtk2 libraries*” ($d = -19$) and “*install gtk2 libraries (32 and 64 bit)*” ($d = +16$) as additional answers (see Fig. 1, data augmentation). Both d and the number of augmented data points per training example are hyper parameters that we tune for the given tasks (see Section 3.2). Note that this augmented data is appended to the original training set.

Intuitively, a model trained on this artificial data (illustrated by the second BERT-MRC block in Figure 1) would be capable of identifying the *approximate* context in which an answer appears, even if it cannot locate the answer exactly [6].²

2.2 Transfer Learning

As shown in Figure 1, we first train the BERT-MRC model on the augmented training data consisting of the fuzzy answer spans. Then, we transfer this knowledge of approximate answer boundaries by continuing to further train our BERT-MRC model on the original answer spans (without augmentation). The resulting model (the third BERT-MRC block in Fig. 1) is the final model used for answer span prediction in the complete task. Note that, while we use BERT-MRC [16, 17] to train these answer extraction components, same idea could be investigated with other answer extraction methods.

Model	DRA@1	DRA@5	F1	Recall
Baselines: Original Settings				
1 BERT Base (OBB)	31.88	45.00	51.73	51.94
2 BERT Large (OBL)	58.13	64.38	55.31	57.81
Data Augmentation (DA)				
3 60% augmentation (aug_60)	46.25	51.25	53.52	53.67
4 80% augmentation (aug_80)	47.53*	53.00*	54.35*	54.73*
5 100% augmentation (aug_100)	46.33	48.75	53.56	52.76
DA+Transfer Learning				
6 aug_60 + Transfer	43.13	50.20	54.62	55.03
7 aug_80 + Transfer	46.87*	52.10*	55.19*	55.68*
8 aug_100 + Transfer	42.58	49.75	54.11	53.39
Document Retrieval + OBL				
9 aug_60 + OBL	69.38	75.63	59.13	59.48
10 aug_80 + OBL	73.25*	78.13*	59.64*	59.84*
11 aug_100 + OBL	71.25	73.13	58.58	58.89

Table 1: Document retrieval accuracies (DRAs) and answer extraction results (F1, Recall) on the development partition of TechQA. We use BERT Base (similar to OBB) for all data augmentation and transfer learning experiments. OBL is a baseline in [2]. ‘Data Augmentation’ shows results of models tuned with augmented data, starting with pretrained weights from BERT Base. ‘Data Augmentation + Transfer Learning’ lists results of data augmented models (DAM) after continuing to tune on the original (no augmentation) data. ‘Document Retrieval + OBL’ uses a hybrid approach where a model from the ‘Data Augmentation’ block is used to retrieve documents that are then fed to OBL for answer extraction. * indicates significant improvement over the respective baselines ($p < 0.02$).

Models	F1	EM
Baselines		
1 OBB pretrained on SQuAD (OBS)	57.2 [†]	27.5 [†]
Data Augmentation		
2 OBS + 3% augmentation (OBS_aug_3)	58.8	29.3
3 OBS + 4% augmentation (OBS_aug_4)	59.8*	30.8*
4 OBS + 5% augmentation (OBS_aug_5)	59.2	30.2
Data Augmentation + Transfer Learning		
5 OBS_aug_3 + Transfer	59.2	29.9
6 OBS_aug_4 + Transfer	59.9*	31.4*
7 OBS_aug_5 + Transfer	59.3	30.6

Table 2: Answer extraction results on PolicyQA development. * indicates significant improvement over the baseline (line 1) with $p < 0.03$. † indicates experiments from a different paper that we reproduced for a fair comparison. These results differ from the original paper [1], which is likely due to the different hardware and library versions used.

2.3 Document Retrieval

Most MRC datasets include a list of N candidate documents for each question [8]. All these documents are scanned for possible answers by typical MRC approaches [14]. Instead, as shown in Fig. 1, we use only documents that are more likely to contain good answer contexts. In particular, we employ the model trained on the augmented data (with fuzzy answer boundaries) to produce the top k answer spans with the highest scores, and keep only the documents associated with these answer spans. For example, on TechQA, by selecting the top documents containing the high score answer spans, our approach narrows down the number of candidate documents to an average of 6.41 per question³ from the initial pool of 50 documents. As shown in Tab.2, feeding a smaller but more relevant pool of candidate documents leads to substantial improvements in the AE performance.

¹<https://github.com/vanh17/techqa>

[†] Work performed while this author was a PhD student at the University of Arizona.

²This is probably how students work towards answering reading comprehension questions as well. See, e.g.: <https://tinyurl.com/khanrc>.

