

Reinforced Dynamic Reasoning for Conversational Question Generation

Boyuan Pan^{1*}, Hao Li¹, Ziyu Yao², Deng Cai^{1,3}, Huan Sun²

¹State Key Lab of CAD&CG, Zhejiang University

²The Ohio State University

³Alibaba-Zhejiang University Joint Institute of Frontier Technologies

{panby, haolics, dcai}@zju.edu.cn

{yao.470, sun.397}@osu.edu

Abstract

This paper investigates a new task named *Conversational Question Generation (CQG)* which is to generate a question based on a passage and a conversation history (i.e., previous turns of question-answer pairs). CQG is a crucial task for developing intelligent agents that can drive question-answering style conversations or test user understanding of a given passage. Towards that end, we propose a new approach named Reinforced Dynamic Reasoning (ReDR) network, which is based on the general encoder-decoder framework but incorporates a reasoning procedure in a *dynamic* manner to better understand what has been asked and what to ask next about the passage. To encourage producing meaningful questions, we leverage a popular question answering (QA) model to provide feedback and fine-tune the question generator using a reinforcement learning mechanism. Empirical results on the recently released CoQA dataset demonstrate the effectiveness of our method in comparison with various baselines and model variants. Moreover, to show the applicability of our method, we also apply it to create multi-turn question-answering conversations for passages in SQuAD.

1 Introduction

In this work, we study a novel task of *conversational question generation (CQG)* which is given a passage and a conversation history (i.e., previous turns of question-answer pairs), to generate the next question.

CQG is an important task in its own right for measuring the ability of machines to lead a question-answering style conversation. It can serve as an essential component of intelligent social bots or tutoring systems, asking meaningful

Shelly is in second grade. She is a new student at her school. Shelly's family has lived in many different places. Shelly was born in Florida. Her family moved to Tennessee when she was two years old. When she was four years old, they moved to Texas. They moved from there to Arizona, where they now live.

Q1: What grade is Shelly in ?

A1: second

R1: Shelly is in second grade.

Q2: Was she a new student ?

A2: Yes

R2: She is a new student at her school.

Q3: Where did she move at 2 years old ?

A2: Tennessee

R3: Her family moved to Tennessee when she was two years old.

Figure 1: An example from the CoQA dataset. Each turn contains a question (Q) and an answer (A). The dataset also provides a rationale (R) (i.e., a text span from the passage) to support each answer.

and coherent questions to engage users or test student understanding about a certain topic. On the other hand, as shown in Figure 1, large-scale high-quality conversational question answering (CQA) datasets such as CoQA (Reddy et al., 2018) and QuAC (Choi et al., 2018) can help train models to answer sequential questions. However, manually creating such datasets is quite costly, e.g., CoQA spent 3.6 USD per passage on crowdsourcing for conversation collection, and automatic CQG can potentially help reduce the cost, especially when there are a large set of passages available.

In recent years, automatic question generation (QG), which aims to generate natural questions based on a certain type of data sources including structured knowledge bases (Serban et al., 2016b; Guo et al., 2018) and unstructured texts (Rus et al.,

*Work done while visiting the Ohio State University.

2010; Heilman and Smith, 2010; Du et al., 2017; Du and Cardie, 2018), has been widely studied. However, previous works mainly focus on generating standalone and independent questions based on a given passage. To the best of our knowledge, we are the first to explore CQG, i.e., generating the next question based on a passage and a *conversation history*.

Comparing with previous QG tasks, CQG needs to take into account not only the given passage, but also the conversation history, and is potentially more challenging as it requires a deep understanding of what has been asked so far and what information should be asked for the next round, in order to make a coherent conversation.

In this paper, we present a novel framework named *Reinforced Dynamic Reasoning* (ReDR) network. Inspired by the recent success of reading comprehension models (Xiong et al., 2017; Seo et al., 2017), ReDR adapts their reasoning procedure (which encodes the knowledge of the passage and the conversation history based on a coattention mechanism) and moreover *dynamically* updates the encoding representation based on a soft decision maker to generate a coherent question. In addition, to encourage ReDR to generate meaningful and interesting questions, ideally, one may employ humans to provide feedback, but as widely acknowledged, involving humans in the loop for training models can be very costly. Therefore, in this paper, we leverage a popular and effective reading comprehension (or QA) model (Chen et al., 2017) to predict the answer to a generated question and use its answer quality (which can be seen as a proxy for real human feedback) as rewards to fine-tune our model based on a reinforcement learning mechanism (Williams, 1992).

