

FCM: A Fine-grained Comparison Model for Multi-turn Dialogue Reasoning

Xu Wang^{1*}, Hainan Zhang^{2†}, Shuai Zhao^{1†},
Yanyan Zou², Hongshen Chen², Zhuoye Ding², Bo Cheng¹, Yanyan Lan³

¹Beijing University of Posts and Telecommunications, Beijing, China

²Data Science Lab, JD.com, China

³Institute for AI Industry Research, Tsinghua University, Beijing, China

wxx@bupt.edu.cn, zhanghainan1990@163.com, {zhaoshuaiby, chengbo}@bupt.edu.cn,
{zouyanyan6, dingzhuoye}@jd.com, ac@chenhongshen.com, lanyanyan@tsinghua.edu.cn

Abstract

Despite the success of neural dialogue systems in achieving high performance on the leaderboard, they cannot meet users' requirements in practice, due to their poor reasoning skills. The underlying reason is that most neural dialogue models only capture the syntactic and semantic information, but fail to model the logical consistency between the dialogue history and the generated response. Recently, a new multi-turn dialogue reasoning task has been proposed, to facilitate dialogue reasoning research. However, this task is challenging, because there are only slight differences between the illogical response and the dialogue history. How to effectively solve this challenge is still worth exploring. This paper proposes a Fine-grained Comparison Model (FCM) to tackle this problem. Inspired by human's behavior in reading comprehension, a comparison mechanism is proposed to focus on the fine-grained differences in the representation of each response candidate. Specifically, each candidate representation is compared with the whole history to obtain a history consistency representation. Furthermore, the consistency signals between each candidate and the speaker's own history are considered to drive a model to prefer a candidate that is logically consistent with the speaker's history logic. Finally, the above consistency representations are employed to output a ranking list of the candidate responses for multi-turn dialogue reasoning. Experimental results on two public dialogue datasets show that our method obtains higher ranking scores than the baseline models.

1 Introduction

Nowadays, the neural dialogue system has achieved high performance and been widely studied in both industry (Nuruzzaman and Hussain, 2018) and academia (Santhanam and Shaikh, 2019). However,

| | |
|------------|--|
| speaker A: | Excuse me. How much is this suit? |
| speaker B: | It's \$750 today. |
| speaker A: | Wow, that is pretty expensive! |
| speaker B: | The material is <i>imported from Italy</i> . If you buy a suit with same material, it may be \$2000. |
| speaker A: | Uh-hah. But I <i>saw a suit just like this one</i> , and it was \$600. I still thought it was expensive. |
| Candidates | 1. <i>No suit has the style as it</i> . It's the style that makes it special. |
| speaker B: | 2. The material of this suit is <i>imported from France</i> . It makes the suit special. |
| | 3. But <i>the color of our suit is very special</i> . |
| | 4. <i>Although the suit you saw is same as it, the material of our suit is imported from Italy</i> . |

Table 1: An example of dialogue reasoning. The logical contradictions are labeled in the same color. The most proper and logically correct response is option 4, labeled in red.

the selected response often contradicts with the dialogue history, such as “*I am a teacher*” as context but “*I work in the factory*” in the response, which greatly affects the user experience. The underlying reason is that existing neural dialogue systems only model the syntactic and semantic relevance but fail to capture the logical consistency between the dialogue history and the generated response.

Recently, a new multi-turn dialogue reasoning task (Cui et al., 2020) has been proposed to facilitate conversation reasoning research. The goal of the dialogue reasoning task is to select the logical response from the extremely similar candidate responses. However, this task is challenging, because there are only slight differences between the illogical response and the dialogue history. For example in Table 1, option 1 and option 3 are in conflict with speaker A who has seen the same suit, and option 2 is in conflict with the imported country. Since the candidate options are only slightly different to the context, traditional dialogue models might tend to select semantically relevant yet illogical candidates, yielding incorrect responses.

As we all know, humans have the ability to make effective and efficient reasoning, because they usu-

*Work done at Data Science Lab, JD.com.

†Corresponding Authors.

ally focus on perceiving fine-grained details and compare the candidates at multiple-granularity levels accordingly. Taking Table 1 as an example, by comparing option 2 and option 4, human can identify that the key difference is “*Italy*” and “*France*”, which can be used as key information to distinguish such two options. Inspired by the human behaviors in reasoning, the fine-grained comparison between response and history should be introduced to improve the reasoning ability for the dialogue reasoning model.

Furthermore, the dialogue history from the same speaker is also critical for modeling the logical consistency (Chen et al., 2017). It is natural that a person holds his logical consistency when speaking. Thus the logical errors for the same speaker might hamper the user experience. For example in Table 1, speaker B says the suit’s material is imported from Italy in the dialogue history, but option 2 for him says the material is imported from France, which is a more obvious and serious logical error. Therefore, a reasoning model needs to consider the speaker’s own historical consistency to distinguish logically incorrect candidate responses.

Inspired by the above analysis, we propose a Fine-grained Comparison Model (FCM) to improve the performance of multi-turn dialogue reasoning. To be specific, we firstly propose a comparison mechanism to compare every candidate response with all other ones. Secondly, we compare each candidate representation with the whole history and the speaker’s own history to obtain the history and the speaker’s consistency representations, respectively. Finally, we utilize the above consistency representations to output a ranking list of the candidate responses for multi-turn dialogue reasoning.

