

Closed-book Question Generation via Contrastive Learning

Xiangjue Dong¹, Jiaying Lu^{2*}, Jianling Wang^{1*}, James Caverlee¹

¹ Texas A&M University, ² Emory University

{xj.dong, jlwang, caverlee}@tamu.edu, jiaying.lu@emory.edu

Abstract

Question Generation (QG) is a fundamental NLP task for many downstream applications. Recent studies on *open-book* QG, where supportive answer-context pairs are provided to models, have achieved promising progress. However, generating natural questions under a more practical *closed-book* setting that lacks these supporting documents still remains a challenge. In this work, we propose a new QG model for this closed-book setting that is designed to better understand the semantics of long-form abstractive answers and store more information in its parameters through contrastive learning and an answer reconstruction module. Through experiments, we validate the proposed QG model on both public datasets and a new WikiCQA dataset. Empirical results show that the proposed QG model outperforms baselines in both automatic evaluation and human evaluation. In addition, we show how to leverage the proposed model to improve existing question-answering systems. These results further indicate the effectiveness of our QG model for enhancing *closed-book* question-answering tasks.

1 Introduction

Question Generation (QG) has a wide range of applications, such as generating questions for exams (Jia et al., 2021; Lelkes et al., 2021; Dugan et al., 2022) or children’s story books (Zhao et al., 2022; Yao et al., 2022), recommending questions for users in a dialogue system (Shukla et al., 2019; Laban et al., 2020), improving visual (Li et al., 2018; Lu et al., 2022) or textual question-answering tasks (Duan et al., 2017; Lewis et al., 2019a; Zhang and Bansal, 2019; Sultan et al., 2020; Lyu et al., 2021), asking clarification questions (Rao and Daumé III, 2019; Yu et al., 2020; Ren et al., 2021), and generating queries for SQL (Wu et al., 2021) or multimodal documents (Kim et al., 2021).

Previous works on QG are mainly under the *open-book* setting, which aims to generate questions based on factoid or human-generated short answers under the assumption that there is access to external knowledge like retrieved documents or passages (Du et al., 2017; Zhao et al., 2018; Kim et al., 2019; Fei et al., 2021). After Roberts et al. (2020) demonstrated that feeding a large pre-trained model input questions alone *without* any external knowledge can lead to competitive results with retrieval-based methods on open-domain question-answering benchmarks, there is an increasing interest in the *closed-book* setting. This closed-book setting is appealing in practice and can be widely applied, e.g., in question suggestion (Laban et al., 2020; Yin et al., 2021), query recommendation (Kim et al., 2021), and other practical settings where extensive external knowledge is unavailable.

However, generating questions without access to such external knowledge is challenging for two key reasons. First, without access to retrieved documents (or passages), simple open-domain strategies like basing the answers on these documents (or passages) are not possible under the closed-book setting. Instead, models must rely on the answers alone. Second, the data used by most of the closed-book works (Lewis et al., 2021; Wang et al., 2021) are variants of existing open-domain datasets, e.g., SQuAD (Rajpurkar et al., 2018), TriviaQA (Joshi et al., 2017), WebQuestions (Berant et al., 2013) that ignore the answer-related passages. These answers in open-book works are usually short, e.g., entities, and easier to be remembered by the language model and stored in the parameters of the model than long-form answers. Thus, this leads to our motivating research question – *How can we empower a QG model to better understand the semantics of long-form abstractive answers and store more information in its parameters?*

To tackle the aforementioned challenges existing

* Equal Contribution

in the closed-book setting, this paper proposes a new QG model with two unique characteristics: (i) a contrastive learning loss designed to better understand the semantics of the answers and the semantic relationship between answers and ground-truth questions at a contextual-level; and (ii) an answer reconstruction loss designed to measure the answerability of the generated question. Contrastive learning has shown promising results in many NLP tasks, e.g., (Giorgi et al., 2021; Gao et al., 2021; Yang et al., 2021) and aligns positive pairs better with available supervised signals (Gao et al., 2021); here we show how to learn question representations by distinguishing features of correct question-answer pairs from features of incorrectly linked question-answer pairs. Further, to ensure the generated questions are of good quality and can be answered by the answer that is used for question generation, we frame the model as a generation-reconstruction process (Cao et al., 2019; Zhu et al., 2020), by predicting the original answers given the generated questions by a pre-trained seq2seq model. In addition, we introduce a new closed-book dataset with long-form abstractive answers – WikiCQA – to complement existing datasets like GooAQ (Khashabi et al., 2021) and ELI5 (Fan et al., 2019) and show how to leverage our model to generate synthetic data to improve closed-book question-answering tasks.

Through experiments, we find that the proposed QG model shows improvement through both automatic and human evaluation metrics on WikiCQA and two public datasets. Compared to the baseline, the proposed QG framework shows an improvement of up to 2.0%, 2.7%, and 1.8% on the ROUGE-L score on WikiCQA, GooAQ-S, and ELI5, respectively, and 1.3% and 2.6% in terms of relevance and correctness. Furthermore, we leverage the QG framework to generate synthetic QA data from WikiHow summary data and pre-train a closed-book QA model on it in both an unsupervised and semi-supervised setting. The performance is evaluated on both seen (WikiCQA) and unseen (GooAQ-S, ELI5) datasets. We find consistent improvements across these datasets, indicating the QG model’s effectiveness in enhancing closed-book question-answering tasks.

In conclusion, our contributions can be summarized as follows:

- We propose a contrastive QG model, which to our knowledge is the first work to explore con-

trastive learning for QG under a closed-book setting.

- The proposed model outperforms baselines on three datasets. The human evaluation also indicates that the questions generated by our model are more informative compared to other baselines.
- We leverage the QG model as a data augmentation strategy to generate large-scale QA pairs. Consistent improvements shown on both seen datasets and unseen datasets indicate the QG model’s effectiveness in enhancing closed-book question-answering tasks.