Our contributions are summarized as follows:

- We introduce a new task of *Conversational Question Generation* (CQG), which is crucial for developing intelligent agents to drive question-answering style conversations and can potentially provide valuable datasets for future relevant research.
- We propose a new and effective framework for CQG, which is equipped with a dynamic reasoning component to generate a conversational question and is further fine-tuned via a reinforcement learning mechanism.
- We show the effectiveness of our method us-

ing the recent CoQA dataset. Moreover, we show its wide applicability by using it to create multi-turn QA conversations for passages in SQuAD (Rajpurkar et al., 2016).

2 Task Definition

Formally, we define the task of *Conversational Question Generation* (CQG) as: Given a passage X and the previous turns of question-answer pairs $\{(q_1, a_1), (q_2, a_2), \dots, (q_{k-1}, a_{k-1})\}$ about X , CQG aims to generate the next question q_k that is related to the given passage and coherent with the previous questions and answers, i.e.,

$$q_k = \arg \max_{q_k} P(q_k | X, q_{<k}, a_{<k}) \quad (1)$$

where $P(q_k | X, q_{<k}, a_{<k})$ is a conditional probability of generating the question q_k .

3 Methodology

We show our proposed framework named *Reinforced Dynamic Reasoning* (ReDR) network in Figure 2. Since a full passage is usually too long and makes it hard to focus on the most relevant information for generating the next question, our method first selects a text span from the passage as the rationale at each conversation turn, and then dynamically models the reasoning procedure for encoding the conversation history and the selected rationale, before finally decoding the next question.

3.1 Rationale Selection

We simply set each sentence in the passage as the corresponding rationale for each turn of the conversation. When experimenting with CoQA, we use the rationale span provided in the dataset. Besides for simplicity and efficiency, another reason that we adopt this rule-based method is that previous research demonstrated that the transition of the dialog attention is smooth (Reddy et al., 2018; Choi et al., 2018), meaning that earlier questions in a conversation are usually answerable by the preceding part of the passage while later questions tend to focus on the ending part of the passage. The selected rationale is then leveraged by subsequent modules for question generation.

3.2 Encoding & Reasoning

At each turn k , we denote the conversation history as a sequence of m tokens, i.e., $c =$

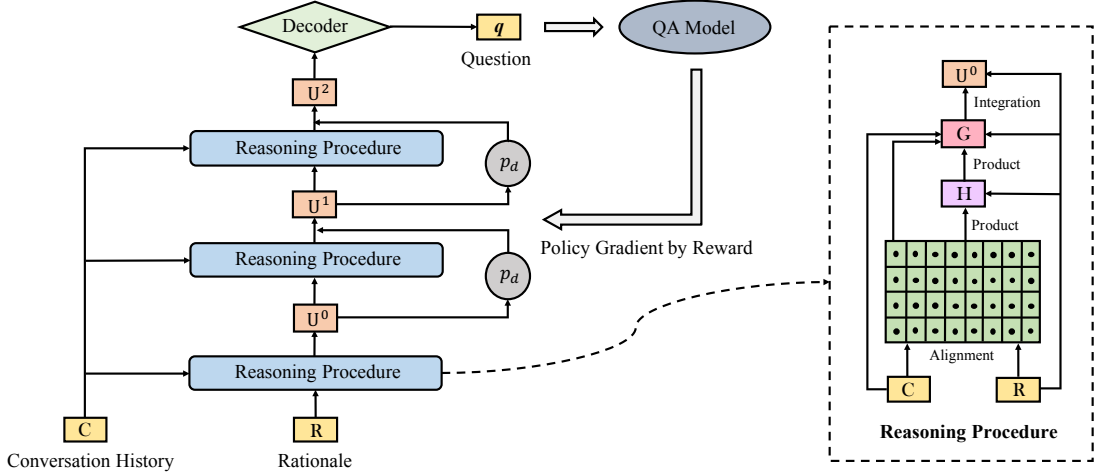


Figure 2: Overview of our Reinforced Dynamic Reasoning (ReDR) network. The reasoning mechanism iteratively reads the conversation history and at each iteration, its output is dynamically combined with the previous encoding representation through a soft decision maker (p_d) as the new encoding representation, which is fed into the next iteration. The model is finally fine-tuned by the reward defined by the quality of the answer predicted from a QA model.