³We also observed that the maximum number of candidate documents is 10

3 EXPERIMENTAL SETUP

We evaluate our approach on two complex MRC tasks: TechQA [2]⁴ and PolicyQA [1]⁵. Both datasets contain questions from *low-resource domains*, and their questions are answered by *long text spans*, which makes these tasks more challenging for RC approaches [2]. Specifically, TechQA is a real-world dataset that requires extraction of long answer spans to industrial technical questions. TechQA questions and answers are substantially longer than those found in common datasets [12]. The average answer span and average question length in the training partition of this dataset are 48.1 and 52.1 tokens, respectively. Further, this task is relatively low resourced: the training partition contains only 600 questions; the development one contains 310. Each question is associated with 50 documents that are likely to contain the answer. Note that some questions (25% in training, 48% in development) are unanswerable i.e., they do not have an answer in the 50 documents provided. Similar to TechQA, PolicyQA is another real-world MRC task that requires the extraction of answers to privacy policy questions from long documents (average document length is 106 tokens). However, each PolicyQA question is associated with just 1 given document. But, similar to TechQA, the average answers spans in PolicyQA are also long: the average answer span and average question length in the training partition of this dataset are 11.2 and 13.3 tokens, respectively. Please note that we selected PolicyQA as the second dataset to demonstrate that our approach works with moderate resourced datasets⁶ as well.

For reproducibility, we use the default settings and hyper parameters recommended by the task creators [1, 2]. Through this, we aim to separate potential improvements of our approaches from those made with improved configurations. Due to computational constraints, we report transfer learning results using BERT Base (uncased) from the HuggingFace library⁷. Additionally, for TechQA, we combine our document retrieval strategy with the answer extraction component based on BERT Large (uncased whole word masking tuned on SQUAD [12]) that is provided by the organizers.

3.1 Evaluation Measures

Answer extraction: we directly followed the evaluation metrics proposed by the original task creators [1, 2]. Specifically, we used the following ancillary metrics to evaluate the answer extraction components: F1 (TechQA, PolicyQA) and exact match (EM) (PolicyQA). For each question, the F1 score is computed based on the character overlap between the predicted answer span and the correct answer span. We also report the Recall score on TechQA, similarly to the task organizers [2].

Document retrieval: we report the original evaluation metric, i.e., document retrieval accuracy (DRA), as proposed by the task organizers [2]⁸. In MRC tasks such as TechQA, DRA is calculated only over *answerable* questions. The score is 1 if the retrieval component retrieves the document containing the correct answer span, and 0

otherwise. We report the performance with DRA@1 and DRA@5, which score the top 1 and 5 retrieved documents, respectively.

3.2 Hyper parameter tuning

We tuned three hyper parameters for data augmentation: percentage of augmented training data points, span displacement size, and number of new spans generated per correct answer span. On TechQA, we tried augmenting 20, 40, 60, 80, 100 % of the questions from training data and found 80% of the training data as the best threshold. For each (randomly) selected question from the training data, we generated 6 augmented data points by moving the gold answer span boundaries. Specifically, we moved answer spans up to 15 characters to left and up to 20 characters to the right for TechQA. On PolicyQA, we tuned smaller movements of answer spans⁹ such as 5 characters on the left and 10 characters on the right of the gold answer spans. For each PolicyQA question, similar to TechQA, we generated 6 additional data points based on the scheme above. We tuned these three hyper parameters using DRA@5 (TechQA) and F1 (PolicyQA),¹⁰ with the aims of improving the ability to localize the context in which an answer appears, and reducing the likelihood of missing correct documents.¹¹

4 RESULTS AND DISCUSSION

We investigate the contribution of our ideas on document retrieval and AE performance as explained below.

4.1 Document Retrieval

Table 1 shows that data augmentation helps with the document retrieval task in TechQA. Our approaches improve DRA scores considerably compared to the BERT Base baseline (OBB) that serves as the starting point for our experiments (row 1 in Table 1). For example, the 'Data Augmentation' models, which were trained on the data augmented with fuzzy answer boundaries (rows 3–5 in Table 1), improve DRA@1 by 14.37–15.65% (absolute) compared to OBB. Similarly, the 'Data Augmentation + Transfer Learning' models (rows 6–8 in Table 1), which continued to train on the original data, improve DRA@1 by 10.70–14.99% (absolute).

Interestingly, 'Data Augmentation' models perform better for document retrieval than their equivalent 'Data Augmentation + Transfer Learning' models. We hypothesize that the transfer learning phase encourages the models to focus on the answer extraction task, ignoring the larger textual context infused in the previous step, which is useful for document retrieval. To further analyze the effects of data augmentation on document retrieval, we also list the results of a hybrid approach, where the documents retrieved by our 'Data Augmentation' models are used as input to the OBL baseline from [2]. DRA@1 increases noticeably by 11.25–15.12% (subtract line 2 from lines 9–11 in Table 1). We conclude that the candidate documents retrieved by data augmented models can boost the document retrieval performance of an external, independent model.