In our experiments, we utilize two public multi-turn dialogue datasets, named MuTual (Cui et al., 2020) and Ubuntu (Lowe et al., 2015), to evaluate our proposed models. The results show that FCM has the ability to rank the candidate responses more accurately than the baseline models. We also conduct some case studies to demonstrate the superiority and soundness of FCM.

The main contributions of this paper include:

- We introduce the response comparison mechanism to enable the dialogue model (e.g., BERT(Devlin et al., 2019)) to have fine-grained detail perception ability, which tackles the difficulty of subtle differences between

candidates and dialogue history in dialogue reasoning.

- We model the speaker’s own logical consistency to further enhance the reasoning ability for dialogue reasoning task.
- We experiment on two public multi-turn dialogue datasets to demonstrate the effectiveness of our proposed model FCM.

2 Related Work

Recently, multi-turn dialogue has gained more attention in both industry (Wu et al., 2020; Zhan et al., 2021) and academia (Cho and May, 2020), compared with single-turn dialogue (Mou et al., 2016; Zhang et al., 2018a; Li et al., 2017a). Serban et al. (2016) proposes a hierarchical recurrent encoder-decoder (HRED) model which uses the hierarchical encoder-decoder framework to model the relevance of the context and response. Wu et al. (2017) uses HRED to model relationships among utterances to enhance the performance of the retrieval-based chatbot. Chen et al. (2018) adds the hierarchical structure and the variable memory network into a neural encoder-decoder network, which can capture both the high-level abstract variations and long-term memories during dialogue tracking. Zhang et al. (2018b) adopts dynamic and static attention to weigh the importance of each utterance in a conversation and then obtain the contextual representation. Zhang et al. (2019) utilizes hierarchical self-attention mechanism to solve the position bias problem of dialogue models.

Although these models have achieved good performance on datasets like DailyDialog (Li et al., 2017b) and ECD (Zhang et al., 2018c), there is still a giant gap between high performance on the leader-board and poor practical user experience (Cui et al., 2020). These models frequently generate responses that are logically incorrect. One possible reason is that the previous models (Young et al., 2018; Clark et al., 2019) only solve the cases through linguistic information matching yet lack of logical reasoning.

Obviously, a dialogue reasoning dataset that can help the model to detect the illogical response is extremely necessary. Recently, an open domain multi-turn dialogue reasoning dataset (MuTual) (Cui et al., 2020) is proposed to facilitate the reasoning capabilities of conversation models. In particular, given a context, it prepares four candidate responses and all of them are relevant to the con-

text, but singly one has the correct logic. It requires fine-grained reasoning ability between context and response to make the correct choice. Current multi-turn dialogue models (Devlin et al., 2019; Zhang et al., 2018c; Liu et al., 2019; Wu et al., 2017), which perform well on existing benchmarks, have declined on this dataset, which proves that the reasoning ability of these models is still insufficient. Therefore, how to enhance the logical reasoning ability of dialogue models is worth discussing.

Logic consistency of history is essential in human-like behavior. For the multi-choice reasoning task (Zhu et al., 2018), there are only slight differences between the illogical response and the dialogue history, so a fine-grained comparison mechanism is required to infer the logic between each candidate response and the dialogue history. Previous models (Young et al., 2018; Devlin et al., 2019) ignore the modeling of fine-grained comparison, which leads to the loss of reasoning ability of their models. In this paper, we overcome this challenge and propose a fine-grained comparison mechanism to perform history consistency and speaker consistency respectively.

3 Model

In this section, we first describe the task definition and then introduce our FCM model in detail, with the architecture shown in Figure 1. Our FCM model consists of three components, i.e., Contextual Encoding, Fine-grained Comparison Module, and Response Prediction. Firstly, we utilize the pre-trained language model BERT to encode each token of context and response into a fixed-length vector, which carries contextual information. Secondly, we utilize our comparison mechanism to obtain the fine-grained response representation and then compare such obtained representation with both the whole history and the speaker’s own history. Finally, we fuse consistency representation and semantic information, and then use a linear layer to obtain the candidate response score for the multi-choice prediction process.

3.1 Task Definition

Given a dialogue context $U = \{u_1, u_2, \dots, u_N\}$ and a candidate response set $R = \{r_1, r_2, \dots, r_M\}$, where $u_i = \{w_{i,1}^u, w_{i,2}^u, \dots, w_{i,l_i}^u\}$ is an utterance with l_i tokens, and $r_i = \{w_{i,1}^r, w_{i,2}^r, \dots, w_{i,j_i}^r\}$ is a candidate response with j_i tokens, the goal of this task is to select the proper and logical response

based on the conditional probability distribution, i.e., $P(r_i|U, R)$, where $r_i \in R$.