$\{c_1, c_2, \dots, c_m\}$, which concatenates the previous questions and answers $\langle q_1, a_1, \dots, q_{k-1}, a_{k-1} \rangle$, and represent the rationale as a sequence of n tokens, i.e., $r = \{r_1, r_2, \dots, r_n\}$. As mentioned earlier, different from previous question generation tasks, we have two knowledge sources (i.e., the conversation history and the rationale) as the inputs. A good encoding of them is crucial for task performance and might involve a reasoning procedure across previous question-answer pairs and the selected rationale for determining the next question. We feed them respectively into a bi-directional LSTM and obtain their contextual representations $\mathbf{C} \in R^{d \times m}$ and $\mathbf{R} \in R^{d \times n}$. Inspired by the coattention reasoning mechanism in previous reading comprehension works (Xiong et al., 2017; Seo et al., 2017; Pan et al., 2017), we compute an alignment matrix of \mathbf{C} and \mathbf{R} to link and fuse the information flow: $\mathbf{S} = \mathbf{R}^\top \mathbf{C} \in R^{n \times m}$. We normalize this alignment matrix column-wise (i.e., $\text{softmax}(\mathbf{S})$) to obtain the relevance degree of each token in the conversation history to the whole rationale. The new representation of the conversation history w.r.t. the rationale is obtained via:

$$\mathbf{H} = \mathbf{R} \cdot \text{softmax}(\mathbf{S}) \in R^{d \times m} \quad (2)$$

Similarly, we compute the attention over the conversation history for each word in the rationale via $\text{softmax}(\mathbf{S}^\top)$ and obtain the context-dependent representation of the rationale by $\mathbf{C} \cdot \text{softmax}(\mathbf{S}^\top)$. In addition, as in (Xiong

et al., 2017), we also consider the above new representation of the conversation history and map it to the space of rationale encodings via $\mathbf{H} \cdot \text{softmax}(\mathbf{S}^\top)$, and finally obtain the co-dependent representation of the rationale and the conversation history:

$$\mathbf{G} = [\mathbf{C}; \mathbf{H}] \cdot \text{softmax}(\mathbf{S}^\top) \in R^{2d \times n} \quad (3)$$

where $[\cdot]$ means concatenation across row dimension. To deeply capture the interaction between the rationale and the conversation history, we feed the co-dependent representation \mathbf{G} combined with the rationale \mathbf{R} into an *integration model* instantiated by a bi-directional LSTM:

$$\mathbf{u}_i^0 = \text{BiLSTM}(\mathbf{u}_{i-1}^0, \mathbf{u}_{i+1}^0, [\mathbf{G}_i; \mathbf{R}_i]) \in R^d \quad (4)$$

We define the reasoning process in our paper as Eqn. (2-4), and now obtain a matrix $\mathbf{U}^0 = [\mathbf{u}_1^0, \mathbf{u}_2^0, \dots, \mathbf{u}_n^0]$ as the encoding representation after one-layer reasoning procedure, which can be fed into the decoder subsequently.

3.3 Dynamic Reasoning

Oftentimes the conversation history is very informative and complicated, and one single layer of reasoning may be insufficient to comprehend the subtle relationship among the rationale, the conversation history, and the to-be-generated question. Therefore, we propose a dynamic reasoning procedure to iteratively update the encoding representation. We regard \mathbf{U}^0 as a new representation

of the rationale and input it to the next layer of reasoning together with \mathbf{C} :

$$\tilde{\mathbf{U}}^1 = F_{reason}(\mathbf{U}^0, \mathbf{C}) \quad (5)$$

where F_{reason} is the reasoning procedure (Eqn. 2-4), and $\tilde{\mathbf{U}}^1$ is the hidden states of the BiLSTM integration model at the next reasoning layer. To effectively learn what information in $\tilde{\mathbf{U}}^1$ and \mathbf{U}^0 is relevant to keep, we use a soft *decision maker* to determine their weights:

$$\begin{aligned} \mathbf{U}^1 &= \mathbf{p}_d \odot \mathbf{U}^0 + (\mathbf{e}_1 - \mathbf{p}_d) \odot \tilde{\mathbf{U}}^1 \\ \mathbf{p}_d &= \sigma(\mathbf{w}_u^\top \mathbf{U}^0 + \mathbf{w}_g^\top \mathbf{G} + \mathbf{w}_r^\top \mathbf{R} + \mathbf{b}) \end{aligned} \quad (6)$$

where \mathbf{e}_1 is an all-ones vector, and $\mathbf{w}_u, \mathbf{w}_g, \mathbf{w}_r, \mathbf{b}$ are trainable parameters. $\mathbf{p}_d \in R^n$ is the decision maker, used as a soft switch to choose between different levels of reasoning. \mathbf{U}^1 is the representation to be used for the next layer of reasoning. This iterative procedure halts when a maximum number of reasoning layers N is reached ($N \geq 1$). The final representation \mathbf{U}^N is fed into the decoder.