2 Related Work

Many previous works on QG are under the open-book setting, which takes factoid short answers (Rajpurkar et al., 2016) or human-generated short answers (Kočíský et al., 2018) with the corresponding passages to generate questions (Zhang et al., 2021). Early approaches for question generation rely on rule-based methods (Labutov et al., 2015; Khullar et al., 2018). To bypass hand-crafted rules and sophisticated pipelines in QG, Du et al. (2017) introduce a vanilla RNN-based sequence-to-sequence approach with an attention mechanism. The recently proposed pre-trained transformer-based frameworks (Lewis et al., 2020; Raffel et al., 2020) also improve the performance of QG. In addition, Sultan et al. (2020) shows that the lexical and factual diversity of QG provides better QA training. However, their success can not directly adapt to the closed-book setting, where the model is supposed to generate questions solely relying on answers. In this work, we explore the widely applicable closed-book QG setting, which is still under-explored.

Contrastive Learning aims to pull semantically similar neighbors close and push non-neighbors apart. It has achieved great success under both supervised and unsupervised settings. In pioneer works, the contrastive loss function (Hadsell et al., 2006; Chopra et al., 2005) has been proposed as a training objective in deep metric learning considering both similar and dissimilar pairs. Recently, Chen et al. (2020) proposes the SimCLR framework to learn useful visual representations. Viewing contrastive learning as dictionary look-up, He et al. (2020) present Momentum Contrast (MoCo) to build dynamic dictionaries for contrastive learning. Some works apply contrastive learning into

the NLP domain to learn better sentence representations (Giorgi et al., 2021; Gao et al., 2021). In addition, contrastive learning has been applied in multilingual neural machine translation (Pan et al., 2021), abstractive summarization (Liu and Liu, 2021), and multi-document question generation (Cho et al., 2021). The recent most relevant work is (Yang et al., 2021), where they design two contrastive losses for paraphrase generation. In this work, we adopt contrastive learning for improving representation learning in question generation under a closed-book setting.

3 Proposed Approach

To answer our research question – How can we empower a QG model to better understand the semantics of long-form abstractive answers and store more information in its parameters? – we propose a closed-book QG model, which generates questions directly without access to external knowledge. Formally, given an answer sentence x , a closed-book QG engine generates a natural question y . Figure 1 illustrates an overview of the proposed QG framework, which consists of three parts: question generation, contrastive learning, and answer reconstruction. The framework is optimized with the joint losses from these three parts simultaneously.




Figure 1: An overview of the proposed closed-book QG framework, which consists of three parts: contrastive learning, question generation, and answer construction. A_i : represents answer i ; Q_i represents question i .

3.1 Question Generation

We first focus on question generation through a sequence-to-sequence architecture which consists of an encoder and a decoder (Sutskever et al., 2014; Vaswani et al., 2017). The encoder takes an input sequence of source words $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and maps it to a sequence of continuous representations $\mathbf{z} = (z_1, z_2, \dots, z_n)$. Then, the decoder takes \mathbf{z} and generates a sequence of target words $\mathbf{y} = (y_1, y_2, \dots, y_m)$ at a time. The closed-book

QG task is defined as finding $\hat{\mathbf{y}}$:

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} P(\mathbf{y}|\mathbf{x}), \quad (1)$$

where $P(\mathbf{y}|\mathbf{x})$ is the conditional likelihood of the predicted question sequence \mathbf{y} given answer \mathbf{x} .

$$P(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^T p(y_i | \mathbf{y}_{<t}, \mathbf{x}), \quad (2)$$

Given the answer-question pairs, the training objective of the generation part in the proposed framework is to minimize the Negative Log-Likelihood (NLL) of the training data,

$$\mathcal{L}_{qg} = - \sum_{i=1}^N \log p(q_i | A), \quad (3)$$

where q_i is the i -th token in the generated question and A is the answer.

A naive question generation model will generate questions based on answers but lacks a rich model of the semantics of answers nor can it guarantee the generated questions have a semantic relationship with the answers. Intuitively, an encoded answer should be similar to its question and dissimilar to others. In addition, the generated question should be able to be answered by the answers. Hence, this motivates the following contrastive learning and answer reconstruction modules.

3.2 Contrastive Learning

Contrastive learning aims to pull positive pairs and push apart negative pairs to learn effective representations. Further, the supervised signals can produce better sentence embeddings by improving alignment between positive pairs (Chen et al., 2020). An effective QG model should be able to understand the semantics of the answers and the semantic relationship with the ground-truth questions. Especially, the encoded answer should have semantic similarity with its ground-truth question and dissimilarity with other questions. Thus, aiming to learn a similarity function that pulls the distance between the answer sequence representation and its ground-truth question sequence representation closer, we design a contrastive loss in the representation space. Specifically, given a positive pair $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, where \mathbf{x}_i and \mathbf{y}_i are semantically related inputs, the other $2(n-1)$ examples

within a mini-batch are treated as negative examples. The training objective for $(\mathbf{x}_i, \mathbf{y}_i)$ is:

$$\mathcal{L}_{cl} = -\log \frac{\exp(\text{sim}(\mathbf{z}_{x_i}, \mathbf{z}_{y_i})/\tau_{cl})}{\sum_{i=1}^{2n} \exp(\text{sim}(\mathbf{z}_{x_i}, \mathbf{z}_{y_i})/\tau_{cl})}, \quad (4)$$

where \mathbf{z}_{x_i} and \mathbf{z}_{y_i} is the representation of input x_i and y_i , $\text{sim}(\mathbf{z}_i, \mathbf{z}_j) = \mathbf{z}_i^\top \mathbf{z}_j / \|\mathbf{z}_i\| \|\mathbf{z}_j\|$ denotes cosine similarity, and τ_{cl} is a temperature parameter.