3.2 Contextual Encoding

Given each input pair (U, r_i) , we concatenate the context and each candidate and then feed them into the pre-trained BERT to obtain the fixed-length vector of each token in the context and response, which is denoted as:

$$[H^U; H^{r_i}] = BERT(\langle U; r_i \rangle), \quad (1)$$

where $BERT(\cdot)$ returns the last layer output of the encoder. $\langle; \rangle$ means concatenation of two sequences. $H^U \in \mathbb{R}^{|U| \times d}$ and $H^{r_i} \in \mathbb{R}^{|r_i| \times d}$ are the token-level vectors of context U and candidate r_i , respectively. d is the dimension of the hidden state. Besides, we obtain the summary vector $h_i^{[cls]} \in \mathbb{R}^d$ for input pair (U, r_i) , which carries the semantic information of the whole input (Devlin et al., 2019).

3.3 Fine-grained Comparison Module

In this section, we introduce our fine-grained comparison module, including three steps: fine-grained response comparison, consistency reasoning with history, and enhancing speaker consistency. Specifically, a fine-grained response comparison mechanism is firstly utilized to imitate human behaviors, which aims to compare the correlation and difference between candidate responses at multi-granularity levels. Then, the history-aware bidirectional matching method is utilized to infer the fine-grained logical consistency between the candidate r_i and the whole history U . Finally, another bidirectional matching model is used to infer the speaker consistency between the response r_i and the speaker’s own history.

3.3.1 Fine-grained Response Comparison

For each response r_i , a fine-grained attention mechanism is used to compare it with all other responses to get the fine-grained comparison information.

Paired Correlation Given the hidden vectors H^{r_i} and H^{r_j} , we calculate the word-level attention between them to obtain the similarity matrix $A^{r_i,j}$, which is defined as:

$$A^{r_i,j} = \left[\frac{\exp(a_{mn}^{r_i,j})}{\sum_n \exp(a_{mn}^{r_i,j})} \right]_{m,n}, \quad (2)$$

$$a_{mn}^{r_i,j} = W_1^T [H_m^{r_i}; H_n^{r_j}; H_m^{r_i} \odot H_n^{r_j}],$$

where \odot is the element multiplication between two matrices. $H_m^{r_i}$ is the hidden representation of m th

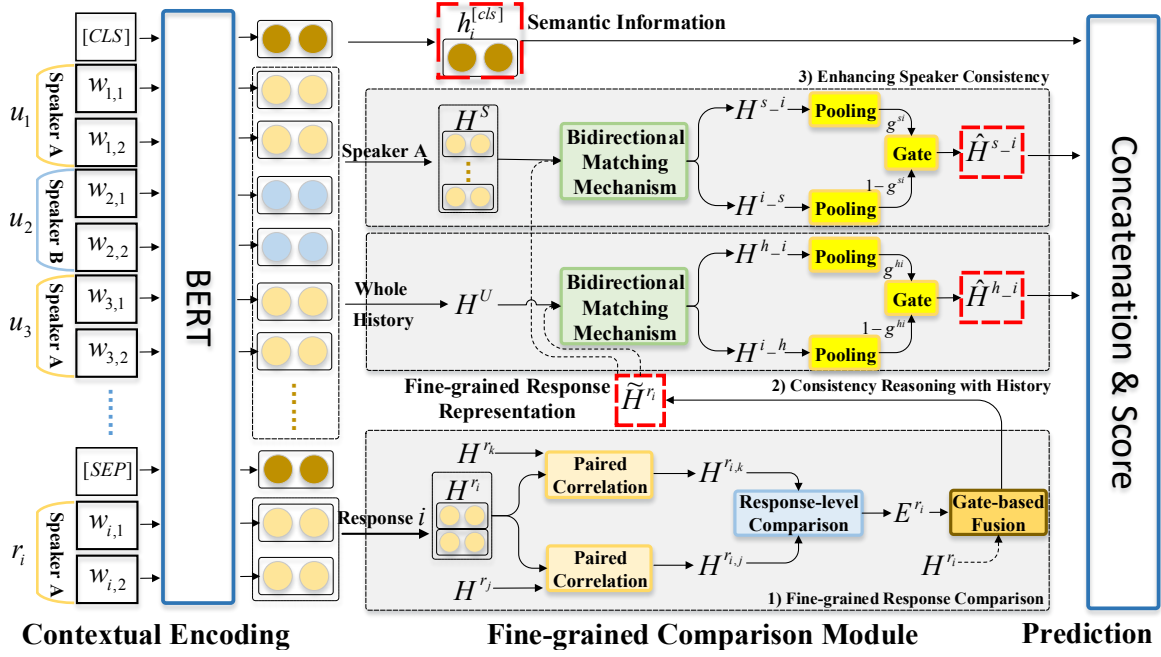


Figure 1: The proposed FCM network for multi-turn dialogue reasoning task. It contains three components: 1) Contextual Encoding, 2) Fine-grained Comparison Module, and 3) Response Prediction.

token in r_i , and $W_1^T \in \mathbb{R}^{1 \times 3d}$ is a learned parameter. $a_{mn}^{r_i, j}$ means the similarity between the m th word of r_i and the n th word of r_j .

Given the similarity matrix $A^{r_i, j}$, the paired correlation information $H^{r_i, j}$ is defined as:

$$\begin{aligned} H^{r_i, j} &= [H^{r_i} - \overline{H}^{r_i, j}; H^{r_i} \odot \overline{H}^{r_i, j}], \\ \overline{H}^{r_i, j} &= A^{r_i, j} H^{r_j}, \end{aligned} \quad (3)$$

where $\overline{H}^{r_i, j}$ highlights the similar part of r_j with r_i , and $H^{r_i, j}$ represents the different part between the candidate r_i and r_j .

