

Can LLMs Augment Low-Resource Reading Comprehension Datasets? Opportunities and Challenges

Vinay Samuel[‡] Houda Aynaou[§] Arijit Ghosh Chowdhury[†]
Karthik Venkat Ramanan[†] Aman Chadha^{♣,◇}

[‡]Carnegie Mellon University, [§]Georgia Institute of Technology,
[†]University of Illinois Urbana-Champaign [♣]Stanford University [◇]Amazon AI
vsamuel@andrew.cmu.edu, haynaou3@gatech.edu, arijit10@gmail.com,
kv16@illinois.edu, hi@aman.ai

Abstract

Large Language Models (LLMs) have demonstrated impressive zero shot performance on a wide range of NLP tasks, demonstrating the ability to reason and apply commonsense. A relevant application is to use them for creating high quality synthetic datasets for downstream tasks. In this work, we probe whether GPT-4 can be used to augment existing extractive reading comprehension datasets. Automating data annotation processes has the potential to save large amounts of time, money and effort that goes into manually labelling datasets. In this paper, we evaluate the performance of GPT-4 as a replacement for human annotators for low resource reading comprehension tasks, by comparing performance after fine tuning, and the cost associated with annotation. This work serves to be the first analysis of LLMs as synthetic data augmenters for QA systems, highlighting the unique opportunities and challenges. Additionally, we release augmented versions of low resource datasets, that will allow the research community to create further benchmarks for evaluation of generated datasets.

1 Introduction

Machine reading comprehension (MRC) is a challenging NLP task where systems are designed to answer questions based on a given context. This task has significant practical value, as it answers user queries in diverse settings, from clinical contexts (Krithara et al., 2023; Pampari et al., 2018; Pappas et al., 2020), to customer support (Castelli et al., 2020) and policy interpretation (Ahmad et al., 2020). BERT-based models (Glass et al., 2020) have achieved state-of-the-art performance when trained with extensive data from datasets like SQuAD (Rajpurkar et al., 2018) and Natu-

ral Questions (Kwiatkowski et al., 2019). However, their effectiveness diminishes in low-resource domains with limited datapoints (Schmidt et al., 2022). This limitation becomes particularly pronounced in newly emerging fields such as COVID-19 (Möller et al., 2020), where substantial annotated instances are often lacking.

Data augmentation has been instrumental in enhancing performance across numerous low-resource NLP tasks (Feng et al., 2021; Wang et al., 2022; Liu et al., 2021). Yet, much of the work on data augmentation for QA (Alberti et al., 2019; Shakeri et al., 2020; Bartolo et al., 2021; Dhingra et al., 2018; Yang et al., 2017), hinges on the availability of unlabeled paragraphs from common sources, such as Wikipedia, to produce new context-question-answer instances. This approach poses a challenge for specialized and mission-critical domains where such unlabeled contexts are scarcely available. Bridging this gap, LLMs (Brown et al., 2020) exhibit a capability to generate texts that closely resemble human-authored content (Brown et al., 2020; Clark et al., 2021). This potential of LLMs can be harnessed to generate both novel contexts and their corresponding question-answer pairs.

Addressing this gap, we introduce a GPT-4 (OpenAI, 2023) based data augmentation technique tailored for low-resource machine reading comprehension, specifically focusing on the extractive setting. Our approach begins by generating supplementary contexts, questions, and answers to augment training sets. To achieve this, we use in-context learning with passages, questions, and answers from the training set, ensuring minimal domain shift between the synthetically generated data and the original datasets

Subsequently, we adopt cycle-consistent filter-

*Work does not relate to position at Amazon.

ing to isolate high-quality training instances. Empirical evaluations conducted on three pertinent real-world low-resource datasets CovidQA (Möller et al., 2020), PolicyQA (Ahmad et al., 2020), and TechQA (Castelli et al., 2020) reveal that our methodology improves the performance of BERT-based MRC on CovidQA by 23% and on PolicyQA by 5% in terms of exact match. Notably, our approach attains state-of-the-art results on CovidQA.

