

A Neural Topical Expansion Framework for Unstructured Persona-oriented Dialogue Generation

Minghong Xu¹ and Piji Li^{2*} and Haoran Yang³ and Pengjie Ren⁴
and Zhaochun Ren^{5*} and Zhumin Chen⁶ and Jun Ma⁷

Abstract. Unstructured Persona-oriented Dialogue Systems (UPDS) has been demonstrated effective in generating persona consistent responses by utilizing predefined natural language user persona descriptions (e.g., “I am a vegan”). However, the predefined user persona descriptions are usually short and limited to only a few descriptive words, which makes it hard to correlate them with the dialogues. As a result, existing methods either fail to use the persona description or use them improperly when generating persona consistent responses. To address this, we propose a neural topical expansion framework, namely Persona Exploration and Exploitation (PEE), which is able to extend the predefined user persona description with semantically correlated content before utilizing them to generate dialogue responses. PEE consists of two main modules: persona exploration and persona exploitation. The former learns to extend the predefined user persona description by mining and correlating with existing dialogue corpus using a variational auto-encoder (VAE) based topic model. The latter learns to generate persona consistent responses by utilizing the predefined and extended user persona description. In order to make persona exploitation learn to utilize user persona description more properly, we also introduce two persona-oriented loss functions: Persona-oriented Matching (P-Match) loss and Persona-oriented Bag-of-Words (P-BoWs) loss which respectively supervise persona selection in encoder and decoder. Experimental results show that our approach outperforms state-of-the-art baselines, in terms of both automatic and human evaluations.

1 Introduction

Persona-oriented dialogue systems have attracted an increasing attention as they can generate persona consistent responses [3, 8, 12, 14]. Existing persona-oriented dialogue systems can be classified into two categories: Structured Persona-oriented Dialogue Systems (SPDS) [19, 32, 33] and Unstructured Persona-oriented Dialogue Systems (UPDS) [24, 31]. The former directly uses structured user persona descriptions in the form of key-value pairs (e.g., $\langle SEX, M \rangle, \langle AGE, 18 \rangle$), whereas the latter mines user persona descriptions from natural language utterances (e.g., “I like music.”, “I like the guitar.”, “I am a vegan.”). In this work, we focus on UPDS.

¹ Shandong University, China, email: minghongxu.sd@outlook.com

² Tencent AI Lab, China, email: pijili@tencent.com

³ The Chinese University of Hong Kong, email: hryang@se.cuhk.edu.hk

⁴ University of Amsterdam, The Netherlands, email: p.ren@uva.nl

⁵ Shandong University, China, email: zhaochun.ren@sdu.edu.cn

⁶ Shandong University, China, email: chenzhumin@sdu.edu.cn

⁷ Shandong University, China, email: majun@sdu.edu.cn

* Piji Li and Zhaochun Ren are corresponding authors.

Table 1. An example of unstructured persona-oriented dialogue system.

| | |
|------------------------|--|
| Personas for Speaker B | 1. I like music. 2. I like to skateboard. 3. I like the guitar. 4. I am a vegan . |
| Dialogue | A (u ₁): Wanna come over and watch the godfather? B (u ₂): I do not have a car, I have a skateboard. A (u ₃): You can skateboard over. I do not live too far. I have candy and soda to share. B (u ₄): No thanks, I do not eat any animal products. A (u ₅): I promise there are no animal products in my candy and soda. B (u ₆): Most candy has some form of dairy . As a vegan I can not have that. |

Recently, there have been some studies which utilize the predefined user persona descriptions to generate persona-oriented responses [10, 24, 31]. However, the given user persona descriptions are mostly short and limited to only a few descriptive words. As a result, the existing methods have a hard time utilizing the user persona descriptions when generating responses. On the one hand, they might fail to use user persona descriptions. For example, the generative profile memory network proposed in [31] simply attends over encoded persona description in decoder. It generates response “I have a lot of candies. I am not sure.” for the case in Table 1 without considering user persona. On the other hand, they cannot use user persona descriptions properly sometimes. For example, the persona-CVAE proposed in [24] uses force decoding strategy to copy persona description. It generates response “I like to skateboard. What are your hobbies?” for the case in Table 1 which use the selected persona improperly and seriously affects its quality. The reason is that with the limited descriptive words, it is hard for these models to understand and correlate the user persona descriptions when generating responses. We argue that this could be alleviated by extending the predefined user persona descriptions with semantically correlated content. As shown in Table 1, the target is to generate the last utterance (u₆) based on the given persona descriptions and historical utterances (u₁-u₅). One of the user persona descriptions for Speaker B is “I am a vegan”. However, only using this user persona description is not enough to generate human-like response (u₆) because “vegan” and “candy” are not directly related. In order to generate u₆, we need to take the following content into consideration simultaneously: (1) the word “vegan” in B’s user persona description; (2) the semantic correlation between “vegan” and “dairy”; (3) speaker A mentioned “animal products” and “candy” in the query utterance; (4) the correlation among “dairy”, “animal products”, and “candy”.