Response-level Comparison Given the correlation information $H^{r_i, j}$ for r_i , the response-level compared information E^{r_i} is defined as:

$$E^{r_i} = \tanh\left(\left[\{H^{r_i, j}\}_{j \neq i}\right] W_2 + b_2\right), \quad (4)$$

where $W_2 \in \mathbb{R}^{2d(M-1) \times d}$ and $b_2 \in \mathbb{R}^d$ are learned parameters. M refers to the total number of candidate responses.

Gate-based Fusion Given the contextual encoding features H^{r_i} and the response-level compared information E^{r_i} of response r_i , an element-wise gating mechanism is utilized to obtain the fine-grained response representation \tilde{H}^{r_i} , which is defined as:

$$\begin{aligned} g^{r_i} &= \sigma([E^{r_i}; H^{r_i}] W_3 + b_3), \\ \tilde{H}^{r_i} &= g^{r_i} \odot E^{r_i} + (1 - g^{r_i}) \odot H^{r_i}, \end{aligned} \quad (5)$$

where $W_3 \in \mathbb{R}^{2d \times d}$ and $b_3 \in \mathbb{R}^d$ are learned parameters. $g^{r_i} \in \mathbb{R}^{|r_i| \times d}$ stands for the element-wise gate value.

3.3.2 Consistency Reasoning with History

Given the context representation $H^U \in \mathbb{R}^{|U| \times d}$ and the fine-grained response representation $\tilde{H}^{r_i} \in \mathbb{R}^{|r_i| \times d}$, we design a bidirectional matching mechanism to obtain the response-aware history representation H^{h-i} and history-aware response representation H^{i-h} respectively, which is defined as:

$$\begin{aligned} A^{h-i} &= \text{SoftMax}(H^U W_4 \tilde{H}^{r_i T}), \\ A^{i-h} &= \text{SoftMax}(\tilde{H}^{r_i} W_5 H^{U T}), \\ H^{h-i} &= \text{Relu}(A^{h-i} \tilde{H}^{r_i} W_6), \\ H^{i-h} &= \text{Relu}(A^{i-h} H^U W_7), \end{aligned} \quad (6)$$

where W_4, W_5, W_6 , and $W_7 \in \mathbb{R}^{d \times d}$ are learned parameters. A^{h-i} and A^{i-h} are the word-level attention matrices between the whole history U and response r_i , which focus on different perspectives.

Given the two representations H^{i-h} and H^{h-i} as input, we utilize a gate mechanism to fuse them to get the history-aware consistency information

\hat{H}^{h-i} , which is defined as:

$$\begin{aligned} E^{h-i} &= \text{MaxPooling}(H^{h-i}), \\ E^{i-h} &= \text{MaxPooling}(H^{i-h}), \\ g^{hi} &= \sigma(E^{h-i}W_8 + E^{i-h}W_9 + b_4), \\ \hat{H}^{h-i} &= g^{hi} \odot E^{h-i} + (1 - g^{hi}) \odot E^{i-h}, \end{aligned} \quad (7)$$

where $W_8, W_9 \in \mathbb{R}^{d \times d}$ and $b_4 \in \mathbb{R}^d$ are learned parameters. *MaxPooling* means row-wise max pooling operation. g^{hi} means the element-wise gate value.

3.3.3 Enhancing Speaker Consistency

Given the speaker’s own context H^S and the compared response representation \tilde{H}^{r_i} of r_i from section 3.3.1, we utilize another bidirectional matching module to obtain the response-aware speaker representation H^{s-i} and speaker-aware response representation H^{i-s} , respectively, which is with the similar definition in section 3.3.2:

$$\begin{aligned} A^{s-i} &= \text{SoftMax}(H^S W_{10} \tilde{H}^{r_i T}), \\ A^{i-s} &= \text{SoftMax}(\tilde{H}^{r_i} W_{11} H^{S T}), \\ H^{s-i} &= \text{Relu}(A^{s-i} \tilde{H}^{r_i} W_{12}), \\ H^{i-s} &= \text{Relu}(A^{i-s} H^S W_{13}), \end{aligned} \quad (8)$$

where A^{s-i} and A^{i-s} are the word-level attention matrices between the speaker’s own context S and response r_i .

