

FlowDelta: Modeling Flow Information Gain in Reasoning for Conversational Machine Comprehension

Yi-Ting Yeh

National Taiwan University
r07922064@csie.ntu.edu.tw

Yun-Nung Chen

National Taiwan University
y.v.chen@ieee.org

Abstract

Conversational machine comprehension requires deep understanding of the dialogue flow, and the prior work proposed FlowQA to implicitly model the context representations in reasoning for better understanding. This paper proposes to explicitly model the *information gain* through dialogue reasoning in order to allow the model to focus on more informative cues. The proposed model achieves state-of-the-art performance in a conversational QA dataset QuAC and sequential instruction understanding dataset SCONE, which shows the effectiveness of the proposed mechanism and demonstrates its capability of generalization to different QA models and tasks ¹

1 Introduction

Machine reading comprehension has been increasingly studied in the NLP area, which aims to read a given passage and then answer questions correctly. However, human usually seeks answers in a conversational manner by asking follow-up questions given the previous answers. Traditional machine reading comprehension (MC) tasks such as SQuAD (Rajpurkar et al., 2016) focus on a single-turn setting, and there is no connection between different questions and answers to the same passage. To address the multi-turn issue, several datasets about conversational question answering (QA) were introduced, such as CoQA (Reddy et al., 2018) and QuAC (Choi et al., 2018).

Most existing machine comprehension models apply single-turn methods and augment the input with question and answer history, ignoring previous reasoning processes in the models. Recently proposed FlowQA (Huang et al., 2018) attempted at modeling such multi-turn reasoning in dialogues in order to improve performance for

¹Our code can be found in <https://github.com/MiuLab/FlowDelta>.

Conversation Flow (over Context)




Figure 1: Illustration of the flow information gain modeled by the FlowDelta mechanism.

conversational QA. However, the proposed FLOW operation is expected to incorporate salient information in an *implicit* manner, because the learned representations captured by FLOW would change during multi-turn questions. It is unsure whether such change correlates well with the current answer or not. In order to *explicitly* model the information gain in FLOW and further relate the current answer to the corresponding context, we present a novel mechanism, FlowDelta, which focuses on modeling the difference between the learned context representations in multi-turn dialogues illustrated in Figure 1. The contributions are 3-fold:

- This paper proposes a simple and effective mechanism to explicitly model information gain in flow-based reasoning for multi-turn dialogues, which can be easily incorporated in different MC models.
- FlowDelta consistently improves the performance on various conversational MC datasets, including CoQA and QuAC.
- The proposed method achieves the state-of-the-art results on QuAC and sequential instruction understanding task (SCONE).

2 Background

Given a document (context), previous conversation history (i.e., question/answer pairs) and the current question, the goal of conversational QA is to find the correct answer. We denote the

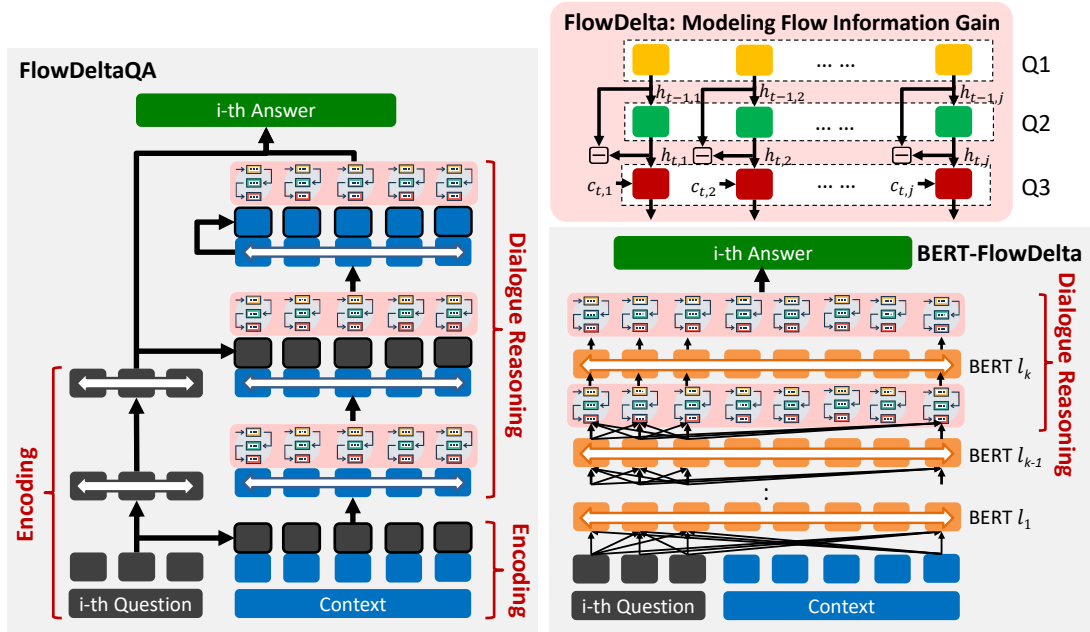


Figure 2: Illustration of the proposed FlowDelta models.

