

Learn to Resolve Conversational Dependency: A Consistency Training Framework for Conversational Question Answering

Gangwoo Kim Hyunjae Kim Jungsoo Park Jaewoo Kang[†]

Korea University

{gangwoo_kim, hyunjae-kim}@korea.ac.kr

{jungsoo_park, kangj}@korea.ac.kr

Abstract

One of the main challenges in conversational question answering (CQA) is to resolve the conversational dependency, such as anaphora and ellipsis. However, existing approaches do not explicitly train QA models on how to resolve the dependency, and thus these models are limited in understanding human dialogues. In this paper, we propose a novel framework, EXCORD (Explicit guidance on how to resolve **C**onversational **D**ependency) to enhance the abilities of QA models in comprehending conversational context. EXCORD first generates self-contained questions that can be understood without the conversation history, then trains a QA model with the pairs of original and self-contained questions using a consistency-based regularizer. In our experiments, we demonstrate that EXCORD significantly improves the QA models' performance by up to 1.2 F1 on QuAC (Choi et al., 2018), and 5.2 F1 on CANARD (Elgohary et al., 2019), while addressing the limitations of the existing approaches.¹

1 Introduction

Conversational question answering (CQA) involves modeling the information-seeking process of humans in a dialogue. Unlike *single-turn* question answering (QA) tasks (Rajpurkar et al., 2016; Kwiatkowski et al., 2019), CQA is a *multi-turn* QA task, where questions in a dialogue are context-dependent;² hence they need to be understood with the conversation history (Choi et al., 2018; Reddy et al., 2019). As illustrated in Figure 1, to answer

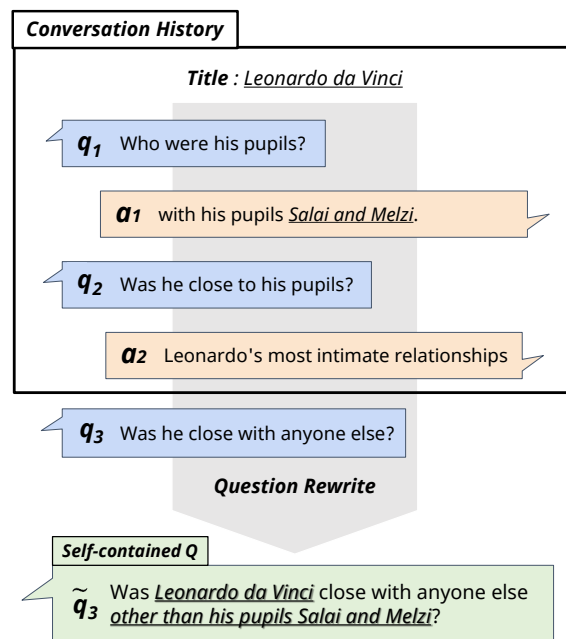


Figure 1: An example of the QuAC dataset (Choi et al., 2018). Owing to linguistic phenomena in human conversations, such as anaphora and ellipsis, the current question q_3 should be understood based on the conversation history: q_1 , a_1 , q_2 , and a_2 . Question q_3 can be reformulated as a self-contained question \tilde{q}_3 via a question rewriting (QR) process.

the current question “*Was he close with anyone else?*,” a model should resolve the conversational dependency, such as anaphora and ellipsis, based on the conversation history.

A line of research in CQA proposes the *end-to-end* approach, where a single QA model jointly encodes the evidence document, the current question, and the whole conversation history (Huang et al., 2018; Yeh and Chen, 2019; Qu et al., 2019a). In this approach, models are required to automatically learn to resolve conversational dependencies. However, existing models have limitations to do so without explicit guidance on how to resolve these dependencies. In the example presented in Figure

[†] Corresponding author

¹Our models and code are available at:
<https://github.com/dmis-lab/excord>

²While the term “context” usually refers to the evidence document from which the answer is extracted, in CQA, it refers to conversational context.

1, models are trained without explicit signals that “he” refers to “Leonardo da Vinci,” and “anyone else” can be more elaborated with “other than his pupils, Salai and Melzi.”

Another line of research proposes a *pipeline* approach that decomposes the CQA task into question rewriting (QR) and QA, to reduce the complexity of the task (Vakulenko et al., 2020). Based on the conversation history, QR models first generate self-contained questions by rewriting the original questions, such that the self-contained questions can be understood without the conversation history. For instance, the current question q_3 is reformulated as the self-contained question \tilde{q}_3 by a QR model in Figure 1. After rewriting the question, QA models are asked to answer the self-contained questions rather than the original questions. In this approach, QA models are trained to answer relatively simple questions whose dependencies have been resolved by QR models. Thus, this limits reasoning abilities of QA models for the CQA task, and causes QA models to rely on QR models.

In this paper, we emphasize that QA models can be enhanced by using both types of questions with explicit guidance on how to resolve the conversational dependency. Accordingly, we propose EXCORD (Explicit guidance on how to Resolve Conversational Dependency), a novel training framework for the CQA task. In this framework, we first generate self-contained questions using QR models. We then pair the self-contained questions with the original questions, and jointly encode them to train QA models with consistency regularization (Laine and Aila, 2016; Xie et al., 2019). Specifically, when original questions are given, we encourage QA models to yield similar answers to those when self-contained questions are given. This training strategy helps QA models to better understand the conversational context, while circumventing the limitations of previous approaches.

