

Other Roles Matter! Enhancing Role-Oriented Dialogue Summarization via Role Interactions

Haitao Lin^{1,2}, Junnan Zhu^{1,2}, Lu Xiang^{1,2}, Yu Zhou^{1,3*},
Jiajun Zhang^{1,2}, Chengqing Zong^{1,2}

¹ National Laboratory of Pattern Recognition, Institute of Automation, CAS, Beijing, China

² School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³ Fanyu AI Laboratory, Zhongke Fanyu Technology Co., Ltd, Beijing, China
{haitao.lin, junnan.zhu, lu.xiang, yzhou, jjzhang, cqzong}@nlpr.ia.ac.cn,

Abstract

Role-oriented dialogue summarization is to generate summaries for different roles in the dialogue, *e.g.*, merchants and consumers. Existing methods handle this task by summarizing each role’s content separately, thus are prone to ignore the information from other roles. However, we believe that other roles’ content could benefit the quality of summaries, such as the omitted information mentioned by other roles. Therefore, we propose a novel role interaction enhanced method for role-oriented dialogue summarization. It adopts cross attention and decoder self-attention interactions to interactively acquire other roles’ critical information. The cross attention interaction aims to select other roles’ critical dialogue utterances, while the decoder self-attention interaction aims to obtain key information from other roles’ summaries. Experimental results have shown that our proposed method significantly outperforms strong baselines on two public role-oriented dialogue summarization datasets. Extensive analyses have demonstrated that other roles’ content could help generate summaries with more complete semantics and correct topic structures.¹

1 Introduction

Dialogue summarization aims at compressing the main content of a long conversation into a short text. With the development of online conversation tools, the amount and length of conversation are growing up rapidly. Since a dialogue often contains complicated structure and ellipsis, it is time-consuming to read the whole dialogue. Dialogue summarization thus becomes valuable since it could extract the key point of a conversation and greatly reduce the time cost. This technique is widely used in customer service (Liu et al., 2019), meeting (McCowan et al., 2005), online chatting (Gliwa et al., 2019), etc.

*Corresponding author.

¹Our codes are available at: <https://github.com/xiaolinAndy/RODS>.

Dialogue	
0 Q: 我的这个原来绑定的手机号没用了, 然后我密码又忘了。(My original phone number was useless, and I forgot my password.)	
1 A: 这边为您转接专员您看可以吗? (Can I transfer you to the commissioner here?)	
2 Q: 好的。(Yes.)	
3 A: 很抱歉转接失败的呢亲, 这边为您升级专员[数字]小时内回电可以吗亲? (I'm sorry, the transfer failed. Can I help to upgrade the commissioner and call you back in [NUM] hours?)	
4 Q: 也可以的啦。(It's OK.)	
5 A: 请问还有其他还可以帮到您的吗? (Is there anything else I can help you with?)	
6 Q: 对了, 京东是可以微信支付的吧。(By the way, Can JD pay via wechat?)	
7 A: 正常是可以的呢亲。(Yes, it is OK normally.)	
User Summary	用户表示绑定手机号不用了, 密码忘记。用户询问京东是否可以微信支付。(The user said that the mobile phone number was useless, and forgot the password. The user asked whether Jingdong can pay via wechat.)
Agent Summary	客服帮助用户转接专员失败后表示[数字]小时内回电。客服表示京东可以微信支付。(The customer service helped the user to call back within [number] hours after the transfer specialist fails. The customer service said JD could pay via wechat.)

Information from the user’s utterances, could enhance agent summary in
Information from the agent’s summaries, could enhance user summary in
Information from the user’s summaries, could enhance agent summary in

Figure 1: An illustration of the role-oriented dialogue summarization. The task will generate *User Summary* and *Agent Summary* for the user (Q) and the agent (A), respectively. The information from other roles could help enhance the summary quality.

In a dialogue, each role has its own opinion and goal, and different roles exchange information or reach a consensus through interactions. Therefore, in addition to summarizing the whole dialogue, we could summarize the main content for each role. Lin et al. (2021) first define the role-oriented dialogue summarization task and provide a related dataset CSDS. They define role-oriented dialogue summarization as grasping the main viewpoint of a given role from dialogue and mention the usage of role-oriented summaries in the customer service domain, *e.g.*, reflecting the user’s requirements and evaluating agent service quality. Besides, role-oriented summarization is beneficial to other dialogue domains such as medical inquiry (Song et al., 2020) and court debate (Duan et al., 2019).

For role-oriented summarization, existing methods simply generate summaries for each role separately (Lin et al., 2021) or generate in a sequence

labeling process (Song et al., 2020). They ignore the strong relativeness among summaries for different roles and thus fail to utilize the information from other roles to enhance the summaries. However, information from other roles is also crucial for this task. We summarize two cases where other roles’ information helps:

(1) Other roles’ dialogue utterances could help enhance the informativeness of summaries.

In Figure 1, utterance 7 (*Yes, it is OK normally.*) is the key utterance of the agent’s content, expressing a confirmation to the user’s question. While only extracting it makes the agent summary ambiguous since it lacks the confirming object (*JD can pay via wechat* in blue). In this case, the agent summary needs to integrate the content from the user’s utterance (utterance 6 in yellow) to enhance its informativeness.

(2) Other roles’ summaries could help judge the key content in the dialogue.

In a dialogue, different roles often discuss the same topic. Therefore, considering the key content of the other role could help grasp the key content of a given role. As shown in Figure 1, the user summary contains a question about the payment (in red), and the agent summary contains the response to the payment question (in blue). If the summary of one role struggles in judging whether the discussion about payment should be contained in the summary, by referring to the summary of the other role, the summarization model could be more confident to include this information in the summary.

Although we notice the importance of other roles’ information, it is difficult to extract the key information from other roles through a simple multi-task framework. The main issue is that it could not judge which information from other roles is useful without modeling the interaction between different roles. Thus, in this work, we propose two interaction methods to obtain the key information from other roles for enhancing role-oriented summarization. First, we apply a cross attention interaction to let each role decoder select the most useful dialogue utterances from other roles. Specifically, we calculate the **Cross Attention** for different roles’ utterances separately and add a new **Attention Divergence Loss** to interactively share the cross attention distributions between different roles. Second, we apply a decoder self-attention interaction to let each role decoder obtain other roles’ summary information. We develop an inter-

active mechanism between decoders to consider other role summary information embedded in the decoder states. A new **Role Attention** module is added to each role decoder, where the attention object is the hidden states of other role decoders. At last, we use the role attention result and multiple context attention results to predict the word probability distribution of the summary. Through these two modules, the model could acquire more precise information from other roles and provide better role-oriented summaries.