In this work, we propose a neural topical expansion framework, namely Persona Exploration and Exploitation (PEE), which is able to extend the predefined user persona descriptions with semanti-

cally correlated content before utilizing them to generate dialogue responses. PEE consists of two main modules: persona exploration and persona exploitation. The former learns to extend the predefined user persona descriptions by mining and correlating with existing dialogue corpus. Specifically, we employ a VAE-based topic model to conduct the unsupervised semantic modeling and extend persona-related words by semantic matching. The latter learns to generate persona consistent responses by utilizing the predefined and extended persona information. Specifically, we design a mutual-reinforcement multi-hop memory retrieval mechanism which retrieves information from two types (predefined and extended) of personas by considering their mutual influence. Furthermore, in order to make persona exploitation learn to utilize user persona descriptions more properly, we also introduce two persona-oriented loss functions: P-Match loss and P-BoWs loss. P-Match loss supervises the choice of predefined persona sentences in encoder. P-BoWs loss supervises to generate more persona-related words in decoder.

The main contributions of this paper are as follows:

- We propose a persona exploration and exploitation (PEE) framework which can explore and exploit persona information to generate informative persona-oriented responses.
- We employ an VAE-based topic model to conduct the efficient unsupervised semantic learning for external persona information mining and distillation.
- We propose two learning strategies for persona exploitation: a mutual-reinforcement multi-hop memory retrieval mechanism and two persona-oriented loss functions.

2 Related work

As a challenging task in the area of natural language processing, open-domain dialogue system has attracted great attention of researchers recently [15, 20, 22, 27]. But there are still some limitations and challenges in this area. Among the many issues, the lack of consistency is one of the most challenging difficulties. Therefore, persona-based dialogue system has been proposed to generate persona consistent and human-like responses [5, 16, 17, 26, 30]. Li et al. [8] learn a user embedding to represent persona implicitly for each user without using explicitly persona information. Then, researchers model user embedding with explicit persona information to generate responses. According to the format of persona information, those methods can be classified into two categories: structured Persona-oriented Dialogue Systems (SPDS) and Unstructured Persona-oriented Dialogue Systems (UPDS).

In SPDS, Wang et al. [28] group users according to the gender attribute and the dialogue features in the same group can be shared. Qian et al. [19] endow the user with explicit structured persona information (a key-value table) and design a profile detection module to select a persona information and inject it to the decoding process. Luo et al. [12] encode user persona description into distributed embeddings and take advantage of conversation history from other users with similar profiles, their model can adopt different recommendation policy based on the user profile. Due to the lack of large scale persona-labelled dataset, Zheng et al. [32] introduce a dataset where persona information is formulated to key-value pairs from dialogue content and they devise two technique to capture and address trait-related information. In UPDS, Zhang et al. [31] contribute a persona-chat dataset with natural sentences persona information and they propose a generative profile memory network to incorporate persona information into responses. Lin et al. [10] model learning different personas as different tasks via meta-learning algorithm

without using explicit persona information since dialogue itself can reflect some persona information. Through this way, their model can generate personalised responses by leveraging only a few dialogue samples instead of human-designed persona descriptions. To generate diverse and sustainable conversations, Song et al. [24] propose a memory-augmented architecture to exploit persona information and utilized a conditional variational autoencoder which can address the one-to-many generation problem.

Prior studies are trained purely on predefined persona corpus, but the limited information leads to generating uninformative responses. Different from them, we employ VAE-based topic model to extend persona information and propose two strategies (a mutual-reinforcement multi-hop memory retrieval mechanism and two persona-oriented loss functions) to integrate persona information to responses.