To demonstrate the effectiveness of EXCORD, we conduct extensive experiments on the three CQA benchmarks. In the experiments, our framework significantly outperforms the existing approaches by up to 1.2 F1 on QuAC (Choi et al., 2018) and by 5.2 F1 on CANARD (Elgohary et al., 2019). In addition, we find that our framework is also effective on a dataset CoQA (Reddy et al., 2019) that does not have the self-contained questions generated by human annotators. This indicates that the proposed framework can be adopted

on various CQA datasets in future work. We summarize the contributions of this work as follows:

- We identify the limitations of previous approaches and propose a unified framework to address these. Our novel framework improves QA models by incorporating QR models, while reducing the reliance on them.
- Our framework encourages QA models to learn how to resolve the conversational dependency via consistency regularization. To the best of our knowledge, our work is the first to apply the consistency training framework to the CQA task.
- We demonstrate the effectiveness of our framework on three CQA benchmarks. Our framework is model-agnostic and systematically improves the performance of QA models.

2 Background

2.1 Task Formulation

In CQA, a single instance is a dialogue, which consists of an evidence document d , a list of questions $\mathbf{q} = [q_1, \dots, q_T]$, and a list of answers for the questions $\mathbf{a} = [a_1, \dots, a_T]$, where T represents the number of *turns* in the dialogue. For the t -th turn, the question q_t and the conversation history $\mathcal{H}_t = [(q_1, a_1), \dots, (q_{t-1}, a_{t-1})]$ are given, and a model should extract the answer from the evidence document as:

$$\hat{a}_t = \arg \max_{a_t} P(a_t | d, q_t, \mathcal{H}_t) \quad (1)$$

where $P(\cdot)$ represents a likelihood function over all the spans in the evidence document, and \hat{a}_t is the predicted answer. Unlike single-turn QA, since the current question q_t is dependent on the conversation history \mathcal{H}_t , it is important to effectively encode the conversation history and resolve the conversational dependency in CQA.

2.2 End-to-end Approach

A naive approach in solving CQA is to train a model in an end-to-end manner (Figure 2a). Since standard QA models generally are ineffective in the CQA task, most studies attempt to develop a QA model structure or mechanism for encoding the conversation history effectively (Huang et al., 2018; Yeh and Chen, 2019; Qu et al., 2019a,b). Although

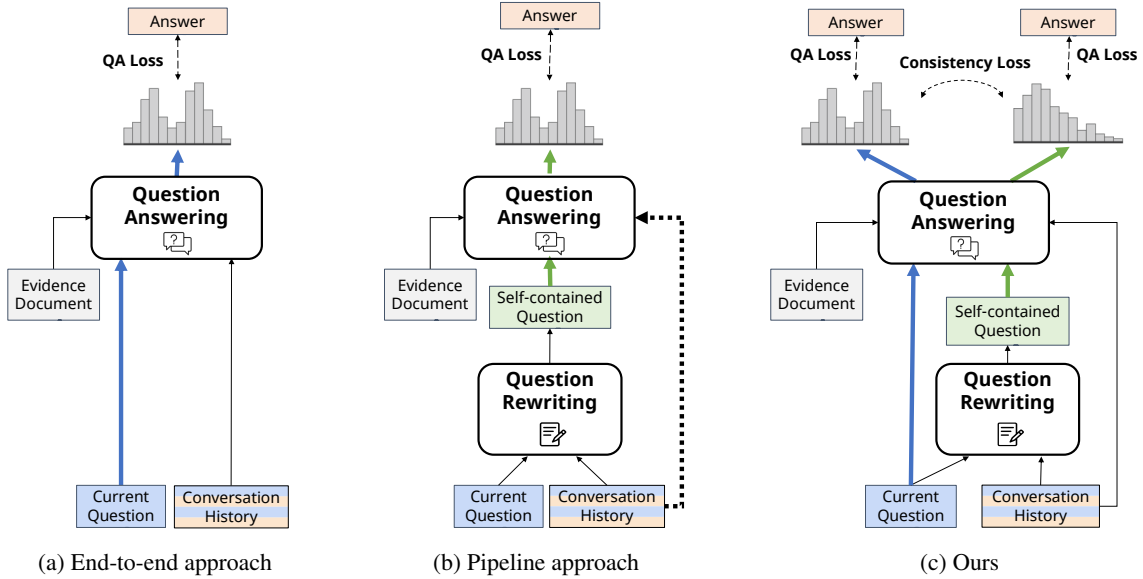


Figure 2: Overview of the end-to-end approach, the pipeline approach, and ours. In the end-to-end approach, QA models are asked to answer the original questions based on the conversation history. In the pipeline approach, the self-contained questions are generated by a QR model, and then QA models answer them. Standard QA models are commonly used in this approach; however conversational QA models that encode the history can be adopted (the dotted line in Figure (b)). In ours, the original and self-contained question are jointly encoded to train QA models with the consistency loss.

these efforts improved performance on the CQA benchmarks, existing models remain limited in understanding conversational context. In this paper, we emphasize that QA models can be further improved with explicit guidance using self-contained questions effectively.

