

Abstract. Neural approaches have become very popular in Question Answering (QA), however, they require a large amount of annotated data. In this work, we propose a novel approach that combines data augmentation via question-answer generation with Active Learning to improve performance in low-resource settings, where the target domains are diverse in terms of difficulty and similarity to the source domain. We also investigate Active Learning for question answering in different stages, overall reducing the annotation effort of humans. For this purpose, we consider target domains in realistic settings, with an extremely low amount of annotated samples but with many unlabeled documents, which we assume can be obtained with little effort. Additionally, we assume a sufficient amount of labeled data from the source domain being available. We perform extensive experiments to find the best setup for incorporating domain experts. Our findings show that our novel approach, where humans are incorporated in a data generation approach, boosts performance in the low-resource, domain-specific setting, allowing for low-labeling-effort question answering systems in new, specialized domains. They further demonstrate how human annotation affects the performance of QA depending on the stage it is performed.

Keywords: Question Answering · Active Learning · Data Generation.

1 Introduction

Machine Reading Question Answering (MRQA) is a challenging and important problem. Facilitating targeted information extraction from documents, it allows users to get fast, easy access to a vast amount of documents available. MRQA models generally need plenty of annotations [31], therefore several methods have been devised for augmenting data by generating new annotated samples, with the ultimate goal of improving quality of predictions. Some of these approaches

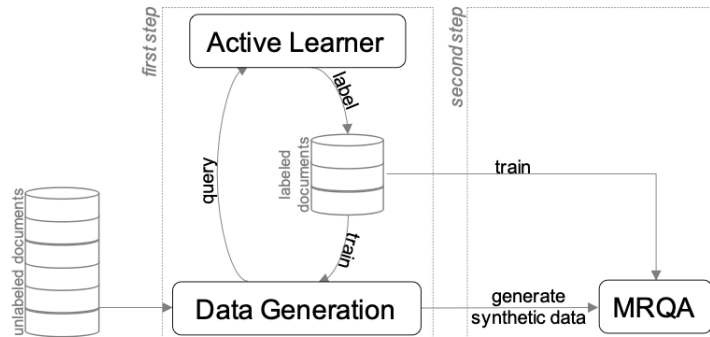


Fig. 1. Our approach combining Active Learning with data generation: In a first step, the data generation model is efficiently trained using Active Learning. Second, this model is then used to generate data for MRQA.

show a real benefit in the downstream MRQA task; however, there is no work employing Language Models (LM) fine-tuned for generating question-answer pairs in low-resource, domain-specific settings, which are often observed in practice.

In these cases, most of the times only few annotated samples are available due to the lack of resources like domain experts [28]. Also, since the specialized domain being usually vastly different from the publicly available labeled data of the source domain, this data is only useful to a limited extent. Moreover, labeling is expensive, as it requires a significant amount of time from domain experts.

One approach to reduce the annotation effort and to employ humans more efficiently is to utilize Active Learning (AL) [33]. For this, samples are carefully selected for annotation to reduce the amount of redundant information.

Another common method for reducing the annotation effort is data generation. By generating synthetic data, several works have shown to successfully counteract the low-resource setting [1,35,34]. Consequently, combining data generation with AL can be another step towards reducing the high annotation effort for MRQA. However, so far it stays unclear how few samples from the target domain, selected based on various criteria, affect the performance of the MRQA model when they are used in the training of the data generation model as well. In this work, we introduce a novel approach that follows this idea, and investigate how employing AL, either at the data generation model or at the MRQA model, affects the overall performance of the MRQA task. To this end, we consider collecting and annotating a small set of samples (e.g., 200), accurately selected to boost performance of the MRQA model, feasible and not too expensive.

More precisely, we aim to answer the following research questions:

RQ1: How can we combine data generation with Active Learning to improve the performance of extractive MRQA?

RQ2: How can we select samples for annotation for extractive MRQA when data generation is used?