context document as a sequence of m words $C = \{c_1, c_2, \dots, c_m\}$, and the i -th question $Q_i = \{q_1, q_2, \dots, q_n\}$ as a sequence of n words. In the extractive setting, the i -th answer A_i is guaranteed to be a span in the context. The main challenge in conversational QA is that current question may depend on the conversation history, which differs from the classic machine comprehension. Therefore, how to incorporate previous history into the QA model is especially important for better understanding. Prior work (Huang et al., 2018) proposes an effective way to model the reasoning in multi-turn dialogues summarized below.

FLOW Operation Instead of only using shallow history like previous questions and answers, Huang et al. (2018) proposed the FLOW operation that feeds the model with entire hidden representations generated during the reasoning process when answering previous questions. FLOW is defined as *a sequence of latent representations based on the context tokens* and is demonstrated effective for conversational QA tasks, because it well incorporates multi-turn information in dialogue reasoning.

Let the context representation for i -th question be $C_i = c_{i,1}, \dots, c_{i,m}$ and the dialogue length is t . When answering questions in the dialogue, there are t context sequences of length m , one for each question. We reshape it to become m sequences of length t , one for each context word, and then pass each sequence into a unidirectional GRU. All context word representation j ($1 \leq j \leq m$) are

processed in parallel in order to model the information via the FLOW direction (vertical direction illustrated in Figure 1).

$$h_{1,j}, \dots, h_{t,j} = GRU(c_{1,j}, \dots, c_{t,j}) \quad (1)$$

Then we reshape the outputs from GRU back and form $F_i = \{h_{i,1}, \dots, h_{i,m}\}$, where F_i is the output of the FLOW layer.

FlowQA The FLOW layer described above is incorporated in FLOWQA for conversational MC, which is built on the single-turn MC model FusionNet (Huang et al., 2017), and the full structure is shown in the left part of Figure 2. Briefly, FLOWQA first performs word-level attention to fuse the information of i -th question Q_i into context C . Then it uses two LSTM cells combined with FLOW layers to integrate the context representations, followed by the context-question attention computation. Finally, FLOWQA performs self-attention (Yu et al., 2018) on the context and predict the answer span. Modeling FLOW is shown effective to improve the performance for conversational MC.

3 Proposed Approaches

This paper extends the concept of FLOW and proposes a flow-based approach, FLOWDELTA, to *explicitly* model information gain in flow during dialogues illustrated in Figure 2. The proposed mechanism is flexible to integrate with different models, including FlowQA and others. To examine such

flexibility and generalization capability, we further apply FLOW and FLOWDELTA to BERT (Devlin et al., 2018), a pretrained language understanding model that shows strong performance in MC tasks, to allow model to grasp dialogue history.

3.1 FlowDeltaQA

In the original FLOW operation in (1), the k -th step computation of GRU is $h_{k,j} = GRU(c_{k,j}, h_{k-1,j})$. We assume that the difference of previous hidden representations $h_{k-1,j}$ and $h_{k-2,j}$ indicates whether the flow change is important, which can be viewed as the information gain through the reasoning process. For example, 3 consecutive questions Q_{k-2}, Q_{k-1}, Q_k . Q_{k-1} and Q_k all discuss the same event described in the span $\{c_j, c_{j+1}, \dots, c_l\}$ of the context, while Q_{k-2} is about another topic. We expect the hidden state $\{h_{k-1,j}, h_{k-1,j+1}, \dots, h_{k-1,l}\}$ of the span in turn $k-1$ is dissimilar to the hidden state in the turn $k-2$, because their topics are different. By explicitly modeling such difference, our model more easily relates the current reasoning process to the corresponding context.

Following the intuition above, we propose FLOWDELTA by modifying the single step computation of FLOW into:

$$h_{k,j} = GRU([c_{k,j}; h_{k-1,j} - h_{k-2,j}], h_{k-1,j}), \quad (2)$$

where $[x; y]$ is the concatenation of the vectors x and y . We also investigate other variants such as Hadamard product ($h_{k-1,j} * h_{k-2,j}$) detailed in Appendix C.

3.2 BERT-FlowDelta

BERT (Devlin et al., 2018) with fine-tuning recently has reached the state-of-the-art in many single-turn MC tasks, such as SQuAD (Rajpurkar et al., 2016, 2018). However, how to extend BERT to the multi-turn setting remains unsolved. We propose to incorporate the FLOWDELTA mechanism to deal with the multi-turn problem, where the FLOW layer automatically integrates multi-turn information instead of tuning the number of QA pairs for inclusion.

