

MCP: Self-supervised Pre-training for Personalized Chatbots with Multi-level Contrastive Sampling

Zhaoheng Huang¹, Zhicheng Dou^{1*}, Yutao Zhu², and Zhengyi Ma¹

¹Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

²University of Montreal, Quebec, Canada

{zhaoheng_huang, dou, zymaa}@ruc.edu.cn, yutao.zhu@umontreal.ca

Abstract

Personalized chatbots focus on endowing the chatbots with a consistent personality to behave like real users and further act as personal assistants. Previous studies have explored generating implicit user profiles from the user’s dialogue history for building personalized chatbots. However, these studies only use the response generation loss to train the entire model, thus it is prone to suffer from the problem of data sparsity. Besides, they overemphasize the final generated response’s quality while ignoring the correlations and fusions between the user’s dialogue history, leading to rough data representations and performance degradation. To tackle these problems, we propose a self-supervised learning framework MCP for capturing better representations from users’ dialogue history for personalized chatbots. Specifically, we apply contrastive sampling methods to leverage the supervised signals hidden in user dialog history, and generate the pre-training samples for enhancing the model. We design three pre-training tasks based on three types of contrastive pairs from user dialogue history, namely response pairs, sequence augmentation pairs, and user pairs. We pre-train the utterance encoder and the history encoder towards the contrastive objectives and use these pre-trained encoders for generating user profiles while personalized response generation. Experimental results on two real-world datasets show a significant improvement in our proposed model MCP compared with the existing methods.

1 Introduction

In recent years, open-domain chatbots have achieved impressive progress in the natural language processing (NLP) field (Li et al., 2016b; Zhou et al., 2018a; Chan et al., 2019; Liu et al., 2020; Zhu et al., 2021b). Given the input utterance, the chatbot responds with an appropriate response. However, for the same input utterance,

chatbots usually provide similar responses for all users. This one-size-fits-all strategy cannot satisfy the varying needs of users and is lean to generate safe but meaningless responses, such as “I don’t know” (Li et al., 2016a; Zhu et al., 2020). To solve this problem, some researchers have begun to endow chatbots with personality and develop personalized chatbots (Zhang et al., 2018; Ma et al., 2021b; Zhong et al., 2022). When equipped with personal information (either given by a predefined profile or learning from dialogue history), personalized chatbots can generate more user-specific and informative responses.

Although the existing methods for personalized chatbots are capable of generating some personalized responses, two major shortcomings hurt user profile building and response generation. **First**, like non-personalized chatbots, they only use the single response generation loss to train the entire model. It has been found that such an optimization way is easy to suffer from the data sparsity scenarios, where users may have limited dialogue history in real-world applications (Song et al., 2019b). **Second**, they overemphasize the quality of the final generated response, while the correlations and the fusion between the user’s historical responses are not captured in the data representations. Such rough data representations are used for building inaccurate user profiles and hurt the generalizability of the model. Recent studies in pre-trained language models have shown that effective data representation is a key factor to improve the performance of various downstream tasks (Kong et al., 2020). Therefore, there is a need to rethink the learning paradigm for developing more effective personalized chatbots.

To address the above issues, we propose to **apply the self-supervised learning framework into personalized chatbots**. We design three novel pre-training tasks to leverage the correlations and self-supervised signals brought by dialogue history.

*Corresponding author.

Self-supervised learning has achieved great success in various NLP (Devlin et al., 2019; Bansal et al., 2021) and Information Retrieval (Guo et al., 2022; Ma et al., 2021a; Zhu et al., 2021a; Zhou et al., 2021) tasks. It aims to let the model learn the intrinsic structure of the raw data without additional supervised signals. In fact, the user dialogue history contains massive supervised signals for capturing personalized information. For example, if one user gave two responses to the same user within a short time, these two responses tend to be semantically relevant, thus the model should build similar representations for these two responses. Such supervised signals hidden in dialogue history can help the model to build more accurate data representations, and we can construct amounts of training samples for pre-training the model in a self-supervised manner. It can be seen that the self-supervised learning framework can perfectly help to solve two major issues: rough data representations and data sparsity.

Based on the above observations, we propose a novel self-supervised learning approach **MCP** for personalized chatbots that is a two-stage neural framework and designs three pre-training tasks based on the dialogue history. To build the user profile from the user’s dialogue history, we design two encoders for learning the single utterance representation and the user profile representation, *i.e.*, **utterance encoder** and **history encoder**. The utterance encoder aims to learn the semantics and representations of user responses, and the history encoder focuses on encoding the user’s historical response sequence and building the user profile representation. We design a two-stage framework for applying self-supervised learning to personalized chatbots. **In the first stage**, we design three pre-training tasks to generate supervised signals from the dialogue history and pre-train the utterance encoder and the history encoder towards these designed objectives. We construct three types of contrastive pairs from the user’s dialogue history as the pre-training samples, namely response pairs, sequence augmentation pairs, and user pairs. Via such a pre-trained method, our model can learn better data representations based on the pre-training data, further adapting data representations to personalized chatbot scenarios. **In the second stage**, we use the parameters of the pre-trained encoders to initialize the encoders in the personalized encoder-decoder model and use the utterance encoder and history encoder to build the user’s profile represen-

tation. The user profiles are fed into the encoder-decoder model to drive the personalized response generation process. To verify the effectiveness of our model, we conduct extensive experiments on two large-scale datasets, *i.e.*, Weibo and Reddit. Experimental results show that MCP achieves state-of-the-art performance compared to a number of competitive methods.

