

An Empirical Study of Content Understanding in Conversational Question Answering

Ting-Rui Chiang, Hao-Tong Ye, Yun-Nung Chen

National Taiwan University, Taipei, Taiwan
{r07922052, r08922065}@csie.ntu.edu.tw, y.v.chen@ieee.org

Abstract

With a lot of work about context-free question answering systems, there is an emerging trend of conversational question answering models in the natural language processing field. Thanks to the recently collected datasets, including QuAC and CoQA, there has been more work on conversational question answering, and recent work has achieved competitive performance on both datasets. However, to best of our knowledge, two important questions for conversational comprehension research have not been well studied: 1) How well can the benchmark dataset reflect models' content understanding? 2) Do the models well utilize the conversation content when answering questions? To investigate these questions, we design different training settings, testing settings, as well as an attack to verify the models' capability of content understanding on QuAC and CoQA. The experimental results indicate some potential hazards in the benchmark datasets, QuAC and CoQA, for conversational comprehension research. Our analysis also sheds light on both what models may learn and how datasets may bias the models. With deep investigation of the task, it is believed that this work can benefit the future progress of conversation comprehension. The source code is available at <https://github.com/MiuLab/CQA-Study>.

1 Introduction

Answering questions in the conversational manner has become an important task for machine reading comprehension. There are two benchmark conversational question answering datasets, QuAC (Choi et al. 2018) and CoQA (Reddy, Chen, and Manning 2019). Different from traditional machine reading comprehension (Rajpurkar et al. 2016; Nguyen et al. 2016; Rajpurkar, Jia, and Liang 2018) whose questions are context-free, questions and answers in QuAC and CoQA are collected in a conversational manner. Same as original machine reading comprehension, questions related to a given passage are asked. However, questions may be also related to the given conversational history, and should be answered accordingly. Such conversational setting is regarded as more practical because people tend to seek for information in a conversational way. QuAC and CoQA also feature many linguistic phenomena unique to conversations, so they are be-

lieved to be important materials for research about conversational question answering.

This work focuses on investigating *how well the performance of a model on these two benchmark datasets reflects its capability of comprehension*. If higher performance on these datasets does not necessarily imply better conversational comprehension, then further investigation must be done when models claim their better understanding performance. However, it has not been well investigated by any of the prior work (Choi et al. 2018; Huang, Choi, and tau Yih 2019; Zhu, Zeng, and Huang 2018; Yeh and Chen 2019).

In this paper, we further analyze *whether the recent models achieving competitive performance rely on content comprehension*. It is motivated by the fact that the position of the answer to the previous question is widely utilized in many of previous conversational question answering models, such as BiDAF-with-ctx (Seo et al. 2016) and FlowQA (Huang, Choi, and tau Yih 2019). Those models leverage the datasets' property that answers or rationals can always be found as a span of the given passage. Thus those models can access the informative content provided by the position information. Models are expected to learn to understand the conversation based on its content. However, it is not clear whether the models rely on the previous conversation content or merely the position information.

The goal of this paper is to investigate the above two questions and to guide the future research on conversational question answering. We focus on recent open-sourced models, including FlowQA (Huang, Choi, and tau Yih 2019), BERT (Devlin et al. 2019), and SDNet (Zhu, Zeng, and Huang 2018). Similar to Sankar et al. (2019), we design a set of experiments consisting of different training settings to address the first question and different testing settings to address the second question. The results express some concerns about conversational QA models: 1) Higher performance on QuAC and CoQA does not necessarily imply better content comprehension. 2) Models trained on QuAC show the tendency of heavily relying on the previous answers' positions rather than their textual content. By pointing out these issues, some potential hazards may be avoided in the future research. Furthermore, our presented experiments can serve as an analysis tool for conversational ques-

tion answering models in the future.