Each layer of BERT is a Transformer block (Vaswani et al., 2017) that consists of multi-head attention (MH) and fully-connected feed forward network (FFN):

$$h_{l+1} = \text{Transformer}(h_l) = \text{LN}(h_l + \text{SA}(h_l)), \\ \text{SA}(h) = \text{FFN}(\text{LN}(h + \text{MH}(h))),$$

where h_l is the hidden representation of the l -th layer, LN is layer normalization (Ba et al., 2016) and SA means self-attention. To utilize L layers from BERT for the extractive question answering task, we feed the hidden representation from last layer h_L to a fully-connected layer (NN) to predict the answer span, written as $P^S, P^E = NN(h_L)$, where P^S and P^E are span start and span end probability for each word respectively.

BERT-FlowDelta incorporates the proposed FLOWDELTA mechanisms for two parts shown in the bottom right corner of Figure 2. First, we add FLOWDELTA layer before the final prediction layer, $P^S, P^E = NN([h_L; \text{FlowDelta}(h_L)])$. Second, we further insert FLOWDELTA into the last BERT layer, considering that modeling dialogue history *within* BERT may be beneficial.

$$h_L = \text{LN}(h_{L-1} + \text{SA}(h_{L-1}) + \text{FlowDelta}(h_{L-1}))$$

These two modifications are called exFlowDelta and inFlowDelta respectively, and the latter also meets the idea from Stickland and Murray who added additional parameters into BERT layers to improve the performance of multi-task learning. In our experiments, we only modify the last BERT layer to avoid largely increasing model size.

4 Experiments

To evaluate the effectiveness of the proposed FLOWDELTA, various tasks that contains dialogue history for understanding are performed in the following experiments.

4.1 Setup

Our models are tested on two conversational MC datasets, CoQA (Reddy et al., 2018) and QuAC (Choi et al., 2018), and a sequential instruction understanding dataset, SCONE (Long et al., 2016). For QuAC, we also report the Human Equivalence Score (HEQ). HEQ-Q and HEQ-D represent the percentage of exceeding the model performance over the human evaluation for each question and dialogue respectively. While CoQA and QuAC both follow the conversational QA setting, SCONE is the task requiring model to understand a sequence of natural language instructions and modify the word state accordingly. We follow Huang et al. (2018) to reduce instruction understanding to machine comprehension. Appendix A contains the example and reduction detail of SCONE for reference.

| Model | CoQA | | | | | | | | QuAC | | | |
|--------------------------|-------------|--|-------------|-------------|-------------|-------------|-------------|--|------|------|--|--|
| | Dev | Test | | | | | | | Dev | Test | | |
| | F1 | Child/Liter/Mid/News/Wiki/Reddit/Sci | F1 | F1 | F1 | HEQ-Q | HEQ-D | | | | | |
| BiDAF++ (N-ctx) | 69.2 | 66.5 65.7 70.2 71.6 72.6 60.8 67.1 | 67.8 | 60.6 | 60.1 | 54.8 | 4.0 | | | | | |
| FlowQA | 76.7 | 73.7 71.6 76.8 79.0 80.2 67.8 76.1 | 75.0 | 63.9 | 64.1 | 59.6 | 5.8 | | | | | |
| SDNet (Zhu et al., 2018) | 78.0 | 75.4 73.9 77.1 80.3 83.1 69.8 76.8 | 76.6 | - | - | - | - | | | | | |
| ConvBERT (unpublished) | - | - | 86.8 | - | 68.0 | 63.5 | 9.1 | | | | | |
| FlowDeltaQA | 77.6 | - | - | 64.8 | - | - | - | | | | | |
| BERT-FlowDelta | 79.4 | 75.9 75.6 80.1 82.1 82.3 69.8 78.8 | 77.7 | 68.6 | 67.8 | 63.6 | 12.1 | | | | | |
| Human | 89.8 | 90.2 88.4 89.8 88.6 89.9 86.7 88.1 | 88.8 | 80.8 | 81.1 | 100 | 100 | | | | | |

Table 1: Conversational QA results on CoQA and QuAC, where (N-ctx) refers to using previous N QA pairs (%).

| Model | CoQA F1 | QuAC F1 |
|----------------|-------------|-------------|
| BERT-FlowDelta | 79.4 | 68.6 |
| - inFlowDelta | 79.0 | 66.2 |
| - exFlowDelta | 78.0 | 64.5 |
| BERT-Flow | 79.2 | 66.8 |

Table 2: The ablation study of BERT-FlowDelta (%).

| Model | Scene | Tangrams | Alchemy |
|-----------------------|-------------|-------------|-------------|
| Long et al. (2016) | 14.7 | 27.6 | 52.3 |
| Guu et al. (2017) | 46.2 | 37.1 | 52.9 |
| Suhr and Artzi (2018) | 66.4 | 60.1 | 62.3 |
| Fried et al. (2017) | 72.7 | 69.6 | 72.0 |
| FusionNet | 58.2 | 67.9 | 74.1 |
| FlowQA | 74.5 | 72.3 | 76.4 |
| FlowDeltaQA | 75.1 | 72.5 | 76.1 |

Table 3: Dialogue accuracy for SCONE test (%).