To achieve this, we consider only a small set of documents from the *target* domain annotated in terms of questions and answers, in addition to a potentially large set of samples from a different, generic *source* domain. We further assume to have access to a large amount of unlabeled data (documents) available from the target domain, which is common in real scenarios (e.g., publications from PubMed⁴ when considering the medical domain). We then employ a state-of-the-art model for generating question-answer pairs on these documents from the target domain, to enrich our training sets with synthetic data. In our research, we study the application of AL in order to label those samples which are most relevant for increasing performance, with the aim of reducing the amount of labeled samples needed. We consider these labeled samples in both the data generation stage as well as for the MRQA model itself.

Our contributions in this paper can be summarized as follows: 1) We introduce a novel approach that combines data augmentation for MRQA via question-answer generation with Active Learning; 2) we identify the most relevant samples for AL by adapting to our setting scoring functions recently used for unsupervised quality assessment of machine translation; 3) we introduce a new sample relevance score specific to MRQA by coupling the generated samples with the downstream task, so that the MRQA model influences the data augmentation process as well; and 4) we perform extensive experiments⁵ to demonstrate that AL is best used in our setup at the stage of data generation for improving MRQA performance in low-resource, domain-specific scenarios.

Our experimental results show great improvements for the extractive MRQA task considering two domain-specific datasets, namely TechQA [3] and BioASQ [38], with few samples specifically annotated for the data generation model. In numbers, combining AL with data generation increases F1 by up to 5.13% compared to applying AL at the MRQA model.

2 Related Work

2.1 Low-Resource MRQA

In the literature, there are several approaches for tackling the low-resource problem, including scenarios where few labeled data are available and others where no labels are available at all.

One approach is to use pre-trained Language Models [1,30,20] backed by Transformers [42], which can be especially useful in cases where little or no labeled data exists and it is costly to generate more. In the best case, the LM can be used without further fine-tuning (when the target task and domain is similar to those used in the pre-training objective). Otherwise, if unlabeled data is available, it may be used for adaptation of the LM in a self-supervised fashion.

If the low-resource domain is accompanied by some annotated samples, weak supervision – where external sources of information are used – becomes relevant

⁴ <https://pubmed.ncbi.nlm.nih.gov>

[12,43]. Moreover, data augmentation [46,41] and LM domain adaptation [27,46] have been shown to improve performance for MRQA.

A different approach is to use domain transfer, where a model is trained on a source domain and then adapted to a different target domain, e.g., by employing adversarial training [19].

2.2 Data Generation

As generative models are becoming increasingly popular, recent work has shifted from domain transfer towards data generation [35,1,29,26,18]. Given a passage of text, the generator model learns to output question-answer pairs. The generation model is trained on a large amount of data from the source domain and is applied to any document in any target domain, commonly employing pre-trained LMs as decoder.

Common generation methods focus only on generating the questions and assess their performance exclusively on the generated questions using automatic metrics [40,22,39,45,5,37]. Only a few approaches are based on generating full question-answer pairs, where the generated data is evaluated by means of the MRQA model. While [15] and [26] only perform in-domain experiments – where the training data and data used for generating new samples come from the same domain – [35] and [18] show the generalizability of their data generation models by also performing out of domain experiments. Recent models even perform well in the few- and zero-shot setting [32]. Since data generation can also be improved using annotated data from the target domain [32], it is of interest to reduce the labeling effort.

2.3 Active Learning

Active Learning [33,6] is a framework aimed to reduce the amount of annotated data required for training machine learning models, based on iteratively selecting a set of specific samples to be labeled by a human with expert knowledge of the domain. AL has been widely used in Deep Learning [23] and NLP [36,8,25,4,21], and it has been shown to be helpful for MRQA as well: While [16] learn whether samples are annotated automatically or manually, others make use of pool based sampling strategies, based on heuristics, to score samples for annotation [21].

