



Bilateral personalized dialogue generation with contrastive learning

Bin Li¹ · Hanjun Deng²

Accepted: 6 September 2022 / Published online: 16 September 2022
© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

Generating personalized responses is one of the major challenges in natural human–robot interaction. Current studies in this field mainly focus on generating responses consistent with the robot’s pre-assigned persona, while ignoring the user’s persona. Such responses may be inappropriate or even offensive, which may lead to the bad user experience. Therefore, we propose a bilateral personalized dialogue generation (BPDG) method for dyadic conversation, which integrates user and robot personas into dialogue generation via designing a dynamic persona-aware fusion method. To bridge the gap between the learning objective function and evaluation metrics, the conditional mutual information maximum (CMIM) criterion is adopted with contrastive learning to select the proper response from the generated candidates. Moreover, a bilateral persona accuracy metric is designed to measure the degree of bilateral personalization. Experimental results demonstrate that, compared with several state-of-the-art methods, the proposed method achieves the improvement on the random personalized test set of 23.99 in bilateral persona accuracy, 1.1 in BLEU, 0.83 in F1, 0.02 in distinct score, and the improvement on the biased personalized test set of 5.56 in bilateral persona accuracy, 7.51 in BLEU, 2.12 in F1, 0.02 in distinct score. On the manual evaluations, the proposed method can generate more fluency, bilateral persona-consistent, and context-consistent responses compared with other state-of-the-art methods.

Keywords Bilateral persona-consistent · Conditional mutual information maximum · Contrastive learning · Personalized dialogue generation

1 Introduction

One of the major challenges in human–robot interaction is to develop an intelligent agent to generate natural, personalized, information-rich, and consistent responses (Adiwardana et al., 2020; Ma, 2021). For this purpose, the dialogue agents have to learn to express personalized information appropriately like humans. Currently, personalized dialogue agents have been widely applied in various human–robot interaction scenarios, such as intelligent personal assistants (Martin and Azvine 2003), public service robots (Tanaka et al. 1997), and wearable devices (Tramontano et al. 2019). The agents with personalization are considered reliable and

trustworthy and can gain the user’s confidence and trust (Roller et al. 2020).

In the past decades, personalization has played an important role in the dialogue system and attracted wide attention (Huang et al. 2020; Zhang et al. 2018a; Song et al. 2019; Qian et al. 2018; Zheng et al. 2019). According to the different ways of personalized information modeling, the existing personalized dialogue systems are mainly categorized into two types: implicit personalization (Zhang et al. 2018a; Song et al. 2019) and explicit personalization (Qian et al. 2018; Zheng et al. 2019).

The implicit personalized dialogue system models personas with unstructured natural language utterances (e.g., “I am a musician.”, “I like to play the guitar.”), where the persona is implicitly mapped into the hidden state vectors. However, these implicit space mapping methods are poor in interpretability and may be over-fitting during the training process. Besides, the given utterances are mostly short and limited to only a few personas; the model may fail to utilize the persona properly when generating responses (Xu et al. 2020). Indeed, implicit personalized dialogue corpus (Zhang

✉ Bin Li
libincn@hnu.edu.cn

Hanjun Deng
931446450@qq.com

¹ College of Electrical and Information Engineering,
Hunan University, Changsha, China

² Experimental High School Affiliated to Beijing Normal
University, Beijing, China

| General GPT2 model with unilateral persona | | | |
|--|-------------------|---------------------------------|------------------|
| User: Where are you from? | | | |
| Robot: I am from Guangzhou. | | | |
| User: Come to my city and play with me? | | | |
| Robot: Yes, I want to go to Guangzhou. | | | |
| Our model with bilateral personas | | | |
| User: Where are you from? | | | |
| Robot: I come from Guangzhou. | | | |
| User: Come to my city and play with me? | | | |
| Robot: Yes, I want to go to Beijing. | | | |
| User's personalized attributes | | Robot's personalized attributes | |
| "Gender" | "Female" | "Gender" | "Male" |
| "Area" | "Beijing" | "Area" | "Guangzhou" |
| "Interests" | "Shopping; Sport" | "Interests" | "Cooking; Movie" |

Fig. 1 Exemplar dialogues with/without bilateral persona-consistent in dyadic conversation. The general GPT2 model with unilateral persona can generate a response that only meets the robot's persona, but ignores the persona of the other party. The proposed method can incorporate bilateral personas and generate a response that matches the personas of both parties

et al. 2018a) reflecting personas in every response is also different from the way of interpersonal conversation.

The explicit personalized dialogue system models the personas with the structured personalized attributes, which are explicitly formatted as different key-value pairs (e.g., <Gender, Female>, <Area, Beijing>). Such explicit persona modeling is more straightforward and interpretable. Specifically, the explicit personalized dialogue corpora (Qian et al. 2018; Zheng et al. 2019) are crawled on a large scale from social networking sites, such as Weibo¹, where people may unintentionally show their personality during the conversation. However, the explicit personalization in Qian et al. (2018) models the robot's persona in the form of a pre-assigned profile and only emphasizes unilateral persona consistency. The latest works (Zheng et al. 2019, 2020) incorporate the structured speaker's profile into the generated response to ensure the persona consistency of the speaker. Although these methods solve the problem of unilateral persona consistency to some extent, the robot may ignore the user's persona during the conversation. As a result, the generated responses may conflict with the user's personalized information.

In the dyadic interpersonal conversations, both two interacting parties know each other's personalized information (Walker et al. 1997). When responding, the speaker should not only focus on their own personalized expression, but also consider the questions and persona of the other party (Isard et al. 2006). As shown in Fig. 1, during the conversation, the robot should generate responses consistent

with the robot's own personalized attributes (i.e., unilateral persona-consistent). Furthermore, the robot also should know the user's persona and generates responses consistent with the user's personalized attributes (i.e., bilateral persona-consistent). Once these factors are ignored, it may annoy the user and reduce the user experience.

To solve the above problem, we propose a bilateral personalized dialogue generation (BPDG) method to generate responses consistent with both personas. Specifically, the BPDG method is based on the structure of language model with multi-task transfer learning. The proposed method optimizes three tasks simultaneously: the language model task, the persona presence prediction task, and the dialogue generation task. In the language model task, the dialogue utterances embedded with the corresponded personalized attributes and relative position are used to train the encoder. In the persona presence prediction task, the dialogue contextual encoding is used to predict the possibilities of the personas' presence in the response. More precisely, the encodings of the dialogue context, bilateral personas and the right-shifted outputs are fused with a dynamic persona-aware fusion module to capture bilateral personas. In the dialogue generation task, the fused encoding is input into the decoder to generate response candidates with the diverse beam search strategy (Vijayakumar et al. 2016). Finally, in order to ensure the generated responses are more personalized and bilateral persona-consistent, we adopt the conditional mutual information maximum (CMIM) criterion with contrastive learning to select the final response from the diversified generated candidates. Thus, the proposed BPDG method can utilize bilateral personalized information to generate personalized and bilateral persona-consistent responses for better user experience in the human–robot interaction.

The main contributions of this article can be summarized as follows.

1. We propose a novel BPDG method, which integrates the bilateral personas to generate responses consistent with both personas. To the best of our knowledge, this is the very first to propose the bilateral persona consistency in the personalized dialogue generation.
2. A dynamic persona-aware fusion module is developed to adaptively control the encodings of the bilateral personalized information, the dialogue context, and the shifted right outputs for decoding to generate bilateral persona-consistent responses.
3. We adopt the criterion of the CMIM with contrastive learning, which bridges the gap between the learning objective and evaluation metrics.
4. Both automatic and manual evaluations show that our method outperforms state-of-the-art methods.

¹ <https://www.weibo.com>.

The remainder of this article is structured as follows: Section 2 reviews the work related to the personalized dialogue system. Section 3 formulates the problem and details the proposed BPDG method. Section 4 fully describes the experimental setups. Automatic and human evaluations are illustrated and analyzed in detail in Sects. 5 and 6, respectively. Finally, the conclusions and some possible future work are pointed out in Sect. 7.

2 Related Work

Our method focuses on the study of bilateral personalization in dyadic interpersonal conversation, where transfer learning and contrastive ranking are covered. In this section, we will elaborate on the related work including personalized dialogue generation, multi-task transfer learning, and contrastive learning.

2.1 Personalized Dialogue Generation

Inspiring by the “Big Five” (Goldberg 1993) in psychology, Mairesse and Walker (2007) take the lead in incorporating the personalities into the framework of dialogue generation, thereby generating responses with recognizable personality. However, the personality of the “Big Five” is extremely implicit and subtle. It is necessary to build rules to capture personality characteristics. Besides, it is a challenge to construct a corpus with limited and laborious collections. With the popularity of deep learning, handcraft rule modeling is gradually replaced by data-driven modeling. Li et al. (2016b) first propose a personalized dialogue generation model, mapping the persona in natural utterance into distributed representation vectors on the seq2seq framework, which is benefited from the neural machine translation (Sutskever et al. 2014). Subsequently, there are other different methods used for personalized dialogue generation modeling; for example, Song et al. (2019) adopt the CVAE method implicitly learns the responses that contain personalized information to generate personalized responses. Madotto et al. (2019) design a personalized dialogue generation model with meta-learning. Yang et al. (2020) describe an empirical survey of personalized dialogue generation via reinforcement learning. The above method is effective, but it also faces the problem of generating general or bilateral-inconsistent responses. Different from the previous work, the proposed BPDG method further integrates personalized information from both parties into the pre-trained decoder-only framework, to generate bilateral persona-consistent responses with multi-task learning and transfer learning.