3.4 Decoding

The decoder generates a word by sampling from the probability $P_{gen}(y_t|y_{<t}, c, r)$ which can be computed via:

$$\begin{aligned} P_{gen}(y_t|y_{<t}, c, r) &= \text{MLP}(\mathbf{o}_t, \mathbf{v}_t) \\ \mathbf{o}_t &= \text{LSTM}(\mathbf{o}_{t-1}, \text{Emb}(y_{t-1}), \mathbf{v}_{t-1}) \end{aligned} \quad (7)$$

where MLP stands for a standard multilayer perceptron network, y_t is the t -th word in the generated question, \mathbf{o}_t is the hidden state of the decoder at time step t , and $\text{Emb}(\cdot)$ indicates the word embedding. \mathbf{v}_t is an attentive read of the encoding representation: $\mathbf{v}_t = \sum_{i=1}^n \alpha_{t,i} \mathbf{u}_i^N$, where the weight $\alpha_{t,i} \in (0, 1)$ is scored by another $\text{MLP}(\mathbf{o}_t, \mathbf{u}_i^N)$ network.

Observing that a question may share common words with the rationale that it is based on and inspired by the widely adopted copy mechanism (Gu et al., 2016; See et al., 2017), we also apply a pointer network for the generator to copy words from the rationale. Now the probability of generating target word y_t becomes:

$$P(y_t|y_{<t}, c, r) = \lambda P_{gen}(y_t) + (1 - \lambda) P_{pt}(y_t) \quad (8)$$

where $P_{gen}(y_t) = P_{gen}(y_t|y_{<t}, c, r)$ is defined earlier, $P_{pt}(y_t) = \sum_{i:r_i=y_t} \alpha_{t,i}$ is the probability of

copying word y_t from r (only if r contains y_t), and λ is the weight to balance the two:

$$\lambda = \sigma(\mathbf{w}_v^\top \mathbf{v}_t + \mathbf{w}_o^\top \mathbf{o}_t + \mathbf{w}_y^\top \text{Emb}(y_{t-1}) + \mathbf{b}_{pt}) \quad (9)$$

where $\mathbf{w}_v^\top, \mathbf{w}_o^\top, \mathbf{w}_y^\top$ and \mathbf{b}_{pt} are to be learnt. To optimize all parameters in ReDR, we adopt the maximum likelihood estimation (MLE) approach, i.e., maximizing the summed log likelihood of words in a target question.

3.5 Reinforcement Learning for Fine-tuning

As shown by recent datasets like CoQA and QuAC, human-created questions tend to be meaningful and interesting. For example, in Figure 1, given the second rationale R2 “She is a new student at her school”, humans tend not to ask “Where is she?”, and similarly given R3, they usually do not create the question “What happened?”. Although both are legitimate questions, they tend to be less interesting and meaningful compared with the human-created ones shown in Figure 1. The interestingness or meaningfulness of a question is subjective and hard to define, automatically measuring which is a difficult problem itself. Ideally, one can involve humans in the loop to judge the generated question and provide feedback, but it can be very costly, if not impossible.

Driven by such observations, we use the REINFORCE (Williams, 1992) algorithm and adopt one of the state-of-the-art reading comprehension models DrQA (Chen et al., 2017) as a substitute for humans to provide feedback to the question generator. DrQA answers a question based on the given passage and has achieved a competitive performance on CoQA (Reddy et al., 2018). During training, we apply DrQA to answer a generated question, and compare its answer with the human-provided answer (which is associated with the same rationale for generating the question)¹. If the answers match well with each other, we regard our generator produces a meaningful question since it asks about the same thing as humans do, and will assign high rewards to such questions.

Formally, we minimize the negative expected reward for a generated question:

$$J_{RL} = -\mathbb{E}_{q \sim \pi(q|r,c)}[R(a, a^*)] \quad (10)$$

where $\pi(q|r, c) = \prod_t P(y_t|y_{<t}, c, r)$ is the action policy defined in Eqn. (8) for producing question

¹We use the CoQA dataset for training and such information is available as shown in Figure 1.

Dataset	Passages	QA Pairs	Turns per Passage
Training	7199	10.8k	15.0
Dev	500	8.0k	15.9

Table 1: Statistics of the CoQA dataset.

q given rationale r and conversation history c , and $R(a, a^*)$ is the reward function defined by the F1 score² between the DrQA predicted answer a and the human-provided answer a^* . For computational efficiency concerns, during training, we make sure that the ground-truth question is in the sampling pool and use beam search to generate 5 more questions.