3 Method

Our novel approach to tackle MRQA in low-resource settings with optimized human effort is based on combining AL with data generation. In this section we describe our method including the models involved, i.e., the data generation and the MRQA model (section 3.1), as well as our iterative sample selection method (section 3.2). The latter is implemented using Active Learning, which adapts and combines various agnostic scores, with the addition of a novel MRQA-tailored one, to improve the annotation efficiency. A high-level overview of our setup is provided in figure 1.

Algorithm 1 Model Training with AL

Require:

- unlabeled target domain documents $\mathcal{D}_{\text{pool}}$
- labeled target domain data \mathcal{D}_{dev}
- scoring function f_{score} taking (sample, models) as input
- iterations r , number of samples n
- 1: Initialize involved models \mathcal{M} with pre-trained weights
- 2: $\mathcal{D}_{\text{train}} = \emptyset$
- 3: **for** iteration = 1, 2, ..., r **do**
- 4: $\mathcal{D}_{\text{selected}} \leftarrow$ select n top scoring samples from $\mathcal{D}_{\text{pool}}$ using $f_{\text{score}}(\cdot, \mathcal{M})$
- 5: $\mathcal{D}_{\text{labeled}} \leftarrow$ label¹ $\mathcal{D}_{\text{selected}}$
- 6: $\mathcal{D}_{\text{train}} \leftarrow \mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{labeled}}$
- 7: $\mathcal{D}_{\text{pool}} \leftarrow \mathcal{D}_{\text{pool}} \setminus \mathcal{D}_{\text{selected}}$
- 8: Train \mathcal{M} using $\mathcal{D}_{\text{train}}$ and \mathcal{D}_{dev}
- 9: **end for**

¹In our experiments we employ an oracle, i.e., we take the label from the dataset

3.1 MRQA with Data Generation

Our methodology refers to the question as q , the answer as a and the corresponding context as c . Regarding data generation, we consider the *QA2S* model proposed by [35], since this model shows best overall performance in their and our experiments. This model makes use of a pre-trained encoder-decoder LM and is fine-tuned to generate the question given the context, followed by generating the answer given the context and the previously decoded question in a subsequent decoding step. More details can be found in the original paper [35].

For the MRQA model, we employ pre-trained BERT [7] for encoding the question concatenated with the context. On top, a span extraction head models the probability for each context token to be the start and the end of the answer span. We refer to [7] for more details.

3.2 Active Learning

We consider pool based sampling for our AL scenario, where samples are iteratively selected for annotation. Therefore, a pool of unlabeled samples is ranked using a scoring function based on context c . Our novel approach, combining data generation together with AL, is briefly described in Algorithm 1. In all cases, we initially fine-tune the models on the generic source domain to have a grounded starting point for AL.

We borrow Sentence Probability (SP), Dropout-based Sentence Probability (D-SP) and Dropout-based Lexical Similarity (LS) from unsupervised quality estimates [9,44] and adapt them as scoring functions for sample selection using the data generation model in our setting. While the first two functions measure the model’s entropy, the idea of the LS scoring function is that the difference of likely answers can also expose the model’s uncertainty. We also introduce round-trip scoring (RT), which incorporates the MRQA model in the sample selection

process: We believe that, during data generation, it is important to link the sample selection process with the eventual downstream task in order to generate high quality samples. Otherwise, the data generation model is not constrained in terms of questions being generated, potentially ending up generating questions which do not take the given context or answer into account or which do not add benefit to the MRQA model in general.

Scoring functions

Sentence Probability (SP) This scoring function is based on the probability distribution of the data generation model. The contexts are scored according to the sentence probability of the answer being generated:

$$\text{SP}(\theta) = \frac{1}{T} \sum_{t=1}^T \log p(a_t | a_{<t}, c, q, \theta). \quad (1)$$

We generate the question and answer using beam search of size 10 without sampling, but only use the generated answer for scoring a sample’s context.