2 Related Work

Language models have played a key role in the creation of synthetic datasets for various NLP tasks. Models such as GPT-2 (Radford et al., 2019) and CTRL (Keskar et al., 2019) have been applied to areas including general language understanding (Meng et al., 2022; He et al., 2022), classification (Kumar et al., 2020; Anaby-Tavor et al., 2019), dialogue tasks (Mohapatra et al., 2021), commonsense reasoning (Yang et al., 2020), and relation extraction (Papanikolaou and Pierleoni, 2020), among others. Recently, large language models have significantly improved the quality and scope of synthetic dataset generation. They have been instrumental in augmenting datasets for tasks such as NLI and sentiment analysis (Dixit et al., 2022), classification (Yoo et al., 2021), and even creating datasets for personalized dialogue generation (Lee et al., 2022), hate speech detection (Hartvigsen et al., 2022), and textual similarity (Schick and Schütze, 2021) to name a few.

Most prior work in synthetic data generation for QA (Riabi et al., 2021; Chakravarti et al., 2020; Du and Cardie, 2018; Alberti et al., 2019) has concentrated on generating questions from Wikipedia passages to produce supplementary training examples. More recently, Kalpakchi and Boye introduced the use of GPT-3 for creating extra training data for Swedish multiple choice questions. Our approach is the first to utilize in-context learning with LLMs for synthesizing contexts, questions, and answers for low-resource MRC.

3 Setup

3.1 Low Resource Datasets

We utilize three reading comprehension datasets in our work: CovidQA, PolicyQA, and TechQA. These datasets cover diverse domains while having relatively small training sizes, making them well-suited for evaluating synthetic data augmentation techniques.

The CovidQA dataset (Möller et al., 2020) focuses on question answering related to the COVID-19 pandemic. It contains 2,019 question-answer pairs on topics such as virus transmission, public health interventions, and social impacts.

PolicyQA (Ahmad et al., 2020) contains 12,102 question-answer pairs about United States immigration and travel policies. The questions require reasoning about specific policy documents to determine the answer.

TechQA (Castelli et al., 2020) provides 1,808 examples related to technical support issues on computer networking, software, and hardware. The goal is to develop QA systems that can resolve technical problems automatically.

In summary, these three datasets cover the domains of healthcare, public policy, and technology, while having relatively small training set sizes between 1-10k examples. This makes them suitable testbeds for studying the effects of augmenting the training data through synthetic example generation.

4 Synthetic Data Generation

We generate synthetic examples for each dataset using the in-context learning capabilities of the GPT-4 model. The data generation process consists of two stages:

4.1 Context Generation

In the first stage, we provide GPT-4 with either 1 example (one-shot) or 2 examples (two-shot) of contexts from the original training set of each dataset. These few-shot examples prime GPT-4 on the style and topics present in the contexts. Providing just one or two examples allows GPT-4 to adapt from demonstrations due to the robust few-shot learning capabilities of LLMs (Reif et al., 2022; Frohberg and Binder, 2022; Wei et al., 2022). We then generate new synthetic paragraph-length contexts by providing a prompt and allowing GPT-4 to complete the paragraph based on the few-shot priming.

4.2 QA Generation

The second stage generates synthetic question-answer pairs conditioned on the synthetic contexts. We again prime GPT-4 with either 1 example (one-shot) or 2 examples (two-shot) of QA pairs from the original dataset. The few-shot priming allows GPT-4 to learn the QA pattern quickly. We then provide the synthetic context from the first stage along with a prompt for GPT-4 to generate a rele-