Our contributions can be summarized as follows: (1) We propose a two-stage framework for applying self-supervised learning to train personalized chatbots for better data representation. To the best of our knowledge, it is the first time that pre-training is leveraged based on a user’s dialogue history for personalized chatbots. (2) We design three pre-training tasks based on the correlations and supervised signals hidden in the user’s dialogue history. We construct response pairs, sequence augmentation pairs, and user pairs from different views of the dialogue history for better pre-training the encoder models. (3) The experimental results on two real-world large-scale datasets show the effectiveness of our proposed model.

2 Related Work

Open-domain Generation-based Chatbots. Here we briefly introduce some generation-based methods, as they are most relevant to ours. In some early studies, dialogue response generation was considered a statistical machine translation problem (Ritter et al., 2011). Nowadays, sequence-to-sequence (Seq2Seq) models are the mainstream method due to the fast development of deep learning and neural network (Serban et al., 2016). Many studies have tried to endow Seq2Seq-based methods with human attributes, such as: (1) introducing additional commonsense knowledge to generate a more reasonable response (Zhou et al., 2018b); (2) tracking the speakers’ emotions to make a more suitable reply (Zhou et al., 2018a); (3) responding with more diverse sentences to avoid a boring user experience (Li et al., 2016a); and (4) generating responses with personas to make the dialogue more consistent (Qian et al., 2018).

Personalized Chatbots. Building a personalized chatbot has attracted public interest in recent years. With a more steady personality, chatbots can have more consistent and informative dialogues with users. Existing methods can be summarized as: (1) Learning with user ID information (Li et al., 2016b; Al-Rfou et al., 2016; Bak and Oh, 2019; Chan et al.,

2019). These methods embed user IDs into vectors and use them to guide the dialogue. However, it is impractical to maintain a huge user ID embedding table for real-world large-scale datasets. (2) Learning with explicit user profiles (Qian et al., 2018; Zhang et al., 2018; Olabiya et al., 2019; Song et al., 2019a). This group of methods aims at using descriptive information (such as persona sentences or attribute tables) to generate personalized responses, but such information is hard to collect, and it cannot keep track of user interest changes. (3) Learning with implicit user profiles (Ma et al., 2021b; Zhong et al., 2022). These methods extract personalized information automatically from a user’s dialogue history to generate personalized responses.

However, user dialogue history is limited and persona-sparse, and it is challenging to learn robust representations of the user’s implicit profile. Therefore, we aim to tackle this problem and enhance representations to generate personalized responses.

3 Methodology

In this section, we first provide the problem statement and an overview of our model. Then, we elaborate the details of each component in the model.

3.1 Problem Statement and Overview

Considering a set of users defined as $U = \{u_1, u_2, \dots, u_l\}$, for a specific user u_i , their dialogue history is denoted as $H^i = [(q_1^i, r_1^i, t_1^i), (q_2^i, r_2^i, t_2^i), \dots, (q_n^i, r_n^i, t_n^i)]$, where n represents the length of user dialogue history, and t_k^i represents the response time of r_k^i . Note that q_j^i is a query issued by other users, while r_j^i is the response given by the user u_i . We only use the historical responses that reflect the personal information of the user u_i , so H^i is simplified as $[r_1^i, r_2^i, \dots, r_n^i]$ henceforth. With the above notations, our task is defined as generating a personalized response r for the user u_i to reply to a new query q with the personalized information extracted from the user’s dialogue history H^i .

The overview of our MCP model is shown in Figure 1. For user dialogue history, we design an utterance encoder to independently represent each utterance into vectors; and a history encoder to model the interactions among utterances and capture their sequential information. The obtained history representation (containing personal information) and the input query are fed into a Transformer-based Seq2Seq structure to generate a personalized re-




Figure 1: The proposed three pre-training tasks based on (1) Response Pairs, (2) Sequence Augmentation Pairs, and (3) User Pairs.

sponse.

As the main contribution of this work, we design three self-supervised pre-training tasks for an utterance encoder and a history encoder. By sampling contrastive samples at utterance-, dialogue history-, and user-level, the two encoders are trained to capture more personalized information (such as user interest) from user dialogue history and create more robust representations. We will elaborate the details of the three pre-training tasks and our model in the following sections.