2.2 Multi-task transfer learning

Multi-task transfer learning aims to extract and transfer the knowledge from the source domain to the target domain (Mo et al. 2016) with different well-designed learning tasks, which has been very popular in the field of the NLP in the past decade (Wang and Zheng 2015). Recent advances in natural language generation rely on pre-training a large generative language model with a large corpus of unsupervised data. It mainly follows the two-stage paradigm of pre-training and fine-tuning. In the field of personalized dialogue generation, Zhang et al. (2018) first introduce transfer learning into the two-stage personalized dialogue generation. Wolf et al. (2019) design a pre-trained dialogue generation model that jointly learns two tasks (e.g., next sentence prediction and language model) when fine-tuning. The experimental results show that multi-task learning can greatly improve the scores in automatic metrics. Golovanov et al. 2020 integrate multi-task learning into the transferred model with shared parameters and design three sub-tasks, including language model task, dialogue generation task, and expected risk task. These tasks are proven to improve the performance in human evaluation. Zheng et al. (2020) leverage target persona information in generating unilateral persona-consistent responses by designing three different tasks, including the language model, persona routing, and dialogue generation. In this article, apart from the language model task and dialogue generation task, we further design a persona prediction task for the dynamic persona-aware fusion module, adaptively fusing the encodings of different information for decoding, to generate responses consistent with bilateral personas.

2.3 Contrastive learning

Contrastive Learning (Chen et al. 2020; Gutmann and Hyvärinen 2012; Hadsell et al. 2006; Dash et al. 2021) uses self-supervised methods to learn the representation of the positive examples and negative examples. The contrastive learning method learns the general features of the corpus without labels by teaching the model which data are similar or different. In the field of natural language processing, contrasted learning has good performance in tasks such as language model task (Baltescu and Blunsom 2015), image captioning (Dai and Lin 2017), and text summarization (Liu and Liu 2021). In the field of human–robot interaction, contrastive learning is conducive to capturing the information implicit in the dialogue (Cai et al. 2020), and it is useful for filling the gap between learning objective function and evaluation metrics (Liu and Liu 2021). Therefore, this paper introduces the conditional mutual information criterion in the bilateral personalized dialogue generation. By ranking the diversified candidate responses through comparative learn-

ing, the final outputs can be rich in bilateral personalized information.

3 Proposed method

In the dyadic interpersonal conversation, both interacting parties have their own personas such as gender, area, and individual interests. Such information may be presented in the response. In the human–robot dialogue, given the user persona U , the robot persona R , the personalized history H , and the user input X , the robot generates a natural, fluent, and personalized response Y , which can be formulated as follows:

$$Y = \arg \max_{Y'} P(Y' | X, H, U, R) \tag{1}$$

where the user persona U and the robot persona R can be represented with the personal profile, which is formatted as a set of attributes composed of key-value pairs. Each attribute in the user persona $U = \{u_1, u_2, \dots, u_m\}$ is a key-value pair $u_i = \langle k_i, v_i \rangle$. The robot persona R is represented likewise. The personalized history is represented as $H = \{\{X_1^U, U\}, \{X_2^R, R\}, \dots, \{X_l^R, R\}\}$, where the superscript indicates the speaker, and the subscript indicates the number of the dialogue rounds. Each sentence is associated with the persona of the corresponded speaker. The user input $X = \{X_{l+1}^U, U\}$ contains the user current input X_{l+1}^U with the user persona U .

Combining the user input X and the personalized history H into the context of the dialogue C , Eq. (1) can be further written as Eq. (2):

$$Y = \arg \max_{Y'} P(Y' | C, U, R) \tag{2}$$

where the dialogue context $C = \langle H, X \rangle$ represents that the personalized history H is concatenated with the current user input X .

Figure 2 is the overview of the proposed BPDG method. The BPDG method consists of the encoder, the dynamic persona-aware fusion module, and the decoder. Following the GPT2 framework, the encoder and decoder share the same weights and act as a backbone to learn the sentence representation. The encoder trains the language model with the dialogue context embedding and encodes the embedding of the user persona and the robot persona independently. The persona-aware fusion module is used for fusing the dialogue context encoding, the bilateral persona encodings, and the shifted right outputs encoding. Afterward, the fused encoding is sent into the decoder for generating several candidate responses with the diverse beam search strategy. Finally, the CMIM criterion is adopted to output a personalized and bilateral persona-consistent response.

3.1 Dialogue context modeling

Dialogue context modeling means that each dialogue utterance embedding is added with the corresponded persona

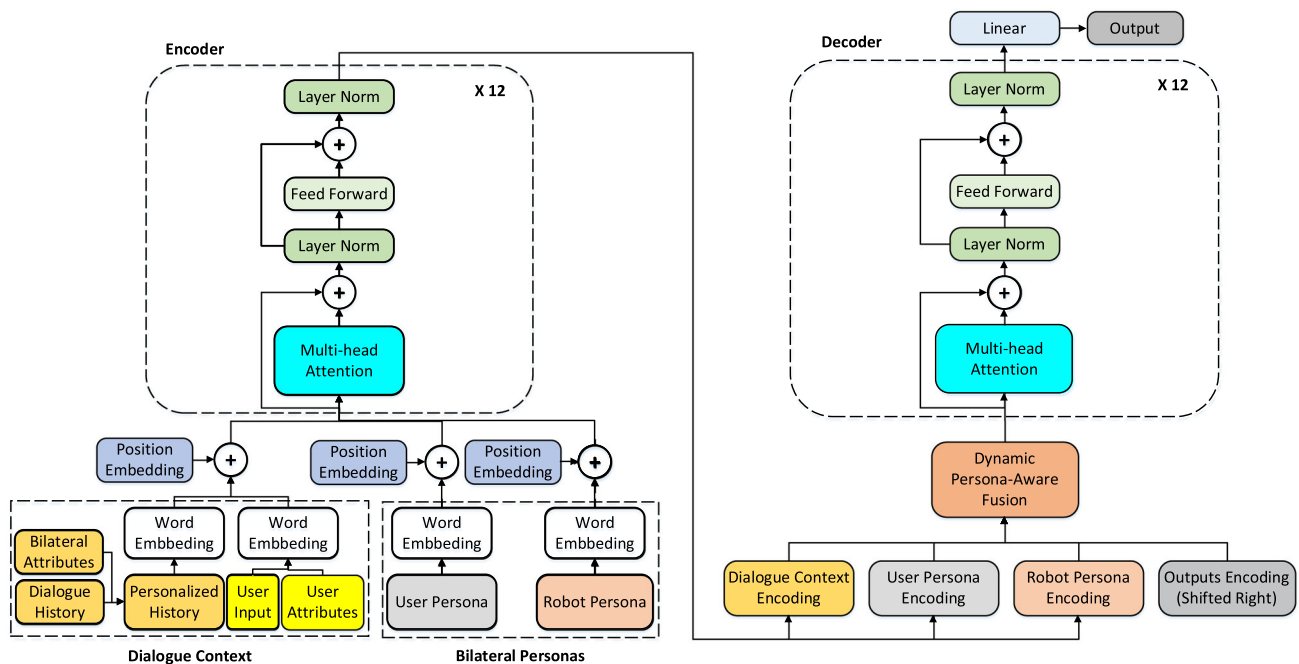


Fig. 2 The overview of the proposed BPDG method

embedding and relative position embedding to obtain the embeddings of personalized history. The dialogue context embedding can be obtained by concatenating the embeddings of the personalized history and the current user input. The dialogue context encoding is obtained with the dialogue context embedding being encoded. The process can be described as follows:

3.1.1 Utterance embedding

The utterances of the user and the robot are first embedded with word embedding, respectively. The X_U represents the embedded user input, and the X_R represents the embedded robot output. Both embeddings are specified with the same length n . If the corresponding length does not reach the specified length, we use $\langle PAD \rangle$ as a placeholder. Otherwise, a truncation operation is taken. The word embedding process is shown as follows:

$$X_U = \{x_1^U, x_2^U, x_3^U, \dots, x_n^U\} \tag{3}$$

$$X_R = \{x_1^R, x_2^R, x_3^R, \dots, x_n^R\} \tag{4}$$

where the X_U is the embedding of the user input, the x_i^U is the word embedding of the i -th token in the sentence input by the user, and the X_R is the embedding of the robot response, the x_i^R is the word embedding of the i -th token in the sentence output by the robot.

3.1.2 Persona embedding

Persona embedding means the utterances embedded with the corresponded personas attributes. As is mentioned before, the profile consists of three attributes: gender, area, and individual interests. The value of the gender is binary (i.e., 0 for male and 1 for female). The value of the area is represented with the index of the corresponded item in the look-up table. The items of the look-up table are sorted by the occurrence frequency of the area in the corpus. The individual interests are represented in a similar way. To take the operation of the user as an example, the process is shown in Eq. (5):

$$\begin{aligned} G_U &= \{g_1^U, g_2^U, \dots, g_j^U, \dots, g_n^U \mid g_j^U = g^U\} \\ A_U &= \{a_1^U, a_2^U, \dots, a_j^U, \dots, a_n^U \mid a_j^U = a^U\} \\ T_U &= \{t_1^U, t_2^U, \dots, t_j^U, \dots, t_n^U \mid t_j^U = t^U\} \end{aligned} \tag{5}$$

where the g^U represents the word embedding of the user's gender extracted from the profile, the g_j^U represents the gender embedding g^U corresponding to the position j in the user input embedding X_U , $j \in [1, n]$. The a^U and t^U represent

the word embedding of the user's area and individual interests tag extracted from the profile, respectively. For multiple individual interests, we take the average of the first-three embeddings of individual interests.

The relative position embedding (Vaswani et al. 2017) is adopted to make the embedded tokens more sensitive to the position in the sentence for further attention operation. The position embedding is written as follows:

$$\begin{aligned} E_i(2k) &= \sin\left(\frac{i}{10000^{\frac{2k}{d_{model}}}}\right) \\ E_i(2k+1) &= \cos\left(\frac{i}{10000^{\frac{2k}{d_{model}}}}\right) \end{aligned} \tag{6}$$

where i is the position of the token in the sentence, k represents the k -th dimension of the word embedding, d_{model} is the fixed embedding dimension.