In this work, aiming to learn better answer representations and force the encoder to drive representations of correct question-answer pairs closer than representations of incorrect question-answer pairs, we take the ground-truth question as the positive instance of an answer and fine-tune the model parameters based on the contrastive loss function (Eq. 4), where \mathbf{z} is the embedding of the special token [CLS] from the transformer encoder, representing the meaning of the entire sentence.

3.3 Answer Reconstruction

The questions that are generated by the model should be of good quality and should also be able to be answered by the answer that is used for question generation. To measure the answerability of the generated question, we design an answer reconstruction module, which uses a pre-trained seq2seq model to predict the original answer given the generated question. The loss is calculated by a negative log-likelihood loss function:

$$\mathcal{L}_{ar} = -\sum_{i=1}^N \log p(a_i|Q) \quad (5)$$

where a_i is the i -th token in the answer and Q is the generated question.

A major challenge is that the generated questions from Section 3.1 are not differentiable. Gradients cannot be back-propagated directly. To solve this challenge, we employ the Straight-Through (ST) Gumbel-Softmax for gradient computation (Jang et al., 2017). The ST Gumbel-Softmax is a discrete version of the Gumbel-Softmax and takes different forward and backward paths. In the forward pass, the embedding is discretized by using the argmax, whereas in the backward pass the gradients are computed by the gumbel-softmax (Qader et al., 2019; Lu et al., 2021).

$$y_i = \frac{\exp((\log(p_i) + g_i)/\tau_{gs})}{\sum_{j=1}^{|V|} \exp((\log(p_j) + g_j)/\tau_{gs})}, \quad (6)$$

where τ_{gs} is a temperature parameter and g_i is the Gumbel noise drawn from a uniform distribution

$(0, 1)$. In this work, the one-hot embedding from ST Gumbel-Softmax is multiplied with the vocabulary embedding and then fed into the encoder of the pre-trained seq2seq model as the representation of the generated question.

3.4 Overall Loss Function

As a result, all three losses are summed together to provide the overall loss function \mathcal{L} as follows:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{gg} + \lambda_2 \mathcal{L}_{cl} + \lambda_3 \mathcal{L}_{ar}, \quad (7)$$

The weights λ_1, λ_2 , and λ_3 are tuneable hyperparameters to balance losses and the final objective is to minimize the overall loss.

The overall structure of the proposed QG framework is presented in Algorithm 1.

Algorithm 1: QG framework.

Input: Pre-trained language model $p(q|a)$,
answer reconstruction model $p(a|q)$
and answer-question pairs

Output: Question generator $p(q|a)$

```

1 for  $i \leftarrow 1$  to  $Epoch$  do
2    $\hat{q} = p(q|a)$ 
3   Compute  $\mathcal{L}_{gg}$  via Eq. 3
4   /* contrastive learning */
5    $a_i = \text{Encoder}([\text{CLS}] \oplus a)[:, 0 :]$ 
6    $a_i^+ = \text{Encoder}([\text{CLS}] \oplus q)[:, 0 :]$ 
7   Get  $\mathcal{L}_{cl}$  to  $(a_i, a_i^+)$  via Eq. 4
8   /* answer reconstruction */
9    $\hat{a} = p(a|\hat{q})$ 
10  Compute  $\mathcal{L}_{ar}$  via Eq. 5
11  Calculate total loss  $\mathcal{L}$  via Eq. 7
12  Update generator  $p(q|a)$  with  $\mathcal{L}$ 
13 end
14 return question generator  $p(q|a)$ 

```

4 Experimental Setup

To evaluate the effectiveness of the proposed QG framework, we aim to answer the following research questions (RQ) via experiments: **RQ1:** Can this proposed QG framework improve the performance of the closed-book QG tasks? **RQ2:** Are the generated questions of good quality? That is, are they fluent and relevant to the answer? **RQ3:** Can this QG framework be leveraged as a good resource to generate synthetic data for the QA task? **RQ4:** How much does each component in the framework contribute?

4.1 Dataset

We conduct the experiments on two public datasets – GooAQ-S (Khashabi et al., 2021) and ELI5 (Fan et al., 2019) – and a new dataset we curate called WikiCQA.

GooAQ-S is a sub-sampled dataset from GooAQ which contains three different sub-tasks: short, snippet (i.e., multi-sentence description), and collection response questions (Khashabi et al., 2021). Under the long-form closed-book setting, we adopt the snippet part (i.e., questions with snippet answers) for our experiments as in the original paper. Furthermore, from the paper, the model performance on the snippet task does not vary much when supervised with 200K and 2M training instances. Thus, to improve the experimental efficiency, we take 200k instances from the snippet set and split them based on their original scripts.¹

ELI5 (Fan et al., 2019) is a widely-used large-scale corpus for long-form question-answering with supporting web documents. In this work, we use the question-answer pairs from the dataset and ignore the supporting context to fit the closed-book setting like (Khashabi et al., 2021). We follow the data split from huggingface.²

WikiCQA is a new closed-book long-form QA dataset (Section 4.2) introduced here. It contains 20,202 question-answer pairs and is collected from a wiki-style website to complement existing datasets. We shuffle the data and split it to train, dev, and test sets by the ratio of 80%/10%/10%. Table 1 shows detailed data statistics and splits of these three datasets.

Datasets	Train	Val	Test
WikiCQA	16,162	2,020	2,020
GooAQ-S	200,000	500	498
ELI5	272,634	9,812	24,512

Table 1: Dataset statistics.