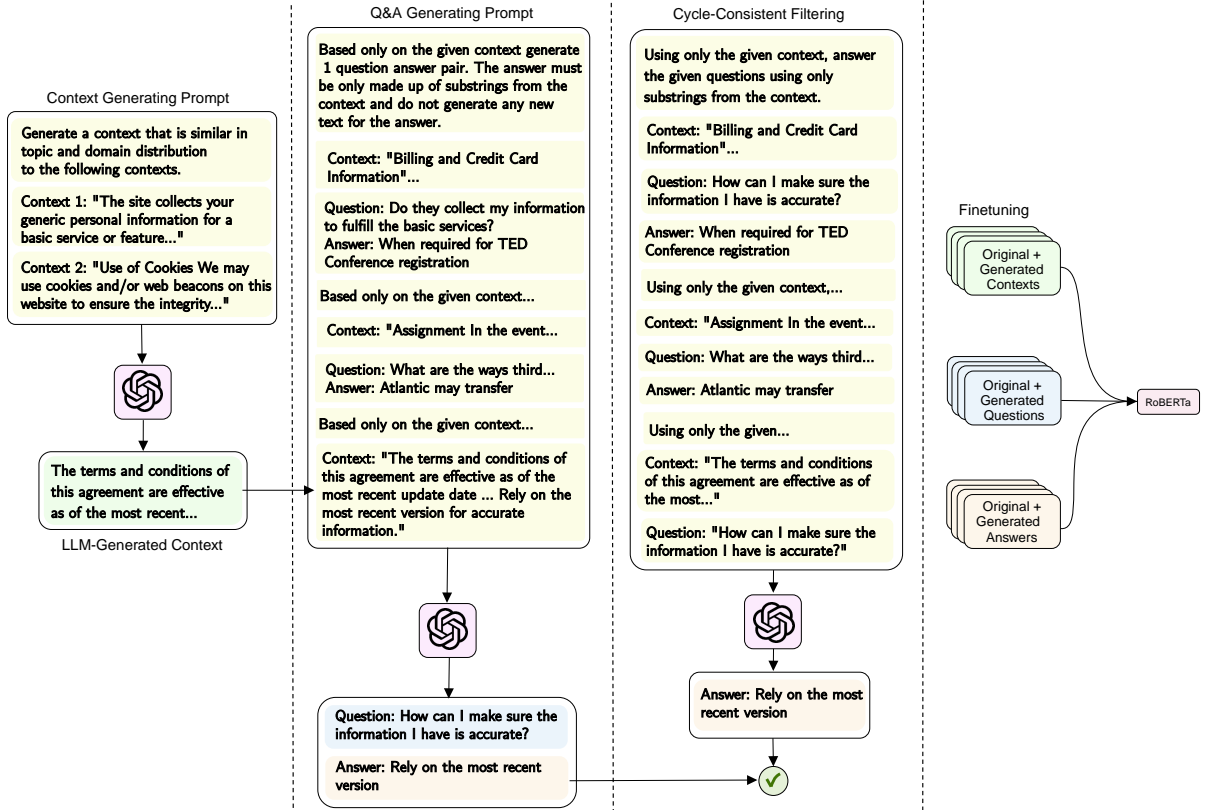


Figure 1: Overview of our methodology using PolicyQA as an example with 2-shot prompts.

vant question and answer pair mimicking the style of the examples.

This two-stage process allows us to leverage the few-shot learning and text generation capabilities of GPT-4 to produce synthetic datasets that mimic the style and semantics of the original data. We generate varying amounts of synthetic data, from 1x to 10x the size of the original training sets, to study the impact on downstream task performance.

4.2.1 Round Trip Filtration

To further improve the quality of the synthetic QA pairs, we implement a round trip filtration technique. After generating a synthetic question and answer using GPT-4, we provide the question back to the model without the answer. We allow GPT-4 to attempt answering the question again based on the context. If the model’s newly generated answer matches the original synthetic answer, we retain this QA pair, as it indicates a high quality question with a consistent answer. If the answers do not match, we discard the synthetic QA pair under the assumption that the question is flawed in some way.

This round trip filtration process provides a mechanism for GPT-4 to self-filter its own generated content. By only keeping QA pairs that

exhibit consistency when answered twice, we obtain higher quality synthetic data for downstream training. The filtration process improves precision at the potential expense of some recall.

4.3 Experiments

We train an extractive reading comprehension model, using the RoBERTa-Base model across all our experiments. We use a learning rate of $3e - 5$, a batch size of 16 and run our experiments for 5 epochs each. We use the implementation provided by HuggingFace, and run our models on a stand-alone Nvidia V100 GPU. For all our experiments, we measure F1 and Exact Match scores.

As a baseline for question-answer generation we use a T5 based question generation model that is trained on the SQUAD dataset, which takes a paragraph as an input and returns a question-answer pair. We use the open source ¹ implementation for this model.