3.2 Utterance-level Contrastive Learning

The first contrastive pre-training task is based on our observation on conversations between two people: the consecutive utterances within a short period are usually about one topic. Intuitively, the representation of two utterances under similar topics in the dialogue history should be closer than other utterances. Following this assumption, we extract contrastive samples from the dataset by the rule: the two utterances should be replied to the

same user and their issued time should be close. Formally, given a user u , we select two triplets (q_i, r_i, t_i) and (q_j, r_j, t_j) . Both queries q_i and q_j are issued by the same user u' , and the response time interval satisfies $|t_i - t_j| \leq \hat{t}$, where \hat{t} is a time threshold. As a result, r_i and r_j construct a pair of contrastive samples. We then apply a contrastive learning objective to pull close their representations and push apart the representation of other utterances.

Specifically, we use an utterance encoder to represent the responses r_i and r_j as:

$$\begin{aligned}\bar{\mathbf{r}}_i &= \text{Mean}(\text{Transformer}_u(r_i)), \\ \bar{\mathbf{r}}_j &= \text{Mean}(\text{Transformer}_u(r_j)),\end{aligned}$$

where $\text{Mean}(\cdot)$ is the mean pooling operation. The general contrastive learning objective aims at reducing the distance of positive pairs and increasing the distance of negative pairs (Gutmann and Hyvärinen, 2010). In our case, for a positive response representation pair $(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}_j)$ in a mini-batch of N pairs, we consider the response $\bar{\mathbf{r}}$ in the other $N - 1$ pairs as negative set R^- . Hence, the loss function can be defined as:

$$\begin{aligned}\mathcal{L}_{\text{utt}} &= -\log \frac{\phi(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}_j)}{\phi(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}_j) + \sum_{\bar{\mathbf{r}}' \in R^-} \phi(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}')}, \\ \phi(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}_j) &= \exp(\text{CosSim}(\bar{\mathbf{r}}_i, \bar{\mathbf{r}}_j)),\end{aligned}$$

where the function $\text{CosSim}(\cdot)$ denotes the cosine-similarity.

3.3 History-level Contrastive Learning

In personalized response generation, the main task is to capture the personalized information (*e.g.*, the user's interests) from the dialogue history. To enhance the history encoder and obtain a more robust history sequence representation, we apply sequence augmentation to construct different views of the dialogue history. By comparing the original sequence and the augmented one, the history encoder is encouraged to highlight the information that best reflects the user's personality. Specifically, we design the following four strategies:

(1) **Session masking.** When chatting with others, a user often shows the same interest to different people. It inspires us that after masking all responses to a single person, the rest of the user's dialogue history can still represent the user's intrinsic interest. Based on this assumption, for each user, we select a series of her responses that respond to the

same user and then mask them in the sequence of her dialogue history. The altered sequence and the original sequence construct a positive pair. With this task, we could improve the generalizability of the user profiling model.

(2) **Sequence random masking.** Generally, a user's long-term interests are stable. This inspires us to randomly mask $k\%$ sentences in sequence to prevent our model from overly focusing on several certain responses. The new sequence and the original one construct a positive pair. The unmasking part of the dialogue sequence doesn't affect the user's long-term interests. Furthermore, the robustness of our model is enhanced with this task.

(3) **Sequence re-ordering.** Although the order of user dialogue history has an effect on the user profile, simply re-ordering two responses from dialogue history should not affect the user's long-term and stable interests. Thus, we randomly choose pairs of responses that respond to different people and swap their positions in the sequence. The augmented sequence and the original one construct a positive pair. The long-term and stable interests are highlighted with this task.

(4) **Short interval sequence masking.** Inspired by research on session-based social recommendation (Zhang et al., 2020b), a pair of consecutive responses with tight time intervals imply more user interests. Thus, we choose pairs from consecutive responses which satisfy the condition that time interval $t \leq \hat{t}$, where \hat{t} denotes a time threshold. By simply masking the latter response in the chosen pairs, the user profile should remain unchanged.

For a specific user dialogue history sequence $[r_1, \dots, r_n]$, we augment it and obtain a new sequence as $[r'_1, \dots, r'_n]$ following the aforementioned four strategies. These two sequences are treated as a positive pair and encoded by the history encoder as:

$$\begin{aligned}\mathbf{U} &= \text{Mean}(\text{Transformer}_h([\bar{\mathbf{r}}_1; \dots; \bar{\mathbf{r}}_n])), \\ \mathbf{U}_{\text{aug}} &= \text{Mean}(\text{Transformer}_h([\bar{\mathbf{r}}'_1; \dots; \bar{\mathbf{r}}'_n])),\end{aligned}$$

where $[\cdot]$ is the concatenation operation. As the obtained vector containing the user's personalized information, we call it *user profile* vector. Thereafter, we get positive augmented pairs $(\mathbf{U}, \mathbf{U}_{\text{aug}})$. Similarly, the negative set S^- contains $N - 1$ augmented sequences sampled from the mini-batch of

N size, and the loss is as follows:

$$\mathcal{L}_{\text{seq}} = -\log \frac{\phi(\mathbf{U}, \mathbf{U}_{\text{aug}})}{\phi(\mathbf{U}, \mathbf{U}_{\text{aug}}) + \sum_{\mathbf{U}' \in \mathcal{S}^-} \phi(\mathbf{U}, \mathbf{U}')}. \quad (1)$$

3.4 User-level Contrastive Learning

Besides, we also notice that user relationships affect the similarity of user profiles. For two users u_i and u_j , the set of people they respond to can be respectively described as U'_i and U'_j . Following research towards a recommendation system on social chatting forum (Ma, 2014), the larger common group of users they chat with, the more common topics and interests they share with. Thus, we only choose pairs which support $|U'_i \cap U'_j| \geq \hat{s}$, where \hat{s} denotes a threshold.