2 Related Work

Machine reading comprehension has attracted lots of interests in recent years. Many datasets (Rajpurkar et al. 2016; Rajpurkar, Jia, and Liang 2018; Nguyen et al. 2016; Trischler et al. 2017; Lai et al. 2017; Shao et al. 2018) have been created, and there has been many models (Wang and Jiang 2016; Weissenborn, Wiese, and Seiffe 2017; Hu et al. 2018; Xiong, Zhong, and Socher 2018; Shen et al. 2017; Liu et al. 2018) crafted for advancing the tasks. The above datasets and models are for context-free question answering. Recently, with the trend of conversational interactions, people started to focus on conversational understanding. Therefore, datasets for conversational question answering were built for promoting this research direction (Saha et al. 2018; Choi et al. 2018; Reddy, Chen, and Manning 2019), and a lot of models (Huang, Choi, and tau Yih 2019; Zhu, Zeng, and Huang 2018) were proposed to tackle this challenge.

Most investigation on what machine comprehension models learn are for those context-free question answering systems. Jia and Liang (2017) proposed a method to generate adversarial examples in order to test the model robustness. Kaushik and Lipton (2018) conducted experiments to verify the reading required to answer questions in the dataset. Weissenborn, Wiese, and Seiffe (2017) found a feature indicating if a word appears in the question important, and suggested that questions can be answered with some rules that rely only on superficial features. Rondeau and Hazen (2018) validated the suggestion by a series of systematic experiments. However, the conversational question answering systems have been rarely explored. Although Yatskar (2019) compared CoQA, SQuAD 2.0 and QuAC qualitatively, there was no investigation on what conversational question answering models capture.

The phenomenon that models does not utilize all useful features is common in diverse areas. For example, Nie, Wang, and Bansal (2019) showed that natural language inference models with high accuracy relied much on the lexical-level features but utilized little compositional semantics. (Sankar et al. 2019) found that not whole conversation history is used in neural dialogue generation systems. Similar phenomena happened in the computer vision area, where Geirhos et al. (2019) indicated that CNN trained on ImageNet relied much on textual information rather than shape information, and Brendel and Bethge (2019) further showed that only textual information can achieve very high accuracy on ImageNet. Our work extends a similar idea to conversational question answering systems and demonstrates that positional information rather than semantic information can be exploited to master some tasks.

3 Conversational Question Answering

There are two main datasets for conversational question answering, QuAC (Choi et al. 2018) and CoQA (Reddy, Chen, and Manning 2019). Questions and answers are collected in a conversational manner, where each conversation includes two participants: a student who asks question about a given

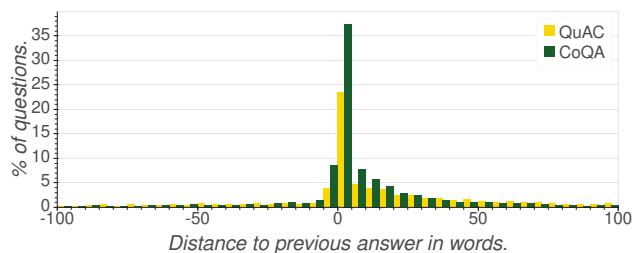


Figure 1: The histogram of distance between answers to consecutive questions of a conversation in words. The bin size is 5 words. The distance is counted as zero if the two answer spans are overlapped, and is positive if the current is after the previous answer, negative otherwise.




Figure 2: Histogram of answer length distribution in QuAC and CoQA. The bin size is 1 word.

passage, and a teacher who answers the question according to the passage. The teacher in both sets may reply “no answer” if the answer cannot be found in the given passage. When evaluating the model for both tasks, the content as well as the position of answers to the previous question is available to the model. In the statistics of two datasets, answers to the consecutive questions tend to be close to each other depicted in Figure 1.

Even though both datasets are collected for conversational question answering, they have several different properties.

- **Answer format**
Answers in QuAC are always the text spans in the given passage, while answers in CoQA are free texts similar to some spans in the passage. Answers in QuAC is generally longer than answers in CoQA shown in Figure 2, where the distribution implies that QuAC is more realistic than CoQA. Answering “yes” or “no” is also allowed in CoQA. Note that the evidence span (span in the passage that supports the answer) is provided in CoQA, so the previous answers’ position information is still available.
- **Dataset collection process**
The Amazon mechanical turkers who generated the CoQA dataset have full access to the passage. On the other hand, turkers who generated questions in QuAC cannot see the passages. The latter setting may be more suitable for practical applications, because real users want to seek for information using questions when not reading passages.