4.2 WikiCQA

WikiCQA contains real user QA pairs collected from WikiHow³. WikiHow is a wiki-style website featuring over 200K how-to articles. Different

from the existing dataset, ELI5, containing questions/answers from the Reddit forum and leveraging evidence queried from the web to help answer the question (Fan et al., 2019) and GooAQ, containing questions from search auto-complete and answers from answer boxes (Khashabi et al., 2021), WikiCQA involves long-form question-answering grounded on WikiHow articles.

Dataset Construction We construct the new dataset by collecting the question-answer pairs from the Q&A section of articles on WikiHow. The questions and answers are related to the specific articles, asked by WikiHow users, and answered by a knowledgeable reader or editor⁴. The questions can not be answered directly by the content of the article. After removing duplicates and question-answer pairs with meaningless answers, 23,037 QA pairs remain. To keep the same format as other existing QA datasets (e.g., ELI5, GooAQ), we discard questions not starting with a question-type word (e.g., what, how). After the dataset processing steps, we arrive at 20,202 question-answer pairs. More details can be found in Appendix A.

In Table 2, we show an example article from WikiHow, which includes the title and summary of an article and question-answer pairs from the corresponding Q&A section. From the example, we can see that answers are abstractive and long-form, written by real users, and not contained within the context passage. These question-answer pairs from the Q&A section are collected in the newly constructed dataset. Thus, this dataset is different from reading comprehension datasets, where the answers are a short text span in context. More comparisons with ELI5 and GooAQ are shown in Appendix B.

Title: How to Prepare a Healthy Meal for Your Pet Dog.
Summary: To prepare a healthy meal for your dog, choose lean meat with the bones and fat removed, like chicken or beef ...

Q: What can I feed my dog if I have run out of dog food?
A: In the short term, any bland human food such as chicken or white fish with rice or pasta is just fine ...
Q: How much homemade dog food do you feed your dog?
A: Great question because it highlights one of the problems of feeding home prepared foods ...
Q: What should I not feed my dog?
A: There are many human foods that are toxic to dogs. Top of the list of foods NOT to give are ...

Table 2: Question-answer pairs from WikiHow Q&A section. **Q:** question; **A:** answer.

¹https://github.com/allenai/gooaq/blob/main/experiments/create_splits.py

²<https://huggingface.co/datasets/eli5>

³<https://www.wikihow.com>, under an Attribution-Noncommercial-Share Alike 3.0 Creative Commons License.

⁴<https://www.wikihow.com/Use-wikiHow#Reading-and-Learning-from-wikiHow>

4.3 Models & Hyper-parameters

We use BART-base (Lewis et al., 2020), a widely used sequence-to-sequence framework with 12 layers and a hidden size of 1024, as the backbone model. Following previous works (Fan et al., 2019; Khashabi et al., 2021), we finetune it using answers as inputs and questions as outputs as the baselines. To get the answer reconstruction loss (Section 3.3), we use a BART-base model which is fine-tuned on the target QA datasets.

Based on the dataset analysis in Appendix B, we set the maximum sequence length to 128 for the question and 256 for the answer sentence to improve calculation efficiency. We train the models for five epochs with a learning rate of 5×10^{-5} and evaluate the checkpoint for each epoch. We select the checkpoint with the highest ROUGE-L score on the validation set and report its corresponding score on the test set. We run each model three times and record the average scores. After performing manual hyper-parameter searching, we set loss parameters λ_1 , λ_2 , and λ_3 in Equation 7 to 1.0, 0.1, 0.1, respectively, which give the best ROUGE-L score on validation set. The temperature of contrastive loss τ_{cl} is set to 0.3. The experiments are run on 1 NVIDIA Tesla V100 GPU. The training time takes about 48 hours.

4.4 Evaluation Metrics

To evaluate the performance of the models, we use **ROUGE** (Lin, 2004) scores, which evaluate the n -grams recall of generated sentences with reference sentences, as automatic evaluation metrics (Fan et al., 2019; Khashabi et al., 2021). We report the F1 for ROUGE-1, ROUGE-2, ROUGE-L, and ROUGE-Lsum. ROUGE-1, 2, L measures the unigram, bigram, and longest common subsequence between the pair of sentences, respectively. The difference between ROUGE-L and ROUGE-Lsum is that ROUGE-Lsum splits text using “\n”. The higher ROUGE score indicates higher similarity between generated questions and references.

We further perform human evaluations for the quality of generated questions in terms of **Fluency**: whether the questions are grammatically correct and fluent; **Relevance**: whether the questions are related to the answers; **Correctness**: whether the questions can be answered by the answers.

5 Experiment Results

In this section, we answer the four experimental research questions in turn.

5.1 Generation Performance (RQ1)

First, does the proposed QG framework have good performance? Table 3 shows the performance of the proposed closed-book QG model (denoted as \mathbf{QG}_{ours}) on the new dataset and two public datasets compared with the baseline model trained only with question generation loss (denoted as \mathbf{QG}_b). We observe that the proposed framework outperforms the baseline on three datasets for all ROUGE scores. For example, the performance increases by up to 2.0%, 2.7%, and 1.8% on the ROUGE-L score, respectively, which means the proposed framework can generate questions having a longer longest common subsequence (LCS) with ground-truth questions. We attribute these improvements to how contrastive learning pulls the answer representations closer to the ground-truth questions and thus, the generated questions have a higher chance to have similar words, phrases, or sentences with the ground-truth questions. The results demonstrate the effectiveness of our proposed QG framework.

Dataset	Model	R-1	R-2	R-L	R-Lsum
WikiCQA	\mathbf{QG}_b	48.39	26.84	46.08	46.16
	\mathbf{QG}_{ours}	49.22	27.79	46.98	47.08
GooAQ-S	\mathbf{QG}_b	44.26	19.73	41.11	41.08
	\mathbf{QG}_{ours}	45.26	20.69	42.20	42.06
ELI5	\mathbf{QG}_b	28.62	10.10	25.93	26.23
	\mathbf{QG}_{ours}	29.15	10.36	26.40	26.69

Table 3: QG results on three datasets. \mathbf{QG}_b : baseline model; \mathbf{QG}_{ours} : our proposed QG framework.