2.3 Pipeline Approach

Recent studies decompose the task into two sub-tasks to reduce the complexity of the CQA task. The first sub-task, question rewriting, involves generating self-contained questions by reformulating the original questions. Neural-net-based QR models are commonly used to obtain self-contained questions (Lin et al., 2020; Vakulenko et al., 2020). The QR models are trained on the CANARD dataset (Elgohary et al., 2019), which consists of 40K pairs of original QuAC questions and their self-contained versions that are generated by human annotators.

After generating the self-contained questions, the next sub-task, question answering, is carried out. Since it is assumed that the dependencies in the questions have already been resolved by QR models, existing works usually use standard QA models (not specialized to CQA); however conversational QA models can also be used (the dotted line in Figure 2b). We formulate the process of

predicting the answer in the pipeline approach as:

$$P(a_t|d, q_t, \mathcal{H}_t) \approx P^{\text{rewr}}(\tilde{q}_t|q_t, \mathcal{H}_t) \cdot P^{\text{read}}(a_t|d, \tilde{q}_t) \quad (2)$$

where $P^{\text{rewr}}(\cdot)$ and $P^{\text{read}}(\cdot)$ are the likelihood functions of QR and QA models, respectively. \tilde{q}_t is a self-contained question rewritten by the QR model.

The main limitation of the pipeline approach is that QA models are never trained on the original questions, which limits their abilities to understand the conversational context. Moreover, this approach makes QA models dependent on QR models; hence QA models suffer from the error propagation from QR models.³ On the other hand, our framework enhances QA models’ reasoning abilities for CQA by jointly utilizing original and self-contained questions. In addition, QA models in our framework do not rely on QR models at inference time and thus do not suffer from error propagation.

3 EXCORD: Explicit Guidance on Resolving Conversational Dependency

We introduce a unified framework that jointly encodes the original and self-contained questions as

³We present an example of the error propagation in Section 5.2.

illustrated in Figure 2c. Our framework consists of two stages: (1) generating self-contained questions using a QR model (§3.1) and (2) training a QA model with the original and self-contained questions via consistency regularization (§3.2).

3.1 Question Rewriting

Similar to the pipeline approach, we utilize a QR model to obtain self-contained questions. We use the obtained questions for explicit guidance in the next stage. As shown in Equation 2, the QR task is to generate a self-contained question given an original question and a conversation history. Following Lin et al. (2020), we adopt a T5-based sequence generator (Raffel et al., 2020) as our QR model, which achieves comparable performance with that of humans in QR.⁴ For training and evaluating the QR model, we use the CANARD dataset following previous works on QR (Lin et al., 2020; Vakulenko et al., 2020). During inference, we utilize the top-k random sampling decoding based on beam search with the adjustment of the softmax temperature (Fan et al., 2018; Xie et al., 2019).

3.2 Consistency Regularization

Our goal is to enhance the QA model’s ability to understand conversational context. Accordingly, we use consistency regularization (Laine and Aila, 2016; Xie et al., 2019), which enforces a model to make consistent predictions in response to transformations to the inputs. We encourage the model’s predicted answers from the original questions to be similar to those from the self-contained questions (§3.1). Our consistency loss is defined as:

$$\mathcal{L}_t^{\text{cons}} = \text{KL}(\text{P}_\theta^{\text{read}}(a_t|d, q_t, \mathcal{H}_t) || \text{P}_{\bar{\theta}}^{\text{read}}(a_t|d, \tilde{q}_t, \tilde{\mathcal{H}}_t)) \quad (3)$$

where $\text{KL}(\cdot)$ represents the Kullback–Leibler divergence function between two probability distributions. θ is the model’s parameters, and $\bar{\theta}$ depicts a fixed copy of θ .

With the consistency loss, QA models are regularized to make consistent predictions, regardless of whether the given question is self-contained or not. In order to output an answer distribution that is closer to $\text{P}_{\bar{\theta}}^{\text{read}}(a_t|d, \tilde{q}_t, \tilde{\mathcal{H}}_t)$, QA models should treat original questions as if they were rewritten into self-contained questions by referring to the

⁴On CANARD, our QR model achieved comparable performance with the human performance in preliminary experiments.

conversation history. Through this process, our consistency regularization method serves as explicit guidance that encourages QA models to resolve the conversational dependency. In our framework, $\text{P}_\theta^{\text{read}}(a_t|\cdot)$ is the answer span distribution over all evidence document tokens. In contrast to Asai and Hajishirzi (2020), by using all probability values in the answer distributions, the signals of self-contained questions can be effectively propagated to the QA model. In addition to using all probability values, we also sharpened the target distribution $\text{P}_{\bar{\theta}}^{\text{read}}(a_t|d, \tilde{q}_t, \tilde{\mathcal{H}}_t)$ by adjusting the temperature (Xie et al., 2019) to strengthen the QA model’s training signal.

Finally, we calculate the final loss as:

$$\mathcal{L} = \mathcal{L}^{\text{orig}} + \lambda_1 \mathcal{L}^{\text{self}} + \lambda_2 \mathcal{L}^{\text{cons}} \quad (4)$$

where λ_1 and λ_2 are hyperparameters. $\mathcal{L}^{\text{orig}}$ and $\mathcal{L}^{\text{self}}$ are calculated by the negative log-likelihood between the predicted answers and gold standards given the original and self-contained questions, respectively.