4 Models

We consider models including FlowQA, BERT, and SDNet, as they are the only publicly available models till now. For FlowQA and SDNet, we use the code released by the authors. Some modifications are made for the following experiments¹. Each model of each setting are trained with 3 different random seeds, and the resulted mean and standard deviation value are reported for reliability.

All models in the experiments are mainly based on those designed for single-turn reading comprehension tasks, so this section focuses on describing the modification for each method in order to handle understanding in conversational question answering. Below three models are detailed.

FlowQA

FlowQA (Huang, Choi, and tau Yih 2019) is the model specifically designed for conversational question answering, which contains a mechanism that can incorporate intermediate representations generated during the process of answering previous questions. This model significantly improved conversational machine comprehension tasks for both QuAC and CoQA data. In FlowQA, the main mechanism designed for the conversational structure is the INTEGRATION-FLOW layer. In the model, there is a question-aware context representation C_i for each question Q_i in the history, and the INTEGRATION process simply applies BiLSTM to each C_i independently. After that, the FLOW process applies BiLSTM to the same word across different context representations in order to capture knowledge of previous questions. Also, an *one-bit* feature is added to the input context word representation indicating whether the word appears in previous answers. The rest of the reasoning procedure is almost identical to FusionNet (Huang et al. 2018) that focuses on single-turn machine comprehension.

BERT

In the current state-of-the-art question answering models, most models leverage the benefits from BERT (Devlin et al. 2019) to advance the task. We apply BERT on QuAC by converting the task into a single-turn machine reading comprehension task such as SQuAD (Rajpurkar et al. 2016). We prepend previous N questions to the current question Q_k , so it becomes $\hat{Q}_k = \{Q_{k-N}, \dots, Q_{k-1}, Q_k\}$. At the embedding layer, beside the original word embedding, segment embedding and position embedding, an additional embedding is also added to represent whether the word appears in previous answer spans. Then we follow the procedure of applying BERT to SQuAD that concatenates the extended and the context to form the input and uses the context output representations from BERT to predict the start and the end of the answer span (Devlin et al. 2019).

SDNet

To incorporate the information from dialogue history, SDNet prepends not only previous questions but also previous answers to the current question Q_k , i.e., $\hat{Q}_k =$

¹For reproducibility, we will release all code, script, and experiment settings.

$\{Q_{k-N}, A_{k-N}, \dots, Q_{k-1}, A_{k-1}, Q_k\}$, and no more additional effort is put to deal with the conversational structure. For the input representation for both context and question words, they used BERT as contextualized embeddings along with GloVe. The rest of the model architecture for reasoning is highly inspired by FusionNet (Huang et al. 2018).

5 How Well the Performance Reflects Content Comprehension?

This section attempts at investigating how well the performance reflects the capability of comprehension on QuAC and CoQA? Both datasets claim rich linguistic phenomena unique to conversations, where QuAC claims to have 61% of questions including coreference referring to entities in the given passage, 44% of coreference referring to entities in previous history, and 11% of questions that ask for more information in the conversation. Also, CoQA claims to have 49.7% of questions with explicit coreference to conversations, and 19.8% with implicit ones. Given such high ratio of questions related to conversations, it is natural to expect that *higher performance implies better understanding in conversational question answering*. Especially, the understanding should be based on the content of the conversation, in which those special linguistic phenomena is embodied. To inspect the expectation, we design experiments based on the premise: If comprehension on the content of conversation is reflected well by the performance, then model trained without the access to conversation content should not achieve high performance.

Experimental Settings

We compare models trained and tested with three different settings:

- **Original**: The model has free access to the previous conversation history, as the setting proposed by the models.
- **- text**: The model has no access to the content of the answer to the previous question, but has access to their position in the provided context. The previous questions are not used either.
- **- conversation**: The model has *no any* access to the previous conversation history.