Dropout-based Sentence Probability (D-SP) In contrast to SP, D-SP makes use of multiple data generation models to compute the sentence probability:

$$\text{D-SP} = \frac{1}{N} \sum_{n=1}^N \text{SP}(\theta_n). \quad (2)$$

We employ dropout at inference time (in addition to during fine-tuning) to realize different subsets of the model, as described by [10], running N forward passes. The question and answer is decoded once similarly to SP and used in all forward passes.

Dropout-based Lexical Similarity (LS) Regarding LS, multiple answers are decoded using multiple models (again realized via dropout), and all of them are compared pairwise at the lexical level using the Meteor [2] metric (with $i \neq j$):

$$\text{LS} = \frac{1}{N * (N - 1)} \sum_{i=1}^N \sum_{j=1}^N \text{Meteor}(a_i, a_j). \quad (3)$$

Decoding is again implemented using beam search of size 10 without sampling.

Round-trip (RT) Since the ultimate goal is to improve the performance of the MRQA model, we propose this novel method integrating the MRQA model in the computation of the ranking score: Question-answer pairs are generated from the context similarly to SP, and then we apply an MRQA model in order to rank the samples according to the F1 score between the predicted answer and the answer generated by the data generation model. We therefore prefer documents for which the data generation model will generate question-answer pairs which the MRQA model cannot predict correctly.

Table 1. Statistics about the train split for each domain obtained with the Bert tokenizer.

Domain	Context tokens			Questions tokens			Answer tokens		
	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean
SQuAD	25	853	155.75	1	61	12.29	1	68	4.23
NQ	10	3143	245.4	7	29	9.76	1	270	5.25
TechQA	38	38925	1484.08	6	561	69.37	5	545	98.76
BioASQ	27	960	337.71	5	36	15.06	1	76	4.42

4 Experimental Setup

In this work, we focus on the low-resource setting, where we assume to be able to use an extremely small amount of 200 annotated samples from the target domain. In addition to a large annotated source domain, we assume to have many unlabeled documents available from the target domain as well, which can be used for generating labeled samples. We consider this assumption to be valid in general, as one of the main goals of MRQA is to provide better access to the vast amount of existing documents.

4.1 Data

We consider SQuAD [31] as source domain and NaturalQuestions (NQ) [17], TechQA [3] and BioASQ [38] as target domains; statistics on tokens are summarized in table 1. The main reason for choosing these datasets is to consider examples of very specialized domains (TechQA, BioASQ) completely different from the source domain, as well as an example of a domain overlapping with the source domain (NQ).

In our experiments with AL, we use SQuAD for initially fine-tuning the data generation model. Since BioASQ only includes a dev set, we create random splits for training, development and testing with 70%, 20% and 10% of samples, respectively.

Regarding the documents needed for generating question-answer pairs, we use the documents from the training data in case of SQuAD and NQ. For TechQA and BioASQ, we use passages from the corpus of technotes containing 800K+ documents and crawled abstracts from PubMed⁶, respectively.

4.2 Data Generation

Training with Active Learning For training the data generation model with AL, we score available unlabeled documents for annotation using each of the scoring functions described in Section 3.2, running 10 forward passes for D-SP and LS. We only score documents for which labels (i.e., questions and answers)

⁶ <https://pubmed.ncbi.nlm.nih.gov>

exist as we simulate the domain expert. That is, for the selected unlabeled documents after scoring, we take the labels from the provided datasets (i.e., we use them as an oracle). We run 4 iterations, starting initially with all available samples from the target domain in the pool, and then selecting the 50 samples the model is least confident about in each step. In each iteration, we remove the selected samples from the query pool and fine-tune the model on all samples selected so far, always re-initializing the model with weights gained when fine-tuning on the source domain in order to not introduce any bias by having the model see some samples more often. Since the MRQA model is used in RT, we also train the MRQA model in each iteration when this scoring functions is used.