3 Method

3.1 Overview

We assume that a conversation is conducted between two users. Given a target user, we denote the user’s persona descriptions as $P = (P_1, P_2, \dots, P_{n_p})$. Each persona sentence P_j is formulated as $P_j = (p_1^j, p_2^j, \dots, p_{l_p}^j)$, where p_i^j refers to a word. Suppose there are already k turns in a dialogue, so we have historical utterances as $X = (X_1, X_2, \dots, X_k)$, where each utterance X_i is depicted as $X_i = (x_1^i, x_2^i, \dots, x_{l_x}^i)$ and x_j^i denotes a word. Accordingly, *unstructured persona-oriented dialogue generation* aims to predict the $(k+1)$ -th utterance, i.e., the response $Y = (y_1, y_2, \dots, y_{l_y})$, according to the predefined persona descriptions P and the historical utterances X :

$$p(Y|X, P) = \prod_{i=1}^{l_y} p(y_i|X, P, y_1, \dots, y_{i-1}). \quad (1)$$

As illustrated in Figure 1, our PEE framework mainly consists of two stages: persona exploration and persona exploitation. Persona exploration employs a VAE-based topic model to conduct the unsupervised semantic modeling and obtains topic-relevant word representations. Then it extends persona-related words by semantic matching based on the predefined persona descriptions. Persona exploitation contains three components: (1) multi-source sequence encoder, which encodes predefined persona descriptions into two kinds of key-value memories and encodes historical utterances into hidden vectors; (2) persona information retrieval, which selects predefined persona descriptions based on historical utterances and considers the impact of personalized information involved in the history; (3) persona-oriented response decoder, which exploits the predefined and explored external persona information to generate responses based on the specially designed mutual-reinforcement multi-hop memory retrieval mechanism. Moreover, two new optimization objectives, persona-oriented matching loss (P-Match) and persona-oriented bag-of-words loss (P-BoWs), are proposed to impel our model to exploit the persona information more precisely. We will introduce the technical details in the following sections.

3.2 Persona Exploration

Based on the predefined persona descriptions, the target of the persona exploration stage is to extend more persona-related words. Therefore the key is to investigate an effective method for the semantic learning of words. The extended persona words must lie in the

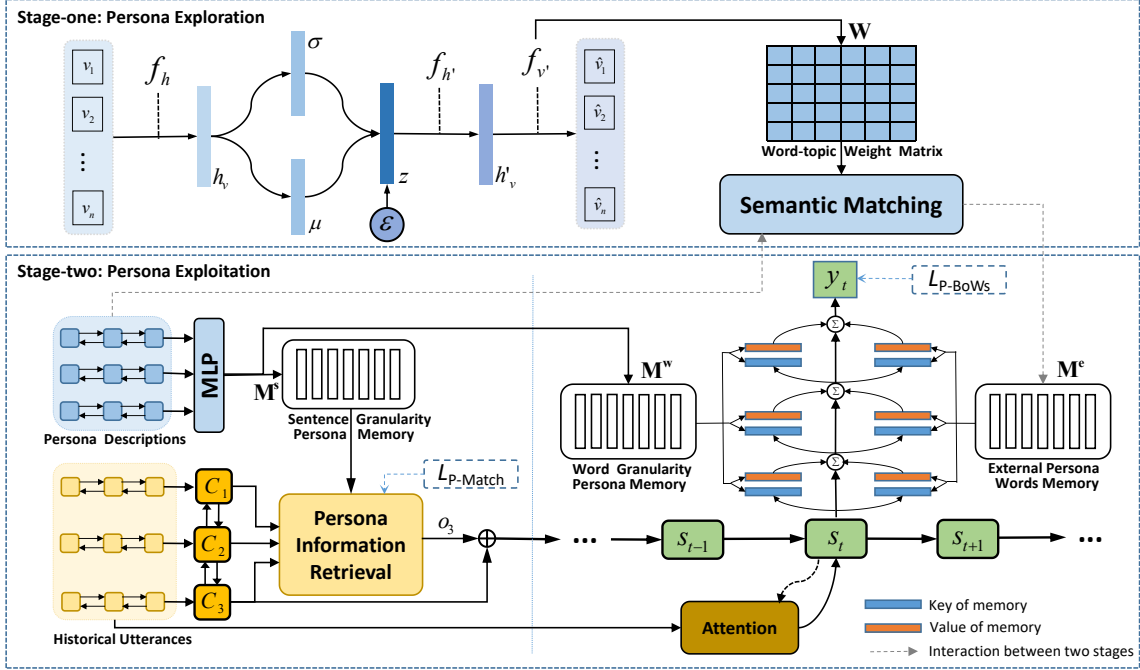


Figure 1. An overview of our PEE framework. It consists of two stages: persona exploration and persona exploitation (multi-source sequence encoder, persona information retrieval and persona-oriented response decoder).