In the **- text** setting, the answer span in the passage is masked with zeros. As questions are to seek for information, the answer content should be highly informative. Therefore, the information loss by answer masking cannot be easily compensated by the surrounding words. Especially, for QuAC, since answer spans are typically as long as sentences, masked language models like BERT can *in no way* recover the masked part. Thus, in this setting, the only information about previous answers is their positions in the passage.

Discussion

We compare the results of models in Table 1. For both QuAC and CoQA, models trained without access to conversation *content* can achieve performance significantly better than models trained without access to any conversation history. Especially for QuAC, **- text** models consistently outperform

Dataset	Model	F1 (stdev)
QuAC	FlowQA	64.4 (.30)
	BERT	63.6 (.16)
	FlowQA - text	62.4 (.20)
	BERT - text	62.2 (.48)
CoQA	FlowQA - conversation	54.1 (.13)
	BERT - conversation	55.3 (.03)
	FlowQA	76.9 (.22)
	SDNet + position	76.4 (.31)
CoQA	FlowQA - text	71.5 (.24)
	SDNet + position - text	74.0 (.19)
	FlowQA - conversation	63.4 (.11)
	SDNet - conversation	68.3 (.45)

Table 1: Model performance on the validation set of QuAC and CoQA. Note that the original SDNet model does not utilize the position information, so *SDNet + position - text* is a modified SDNet with additional one dimension feature indicating the previous answer as the input.

- conversation models by up to absolute 7% F1 score. As for CoQA, though not as consistently, but similar comparison can also be observed.

The above results indicate that better understanding in the dataset is not well reflected by the performance on these two datasets. Undoubtedly, better comprehension should be based on semantic understanding specific to the textual content. However, even no content is provided to - text models, - text models can still achieve performance higher than - conversation models. It indicates that *higher performance* does not necessarily imply *better content understanding*. Therefore, future models may need to further verify whether they indeed focus on semantic understanding instead of utilizing the position information only.

6 Do Models Understand Conversation Content?

The results in §5 do not answer the question that if the full models understand content of conversations well. As shown in §5, the full models can achieve better performance than models use only the answers' position information. However, §5 also shows the usefulness of the previous position information. Since the full model has access to both the position and content of the previous answers, it is not clear whether the better performance is contributed by understanding conversation content. To answer this question, we analyze the trained models by a series of testing settings.

Repeat Attack

We propose *repeat attack* that increases the distance between answers in the context. To do so, a text span is repeated between the answer spans in the passage. For QuAC, as most answers are sentences, we repeat the answer sentence for each answer spans in the passage. For CoQA, since

Repeat Attack Example in QuAC

Passage: In 2004, Oldman returned to prominence when he landed a significant role in the Harry Potter film series, playing Harry Potter's godfather Sirius Black.^{ans1} *In 2004, Old man returned to prominence when he landed a significant role in the Harry Potter film series, playing Harry Potter's godfather Sirius Black.* The following year, he starred as James Gordon in Christopher Nolan's commercially and critically successful *Batman Begins*.^{ans2} starred as James Gordon in Christopher Nolan's commercially and critically successful *Batman Begins*, ...

Question 1: What was his resurgence or comeback role?

Answer 1: In 2004, Oldman returned to prominence when he landed a significant role in the Harry Potter film series, playing Harry Potter's godfather Sirius Black.

Question 2: Are there any other interesting aspects about this article?

Answer 2: The following year, he starred as James Gordon in Christopher Nolan's commercially and critically successful *Batman Begins*,

Origin FlowQA prediction: (F1 0.83) The following year, he starred as James Gordon in Christopher Nolan's commercially and critically successful *Batman Begins*,

Origin BERT prediction: (F1 0.83) The following year, he starred as James Gordon in Christopher Nolan's commercially and critically successful *Batman Begins*

Attacked FlowQA prediction: (F1 0.20) In 2004, Oldman returned to prominence when he landed a significant role in the Harry Potter film series, playing Harry Potter's godfather Sirius Black.