After sampling among users, we get a positive pair of similar user profiles ($\mathbf{U}, \mathbf{U}_{\text{pos}}$), and a negative set \mathbf{U}^- which consists of other user profiles in the mini-batch. The loss learned by this task is as follows:

$$\mathcal{L}_{\text{user}} = -\log \frac{\phi(\mathbf{U}, \mathbf{U}_{\text{pos}})}{\phi(\mathbf{U}, \mathbf{U}_{\text{pos}}) + \sum_{\mathbf{U}' \in \mathbf{U}^-} \phi(\mathbf{U}, \mathbf{U}')}. \quad (2)$$

3.5 Personalized Response Generation

We apply an encoder-decoder architecture based on Transformer (Vaswani et al., 2017) to generate personalized responses. First, the Transformer encoder encodes both user profile \mathbf{U} and current query p combined together. Then, the decoding module decodes the outputs from the first step, added with a linear layer so as to project them into a space with a vocabulary-size dimension. The process can be defined as:

$$\mathbf{M} = \text{TransformerEnc}([\mathbf{U}; p]), \quad (3)$$

$$\hat{y} = \text{Linear}(\text{TransformerDec}(\mathbf{M})), \quad (4)$$

where $\hat{y} \in \mathbb{R}^v$, and v is the size of the vocabulary. Later, \hat{y} is normalized by a softmax layer and represents the generation probability:

$$\hat{y}_{\text{prob}} = \text{softmax}(\hat{y}). \quad (5)$$

3.6 Training and Optimization

At the pre-training stage, our goal is to optimize the loss of the proposed tasks at three levels:

$$\mathcal{L}^p = \mathcal{L}_{\text{utt}} + \mathcal{L}_{\text{seq}} + \mathcal{L}_{\text{user}}. \quad (6)$$

At the fine-tuning stage, the model is trained to maximize the generation probability of the ground-truth response y , which can be optimized by cross-entropy loss as:

$$\mathcal{L}^f = -y \log \hat{y}_{\text{prob}}. \quad (7)$$

4 Experiments

4.1 Datasets

Following previous study (Ma et al., 2021b), we conduct experiments on two datasets from Chinese Weibo (Qian et al., 2021) and English Reddit (Zhang et al., 2020a). Both datasets are collected from open-domain social media platforms, where users can post with various interests, and respond to other users in who they are interested. Each utterance in the dialogue has its own timestamp and user ID. We consider that a training sample contains three parts: (1) a query; (2) a corresponding response; and (3) a dialogue history containing several responses issued before the current response.

Weibo Dataset Following (Ma et al., 2021b), we use a subset of the PChatbotW dataset (Qian et al., 2021) containing 300K users. The dataset is collected from the Weibo online platform¹. We construct the corresponding query and response to a query-response pair. Moreover, we compare the response timestamps and consider the replies before the current response as the dialogue history. Also, we refer to the method introduced in (Qian et al., 2021). Specifically, we remove hashtags, URLs, emojis and swear words from sentences. Then, we remove sentences that include more than 100 words or less than 5 words or contain multi-languages.

Reddit Dataset To measure model performance in different language, we also introduce the English dataset, which is collected from the English social platform Reddit². In terms of the specific tree structure of Reddit, we consider the pair of parent node and its child node as a query-response pair. We use the same cleaning process applied to the Weibo dataset to handle the Reddit dataset. It comprises 315.34K users.

4.2 Baseline Methods

We evaluate our model with four groups of relevant and typical baseline methods:

¹<https://www.weibo.com>

²<https://www.reddit.com>

Table 1: The results of automatic evaluations on Weibo and Reddit datasets. We categorize the baselines into four groups: (1) non-personalized; (2) using user-id; (3) using explicit user profiles; and (4) using dialogue history. The best results are in bold. “†” indicates that our model achieves significant improvement over the state-of-the-art result in paired t-test with p -value < 0.05 .