Comparison with previous works Consistency training has mainly been studied as a method for regularizing model predictions to be invariant to small noises that are injected into the input samples (Sajjadi et al., 2016; Laine and Aila, 2016; Miyato et al., 2016; Xie et al., 2019). The intuition behind consistency training is to push noisy inputs closer towards their original versions. Therefore, only the original parameters (i.e., θ) are updated, while the copied model parameters (i.e., $\bar{\theta}$) are fixed.

In contrast to the original concept of consistency training, our goal is to go in the opposite direction and update the original parameters. Thus, we fix the parameters $\bar{\theta}$ with self-contained questions, and solely update θ for each training step as shown in Equation 3.

4 Experiments

In this section, we describe our experimental setup and compare our framework to baseline approaches (i.e., the end-to-end and pipeline approaches).

4.1 Datasets

QuAC QuAC (Choi et al., 2018) comprises 100k QA pairs in information-seeking dialogues, where a student asks questions based on a topic with background information provided, and a teacher provides the answers in the form of text spans in

Wikipedia documents. Since the test set is only available in the QuAC challenge, we evaluate models on the development set.⁵ For validation, we use a subset of the original training set of QuAC, which consists of questions that correspond to the self-contained questions in CANARD’s development set. The remaining data is used for training.

CANARD CANARD (Elgohary et al., 2019) consists of 31K, 3K, and 5K QA pairs for training, development, and test sets, respectively. The questions in CANARD are generated by rewriting a subset of the original questions in QuAC. We use the training and development sets for training and validating QR models, and the test set for evaluating QA models.

CoQA CoQA (Reddy et al., 2019) consists of 127K QA pairs and evidence documents in seven domains. In terms of the question distribution, CoQA significantly differs from QuAC (see §5.3). We use CoQA to test the transferability of EXCORD, where a QR model trained on CANARD generates the self-contained questions in a zero-shot manner. Subsequently, we train a QA model by using the original and synthetic questions. Similar to QuAC, the test set of CoQA is solely available in the CoQA challenge.⁶ Therefore, we randomly sample 5% of the QA dialogues in the training set and adopt them as our development set.

4.2 Metrics

Following Choi et al. (2018), we use the F1, HEQ-Q, and HEQ-D for QuAC and CANARD. HEQ-Q measures whether a model finds more accurate answers than humans (or the same answers) in a given question. HEQ-D measures the same thing, but in a given dialog instead of a question. For CoQA, we report the F1 scores for each domain (children’s story, literature from Project Gutenberg, middle and high school English exams, news articles from CNN, Wikipedia) and the overall F1 score, as suggested by Reddy et al. (2019).

4.3 QA models

Note that the baseline approaches and our framework do not limit the structure of QA models. For a fair comparison of the baseline approaches and EXCORD, we test the same QA models in all approaches. The selected QA models are commonly used and have been proven to be effective in CQA.

⁵<https://quac.ai/>

⁶<https://stanfordnlp.github.io/coqa/>

BERT BERT (Devlin et al., 2019) is a contextualized word representation model that is pretrained on large corpora. BERT also works well on CQA datasets, although it is not designed for CQA. It receives the evidence document, current question, and conversation history of the previous turn as input.

BERT+HAE BERT+HAE is a BERT-based QA model with a CQA-specific module. Following Qu et al. (2019a), we add the history answer embedding (HAE) to BERT’s word embeddings. HAE encodes the information of the answer spans from the previous questions.

RoBERTa RoBERTa (Liu et al., 2019) improves BERT by using pretraining techniques to obtain the robustly optimized weights on larger corpora. In our experiments, we found that RoBERTa performs well in CQA, achieving comparable performance with the previous SOTA model, HAM (Qu et al., 2019b), on QuAC. Thus, we adopt RoBERTa as our main baseline model owing to its simplicity and effectiveness. It receives the same input as BERT, otherwise specified.

4.4 Implementation Details

The CANARD training set provides 31,527 self-contained questions from the original QuAC questions. Therefore, we can obtain 31,527 pairs of original and self-contained questions without question rewriting. For the rest of the original questions, we automatically generate self-contained questions by using our QR model. Finally, we obtain 83,568 question pairs and use them in our consistency training. We denote the original questions, self-contained questions generated by humans, and self-contained questions generated by a QR model as Q , \tilde{Q}^{human} , and \tilde{Q}^{syn} , respectively. Additional implementation details are described in Appendix B

4.5 Results

Table 1 presents the performance comparison of the baseline approaches to our framework on QuAC and CANARD. Compared to the end-to-end approach, EXCORD consistently improves the performance of QA models on both datasets. Also, these improvements are significant: EXCORD improves the performance of the RoBERTa by absolutely 1.2 and 2.3 F1 scores and BERT by 1.2 and 5.2 F1 scores on QuAC and CANARD, respectively. From these results, we conclude that the consistency training with original and self-contained