Note that besides providing rewards for fine-tuning our generator, DrQA model also serves another purpose: When applying our framework to any passage, we can use DrQA to produce an answer to the currently generated question so that the conversation history can be updated for the next-turn of question generation. In addition, our framework is not limited to DrQA and other more advanced QA models can apply as well.

4 Experiments

4.1 Dataset

We use the CoQA dataset³ (Reddy et al., 2018) to experiment with our ReDR and baseline methods. CoQA contains text passages from diverse domains, conversational questions and answers developed for each passage, as well as rationales (i.e., text spans extracted from given passages) to support answers. The dataset consists of 108k questions in the training set and 8k questions in the development (dev) set with a large hidden test set for competition purpose, and our results are shown on the dev set.

4.2 Baselines

As discussed earlier, CQG has been under-investigated so far, and there are few existing baselines for our comparison. Because of their high relevance with our task as well as their superior performance demonstrated by previous works, we choose to compare with the following models:

²F1 score is the common evaluation metric for QA and is defined as the harmonic mean of precision and recall.

³<https://stanfordnlp.github.io/coqa/>

Seq2Seq (Sutskever et al., 2014) is a basic encoder-decoder sequence learning system, which has been widely used for machine translation (Luong et al., 2015) and dialogue generation (Wen et al., 2017). We concatenate the rationale and the conversation history as the input sequence in our setting.

NQG (Du et al., 2017) is a strong attention-based neural network approach for question generation task. The input is the same as the above Seq2Seq model.

4.3 Implementation Details

Our word embeddings are initialized by `glove.840B.300d` (Pennington et al., 2014). We set the LSTM hidden unit size to 500 and set the number of layers of LSTMs to 2 in both the encoder and the decoder. Optimization is performed using stochastic gradient descent (SGD), with an initial learning rate of 1.0. The learning rate starts decaying at the step 15000 with a decay rate of 0.95 for every 5000 steps. The mini-batch size for the update is set at 64. We set the dropout (Srivastava et al., 2014) ratio as 0.3 and the beam size as 5. The maximum number of iterations for the dynamic reasoning is set to be 3. Since the CoQA contains abstractive answers, we apply DrQA as our question answering model and follow Yatskar (2018) to separately train a binary classifier to produce “yes” or “no” for yes/no questions⁴. Code is available at <https://github.com/ZJULearning/ReDR>.

4.4 Automatic Evaluation

Metrics We follow previous question generation work (Xu et al., 2017; Du et al., 2017) to use BLEU⁵ (Papineni et al., 2002) and ROUGE-L (Lin, 2004) to measure the *relevance* between the generated question and the ground-truth one. To evaluate the *diversity* of the generated questions, we follow (Li et al., 2016a) to calculate Dist-n ($n=1,2$), which is the proportion of unique n-grams over the total number of n-grams in the generated questions for all passages, and (Zhang et al., 2018) to use the Ent-n ($n=4$) metric, which reflects how evenly the n-gram distribution is over all generated questions. For all the metrics, the larger they are,

⁴Our modified DrQA model achieves 68.8 F1 scores on the CoQA dev set.

⁵We adopt the 4th smoothing technique as proposed in (Chen and Cherry, 2014) for short text generation.

Models	Relevance		Diversity		
	BLEU	RG-L	Dist-1	Dist-2	Ent-4
Vanilla Seq2Seq Model	7.64	26.68	0.010	0.034	3.370
NQG (Du et al., 2017)	13.97	31.75	0.017	0.068	6.518
With 1 Layer Reasoning, no RL	16.13	32.24	0.053	0.171	7.862
With 2 Layer Reasoning, no RL	17.85	33.06	0.062	0.216	8.285
With 3 Layer Reasoning, no RL	17.42	32.88	0.061	0.205	8.247
With Dynamic Reasoning, no RL	19.10	33.57	0.064	0.220	8.304
Reinforced Dynamic Reasoning (ReDR)	19.69	34.05	0.069	0.225	8.367

Table 2: Quantitative evaluation for conversational question generation using CoQA dataset.

the more relevant or diverse the generated questions are.

Results and Analysis Table 2 shows the performance of various models on the CoQA dataset. As we can see, our model ReDR and its variants perform much better than the baselines, which indicates that the reasoning procedure can significantly boost the quality of the encoding representations and thus improve the question generation performance.