Given the two representations H^{i-s} and H^{s-i} as input, the speaker-aware consistency information \hat{H}^{s-i} is defined as:

$$\begin{aligned} E^{s-i} &= \text{MaxPooling}(H^{s-i}), \\ E^{i-s} &= \text{MaxPooling}(H^{i-s}), \\ g^{si} &= \sigma(E^{s-i}W_{14} + E^{i-s}W_{15} + b_5), \\ \hat{H}^{s-i} &= g^{si} \odot E^{s-i} + (1 - g^{si}) \odot E^{i-s}. \end{aligned} \quad (9)$$

3.4 Response Prediction

Given the semantic information $h_i^{[cls]}$, history-aware consistency \hat{H}^{h-i} and speaker-aware consistency \hat{H}^{s-i} , we concatenate them to get the reasoning information H^i for response r_i , which is defined as:

$$H^i = [h_i^{[cls]}, \hat{H}^{h-i}, \hat{H}^{s-i}]. \quad (10)$$

The score $P(r_i|U, R)$ of the i th candidate response is computed as follows:

$$P(r_i|U, R) = \frac{\exp(W_{16}H^i + b_6)}{\sum_{i=0}^M \exp(W_{16}H^i + b_6)}, \quad (11)$$

where $W_{16} \in \mathbb{R}^{1 \times 3d}$ and $b_6 \in \mathbb{R}^1$ are learned parameters.

The loss function is defined as:

$$J(\theta) = -\frac{1}{N} \sum \log P(\hat{r}_i|U, R) + \lambda \|\theta\|_2^2, \quad (12)$$

where λ is a hyperparameter, θ are all trainable parameters, N is the size of training examples in the dataset, and \hat{r}_i is the ground-truth response.

4 Experiments

In this section, we conduct experiments on MuTual reasoning and Ubuntu dialogue datasets to evaluate our proposed method.

4.1 Experimental Settings

We first introduce some empirical settings, including datasets, baseline methods, parameter settings, and evaluation measures.

4.1.1 Datasets

We test our model on two public multi-turn dialogue datasets: MuTual and Ubuntu.

MuTual (Cui et al., 2020) consists of 8860 manually annotated dialogues, which is based on Chinese student English listening comprehension exams. Each dialogue has two speakers speaking in turn and contains four candidate responses. The goal of this task is to select the correct and logical response according to the historical contexts. The training, validation, and testing sets contain 7088, 886, and 886 pairs, respectively.

Ubuntu (Lowe et al., 2015) consists of English multi-turn dialogues about technical support collected from chat logs on Ubuntu forum. The dataset contains 1 million context-response training pairs, 0.5 million validation pairs, and 0.5 million testing pairs. Each pair has one positive response and nine negative responses. Because there are some sessions in Ubuntu dataset that require reasoning, we utilize Ubuntu dataset to verify the reasoning ability of FCM.

4.1.2 Baseline Methods

We use 11 baselines for comparison, including the traditional TF-IDF (Paik, 2013), Dual LSTM (Lowe et al., 2015), SMN (Wu et al., 2017), DAM (Zhou et al., 2018), BERT (Devlin et al., 2019), BERT-MC (Cui et al., 2020), GPT-2 (Radford et al., 2019), Deep Utterance Aggregation (DUA) (Zhang et al., 2018c), Interaction-over-Interaction (IoI) (Tao et al., 2019b), Multi-hop Selector (MSN)

| Model | Dev | | | Test | | |
|-------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | $R_4@1$ | $R_4@2$ | MRR | $R_4@1$ | $R_4@2$ | MRR |
| TF-IDF (Paik, 2013) | 0.276 | 0.541 | 0.541 | 0.279 | 0.536 | 0.542 |
| Dual LSTM (Lowe et al., 2015) | 0.266 | 0.528 | 0.538 | 0.260 | 0.491 | 0.743 |
| SMN (Wu et al., 2017) | 0.274 | 0.524 | 0.575 | 0.299 | 0.585 | 0.595 |
| DAM (Zhou et al., 2018) | 0.239 | 0.463 | 0.575 | 0.241 | 0.465 | 0.518 |
| GPT-2(Radford et al., 2019) | 0.335 | 0.595 | 0.586 | 0.332 | 0.602 | 0.584 |
| BERT (Devlin et al., 2019) | 0.657 | 0.867 | 0.803 | 0.648 | 0.847 | 0.795 |
| BERT-MC (Devlin et al., 2019) | 0.661 | 0.871 | 0.806 | 0.667 | 0.878 | 0.810 |
| FCM | 0.696 | 0.884 | 0.824 | 0.692 | 0.884 | 0.823 |

Table 2: The metric-based evaluation on MuTual dataset.

(Yuan et al., 2019) and Multi-Representation Fusion (MRFN) (Tao et al., 2019a).

4.1.3 Parameter Settings

We utilize the open-source pre-trained model BERT_{base}* for the dialogue reasoning task. BERT_{base} has 12-layer transformer blocks, 768 hidden-size, and 12 self-attention heads. It totally contains 110M parameters. In order to make a fair comparison between our model and baselines, 1) for MuTual dataset, we refer to (Cui et al., 2020) and set the max input sequence length to 350. We set the dropout rate to 0.2. The L_2 weight decays λ is set to 0.01. We employ Adam (Kingma and Ba, 2015) to optimize the model with a learning rate $1e-6$. We run the experiments on two TITAN XP GPUs with 12G memory and train for 10 epochs with batch size of 4; 2) for Ubuntu dataset, we use the same evaluation metrics which are used in previous works (Gu et al., 2020; Zhang et al., 2018c).

4.1.4 Evaluation Measures

We consider the dialogue reasoning task as a retrieval-based response selection task and apply traditional information retrieval metrics. On MuTual, we display the recall and Mean Reciprocal Rank measures (Voorhees and Tice, 2000), i.e., $R_4@1$, $R_4@2$, and MRR. On Ubuntu, we use $R_{10}@1$, $R_{10}@2$, and $R_{10}@5$ for evaluation.