To examine the effectiveness of our method, we conduct experiments on two dialogue summarization datasets (Lin et al., 2021; Song et al., 2020) with role-oriented summaries in different domains (customer service, medical inquiry). We apply our method on two widely-used summarization frameworks (RNN-based and Transformer-based). The results have shown that, compared with baseline systems and naive multi-task approaches, applying role interactions could significantly improve the quality of role-oriented summaries. Further analyses verify that our proposed method can help the model correctly attend to other roles’ key information and generate summaries with more complete semantic and correct topic structures.

The main contributions of this paper include: (1) We are the first to enhance role-oriented dialogue summarization by focusing on other roles’ key information. (2) We innovatively design two role interaction methods to obtain other roles’ key information useful for generating summaries. (3) Experimental results on two datasets have shown that our method could lead to considerable improvements. Besides, our method has good generalizability since it works on multiple baseline frameworks.

2 Related Work

2.1 Dialogue Summarization

Dialogue summarization has been studied in various domains, *e.g.*, meeting (McCowan et al., 2005; Janin et al., 2003), daily chatting (Gliwa et al., 2019; Chen et al., 2021), customer service (Liu et al., 2019; Zou et al., 2021), and medical inquiry (Song et al., 2020; Krishna et al., 2021). Considering the particularity of dialogue, many studies try to improve the dialogue summarization performance by focusing on dialogue-specific features (Feng et al., 2021), such as topic information (Chen and Yang, 2020), discourse structure (Chen and Yang, 2021), coreference information (Liu et al., 2021)

and speaker information (Lei et al., 2021; Zhu et al., 2020).

However, all the above studies focus on summarizing the whole dialogue. Only a few studies pay attention to role-oriented summarization, which aims to summarize the main content of a single role in the dialogue. A relative task is focused meeting summarization (Wang and Cardie, 2013; Mehdad et al., 2014; Zhong et al., 2021). It aims to summarize a specific part of the meeting dialogue, while role-oriented summarization focuses on a single role, and the relationship between different roles is much closer. Tamura et al. (2011) focus on contact center dialogue summarization, but they only extract salient sentences from the dialogue and do not summarize for different roles.

Due to the lack of labeled data, Zhang et al. (2021) propose an unsupervised method to generate summaries for the customer and the agent under a variational auto-encoder framework. As for supervised methods, there are only two datasets available for training. Lin et al. (2021) propose a customer service domain dataset named CSDS, where each dialogue has an overall summary and two role-oriented summaries for user and agent. They train two separate models for generating user summary and agent summary. Song et al. (2020) provide a medical inquiry dialogue summarization dataset where each dialogue has two extractive summaries for the patient and the doctor. They train a sequence labeling model to extract summaries for these two roles. Compared with these approaches, to the best of our knowledge, we are the first to enhance role-oriented summarization by explicitly considering other roles' critical information.

2.2 Interactive Decoding

Interactive decoding is a mechanism to share information between different decoders in the decoding process. Zhou et al. (2019) propose this mechanism and use it on machine translation to simultaneously decode from both left-to-right and right-to-left. Wang et al. (2019) and Liu et al. (2020) further utilize it on more complex machine translation tasks, including multilingual translation and speech translation. In this work, we first use the interactive decoding mechanism on the summarization task to decode summaries for different roles, aiming at utilizing other roles' summary information for summarization. Besides, we also propose an interaction method on cross attention to utilize

other roles' critical dialogue utterance information.

3 Our Approach

3.1 Task Definition

Given a dialogue D containing m utterances $\{u_1, \dots, u_m\}$ and p speakers $S = \{s_1, \dots, s_p\}$, the role-oriented summarization task aims to generate a summary y^k for each speaker s_k . Each utterance u_k consists of a speaker role $r_k \in S$ and related content. By concatenating all the utterances and related speaker roles, we achieve the final input $\{x_1, \dots, x_n\}$. Note that since both datasets used in this work have two speakers, one asking questions and one giving answers, we thus use y^{user} and y^{agent} to represent two role-oriented summaries in the following illustration².

3.2 Role Interactions

In a traditional encoder-decoder framework for dialogue summarization, the encoder hidden states represent the semantic information of input dialogue utterances, and the decoder hidden states contain the information used to generate summaries. To fully exploit the information from other roles, we apply two role interactions on the attention module of both hidden states. We present the structure of our method in Figure 2 and introduce the detail of interactions in the following paragraphs.

3.2.1 Cross Attention Interaction

Our method is constructed based on a multi-task framework where an encoder is used to encode dialogue utterances and two role decoders (user decoder and agent decoder) are used to decode user summary and agent summary. First, the input $\{x_1, \dots, x_n\}$ is sent to an encoder (omitted in the figure for simplicity) and the encoder outputs the context hidden representation $\{h_1, \dots, h_n\}$. In the decoding phase, to calculate the cross attention results for different roles separately, we use **User Mask** and **Agent Mask** to split the context information into user context H_u^{enc} and agent context H_a^{enc} . H_u^{enc} contains the hidden representation of all user utterances, and H_a^{enc} contains the hidden representation of all agent utterances.

The cross attention module extracts the most useful information from the context according to the temporary decoder state. Here we modify the module to attend to different role contexts separately.

²Here we need to point out that our method could also apply for more than two speakers.