To investigate the effect of the reasoning procedure and fine-tuning in our model design, we also conduct an ablation study: (1) We first test our model with only one layer of reasoning, i.e., directly feeding the encoding representation U^0 into the decoder. The results drop a lot on all the metrics, which indicates that there is abundant semantic information in the input text so the multi-layer reasoning is necessary. (2) We then augment our model with two or three layers of reasoning but without the decision maker p_d . In other words, we directly use the hidden states of the integration LSTM as the input to the next reasoning layer (formally, $U^j = \tilde{U}^j$). We can see that the performance of our model increases with a two-layer reasoning while decreases with a three-layer reasoning. We conjecture that the two-layer reasoning network is saturated for most of the input text sequences, thus directly adding a layer of network for all the input text seems not optimal. (3) When we add the decision maker to dynamically compute the encoding representations, the results are greatly improved, which demonstrates that using a dynamic procedure can distribute proper weight of each layer to the input sequences in different lengths and amount of information. (4) Finally, we fine-tune the model with the reinforcement learning framework, and the results show that using the

	NQG	ReDR	Human
Naturalness	1.94	1.92	2.14
Relevance	1.16	2.02	2.82
Coherence	1.12	1.94	2.94
Richness	1.16	2.30	2.54
Answerability	1.18	1.86	2.96

Table 3: Human evaluation results on CoQA. “Human” in the table means the original human-created questions in CoQA.

answer quality as the reward is helpful for generating better questions.

4.5 Human Evaluation

We conduct human evaluation to measure the quality of generated questions. We randomly sampled 50 questions along with their conversation history and the passage, and consider 5 aspects: *Naturalness*, which indicates the grammaticality and fluency; *Relevance*, which indicates the connection with the topic of the passage; *Coherence*, which measures whether the generated question is coherent with the conversation history; *Richness*, which measures the amount of information contained in the question. *Answerability*, which indicates whether the question is answerable based on the passage. For each sample, 5 people⁶ are asked to rank three questions (the ReDR question, the NQG question and the human-created question) by assigning each a score from $\{1,2,3\}$ (the higher, the better). For each aspect, we show the average score across the five annotators on all samples.

Table 3 shows the results of human evaluation. We can see that our method almost outperforms NQG in all aspects. For *Naturalness*, the three

⁶All annotators are native English speakers.

Category	NQG	ReDR	Human
Question Type			
“what” Question	0.45	0.42	0.35
“which” Question	0.01	0.01	0.02
“when” Question	0.07	0.05	0.04
“where” Question	0.08	0.06	0.07
“who” Question	0.06	0.22	0.15
“why” Question	0.15	0.03	0.03
yes/no Question	0.08	0.07	0.21
Linguistic Feature			
Question Length	4.05	5.34	6.48
Explicit Coref.	0.51	0.53	0.47
Implicit Coref.	0.32	0.19	0.19

Table 4: Linguistic statistics for the generated questions and the human annotated questions in CoQA.

methods obtain the similar scores, which is probably because that the most generated questions are short and fluent, makes them have no significant difference on this aspect. We also observe that on the *Relevance*, *Coherence* and *Answerability* aspects, there is an obvious gap between the generative models and human annotation. This indicates that the contextual understanding is still a challenging problem for the task of the conversational question generation.

4.6 Linguistic Analysis

We further analyze the generated questions in terms of their linguistic features and constitutions in Table 4, from which we draw three observations: (1) Overall, the distribution of the major types of questions generated by ReDR is closer to human-created questions, in comparison with NQG. For example, ReDR generates a large portion of “what” and “who” questions, similarly as humans. (2) We observe that NQG tends to generate many single-word questions such as “Why?” while our method successfully alleviates this problem. (3) Both ReDR and NQG generate fewer yes/no questions than humans, as a result of generating more “wh”-type of questions.

For the relationship between a question and its conversation history, following the analysis in CoQA, we randomly sample 150 questions respectively from each method and observe that about 50% questions generated by ReDR contain explicit coreference markers such as “he”, “she” or “it”, which is similar to the other two methods.

Once upon a time, in a barn near a farm house, there lived a little white kitten named Cotton. Cotton lived high up in a nice warm place above the barn where all of the farmer's horses slept. But Cotton wasn't alone in her little home above the barn, oh no. She shared her hay bed with her mommy and 5 other sisters...
OQ1: What color was cotton? A1: white NQG: What type of animal was it ? ReDR: What was the animal 's name ?
OQ2: Where did she live ? A2: in a barn NQG: What was it ? ReDR: What kind of house did she live ?
OQ3: Did she live alone ? A3: no NQG: Why ? ReDR: Was she alone ?
OQ4: Who did she live with? A4: with her mommy and 5 sisters NQG: What does she do ? ReDR: Who else ?