5.2 Quality of Generated Questions (RQ2)

Next, are the generated questions of good quality? To answer this question, we perform human evaluation to measure the quality of generated questions in terms of *fluency*, *relevance*, and *correctness*. We randomly sample 100 answer-question pairs from the WikiCQA dataset, which contains answers, ground-truth questions, and generated questions from the baseline model and our best model. Then, we ask three annotators to rate the generated question pairs, comparing them with the ground-truth questions. The generated questions are rated on a 1-5 scale (5 for the best) from the aspects of the three aspects above. Further, we calculate the percentage of questions that have higher ratings in each group.

Table 4 shows the average scores and the percentages of preferred questions on the three criteria. Both models are equally good at generating fluent questions. In terms of relevance and correctness, our approach shows 1.3% and 2.6% higher scores than the baseline method. These are consistent with what we expect: the contrastive learning part pulls the answer representation closer to the ground-truth question representation and generates more relevant questions; the reconstruction part can ensure that generated questions are of good quality and answerable. In addition, among 100 answer-question pairs, 24.7% and 23.3% of generated questions from our best model are rated higher while 56.3% and 57.4% of them have the same rating as those from baselines, which indicates that there is still substantial room to improve. To get a sense of the stability of the human evaluation results, we measure the inter-agreement among annotators using the Correlation Coefficient, which is 0.998, 0.944, and 0.976 in terms of these three aspects, showing excellent reliability.

Model	fluency		relevance		correctness	
	score	%	score	%	score	%
QG _b	4.83	8.3	3.84	19.0	3.49	19.3
QG _{ours}	4.83	7.3	3.89	24.7	3.58	23.3

Table 4: Results of human evaluation for baseline and our best model. **score**: is the average score from raters; **%**: represents the proportion of generated questions from one model that were rated higher than those from its counterpart.

5.3 Synthetic Data Generation (RQ3)

Further, can this QG framework be leveraged to generate effective synthetic data that can improve the closed-book QA task? Data augmentation is one of the main directions that question generation has been used for previously, with several studies finding improvements on the QA task (Lewis et al., 2019b; Alberti et al., 2019). Here, we show how to leverage the proposed QG framework to improve closed-book QA tasks on seen data (WikiCQA) and unseen data (GooAQ and ELI5). Since freely available summary data is a good resource to generate synthetic data (Lyu et al., 2021), we use WikiSum (Cohen et al., 2021), which contains 39,775 coherent-paragraph summaries written by the article’s authors on the WikiHow website. We take each sentence from the article summary as an answer and pass it into the best QG model, described in Section 3.1, to generate a question. Then,

we train an unsupervised QA model based on the synthetic QA pairs and optimize it by the following negative log-likelihood loss function:

$$\mathcal{L}_{qa} = - \sum_{i=1}^N \log p(a_i|Q), \quad (8)$$

where a_i is the i -th token in the generated answer and Q is the question.

In this work, we pre-train a BART-base model on the 200K synthetic QA pairs that are generated through the best QG model (denoted as QA_s) and evaluate it on the test set of the seen dataset (WikiCQA) and unseen datasets (GooAQ-S and ELI5). This approach is unsupervised since the model is trained on no labeled question-answer pairs. The results are summarized in Table 5a, showing 22.4%, 12.5%, and 11.4% improvement than the BART-base model without any synthetic pre-training (denoted as QA_b), on WikiCQA, GooAQ-S, and ELI5, respectively. This shows that pre-training on the generated question-answer pairs derived from our QG model leads to significant improvements, echoing findings in previous works that find synthetic data can be helpful (Khashabi et al., 2021; Lewis et al., 2021; Ding et al., 2021). After further fine-tuning on the target training set (denoted as QA_{s+f}), from Table 5b we can see that QA_{s+f} also achieves better results than fine-tuned baselines model QA_{b+f} by 3.6%, 0.4%, and 4.3% on the same three datasets.

Dataset	Model	R-1	R-2	R-L	R-Lsum
WikiCQA	QA _b	18.41	5.53	14.72	15.98
	QA _s	24.10	7.07	18.02	20.31
GooAQ-S	QA _b	18.58	5.80	14.77	15.66
	QA _s	21.93	6.12	16.62	18.35
ELI5	QA _b	12.38	2.19	8.98	10.63
	QA _s	13.97	2.29	10.00	11.77

(a) Unsupervised QA results.

Dataset	Model	R-1	R-2	R-L	R-Lsum
WikiCQA	QA _{b+f}	27.16	7.32	19.85	23.49
	QA _{s+f}	28.44	8.14	20.56	24.77
GooAQ-S	QA _{b+f}	27.72	7.68	20.08	23.64
	QA _{s+f}	28.44	7.73	20.17	24.12
ELI5	QA _{b+f}	22.06	3.93	13.93	19.72
	QA _{s+f}	23.37	4.27	14.53	20.88

(b) Semi-supervised QA results.

Table 5: Evaluation on QA tasks. QA_{s/b}: QA model w/o pre-training on synthetic data; QA_{*+f}: QA* fine-tuned on target dataset.

5.4 Ablation Study (RQ4)

Finally, we investigate the contribution of each component in the proposed closed-book QG framework. Table 6 shows the results on the WikiCQA

dataset. As described in Section 3.2, the contrastive learning module takes ground-truth questions as positive pairs to the input answers (denoted as \mathbf{CL}_t). We also explore its variant – in which we feed the same input (answers) to the encoder with different dropouts to obtain the positive pairs (Gao et al., 2021) (denoted as \mathbf{CL}_s). We find that the choice of \mathbf{CL}_t is slightly better than \mathbf{CL}_s in most evaluation metrics. In addition, by adding the answer reconstruction (**AR**) module, the model performance can be further improved. Thus, we can observe the effectiveness of the proposed contrastive learning module and the overall design of the framework.