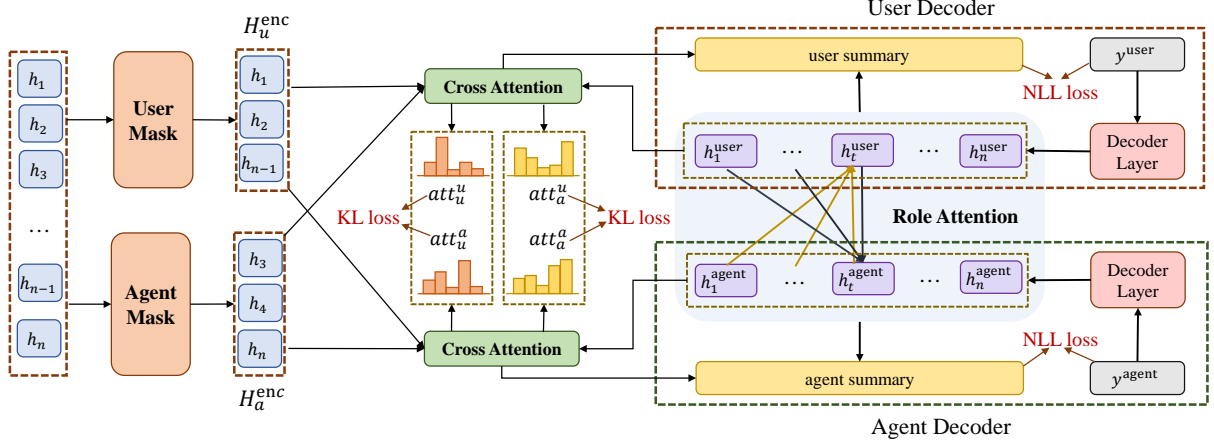


Figure 2: The model structure of our proposed method with role interactions.

Taking user decoder as example, at step k , we use the hidden state of user decoder h_k^{user} to attend to user context H_u^{enc} and agent context H_a^{enc} , obtaining two attention distributions $att_{u,k}^u$, $att_{a,k}^u$ and context attention results $c_{u,k}^u$, $c_{a,k}^u$. Both context results involve generating summaries. The process is the same with agent decoder, where two attention distributions are noted as $att_{u,k}^a$, $att_{a,k}^a$.

Since existing models are poor at extracting important information from other roles, it reflects in incorrect cross-role attentions $att_{u,k}^a$ (agent decoder to user context) and $att_{a,k}^u$ (user decoder to agent context). Meanwhile, the same-role attentions $att_{u,k}^u$ (user decoder to user context) and $att_{a,k}^a$ (agent decoder to agent context) are learned better since most information of role-oriented summaries comes from the given role's utterances. Thus we want to use the same-role attention to guide the cross-role attention. As different roles often discuss the same topic in one dialogue, the accumulated cross attention distribution for user decoder and agent decoder on the same role's utterances should be similar. A new **Attention Divergence Loss** is added to constrain this attention similarity as:

$$\mathcal{L}_{\text{att-user}} = \text{KL}(\text{Avg}(att_u^a) || \text{Avg}(att_u^u))$$

$$\mathcal{L}_{\text{att-agent}} = \text{KL}(\text{Avg}(att_a^u) || \text{Avg}(att_a^a))$$

By minimizing these two losses, the agent decoder attends to user utterances as the user decoder does and the user decoder attends to agent utterances as the agent decoder does. Two role decoders interactively learn to focus on the key information of the other role in dialogue utterances.

3.2.2 Decoder Self-Attention Interaction

Since the decoder calculates the hidden states that could help predict summaries, the hidden states must contain much important information of summaries. We thus try to exploit the information embedded in other role decoders. Specifically, for user decoder, at time step t , we achieve the decoder hidden states h_t^{user} and use a **Role Attention** module to weigh the last t hidden states of agent decoder $\{h_1^{\text{agent}}, \dots, h_t^{\text{agent}}\}$ ³. The role context information r_t^{user} is obtained by adding all the agent hidden states with their weights, and it helps generate the probability of next word \hat{y}_t^{user} for user summary. The calculation formulas are given as:

$$r_t^{\text{user}} = \text{Attn}(h_t^{\text{user}}, h_{1:t}^{\text{agent}})$$

$$p(\hat{y}_t^{\text{user}}) = \mathcal{F}(h_t^{\text{user}}, r_t^{\text{user}}, c_{u,k}^u, c_{a,k}^u)$$

The function \mathcal{F} includes an MLP layer to fuse different information and a softmax layer to predict the vocabulary probability distribution. The process is the same with the agent decoder, and two decoders decode interactively.

3.2.3 Training and Inference

In the training phase, we use the teacher-forcing method to jointly train two role decoders and use the Negative Log-Likelihood loss to optimize. The NLL loss for a single sample is formulized as:

$$\mathcal{L}_{\text{nll}} = -(\alpha \cdot \sum_{i=1}^{|y^{\text{user}}|} \log P(y_i^{\text{user}} | y_{<i}^{\text{user}}, y_{<i}^{\text{agent}}, D) + (1 - \alpha) \cdot \sum_{i=1}^{|y^{\text{agent}}|} \log P(y_i^{\text{agent}} | y_{<i}^{\text{agent}}, y_{<i}^{\text{user}}, D))$$

³Since two decoders decode simultaneously, at step t , the other decoder could only provide the states from step 1 to t .

α is a hyper-parameter for balancing the weights of different summarization tasks. Besides, we add the attention divergence loss to constrain the attention distribution, and the total loss is calculated as:

$$\mathcal{L} = \mathcal{L}_{\text{nll}} + \beta(\mathcal{L}_{\text{att-user}} + \mathcal{L}_{\text{att-agent}})$$

β is a hyper-parameter for balancing the weights of different loss functions.

In the inference phase, we also make some adjustments on beam search for our proposed method. We maintain two beams, one for user summary and one for agent summary. At each step of decoding, the k^{th} sequence of the user summary beam should consider the states in the k^{th} sequence of the agent summary beam for role attention. Once one beam has finished decoding, we keep the beam fixed and search for the other one. The beam search will finish if both two beams have finished searching.

4 Experiments

4.1 Datasets

There are two dialogue summarization datasets with role-oriented summarization tasks. Thus we evaluate the effectiveness of our proposed method on both datasets. First, we experiment on a Chinese fine-grained customer service summarization dataset CSDS⁴ (Lin et al., 2021). It provides separate summaries for the user and the agent, and both may contain multiple topics. The other one is a Chinese medical inquiry summarization dataset MC⁵ (Song et al., 2020). Each dialogue has a summary of the patient’s description and a summary of the doctor’s suggestion. We note them as user summary and agent summary as well. Most of the summaries in MC are extractive, and only some are different from dialogue scripts. Moreover, most dialogues in MC have only one topic. Comparing two datasets, MC is easier to summarize while CSDS is more specific for role-oriented summarization and more challenging. The detailed statistics of the two datasets are given in Table 1.