Furthermore, as D-SP, LS and RT are computationally expensive, with LS actually decoding generated text after each forward pass, performing AL with LS and RT on NQ was not feasible for us. Therefore, we selected 10000 random samples from this domain for use with AL, denoted as NQ#10000 in our experiments.

Regarding the models, we use `bart-large` [20] with 406M parameters for the data generation model as this showed better results in our preliminary experiments compared to `gpt2-large` [30] and is similar to previous work [35] from where we also take hyperparameters. Therefore, the data generation model is fine-tuned using cross entropy loss with AdamW Optimizer [24], a learning rate of $3e-5$ with 10% warm-up and a batch size of 24. The best model is selected on the evaluation data w.r.t. loss. Fine-tuning of the data generation model is performed for 5 epochs on source domain data first, followed by fine-tuning on the target domain. As we found that 5 epochs is too few for fine-tuning in the low-resource setting, on the target domain we fine-tune it for 10 epochs instead. We use this type of large language model as it can be trained with rather low hardware constraints, and iterative training is crucial to AL. For details of the MRQA model and its training see section 4.3.

To fit the context concatenated with the question into the models, we split the context into several chunks. We only consider those chunks where the answer occurs in the context. Additionally, since TechQA has long questions, we truncate the questions to the first 200 tokens, to allow for sufficient space in the input for the context as well.

Synthetic Data Generation Regarding question generation (i.e., the first decoding step), we allow contexts to have a size of up to 724 tokens; this allows the generated questions to be composed of up to 300 tokens, as both will be fed into the model in the second decoding step and their sum cannot exceed the total of 1024 tokens (including special tokens). Chunking this input would be difficult, since aggregating the generated answer is not simple (especially if outputs do not overlap). We generate 10 questions per context, followed by one answer for each question-context pair, yielding up to 10 samples per context. We do this to increase the variance of generated data for which different questions are crucial, but there can only be a single answer for each question-context pair.

Question generation is performed using nucleus sampling [13], considering 95% of the probability mass (top-p) and top 20 tokens (top-k) in each step. The answer is then decoded using beam search of size 10 without sampling similar to prior work [35]. Samples for which the answer does not occur in the context are discarded since we focus on extractive MRQA. Furthermore, only samples for which the corresponding end tokens are predicted correctly are considered as valid.

In total, we choose 100000 documents from the target domain, ignoring contexts with less than 100 tokens. In case of NQ, only 50535 documents remain.

Filtering Finally, only a subset of the generated samples is kept. Therefore, they are filtered using two sample filtering approaches to get rid of noisy samples, namely the LM score filtering [35] and the round-trip consistency (RTcons) [1]. In LM score filtering, the generated samples are sorted according to the probability $p(y|y_{<t}, x)$ as given by the generation model and the top n samples are kept (we use $n = 5$). Regarding the RTcons filtering, an MRQA model is used to assess the generated question-answer pair: The generated question along with the context is fed into the MRQA model to predict the answer. The sample is then discarded if the predicted answer does not match the generated one. Since RTcons might discard all samples generated from a context, we fine-tune the MRQA model (previously fine-tuned on the source domain) used for RTcons filtering also on the available samples from the target domain.

4.3 MRQA Model

In all cases (i.e., for scoring with RT, for filtering with RTcons and for the eventual downstream task), we use `bert-base-uncased` [7] for the MRQA model similar to the setup in [35]. For both, training and inference, the following configuration applies: A maximum input length of 512 tokens and a context stride of 128 is used. Also, similar to the data generation model, when considering the TechQA dataset, we truncate the questions in the input.

Similar to the data generation model, we borrow hyperparameters from previous work [35], therefore training the MRQA model with AdamW Optimizer [24], a learning rate of $3e-5$ and a batch size of 24. Warm-up was disabled for fine-tuning the MRQA model.