Models	R-1	R-2	R-L	R-LSum
QG _b	48.39	26.84	46.08	46.16
QG _b +CL _s	48.67	27.34	46.43	46.46
QG _b +CL _t	48.72	27.26	46.33	46.41
QG _b +CL _s +AR	49.26	27.77	46.82	46.88
QG _b +CL _t +AR	49.22	27.79	46.98	47.08

Table 6: Ablation study of our proposed QG framework on WikiCQA. \mathbf{CL}_* : contrastive learning module; **AR**: answer reconstruction module.

6 Case Study

Finally, we showcase some examples of the generated questions on the WikiCQA dataset. For the first example in Table 7a, we can see our model generates “how many” for the statistics-type answer, which is closer to the reference “what percentage” than the baseline model, generating a “yes/no” question. For the second one, although the reference contains the detailed information “textured surface” which is not in the answer, our model captures the information “without nails”, which is more informative than the baseline model. Thus, the proposed model can produce better questions that are more relevant to the answers and contain more detailed information than questions generated from baseline models. There are also some cases where the proposed model fails to generate good questions shown in Table 7b. Compared with the reference question “what food can be included on a clear liquid diet”, our model generates “what can I eat to lose weight”, which fails in automatic evaluation metrics and is hard for annotators to justify in terms of relevance. For the second one, although the “broken door” generated from our model is relevant to the answer, it’s still far from the reference “door sticking to the weather strip”. The question generated from the baseline model related to “bedroom” is even less relevant. These examples encourage us to explore how to capture more

A	The most recent statistics show that around 2.5% of small businesses are audited by the irs.
Q _B	Are small businesses audited by the irs?
Q _G	How many small businesses are audited by the irs?
Q _R	What percentage of small businesses are audited?
A	Try using poster putty to secure the images in your collage. if that doesn’t work, you may have to use nails.
Q _B	How do i put pictures in a collage?
Q _G	How do i make a collage without nails?
Q _R	How can i make a collage on a wall with a textured surface?

(a) Good examples.

A	Foods such as popsicles, hard candy, and gelatin can be eaten on a clear liquid diet.
Q _B	what can you eat on a clear liquid diet?
Q _G	What can I eat to lose weight?
Q _R	What food can be included on a clear liquid diet?
A	You’ll have to remove the door and sand, prime, and repaint it.
Q _B	What do i have to do to make my bedroom look nice?
Q _G	How do i fix a broken door?
Q _R	What do i do if my door is sticking to the weather strip?

(b) Bad examples.

Table 7: Generated question examples. \mathbf{Q}_B : generated questions from baseline model; \mathbf{Q}_G : generated questions from proposed QG model; \mathbf{Q}_R : reference questions.

semantics in the answers and generate questions that have more detailed information in the future.

7 Conclusion

In this work, aiming to empower a QG model to better understand the semantics of long-form abstract answers and store more information in its parameters, we propose a closed-book QG model empowered by a contrastive learning module and an answer reconstruction module. We present a new *closed-book* long-form QA dataset – WikiCQA involving more than 20K real user QA pairs and show that WikiCQA is a valuable training resource, complementing public datasets. Through the experiments, the proposed QG model shows better performance than baselines through automatic and human evaluations. Moreover, we show how to leverage the proposed model as a data augmentation strategy to improve existing *closed-book* QA systems. The closed-book QA model, pre-trained on our generated synthetic QA pairs, achieves better performance on the seen dataset (WikiCQA). In addition, it shows strong generalization on unseen datasets (GooAQ and ELI5), which further demonstrates the effectiveness of our QG framework for enhancing closed-book QA performance.

Limitations

While our model shows promising results on English datasets, its efficiency and performance in other languages require further investigation. Moreover, our evaluation is limited to datasets with one-to-one QA pairs, without considering situations where multiple answers correspond to a single question. Before feeding the input into the model, we ignore the tokens that exceed the maximum length we set, which may sometimes bring in an information loss for the input corpus. Additionally, our model is built on transformer architecture and therefore requires huge computation resources, especially for large training data sizes. Our focus in this paper is solely on the closed-book setting and our experiments center around four research questions to evaluate the effectiveness of the proposed framework, which we believe is an essential and emerging area deserving of specialized research investigation. In future work, we aim to explore the feasibility of adapting our framework for other QG scenarios.

Acknowledgements

We thank Zhuoer Wang, Yun He, Ziwei Zhu, anonymous reviewers, and the action editor for providing valuable feedback.

References

- 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.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Ruisheng Cao, Su Zhu, Chen Liu, Jieyu Li, and Kai Yu. 2019. [Semantic parsing with dual learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 51–64, Florence, Italy. Association for Computational Linguistics.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Woon Sang Cho, Yizhe Zhang, Sudha Rao, Asli Celikyilmaz, Chenyan Xiong, Jianfeng Gao, Mengdi Wang, and Bill Dolan. 2021. [Contrastive multi-document question generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 12–30, Online. Association for Computational Linguistics.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE.
- Nachshon Cohen, Oren Kalinsky, Yftah Ziser, and Alessandro Moschitti. 2021. [WikiSum: Coherent summarization dataset for efficient human-evaluation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 212–219, Online. Association for Computational Linguistics.
- Kaize Ding, Dingcheng Li, Alexander Hanbo Li, Xing Fan, Chenlei Guo, Yang Liu, and Huan Liu. 2021. Learning to selectively learn for weakly-supervised paraphrase generation. In *2021 Conference on Empirical Methods in Natural Language Processing*, pages 5930–5940.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. [Question generation for question answering](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 866–874, Copenhagen, Denmark. Association for Computational Linguistics.
- Liam Dugan, Eleni Miltsakaki, Shriyash Upadhyay, Etan Ginsberg, Hannah Gonzalez, DaHyeon Choi, Chuning Yuan, and Chris Callison-Burch. 2022. A feasibility study of answer-unaware question generation for education. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1919–1926.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.