4.2 Results

Table 1 reports model performance on CoQA and QuAC. It can be found that FlowDeltaQA yields substantial improvement over FlowQA on both datasets (+ 0.9 % F1 on both CoQA and QuAC), showing the usefulness of explicitly modeling the information gain in the FLOW layer. Furthermore, BERT-FlowDelta outperforms the published models on CoQA and achieves the state-of-the-art scores on the QuAC leaderboard on Apr 24, 2019. Specifically, BERT-FlowDelta outperforms ConvBERT by a large margin in HEQ-D on QuAC, showing the superiority of our model in modeling whole dialogue. Note that FLOWDELTA actually introduced few additional parameters compared to FLOW, since it only augments the input dimension of GRU. The consistent improvement from both data demonstrates the generalization capability of applying the proposed mechanism to various models.

Table 2 shows the ablation study of BERT-FlowDelta, where two proposed modules are both important for achieving such results. It is interesting that the proposed inFlowDelta and exFlowDelta boost the performance more on QuAC. As Yatskar (2018) mentioned, the topics in a dialogue shift more frequently on QuAC than

on CoQA, and we can see vanilla BERT also performs well on CoQA in the ablation of FLOW which provides long term dialog history information. Therefore, we can conclude that while FLOWDELTA improves the ability to grasp information gain in the dialog, it brings less performance improvement in the setting we do not need much contexts to answer the question.

Table 3 shows the performance of our FlowDeltaQA on the SCONE². Our model outperforms FlowQA and achieves the state-of-the-art in SCENE and TANGRAMS domains. The small performance drop in ALCHEMY aligns well with the statement in the ablation study. Because experiments show that removing FLOW affects performance in ALCHEMY less when comparing between FlowQA and FusionNet (Huang et al., 2017) (same models except FLOW), we claim that the previous dialogue history is less important in this domain. Thus replaying FLOW with FlowDelta does not bring any improvement in the ALCHEMY domain. The detailed qualitative study can be found in Appendix D.

5 Conclusion

This paper presents a simple and effective extension of FLOW named FLOWDELTA, which is capable of explicitly modeling the dialogue history in reasoning for better conversational machine comprehension. The proposed FlowDelta is flexible to apply to other machine comprehension models including FlowQA and BERT. The experiments on three datasets show that the proposed mechanism can model the information flow in the multi-turn dialogues more comprehensively, and further boosts the performance consistently. In the future, we will investigate more efficient ways to model the dialogue flow for conversational tasks.

²The results of BERT-FlowDelta are not shown, since SCONE is a relatively small and synthetic dataset.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Daniel Fried, Jacob Andreas, and Dan Klein. 2017. Unified pragmatic models for generating and following instructions. *arXiv preprint arXiv:1711.04987*.
- Kelvin Guu, Panupong Pasupat, Evan Zheran Liu, and Percy Liang. 2017. From language to programs: Bridging reinforcement learning and maximum marginal likelihood. *arXiv preprint arXiv:1704.07926*.
- Hsin-Yuan Huang, Eunsol Choi, and Wen-tau Yih. 2018. Flowqa: Grasping flow in history for conversational machine comprehension. *arXiv preprint arXiv:1810.06683*.
- Hsin-Yuan Huang, Chenguang Zhu, Yelong Shen, and Weizhu Chen. 2017. Fusionnet: Fusing via fully-aware attention with application to machine comprehension. *arXiv preprint arXiv:1711.07341*.
- Reginald Long, Panupong Pasupat, and Percy Liang. 2016. Simpler context-dependent logical forms via model projections. *arXiv preprint arXiv:1606.05378*.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*.
- 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. 2018. Coqa: A conversational question answering challenge. *arXiv preprint arXiv:1808.07042*.
- Asa Cooper Stickland and Iain Murray. 2019. Bert and pals: Projected attention layers for efficient adaptation in multi-task learning. *arXiv preprint arXiv:1902.02671*.
- Alane Suhr and Yoav Artzi. 2018. Situated mapping of sequential instructions to actions with single-step reward observation. *arXiv preprint arXiv:1805.10209*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Mark Yatskar. 2018. A qualitative comparison of coqa, squad 2.0 and quac. *arXiv preprint arXiv:1809.10735*.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.
- Chenguang Zhu, Michael Zeng, and Xuedong Huang. 2018. Sdnet: Contextualized attention-based deep network for conversational question answering. *arXiv preprint arXiv:1812.03593*.