5 Results

Table 1 highlights results across the three datasets. For the CovidQA dataset, we observed steady im-

¹https://github.com/patil-suraj/question_generation

CovidQA		
Setup	Exact Match	F1 Score
Original Trainset	25.81	50.91
Baseline	19.71	44.18
One Shot	30.82	57.87
Two Shot	31.18	55.64
One Shot (CC)	31.90	58.66
Two Shot (CC)	30.82	53.40
PolicyQA		
Setup	Exact Match	F1 Score
Original Trainset	30.56	58.15
Baseline	30.08	57.65
One Shot	32.18	59.61
Two Shot	30.97	59.12
One Shot (CC)	30.76	58.71
Two Shot (CC)	30.47	58.46
TechQA		
Setup	Exact Match	F1 Score
Original Trainset	11.11	39.45
Baseline	44.44	59.92
One Shot	22.22	36.91
Two Shot	11.11	36.50
One Shot (CC)	22.22	41.76
Two Shot (CC)	22.22	44.73

Table 1: Experimental Results for MRC Across Various Datasets and Settings.

provements in question answering performance as we augmented the original training set with increasing amounts of synthetic data generated by GPT-4. Using just the original training examples, our model achieved baseline exact match (EM) and F1 scores on the validation set. Adding one-shot synthetic examples improved both the EM and F1 metrics over the baseline. We observed further gains when using two-shot synthetic data, achieving higher EM and F1 compared to one-shot.

The best validation results on CovidQA were obtained by using the one-shot synthetic dataset combined with the round trip filtration process. This achieved the highest EM and F1 scores, significantly improving over the original training distribution. We hypothesize that the round trip filtration allows for higher precision synthetic data, while the one-shot generation provides greater diversity compared to two-shot. The balance of quality and variety in this one-shot filtered dataset appears optimal for augmenting the limited original examples in the CovidQA training set.

In summary, for the CovidQA task we find that synthetic data augmentation uniformly improves performance as more examples are added. The best

results come from combining one-shot generation with round trip filtration, which improves exact match and F1 score over the baseline set using just the original dataset.

With over 12,000 examples, PolicyQA was the largest dataset we utilized. For this task, augmenting the original training set with one-shot synthetic data without filtration achieved the best question answering performance. This improved exact match by 1.6 points and F1 score by 1.5 points compared to using just the original examples. The one-shot augmentation outperformed both two-shot and cycle filtered variations.

Overall for PolicyQA, we find that synthetic data augmentation consistently improves upon the baseline set using just the original training examples. The best configuration utilizes unfiltered one-shot generation, likely due to the greater diversity of examples compared to two-shot or filtered versions. While the domain of US immigration policies has high complexity, the large size of the PolicyQA dataset reduces the need for precision-enhancing filtration. The additional synthetic examples provide useful variability when training the model.

With only 1,808 examples, TechQA was the smallest dataset in our experiments. The tiny test set of just 9 examples also made evaluation challenging. On this task, augmenting with synthetic data did not lead to clear improvements in question answering accuracy over the original training set. The baseline model trained on just the 1,808 TechQA examples achieved the highest exact match score, with the two-shot cycle filtered, one-shot filtered, and one-shot unfiltered configurations performing second best in terms of EM. For F1, two-shot cycle filtered data obtained the second highest score after the baseline.

The lack of consistent gains from synthetic data augmentation on TechQA can likely be attributed to the very small data size. With fewer than 2,000 training examples, there is insufficient context for the language model to learn effective generalization. The technical support domain also exhibits diversity that may not be captured from only 1-2 conditioning examples. Furthermore, the small test set provides high variance in evaluation.

6 Opportunities

Our experiments demonstrate the significant potential of leveraging large language models (LLMs) like GPT-3 for synthetic data generation. In the

CovidQA and PolicyQA domains where a moderate amount of training data was available, augmenting with LLM-produced synthetic examples consistently improved performance over the baseline trained on just the original dataset. This confirms the few-shot generalization abilities of modern LLMs in producing varied, high-quality synthetic data when primed with only a handful of real examples. Indeed, the one-shot synthetic data augmented models achieved the best results on both CovidQA and PolicyQA, surpassing two-shot and other configurations.

The natural language generation capabilities of LLMs afford great opportunity to increase the diversity and size of limited training sets for downstream tasks. By prompting the models to produce synthetic examples mimicking the patterns in the data, we can expand datasets to be orders of magnitude larger with plausible, human-like samples. This data augmentation approach can be applied to many NLP tasks suffering from small training sizes like reading comprehension, summarization, translation, and more. High-quality synthetic data translates into better task performance without the expense of human labeling efforts.