Dataset	Model	Overlap-based Metric			Embedding Similarity			Personalization	
		BLEU-1	BLEU-2	ROUGE-L	Average	Extreme	Greedy	P-F1(%)	P-Cover
Weibo	(1) Seq2SeqWA	3.335	0.294	8.740	0.321	0.266	0.254	1.698	0.041
	(1) MMI	3.632	0.095	5.317	0.477	0.695	0.305	1.874	0.054
	(2) Speaker	4.997	0.113	7.993	0.492	0.712	0.311	2.119	0.082
	(2) PersonaWAE	3.503	0.155	11.305	0.513	0.724	0.307	5.108	0.093
	(3) GPMN	4.901	0.696	8.090	0.353	0.391	0.301	4.512	0.084
	(3) PerCVAE	5.115	0.299	7.952	0.469	0.659	0.299	3.817	0.086
	(4) VHRED-P	6.992	0.709	10.695	0.437	0.560	0.307	5.459	0.065
	(4) ReCoSa-P	7.266	0.844	11.469	0.419	0.510	0.312	5.717	0.061
	(4) DHAP	9.324	0.894	14.122	0.523	0.747	0.313	7.013	0.144
	(4) MCP (Ours)	9.714 †	0.952 †	14.601 †	0.542 †	0.783 †	0.319 †	7.315 †	0.149 †
Reddit	(1) Seq2SeqWA	1.819	0.023	4.069	0.545	0.554	0.472	0.516	0.029
	(1) MMI	2.065	0.011	3.784	0.543	0.607	0.454	0.828	0.038
	(2) Speaker	2.642	0.054	4.469	0.538	0.606	0.457	1.455	0.031
	(2) PersonaWAE	2.637	0.113	8.199	0.629	0.685	0.442	3.392	0.032
	(3) GPMN	2.686	0.376	4.776	0.406	0.331	0.358	3.026	0.037
	(3) PerCVAE	5.933	0.576	8.112	0.637	0.649	0.499	3.456	0.040
	(4) VHRED-P	5.802	0.648	8.345	0.558	0.635	0.472	3.772	0.047
	(4) ReCoSa-P	6.113	0.686	8.899	0.574	0.632	0.510	3.998	0.044
	(4) DHAP	6.858	0.737	11.720	0.709	0.721	0.539	4.639	0.111
	(4) MCP (Ours)	6.590	0.765 †	12.068 †	0.712	0.749 †	0.541	4.919 †	0.117 †

(1) *Non-personalized response generation models*: **Seq2SeqWA** (Bahdanau et al., 2015) applies attention (Luong et al., 2015) module to the GRU-based Seq2Seq model. **MMI** (Li et al., 2016a) optimizes the maximum mutual information loss to improve the diversity of generated responses.

(2) *Personalized models using user ID*: **Speaker** (Li et al., 2016b) improves the Seq2SeqWA model by incorporating user ID embeddings into the input of the decoder. **PersonaWAE** (Chan et al., 2019) generates personalized responses with a Wasserstein autoencoder. The user ID embeddings are used to build a Gaussian mixture distribution for personalization.

(3) *Personalized models using explicit user profiles*: **GPMN** (Zhang et al., 2018) designs a memory module to encode and store the persona profile to enhance the Seq2Seq model. **PerCVAE** (Zhao et al., 2017) applies a conditional variational autoencoder to improve the diversity of response.

(4) *Personalized models using implicit user profiles*: **VHRED-P** (Serban et al., 2017) is a multi-turn response generation method that models dependencies among user’s dialogue context. We replace the dialogue context with the user’s historical post-response pairs to achieve personalization. **ReCoSa-P** (Zhang et al., 2019) is also a multi-turn response generation model that measures the con-

text with an attention mechanism. Similarly, we simulate the context through the historical post-response pairs. **DHAP** (Ma et al., 2021b) constructs general user profile from user response history, and establishes a key-value memory network to build a dynamic query-aware user profile. Then, the model generates personalized responses with a personalized encoder. This is the state-of-the-art method in personalized response generation.

4.3 Implementation Details

For a user, we use the former 80% dialogue history to construct the training set, the middle 10% part for the validation set, and the last 10% for the test set. For each response, we use the corresponding query as the current query and the previous responses as the dialogue history. The mentioned thresholds \tilde{t} , \hat{t} , and \hat{s} are set respectively to 25% and 5% quantile time of all consecutive responses on datasets, and 1. The history length is set as 20, and the hyper-parameter k is set as 30. We initialize the utterance encoder by the parameters of `bert-base-chinese`³ and `bert-base-uncased`⁴ checkpoint respectively.

³<https://huggingface.co/bert-base-chinese>

⁴<https://huggingface.co/bert-base-uncased>

The number of Transformer layers is 6. The hidden size and number of heads of the Transformer are set respectively to 768 and 8. We use beam search with a beam width of 12 to decode words.

4.4 Evaluation Metrics

Automatic Evaluation. We introduce three groups of automatic evaluation metrics to measure the performance of our model.

(1) **Overlap-based Metric.** We use BLEU-1/2 (Papineni et al., 2002) and ROUGE-L (Lin and Och, 2004) to measure the n -gram overlap between the generated response and ground-truth response. (2) **Embedding Similarity.** We employ embedding-based metrics (Chan et al., 2019) to measure the semantic similarity between the generated response and the ground-truth one. (3) **Personalization.** Following previous studies (Ma et al., 2021b), we employ Persona-F1 (Lian et al., 2019) to measure the unigram F1 between the generated responses and historical responses, and Persona Coverage (Song et al., 2019a) to calculate IDF-weighted word overlapping score between the generated responses and dialogue response history. **Human Evaluation.** Considering the diversity of utterances in the real world, a response different from the ground-truth may also be appropriate. Therefore, we randomly select 100 test triplets (*i.e.*, query, response, and user dialogue history) and manually evaluate their quality in terms of three aspects: readability, informativeness, and personalization, which are suggested by Chan et al. (2019). The detailed evaluation criteria are given in Appendix A.