3.1.3 Personalized history embeddings

Figure 3 shows the structure of personalized history embeddings. The personalized history embeddings are a combination of the aforementioned three types of embeddings, i.e., the embeddings of the utterance, the persona embeddings, and the position embeddings, with the $\langle SEP \rangle$ being used as the separator. Specifically, the personalized history embeddings are formatted utterance by utterance with concatenation, which can be written as Eq. (7).

$$\begin{aligned} H &= \text{Concat} \{h_1, h_2, \dots, h_j, \dots, h_l\} \\ h_j &= \begin{cases} D_U, & \text{if } \text{mod}(j, 2) = 0 \\ D_R, & \text{if } \text{mod}(j, 2) = 1 \end{cases}, j \in [1, l] \end{aligned} \tag{7}$$

where the $\text{Concat} \{\}$ represents the operation of concatenation, l represents the total number of rounds of the personalized history, h_j represents the personalized history of the j round, $j \in [1, l]$. For each utterance, the personalized history embeddings are calculated via aligning the embed-

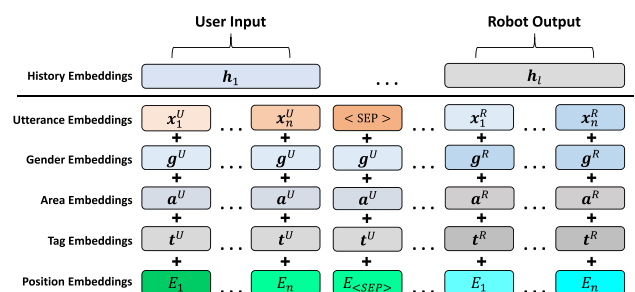


Fig. 3 The structure of personalized history embeddings

dings by token and performing token-wise aggregation. This process can be expressed as follows:

$$D_U = \text{Add}(X_U, G_U, A_U, T_U, E) \tag{8}$$

$$D_R = \text{Add}(X_R, G_R, A_R, T_R, E) \tag{9}$$

where the $\text{Add}()$ represents the token-wise addition operation of the different embeddings with the same embedded length.

3.1.4 Dialogue context embedding

The personalized history embeddings and the user current input at the $l + 1$ round are concatenated into the dialogue context embedding C , which can be expressed as follows:

$$C = \text{Concat}\{H, h_{l+1}\} \tag{10}$$

Finally, the dialogue context encoding E_C is obtained after the dialogue context embedding C is encoded.

3.2 Bilateral profile modeling

To take advantage of the bilateral personas in the dialogue generation, the explicit form of persona, i.e., the profile, is used in the proposed method. Word embedding is performed on the profile text to represent the semantic information in the same way as the utterance, which will benefit the further processing. Specifically, the word embedding of the user persona U and the robot persona R can be written as follows:

$$U = \{u_1, u_2, u_3 \mid u_i = \{s, v\}, i = 1, 2, 3\} \tag{11}$$

$$R = \{r_1, r_2, r_3 \mid r_i = \{s', v'\}, i = 1, 2, 3\} \tag{12}$$

where each attribute u_i in the embedded user persona U is the word embedding of the key-value pair. The embedded user persona U is the concatenation of the three attributes corresponding to gender, area, and individual interests, respectively. The comma is used as the separator to concatenate each key-value pair. The embedded robot persona R is formatted likewise.

Further, the embedded user persona U with relative position embedding E is input into the encoder to obtain the user persona encoding E_U , while the embedded robot persona R turns into the E_R that is in the same way. The above process is implemented independently, which means that the E_U and E_R do not participate in the training of the encoder.

3.3 Persona-aware fusion module

In the bilateral personalized dialogue generation, two critical problems have to be addressed for appropriate persona

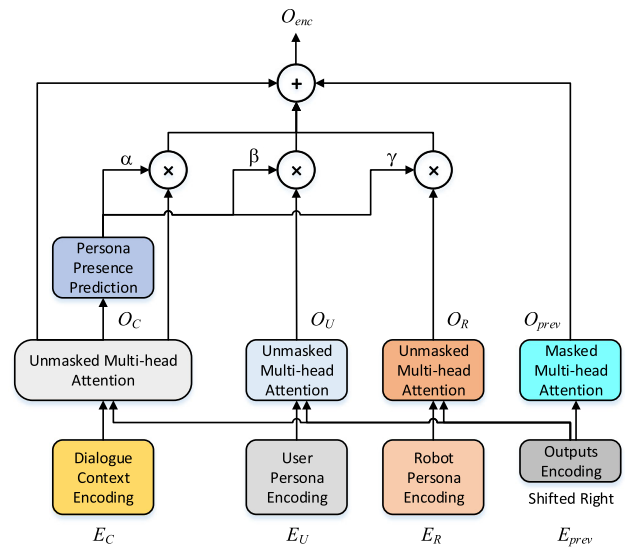


Fig. 4 The structure of the dynamic persona-aware fusion module

expression: (1) when to express persona and (2) whose persona should be expressed. Therefore, we propose dynamic persona-aware fusion to predict the presence of the bilateral personas and adaptively fuse them into the encodings for the further personalized response generation. Figure 4 shows the structure of the dynamic persona-aware fusion module. The persona-aware means that the presence of the persona in the generated response can be predicted with the dialogue contextual encoding O_C obtained from the attention operation. The prediction probability is used to dynamically weighted to the corresponded attention encoding for fusion.

3.3.1 Encoding attention mechanism

In order to effectively utilize the information of the encodings, we design different encoding attention mechanisms. Each encoding from the encoder participates in the unmasked multi-head attention mechanism. The masked multi-head attention mechanism is designed to avoid feeding the shifted-right ground-truth tokens when training. The $prev$ represents the previously decoded output word, which turns into the outputs encoding E_{prev} with word embedding and position embedding. The E_U is input into the unmasked multi-head attention network to obtain the user personalized encoding O_U and the robot personalized encoding O_R that is obtained in the same way. The unmasked multi-head attention process is shown as follows:

$$O_U = \text{Multi-head}(E_{prev}, E_U, E_U) \tag{13}$$

$$O_R = \text{Multi-head}(E_{prev}, E_R, E_R) \tag{14}$$

where the E_{prev} is the query, the E_U is both the key and the value in the unmasked multi-head mechanism, and the operation of the robot personalized encoding O_R is the same.

The context encoding E_C and the outputs encoding at the previous moment E_{prev} are used to obtain the dialogue contextual encoding O_C with the unmasked multi-head attention network, as shown in Eq. (15):

$$O_C = \text{Multi-head}(E_{prev}, E_C, E_C) \tag{15}$$

where the E_{prev} is the query; the E_C is both the key and the value in the unmasked multi-head mechanism.

Furthermore, a masked multi-head attention network is used to obtain the previous outputs encodings O_{prev} , as shown in Eq. (16):

$$O_{prev} = \text{MaskedMulti-Head}(E_{prev}, E_{prev}, E_{prev}) \tag{16}$$

where the E_{prev} is the query, the key, and the value in the masked multi-head mechanism.

3.3.2 Persona presence prediction

The presence of the bilateral personas in the response is predicted for the dynamic persona-aware fusion of different encodings. To train a subnetwork for this task, we construct a heuristic script to label the utterance with three labels based on the presence of bilateral personas. The dialogue contextual encoding O_C is used to predict the probability of three types of information, which is presented in the response sentence. The loss function is designed as follows:

$$L_P(\theta) = - \sum_{j=1}^3 l_j \log P_\theta(l = j | O_C) \tag{17}$$

where the l_j represents the label of different persona type, $\log P_\theta(l = j | O_C)$ represents the probability of the persona type predicted in the generated response based on the dialogue contextual encoding O_C .

3.3.3 Persona encoding fusion

To utilize personalized information of different encodings, the dynamic encoding fusion is designed to adaptively control the persona presented in the generated response. The probability of three categories is used as the persona-aware weight for dynamic encoding fusion. Specifically, each category is operated with the *softmax* operation, which can be shown in Eq. (18):

$$P_\theta(l = j | O_C) = \frac{\exp O_C^{(j)}}{\sum_{i=0}^2 \exp O_C^{(i)}}, j = 0, 1, 2 \tag{18}$$

where the $O_C^{(j)}$ represents the dialogue contextual encoding O_C corresponding to the j -th label, which is obtained with a two-layer perception network with global and average pooling.

Each prediction probability is defined as the persona-aware weight as follows:

$$\alpha = P_\theta(l = 0 | O_C) \tag{19}$$

$$\beta = P_\theta(l = 1 | O_C) \tag{20}$$

$$\gamma = P_\theta(l = 2 | O_C) \tag{21}$$

where the α represents the probability of the user personalized information presented in the response, the β represents the probability of the robot personalized information presented in the response, and the γ represents the probability that personalized information does not present in the response, which means the context-related. Three different encodings are dynamically weighted and fused, with the dialogue contextual encoding O_C and the previous outputs encodings O_{prev} . These encodings together form the fused encoding O_{enc} , as shown in Eq. (22):

$$O_{enc} = \alpha O_U + \beta O_R + (\gamma + 1) O_C + O_{prev} \tag{22}$$

where $\alpha + \beta + \gamma = 1$.

After fusing the different encodings with the dynamic persona-aware fusion module, the fused encoding O_{enc} is input into the decoder for dialogue generation.

3.4 Multi-task learning for dialogue generation

To train the proposed BPDG model, three different tasks have to be accomplished including language model task, persona prediction task and dialogue generation task. These tasks will be described below.

3.4.1 Language model task

A pre-trained model is first utilized to initialize the parameters of the GPT2 framework. In order to bridge the gap between the data utilized in the pre-training and fine-tuning stage, the language model is then adopted to fine-tune with the bilateral personalized dialogue dataset mentioned in Section IV-A. The language model is trained by optimizing the standard maximum log-likelihood loss, as shown in Eq. (23):

$$L_{LM}(\varphi) = - \sum_i \log P_\varphi(x_i | x_{i-k}, \dots, x_{i-1}) \tag{23}$$

where φ represents the parameters of language model, k is the size of the context window, and $x_{i-k}, \dots, x_{i-1}, x_i$ is sequence of tokens sampled from the training corpus.

3.4.2 Persona prediction task

The persona prediction task is to predict the persona presence according to the contextual encoding O_C . The loss function is shown in Eq. (17). As a result, the prediction probability is used to dynamically weighted the different encodings to get the fused encoding O_{enc} . Finally, the O_{enc} is input into the decoder for dialogue generation.