same topic with the predefined persona information, so as to guarantee the topic consistency of the conversations. Topic modeling methods [25] are appropriate. Therefore, inspired by [23], we employ a topic model based on variational auto-encoder (VAE) [7] and make adjustments according to our task to conduct the unsupervised global semantic modeling. Compared to the traditional methods such as Latent Dirichlet Allocation (LDA) [2], VAE-based topic model is less time-consuming for training and more flexible for inferring latent representations for new documents.

As shown in the upper side of Figure 1, the input of VAE-based topic model is a document presentation v and the output v' is the reconstruction of input. We regard each conversation as a document and represent it by *tf-idf* features. The encoding process can be formalized as:

$$\begin{aligned} h_v &= f_h(v), \\ \mu &= f_\mu(h_v), \quad \log(\sigma^2) = f_\sigma(h_v), \\ z &= \mu + \sigma \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1), \end{aligned} \quad (2)$$

where $f_*(\cdot)$ denotes non-linear transformation, μ and σ are mean and standard deviation vectors of multivariate normal distribution respectively, z is a latent vector sampled from the multivariate normal distribution by reparameterization trick. We use the latent vector z to reconstruct the document:

$$\begin{aligned} h'_v &= f_{h'}(z), \\ v' &= f_{v'}(h'_v). \end{aligned} \quad (3)$$

We learn all parameters via optimizing the evidence lower bound (ELBO) [7]. After the training, we draw a word-topic weight matrix $\mathbf{W} \in \mathbb{R}^{K \times |V'|}$ from the output layer $f_{v'}$. The matrix represents the topical saliency for each word, where K is the number of topics, V' is the vocabulary of topic model and $|V'|$ is the vocabulary size. Each column $u \in \mathbb{R}^K$ in \mathbf{W} can be regarded as a topic-based representation for the corresponding word.

Given topic-relevant word representations, we extend words for every dialogue. After removing stop-words in the dialogue, we filter

a vocabulary set $V^P \subset V'$ from predefined persona descriptions. For each word $w \in V^P$, we select most relevant m external words based on cosine similarities of topic-relevant word representations. Then, we re-rank all external persona words of V^P according to the cosine similarity score. If an external word is selected more than once, we just record the highest score. Thereafter, we select the top n_w words of them. Finally, we convert each extended persona words into key-value representation by two multi-layer perceptron neural networks and store these representations in an **external persona words memory** M^e .

3.3 Persona Exploitation

Given predefined persona descriptions and extended persona-relevant words, persona exploitation aims to integrate them to generate informative responses. In this section, we detail three components of persona exploitation stage: multi-source sequence encoder, persona information retrieval and persona-oriented response decoder.

Multi-Source Sequence Encoder. The input contains persona descriptions and historical utterances, we design two independent encoders for them.

Persona memory encoder. We encode the predefined persona information into sentence and word granularity presentations and store them in two memories respectively. For each sentence P_i , we obtain a sentence representation e_{P_i} by a bidirectional Gated Recurrent Networks (Bi-GRU) [4]. Then we convert e_{P_i} into key m_i^S and value c_i^S by two multi-layer perceptron neural networks, and store them in the sentence granularity persona memory M^s . Simultaneously, for word p_j^i , we obtain word representation $e_{p_j^i}$ from the j -th step of Bi-GRU for the i -th sentence. Same as above, we convert each word representation into key and value, and store them in the word granularity persona memory M^w .




Figure 2. An overview of Persona Information Retrieval.

Historical utterances encoder. In order to capture the relationship among the historical utterances X , we use a hierarchical recurrent encoder [21] to conduct the semantic modeling. From the second level of the hierarchical Bi-GRU, we obtain the final representations e_X for the whole historical utterances and a sentence vector C_i for each utterance X_i .