Attacked BERT prediction: (F1 0.21) In 2004, Oldman returned to prominence when he landed a significant role in the Harry Potter film series, playing Harry Potter's godfather Sirius

Figure 3: An example of repeat attack on QuAC and the model results. The red italic text is the inserted attack.

Repeat Attack Example in CoQA

Passage: Once upon a time, in a barn.^{ans2} near a farm house, there lived a little white^{ans1} kitten named Cotton. *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**.^{ans3} in her little home above the barn, oh no. ...

Question 1: What color was Cotton?

Answer 1: white

Question 2: Where did she live?

Answer 2: in a barn

Question 3: Did she live alone?

Answer 3: Cotton wasn't alone

Figure 4: An example of repeat attack on CoQA. The red italic text is the inserted attack. In this example, the attack increases the distance between answer 2 and 3.

its answer span is generally much shorter than QuAC, we repeat sentences that contain the answer span. The attack




Figure 5: Histogram of distance between answers to consecutive questions in words *after* the attack. The bin size is 5.

examples for both datasets are shown in Table 3 and Table 4. By repeating part of the passage, the meaning delivered should remain the same. When evaluating the models, the previous answer positions provided to models contain only the text as in original answer. Due to the repeated text, the distance between consecutive answer spans is lengthened by this attack. The distribution of the distance is visualized in Figure 5, which is much smoother than the one before the attack (Figure 1).

We use *repeat attack* to investigate models’ understanding of conversation content. It is motivated by the high ratio of answers close to answers to the previous questions (as shown in Figure 1). It is possible that positions of previous answers leak the position information of the current answer. The model may thus learn to take as the answer candidates the sentences close to the previous answer span. On the other hand, if a model does answer the question by understanding answer content, then the model should be robust to this attack. Therefore, by using the attacked data to test models that are trained on the normal data, the models’ capability of understanding can be well investigated.

For QuAC, the results shown in Table 2 indicates that both FlowQA and BERT are sensitive to the distance between consecutive answers. Their performance drops significantly when applying the repeat attack. The F1 score under attack is roughly the F1 score of models trained without any conversational information. Furthermore, we conduct the same experiments on FlowQA and BERT trained without using position information of the previous answers, and find that although they perform much worse than the full models, they are more robust against the attack.

To better investigate how the answer position is related to the model robustness, we plot the relation between F1 scores before/after attack in terms of the distance to previous answer in Figure 6. We find that both BERT and FlowQA predict answers more accurately when the current answer has short distance to the previous answer (green and red lines). Also, our attack is more effective when the distance is short (yellow line). These findings imply that models trained on QuAC with the position information may rely less on the answer content but rely too much on the answer position. In addition, models trained without the position information may rely more on the semantic information in the conversations.

We also investigate the performance affected by the at-

Dataset	Model	Repeat Attack	
		w/o	w/
QuAC	FlowQA	64.4 (0.30)	53.3 (0.75)
	BERT	63.6 (0.16)	55.3 (0.74)
	FlowQA - position	59.3 (0.37)	56.9 (0.37)
	BERT - position	58.0 (0.18)	56.5 (0.22)
CoQA	FlowQA	76.9 (0.22)	72.0 (0.20)
	SDNet	76.4 (0.31)	71.8 (0.48)
	FlowQA - position	76.8 (0.42)	72.4 (0.24)
	SDNet - position	75.7 (0.76)	70.9 (0.10)

Table 2: F1 score on the validation of QuAC and CoQA with or without attack. ”- position” indicates training without the position information of answers to previous questions.

tack based on if the question is a followup of the previous question, which is annotated in QuAC dataset. The followup questions are questions more likely to depend on the previous question, and therefore are expected to require understanding of conversation content. However, results shown in table 3 indicate that they are more vulnerable to our attack.

QuAC models’ reliance on position information can be further shown by the qualitative error analysis on the attacked validation set. We inspect the questions that the models originally predict the answer with a high F1 score, but predict the answer with a low F1 scores when applying attack. One sample is shown in Table 3, where two models under attack directly predict the sentence after the previous answer span regardless of the content. On other cases, we find FlowQA under attack makes mistakes by selecting the next span more consistently. In contrast, BERT under attack is prone to predict a much implausible answer or reply “no answer”. The reason may be that BERT is specialized to search answers in the region near the previous answer, so consequently, when the next sentence after the previous answer is incorrect, it will fail to answer the question correctly.