4.2 Baselines and Experiment Settings

We apply the role interaction methods on two widely-used seq2seq models in the summarization

	CSDS	MC
Train Size	9,101	29,324
Val Size	800	3,258
Test Size	800	8,146
Turns	25.92	18.48
Dial. Length	321.92	292.21
User Sum. Length	37.28	22.37
Agent Sum. Length	48.08	95.32

Table 1: Statistics of CSDS and MC. All the lengths are counted on Chinese characters.

community, including PGN (See et al., 2017) and BERTAbs (Liu and Lapata, 2019). Therefore, we will introduce these two backbone models and how we apply **Role Interactions** to them.

4.2.1 PGN-based Methods

PGN is an LSTM-based seq2seq model with a copy mechanism to copy words from the input and a coverage mechanism for constraining context attention. We set two PGN-based baselines for comparison. **PGN-single** is to separately train two PGN models for generating user summary and agent summary, while **PGN-multi** tries to jointly train two PGN models by sharing the same encoder. Both baselines adopt all the dialogue context as input.

To apply role interactions, we choose the output of the LSTM cell in the decoder as the query to calculate cross attention and role attention. The attention object in role attention is the output of the LSTM cell from the other decoder. Since we calculate the cross attention for different roles separately, we use a learnable gate p_{role} to control the weight of different cross attentions and add them together according to their weights to achieve the overall context attention distribution. It is also used for the copy and coverage mechanism. We set **PGN-cross** as adding cross attention interaction, **PGN-self** as adding decoder self-attention interaction, and **PGN-both** as adding both interactions.

4.2.2 BERTAbs-based Methods

Transformer has been widely used in language understanding and generation models due to its strong representation ability and concurrency, especially in pretrained models (Devlin et al., 2019; Lewis et al., 2020). Here we choose BERTAbs (Liu and Lapata, 2019) as the backbone structure since it performs well on many summarization datasets and is available for non-English languages such as Chinese. It adopts a pretrained BERT model as encoder and a transformer decoder structure to decode summaries. Both the encoder and the decoder contain

⁴<https://github.com/xiaolinAndy/CSDS>.

⁵<https://github.com/cuhksz-nlp/HET-MC>. We use the official crawling script to acquire the dataset and divide some data from the training set as the validation set. Due to the update of the website, the data may have a slight difference compared with the version in the original paper.

six layers, and each layer contains three sub-layers (self-attention, encoder-decoder attention, feedforward). Similar with PGN-based methods, we set **BERT-single** and **BERT-multi** as two baselines.

We apply both interactions to each layer in BERTAbs. For cross attention interaction, we change the encoder-decoder attention sub-layer into two separate cross attention modules; for decoder self-attention interaction, we add the role attention module parallel with the encoder-decoder attention module. The query, key, and value of the role attention module are all the output from the self-attention sub-layer. **BERT-cross**, **BERT-self**, and **BERT-both** are kept the same with the settings in PGN-based methods.

4.2.3 Other Experiment Settings

We add the role information to the front of the utterance in each turn and concatenate all the utterances in the dialogue sequentially as the input of the model. Both PGN⁶ and BERTAbs⁷ baseline methods are adopted from publicly available codes. For PGN-based methods, we use pretrained Chinese word vectors provided by Tencent⁸, and the vocabulary size is 10,000. While for BERTAbs-based methods, we use Chinese-BERT-wwm⁹ pretrained embeddings. The best checkpoint is chosen based on validation set loss, and we use beam search to decode summaries on the best checkpoint with beam size 5. For choosing hyper-parameters, α is set to be 0.5 for PGN and 0.25 for BERTAbs. Since the agent summary is more complex than the user summary in MC, we set β to be 0.2 to give the agent summary more weight. It is set to be 0.5 for CSDS. The hyper-parameter settings are chosen by experimenting on the validation set. More details are given in Appendix A.

4.3 Evaluation Metrics

We adopt six common automatic evaluation metrics to evaluate the summary quality. The metrics include traditional n-gram overlapping metrics, such as **ROUGE** (Lin and Hovy, 2002), **BLEU** (Papineni et al., 2002), and distributed representation matching metrics, including **BERTScore** (Zhang et al., 2020) and **MoverScore** (Zhao et al., 2019). We use files2rouge toolkit to calculate the F1 score of ROUGE-1, ROUGE-2, ROUGE-L. More details

of evaluation scripts are given in Appendix A.

In addition to automatic metrics, we also compare the summary quality at a fine-grained level through human evaluation. Following the human evaluation process in Lin et al. (2021), we recruit several volunteers and let them evaluate the summaries in the following aspects: (1) **Informativeness**: Does the generated summary correctly cover the information in the ground truth summary? (2) **Non-redundancy**: Does the generated summary not contain repeated, meaningless or unnecessary information? (3) **Fluency**: Is the generated summary well-formed, semantically complete, and easy to understand? All three aspects are evaluated at the sub-summary level¹⁰ on a three-point scale, 0 for the worst, 1 for the medium, and 2 for the best.

5 Results and Analysis

5.1 Automatic Evaluation Results

First, we present the results of automatic metrics with Student’s t-test as significance test in Table 2 and 3. The results are similar on the two datasets. First, the multi-task mechanism could bring some improvement than separately training on most of the metrics. However, the improvement is limited, especially for the PGN model on CSDS. After adding the enhancement of other roles’ information, the performance is significantly boosted.

On CSDS, PGN-single and BERT-single are two strong baselines provided in Lin et al. (2021)¹¹. For PGN-based methods, the best method PGN-both utilizes two interactions and achieves 2.84 and 1.53 higher points on ROUGE-L for user summary and agent summary than PGN-single. For BERTAbs-based methods, the improvements are even greater, which are 4.73 and 2.69. We also conduct ablation studies by only applying one interaction (*-cross* or *-self*). Both settings show promising improvement than *single* and *multi* baselines on nearly all the metrics, demonstrating the effectiveness of each interaction method. In comparison, applying two interactions together yields the best result on the majority of metrics.

The circumstance is similar on MC. User sum-

¹⁰We split summaries into different topic segments, and each segment is a sub-summary, same as the process in Lin et al. (2021).

¹¹Note that we do not mention the baseline Fast-RL (Chen and Bansal, 2018) in Lin et al. (2021). It first extracts salient utterances and generates summary sentences from each utterance separately, which is not available to add our proposed interaction methods.