QA Model	Approach	QuAC			CANARD		
		F1	HEQ-Q	HEQ-D	F1	HEQ-Q	HEQ-D
BERT	End-to-end	61.5	57.1	5.0	57.4	52.9	3.2
	Pipeline	61.2 (-0.3)	56.8 (-0.3)	5.0 (-)	62.2 (+4.8)	57.8 (+4.9)	6.0 (+2.8)
	Ours	62.7 (+1.2)	58.4 (+1.3)	6.0 (+1.0)	62.6 (+5.2)	58.2 (+5.3)	6.4 (+3.2)
BERT+HAE	End-to-end	62.0	57.3	5.5	58.2	53.5	5.5
	Pipeline	61.1 (-0.9)	56.3 (-1.0)	5.0 (-0.5)	62.4 (+4.2)	57.8 (+4.3)	6.0 (+0.5)
	Ours	63.2 (+1.2)	58.9 (+1.6)	5.7 (+0.2)	63.1 (+4.9)	58.4 (+4.9)	5.7 (+0.2)
RoBERTa	End-to-end	66.5	62.4	7.2	65.8	62.2	7.1
	Pipeline	65.2 (-1.3)	60.9 (-1.5)	7.1 (-0.1)	66.9 (+1.1)	63.2 (+1.0)	7.3 (+0.2)
	Ours	67.7 (+1.2)	64.0 (+1.6)	9.3 (+2.1)	68.1 (+2.3)	64.2 (+2.0)	8.4 (+1.3)

Table 1: Comparison in performance of the baseline approaches and our framework on QuAC and CANARD. The best scores are highlighted in bold.

questions enhances ability of QA models to understand the conversational context.

On QuAC, the pipeline approach underperforms the end-to-end approach in all baseline models. This indicates that training a QA model solely with self-contained questions is ineffective when human rewrites are not given at the inference phase. On the other hand, EXCORD improves QA models by using both types of questions. As presented in Table 1, our framework significantly outperforms the baseline approaches on QuAC.

On CANARD, the pipeline approach is significantly more effective than the end-to-end approach. Since QA models are trained with self-contained questions in the pipeline approach, they perform well on CANARD questions. Nevertheless, EXCORD still outperforms the pipeline approach in most cases. Compared to the pipeline approach, our framework improves the performance of RoBERTa by absolutely 1.2 F1 score.

5 Analysis and Discussion

We elaborate on analyses regarding component ablation and transferability. We also describe a case study carried out to highlight such differences between our and baseline approaches.

5.1 Ablation Study

In this section, we comprehensively explore the factors contributing to this improvement in detail: (1) using self-contained questions that are rewritten by humans (\tilde{Q}^{human}) as additional data, (2) using self-contained questions that are synthetically generated by the QR model (\tilde{Q}^{syn}), and (3) training a QA model with our consistency framework. In Table 2, we present the performance gaps when each component is removed from our framework. We use RoBERTa on QuAC in this experiment.

Method	QuAC F1	CANARD F1
EXCORD	67.7	68.1
- \tilde{Q}^{syn}	67.5	67.7
- \tilde{Q}^{human}	67.3	67.2
Question Augment. (w/o. EXCORD)	65.9	66.2
- \tilde{Q}^{syn}	66.1	66.5
- \tilde{Q}^{human}	65.3	66.0
- $\tilde{Q}^{\text{syn}}, \tilde{Q}^{\text{human}}$ (End-to-end)	66.5	65.8

Table 2: Effect of self-contained questions and our consistency framework. We use RoBERTa in this experiment.

We first explore the effects of \tilde{Q}^{human} and \tilde{Q}^{syn} . As shown in Table 2, excluding \tilde{Q}^{human} degrades the performance of RoBERTa in our framework. Although automatically generated, \tilde{Q}^{syn} contributes to the performance improvement. Therefore, both types of self-contained questions are useful in our framework.

To investigate the effect of our framework, we simply augment \tilde{Q}^{human} and \tilde{Q}^{syn} to Q^{orig} , which is called Question Augment (question data augmentation). We find that Question Augment slightly improves the performance of RoBERTa on CANARD, whereas it degrades the performance on QuAC. This shows that simply augmenting the questions is ineffective and does not guarantee improvement. On the other hand, our consistency training approach significantly improves performance, showing that EXCORD is a more optimal way to utilizing self-contained questions.

5.2 Case Study

We analyze several cases that the baseline approaches answered incorrectly, but our framework answered correctly. We also explore how our framework improves the reasoning ability of QA models, compared to the baseline approaches. These cases

Error case # 1**Title :** Montgomery Clift **Section Title :** Film career**Document d :**

...

His second movie was The Search. Clift was unhappy with the quality of the script, and edited it himself. The movie was awarded a screenwriting Academy Award for the credited writers.

...

 q_1 : When did Clift start his film career? a_1 : His first movie role was opposite John Wayne in Red River, which was shot in 1946 and released in 1948.

Current Question q_2 : Was the film a success?**Human Rewrite r_2 :** Was Montgomery Clift’s film Red River a success?**Golden Answer :** CANNOTANSWER**Prediction of End-to-End :** The movie was awarded a screenwriting Academy Award for the credited writers.**Prediction of Ours :** CANNOTANSWER

Error case # 2**Title :** Train (band) **Section Title :** 2003-2004: My Private Nation

...

 q_5 : Did my private nation do any other features? a_5 : CANNOTANSWER

Current Question q_6 : Did my private nation have any good singles?**Generated Question \hat{q}_6 :** Did Train’s private nation have any good singles?**Golden Answer :** “Get to Me” (written by Rob Hotchkiss and Pat Monahan) reached number nine on the Billboard Adult Top 40.**Prediction of Pipeline :** CANNOTANSWER**Prediction of Ours :** “Get to Me” (written by Rob Hotchkiss and Pat Monahan) reached number nine on the Billboard Adult Top 40.