On the other hand, it is less clear if the drop of performances on CoQA dataset is due to insufficient of conversation content understanding. Unlike on QuAC, - position models are not more robust than the original models as obviously as on QuAC. By plotting figures for CoQA in the same setting, we find that questions in CoQA seem to be susceptible to the attack equally regardless of the distance to previous answers. Conversation content understanding of the models trained on CoQA may require further investigation.

Predict without Previous Answer Text

To investigate if the content of previous answers is used by the models trained with position information, we measure the performance on the validation set predicted without the content of previous answers. To remove the content information of previous answers, we mask the previous answer spans with zeros. Different to previous - text settings, models here has access to previous questions. Particularly, for FlowQA, the content information of the previous answer may be flowed along with the flow structure, and the RNN




Figure 6: F1 change between before and after repeat attack in terms of the answer’s distance to its previous answers on QuAC (left: FlowQA; right: BERT). The x-coordinate is the answer’s distance to the previous answer in words, and the meaning of zero, positive and negative is same as in Figure 1. The y-coordinate is the average F1 score. “no answer” sample are ignored here. The dotted line is the average F1 score of all samples.

Model		Followup	No Followup
FlowQA	no attack	64.0 (0.59)	64.8 (0.11)
	attack	49.5 (0.59)	57.6 (0.94)
	Δ	14.6 (1.17)	7.2 (1.00)
BERT	no attack	63.4 (0.32)	64.0 (0.57)
	attack	52.4 (1.37)	58.5 (0.18)
	Δ	11.0 (1.12)	5.5 (0.41)

Table 3: Impact on F1 score by the attack.

Dataset	Model	Ans. Mask	
		w/o	w/
QuAC	FlowQA	64.4 (0.30)	60.5 (0.54)
	BERT	63.6 (0.16)	52.6 (1.76)
	FlowQA - position	59.3 (0.37)	55.0 (1.80)
	BERT - position	58.0 (0.18)	50.1 (0.45)
CoQA	FlowQA	76.9 (0.22)	71.2 (0.41)
	SDNet	76.4 (0.31)	73.5 (0.51)
	FlowQA - position	76.8 (0.42)	68.2 (0.28)
	SDNet - position	75.7 (0.76)	75.6 (0.69)

Table 4: F1 results on the validation sets of QuAC and CoQA. Models are inferred without/with applying masks on the previous answers. Models without suffix are trained with full access to conversation, while “- position” indicates that the models are trained without the previous answer position information. Note that in both training settings, masks are not applied.

memory of the previous answer spans is also reset to zeros.

The more performance drops when masking previous answers implies that the model relies more on the content information of those answers. According to the results in Table 4, it seems that all models more or less rely on the text of previous answers. Meanwhile, as expected, the models except SDNet - position trained without the position information almost drop to the performance of ones trained without any conversation history information. SDNet - position is an exception because answers to previous questions are prepended to the question when training and testing, so masking answers in the passage does not remove all the content of previous answers. However, it is surprising that FlowQA on QuAC can still keep the performance up to 60% F1, implying that FlowQA may rely on position information much more than the semantic information.

Predict without Previous Answer Position

To directly test to what extent the models rely on position information of previous answers, we conduct the experiments of predicting answers without position information. The results are shown in Table 5, where both results of FlowQA and BERT on QuAC drop significantly if not using position information. Among them, FlowQA drops even more, regardless of the flow structure that models the dialog flow. The performance is even lower than the models trained with-

out using any conversation history. It indicates that although the models on QuAC may rely on the semantics of previous answers, the position information is indeed exploited. On the other hand, it is interesting to see that FlowQA trained on CoQA rely little position information of the previous answers.