4.5 Experimental Results

Automatic Evaluation . The metric-based evaluation results are shown in Table 1. We can observe that our MCP method achieves the best performance on both datasets in terms of most metrics. The improvement is statistically significant (t-test with p -value < 0.05), which demonstrates the effectiveness of applying self-supervised tasks to pre-train the encoder. We also have the following observations:

(1) In general, the models using dialogue history perform better than those using explicit user profiles or user-id. This indicates that the user’s dialogue history contains sufficient personalized information and is a more suitable way for building personalized chatbots. Furthermore, using any kind of personalized information can improve

Table 2: Human evaluation results on Weibo dataset.

Model	Readabi.	Informat.	Persona.
(1) Seq2SeqWA	2.10	1.85	0.19
(1) MMI	2.06	1.88	0.23
(2) PersonaWAE	2.07	1.99	0.36
(3) GPMN	2.12	1.92	0.35
(3) PerCVAE	2.04	2.01	0.39
(4) DHAP	2.26	2.09	0.56
(4) MCP (Ours)	2.32	2.15	0.59
Ground-Truth	2.69	2.35	0.84




Figure 2: Experiments with different lengths of user dialogue history on the Weibo dataset.

the quality of the response. This confirms the effectiveness and necessity of applying personalization to response generation. (2) Compared with DHAP, MCP achieves better performance on all three groups of metrics, indicating that our model is capable of generating more fluent and personalized responses. The improvement obtained by our effective method can demonstrate the superiority of the proposed self-supervised pre-training tasks.

Human Evaluation. Table 2 shows the result of human evaluation on Weibo dataset. The Fleiss Kappa is around 0.51, which shows a moderate agreement by the annotators. Specifically, MCP outperforms DHAP by 2.65%/2.87%/5.36% respectively on the three aspects. All these results demonstrate the superiority of MCP in generating more fluent, informative, and personalized responses, and the potential of contrastive pre-training tasks on personalized response generation.

5 Further Analysis

We further conduct a series of experiments to make an in-depth analysis of our model. The analyses are based on the results of the Weibo dataset (if not explicitly stated), and similar results can be observed from the Reddit dataset.




Figure 3: The response representation distribution after PCA with scaling (before and after self-supervised pre-training) of user #60 on the Reddit dataset.

Table 3: Ablation study results on Weibo dataset.

Models	BLEU-1	Avg. Emb.	P-F1
MCP (Full)	9.714	0.542	7.315
<i>w/o</i> Utter. Enc	6.656	0.441	5.169
<i>w/o</i> His. Enc	7.548	0.489	5.742
<i>w/o</i> Pre. Utter.	9.144	0.519	6.885
<i>w/o</i> Pre. His.	7.893	0.497	5.969
<i>w/o</i> Pre. User	9.438	0.534	7.146




Figure 4: Distribution of similarity between users on the Weibo dataset.

Influence of History Length Since the personalized information is extracted from the dialogue history, the length of the history will influence the final performance. To explore such an influence, we conduct an experiment with MCP and two other strong baselines (DHAP and ReCoSa-P) by using various numbers of historical responses. The results are shown in Figure 2. We have the following observations: (1) In general, the performance of all three methods is increasing along with more historical responses used. This is consistent with our assumption that more dialogue history can provide more sufficient information. (2) MCP performs

better with different lengths of dialogue history, indicating its robustness and generalizability. (3) When no dialogue history is provided, ReCoSa-P performs much better than both MCP and DHAP. The potential reason is that ReCoSa-P has a more complex structure to encode the user’s input query and generate the response.

Ablation Study To investigate the effect of two encoders in MCP, we conduct an ablation study as: (1) We remove the utterance encoder (*w/o* Utter. Enc), and the utterance embedding is replaced by the mean pooling over all word embeddings. (2) We remove the history encoder (*w/o* His. Enc), and the user profile representation is replaced by the mean pooling over all utterance embeddings. The results are shown on the upper side of Table 3. We can see that both encoders are important for our model since removing either of them leads to performance degradation. Concretely, the utterance encoder has a greater influence on the final performance. This is because the utterance encoder works at a lower level and is designed to capture fine-grained information from the utterance. If it is removed from the model, the poor utterance representation will also influence the history encoder. On the contrary, the history encoder aims at aggregating information from each historical utterance. It has relatively less influence when valuable information has already been refined into the utterance representation.

We further study the effectiveness of our proposed three pre-training tasks, and the results are shown in the lower side of Table 3. We denote the model without pre-training on utterance-/history-/user-level as *w/o* Pre. Utter./His./User., respectively. It is clear to see that all three tasks are beneficial to our model. This demonstrates the

great potential of contrastive learning on optimizing the representation of the utterance and dialogue history.