⁴<https://github.com/IBM/techqa>

⁵<https://github.com/wasiahmad/PolicyQA>

⁶PolicyQA has 17K questions in the training partition.

⁷<https://github.com/huggingface>

⁸We do not use Precision@N score for retrieval evaluation to keep our results comparable with the other original methods evaluated on these datasets.

⁹because PolicyQA answer spans are smaller than those in TechQA.

¹⁰PolicyQA does not have document retrieval components because each question is associated with only one long document.

¹¹As expected, tuning these hyper parameters is a balancing act between the desire to add more training data and introducing noise. We found that there is a "Goldilocks" zone for key parameters such as d , and the percentage of augmented training data.

Question	Context
1 Will you take my consent before	E-mail and newsletters: "If you signed up" for one of our email collecting or using my data? newsletters, we will send you ... We may also use the information ... (#if you have opted in for this information.)
2 Does the company use any form of aggregated or anonymized information?	We reserve the right to use, transfer, and *share #aggregated#, anonymous data* about our users as a group for any lawful business purpose, such as analyzing usage trends and seeking compatible advertisers, sponsors, clients and customers.

Table 3: Qualitative comparison of the outputs from our approach (Data Augmentation + Transfer Learning) and the respective baseline (Original settings with BERT base pretrained on SQUAD) on PolicyQA. The text in bold font indicates the gold answer for the corresponding question. The text between # and * indicate the answers from the baseline and our approach respectively for the given question.

Gold Data	Our Approach	Baseline
1 Q: What is the difference between Advanced and AdvancedOnly options Doc: None, Ans:	Doc: None, Ans: None	Doc: swg21672885 Ans: The bootstrapProcessServerData ...
2 Q: Is using a monitored JBoss a Windows server with ITCAM supported in Service? Doc: swg21967756 Ans: The JBoss service is not available to run with the ITCAM ... involves changes ... files currently not supported.	Doc: swg21967756 Ans: The JBoss service is not available to run as a Windows service when configured with the ITCAM for J2EE agent/DC because this involves changes	Doc: swg21967756 Ans: The JBoss service is not available
3 Q: Why do we get server error message when running BIRT reports after upgrading to Atlas 6.0.3.3? Doc: swg21982354 Ans: This happens when the BIRT Reports is running in Standalone mode and happens due to a new configuration - report.standalone.userid 1. Navigate to Atlas Properties Folder 2. Edit ... #report.standalone.userid=1	Doc: swg21982354 Ans: This happens when the BIRT Reports is running in Standalone mode and happens due to a new configuration - report.standalone.userid 1. Navigate to Atlas Properties Folder 2. Edit ... #report.standalone.userid=1	Doc: swg21981881 Ans: In Atlas 6.0.3.3, the Atlas Extensions logging configuration has moved to log4j.properties file. 1. Navigate to Atlas_Install_folder/Atlas/Properties 2. Edit log4j.properties file 3. Update the path

Table 4: Qualitative comparison of the outputs from our approach (Data Augmentation + Transfer Learning) and the respective baseline (OBB) on TechQA.

To illustrate the benefits of additional augmented data, we provide a qualitative comparison between the outputs of our approach and the respective baseline in Table 4. With the transferred knowledge of larger context, our model successfully retrieves the correct document (row 3 in Table 4). Further, our method correctly detects if a question has no answer in the given documents (row 1 in Table 4). The baseline struggles in both scenarios.

4.2 Answer Extraction

Data augmented models: As shown in Table 1, our data-augmented models (DAMs) yield improved AE performance on TechQA. Compared to the OBB baseline (row 1 in Table 1), F1 scores for answer extraction increase when data augmentation is added (see rows 3–5 in Table 1). There are absolute increases of 1.79–2.62% in F1 for DAMs. Table 4 (row 2) shows a qualitative example where the additional knowledge from the larger context helps the answer extraction component provide better (but not perfect) answer boundaries. Further, there are absolute increases of 3.82%, 4.33%, and 3.27% in F1 scores, respectively, when the DAMs are used to retrieve candidate documents as inputs to OBL for TechQA (rows 9–11 in Table 1). These results confirm that the better documents retrieved by our method lead to better answers extracted by OBL. Importantly, even though our document retrieval component produces much fewer documents than the original of 50, answer recall increases across the board, further confirming the improved quality of the documents retrieved.

Similarly, on PolicyQA, data augmentation improves AE performance (see Table 2). Compared to the OBS baseline (row 1), evaluation scores for answer extraction increase when data augmentation is added (absolute increases of 1.6–2.6% and 1.8–3.3% in F1 and EM, respectively, see rows 2–4)¹². This suggests that our simple approach is not dataset specific, which is encouraging. Similar to other previous works, we conjecture that our proposed data augmentation approach contributes substantially in very low

resource tasks (e.g., TechQA), and its impact decreases as more training resources are available (e.g., PolicyQA).