Persona Information Retrieval. After obtaining the representations of historical utterances and sentence granularity persona memory via the previous component, we use historical utterances to select persona information for response. Considering the key-value memory retrieval is a frequent component in the following modules, we provide the general definition here. Assume that query vector is q and memory \mathbf{M} contains key m and value c , the retrieval operation $\text{retri}(q, \mathbf{M}) = o$ is defined as:

$$\begin{aligned} o &= \sum_i a_i c_i, \\ a_i &= \frac{\exp(s_i)}{\sum_j \exp(s_j)}, \\ s_j &= q^T m_j, \end{aligned} \quad (4)$$

where the output vector o is a weighted sum of values in memory and represents retrieved information.

As shown in Figure 2, we use each historical utterances to retrieve user’s persona information in turn. During the chat process, some of persona information used in the history has an impact on the choice of persona information for response. In order to take advantage of this impact, in the i -th step, we combine historical utterance presentation C_i and result of previous retrieval step o_{i-1} as query vector q_i , which can be formalized as:

$$q_i = \begin{cases} C_i, & i = 1; \\ C_i + o_{i-1}, & i > 1. \end{cases} \quad (5)$$

Then we retrieve the sentence granularity persona memory \mathbf{M}^s by query vector q_i :

$$o_i = \text{retri}(q_i, \mathbf{M}^s). \quad (6)$$

Finally, we concatenate the result of last retrieval step o_k and the whole historical utterances representations e_X :

$$s_0 = [e_X; o_k], \quad (7)$$

where s_0 is a merged vector used as the initial state of decoder.

Persona-Oriented Response Decoder. The decoder is a GRU based sequence prediction framework with an attention mechanism on the historical utterances and a mutual-reinforcement multi-hop

memory retrieval mechanism. Given the current input y_{t-1} as well as the previous hidden state s_{t-1} , the recurrent calculation of GRU is defined as:

$$s_t = \text{GRU}(y_{t-1}, s_{t-1}). \quad (8)$$

Then we design an attention mechanism to absorb relevant information from historical utterances and a mutual-reinforcement multi-hop memory retrieval mechanism to obtain the relevant persona information from predefined and explored external persona information.

Attention on the historical utterances. produces a historical utterances vector u^X at each decoding step by attending to historical utterances. We formalize it as:

$$\begin{aligned} u^X &= \sum_{i=1}^n a_i h_i^X, \\ a_i &= \frac{\exp(s_i)}{\sum_{j=1}^n \exp(s_j)}, \\ s_j &= v^T \tanh(W_s s_t + W_t h_j^X + b), \end{aligned} \quad (9)$$

where h_j^X ($j = 1, 2, \dots, n$) is the j -th word hidden state of historical utterances obtained from the first level of the hierarchical encoder for X .

Mutual-reinforcement multi-hop memory retrieval. Recall that we build an external persona words memory \mathbf{M}^e for persona exploration and a word granularity persona memory \mathbf{M}^w in encoder. There is an association between the two memories. For example, if we retrieve a word in the predefined persona descriptions which is related to current conversation, the information in external persona memory related to this word will be more likely to be applied, vice versa. Therefore, the results of the two types of persona information retrievals are mutually influential and we propose a mutual-reinforcement multi-hop memory retrieval mechanism to model this influence.

First, we use the current hidden state s_t as query vector q to retrieve \mathbf{M}^w and \mathbf{M}^e respectively:

$$\begin{aligned} o^w &= \text{retri}(q, \mathbf{M}^w), \\ o^e &= \text{retri}(q, \mathbf{M}^e). \end{aligned} \quad (10)$$

Considering that the result of one memory retrieval (e.g. o^w) will affect the next retrieval of another memory (e.g. \mathbf{M}^e), we update query vector by adding the two retrieved results o^w and o^e :

$$q^{new} = q^{old} + o^w + o^e. \quad (11)$$

This update means that the results of the two retrievals will affect each other in the next hop. In our experiment, we use three hops unless otherwise stated.

Finally, based on the exploitation of predefined and extended persona information, the output word distribution p_{y_t} at time step t of the decoder is produced by:

$$\begin{aligned} \tilde{s}_t &= f_o([s_t; u^X; o^w; o^e]), \\ p_{y_t} &= \text{softmax}(\tilde{s}_t), \end{aligned} \quad (12)$$

where f_o is the neural non-linear operation on the output layer.