Effect of Self-supervised Learning To investigate the effect of our proposed self-supervised learning, we measure the representation differences before and after pre-training.

(1) *Utterance Representation Changes.* The purpose of the self-supervised pre-training task is to close the distance between responses sharing the same interests. Thus, we randomly select a user from the Reddit dataset, extract historical responses, and map their embeddings to a two-dimensional space through PCA with scaling. As shown in Figure 3, topic-related responses are marked with the same color and categorized into four groups. In the initialized distribution, responses with the same color are dispersed, but they are obviously closer after self-supervised pre-training. Moreover, the boundary between the four groups is wide and clear, which shows that the utterance encoder has the ability to sense different topic-related information in each response.

(2) *History Representation Changes.* The history encoder is designed to close the distance between representations of similar users. The ability to distinguish users affects the personalization of generated responses. Thus, we randomly select 1000 users from the Weibo dataset and calculate the distance between representations of every two users. Figure 4 shows that the initialized model cannot distinguish users well. After training with self-supervised tasks, the differences between users are magnified, and the distance between similar users is reduced. An interesting finding is that the group of dissimilar users is much larger than the group of similar ones. This may benefit from our user-level contrastive learning task, which samples out similar users and closes their distance of representation.

6 Conclusion

In this paper, we proposed an MCP model for personalized response generation. Different from previous personalized work, first, we proposed a personalized responses generator that consists of an utterance encoder and a history encoder. Next, we designed three self-supervised tasks from two levels to pre-train the two encoders. Experimental results on two real-world datasets confirm the effectiveness of our model in generating informative and personalized responses.

Limitation

Though our method achieves promising results on two real-world datasets, there are still many limitations: (1) The architecture of our method is relatively simple compared with baseline methods. Other advanced structure for modeling the dialogue history may bring more improvement. (2) We have noticed that pre-trained language models have been applied to dialogue generation (e.g., DialoGPT (Zhang et al., 2020a)). However, our method is currently based on a standard Transformer without pre-training on large-scale text datasets. Therefore, how to combine our proposed pre-training task with existing pre-trained language models is still under exploration.

Ethical Statement

The datasets used in this paper are available publicly online. This paper does not contain any data collection or release, so there are no privacy issues. The research will not pose any ethical issues. We hired three well-educated and part-time annotators to conduct human evaluation with a pay of 100 CNY/hour during their evaluation.

Acknowledgments

Zhicheng Dou is the corresponding author. This work was supported by the National Natural Science Foundation of China No. 61872370, Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098, the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China NO. 22XNKJ34, Public Computing Cloud, Renmin University of China, and Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China. The work was partially done at Beijing Key Laboratory of Big Data Management and Analysis Methods, and Key Laboratory of Data Engineering and Knowledge Engineering, MOE.

References

Rami Al-Rfou, Marc Pickett, Javier Snaider, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2016. *Conversational contextual cues: The case of personalization and history for response ranking*. *CoRR*, abs/1606.00372.

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- JinYeong Bak and Alice Oh. 2019. [Variational hierarchical user-based conversation model](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1941–1950. Association for Computational Linguistics.
- Trapti Bansal, Karthick Prasad Gunasekaran, Tong Wang, Tsendsuren Munkhdalai, and Andrew McCallum. 2021. [Diverse distributions of self-supervised tasks for meta-learning in NLP](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 5812–5824. Association for Computational Linguistics.
- Zhangming Chan, Juntao Li, Xiaopeng Yang, Xiuying Chen, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. [Modeling personalization in continuous space for response generation via augmented wasserstein autoencoders](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 1931–1940. 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Yu Guo, Zhengyi Ma, Jiaxin Mao, Hongjin Qian, Xinyu Zhang, Hao Jiang, Zhao Cao, and Zhicheng Dou. 2022. [Webformer: Pre-training with web pages for information retrieval](#). In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pages 1502–1512. ACM.
- Michael Gutmann and Aapo Hyvärinen. 2010. [Noise-contrastive estimation: A new estimation principle for unnormalized statistical models](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 297–304. JMLR.org.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. [A mutual information maximization perspective of language representation learning](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. [A diversity-promoting objective function for neural conversation models](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 110–119. The Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016b. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. [Learning to select knowledge for response generation in dialog systems](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5081–5087. ijcai.org.
- Chin-Yew Lin and Franz Josef Och. 2004. [Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics, 21-26 July, 2004, Barcelona, Spain*, pages 605–612. ACL.
- Ya-Fei Liu, Hongjin Qian, Hengpeng Xu, and Jinmao Wei. 2020. [Speaker or listener? the role of a dialogue agent](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4861–4869. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1412–1421. The Association for Computational Linguistics.
- Hao Ma. 2014. [On measuring social friend interest similarities in recommender systems](#). In *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014*, pages 465–474. ACM.
- Zhengyi Ma, Zhicheng Dou, Wei Xu, Xinyu Zhang, Hao Jiang, Zhao Cao, and Ji-Rong Wen. 2021a. [Pre-training for ad-hoc retrieval: Hyperlink is also you](#)