Transfer learning models: Compared to OBB, our 'Data Augmentation + Transfer Learning' (DA+TL) models perform considerably better for answer extraction in TechQA (rows 6–8 in Table 1). Transfer learning models also outperform data augmented models (0.55–1.1% absolute increases in F1). Comparing the best configurations between the two (rows 4 vs. 7 in Table 1), we observe that transfer learning models extract answer spans more effectively (absolute increases of 0.84% in F1 score). This confirms that transferring to the exact answer spans is beneficial for answer extraction components.

On PolicyQA, our DA+TL models improve AE performance (2.0–2.7% and 2.4–3.9% absolute increases in F1 and EM, subtract row 1 from rows 5–7 in Table 2). Also, Transfer Learning further improves AE performance (absolute increases of 0.1% and 0.6% in F1 and EM, compare the best configurations between the two: rows 4 vs. 7 in Table 2). Table 3 shows qualitative examples comparing DA+TL with the OBS baseline. Our qualitative analysis shows that our DA+TL approach more efficiently: (a) locates the correct answers spans (see row 1 in Table 3), and (b) captures necessary information by expanding the answer span (row 2 in Table 3).

5 CONCLUSION

We introduced a simple strategy for data augmentation for MRC, where we generate artificial training data points where answer spans are moved to the left/right of the correct spans. We demonstrated that this simple idea is useful for document retrieval (because it was able to capture the larger context in which the answer appears), and also as an effective pretraining step for answer extraction (because of the additional training data, which helps the extraction model narrow its search space). Our combined contributions yield the best overall performances for TechQA and PolicyQA.

6 ACKNOWLEDGEMENTS

This work was supported by the Defense Advanced Research Projects Agency (DARPA) under the World Modelers program, grant number W911NF1810014. Mihai Surdeanu declares a financial interest in lum.ai. This interest has been properly disclosed to the University of Arizona Institutional Review Committee and is managed in accordance with its conflict of interest policies.

¹²Please note that since PolicyQA contains 17,056 questions in its training partition (which is 30 times more than the TechQA training set), we observed the best improvements on PolicyQA when only 4% of the training data were used to generate new QA pairs with approximate boundaries.

REFERENCES

- [1] Wasi Ahmad, Jianfeng Chi, Yuan Tian, and Kai-Wei Chang. 2020. PolicyQA: A Reading Comprehension Dataset for Privacy Policies. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*. 743–749.
- [2] Vittorio Castelli, Rishav Chakravarti, Saswati Dana, Anthony Ferritto, Radu Florian, Martin Franz, Dinesh Garg, Dinesh Khandelwal, Scott McCarley, Mike McCawley, et al. 2019. The techqa dataset. *arXiv preprint arXiv:1911.02984* (2019).
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [4] Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179* (2017).
- [5] Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. MRQA 2019 shared task: Evaluating generalization in reading comprehension. *arXiv preprint arXiv:1910.09753* (2019).
- [6] David Golub, Po-Sen Huang, Xiaodong He, and Li Deng. 2017. Two-Stage Synthesis Networks for Transfer Learning in Machine Comprehension. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 835–844.
- [7] Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551* (2017).
- [8] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics* 7 (2019), 453–466.
- [9] Dayiheng Liu, Yeyun Gong, Jie Fu, Yu Yan, Jiusheng Chen, Jiancheng Lv, Nan Duan, and Ming Zhou. 2020. Tell Me How to Ask Again: Question Data Augmentation with Controllable Rewriting in Continuous Space. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 5798–5810.
- [10] Yinhan Liu, Mylène Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [11] Shayne Longpre, Yi Lu, Zhucheng Tu, and Chris DuBois. 2019. An Exploration of Data Augmentation and Sampling Techniques for Domain-Agnostic Question Answering. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*. 220–227.
- [12] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 784–789.
- [13] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603* (2016).
- [14] Swabha Swayamdipta, Ankur P Parikh, and Tom Kwiatkowski. 2018. Multi-Mention Learning for Reading Comprehension with Neural Cascades. In *International Conference on Learning Representations*.
- [15] Alon Talmor and Jonathan Berant. 2019. MultiQA: An Empirical Investigation of Generalization and Transfer in Reading Comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 4911–4921.
- [16] Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019. Multi-passage bert: A globally normalized bert model for open-domain question answering. *arXiv preprint arXiv:1908.08167* (2019).
- [17] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771* (2019).
- [18] Vikas Yadav, Steven Bethard, and Mihai Surdeanu. 2020. Unsupervised Alignment-based Iterative Evidence Retrieval for Multi-hop Question Answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4514–4525.