⁶https://github.com/atulkum/pointer_summarizer

⁷<https://github.com/nlpyang/PreSumm>

⁸<https://ai.tencent.com/ailab/nlp/en/embedding.html>

⁹<https://github.com/ymcui/Chinese-BERT-wwm>

CSDS	ROUGE-1		ROUGE-2		ROUGE-L		BLEU		BERTScore		MoverScore	
	user	agent	user	agent	user	agent	user	agent	user	agent	user	agent
PGN-single	53.55	50.20	37.06	35.12	51.05	47.59	29.64	28.25	78.68	76.13	26.68	25.13
PGN-multi	54.01	49.94	37.38	34.78	51.95	48.20	30.04	29.09	78.78	75.95	27.16	24.90
PGN-cross	54.34	50.80	37.75	35.89	51.95	48.20	31.19	30.40	78.97	76.51	27.89	25.60
PGN-self	55.49	51.00	38.75	35.70	53.08	48.52	31.84	30.47	79.37	76.48	27.74	25.55
PGN-both	56.08*	51.62*	39.10*	36.50*	53.89*	49.12*	33.54*	29.78*	79.52*	76.74*	28.28*	26.25*
BERT-single	52.72	49.57	36.39	33.82	50.44	46.83	30.17	26.99	79.23	76.39	24.96	23.87
BERT-multi	56.09	50.49	39.91	35.17	54.02	48.08	26.91	25.39	80.50	76.65	27.19	23.71
BERT-cross	57.29	50.35	41.03	35.27	55.29	48.09	30.70	24.19	80.90	76.65	28.55	23.70
BERT-self	56.94	50.96	40.37	35.24	54.85	48.40	30.61	27.13	80.53	76.80	28.24	24.83
BERT-both	57.36*	51.92*	40.70	36.37*	55.17*	49.52*	32.04*	29.23*	80.70	77.23*	28.66*	25.48*

Table 2: The automatic metric results for CSDS. * indicates that the improvement of applying two interactions (PGN-both, BERT-both) over *single* and *multi* are both statistically significant ($p < 0.01$).

MC	ROUGE-1		ROUGE-2		ROUGE-L		BLEU		BERTScore		MoverScore	
	user	agent	user	agent	user	agent	user	agent	user	agent	user	agent
(Song et al., 2020)	92.80	83.31	88.97	75.48	92.80	83.29	-	-	-	-	-	-
PGN-single	94.83	82.63	94.32	77.83	94.78	81.51	87.66	68.10	97.60	91.74	90.28	67.95
PGN-multi	94.58	83.16	93.98	78.33	94.53	81.96	87.23	69.96	97.49	91.92	89.87	68.42
PGN-cross	95.12	83.40	94.63	78.60	95.07	82.18	87.99	69.61	97.75	92.07	90.73	69.06
PGN-self	95.08	83.17	94.59	78.48	95.04	82.00	87.90	69.29	97.70	91.99	90.64	68.54
PGN-both	95.11*	83.48*	94.59*	78.73*	95.06*	82.28*	87.82*	69.63	97.71*	92.15*	90.66*	69.24*
BERT-single	95.13	81.66	94.50	76.73	95.08	80.42	87.20	64.09	97.86	91.71	90.31	68.29
BERT-multi	95.18	81.20	94.61	76.37	95.13	79.97	87.38	64.83	97.90	91.51	90.71	67.55
BERT-cross	95.18	81.75	94.61	77.04	95.13	80.55	87.40	65.63	97.89	91.70	90.67	68.28
BERT-self	95.18	81.61	94.61	77.01	95.13	80.49	87.37	65.01	97.89	91.72	90.69	68.37
BERT-both	95.19	82.11*	94.63	77.49*	95.14	80.92*	87.40	65.40*	97.90	91.91*	90.72	68.95*

Table 3: The automatic metric results for MC. * represents the same with the one in Table 2.

marization is relatively simple on MC, and the baseline methods could achieve high performance (5.35 points of ROUGE-2 higher than the best performance in the original paper (Song et al., 2020)). Despite this, both cross attention interaction and decoder self-attention interaction could still increase the performance of user summary a bit. Additionally, the improvement on agent summary is more significant. PGN-both method achieves 0.90 points of ROUGE-2 and 1.29 points of MoverScore improvement, while BERT-both achieves 0.76 points of ROUGE-2 and 0.66 points of MoverScore improvement. PGN-both also beats the best result in the original paper on most of the metrics, which uses additional information such as hospital department and disease name. In conclusion, our proposed two interaction methods could bring remarkable improvement on different backbone structures and different datasets.

5.2 Human Evaluation Results

To evaluate the summary quality at a more fine-grained level, we compare the summaries from different models according to the pre-defined three aspects: informativeness, non-redundancy, fluency. Since the multi-task framework works better than the single baseline, we directly compare it with applying both interactions. As CSDS is more chal-

CSDS	Info	Non-Red	Flu	Overall
PGN-multi	0.69/0.65	0.54/0.55	0.70/0.79	0.64/0.66
PGN-both	0.66/ 0.69	0.58/0.59*	0.73/0.81	0.66/0.70*
BERT-multi	0.58/0.56	0.66/0.61	0.84/ 0.87	0.69/0.68
BERT-both	0.62*/0.60*	0.62/0.60	0.85/0.87	0.70/0.69

Table 4: The human evaluation results for CSDS. Two values in each block represents user summary and agent summary. All the values are in range 0 to 1. * indicates that the improvement of applying two interactions over the *multi* baseline is statistically significant ($p < 0.05$).

lenging for this task, we randomly select 100 samples from the test set and obtain the outputs of two baseline methods (PGN-multi and BERT-multi) and two interaction methods (PGN-both and BERT-both). We recruit three volunteers and train them on the evaluation rules¹². Then we let them evaluate the generated summaries according to the ground truth and the original dialogue in the three aspects. We run the inter-annotator agreement study on three volunteers’ scores, and obtain a reasonable kappa score, 0.48 on average. We also calculate an “Overall” metric by averaging the results of all three aspects to represent the summary quality in general. We normalize the result into 0 to 1 and present it in Table 4.

The result shows different trends on two backbone structures. For the PGN model, applying

¹²More details are in Appendix C with ethical concerns.