- need. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, pages 1212–1221. ACM.
- Zhengyi Ma, Zhicheng Dou, Yutao Zhu, Hanxun Zhong, and Ji-Rong Wen. 2021b. [One chatbot per person: Creating personalized chatbots based on implicit user profiles](#). In *Proceedings of the SIGIR 2021*. ACM.
- Oluwatobi Olabiyi, Anish Khazane, Alan Salimov, and Erik T. Mueller. 2019. [An adversarial learning framework for a persona-based multi-turn dialogue model](#). *CoRR*, abs/1905.01992.
- 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, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Hongjin Qian, Xiaohe Li, Hanxun Zhong, Yu Guo, Yueyuan Ma, Yutao Zhu, Zhanliang Liu, Zhicheng Dou, and Ji-Rong Wen. 2021. [Pchatbot: A large-scale dataset for personalized chatbot](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 2470–2477. ACM.
- Qiao Qian, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. [Assigning personality/profile to a chatting machine for coherent conversation generation](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4279–4285. ijcai.org.
- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. [Data-driven response generation in social media](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 583–593. ACL.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784. AAAI Press.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3295–3301. AAAI Press.
- Haoyu Song, Weinan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019a. [Exploiting persona information for diverse generation of conversational responses](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5190–5196. ijcai.org.
- Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. 2019b. [Autoint: Automatic feature interaction learning via self-attentive neural networks](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, pages 1161–1170. ACM.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Hainan Zhang, Yanyan Lan, Liang Pang, Jiafeng Guo, and Xueqi Cheng. 2019. [Recosa: Detecting the relevant contexts with self-attention for multi-turn dialogue generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3721–3730. Association for Computational Linguistics.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2204–2213. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020a. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL 2020, Online, July 5-10, 2020*, pages 270–278. Association for Computational Linguistics.
- Youjie Zhang, Ting Bai, Bin Wu, and Bai Wang. 2020b. [A time interval aware approach for session-based social recommendation](#). In *Knowledge Science, Engineering and Management - 13th International Conference, KSEM 2020, Hangzhou, China, August 28-30, 2020, Proceedings, Part II*, volume 12275 of *Lecture Notes in Computer Science*, pages 88–95. Springer.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#).

- In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 654–664. Association for Computational Linguistics.
- Hanxun Zhong, Zhicheng Dou, Yutao Zhu, Hongjin Qian, and Ji-Rong Wen. 2022. [Less is more: Learning to refine dialogue history for personalized dialogue generation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 5808–5820. Association for Computational Linguistics.
- Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018a. [Emotional chatting machine: Emotional conversation generation with internal and external memory](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 730–739. AAAI Press.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018b. [Commonsense knowledge aware conversation generation with graph attention](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4623–4629. ijcai.org.
- Yujia Zhou, Zhicheng Dou, Yutao Zhu, and Ji-Rong Wen. 2021. [PSSL: self-supervised learning for personalized search with contrastive sampling](#). In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, pages 2749–2758. ACM.
- Yutao Zhu, Zhicheng Dou, Jian-Yun Nie, and Ji-Rong Wen. 2020. [Reboost: a retrieval-boosted sequence-to-sequence model for neural response generation](#). *Inf. Retr. J.*, 23(1):27–48.
- Yutao Zhu, Jian-Yun Nie, Zhicheng Dou, Zhengyi Ma, Xinyu Zhang, Pan Du, Xiaochen Zuo, and Hao Jiang. 2021a. [Contrastive learning of user behavior sequence for context-aware document ranking](#). In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, pages 2780–2791. ACM.
- Yutao Zhu, Jian-Yun Nie, Kun Zhou, Pan Du, and Zhicheng Dou. 2021b. [Content selection network for document-grounded retrieval-based chatbots](#). In *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part I*, volume 12656 of *Lecture Notes in Computer Science*, pages 755–769. Springer.

Appendix

A Human Evaluation Criteria

We randomly select 100 test samples to conduct human evaluations. The three well-educated annotators are shown with the generated responses, the corresponding query, and the user’s historical query-response pairs. Then the annotators will evaluate the generated responses from three aspects in a double-blind fashion. The human evaluation criteria include: (1) Readability, which measures the grammatical correctness and fluency of the generated response; (2) Informativeness, which measures if the generated response is informative or trivial; and (3) Personalization, which measures whether the generated response shares some persona-related information with the dialogue history of the user. The range of the first two factors is [1, 3], while the third is scored on a scale of [0, 1]. The specific scoring criteria for human annotation are shown in Table 4.

Readability

- 1: Hard to read or syntactically incorrect
- 2: Good grammatical format
- 3: Fluent and easy to understand

Informativeness

- 1: Meaningless responses
- 2: Ambiguous or contain insufficient informative words
- 3: Responses with clear meaning or purpose

Personalization

- 0: Contain no personal information in user history
 - 1: Reflect user’s personal information in history
-

Table 4: Specific scoring criteria of human annotation.