Table 3: Error analysis for predictions of RoBERTa that are trained with the baseline approaches and EXCORD. In the first case, the QA model trained with the end-to-end approach fails to resolve the conversational dependency. The QR model in the second case misunderstands the “my,” and generates an unnatural question, triggering an incorrect prediction.

are obtained from the development set of QuAC.

The first case in Table 3 shows the predictions of the two RoBERTa models trained in the end-to-end approach and our framework, respectively. Note that “the film” in the current question does not refer to “The Search” (red box) in the document d , but “Red River” (blue box) in a_1 . When trained in the end-to-end approach, the model failed to comprehend the conversational context and misunderstood what “the film” refers to, resulting in an incorrect prediction. On the other hand, when trained in EXCORD, the model predicted the correct answer because it enhances the ability to resolve conversational dependency.

In the second case, we compare the pipeline approach to EXCORD. In this case, the QR model misunderstood “my” in the current question as a pronoun and replaced it with the band’s name, “Train’s.” Consequently, the QA model received the erroneous self-contained question, resulting in an incorrect prediction. On the other hand, the QA model trained in our framework predicted the

correct answer based on the original question q_6 .

5.3 Transferability

We train a QR model to rewrite QuAC questions into CANARD questions. Then, self-contained questions can be generated for the samples that do not have human rewrites. This results in the improvement of QA models’ performance on QuAC and CANARD (§4.5). However, it is questionable whether the QR model can successfully rewrite questions when the original questions significantly differ from those in QuAC. To answer this, we test our framework on another CQA dataset, CoQA. We first analyze how the question distributions of QuAC and CoQA differ. We found that question types in QuAC and CoQA are significantly different, such that QR models could suffer from the gap of question distributions between two datasets. (See details in Appendix A).

To test the transferability of EXCORD, we compare the end-to-end approach to our framework on the CoQA dataset. Using a QR model trained on

QA model	CoQA (F1)					
	Overall	Child.	Liter.	M&H	News	Wiki.
BERT						
End-to-End	78.3	77.9	73.9	76.4	80.6	82.7
Pipeline	76.1	75.7	73.2	74.1	78.0	79.6
Ours	78.8	78.2	75.8	75.5	81.3	83.2
RoBERTa						
End-to-End	82.8	82.5	80.2	80.1	84.3	87.0
Pipeline	81.1	81.9	78.2	78.3	82.4	85.2
Ours	83.4	84.4	81.2	79.8	84.6	87.0

Table 4: Effect of our framework on the CoQA dataset that do not have human rewrites. We exclude BERT+HAE for simplification in this experiment.

CANARD, we generate the self-contained questions for CoQA and train QA models with our framework. As presented in Table 4, our framework performs well on CoQA. The improvement in BERT is 0.5 based on the overall F1, and the performance of RoBERTa is also improved by an overall F1 of 0.6. Improvements are also consistent in most of the documents’ domains. Therefore, we conclude that our framework can be simply extended to other datasets and improve QA performance even when question distributions are significantly different. We plan to improve the transferability of our framework by fine-tuning QR models on target datasets in future work.

6 Related Work

Conversational Question Answering Recently, several works introduced CQA datasets such as QUAC (Choi et al., 2018) and COQA (Reddy et al., 2019). We classified proposed methods to solve the datasets into two approaches: (1) end-to-end and (2) pipeline. Most works based on the end-to-end approach focused on developing a model structure (Zhu et al., 2018; Ohsugi et al., 2019; Qu et al., 2019a,b) or training strategy such as multi-task with rationale tagging (Ju et al., 2019) that are specialized in the CQA task or datasets. Several works demonstrated the effectiveness of the flow mechanism in CQA (Huang et al., 2018; Chen et al., 2019; Yeh and Chen, 2019).

With the advent of a dataset consisting of self-contained questions rewritten by human annotators (Elgohary et al., 2019), the pipeline approach has drawn attention as a promising method for CQA in recent days (Vakulenko et al., 2020). The approach is particularly useful for the open-domain CQA or passage-retrieval (PR) tasks (Dalton et al., 2019; Ren et al., 2020; Anantha et al., 2020; Qu et al., 2020) since self-contained questions can be

fed into existing non-conversational search engines such as BM25. Note that our framework can be used jointly with the pipeline approach in the open-domain setting because our framework can improve QA models’ ability to find the answers from the retrieved documents. We will test our framework in the open-domain setting in future work.

Question Rewriting QR has been studied for augmenting training data (Buck et al., 2018; Sun et al., 2018; Zhu et al., 2019; Liu et al., 2020) or clarifying ambiguous questions (Min et al., 2020). In CQA, QR can be viewed as a task of simplifying difficult questions that include anaphora and ellipsis in a conversation. Elgohary et al. (2019) first proposed the question rewriting task as a sub-task of CQA and the CANARD dataset for the task, which consists of pairs of original and self-contained questions that are generated by human annotators. Vakulenko et al. (2020) used a coreference-based model (Lee et al., 2018) and GPT-2 (Radford et al., 2019) as QR models and tested the models in the QR and PR tasks. Lin et al. (2020) conducted the QR task using T5 (Raffel et al., 2020) and achieved on performance comparable to humans on CANARD. Following Lin et al. (2020), we use T5 in our experiments to generate high-quality questions for enhancing QA models.