CSDS	ROUGE-1 Type A/B	ROUGE-2 Type A/B	ROUGE-L Type A/B	BLEU Type A/B	BERTScore Type A/B	MoverScore Type A/B
PGN-multi	55.13/59.45	37.76/41.22	52.73/56.20	30.66/28.29	76.40/77.64	23.74/25.47
PGN-both	56.00/62.28	38.58/43.88	53.66/58.99	31.06/29.14	76.84/78.59	24.41/27.15
BERT-multi	46.59/50.07	32.33/34.59	44.49/47.65	23.45/26.37	75.03/75.64	22.47/24.34
BERT-both	50.96/54.62	35.72/37.93	48.82/51.93	27.47/30.10	76.27/76.94	24.19/26.17

Table 5: The performance on different types of samples. Type A represents agent summaries that need to be integrated, and Type B represents for those that do not. Here all the metrics here are recall scores except for BLEU and MoverScore since they do not have a recall version. We use their available results instead.

interactions could largely reduce the redundancy of both user and agent summary, with a comparable performance of informativeness. Besides, it also improves the fluency of the two summaries. While for the BERTAbs model, the interaction method significantly improves the informativeness while the redundancy also increases a bit. The difference exists because BERTAbs prefers to generate short summaries. Thus, considering information from other roles could help generate more useful information but also induce some redundant texts. In contrast, PGN tends to generate lengthy summaries. When considering information from other roles, it first tries to discard the redundant texts and only retains more important ones. The fluency improvement on both methods proves that other roles’ information helps generate more semantically complete summaries. Considering the overall metric, we conclude that our proposed interaction method is also effective through human evaluation.

5.3 Further Analysis

Agent Summary Completeness Analysis The agent summary often suffers semantic incompleteness due to missing key information from other roles (Lin et al., 2021). Since our proposed role interactions aim at extracting other roles’ key information, we wonder whether it works on these incomplete cases. Following the settings in Lin et al. (2021), we compare the summary quality of samples that need to integrate other roles’ information and those that do not need separately¹³. The result in Table 5 shows that the interaction method could actually help improve the performance on samples that need to integrate. Besides, samples that do not need also get improved. We believe that it is because considering other roles’ information could also help extract critical content from the role’s own utterances as well.

¹³It is judged by considering whether the summary needs to refer to other roles’ utterances, which is already labeled in CSDS.

Methods	Precision	Recall	F1
PGN-multi	28.61/18.86	28.87/19.67	28.74/19.27
PGN-both	31.79/21.06	30.85/21.58	31.31/21.32
BERT-multi	40.16/23.99	30.26/18.81	34.51/21.09
BERT-both	37.37/22.09	32.17/20.66	34.57/21.35

Table 6: Sub-summary matching ratio for baselines and our methods. Two values in each block represents user summary and agent summary.

Topic Structural Summary Analysis Since we assume that role interactions could help generate better summaries by sharing the same discussion topic, we wonder whether the summaries generated by our methods could include the correct topic structure. More specifically, we want to find out the performance of our methods on summarizing each topic. Following the evaluation method in (Lin et al., 2021), we treat each sentence in the summary as a sub-summary for a single topic and calculate the number of matching sub-summaries with the reference by a ROUGE-L-based matching algorithm. We calculate the precision, recall, and F1 scores of correctly matched sub-summary ratios and present them in Table 6. The result shows that two role interaction methods achieve higher recall and F1 scores on sub-summary matching. It proves that role interactions could help the model grasp the discussion topic in the dialogue and generate a more accurate summary for each topic.

We also present an example in Appendix B to prove the effectiveness of our proposed role interaction method.

6 Conclusion and Future Work

In this paper, we focus on the role-oriented dialogue summarization task. To fully exploit the information from other roles, we propose two role interaction methods on cross attention and decoder self-attention. The cross attention interaction calculates the context information for different roles separately and uses same-role attention to guide cross-role attention. The decoder self-attention in-

teraction adds a role attention module to attend to other role decoder states interactively. Experiments on two dialogue summarization datasets prove that both interactions perform significantly better than strong baseline methods. Adding role interactions also helps generate summaries with complete semantics and correct topic structure. In the future, we will try to apply this method to other dialogue-related tasks and conduct more experiments on stronger summarization methods.

Acknowledgements

First, We thank all anonymous reviewers for giving useful suggestions. Second, we thank all the volunteers for doing the human evaluation. The research work has been supported by the Natural Science Foundation of China under Grant No. 62106263.

References

- Jiaao Chen and Diyi Yang. 2020. [Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4106–4118, Online. Association for Computational Linguistics.
- Jiaao Chen and Diyi Yang. 2021. [Structure-aware abstractive conversation summarization via discourse and action graphs](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1380–1391, Online. Association for Computational Linguistics.
- Yen-Chun Chen and Mohit Bansal. 2018. [Fast abstractive summarization with reinforce-selected sentence rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021. [DialogSum: A real-life scenario dialogue summarization dataset](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xinyu Duan, Yating Zhang, Lin Yuan, Xin Zhou, Xiaozhong Liu, Tianyi Wang, Ruocheng Wang, Qiong Zhang, Changlong Sun, and Fei Wu. 2019. [Legal summarization for multi-role debate dialogue via controversy focus mining and multi-task learning](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 1361–1370, New York, NY, USA. Association for Computing Machinery.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2021. [A survey on dialogue summarization: Recent advances and new frontiers](#). *arXiv preprint arXiv:2107.03175*.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. [SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al. 2003. [The icsi meeting corpus](#). In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)*, volume 1, pages I–I. IEEE.
- Kundan Krishna, Sopan Khosla, Jeffrey Bigham, and Zachary C. Lipton. 2021. [Generating SOAP notes from doctor-patient conversations using modular summarization techniques](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4958–4972, Online. Association for Computational Linguistics.
- Yuejie Lei, Yuanmeng Yan, Zhiyuan Zeng, Keqing He, Ximing Zhang, and Weiran Xu. 2021. [Hierarchical speaker-aware sequence-to-sequence model for dialogue summarization](#). In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7823–7827.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin and Eduard Hovy. 2002. [Manual and automatic evaluation of summaries](#). In *Proceedings of the ACL-02 Workshop on Automatic Summarization*, pages 45–51, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Haitao Lin, Liqun Ma, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2021. [CSDS:](#)