In case of fine-tuning on target domain data, model training turned out to be quite fluctuating in terms of the evaluation score. Therefore, we performed a hyperparameter search w.r.t. F1 score including learning rate, warm-up steps (using linear scheduler), L2 regularization, pre-trained weight decay for the encoder and for the output layer separately, freezing the encoder or its embedding, training of only top-n layers and re-initializing top-n layers. As a result of manual tuning, only pre-trained weight decay on the encoder with $\lambda = 1e-7$ was employed, while all layers were fine-tuned without re-initialization.

Since chunking is applied to the inputs and predictions are aggregated afterwards (by choosing the best span over all chunks' predictions), it may happen

that the answer of a sample is not completely included in a chunk’s context. We especially observe this for TechQA where answers tend to be longer, as the performance on the evaluation data is greatly degraded as the model never has the chance to predict the answer correctly for some samples.

For the downstream task, we have seen better results in preliminary experiments when first fine-tuning on generated data followed by fine-tuning on labeled target domain data in a subsequent step.

4.4 Baselines

In order to assess the performance of our approach, we compare it against several baselines. Since we aim at comparing the use of humans at different stages within the process of model development instead of solely improving MRQA performance, we compare our approach to other, effective approaches utilizing humans (b_3 , b_5) while baselines b_1 , b_2 and b_4 serve as a measure of human impact.

In absence of AL, our first baseline (b_1) consists of an MRQA model trained only on randomly selected labeled target domain data (again 200 samples). In addition, b_2 is additionally trained on SQuAD first.

The last baseline without data generation (b_3) differs from b_2 in that it does not use random target domain data but rather employs AL. For this purpose, we consider Bayesian Active Learning by Disagreement (BALD) [14,11] to score target domain samples. BALD measures the uncertainty of the model’s predictions, using entropy H given training data \mathcal{D} and model weights θ :

$$H[y|x, \mathcal{D}] - \mathbb{E}_{\theta \sim p(\theta|\mathcal{D})} H[y|x, \theta]. \quad (4)$$

In order to make computation feasible, we calculate the BALD score (using dropout) independently for start and end probabilities. The sum of both scores is then used to rate and order the unlabeled samples where higher scores are preferred.

To research our approach employing AL at the data generation model, we also consider two baselines where data generation is used. While b_4 selects samples randomly, b_5 again selects samples using BALD after synthetic data generation.

For reference, we also report results of MRQA models only trained on synthetic data involving SQuAD training data only (r_1) as well as on the full set of target domain training samples (r_2).

In all cases (but r_2), if target domain data is used, a total of 200 labeled target domain samples is selected.

5 Results & Discussion

6 Analysis

To better understand the performance of the AL methods considered in our data generation setup, we analyze different aspects of the selected samples.

7 Conclusion

Although data scarcity is ubiquitous in practice, there is a lack of approaches that address this issue for the MRQA task in specialized domains. Therefore, appropriate methods applicable to realistic low-resource domains are of great importance since deep learning models usually work well when plenty of annotated, in-domain samples are available. Our work shows how to utilize Active

Learning to boost performance when addressing the challenging MRQA problem in low-resource, domain-specific scenarios relevant in practice. To this end, we introduced a novel approach that combines data generation with AL, resulting in a boost in performance with low labeling effort. We further analyzed the performance of AL when applied to the MRQA model directly compared to when applied to the data generation process, showing evidenced improvements in predictive performance in the latter scenario, especially when using our newly introduced scoring function tailored to the MRQA task.