3.4.3 Dialogue generation task

The dialogue generation task is designed to generate the bilateral personalized responses; the loss function of the dialogue generation task is shown as Eq. (24):

$$\begin{aligned}
 L_D(\varphi) &= - \sum_i \log P_\varphi(y_i | y_0, \dots, y_{i-1}, E_C, E_U, E_R) \\
 &= - \sum_i \log P_\varphi(y_i | O_{enc})
 \end{aligned}
 \tag{24}$$

where y_i represents the i -th word generated by the decoder, and y_0, \dots, y_{i-1} is a sequence of previously generated words. Identically, the input of the decoder also can be written as the fused encoding.

Finally, the joint loss function of the entire model is presented in Eq. (25):

$$L(\varphi, \theta) = L_D(\varphi) + \lambda_1 L_{LM}(\varphi) + \lambda_2 L_P(\theta)
 \tag{25}$$

where the λ_1 and λ_2 are the balance weights of the loss function of the language model task and the loss function of the persona prediction task, respectively.

3.5 Candidate selection with CMIM

After the dialogue generation via dynamic persona-aware fusion, the response is output with the decoding strategy. However, the top-ranked candidates with the beam search strategy are usually general, short, or even unrelated (Kulikov et al. 2019), so that responses related to both personas and history conditions often fail to achieve high ranking scores. To remedy this, the criterion of CMIM (Fleuret 2004) is adopted to constrain the personalized and history information that reflects in the response. Specifically, the BPDG method utilizes the diverse beam search strategy to generate the best diversified top-20 candidate list and adopts the CMIM criterion to select the response with the largest conditional mutual information value as the final response.

3.5.1 Conditional mutual information modeling

In order to simplify the modeling process, the user persona U , the robot persona R , and personalized history information




Fig. 5 Illustration of conditional mutual information. The circles represent the information entropies of the different variables. The dashed circle represents the information entropy of the generated responses

H can be regarded as the condition Z . The illustration of conditional mutual information is shown in Fig. 5. Given the different conditions, i.e., H, U, R in the same dialogue, the value of conditional mutual information $CMIM_v$ of the user input X and the robot-generated candidate response Y_i can be expressed as Eq. (26):

$$\begin{aligned}
 CMIM_v(Y_i) &\equiv I(Y_i; X | H, U, R) \\
 &= \underbrace{H(Y_i | H, U, R)}_{\text{Relevance Ranking}} - \underbrace{H(Y_i | X, H, U, R)}_{\text{Dialogue Generation}}
 \end{aligned}
 \tag{26}$$

where the CMIM criterion can be modeled with two terms, i.e., the dialogue generation item and the relevance ranking item.

According to the definition of the CMI (Fleuret 2004), the maximum of Eq. (26) can be achieved by solving the following optimization problem:

$$Y^* = \arg \max_{Y_i} \log \frac{P(Y_i | X, Z)}{P(Y_i | Z)}
 \tag{27}$$

where the Y^* represents the final response in the top-20 candidate list. The $P(Y_i|X, Z)$ and $P(Y_i|Z)$ are corresponded to the dialogue generation term and relevance ranking term in Eq. (26), respectively.

The $P(Y_i|X, Z)$ is the probability of the generated response conditioned on the input and the context with the word granularity, while the $P(Y_i|Z)$ is the relevance of the response to the contextual content with the sentence granu-

larity. Therefore, the $P(Y_i|X, Z)$ and $P(Y_i|Z)$ of Eq. (27) are not optimized jointly.

3.5.2 Dialogue generation

The $P(Y_i|X, Z)$ can be modeled with the BPDG model and calculated with the diversified beam search score. By substituting the Z with H, U, R , the $P(Y_i|X, H, U, R)$ can be written as Eq. (28):

$$\begin{aligned} \log P(Y_i | X, H, U, R) &\equiv \log P_\psi(Y_i | X, E_H, E_U, E_R) \\ &= \log P_\psi(Y_i | O_{enc}) \end{aligned} \quad (28)$$

where the ψ represents the parameters of the trained BPDG model, containing all the parameters in (25).

3.5.3 Relevance ranking with contrastive learning

After the candidate list is generated with the diverse beam search strategy, each candidate can be ranked with relevance ranking. Given the condition Z , i.e., the user persona U , the robot persona R , and the personalized history H , the relevance probability is calculated as:

$$\begin{aligned} \log P(Y_i | H, U, R) &= \frac{\log P_\phi(Y_i, H, U, R)}{\log P_\phi(H, U, R)} \\ &\propto \log P_\phi(Y_i, H, U, R) \end{aligned} \quad (29)$$

where the ϕ represents the parameters of the content relevance classifier model trained on the corpus, the co-occurrence probability $P_\phi(H, U, R)$ is not related to Y_i , which can be omitted, and the $P_\phi(Y_i, H, U, R)$ represents the co-occurrence probability of Y_i, H, U and R in the same dialogue.

Therefore, the relevance probability of each candidate can be modeled with the content relevance classifier $P_\phi(Y_i, H, U, R)$; we adopt the contrastive learning training method (Cai et al. 2020) to perform the relevance ranking step. To construct the training corpus for content relevance classifier, the Y, H, U , and R from the corpus \mathbb{D} are used as positive training samples, which has marked, while the Y', H, U , and R from different corpus are sampled as negative samples, which is inspired by the practice in Lan et al. (2019). The cross-entropy loss function used to train content relevance classifier ϕ is as follows:

$$\begin{aligned} L_\phi &= -\frac{1}{N} \sum_{(Y, H, U, R) \in \mathbb{D}} \log P((Y, H, U, R)^+; \phi) \\ &\quad - \frac{1}{N} \sum_{(Y', H, U, R) \in \mathbb{D}} \log [1 - P(Y', H, U, R)^-; \phi]. \end{aligned} \quad (30)$$

3.5.4 Candidate selection

With the BPDG model and the content relevance classifier, the optimization problem in (27) can be written as follows:

$$Y^* = \arg \max_{Y_i} \log \frac{P_\psi(Y_i | O_{enc})}{P_\phi(Y_i, H, U, R)} \quad (31)$$

where Y_i represents the response candidates.

Thus, the calculation of the response candidates can be selected by Eq. (32):

$$Y^* = \arg \max_{Y_i} \log P_\psi(Y_i | O_{enc}) - \log P_\phi(Y_i, H, U, R) \quad (32)$$

4 Experiments

In this section, we will introduce the experimental data sets and elaborate on the bilateral persona and content relevance classifiers. Moreover, we will introduce the implementation and compared methods in detail for further experiments.

4.1 Data set description

To evaluate the effectiveness of the BPDG method, extensive experiments are conducted based on the PersonalDialog dataset (Zheng et al. 2019). This corpus contains sparse personas of multi-party, where the personalized responses in dyadic dialogues involve bilateral personalized information. It is very challenging to choose which persona to generate, so we pick dyadic dialogues from the original corpus for our research. This dataset provides personalized profiles of both speakers, including three personal attributes, i.e., ‘‘Gender’’, ‘‘Area’’ and ‘‘Individual interests’’.

Since, in some cases of the original corpus, the personalized profiles are missing, we construct a heuristic script to select the data with complete personalized information of both parties. The constructed dialogue dataset is referred to as the bilateral personalized dataset in this article. The bilateral personalized dataset consists of 410K dialogues in total, where 400K is randomly sampled as the training set, and the rest 10K data as the validation set. The average length of each dialogue is about 3.5 rounds, and the average length of each sentence is about 7.45 characters.

The evaluation settings of the ECDT² are adopted, to test the performance of different methods in different contexts. Specifically, two test sets³ (i.e., a random set and a biased set) are constructed for the evaluation. The random set is a collection of dialogues between both parties, most of which

² <http://conference.cipscc.org.cn/smp2019/evaluation.html>.

³ <https://worksheets.codalab.org/worksheets/0x8f68b61a8b2249d7b314c6e800e2dace>.

do not contain personas. It is constructed for testing the performance of different methods in a context where the two interacting parties do not intentionally show their personas. The biased test set is the dialogue set between both parties with personalized information, where the speaker tends to reveal personalized information about both parties during the conversation. It is used for testing the performance of different methods in the context where the speakers intentionally show their personas.

4.2 Bilateral persona classifier

To better evaluate whether the response is bilateral persona-consistent or not, we design the bilateral persona classifier P_π as an objective metric, which is trained with the aforementioned personalized labels. Each sentence is labeled with one of the three labels: 0 for the sentence related to the persona of the user, 1 for the sentence related to the persona of the robot, and 2 for the sentence that does not contain the persona.

The bilateral persona classifier is used to evaluate whether the response Y contains the user persona U or the robot persona R . To calculate each probability of the respective category, the response Y with bilateral personas is concatenated with $\langle SEP \rangle$. After calculating each probability, the probability of category 0 and category 1 is added together as the probability of the bilateral personalized response. About 10K rounds of dialogues containing bilateral personas are randomly sampled from the bilateral personalized dataset, where the category ratio is 1:1:3. Then, we divide the above data into training, validation, and test sets at a ratio of 8:1:1 to train the bilateral persona classifier. The accuracy of the classifier on the test set reaches 90.2% in a fivefold cross-validation setting.

4.3 Content relevance classifier

The content relevance classifier is used for ranking the candidates under the criterion of the CMIM with contrastive learning. After the candidate list is generated by the BPDG model, we calculate the content relevance probability of each generated response co-occurring in the current dialogues under the conditions of the personalized history H , the user persona U , and the robot persona R . These conditions and each generated response are concatenated with $\langle SEP \rangle$ for calculating the content relevance probability. After the probability of each generated response is calculated, the final response is selected to output. Specifically, the content relevance classifier is trained on the bilateral personalized dataset, using the ERNIE-base model (Sun et al. 2019) to fine-tune in the labeled dialogues. In the fivefold cross-validation setting, the accuracy reaches 80.4%.