Implication of Above Experiments

Our results show that the models trained on QuAC have high tendency to rely heavily on the position of previous answers. The proposed attack and experiment settings can serve as a diagnosis tool in the future.

7 Dataset and Model Analysis

To further investigate the difference between models trained on QuAC and CoQA, two questions are focused here.

Dataset	Model	Position Info.	
		w/	w/o
QuAC	FlowQA	64.4 (0.30)	48.1 (0.72)
	BERT	63.6 (0.16)	54.9 (0.08)
CoQA	FlowQA	76.9 (0.22)	76.7 (0.36)
	SDNet	76.4 (0.31)	73.7 (0.18)

Table 5: F1 score on the validation sets of QuAC and CoQA. Models are inferred with/without access to position information, but trained with access to position information.

Dataset	Model	Shuffle	
		w/o	w/
CoQA	FlowQA - position	76.8 (0.42)	71.7 (0.25)
	SDNet - position	75.7 (0.76)	70.1 (0.59)

Table 6: F1 score of models on shuffled validation set of CoQA. Models here do not use the position information of the previous answers.

Why do CoQA models rely less on position information?

It is unclear why the models trained on CoQA rely less on the position information of the previous answer as shown in the previous sections. Figure 2 shows that answers in CoQA are much shorter than in QuAC, so if we normalize the distance to the previous answer related to the length of the answers, the average distance to the previous answer in CoQA would be much longer than QuAC. Furthermore, short answers also imply that an answer can be identified as the sentence containing the answer. To verify the above claim, we randomly shuffle the sentences in the passage for each CoQA example, so that the cross-sentence information and the sentence order are removed from the passage. Then we use CoQA models trained without position information for answer prediction. In Table 6, roughly speaking, up to 70% questions can still be answer correctly. It thus supports our claim and partly explains why models trained on CoQA rely less on the position information.

Why do QuAC models rely more position information?

We describe the potential reason why the models trained on QuAC rely much on the position information. According to the analysis on QuAC (Choi et al. 2018), 11% of the questions is of the type “*Anything else?*”. It is common in the general article where the information of the same type is written in near contexts. Because the question “*Anything else?*” is asked to seek for more information similar to the previous answer, it is very likely that the position of the previous answer provides a strong hint for the current answer. Though there is a high percentage of the questions containing pronouns in QuAC, they do not necessarily force the model to learn coreference resolution either. For the pronouns in questions, they often refer to entities in the previous answer. Because entities in consecutive sentences seldom change much, simply looking for the answer near the previous answer location may be sufficient to answer the

question. Especially, we find personal pronouns in QuAC questions of a conversation often refer to only one same person. Also, the referred person is often the main role in the passage. This further removes the necessity to understand the conversation history.

8 Conclusion

This paper investigates content understanding of different models learned from different datasets. The experiments shows concerns: 1) Performance on QuAC and CoQA does not well reflect the model’s comprehension on conversation content. 2) The model trained on QuAC does not necessarily learn conversation comprehension. 3) In CoQA, cross-sentence information is not that important for the current model. By pointing out these concerns, we suggest future directions in both aspects of data collection and model development.

In terms of data collection, more realistic collection strategies should be adopted. In the collection process of QuAC, for example, workers are encouraged to create *long* conversations with *few* unanswerable questions, causing the worker to ask questions conservatively and passively ask questions very likely answerable. That may consequently result in the shorter distance between consecutive answers. In addition, it is unknown that if such setting is realistic. Is such information seeking pattern practical in real scenarios? Therefore, inventing a more realistic data collection process may be an important future direction.

In terms of model development, conversational question answering systems should address more about semantic comprehension. As we have shown, better performance does not necessarily imply better content comprehension, so proving the model’s conversational understanding still remains challenging. It is also unclear how to design a model that learns conversation comprehension naturally. The future work could focus more studies on those directions, including generalization and robustness to diverse scenarios.

Acknowledgements

We thank reviewers for their insightful comments. This work was financially supported from the Young Scholar Fellowship Program by Ministry of Science and Technology (MOST) in Taiwan, under Grant 108-2636-E002-003 and 108-2634-F-002-019.