- A fine-grained Chinese dataset for customer service dialogue summarization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4436–4451, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Chunyi Liu, Peng Wang, Jiang Xu, Zang Li, and Jieping Ye. 2019. Automatic dialogue summary generation for customer service. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1957–1965.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yuchen Liu, Jiajun Zhang, Hao Xiong, Long Zhou, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2020. Synchronous speech recognition and speech-to-text translation with interactive decoding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8417–8424.
- Zhengyuan Liu, Ke Shi, and Nancy Chen. 2021. Coreference-aware dialogue summarization. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 509–519, Singapore and Online. Association for Computational Linguistics.
- I McCowan, J Carletta, W Kraaij, S Ashby, S Bourban, M Flynn, M Guillemot, T Hain, J Kadlec, V Karaiskos, et al. 2005. The ami meeting corpus. In *Proceedings of Measuring Behavior 2005, 5th International Conference on Methods and Techniques in Behavioral Research*, pages 137–140. Noldus Information Technology.
- Yashar Mehdad, Giuseppe Carenini, and Raymond T. Ng. 2014. Abstractive summarization of spoken and written conversations based on phrasal queries. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1220–1230, Baltimore, Maryland. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Yan Song, Yuanhe Tian, Nan Wang, and Fei Xia. 2020. Summarizing medical conversations via identifying important utterances. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 717–729, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Akihiro Tamura, Kai Ishikawa, Masahiro Saikou, and Masaaki Tsuchida. 2011. Extractive summarization method for contact center dialogues based on call logs. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 500–508, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- Lu Wang and Claire Cardie. 2013. Domain-independent abstract generation for focused meeting summarization. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1395–1405, Sofia, Bulgaria. Association for Computational Linguistics.
- Yining Wang, Jiajun Zhang, Long Zhou, Yuchen Liu, and Chengqing Zong. 2019. Synchronously generating two languages with interactive decoding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3350–3355, Hong Kong, China. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Xinyuan Zhang, Ruiyi Zhang, Manzil Zaheer, and Amr Ahmed. 2021. Unsupervised abstractive dialogue summarization for tete-a-tetes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14489–14497.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 563–578, Hong Kong, China. Association for Computational Linguistics.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 2021. QMSum: A new benchmark for query-based multi-domain meeting summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5905–5921, Online. Association for Computational Linguistics.

Long Zhou, Jiajun Zhang, and Chengqing Zong. 2019. [Synchronous Bidirectional Neural Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 7:91–105.

Chenguang Zhu, Ruochen Xu, Michael Zeng, and Xuedong Huang. 2020. [A hierarchical network for abstractive meeting summarization with cross-domain pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 194–203, Online. Association for Computational Linguistics.

Yicheng Zou, Lujun Zhao, Yangyang Kang, Jun Lin, Minlong Peng, Zhuoren Jiang, Changlong Sun, Qi Zhang, Xuanjing Huang, and Xiaozhong Liu. 2021. [Topic-oriented spoken dialogue summarization for customer service with saliency-aware topic modeling](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14665–14673.

Appendix

A Experiment Details

Here, we will introduce some detailed settings for our experiments on two datasets.

- **PGN-based methods:** We construct the vocabulary by choosing the top 10,000 most frequent words in the training data. The settings of PGN are the same as the original setting with hidden size 256. The optimizer is Adagrad and the learning rate is 0.15.

For CSDS dataset, we use the given word split result to construct the vocabulary. The maximum input length is set as 500. The maximum output length is 100, and the minimum is 15. We train 40 epochs without coverage mechanism or KL divergence loss (if needed) and 10 epochs with coverage mechanism and KL divergence loss. Then we choose the best checkpoint by comparing the loss on the validation set and use it to decode summaries.

For MC dataset, we use jieba¹⁴ tool to split sentences into words for constructing the vocabulary. The maximum input and output length are the same with CSDS, while the minimum output length is 10 since the user summary in MC is relatively shorter. We train 30 epochs without coverage mechanism and do not finetune with coverage mechanism since we found that it makes the performance worse. The KL divergence loss is added to the training loss for PGN-cross and PGN-both in all the training process.

¹⁴<https://pypi.org/project/jieba/>

- **BERTAbs-based methods:** Since the BERT model is already finetuned, there is no need to reconstruct the vocabulary. The Chinese BERT model works on character-level. Thus we set the length limit larger. The dimension and optimizer settings of BERTAbs are also the same as the original settings.

For CSDS dataset, the maximum input and output length are 1,000 and 200, respectively. The minimum output length is 15. We train the model for 4000 steps and save the checkpoint for every 400 steps. We use Adam optimizer with a warmup of 1000 steps. The KL divergence loss is added by finetuning 1000 more steps. During the inference time, we control the maximum non-repeat n-gram length as 5.

For MC dataset, the maximum input and output length are kept the same as in CSDS, and the minimum output length is 10. We train the model for 8000 steps and add the KL divergence loss in all the training process.

All the PGN-based models are run on an NVIDIA TITAN Xp, and all the BERTAbs-based models are run on an NVIDIA RTX3090. The whole running time is less than a week.

We also provide the running scripts of auto evaluation metrics for better reproduction. For ROUGE metrics, we use the files2rouge¹⁵ toolkit with the default parameters. All the Chinese characters are transferred into number ids for calculation, and the period is used to split each summary into several sentences for ROUGE-L calculation. For BERTScore, we use the official code¹⁶. For MoverScore, we use moverscore-v2¹⁷ and the bert-base-chinese pretrained model for obtaining representations.

B Case Study

Here we use the same example illustrated in the main paper to prove the effectiveness of our proposed method. The outputs of different methods are given in Figure 3. Comparing the outputs of user summary, only PGN-both correctly summarizes the key issue “*The user asked whether wechat payment is available.*”, while other baselines omit it. This could be contributed to decoder self-attention

¹⁵<https://github.com/pltrdy/files2rouge>

¹⁶<https://pypi.org/project/bert-score/0.2.1/>

¹⁷<https://github.com/AIPHES/emnlp19-moverscore>

interaction by attending to the information in the decoded agent summary. As for the agent summary, PGN omits the keyword “JD” and PGN-multi generates a redundant clause “*The customer service replied that it can be paid by wechat*”. Compared with them, PGN-both generates the agent summary with both informativeness and preciseness.