Critical research directions include developing more advanced filtering techniques to distill only the most useful synthetic samples, as well as integrating external knowledge sources to improve few-shot priming. But the overarching opportunity is clear – properly harnessed, LLMs have enormous potential to ameliorate the limited data problem through strategic synthetic generation.

7 Challenges

However, our experiments on the extremely small TechQA dataset also reveal current limitations in using LLMs for robust synthetic data generation. When provided with only around 1,000 original training examples, the LLM-augmented models performed no better than baseline. The models failed to learn adequate representations from such scarce data for producing useful synthetic examples. This highlights how modern LLMs, despite their progress, still struggle in low-data regimes where broad generalization capabilities are required.

Critical challenges remain in improving LLMs' few-shot learning to make them reliable across diverse domains. Environments with limited data require synthesizing examples from broader con-

ceptual knowledge, not just mimicking surface patterns. Integrating external knowledge into LLMs is an active area of research, but effectively utilizing such knowledge in few-shot scenarios remains difficult. There are also challenges in filtering large volumes of synthetic data to maximize diversity while maintaining precision and quality.

In summary, while LLMs offer promise for alleviating limited training data, substantial challenges persist. Robustness to low-data regimes, integration of world knowledge, and advanced content filtering mechanisms are needed to make synthetic data generation truly effective for any NLP task. This is an exciting and rapidly evolving area of research that will determine whether LLMs can deliver on their potential to mitigate limited datasets through strategic synthetic example construction.

References

- Wasi Ahmad, Jianfeng Chi, Yuan Tian, and Kai-Wei Chang. 2020. [PolicyQA: A reading comprehension dataset for privacy policies](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 743–749, Online. Association for Computational Linguistics.
- Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. 2019. [Synthetic QA corpora generation with roundtrip consistency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6168–6173, Florence, Italy. Association for Computational Linguistics.
- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, N. Teppe, and Naama Zwerdling. 2019. [Do not have enough data? deep learning to the rescue!](#) In *AAAI Conference on Artificial Intelligence*.
- Max Bartolo, Tristan Thrush, Robin Jia, Sebastian Riedel, Pontus Stenetorp, and Douwe Kiela. 2021. [Improving question answering model robustness with synthetic adversarial data generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8830–8848, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020.

- Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Vittorio Castelli, Rishav Chakravarti, Saswati Dana, Anthony Ferritto, Radu Florian, Martin Franz, Dinesh Garg, Dinesh Khandelwal, Scott McCarley, Michael McCawley, Mohamed Nasr, Lin Pan, Cezar Pendus, John Pitrelli, Saurabh Pujar, Salim Roukos, Andrzej Sakrajda, Avi Sil, Rosario Uceda-Sosa, Todd Ward, and Rong Zhang. 2020. [The TechQA dataset](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1269–1278, Online. Association for Computational Linguistics.
- Rishav Chakravarti, Anthony Ferritto, Bhavani Iyer, Lin Pan, Radu Florian, Salim Roukos, and Avi Sil. 2020. [Towards building a robust industry-scale question answering system](#). In *Proceedings of the 28th International Conference on Computational Linguistics: Industry Track*, pages 90–101, Online. International Committee on Computational Linguistics.
- Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A. Smith. 2021. [All that’s ‘human’ is not gold: Evaluating human evaluation of generated text](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7282–7296, Online. Association for Computational Linguistics.
- Bhuvan Dhingra, Danish Danish, and Dheeraj Rajagopal. 2018. [Simple and effective semi-supervised question answering](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 582–587, New Orleans, Louisiana. Association for Computational Linguistics.
- Tanay Dixit, Bhargavi Paranjape, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [CORE: A retrieve-then-edit framework for counterfactual data generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2964–2984, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xinya Du and Claire Cardie. 2018. [Harvesting paragraph-level question-answer pairs from Wikipedia](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1907–1917, Melbourne, Australia. Association for Computational Linguistics.
- Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edward Hovy. 2021. [A survey of data augmentation approaches for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 968–988, Online. Association for Computational Linguistics.
- Jörg Froberg and Frank Binder. 2022. [CRASS: A novel data set and benchmark to test counterfactual reasoning of large language models](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2126–2140, Marseille, France. European Language Resources Association.
- Michael Glass, Alfio Gliozzo, Rishav Chakravarti, Anthony Ferritto, Lin Pan, G P Shrivatsa Bhargav, Dinesh Garg, and Avi Sil. 2020. [Span selection pre-training for question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2773–2782, Online. Association for Computational Linguistics.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. [Toxigen: A large-scale machine-generated dataset for implicit and adversarial hate speech detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*.
- Xuanli He, Islam Nassar, Jamie Kiros, Gholamreza Hafari, and Mohammad Norouzi. 2022. [Generate, annotate, and learn: NLP with synthetic text](#). *Transactions of the Association for Computational Linguistics*, 10:826–842.
- Dmytro Kalpakchi and Johan Boye. 2023. [Quasi: a synthetic question-answering dataset in Swedish using GPT-3 and zero-shot learning](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 477–491, Tórshavn, Faroe Islands. University of Tartu Library.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. [Ctrl: A conditional transformer language model for controllable generation](#). *ArXiv*, abs/1909.05858.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. [Bioasqqa: A manually curated corpus for biomedical question answering](#). *Scientific Data*, 10(1):170.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. [Data augmentation using pre-trained transformer models](#). In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*, pages 18–26, Suzhou, China. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.