References

1. Alberti, C., Andor, D., Pitler, E., Devlin, J., Collins, M.: Synthetic QA Corpora Generation with Roundtrip Consistency (Jun 2019)
2. Banerjee, S., Lavie, A.: METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In: Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization. pp. 65–72. Association for Computational Linguistics, Ann Arbor, Michigan (Jun 2005)
3. Castelli, V., Chakravarti, R., Dana, S., Ferritto, A., Florian, R., Franz, M., Garg, D., Khandelwal, D., McCarley, S., McCawley, M., Nasr, M., Pan, L., Pendus, C., Pitrelli, J., Pujar, S., Roukos, S., Sakrajda, A., Sil, A., Uceda-Sosa, R., Ward, T., Zhang, R.: The TechQA Dataset (Nov 2019)
4. Chang, H.S., Vembu, S., Mohan, S., Uppaal, R., McCallum, A.: Using Error Decay Prediction to Overcome Practical Issues of Deep Active Learning for Named Entity Recognition. *Machine Learning* **109**(9-10), 1749–1778 (Sep 2020)

5. Chen, Y., Wu, L., Zaki, M.J.: Reinforcement Learning Based Graph-to-Sequence Model for Natural Question Generation (Aug 2020)
6. Cohn, D.A., Ghahramani, Z., Jordan, M.I.: Active Learning with Statistical Models (Feb 1996)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding (May 2019)
8. Fang, M., Li, Y., Cohn, T.: Learning how to Active Learn: A Deep Reinforcement Learning Approach (Aug 2017)
9. Fomicheva, M., Sun, S., Yankovskaya, L., Blain, F., Guzmán, F., Fishel, M., Altraras, N., Chaudhary, V., Specia, L.: Unsupervised Quality Estimation for Neural Machine Translation (Jul 2020)
10. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning (Oct 2016)
11. Gal, Y., Islam, R., Ghahramani, Z.: Deep Bayesian Active Learning with Image Data (Mar 2017)
12. Hedderich, M.A., Lange, L., Adel, H., Strötgen, J., Klakow, D.: A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios (Mar 2021)
13. Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The Curious Case of Neural Text Degeneration (Feb 2020)
14. Houthby, N., Huszár, F., Ghahramani, Z., Lengyel, M.: Bayesian Active Learning for Classification and Preference Learning (Dec 2011)
15. Klein, T., Nabi, M.: Learning to Answer by Learning to Ask: Getting the Best of GPT-2 and BERT Worlds (Nov 2019)
16. Kratzwald, B., Feuerriegel, S., Sun, H.: Learning a Cost-Effective Annotation Policy for Question Answering (Nov 2020)
17. Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Kelcey, M., Devlin, J., Lee, K., Toutanova, K.N., Jones, L., Chang, M.W., Dai, A., Uszkoreit, J., Le, Q., Petrov, S.: Natural Questions: a Benchmark for Question Answering Research. *Transactions of the Association of Computational Linguistics* (2019)
18. Lee, D.B., Lee, S., Jeong, W.T., Kim, D., Hwang, S.J.: Generating Diverse and Consistent QA pairs from Contexts with Information-Maximizing Hierarchical Conditional VAEs (Jun 2020)
19. Lee, S., Kim, D., Park, J.: Domain-agnostic Question-Answering with Adversarial Training (Oct 2019)
20. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension (Oct 2019)
21. Lin, X., Parikh, D.: Active Learning for Visual Question Answering: An Empirical Study (Nov 2017)
22. Liu, B., Wei, H., Niu, D., Chen, H., He, Y.: Asking Questions the Human Way: Scalable Question-Answer Generation from Text Corpus. *Proceedings of The Web Conference 2020* pp. 2032–2043 (Apr 2020)
23. Liu, P., Wang, L., Ranjan, R., He, G., Zhao, L.: A Survey on Active Deep Learning: From Model Driven to Data Driven. *ACM Computing Surveys* **54**(10s), 1–34 (Jan 2022)
24. Loshchilov, I., Hutter, F.: Decoupled Weight Decay Regularization (Jan 2019)
25. Lowell, D., Lipton, Z.C., Wallace, B.C.: Practical Obstacles to Deploying Active Learning (Nov 2019)