3.4 Persona-Oriented Loss.

In order to impel the model to exploit the persona information more precisely, besides the Negative Log-Likelihood loss (NLL), we propose two new persona-oriented loss functions: Persona-oriented Matching loss (P-Match) and Persona-oriented Bag-of-Words loss (P-BoWs). P-Match loss supervises the choice of predefined persona sentences in persona information retrieval module and P-BoWs loss supervises to generate more persona-related words in decoder.

P-Match Loss. Recall that in the persona information retrieval module (Eq. 6), we can get a match weight over the sentence granularity persona memory \mathbf{M}^s in every step. Assume that the match weight in the last step is $a^s \in \mathbb{R}^{|P|}$. Intuitively, if the ground truth response contains the information from persona sentence P_i , then a_i^s should obtain a large value. Is it possible to employ the relation between the ground truth response and the persona sentences to improve the modeling of persona information retrieval? To tackle this, we design the persona-oriented matching loss (P-Match). The 0-1 label $a \in \mathbb{R}^{|P|}$ is decided based on a threshold θ_a of the similarity between the persona sentences and the ground truth response. Jaccard Index⁸ is employed for the similarity calculation. The P-Match loss is defined as:

$$\mathcal{L}_{P-Match} = - \sum_{i=1}^{|P|} a_i \log a_i^s. \quad (13)$$

P-BoWs Loss. Inspired by [13], we design a persona-oriented Bag-of-Words loss function to enhance the ability of persona information capturing. Specifically, We label each response with a vocabulary-size vector $b \in \mathbb{R}^{|V|}$, where the non-stop words in the current response will get values 1. If words are persona-based information, we increase the weight to $1 + \lambda$, where λ is a positive value. We use a multi-label classifier to generate BoWs representation P_b (sentence-level probability) by summing the scores of all positions of the generated sentence in decoder: $p_b = \text{sigmoid}(\sum_{t=1}^{|Y|} \tilde{s}_t)$. We define P-BoWs loss using cross entropy:

$$\mathcal{L}_{P-BoWs} = - \frac{1}{|V|} \sum_{i=1}^{|V|} [b_i \log p_{b_i} + (1 - b_i) \log (1 - p_{b_i})]. \quad (14)$$

3.5 Joint Training

Negative log-likelihood loss (NLL) is employed as the basic optimization objective:

$$\mathcal{L}_{NLL} = - \frac{1}{|Y|} \sum_{t=1}^{|Y|} y_t \log p_{y_t}. \quad (15)$$

Finally, a unified optimization objective is designed by integrating P-Match loss, P-BoWs loss and the NLL loss:

$$\mathcal{L} = \mathcal{L}_{NLL} + \gamma_1 \mathcal{L}_{P-Match} + \gamma_2 \mathcal{L}_{P-BoWs}, \quad (16)$$

where γ_1 and γ_2 are trade-off parameters controlling the balance between three loss functions.

4 Experiments

In this section, we first introduce two datasets used in our experiment and list setups and baseline models. Next, we evaluate the performance of various models by automated evaluation and human evaluation.

4.1 Datasets

Our experiments use two public multi-turn dialogue datasets: Persona-Chat⁹ [31] and DailyDialog¹⁰ [9]. The Persona-Chat

dataset contains 10907 dialogues between pairs of speakers, where 968 dialogues are set aside for validation and 1000 for testing. Each speaker is described by 3-5 persona sentences. (e.g. “I like reading.” or “I am a nurse.”, etc). The total number of personas is 1155, and 100 personas for validation and 100 for testing. The DailyDialog dataset is constructed by raw data crawled from various websites, which serves for English learners to practice English dialog in daily life. It contains 13,118 multi-turn dialogues without persona descriptions, the number of turns are roughly 8 and the average number of tokens of each utterance is about 15.

Our experiments are performed on the persona-chat dataset. In order to expand the knowledge space, we merge DailyDialog and the training set of Persona-Chat as the basic knowledge source to pre-train the topic model for persona exploration.

4.2 Baselines

We consider the following comparison methods and their inputs consist of predefined persona descriptions, historical conversation utterances and the current query utterance.