- Young-Jun Lee, Chae-Gyun Lim, Yunsu Choi, Ji-Hui Lm, and Ho-Jin Choi. 2022. [PERSONACHATGEN: Generating personalized dialogues using GPT-3](#). In *Proceedings of the 1st Workshop on Customized Chat Grounding Persona and Knowledge*, pages 29–48, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Linlin Liu, Bosheng Ding, Lidong Bing, Shafiq Joty, Luo Si, and Chunyan Miao. 2021. [MulDA: A multilingual data augmentation framework for low-resource cross-lingual NER](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5834–5846, Online. Association for Computational Linguistics.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. In *Advances in Neural Information Processing Systems*.
- Biswesh Mohapatra, Gaurav Pandey, Danish Contractor, and Sachindra Joshi. 2021. [Simulated chats for building dialog systems: Learning to generate conversations from instructions](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1190–1203, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Timo Möller, Anthony Reina, Raghavan Jayakumar, and Malte Pietsch. 2020. [COVID-QA: A question answering dataset for COVID-19](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*, Online. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. 2018. [emrQA: A large corpus for question answering on electronic medical records](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2357–2368, Brussels, Belgium. Association for Computational Linguistics.
- Yannis Papanikolaou and Andrea Pierleoni. 2020. [Dare: Data augmented relation extraction with gpt-2](#). *ArXiv*, abs/2004.13845.
- Dimitris Pappas, Petros Stavropoulos, Ion Androutsopoulos, and Ryan McDonald. 2020. [BioMRC: A dataset for biomedical machine reading comprehension](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 140–149, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- 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)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. [A recipe for arbitrary text style transfer with large language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, Dublin, Ireland. Association for Computational Linguistics.
- Arij Riabi, Thomas Scialom, Rachel Keraron, Benoît Sagot, Djamé Seddah, and Jacopo Staiano. 2021. [Synthetic data augmentation for zero-shot cross-lingual question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7016–7030, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Generating datasets with pretrained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6943–6951, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Maximilian Schmidt, A. Bartezzaghi, Jasmina Bogojeska, Adelmo Cristiano Innocenza Malossi, and Thang Vu. 2022. [Improving low-resource question answering using active learning in multiple stages](#). *ArXiv*, abs/2211.14880.
- Siamak Shakeri, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. [End-to-end synthetic data generation for domain adaptation of question answering systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5445–5460, Online. Association for Computational Linguistics.
- Yufei Wang, Can Xu, Qingfeng Sun, Huang Hu, Chongyang Tao, Xiubo Geng, and Daxin Jiang. 2022. [PromDA: Prompt-based data augmentation for low-resource NLU tasks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4242–4255, Dublin, Ireland. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. [Generative data augmentation for common-sense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages

1008–1025, Online. Association for Computational Linguistics.

Zhilin Yang, Junjie Hu, Ruslan Salakhutdinov, and William Cohen. 2017. [Semi-supervised QA with generative domain-adaptive nets](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1040–1050, Vancouver, Canada. Association for Computational Linguistics.

Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. 2021. [GPT3Mix: Leveraging large-scale language models for text augmentation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic. Association for Computational Linguistics.