References

- Brendel, W., and Bethge, M. 2019. Approximating CNNs with bag-of-local-features models works surprisingly well on imagenet. In *International Conference on Learning Representations*.
- Choi, E.; He, H.; Iyyer, M.; Yatskar, M.; Yih, W.-t.; Choi, Y.; Liang, P.; and Zettlemoyer, L. 2018. Quac: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2174–2184.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for lan-

- guage understanding. 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)*, 4171–4186.
- Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2019. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*.
- Hu, M.; Peng, Y.; Huang, Z.; Qiu, X.; Wei, F.; and Zhou, M. 2018. Reinforced mnemonic reader for machine reading comprehension. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 4099–4106. AAAI Press.
- Huang, H.-Y.; Zhu, C.; Shen, Y.; and Chen, W. 2018. Fusionnet: Fusing via fully-aware attention with application to machine comprehension. In *International Conference on Learning Representations*.
- Huang, H.-Y.; Choi, E.; and tau Yih, W. 2019. FlowQA: Grasping flow in history for conversational machine comprehension. In *International Conference on Learning Representations*.
- Jia, R., and Liang, P. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2021–2031.
- Kaushik, D., and Lipton, Z. C. 2018. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 5010–5015.
- Lai, G.; Xie, Q.; Liu, H.; Yang, Y.; and Hovy, E. 2017. Race: Large-scale reading comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 785–794.
- Liu, X.; Shen, Y.; Duh, K.; and Gao, J. 2018. Stochastic answer networks for machine reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1694–1704.
- Nguyen, T.; Rosenberg, M.; Song, X.; Gao, J.; Tiwary, S.; Majumder, R.; and Deng, L. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*.
- Nie, Y.; Wang, Y.; and Bansal, M. 2019. Analyzing compositionality-sensitivity of nli models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6867–6874.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2383–2392.
- Rajpurkar, P.; Jia, R.; and Liang, P. 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)*, 784–789.
- Reddy, S.; Chen, D.; and Manning, C. D. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics* 7:249–266.
- Rondeau, M.-A., and Hazen, T. J. 2018. Systematic error analysis of the stanford question answering dataset. In *Proceedings of the Workshop on Machine Reading for Question Answering*, 12–20.
- Saha, A.; Pahuja, V.; Khapra, M. M.; Sankaranarayanan, K.; and Chandar, S. 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*.
- Sankar, C.; Subramanian, S.; Pal, C.; Chandar, S.; and Bengio, Y. 2019. Do neural dialog systems use the conversation history effectively? an empirical study. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 32–37. Florence, Italy: Association for Computational Linguistics.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2016. Bidirectional attention flow for machine comprehension. In *Proceedings of ICLR*.
- Shao, C. C.; Liu, T.; Lai, Y.; Tseng, Y.; and Tsai, S. 2018. Drcd: a chinese machine reading comprehension dataset. *arXiv preprint arXiv:1806.00920*.
- Shen, Y.; Huang, P.-S.; Gao, J.; and Chen, W. 2017. Reasonet: Learning to stop reading in machine comprehension. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1047–1055. ACM.
- Trischler, A.; Wang, T.; Yuan, X.; Harris, J.; Sordani, A.; Bachman, P.; and Suleman, K. 2017. Newsqa: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, 191–200.
- Wang, S., and Jiang, J. 2016. Machine comprehension using match-lstm and answer pointer.
- Weissenborn, D.; Wiese, G.; and Seiffe, L. 2017. Making neural qa as simple as possible but not simpler. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, 271–280.
- Xiong, C.; Zhong, V.; and Socher, R. 2018. DCN+: Mixed objective and deep residual coattention for question answering. In *International Conference on Learning Representations*.
- Yatskar, M. 2019. A qualitative comparison of coqa, squad 2.0 and quac. 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)*, 2318–2323.
- Yeh, Y.-T., and Chen, Y.-N. 2019. Flowdelta: Modeling flow information gain in reasoning for conversational machine comprehension. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, 86–90.
- Zhu, C.; Zeng, M.; and Huang, X. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.