4.2 Experimental Results

Now we demonstrate our experimental results on the two public datasets.

4.2.1 Metric-based Evaluation

The metric-based evaluation results on MuTual and Ubuntu are shown in Table 2 and Table 3. From

*<https://github.com/huggingface/transformers>

| Model | $R_{10}@1$ | $R_{10}@2$ | $R_{10}@5$ |
|----------------------------|--------------|--------------|--------------|
| SMN (Wu et al., 2017) | 0.726 | 0.847 | 0.961 |
| DUA (Zhang et al., 2018c) | 0.752 | 0.868 | 0.962 |
| DAM (Zhou et al., 2018) | 0.767 | 0.874 | 0.969 |
| IoI (Tao et al., 2019b) | 0.796 | 0.894 | 0.974 |
| MSN (Yuan et al., 2019) | 0.800 | 0.899 | 0.978 |
| MRFN (Tao et al., 2019a) | 0.786 | 0.886 | 0.976 |
| BERT (Devlin et al., 2019) | 0.808 | 0.897 | 0.975 |
| FCM | 0.816 | 0.908 | 0.983 |

Table 3: The metric-based evaluation on Ubuntu.

the results, we can see that the performance of well-designed RNN-based networks, such as Dual LSTM and SMN, is relatively poor, which demonstrates that such traditional models cannot deal with the dialogue reasoning task. Although the performance of BERT is better than other baseline models, merely using self-attention between context and candidate responses still misses fine-grained consistency information. With the introduction of fine-grained comparison information, our FCM model outperforms all baseline models. Take the $R_4@1$ and $R_4@2$ on the MuTual dev set as an example, the $R_4@1$ and MRR of our FCM model are 69.6% and 82.4%, respectively, which is significantly better than that of BERT-MC, i.e., 3.5% and 1.8%. On Ubuntu, we find that FCM also outperforms the BERT model on three metrics. In summary, our FCM model has the ability to select a more logically consistent response than baselines.

4.2.2 Case Study

To facilitate a better understanding of our model, we present the examples on MuTual in Table 6. From the results, we can see that our model is more accurate in multi-choice prediction than the traditional baselines and transformer models. In this example, both response 1 and response 4 state that “*Sichuan food*” is speaker B’s favorite food.

| MuTual Dev | | | |
|------------------------------|--------------|--------------|--------------|
| Model | $R_4@1$ | $R_4@2$ | MRR |
| FCM | 0.696 | 0.884 | 0.824 |
| w/o Response Comparison | 0.680 | 0.875 | 0.815 |
| w/o Consistency with History | 0.690 | 0.878 | 0.820 |
| w/o Speaker Consistency | 0.685 | 0.880 | 0.816 |
| Ubuntu | | | |
| Model | $R_{10}@1$ | $R_{10}@2$ | $R_{10}@5$ |
| FCM | 0.816 | 0.908 | 0.983 |
| w/o Response Comparison | 0.814 | 0.904 | 0.980 |
| w/o Consistency with History | 0.812 | 0.903 | 0.978 |
| w/o Speaker Consistency | 0.812 | 0.902 | 0.979 |

Table 4: Ablation experimental results of our FCM model on MuTual Dev and Ubuntu datasets.

BERT and BERT-MC agree with this statement. However, from speaker A’s historical utterance “*I know Guangdong food is your favorite kind of Chinese food.*”, we can guess that speaker B’s favorite is “*Guangdong food*”. According to the above analysis, we argue that the two baselines lack the modeling of fine-grained consistency reasoning, especially for the speaker’s own consistency, so they predict the wrong response, which proves the effectiveness of logical consistency.

Given the wrong response 2 and the correct response 3, the biggest difference between them is the utterance “*I know where it is.*”. Using the consistency reasoning mechanism, we can infer from the speaker A’s historical utterance “*I do not know where it is.*” that this statement is wrong. However, baselines do not use the fine-grained response comparison mechanism and can not focus on the fine-grained differences between response 2 and response 3, and then it predicts the wrong response 2, which illustrates the importance of fine-grained response comparison in logic consistency. In summary, compared with baseline models, our proposed model FCM, which carries the fine-grained comparison ability, is capable of inferring logic consistency more accurately for the multi-turn dialogue reasoning task.

4.3 Ablation Study

To study the contributions of the main components in FCM, we conduct ablation experiments, mainly including removing the **Fine-grained Response Comparison** module, **Consistency Reasoning with History** module, and **Enhancing Speaker Consistency** module from our proposed model, respectively.