Figure 3: Example questions generated by human (i.e., original questions denoted as OQ), NQG and our ReDR on CoQA.

However, NQG generates much more questions consisting of implicit coreference markers like “Where?” or “Who?”, which can be less meaningful or not answerable as also verified in Table 3.

4.7 Case Study

In Figure 3, we show the output questions of our ReDR and NQG on an example from CoQA dataset. For the first turn, both ReDR and NQG generate a meaningful and answerable question. For the second turn, NQG generates “What was it?”, which is answerable and related to the conversation history but simpler than our question “What kind of house did she live?”. For the third turn, NQG generates a coherent but less meaningful question “Why?”, while our method generates “Was she alone?”, which is very similar to the human-created question. For the last turn, NQG produces a question that is neither coherent nor answerable, while ReDR asks a much better question “Who else?”.

To show the applicability of ReDR to generate QA style conversations on any passages, we apply it to passages in the SQuAD reading comprehension dataset (Rajpurkar et al., 2016) and show an example in Figure 4. Since there are no rationales

The game's Media Day, which was typically held on the Tuesday afternoon prior to the game, was moved to the Monday evening and rebranded as super bowl opening night. The event was held on February 1, 2016 at Sap Center in San Jose. Alongside the traditional media availabilities, the event featured an opening ceremony with player introductions on a replica of the golden gate bridge ...

Q1: What was held on Monday ?

A1: game's Media Day

Q2: Where ?

A2: Sap Center

Q3: What was the opening ceremony for ?

A3: player introductions

Figure 4: Our generated conversation on a SQuAD passage. The questions are generated by our ReDR and the answers are predicted by DrQA.

provided in the dataset for generating consecutive questions, we first apply our rule-based rationale selection as introduced in Section 3.1 and then generate a question based on the selected rationale and the conversation history. The answers are predicted by our modified DrQA. Figure 4 shows that our generated questions are closely related to the passage, e.g., the first question contains “Monday” and the third one mentions “opening ceremony”. Moreover, we can also generate interesting questions such as “Where?” which connects to previous questions and makes a coherent conversation.

5 Related Work

Question Generation. Generating questions from various kinds of sources, such as texts (Rus et al., 2010; Heilman and Smith, 2010; Mitkov and Ha, 2003; Du et al., 2017), search queries (Zhao et al., 2011), knowledge bases (Serban et al., 2016b) and images (Mostafazadeh et al., 2016), has attracted much attention recently. Our work is most related to previous work on generating questions from sentences or paragraphs. Most early approaches are based on rules and templates (Heilman and Smith, 2010; Mitkov and Ha, 2003), while Du et al. (2017) recently proposed to generate a question by a Sequence-to-Sequence neural network model (Sutskever et al., 2014) with attention (Luong et al., 2015). Other approaches such as (Zhou et al., 2017; Subramanian et al., 2017) take into account the answer information in addition to the given sentence or paragraph. (Du and Cardie,

2018; Song et al., 2018) further modeled the surrounding paragraph-level information of the given sentence. However, most of the work focused on generating standalone questions solely based on a sentence or a paragraph. In contrast, this work explores *conversational* question generation and has to additionally consider the conversation history in order to generate a coherent question, making the task much more challenging.

Conversation Generation. Building chatbots and conversational agents has been pursued by many previous work (Ritter et al., 2011; Vinyals and Le, 2015; Sordoni et al., 2015; Serban et al., 2016a; Li et al., 2016a,b). Vinyals and Le (2015) used a Sequence-to-Sequence neural network (Sutskever et al., 2014) for generating a response given the dialog history. Li et al. (2016a) further optimized the response diversity by maximizing the mutual information between inputs and output responses. Different from these work where the response can be in any form (usually a declarative statement) and is generated solely based on the dialog history, our task is potentially more challenging as it additionally restricts the generated response to be a *follow-up question* about a given *passage*.

Conversational Question Answering (CQA). CQA aims to automatically answer a sequence of questions. It has been studied in the knowledge base setting (Saha et al., 2018; Iyyer et al., 2017) and is often framed as a semantic parsing problem. Recently released large-scale datasets (Reddy et al., 2018; Choi et al., 2018) enabled studying it in the textual setting where the information source used to answer questions is a given passage, and they inspired many significant work (Zhu et al., 2018; Huang et al., 2018; Yatskar, 2018). However, collecting such datasets has heavily relied on human efforts and can be very costly. Based on one of the most popular datasets CoQA (Reddy et al., 2018), we examine the possibility of automatically *generating* conversational questions, which can potentially reduce the data collection cost for CQA.