Consistency Training Consistency regularization (Laine and Aila, 2016; Sajjadi et al., 2016) has been mainly explored in the context of semi-supervised learning (SSL) (Chapelle et al., 2009; Oliver et al., 2018), which has been adopted in the textual domain as well (Miyato et al., 2016; Clark et al., 2018; Xie et al., 2020). However, the consistency training framework is also applicable when only the labeled samples are available (Miyato et al., 2018; Jiang et al., 2019; Asai and Hajishirzi, 2020). The consistency regularization requires adding noise to the sample, which can be either discrete (Xie et al., 2020; Asai and Hajishirzi, 2020) or continuous (Miyato et al., 2016; Jiang et al., 2019). Existing works regularize the predictions of the perturbed samples to be equivalent to be that of the originals’. On the other hand, our method encourages the models’ predictions for the original answers to be similar to those from the rewritten questions, i.e., synthetic ones.

7 Conclusion

We propose a consistency training framework for conversational question answering, which enhances QA models' abilities to understand conversational context. Our framework leverages both the original and self-contained questions for explicit guidance on how to resolve conversational dependency. In our experiments, we demonstrate that our framework significantly improves the QA model's performance on QuAC and CANARD, compared to the existing approaches. In addition, we verified that our framework can be extended to CoQA. In future work, the transferability of our framework can be further improved by fine-tuning the QR model on target datasets. Furthermore, future work would include applying our framework to the open-domain setting.

Acknowledgements

We thank Sean S. Yi, Miyoung Ko, and Jinhyuk Lee for providing valuable comments and feedback. This research was supported by the MSIT (Ministry of Science and ICT), Korea, under the ICT Creative Consilience program (IITP-2021-2020-0-01819) supervised by the IITP (Institute for Information communications Technology Planning Evaluation). This research was also supported by National Research Foundation of Korea (NRF-2020R1A2C3010638).

References

- Raviteja Anantha, Svitlana Vakulenko, Zhucheng Tu, Shayne Longpre, Stephen Pulman, and Srinivas Chappidi. 2020. Open-domain question answering goes conversational via question rewriting. *arXiv preprint arXiv:2010.04898*.
- Akari Asai and Hannaneh Hajishirzi. 2020. Logic-guided data augmentation and regularization for consistent question answering. *arXiv preprint arXiv:2004.10157*.
- Christian Buck, Jannis Bulian, Massimiliano Ciaramita, Wojciech Gajewski, Andrea Gesmundo, Neil Houlsby, and Wei Wang. 2018. Ask the right questions: Active question reformulation with reinforcement learning. In *International Conference on Learning Representations*.
- Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. 2009. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2019. Graphflow: Exploiting conversation flow with graph neural networks for conversational machine comprehension. *arXiv preprint arXiv:1908.00059*.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. *arXiv preprint arXiv:1808.07036*.
- Kevin Clark, Minh-Thang Luong, Christopher D Manning, and Quoc V Le. 2018. Semi-supervised sequence modeling with cross-view training. *arXiv preprint arXiv:1809.08370*.
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2019. Cast 2019: The conversational assistance track overview. In *Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC*, pages 13–15.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Ahmed Elgohary, Denis Peskov, and Jordan Boyd-Graber. 2019. Can you unpack that? learning to rewrite questions-in-context. *Can You Unpack That? Learning to Rewrite Questions-in-Context*.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.
- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018. Flowqa: Grasping flow in history for conversational machine comprehension. *arXiv preprint arXiv:1810.06683*.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019. Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. *arXiv preprint arXiv:1911.03437*.
- Ying Ju, Fubang Zhao, Shijie Chen, Bowen Zheng, Xuefeng Yang, and Yunfeng Liu. 2019. Technical report on conversational question answering. *arXiv preprint arXiv:1909.10772*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Samuli Laine and Timo Aila. 2016. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*.

- Kenton Lee, Luheng He, and Luke Zettlemoyer. 2018. Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 687–692.
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2020. Conversational question reformulation via sequence-to-sequence architectures and pretrained language models. *arXiv preprint arXiv:2004.01909*.
- Dayiheng Liu, Yeyun Gong, Jie Fu, Yu Yan, Jiusheng Chen, Jiancheng Lv, Nan Duan, and Ming Zhou. 2020. Tell me how to ask again: Question data augmentation with controllable rewriting in continuous space. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5798–5810.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. Ambigqa: Answering ambiguous open-domain questions. *arXiv preprint arXiv:2004.10645*.
- Takeru Miyato, Andrew M Dai, and Ian Goodfellow. 2016. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993.
- Yasuhito Ohsugi, Itsumi Saito, Kyosuke Nishida, Hisako Asano, and Junji Tomita. 2019. A simple but effective method to incorporate multi-turn context with bert for conversational machine comprehension. *arXiv preprint arXiv:1905.12848*.
- Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. 2018. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems*, 31:3235–3246.
- Chen Qu, Liu Yang, Cen Chen, Minghui Qiu, W Bruce Croft, and Mohit Iyyer. 2020. Open-retrieval conversational question answering. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 539–548.
- Chen Qu, Liu Yang, Minghui Qiu, W Bruce Croft, Yongfeng Zhang, and Mohit Iyyer. 2019a. Bert with history answer embedding for conversational question answering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1133–1136.
- Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W Bruce Croft, and Mohit Iyyer. 2019b. Attentive history selection for conversational question answering. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1391–1400.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- 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(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Pengjie Ren, Zhumin Chen, Zhaochun Ren, Evangelos Kanoulas, Christof Monz, and Maarten de Rijke. 2020. Conversations with search engines. *arXiv preprint arXiv:2004.14162*.
- Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2018. Improving machine reading comprehension with general reading strategies. *arXiv preprint arXiv:1810.13441*.
- Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2020. Question rewriting for conversational question answering. *arXiv preprint arXiv:2004.14652*.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. 2020. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33.