The results on MuTual and Ubuntu are shown

| Model | $R_4@1$ | $R_4@2$ | MRR |
|----------------|---------------|---------------|---------------|
| FCM | 0.6964 | 0.8841 | 0.8236 |
| coarse-grained | 0.6817 | 0.8772 | 0.8161 |
| simple-add | 0.6839 | 0.8837 | 0.8192 |
| no-source | 0.6884 | 0.8818 | 0.8202 |
| no-gate | 0.6871 | 0.8795 | 0.8189 |

Table 5: Analysis of fine-grained response comparison.

in Table 4. We can see that without response comparison, the performance of the model drops in all three metrics. Taking the Mutual dev set as an example, the model decreased by 1.6%, 0.9%, and 0.9% on $R_4@1$, $R_4@2$, and MRR, respectively, which proves the importance of fine-grained response comparison in the reasoning process. When without whole history consistency reasoning, we can see that the performance is reduced, which demonstrates the necessity of modeling consistency between each candidate and the history. When without speaker consistency reasoning, we can see that the performance is also reduced. Taking the MuTual dev set as an example, the measures decreased by 1.1% and 0.8% on $R_4@1$ and MRR, respectively. This reduction proves that enhanced speaker consistency is helpful for the model to calculate the score of each candidate response.

4.4 Analysis of Fine-grained Response Comparison

To prove the effectiveness of our fine-grained response comparison module, we have conducted experiments on each operation of this module, mainly focusing on the operations of paired correlation calculation and gate-based fusion.

Analysis of paired correlation calculation To check the validity of these operations, we select a simple method to get the coarse-grained correlation, instead of Equation 2 and Equation 3, which is defined as follows:

$$\begin{aligned}
 A^{r_{i,j}} &= \text{SoftMax}(H^{r_i}(H^{r_j})^T), \\
 H^{r_{i,j}} &= A^{r_{i,j}} H^{r_j}.
 \end{aligned}
 \tag{13}$$

From the results in Table 5, with (coarse-grained), we observe that the performance of the FCM decreases in $R_4@1$, $R_4@2$, and MRR, which proves the effectiveness of our designed fine-grained method for paired correlation.

| The Example of MuTual Dataset | | | | |
|-------------------------------|----------------------|---|-------|------------|
| Context | speaker A | It’s already 30. How about preparing supper now? | | |
| | speaker B | But I don’t want to cook today. I’m tired of cooking every day. | | |
| | speaker A | How about having supper out tonight? There is a new Chinese restaurant on the third street. Tom went there yesterday, and he said it was great. | | |
| | speaker B | Really? What kinds of food does it have? You know, I don’t like food that’s too spicy. | | |
| | speaker A | Don’t worry. One of the chefs is from Guangdong. | | |
| | speaker B | <i>I know Guangdong food is your favorite kind of Chinese food.</i> | | |
| | speaker A | That’s great. Do you know how to get there? | | |
| | speaker B | <i>I do not know where it is. I just know it’s on the third street.</i> | | |
| | speaker A | Don’t worry. I’m sure we will find it. | | |
| | speaker B | But I don’t feel like walking now. It’s still so hot outside. | | |
| Candidates | speaker A | Then how about asking Tom to pick us up? We can treat him to supper. | | |
| | speaker B | That’s a good idea. | | |
| | 1) speaker A: | So since <i>Sichuan food is your favorite kind</i> of Chinese food, why don’t we go there after work? | | |
| | 2) speaker A: | Ok, dear, let’s go. <i>I know where is it.</i> | | |
| | 3) speaker A: | Ok, dear, let’s go! | | |
| | 4) speaker A: | Great. After school, we can go there to eat <i>your favorite Sichuan food.</i> | | |
| Model | BERT | BERT-MC | GPT-2 | FCM |
| Predictions | 1) | 4) | 2) | 3) |

Table 6: The selected candidate response from our FCM model and baselines on MuTual.

Analysis of gate-based fusion In order to prove the effectiveness of the gate-based fusion, we design the following experiments: 1) (no-source): At the calculation of gate value g^{r_i} in Equation 5, we remove the original response representation H^{r_i} and only utilize the updated response representation E^{r_i} ; 2) (simple-add): In Equation 5, we remove the operation of the weight-based summarization; 3) (no-gate): We directly remove the gating mechanism and use the output of Equation 4 to represent the response-level compared information.

The results are shown in Table 5, and we obtain the following conclusions: 1) With (no-source), the performance of our model decreases, which shows that the retention of the original information is necessary to calculate gate value g^{r_i} ; 2) With (simple-add), the updated response information and the original response information cannot be treated equally; 3) With (no-gate), we still get a lower performance, which means that after obtaining new knowledge, the response-level compared information may lose the original information.

4.5 Generality of FCM

We test the generality of FCM in the pre-trained language models. Specifically, we apply FCM to the widely used models: BERT_{base}, BERT_{large}, RoBERTa_{base}, RoBERTa_{large}, ELECTRA_{base} and ELECTRA_{large}, respectively. The experiment results are shown in Figure 2. From the results, we can discover that after applying our model to differ-

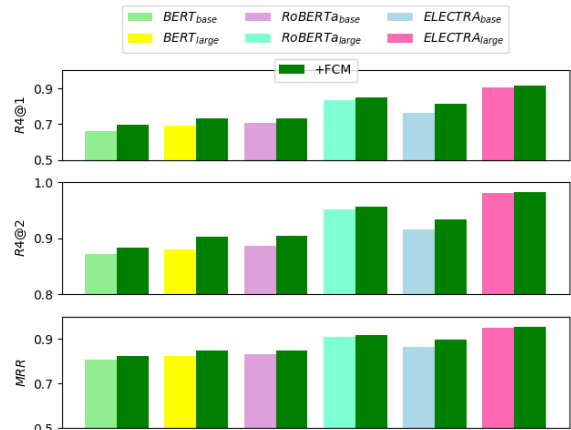


Figure 2: The $R_4@1$, $R_4@2$ and MRR performance of different pre-trained language models with FCM on MuTual.

ent pre-trained language models, their performance can all be improved, which proves that FCM is generally effective to the widely used pre-trained language models.