6 Conclusion

In this paper, we introduce the task of *Conversational Question Generation* (CQG), and propose a novel framework which achieves promising performance on the popular dataset CoQA. We in-

corporate a dynamic reasoning procedure to the general encoder-decoder model and dynamically update the encoding representations of the inputs. Moreover, we use the quality of the answers predicted by a QA model as rewards and fine-tune our model via reinforcement learning. In the future, we would like to explore how to better select the rationale for each question. Besides, it would also be interesting to consider using linguistic knowledge such as named entities or part-of-speech tags to improve the coherence of the conversation.

7 Acknowledgments

This research was sponsored in part by the Army Research Office under grant W911NF-17-1-0412, NSF Grant IIS-1815674, the National Nature Science Foundation of China (grant No. 61751307), and Ohio Supercomputer Center (Center, 1987). The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notice herein.

References

- Ohio Supercomputer Center. 1987. [Ohio supercomputer center](#).
- Boxing Chen and Colin Cherry. 2014. A systematic comparison of smoothing techniques for sentence-level bleu. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 362–367.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1870–1879.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. Quac: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184.
- Xinya Du and Claire Cardie. 2018. Harvesting paragraph-level question-answer pairs from wikipedia. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1907–1917.
- 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)*, volume 1, pages 1342–1352.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1631–1640.
- Daya Guo, Yibo Sun, Duyu Tang, Nan Duan, Jian Yin, Hong Chi, James Cao, Peng Chen, and Ming Zhou. 2018. Question generation from sql queries improves neural semantic parsing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1597–1607.
- Michael Heilman and Noah A Smith. 2010. Good question! statistical ranking for question generation. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617. Association for Computational Linguistics.
- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018. Flowqa: Grasping flow in history for conversational machine comprehension. *arXiv preprint arXiv:1810.06683*.
- Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. 2017. Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1821–1831.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016b. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out*.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.

- Ruslan Mitkov and Le An Ha. 2003. Computer-aided generation of multiple-choice tests. In *Proceedings of the HLT-NAACL 03 workshop on Building educational applications using natural language processing-Volume 2*, pages 17–22. Association for Computational Linguistics.
- Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. 2016. Generating natural questions about an image. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1802–1813.
- Boyuan Pan, Hao Li, Zhou Zhao, Bin Cao, Deng Cai, and Xiaofei He. 2017. Memen: Multi-layer embedding with memory networks for machine comprehension. *arXiv preprint arXiv:1707.09098*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- 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 (EMNLP)*, pages 2383–2392.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.
- Vasile Rus, Brendan Wyse, Paul Piwek, Mihai Lintean, Svetlana Stoyanchev, and Christian Moldovan. 2010. The first question generation shared task evaluation challenge. In *Proceedings of the 6th International Natural Language Generation Conference*.
- Amrita Saha, Vardaan Pahuja, Mitesh M Khapra, Karthik Sankaranarayanan, and Sarath Chandar. 2018. Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. *ICLR*.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016a. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Iulian Vlad Serban, Alberto García-Durán, Caglar Gulcehre, Sungjin Ahn, Sarath Chandar, Aaron Courville, and Yoshua Bengio. 2016b. Generating factoid questions with recurrent neural networks: The 30m factoid question-answer corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 588–598.
- Linfeng Song, Zhiguo Wang, Wael Hamza, Yue Zhang, and Daniel Gildea. 2018. Leveraging context information for natural question generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 569–574.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A neural network approach to context-sensitive generation of conversational responses. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Sandeep Subramanian, Tong Wang, Xingdi Yuan, Saizheng Zhang, Yoshua Bengio, and Adam Trischler. 2017. Neural models for key phrase detection and question generation. *arXiv preprint arXiv:1706.04560*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 438–449.

- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256.
- Caiming Xiong, Victor Zhong, and Richard Socher. 2017. Dynamic coattention networks for question answering. *ICLR*.
- Zhen Xu, Bingquan Liu, Baoxun Wang, SUN Chengjie, Xiaolong Wang, Zhuoran Wang, and Chao Qi. 2017. Neural response generation via gan with an approximate embedding layer. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 617–626.
- Mark Yatskar. 2018. A qualitative comparison of coqa, squad 2.0 and quac. *arXiv preprint arXiv:1809.10735*.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In *Advances in Neural Information Processing Systems*, pages 1815–1825.
- Shiqi Zhao, Haifeng Wang, Chao Li, Ting Liu, and Yi Guan. 2011. Automatically generating questions from queries for community-based question answering. In *Proceedings of 5th international joint conference on natural language processing*, pages 929–937.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671. Springer.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.