4.4 Implementation details

We implement all the experiments of the bilateral personalized dialogue generation with the pre-train model called LCCC-base (Wang et al. 2020), which is a Chinese pre-trained model based on the GPT2 framework with a vocab of 13088 characters, is used to initialize the parameters of the encoder and decoder with transfer learning. According to He et al. (2016), the shared weights of the encoder and decoder are adopted in this article, as it is beneficial for improving the quality of generated responses. The encoder and decoder include 12 transformer blocks, among which the self-attention heads are 12. The size of the token embedding is 768 and the context window is 512. The parameter $d_{model} = 512$, $n = 64$, $\lambda_1 = 0.2$, $\lambda_2 = 0.5$. The diverse beam search strategy adopted in the proposed method is to generate the candidate list with the BPDG model, where the beam size is set to 20 and the group size is set to 4. The content relevance classifier is to calculate the relevance probability for each sentence in the candidate list under the criterion of the CMIM. The final generated response Y^* is selected to output. The BPDG model is fine-tuned directly on the bilateral personalized dataset for 30 epochs, where the batch size is 64 with gradient accumulation, using the Noam optimization scheduler (Rush 2018) with 2000 warm-up steps on two NVIDIA 2080 Ti GPUs. All the experimental codes are released at <https://github.com/Lirea-nstar/BPDG>⁴.

4.5 Compared methods

Several state-of-the-art baseline methods are compared with ours. These methods are described below:

1. *S2S + Atten.*: This method applies a three-layer Bi-GRU to project the input text into embeddings with a fixed size. Another three-layer GRU utilizes an attention mechanism (Luong et al. 2015) for response generation. The word embedding parameters of encoder and decoder are initialized by the pre-trained word vector.⁵ The parameter weights of the GRU network are initialized with a uniform distribution $[-0.05, 0.05]$. The model is optimized by implementing the Adam optimization scheduler.
2. *Trans.*: The Trans. employs the original transformer (Vaswani et al. 2017) using the self-attention mechanism to generate responses. The model is initialized with the uniform distribution $[-0.02, 0.02]$ and takes the concatenated dialogue history as input without using personas. We optimize the model by implementing the Noam (Rush 2018) optimization scheduler.

⁴ Code and data will be publicly available.

⁵ <https://github.com/Embedding/Chinese-Word-Vectors>.

3. *TTransfo.*: The *TTransfo.* is introduced by Wolf et al. (2019) optimizing a multi-task object for training. This model is initialized by the LCCC-base pre-trained model and fine-tunes directly on the bilateral personalized dataset only with the concatenated history. The Norm optimization scheduler is used for training the model with gradient accumulation (with batch size 64).
4. *LConv.*: The *LConv.* represents the multi-input model proposed in Golovanov et al. (2019). This model is initialized with the LCCC-base pre-trained model, which shares the parameters of the encoder and decoder. The model fine-tunes directly on the bilateral personalized dataset with the concatenated dialogue history. The Norm optimization scheduler is used for training the model with gradient accumulation (with batch size 64).
5. *TTransfo.+P*: It extends the *TTransfo.* by incorporating the speaker's persona. When fine-tuning, the contextual dialogues concatenated with the speaker's personalized information are input into the model. The Norm optimization scheduler is implemented for training, where the batch size is set to 64 with gradient accumulation.
6. *LConv.+P*: It extends the *LConv.* by incorporating the speaker's persona. When fine-tuning the contextual dialogues concatenated with the speaker's personalized information are input into the model. The Norm optimization scheduler is implemented for training, where the batch size is set to 64 with gradient accumulation.
7. *PWDWP*: The *PWDWP* (Zheng et al. 2020) is initialized by the LCCC-base pre-trained model and fine-tunes on the bilateral personalized dataset. This model incorporates personalized attributes embedding in the dialogue context for each speaker and devises a persona routing to weigh the persona-related encodings that are input into the decoder. The Norm optimization scheduler is implemented for training, where the batch size is set to 64 with gradient accumulation. This model is the strong baseline method in the explicit personalized dialogue system.

5 Automatic evaluation

In order to fully evaluate the effectiveness of the proposed method compared with the baseline methods, we choose various metrics for the automatic evaluation. In this section, we introduce these metrics and give a detailed analysis of the results.

5.1 Objective metrics introduction

(1) Bi-Persona Acc

The Bi-Persona Acc (BPAcc) is used to measure the degree of personalization in the response. We extend the unilateral persona-consistent (Zheng et al. 2020), which represents that

the persona is consistent with the speaker, to the bilateral persona-consistent. The Bi-Persona Acc represents the bilateral persona classification accuracy of the sentence, which is not only consistent with the speaker's persona but also with the persona of the other party. Each generated response and the bilateral personas are input into the bilateral persona classifier to obtain the Bi-Persona Acc. Therefore, we add the user and robot persona classification accuracy together to obtain the possibility of the response that contains bilateral personalized information. The higher Bi-Persona Acc score means that the generated response is more personalized and more likely to be bilateral persona-consistent.

$$BPAcc = \frac{P_{\pi}(1) + P_{\pi}(2)}{P_{\pi}(0) + P_{\pi}(1) + P_{\pi}(2)} \tag{33}$$

(2) BLEU

The BLEU (bilingual evaluation understudy) (Papineni et al. 2002) is utilized to evaluate the quality of the text in translation. In dialogue generation, the BLEU is calculated with the weighted n-gram overlap between the ground-truth response \hat{Y} and generated responses Y^* . The n-gram calculation is shown in Eq. (34):

$$P_n(\hat{Y}, Y^*) = \frac{\sum_k \min(\text{Cnt}_{\text{clip}}(k, \hat{Y}), \text{Cnt}_{\text{clip}}(k, Y^*))}{\sum_k \text{Cnt}(k, \hat{Y})} \tag{34}$$

where k traverses all the n-grams candidates, the \hat{Y} and the Y^* represent the ground-truth response and the generated response, respectively, $\text{Cnt}_{\text{clip}}(k, Y^*)$ represents the clipped n-grams number in the generated response Y^* , $\text{Cnt}(k, \hat{Y})$ represents n-grams number in the ground-truth response \hat{Y} . The weight $BP(\hat{Y}, Y^*)$ can be calculated as Eq. (35):

$$BP(\hat{Y}, Y^*) = \begin{cases} 1, & \text{if } |Y^*| > |\hat{Y}| \\ e^{(1-|\hat{Y}|/|Y^*|)}, & \text{if } |Y^*| \leq |\hat{Y}| \end{cases} \tag{35}$$

where $|Y^*|$ represents the length of the generated response, $|\hat{Y}|$ represents the length of the ground-truth response. The BLEU is calculated as follows:

$$BLEU = BP(\hat{Y}, Y^*) \cdot \exp\left(\sum_{n=1}^N w_n \log P_n(\hat{Y}, Y^*)\right) \tag{36}$$

where N is set to 2 and the weighted factor w_n is set to $1/N$, the percentile fraction we use is set to 1000, which is the same settings as the NLTK⁶. The higher the BLEU score, the better the quality of the generated response.

(3) F1

The F1 score is implemented to measure the accuracy of the model on the data set compared to the ground truth, which

⁶ <https://github.com/nltk/nltk>.

includes two parts: precision and recall. The precision means the proportion of words in the generated response contained in the ground-truth response, and the recall means the proportion of words in the ground-truth response contained in the generated response. The calculation of F1 score is the same as Dinan et al. (2019) and can be written as Eq. (37):

$$F1 = 2 \times \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \tag{37}$$

(4) Distinct

The Distinct (Li et al. 2016a) is adopted to measure the average score of the sum of unique unigrams and bigrams contained in the generated responses, which is divided by the total number of generated words. The equation can be written as follows:

$$\text{Distinct} = \frac{1}{2} \times \frac{\text{Cnt}(U_{uni}) + \text{Cnt}(U_{bi})}{\text{Num}_{tokens}} \tag{38}$$

where the $\text{Cnt}(U_{uni})$ represents the number of unigrams that are not repeated in the response compared with the ground-truth response, the Num_{tokens} represents the total number of generated words, the higher the distinct score, the more specific and diverse the response generated.

(5) PPL

The PPL (perplexity) (Huang et al. 2020a) is widely used to measure the performance that the model predicts different utterances in the test set. For the ground-truth response $\hat{Y} = \{y_1, y_2, \dots, y_m\}$, the perplexity is calculated by the trained model and can be calculated as Eq. (39):

$$\text{Perplexity} = \exp\left(-\frac{1}{N} \sum_{i=1}^m t_i\right) \tag{39}$$

$$t_i = \begin{cases} \log P(y_i) + \varepsilon, & \text{if } y_i \in F \\ \log(P(\text{unk})/|R|) + \varepsilon, & \text{if } y_i \in R \end{cases} \tag{40}$$

where the F represents the set of words in the frequent vocabulary and the R represents the set of words that are in the rare vocabulary, $P(\text{unk})$ represents the logits of unknown token predicted by the model. $|R|$ is the number of words that are in the rare vocabulary, the ε is set to 10^{-8} , which is used to ensure that logits are not zero.