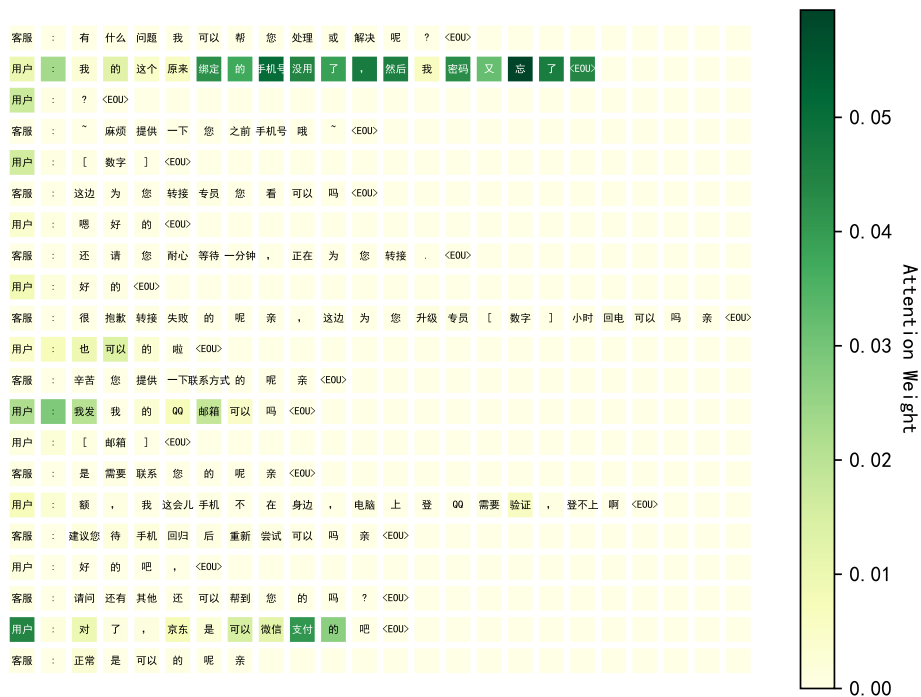
As given in Figure 4, we also present the average attention distribution for cross attention module in the PGN-both method. Although the attention distributions for user summary and agent summary are different, they also show some similarities, such as both focusing on the second line from the bottom, which is the key utterance for both summaries. This could be benefited by the cross attention interaction since we close the gap between the attention distributions of different role decoders on the same role’s utterances.

C Ethical Concerns

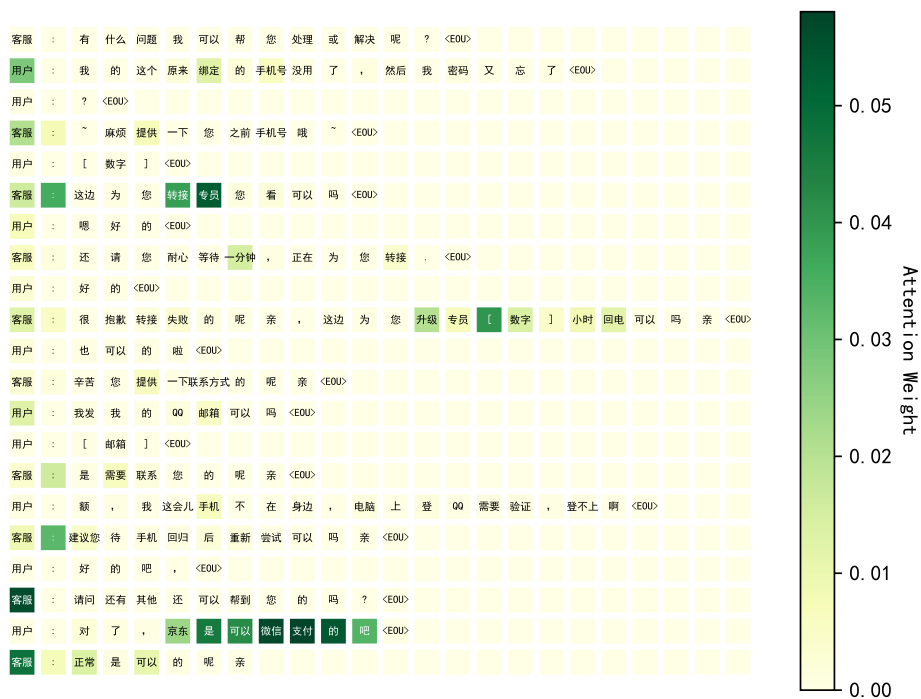
We only use the data provided by two datasets for training. The private information in CSDS has already been anonymized, such as replacing all numbers with special token <NUM> and all order IDs with <ORDER-ID>. There is no personal information available in CSDS. The circumstance is the same for MC, where all the dialogues do not contain detailed personal information. Thus the methods provided in our experiment do not have any issues with privacy disclosure. As for human evaluation, all the participants are Chinese graduate students who volunteer to make the evaluation, and they are all proficient in Chinese. We first let them read the evaluating instructions and let them evaluate on ten samples without showing the model name for each summary. After confirming that the results provided by three volunteers attain a moderate level of agreement on all the aspects, we allow them to examine the remaining samples.

User Summary	
PGN	用户询问原来绑定的手机号没用了，然后密码又忘了吗。(The user said that the original mobile phone number was useless, and forgot the password.)
PGN-multi	用户询问原来绑定的手机号没用了，然后我密码又忘了怎么办。(The user said that the original mobile phone number was useless, and how to deal with since I forgot the password.)
PGN-both	用户表示绑定的手机号没用了，然后密码又忘了。 用户询问是否可以微信支付。 (The user said that the original mobile phone number was useless, and forgot the password. The user asked whether wechat payment is available.)
Refer-ence	用户表示绑定手机号不用了，密码忘记。 用户询问京东是否可以微信支付。 (The user said that the mobile phone number was useless, and forgot the password. The user asked whether Jingdong can pay via wechat.)
Agent Summary	
PGN	客服回答这边为用户升级专员[数字]小时回电。客服回答 正常是可以微信支付。 (The customer service helped the user upgrade the commissioner and called back in [NUM] hours. The customer service replied that it can be paid by wechat normally.)
PGN-multi	客服帮助用户转接专员。 客服表示可以微信支付。 客服表示京东是可以微信支付。(The customer service helped the user transfer the specialist. The customer service replied that it can be paid by wechat. The customer service replied that JD can be paid by wechat.)
PGN-both	客服帮助用户转接专员处理。客服表示 京东是可以微信支付的。 (The customer service helped the user transfer the specialist to deal with. The customer service replied that JD can be paid by wechat.)
Refer-ence	客服帮助用户转接专员失败后表示[数字]小时内回电。客服表示 京东可以微信支付。 (The customer service helped the user to call back within [number] hours after the transfer specialist fails. The customer service said JD could pay via wechat.)

Figure 3: The generated summaries and ground truth for the example dialogue.



(a) The average attention distribution for user summary



(b) The average attention distribution for agent summary

Figure 4: The average attention distribution for decoding summaries using the PGN-both method.