26. Luo, H., Li, S.W., Yu, S., Glass, J.: Cooperative Learning of Zero-Shot Machine Reading Comprehension (Mar 2021)
27. Nishida, K., Nishida, K., Saito, I., Asano, H., Tomita, J.: Unsupervised Domain Adaptation of Language Models for Reading Comprehension (May 2020)
28. Otegi, A., Agirre, A., Campos, J.A., Soroa, A., Agirre, E.: Conversational Question Answering in Low Resource Scenarios: A Dataset and Case Study for Basque. In: Proceedings of the Twelfth Language Resources and Evaluation Conference. pp. 436–442. European Language Resources Association, Marseille, France (May 2020)
29. Puri, R., Spring, R., Patwary, M., Shoeybi, M., Catanzaro, B.: Training Question Answering Models From Synthetic Data (Feb 2020)
30. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language Models are Unsupervised Multitask Learners. undefined (2019)
31. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: SQuAD: 100,000+ Questions for Machine Comprehension of Text (Oct 2016)
32. Schmidt, M., Bartezzaghi, A., Vu, N.T.: Prompting-based synthetic data generation for few-shot question answering. In: International Conference on Language Resources and Evaluation (2024)
33. Settles, B.: Active Learning. Synthesis Lectures on Artificial Intelligence and Machine Learning Series, Morgan & Claypool (2012)
34. Shakeri, S., Constant, N., Kale, M.S., Xue, L.: Towards Zero-Shot Multilingual Synthetic Question and Answer Generation for Cross-Lingual Reading Comprehension (May 2021)
35. Shakeri, S., Santos, C.N.d., Zhu, H., Ng, P., Nan, F., Wang, Z., Nallapati, R., Xiang, B.: End-to-End Synthetic Data Generation for Domain Adaptation of Question Answering Systems (Oct 2020)
36. Siddhant, A., Lipton, Z.C.: Deep Bayesian Active Learning for Natural Language Processing: Results of a Large-Scale Empirical Study (Sep 2018)
37. Sun, X., Liu, J., Lyu, Y., He, W., Ma, Y., Wang, S.: Answer-focused and Position-aware Neural Question Generation. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. pp. 3930–3939. Association for Computational Linguistics, Brussels, Belgium (Oct 2018)
38. Tsatsaronis, G., Balikas, G., Malakasiotis, P., Partalas, I., Zschunke, M., Alvers, M.R., Weissenborn, D., Krithara, A., Petridis, S., Polychronopoulos, D., Almirantis, Y., Pavlopoulos, J., Baskiotis, N., Gallinari, P., Artières, T., Ngomo, A.C.N., Heino, N., Gaussier, E., Barrio-Alvers, L., Schroeder, M., Androutsopoulos, I., Paliouras, G.: An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics* **16**(1), 138 (Apr 2015)
39. Tuan, L.A., Shah, D.J., Barzilay, R.: Capturing Greater Context for Question Generation (Oct 2019)
40. Ushio, A., Alva-Manchego, F., Camacho-Collados, J.: Generative Language Models for Paragraph-Level Question Generation (Jan 2023)
41. Van, H., Yadav, V., Surdeanu, M.: Cheap and Good? Simple and Effective Data Augmentation for Low Resource Machine Reading. Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval pp. 2116–2120 (Jul 2021)
42. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need (Dec 2017)
43. Wang, Y., Yao, Q., Kwok, J., Ni, L.M.: Generalizing from a Few Examples: A Survey on Few-Shot Learning (May 2019)

44. Xiao, T.Z., Gomez, A.N., Gal, Y.: Wat zei je? Detecting Out-of-Distribution Translations with Variational Transformers (Jun 2020)
45. Yin, X., Zhou, L., Small, K., May, J.: Summary-Oriented Question Generation for Informational Queries (Jul 2021)
46. Zhang, R., Reddy, R.G., Sultan, M.A., Castelli, V., Ferritto, A., Florian, R., Kayi, E.S., Roukos, S., Sil, A., Ward, T.: Multi-Stage Pre-training for Low-Resource Domain Adaptation (Oct 2020)