5.2 Results and analysis

Tables 1 and 2, respectively, show the comparison results of the proposed method and different baseline methods on five metrics, and also present the performance of our method with different persona-aware weights. It can be seen from the results that, compared with the baseline methods, our method is superior to all metrics except the PPL. Noted that the ppl. score is inconsistent with Zheng et al. (2020), because they have used external personalized corpus for pre-training,

Table 1 Experimental results of five metrics on random test set

| Method | BPAcc | BLEU | F1 | Distinct | PPL |
|--------------------|--------------|--------------|--------------|--------------|--------------|
| S2S + Atten. | 9.71 | 3.18 | 4.32 | 0.098 | 87.41 |
| Trans. | 11.84 | 3.27 | 7.48 | 0.187 | 77.55 |
| TTransfo. | 15.43 | 6.33 | 10.40 | 0.255 | 47.48 |
| TTransfo.+P | 44.68 | 6.27 | 10.11 | 0.251 | 52.14 |
| LConv. | 12.41 | 3.38 | 7.52 | 0.243 | 51.32 |
| LConv.+P | 51.36 | 6.02 | 8.81 | 0.234 | 55.04 |
| PWDWP | 63.13 | 10.56 | 10.47 | 0.280 | 60.33 |
| Ours | 64.63 | 11.66 | 11.30 | 0.282 | 61.93 |
| Ours, $\alpha = 1$ | 87.12 | 10.64 | 10.97 | 0.279 | 60.55 |
| Ours, $\beta = 1$ | 82.42 | 10.32 | 10.52 | 0.274 | 60.12 |
| Ours, $\gamma = 1$ | 52.11 | 8.34 | 10.22 | 0.272 | 59.92 |

Table 2 Experimental results of five metrics on biased test set

| Method | BPAcc | BLEU | F1 | Distinct | PPL |
|--------------------|--------------|--------------|--------------|--------------|--------------|
| S2S + Atten. | 22.42 | 4.32 | 4.71 | 0.116 | 97.53 |
| Trans. | 29.14 | 6.89 | 8.07 | 0.251 | 79.23 |
| TTransfo. | 43.12 | 15.22 | 15.93 | 0.268 | 49.59 |
| TTransfo.+P | 60.25 | 15.55 | 16.27 | 0.260 | 57.63 |
| LConv. | 42.73 | 10.03 | 9.09 | 0.270 | 56.41 |
| LConv.+P | 57.43 | 12.13 | 10.02 | 0.263 | 58.74 |
| PWDWP | 88.19 | 17.21 | 16.40 | 0.285 | 60.81 |
| Ours | 92.14 | 24.64 | 18.05 | 0.287 | 58.59 |
| Ours, $\alpha = 1$ | 93.75 | 23.51 | 17.97 | 0.286 | 61.77 |
| Ours, $\beta = 1$ | 91.12 | 24.72 | 18.52 | 0.285 | 60.11 |
| Ours, $\gamma = 1$ | 57.62 | 17.51 | 17.01 | 0.282 | 62.66 |

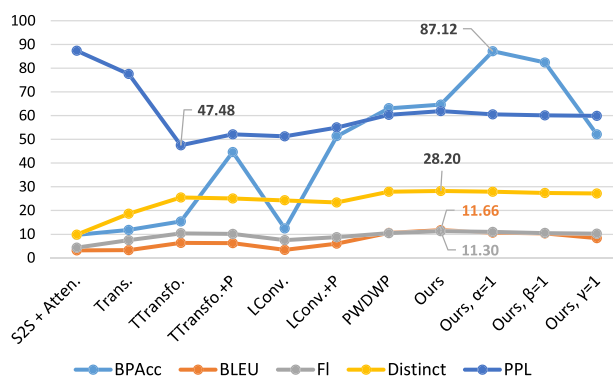


Fig. 6 Visualization of the results of five metrics on random test set, where the distinct score is multiplied with 100 to facilitate reading

while this pre-training corpus is not open source. Tables 1 and 2 show that we have used the open-source LCCC model for initializing all the baseline models.

The visualization of both random and biased test sets is shown in Figs. 6 and 7. Under the same experimental conditions, further conclusions are that: (1) under the same

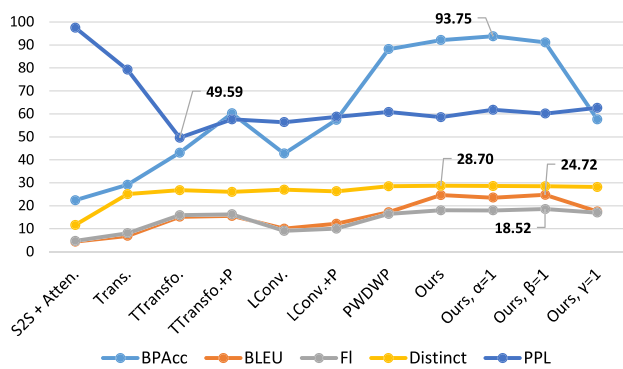


Fig. 7 Visualization of the results of five metrics on biased test set, where the distinct score is multiplied with 100 to facilitate reading

automatic weighting setting, our method is better than the strong baseline method (i.e., PWD-WP). On the random set, it outperforms with 1.5% in BPAcc, 1.1% in BLEU, 0.83% in F1, and 0.2% in Distinct. While on the biased set, our method outperforms with 3.95% in BPAcc, 7.43% in BLEU, 1.65% in F1, and 0.2% in Distinct. Especially on the biased set, our method is superior to the compared baseline methods. This shows that our method can generate more personalized and better responses. (2) It can be found that both in Tables 1 and 2 the PPL scores in bold (i.e., 47.48 and 49.59) show that the best results of the PPL appear on the TTransfo, which is the method without incorporating the personalized information. However, the methods with personalized information (i.e., TTransfo.+P, LConv.+P, PWDWP, and our method) all obtain the higher PPL score. This indicates that generating responses with personalized information will hurt the PPL score. It occurs because the words involving the persona in social conversation are relatively rare. Such words may bring bias and lead to the worse perplexity score, which is in line with the results in Zheng et al. (2020); Dinan et al. (2019). The baseline methods with a lower perplexity score tend to generate more general responses; thus, they cannot generate responses that match the bilateral personas. As a result, the BPAcc scores of these baseline methods are relatively low. (3) Compared with the methods without personalized information (i.e., S2S + Atten., TTransfo. and LConv.), the methods with unilateral personalized information (i.e., TTransfo.+P, LConv.+P, and PWD-WP) on the two test sets get higher BPAcc scores. Moreover, the method with bilateral personalized information (i.e., our method) has a higher BPAcc score on the two test sets than the strong baseline method with unilateral personalized information (i.e., PWDWP). It indicates the effectiveness of the proposed bilateral persona classifier to evaluate the degree of personalization and bilateral-consistent. (4) On the random set, the proposed method outperforms the other baseline methods that only incorporate the unilateral persona in BPAcc (i.e., 87.12 in bold). Similar trends are observed on the biased set

(i.e., 93.75 in bold), which indicates that incorporating the other party’s personalized information in the decoding process is beneficial to generate more personalized and more bilateral persona-consistent responses. (5) The proposed different persona-aware weights (i.e., α , β , and γ) can be used to control the persona presented in the generated response. The results of the two test sets show that under different context settings, it will improve the effect of personalized response generation with different persona-aware weights. This indicates that the proposed dynamic persona-aware fusion module is beneficial to generate diversified dialogue responses rich in bilateral personalized information.

5.3 Ablation study

In order to test the performance of different modules on the proposed method, several ablation experiments are implemented as follows. (1) Each module of multi-task settings is deleted, respectively, including the language model (w/o LM) and the dynamic persona-aware fusion module (w/o PAF). (2) The pre-trained model is also deleted (w/o PreT) to test the performance of transfer learning. (3) The dialogue utterance with corresponded personas embedding (w/o PEmb) and the conditional mutual information maximum criterion (w/o CMIM) are deleted, respectively, to test the effect of different strategies on the BPDG method.

Tables 3 and 4 show the ablation results. The visualization of the ablation study on both random and biased test sets is shown in Figs. 8 and 9. From the results, the further conclusion can be drawn that: (1) the LM module learns the language’s semantics from the dialogue context. Without the

Table 3 Ablation results of our proposed method on random test set

| Method | BPAcc | BLEU | F1 | Distinct | PPL |
|----------|-------|-------|-------|--------------|-------|
| BPDG | 64.63 | 11.66 | 11.30 | 0.282 | 61.93 |
| w/o LM | 40.42 | 9.21 | 10.27 | 0.273 | 60.17 |
| w/o PAF | 47.12 | 8.06 | 10.51 | 0.267 | 57.34 |
| w/o PreT | 42.19 | 4.71 | 8.74 | 0.247 | 88.89 |
| w/o PEmb | 43.24 | 9.12 | 11.01 | 0.270 | 63.12 |
| w/o CMIM | 60.14 | 8.13 | 8.83 | 0.301 | 61.93 |

Table 4 Ablation results of our proposed method on biased test set

| Method | BPAcc | BLEU | F1 | Distinct | PPL |
|----------|-------|-------|-------|--------------|-------|
| BPDG | 92.14 | 24.64 | 18.05 | 0.287 | 58.59 |
| W/o LM | 69.54 | 20.96 | 17.87 | 0.254 | 62.50 |
| W/o PAF | 77.22 | 16.95 | 17.15 | 0.240 | 60.34 |
| W/o PreT | 74.57 | 13.56 | 15.12 | 0.232 | 77.52 |
| W/o PEmb | 74.41 | 18.79 | 17.44 | 0.249 | 61.19 |
| W/o CMIM | 87.11 | 17.22 | 16.12 | 0.312 | 58.59 |




Fig. 8 Visualization of the ablation results of our proposed method on random test set, where the distinct score is multiplied with 100 for easy reading




Fig. 9 Visualization of the ablation results of our proposed method on biased test set, where the distinct score is multiplied with 100 for easy reading

LM module, it will hurt the dynamic persona-aware fusion on the BPDG method. As a result, the BPAcc score will be decreased most. (2) The PAF module is beneficial to generate more personalized and diversified responses. The above different modules of multi-task learning prove to improve the total effect of personalized dialogue generation. (3) The pre-trained language model provides a good parameter initialization for the BPDG method, which helps to improve training efficiency by transferring the knowledge of the original domain to the target domain. (4) The PEmb strategy improves the final performance by embedding the personalized attributes to the corresponded dialogue utterances. (5) More importantly, the CMIM criterion is effective in improving the BPAcc, BLEU, and F1 scores, but it may decrease the Distinct scores, which are bolded in Tables 3 and 4. This is because the sorting and selection steps from the candidates may hurt the diversity of the generated responses.