5 Conclusion

In this paper, we focus on multi-turn dialogue reasoning tasks and propose the FCM model. The motivation comes from the fact that the widely used dialogue models only focus on the syntactic and semantic relevance but fail to model the logical consistency between the dialogue history and the generated response. This task is challenging because there are only slight differences between the illogical response and the dialogue history. Our

core idea is to propose a fine-grained comparison mechanism to focus on the fine-grained differences in the representation of each response candidate, and then each candidate representation is compared with the history to obtain a consistency score. In the future, we plan to further investigate the fine-grained correlation between different speakers, and utilize this information to help improve our model.

Acknowledgements

This work is supported by Beijing Nova Program of Science and Technology (Grant No.Z191100001119031), National Key Research and Development Program of China (No. 2018YFB1003804), Guangxi Key Laboratory of Cryptography and Information Security (No.GCIS202111), the Open Program of Zhejiang Lab (Grant No.2019KE0AB03), the Beijing Academy of Artificial Intelligence (BAAI), and the National Natural Science Foundation of China (NSFC) (No.61773362). Xu Wang is supported by BUPT Excellent Ph.D. Students Foundation under grant CX2019137. We thank the anonymous reviewers for their constructive comments and suggestions.

References

- Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *SIGKDD*.
- Hongshen Chen, Zhaochun Ren, Jiliang Tang, Yihong Eric Zhao, and Dawei Yin. 2018. Hierarchical variational memory network for dialogue generation. *WWW*.
- Hyundong Cho and Jonathan May. 2020. Grounding conversations with improvised dialogues. In *ACL*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does BERT look at? an analysis of bert’s attention. *CoRR*.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. Mutual: A dataset for multi-turn dialogue reasoning. *ACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. *NAACL*.
- Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. Speaker-aware BERT for multi-turn response selection in retrieval-based chatbots. *CIKM*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. Adversarial learning for neural dialogue generation. *EMNLP*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. Dailydialog: A manually labelled multi-turn dialogue dataset. *IJCNLP*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. *SIGDIAL Conference*.
- Lili Mou, Yiping Song, Rui Yan, Ge Li, Lu Zhang, and Zhi Jin. 2016. Sequence to backward and forward sequences: A content-introducing approach to generative short-text conversation. *COLING*.
- M. Nuruzzaman and O. K. Hussain. 2018. A survey on chatbot implementation in customer service industry through deep neural networks. In *ICEBE*.
- Jiaul H. Paik. 2013. A novel TF-IDF weighting scheme for effective ranking. *SIGIR*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*.
- Sashank Santhanam and Samira Shaikh. 2019. A survey of natural language generation techniques with a focus on dialogue systems - past, present and future directions. *CoRR*.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019a. Multi-representation fusion network for multi-turn response selection in retrieval-based chatbots. *WSDM*.
- Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019b. One time of interaction may not be enough: Go deep with an interaction-over-interaction network for response selection in dialogues. In *ACL*.
- Ellen M. Voorhees and Dawn M. Tice. 2000. The TREC-8 question answering track. European Language Resources Association.

- Sixing Wu, Ying Li, Dawei Zhang, Yang Zhou, and Zhonghai Wu. 2020. Diverse and informative dialogue generation with context-specific common-sense knowledge awareness. *ACL*.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. *ACL*.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. Augmenting end-to-end dialogue systems with common-sense knowledge. *AAAI*.
- Chunyuan Yuan, Wei Zhou, Mingming Li, Shangwen Lv, Fuqing Zhu, Jizhong Han, and Songlin Hu. 2019. Multi-hop selector network for multi-turn response selection in retrieval-based chatbots. *EMNLP*.
- Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao, and Yanyan Lan. 2021. Augmenting knowledge-grounded conversations with sequential knowledge transition. *NAACL*.
- Hainan Zhang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. 2018a. Reinforcing coherence for sequence to sequence model in dialogue generation. *IJCAI*.
- Hainan Zhang, Yanyan Lan, Liang Pang, Jiafeng Guo, and Xueqi Cheng. 2019. Recosa: Detecting the relevant contexts with self-attention for multi-turn dialogue generation. *ACL*.
- Weinan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018b. Context-sensitive generation of open-domain conversational responses. *COLING*.
- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018c. Modeling multi-turn conversation with deep utterance aggregation. *COLING*.
- Xiangyang Zhou, Lu Li, Daxiang Dong, Yi Liu, Ying Chen, Wayne Xin Zhao, Dianhai Yu, and Hua Wu. 2018. Multi-turn response selection for chatbots with deep attention matching network. *ACL*.
- Haichao Zhu, Furu Wei, Bing Qin, and Ting Liu. 2018. Hierarchical attention flow for multiple-choice reading comprehension. In *AAAI*.