Yi-Ting Yeh and Yun-Nung Chen. 2019. Flowdelta: Modeling flow information gain in reasoning for conversational machine comprehension. *arXiv preprint arXiv:1908.05117*.

Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.

Haichao Zhu, Li Dong, Furu Wei, Wenhui Wang, Bing Qin, and Ting Liu. 2019. Learning to ask unanswerable questions for machine reading comprehension. *arXiv preprint arXiv:1906.06045*.

QuAC	CoQA
Title : Scott Walker (politician) Section Title : Education q_1 : What kind of education did Scott Walker have? a_1 : CANNOTANSWER q_2 : Are there any other interesting aspects about this article? a_2 : signed a law to fund evaluation of the reading skills of kindergartners as part of an initiative to ensure that students are reading at or above grade level	q_1 : Is the US dollar on a decimal system? a_1 : U.S. dollar is based upon a decimal system of values. I q_2 : What country’s dollar is not? a_2 : Unlike the Spanish milled dollar the U.S. dollar is based upon a decimal system of values. q_3 : What is a mill? a_3 : n addition to the dollar the coinage act officially established monetary units of mill or one-thousandth of a dollar
Current Question q_3 : What other programs did he sign? Self-contained Question \tilde{q}_3 :What other programs did Scott Walker sign other than a law to fund evaluation	Current Question q_4 : And a cent? Self-contained Question \tilde{q}_4 : What is a cent? -

Table 5: Comparison of questions in QuAC and CoQA. In the left side, we can observe several question types that are frequently used in QuAC: unanswerable question (q_1) and “Anything else?” question (q_2). The current question q_3 refers to the previous answer (green box) and the background information (blue box). On the other hand, in the right side, the current question q_4 omits the question word that are used in the previous question (yellow box).

Question Type	QuAC	CoQA
Non-factoid	54 %	38 %
Anything else?	11 %	3 [†] %
Unanswerable	20 %	1 %

Table 6: Statistics of question types for QuAC and CoQA. All values can be found in the QuAC and CoQA papers (Choi et al., 2018; Reddy et al., 2019) except for those with the dagger [†]. We randomly sampled 106 questions and manually labeled for obtaining the number with the dagger [†].

A Comparison of Questions in QuAC and CoQA

Before testing the transferability of EXCORD (§5.3), we compare the question distribution of QuAC to that of CoQA. The types of questions are significantly different due to the difference in task setups. When questions were generated in QuAC, evidence documents were solely provided to answerers, but not to questioners. This setup prevented questioners from referring to the evidence documents, which encouraged the questioners to ask natural and information-seeking questions. By contrast, when creating CoQA, questioners and answerers shared the same evidence documents.

Examples of QuAC and CoQA are presented in Table 5 and the categorization of question types in Table 6. The results are as follows: (1) QuAC has more non-factoid questions. Approximately half of QuAC questions are non-factoid, whereas more than 60% of questions in CoQA can be answered with either entities or noun phrases. (2) “Anything else?” questions are more frequently observed in QuAC. When questioners cannot find what to ask, they use “Anything else?” questions to seek new topics and continue the conversation. In CoQA,

questioners rarely used the “Anything else?” question (2.8%) since they did not need to seek new topics. This type of question is observed in Table 5 (q_2 in the left side). (3) CoQA has few unanswerable questions. Since questioners and answerers share the evidence documents when creating CoQA, only 1.3% of unanswerable questions are asked. However, approximately 20% of questions in QuAC are unanswerable.

B Hyperparameters

Our implementation is based on PyTorch.⁷ We implemented BERT using the Transformers library.⁸ We implemented the T5-based QR model using the Transformers library and adopted the same QR model in the pipeline approach and EXCORD. We use a single 24GB GPU (RTX TITAN) for the experiments.

We measured the F1 scores on the development set for each 4k training step, and adopted the best-performing models. We trained QA models based on the AdamW optimizer with a learning rate of $3e-5$. We use the maximum input sequence length as 512 and the maximum answer length as 30. We set the maximum query length to 128 for all approaches since self-contained questions are usually longer than original questions. We use a batch size 12 for BERT and RoBERTa in all baseline approaches. For EXCORD, we set the coefficient λ_1 for QA loss for rewritten questions to 0.5. Also we search the coefficient λ_2 for consistency loss within the range of [0.7, 0.5] and the softmax temperature within the range of [1.0, 0.9] (Xie et al., 2019).

⁷<https://pytorch.org/>

⁸<https://github.com/huggingface>