5.4 Optimal parameter selection

As shown in Tables 5 and 6, the ablation study of the optimal parameter selection is presented. The experiments are implemented in the random personalized test set, where we

Table 5 Ablation study of hyper-parameter λ_1

| Value | BPAcc | BLEU | F1 | Distinct | PPL |
|-------|--------------|-------------|-------------|--------------|--------------|
| 0 | 33.74 | 6.21 | 8.54 | 0.188 | 59.13 |
| 0.1 | 36.88 | 6.87 | 8.48 | 0.199 | 56.98 |
| 0.2 | 37.48 | 7.63 | 8.81 | 0.227 | 56.66 |
| 0.3 | 36.45 | 7.46 | 8.64 | 0.214 | 56.34 |
| 0.4 | 36.11 | 7.38 | 8.54 | 0.210 | 56.08 |
| 0.5 | 35.84 | 7.24 | 8.34 | 0.197 | 55.95 |
| 0.6 | 34.34 | 7.17 | 8.17 | 0.187 | 55.76 |
| 0.7 | 34.01 | 7.16 | 8.10 | 0.185 | 55.61 |
| 0.9 | 33.66 | 7.09 | 8.07 | 0.181 | 55.39 |
| 1.0 | 33.15 | 7.01 | 8.02 | 0.177 | 55.07 |

Table 6 Ablation study of hyper-parameter λ_2 , where the λ_1 is fixed for optimal

| Value | BPAcc | BLEU | F1 | Distinct | PPL |
|------------|--------------|-------------|-------------|--------------|--------------|
| 0 | 37.48 | 7.63 | 7.64 | 0.227 | 56.66 |
| 0.1 | 45.58 | 7.91 | 7.91 | 0.243 | 58.44 |
| 0.2 | 50.58 | 7.96 | 7.99 | 0.266 | 59.60 |
| 0.3 | 54.99 | 8.04 | 8.32 | 0.287 | 59.97 |
| 0.4 | 57.87 | 8.10 | 8.65 | 0.293 | 60.44 |
| 0.5 | 60.14 | 8.13 | 8.83 | 0.301 | 61.93 |
| 0.6 | 60.01 | 8.11 | 8.81 | 0.299 | 62.94 |
| 0.7 | 59.92 | 8.08 | 8.79 | 0.297 | 63.10 |
| 0.9 | 59.81 | 8.04 | 8.67 | 0.294 | 63.99 |
| 1.0 | 58.58 | 8.01 | 8.48 | 0.284 | 64.84 |

first-tune the λ_1 to find the optimal value. Then, we try different hyper-parameters of λ_2 to select the best value, while fixing the value of λ_1 . From the results, we can conclude that the best hyper-parameter λ_1 equals 0.2, where the optimal metric scores are in bold. The optimal hyper-parameter λ_2 is 0.5. What's more, with the increase of λ_1 , we can find that the PPL value shows a downward trend, which indicates that language modeling can alleviate the language generation perplexity. With the increase of λ_2 , the PPL score tends to increase, which indicates that adding persona prediction will lead to more perplexity. As a result, it still improves the performance of the final response.

6 Human evaluation

We also perform a human evaluation to test the quality of responses generated by different methods. In this section, we introduce these metrics and give a comprehensive analysis of the results.

Table 7 Human evaluation on the random and biased test set

| Method | Sentence Fluency | | Bilateral Persona Consistency | | Context Consistency | |
|------------------|------------------|--------------|-------------------------------|--------------|---------------------|-------------|
| | Rand | Biasd | Rand | Biasd | Rand | Biasd |
| S2S + Atten. | 1.24* | 1.14* | 0.72* | 0.91* | 0.95* | 1.12* |
| Trans. | 1.39* | 1.42* | 0.88* | 1.07* | 1.12* | 1.38* |
| TTransfo. | 1.41* | 1.37* | 0.92* | 1.19* | 1.04* | 1.41 |
| TTransfo.+P | 1.42* | 1.43* | 0.96* | 1.27* | 1.05* | 1.39* |
| LConv. | 1.55* | 1.54* | 1.08* | 1.42 | 1.01* | 1.58 |
| LConv.+P | 1.70* | 1.66* | 1.12* | 1.45 | 1.23 | 1.60* |
| PWDWP | 1.79 | 1.82 | 1.35 | 1.63* | 1.25* | 1.71 |
| Ours | 1.82 | 1.87* | 1.41 | 1.74 | 1.27 | 1.75 |
| Ours, $\alpha=1$ | 1.79* | 1.89* | 1.42 | 1.80* | 1.18* | 1.71* |
| Ours, $\beta=1$ | 1.81 | 1.85 | 1.47 | 1.72* | 1.16* | 1.67 |
| Ours, $\gamma=1$ | 1.76* | 1.78* | 1.09* | 1.32* | 1.24 | 1.69 |
| Human Resp | 1.89 | 1.90 | 1.07* | 1.78 | 1.67 | 1.88 |

With * indicate the significant difference with the best result (t -test and p -value < 0.05)

6.1 Subjective metrics introduction

The evaluation metrics we choose are from three aspects, as is shown below.

(1) Sentence fluency

Sentence fluency represents the fluency of responses generated by different methods.

(2) Bilateral persona consistency

Bilateral persona consistency indicates whether the information is consistent with the user's or the robot's personalized information when generating a response by different methods.

(3) Context consistency

Context consistency means whether the response generated by different methods is consistent with the dialogue context.

Three annotators are required to rate the quality of the responses according to the following three rating criteria: (1) +2: the response is not only semantically and grammatically related, but also bilateral persona-consistent. (2) +1: the response satisfies the grammatical rules and can be used as a response, but is too general and trivial. (3) +0: the response is semantically irrelevant, ungrammatical, or conflicts with the personalized information.

6.2 Results and analysis

We sample 100 dialogue sessions from the original random and biased test set, respectively, for the human evaluation. The inter-annotator agreement is measured with Fleiss's kappa κ (Randolph 2005). Particularly, the κ value for sentence fluency, bilateral persona consistency, and context

consistency is 0.81, 0.71, 0.64 on the random test set, respectively, and 0.75, 0.67, 0.61 on the biased test set, respectively. The results indicate that the sentence fluency, the bilateral persona consistency, and the context coherency of two test sets achieve substantial annotation agreement.

Table 7 shows the results of the human evaluation that the proposed method outperforms all baseline methods in all human metrics (t -test and p -value < 0.05). Further observations indicate that (1) incorporating bilateral personas into the generated response will impair the sentence fluency and the context consistency, which corresponds to the high BPAcc score and the low PPL score in the automatic evaluation. Despite this, our method has achieved significant advantages in fluency and context consistency in two test sets compared with other methods. (2) The proposed dynamic persona-aware fusion module is designed to control different persona-aware weights for the personalized response generation. This module contributes to better bilateral persona consistency. At the same time, the bilateral persona consistency outperforms the human in the random test set and the test set. This shows that the proposed dynamic persona-aware fusion module is conducive to generating more personalized responses in both dialogue contexts. This observation is also in line with the BPAcc in automatic evaluation shown in Tables 1 and 2. (3) Compared with the PWDWP method, the proposed BPDG has a great improvement in context consistency. This is due to the effect of the CMIM criterion, which selects the response from the generated the candidate list under the condition of the bilateral personas and the context. This observation also corresponds with the automatic evaluation results of BLEU and F1 metrics shown in Tables 3 and 4.

Fig. 10 Sampled responses generated by baseline methods and our method

| |
|--|
| All profiles: User: {"Gender": "Male", "Area": " Shanghai ", "Interests": "Sports; Movie"} Robot: {"Gender": "Female", "Area": " Guangzhou ", "Interests": "Shopping; Makeups"} |
| User profile: {"Gender": "Male", "Area": " Shanghai ", "Interests": "Sports; Movie"} Robot profile: {"Gender": "Female", "Area": " Guangzhou ", "Interests": "Shopping; Makeups"} |
| History: User: What are your recent arrangements? Robot: I plan to travel to Nanjing recently. |
| User's question: Nanjing is not fun, come to visit me? |
| Ground-truth response: Are there any interesting attractions in Shanghai ? |
| S2S + Atten.: Ok, it's fun. Trans.: Ok, let's play fun. LConv.: Ok, let's play together. LConv. + P: Go to fun attractions together? TTransfo.: I have a chance to come to play with you! TTransfo. + P: I will come to Guangzhou . PWDWP: I will see you in Guangzhou , and play with you! |
| Ours: What's fun in Shanghai ? I'm in Guangzhou . Ours, $\alpha = 1$: What's fun in Shanghai ? Ours, $\beta = 1$: Why are you not so interesting about Nanjing? I'm in Guangzhou . Ours, $\gamma = 1$: Are there any fun places to travel? |

6.3 Case study

The case study is shown in Fig. 10. The proposed method can generate a response consistent with the personas of both parties in the conversation. As we can see, the response generated by the TTransfo.+P and the PWDWP methods may be unilateral persona-consistent without incorporating the persona of the other party. The other baseline methods (i.e., S2S + Atten., TTrans., TTransfo., LConv., LConv.+P) may also generate a general response that lacks personalized information. The proposed BPDG method utilizes bilateral personalized information to generate responses that are in line with human cognition while constraining the contents of the generated responses with the CMIM criterion. Specifically, given the user input and the bilateral personas, our method can control the generated response content with different persona-aware weights. The $\alpha = 1$ means that the user's personalized information is presented in the response, such as Shanghai. The $\beta = 1$ means that the robot's personalized information presents in the response such as Guangzhou. The $\gamma = 1$ means that the personalized information does not present in the response, but it is relevant to the context, such as travel.

7 Conclusion

This article proposed the bilateral personalized dialogue generation (BPDG) method to generate more personalized and bilateral persona-consistent responses. Specifically, our method first utilized transfer learning to initialize the param-

eters of the pre-trained model. Then, dialogue context and bilateral personas were encoded through the encoder. Next, the dynamic persona-aware fusion module was designed to control the persona presented in the generated response adaptively. Finally, the encoder, the dynamic persona-aware fusion module, and the decoder were jointly trained with multi-task learning. The multi-tasks contained the language model, persona prediction, and dialogue generation. The conditional mutual information maximum (CMIM) criterion was adopted with contrastive learning to select the proper response from the generated candidates to bridge the gap between the learning objective function and evaluation metrics. Experiments showed that the transfer learning and multi-task learning method were conducive to improving the performance of dialogue generation in metrics of bilateral persona accuracy. In addition, the generated candidate responses were selected with the CMIM criterion through contrastive learning, which showed that the quality of the final response could be significantly improved. Extensive experiments in the random and biased personalized dialogue test sets were conducted to measure the effectiveness of the BPDG method, which showed that the BPDG method had advantages in four metrics, including the bilateral persona accuracy, F1, BLEU, and distinct scores. The human evaluation results proved that the BPDG method generated more fluent, context-consistent, and bilateral persona-consistent responses than several state-of-the-art methods.

It is worth noting that in open-domain dialogue, the human response is one-to-many, and the open-domain corpus cannot contain all the situations. Moreover, people will respond and reason based on existing information during the conversation.