- Zichu Fei, Qi Zhang, and Yaqian Zhou. 2021. [Iterative GNN-based decoder for question generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2573–2582, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- John Giorgi, Osvald Nitski, Bo Wang, and Gary Bader. 2021. [DeCLUTR: Deep contrastive learning for unsupervised textual representations](#). 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 879–895, Online. Association for Computational Linguistics.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *Conference ICLR*.
- Xin Jia, Wenjie Zhou, Xu Sun, and Yunfang Wu. 2021. [Eqg-race: Examination-type question generation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 13143–13151.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Daniel Khashabi, Amos Ng, Tushar Khot, Ashish Sabharwal, Hannaneh Hajishirzi, and Chris Callison-Burch. 2021. [GooAQ: Open question answering with diverse answer types](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 421–433, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Payal Khullar, Konigari Rachna, Mukul Hase, and Manish Shrivastava. 2018. [Automatic question generation using relative pronouns and adverbs](#). In *Proceedings of ACL 2018, Student Research Workshop*, pages 153–158, Melbourne, Australia. Association for Computational Linguistics.
- Kyungho Kim, Kyungjae Lee, Seung-won Hwang, Young-In Song, and Seungwook Lee. 2021. [Query generation for multimodal documents](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 659–668, Online. Association for Computational Linguistics.
- Yanghoon Kim, Hwanhee Lee, Joongbo Shin, and Kyomin Jung. 2019. Improving neural question generation using answer separation. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, pages 6602–6609.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Philippe Laban, John Canny, and Marti A. Hearst. 2020. [What’s the latest? a question-driven news chatbot](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 380–387, Online. Association for Computational Linguistics.
- Igor Labutov, Sumit Basu, and Lucy Vanderwende. 2015. [Deep questions without deep understanding](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 889–898, Beijing, China. Association for Computational Linguistics.
- Adam D Lelkes, Vinh Q Tran, and Cong Yu. 2021. [Quiz-style question generation for news stories](#). In *Proceedings of the Web Conference 2021*, pages 2501–2511.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ludovic Denoyer, and Sebastian Riedel. 2019a. Unsupervised question answering by cloze translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4896–4910.
- Patrick Lewis, Ludovic Denoyer, and Sebastian Riedel. 2019b. [Unsupervised question answering by cloze](#)

- translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4896–4910, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. [PAQ: 65 million probably-asked questions and what you can do with them](#). *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Yikang Li, Nan Duan, Bolei Zhou, Xiao Chu, Wanli Ouyang, Xiaogang Wang, and Ming Zhou. 2018. Visual question generation as dual task of visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yixin Liu and Pengfei Liu. 2021. [SimCLS: A simple framework for contrastive learning of abstractive summarization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 1065–1072, Online. Association for Computational Linguistics.
- Jiaying Lu, Xiangjue Dong, and Carl Yang. 2021. Weakly supervised concept map generation through task-guided graph translation. *arXiv preprint arXiv:2110.15720*.
- Jiaying Lu, Xin Ye, Yi Ren, and Yezhou Yang. 2022. Good, better, best: Textual distractors generation for multiple-choice visual question answering via reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (O-DRUM Workshop)*, pages 4921–4930.
- Chenyang Lyu, Lifeng Shang, Yvette Graham, Jennifer Foster, Xin Jiang, and Qun Liu. 2021. [Improving unsupervised question answering via summarization-informed question generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4134–4148, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiao Pan, Mingxuan Wang, Liwei Wu, and Lei Li. 2021. [Contrastive learning for many-to-many multilingual neural machine translation](#). 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 244–258, Online. Association for Computational Linguistics.
- Raheel Qader, François Portet, and Cyril Labbé. 2019. [Semi-supervised neural text generation by joint learning of natural language generation and natural language understanding models](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 552–562, Tokyo, Japan. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67.
- 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.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Sudha Rao and Hal Daumé III. 2019. [Answer-based Adversarial Training for Generating Clarification Questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 143–155, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xuhui Ren, Hongzhi Yin, Tong Chen, Hao Wang, Zi Huang, and Kai Zheng. 2021. [Learning to ask appropriate questions in conversational recommendation](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’21*, page 808–817, New York, NY, USA. Association for Computing Machinery.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Pushkar Shukla, Carlos Elmadjian, Richika Sharan, Vivek Kulkarni, Matthew Turk, and William Yang Wang. 2019. [What should I ask? using conversationally informative rewards for goal-oriented visual dialog](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6442–6451, Florence, Italy. Association for Computational Linguistics.
- Md Arafat Sultan, Shubham Chandel, Ramón Fernández Astudillo, and Vittorio Castelli. 2020. [On the importance of diversity in question generation for QA](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages

- 5651–5656, Online. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Cunxiang Wang, Pai Liu, and Yue Zhang. 2021. [Can generative pre-trained language models serve as knowledge bases for closed-book QA?](#) 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 3241–3251, Online. Association for Computational Linguistics.
- Kun Wu, Lijie Wang, Zhenghua Li, Ao Zhang, Xinyan Xiao, Hua Wu, Min Zhang, and Haifeng Wang. 2021. Data augmentation with hierarchical sql-to-question generation for cross-domain text-to-sql parsing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8974–8983.
- Haoran Yang, Wai Lam, and Piji Li. 2021. [Contrastive representation learning for exemplar-guided paraphrase generation.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4754–4761, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Bingsheng Yao, Dakuo Wang, Tongshuang Sherry Wu, Zheng Zhang, Toby Jia-Jun Li, Mo Yu, and Ying Xu. 2022. It is ai’s turn to ask humans a question: Question-answer pair generation for children’s story books. In *ACL*.
- Xusen Yin, Li Zhou, Kevin Small, and Jonathan May. 2021. Summary-oriented question generation for informational queries. In *Proceedings of the 1st Workshop on Document-grounded Dialogue and Conversational Question Answering (DialDoc 2021)*, pages 81–97.
- Lili Yu, Howard Chen, Sida I. Wang, Tao Lei, and Yoav Artzi. 2020. [Interactive classification by asking informative questions.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2664–2680, Online. Association for Computational Linguistics.
- Ruqing Zhang, Jiafeng Guo, Lu Chen, Yixing Fan, and Xueqi Cheng. 2021. A review on question generation from natural language text. *ACM Transactions on Information Systems (TOIS)*, 40(1):1–43.
- Shiyue Zhang and Mohit Bansal. 2019. Addressing semantic drift in question generation for semi-supervised question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2495–2509.
- Yao Zhao, Xiaochuan Ni, Yuanyuan Ding, and Qifa Ke. 2018. [Paragraph-level neural question generation with maxout pointer and gated self-attention networks.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3901–3910, Brussels, Belgium. Association for Computational Linguistics.
- Zhenjie Zhao, Yufang Hou, Dakuo Wang, Mo Yu, Chengzhong Liu, and Xiaojuan Ma. 2022. [Educational question generation of children storybooks via question type distribution learning and event-centric summarization.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5073–5085, Dublin, Ireland. Association for Computational Linguistics.
- Su Zhu, Ruisheng Cao, and Kai Yu. 2020. Dual learning for semi-supervised natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1936–1947.

Appendix A Data Filtering

First, we choose the question-answer pairs where the questions start with a question-word in the list [‘how’, ‘what’, ‘can’, ‘is’, ‘do’, ‘why’, ‘are’, ‘does’, ‘where’, ‘when’, ‘should’, ‘will’, ‘did’, ‘which’, ‘who’, ‘would’, ‘if’, ‘about’, ‘for’, ‘as’, ‘could’, ‘in’, ‘after’, ‘at’, ‘while’, ‘to’, ‘am’, ‘has’, ‘any’] and end with a question mark. Then, we filter out the pairs where the length of answers is less than 8 tokens or the answers are meaningless. For the “meaningless” answers we mean, answers are “please refer to certain website <http://xxx>” or “please refer to the article xxx”.

Appendix B Dataset Analysis

To better understand the content of WikiCQA in comparison with existing QA datasets, Figure 2 shows the distributions of question length, answer length, and common question word. The questions and answers are tokenized by BART (Lewis et al., 2020). From Figure 2a and 2b we could see more than 80% of the questions in WikiCQA and GooAQ-S have less than 15 tokens and more than 90% of answers of them lie in the range of 16-127 tokens while both questions and answers

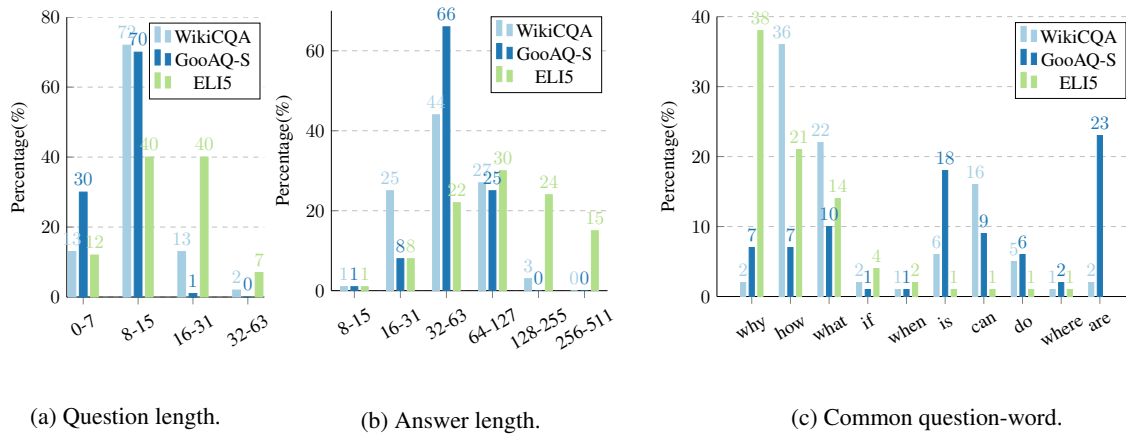


Figure 2: Comparison of the distribution of QA length and question types among three datasets.

in ELI5 have a broader range of lengths. By measuring the proportion of the leading unigram that starts a question (Figure 2c), the dominant pattern is “how” questions in WikiCQA, while ELI5 has many “why” questions. We think WikiCQA has the potential to benefit the community in three ways. Firstly, its source is unique compared to the current two closed-book data sources, as it is collected from the Q&A section of articles on WikiHow. The questions and answers are related to the specific articles, asked by WikiHow users who read the articles and answered by a knowledgeable reader or editor. Secondly, compared to large-scale QA datasets, WikiCQA has longer answers and more open-ended questions with a wider range of question types. Additionally, it has a different data distribution compared to the other two datasets, which is crucial for testing the generalization ability of question-answering models. Finally, the question-answer pairs within relevant articles in WikiCQA can also be used as a valuable resource for conversational question-answering tasks. Our code and data can be found at <https://github.com/dongxiangjue/Closed-book-Question-Generation-via-Contrastive-Learning>.