In the future, we will explore other fusion strategy-based dialogue generation methods with comprehensive reasoning of the existing information to improve the generated response's quality.

Funding The authors have not disclosed any funding.

Data Availability Enquiries about data availability should be directed to the authors.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

References

- Adiwardana D, Luong, et al (2020) Towards a human-like open-domain chatbot. arXiv preprint [arXiv:2001.09977](https://arxiv.org/abs/2001.09977)
- Baltescu P, Blunsom P (2015) Pragmatic neural language modelling in machine translation. In: Proceedings of the 2015 conference of the north american chapter of the association for computational linguistics: human language technologies, pp 820–829
- Cai H, Chen H, Song Y, Ding Z, Bao Y, Yan W, Zhao X (2020) Group-wise contrastive learning for neural dialogue generation. In: Proceedings of the 2020 conference on empirical methods in natural language processing: findings, pp 793–802
- Chen T, Kornblith S, Norouzi M, Hinton G (2020) A simple framework for contrastive learning of visual representations. In: International conference on machine learning, PMLR, pp 1597–1607
- Dai B, Lin D (2017) Contrastive learning for image captioning. In: Proceedings of the 31st international conference on neural information processing systems, pp 898–907
- Dash PB, Naik B, Nayak J, Vimal S (2021) Deep belief network-based probabilistic generative model for detection of robotic manipulator failure execution. *Soft Comput* pp 1–13
- Dinan E, Logacheva V, Malykh V, Miller A, Shuster K, Urbanek J, Kiela D, Szlam A, Serban I, Lowe R, et al (2019) The second conversational intelligence challenge (convai2). arXiv preprint [arXiv:1902.00098](https://arxiv.org/abs/1902.00098)
- Fleuret F (2004) Fast binary feature selection with conditional mutual information. *J Mach Learn Res* 5:1531–1555
- Goldberg LR (1993) The structure of phenotypic personality traits. *Am Psychol* 48(1):26–34
- Golovanov S, Kurbanov R, Nikolenko S, Truskovskiy K, Tselousov A, Wolf T (2019) Large-scale transfer learning for natural language generation. In: Proceedings of the 57th annual meeting of the association for computational linguistics, pp 6053–6058
- Golovanov S, Tselousov A, Kurbanov R, Nikolenko SI (2020) Lost in conversation: A conversational agent based on the transformer and transfer learning. In: The NeurIPS'18 competition, Springer, pp 295–315
- Gutmann MU, Hyvärinen A (2012) Noise-contrastive estimation of unannotated statistical models, with applications to natural image statistics. *J Mach Learn Res* 13(2)
- Hadsell R, Chopra S, LeCun Y (2006) Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06), IEEE, vol 2, pp 1735–1742
- He D, Xia Y, Qin T, Wang L, Yu N, Liu TY, Ma WY (2016) Dual learning for machine translation. In: Advances in neural information processing systems, pp 820–828
- Huang F, Wan D, Shao Z, Ke P, Guan J, Niu Y, Zhu X, Huang M (2020a) Cotk: An open-source toolkit for fast development and fair evaluation of text generation. arXiv preprint [arXiv:2002.00583](https://arxiv.org/abs/2002.00583)
- Huang M, Zhu X, Gao J (2020) Challenges in building intelligent open-domain dialog systems. *ACM Trans Inf Syst (TOIS)* 38(3):1–32
- Isard A, Brockmann C, Oberlander J (2006) Individuality and alignment in generated dialogues. In: Proceedings of the fourth international natural language generation conference, pp 25–32
- Kulikov I, Lee J, Cho K (2019) Multi-turn beam search for neural dialogue modeling. arXiv preprint [arXiv:1906.00141](https://arxiv.org/abs/1906.00141)
- Lan Z, Chen M, Goodman S, Gimpel K, Sharma P, Soricut R (2019) Albert: A lite bert for self-supervised learning of language representations. In: International conference on learning representations
- Li J, Galley M, Brockett C, Gao J, Dolan B (2016a) A diversity-promoting objective function for neural conversation models. In: Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: human language technologies, pp 110–119
- Li J, Galley M, Brockett C, Spithourakis G, Gao J, Dolan B (2016b) A persona-based neural conversation model. In: Proceedings of the 54th annual meeting of the association for computational linguistics (Volume 1: Long Papers), pp 994–1003
- Liu Y, Liu P (2021) Simcls: A simple framework for contrastive learning of abstractive summarization. arXiv preprint [arXiv:2106.01890](https://arxiv.org/abs/2106.01890)
- Luong MT, Pham H, Manning CD (2015) Effective approaches to attention-based neural machine translation. In: Proceedings of the 2015 conference on empirical methods in natural language processing, pp 1412–1421
- Ma WJLHH et al (2021) Hierarchical matching network for multi-turn response selection in retrieval-based chatbots. *Soft Comput* 9:9609–9624
- Madotto A, Lin Z, Wu CS, Fung P (2019) Personalizing dialogue agents via meta-learning. In: Proceedings of the 57th annual meeting of the association for computational linguistics, pp 5454–5459
- Mairesse F, Walker M (2007) Personage: Personality generation for dialogue. In: Proceedings of the 45th annual meeting of the association of computational linguistics, pp 496–503
- Martin TP, Azvine B (2003) Adaptive user modelling in intelligent telephone and email assistants. *Soft Comput* 8(2):93–101
- Mo K, Li S, Zhang Y, Li J, Yang Q (2016) Personalizing a dialogue system with transfer reinforcement learning. arXiv preprint [arXiv:1610.02891](https://arxiv.org/abs/1610.02891)
- Papineni K, Roukos S, Ward T, Zhu WJ (2002) Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting of the association for computational linguistics, pp 311–318
- Qian Q, Huang M, Zhao H, Xu J, Zhu X (2018) Assigning personality/profile to a chatting machine for coherent conversation generation. In: Proceedings of the 27th international joint conference on artificial intelligence, pp 4279–4285
- Randolph JJ (2005) Free-marginal multirater kappa (multirater k [free]): An alternative to fleiss' fixed-marginal multirater kappa. Online submission
- Roller S, Boureau YL, Weston J, Bordes A, Dinan E, Fan A, Gunning D, Ju D, Li M, Poff S, et al. (2020) Open-domain conversational agents: Current progress, open problems, and future directions. arXiv preprint [arXiv:2006.12442](https://arxiv.org/abs/2006.12442)
- Rush AM (2018) The annotated transformer. In: Proceedings of workshop for NLP open source software (NLP-OSS), pp 52–60
- Song H, Zhang WN, Cui Y, Wang D, Liu T (2019) Exploiting persona information for diverse generation of conversational responses. In:

- Proceedings of the 28th international joint conference on artificial intelligence, AAAI Press, pp 5190–5196
- Sun Y, Wang S, Li Y, Feng S, Chen X, Zhang H, Tian X, Zhu D, Tian H, Wu H (2019) Ernie: Enhanced representation through knowledge integration. arXiv preprint [arXiv:1904.09223](https://arxiv.org/abs/1904.09223)
- Sutskever I, Vinyals O, Le QV (2014) Sequence to sequence learning with neural networks. In: Advances in neural information processing systems, pp 3104–3112
- Tanaka T, Ohwi J, Litvintseva LV, Yamafuji K, Ulyanov SV (1997) Soft computing algorithms for intelligent control of a mobile robot for service use. *Soft Comput* 1(2):88–98
- Tramontano A, Scala M, Magliulo M (2019) Wearable devices for health-related quality of life evaluation. *Soft Comput* 23(19):9315–9326
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: Advances in neural information processing systems, pp 5998–6008
- Vijayakumar AK, Cogswell M, Selvaraju RR, Sun Q, Lee S, Crandall D, Batra D (2016) Diverse beam search: Decoding diverse solutions from neural sequence models. arXiv preprint [arXiv:1610.02424](https://arxiv.org/abs/1610.02424)
- Walker MA, Cahn JE, Whittaker SJ (1997) Improvising linguistic style: Social and affective bases for agent personality. In: Proceedings of 1st international conference automation agents, pp 96–105
- Wang D, Zheng TF (2015) Transfer learning for speech and language processing. In: 2015 Asia-Pacific signal and information processing association annual summit and conference (APSIPA), IEEE, pp 1225–1237
- Wang Y, Ke P, Zheng Y, Huang K, Jiang Y, Zhu X, Huang M (2020) A large-scale chinese short-text conversation dataset. arXiv preprint [arXiv:2008.03946](https://arxiv.org/abs/2008.03946)
- Wolf T, Sanh V, Chaumond J, Delangue C (2019) Transfertransfo: A transfer learning approach for neural network based conversational agents. arXiv preprint [arXiv:1901.08149](https://arxiv.org/abs/1901.08149)
- Xu M, Li P, Yang H, Ren P, Ren Z, Chen Z, Ma J (2020) A neural topical expansion framework for unstructured persona-oriented dialogue generation. arXiv preprint [arXiv:2002.02153](https://arxiv.org/abs/2002.02153)
- Yang M, Huang W, Tu W, Qu Q, Shen Y, Lei K (2020) Multitask learning and reinforcement learning for personalized dialog generation: An empirical study. *IEEE transactions on neural networks and learning systems* pp 1–14, 10.1109/TNNLS.2020.2975035
- Zhang S, Dinan E, Urbanek J, Szlam A, Kiela D, Weston J (2018a) Personalizing dialogue agents: I have a dog, do you have pets too? In: Proceedings of the 56th annual meeting of the association for computational linguistics (Volume 1: Long Papers), pp 2204–2213
- Zhang WN, Zhu Q, Wang Y, Zhao Y, Liu T (2018) Neural personalized response generation as domain adaptation. *World Wide Web* 22(4):1427–1446
- Zheng Y, Chen G, Huang M, Liu S, Zhu X (2019) Personalized dialogue generation with diversified traits. arXiv preprint [arXiv:1901.09672](https://arxiv.org/abs/1901.09672)
- Zheng Y, Zhang R, Huang M, Mao X (2020) A pre-training based personalized dialogue generation model with persona-sparse data. AAAI Press, pp 9693–9700

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.