Seq2Seq [1]: the standard Sequence-to-Sequence Model with attention. We concatenate persona descriptions and historical utterances as a sequence input and generate the response.

HRED [21]: Hierarchical Recurrent Encoder-Decoder model with attention. The input contains all sentences in persona and history conversation.

Profile Memory [31]: Generative Profile Memory network is a generative model that encodes each of persona descriptions as a individual memory representation in a memory network.

Per-CVAE [24]: Persona-CVAE is a memory-augmented architecture which focus on the diverse generation of conversational responses based on chatbot’s persona. In our experiment, we sample one time from the latent z to generate a response.

PED: Persona-oriented Encoder-Decoder Model. i.e., our PEE framework without the persona exploration, the P-BoWs loss and the P-Match loss. Without external persona words memory, mutual-reinforcement multi-hop memory retrieval mechanism is equivalent to normal multi-hop memory retrieval mechanism.

PED+PE: Our PEE framework without the P-BoWs loss and P-Match loss.

PED+PE+P-BoWs: Our PEE framework without the P-Match loss.

PED+PE+P-Match: Our PEE framework without the P-BoWs loss.

PEE: PED + PE + P-BoWs + P-Match, i.e., our proposed PEE framework.

4.3 Experimental Settings

We treat each complete dialog (including personas) as a document, remove the stop words and select the top 10,000 frequent words to train the VAE-based topic model. For the number of topics, we follow previous settings [29], [23] to set $K = 50$. In our experiments, we use GloVe [18] for word embedding and employ bi-directional GRU for encoders, and we set hidden states size is 512 and batch size is 64. We use Adam optimizer [6] to train the model and set the learning rate is 0.0001. For testing, we use beam search with beam size 2. For all the other hyperparameters, we tune them on the development set by grid search. The number of extended persona-related words for each dialogue n_w is 100. The additional weight λ in the P-BoWs target is 1 and threshold θ_a in labeling P-Match target process is 0.03. During training, trade-off parameter γ_1 is 0.1 and γ_2 is 0.1.

⁸ http://en.wikipedia.org/wiki/Jaccard_index

⁹ <https://github.com/facebookresearch/ParlAI/tree/master/projects/personachat>

¹⁰ <http://yanran.li/dailydialog>

Table 2. Automatic evaluation results. The best results are bold.

| Model | BLEU1 | BLEU2 | BLEU3 | BLEU4 | F1 | Average | Extrema | Greedy |
|----------------|----------------|----------------|---------------|---------------|----------------|---------------|---------------|---------------|
| Seq2Seq | 20.1381 | 9.9395 | 5.2887 | 2.9840 | 17.7972 | 0.8551 | 0.4980 | 0.6751 |
| HRED | 19.0920 | 9.5668 | 5.0191 | 2.7779 | 17.9184 | 0.8531 | 0.4882 | 0.6714 |
| Profile Memory | 20.8713 | 9.8526 | 4.9942 | 2.6852 | 17.1553 | 0.8675 | 0.4835 | 0.6752 |
| Per.-CVAE | 17.2315 | 7.2602 | 3.2081 | 1.4541 | 14.6121 | 0.8458 | 0.4688 | 0.6516 |
| PED | 21.4611 | 10.6992 | 5.7845 | 3.3344 | 18.4759 | 0.8593 | 0.4993 | 0.6838 |
| PED+PE | 21.8970 | 10.9987 | 5.9965 | 3.5334 | 18.4140 | 0.8643 | 0.4999 | 0.6856 |
| PED+PE+P-BoWs | 21.9768 | 11.0710 | 6.0154 | 3.5574 | 18.2781 | 0.8626 | 0.4986 | 0.6822 |
| PED+PE+P-Match | 22.4668 | 11.2560 | 5.9846 | 3.3031 | 18.2615 | 0.8592 | 0.4940 | 0.6803 |
| PEE | 23.1926 | 11.5166 | 6.1248 | 3.4977 | 18.4130 | 0.8691 | 0.5010 | 0.6906 |

Table 3. Human evaluation on four aspects: Fluency, Engagingness, Consistency and Persona Detection (PD). The value in parentheses is standard deviation.

| Model | Fluency | Engagingness | Consistency | PD(%) |
|----------------|-------------------|-------------------|-------------------|--------------------|
| Seq2Seq | 4.08(0.71) | 3.02(0.96) | 3.00(1.03) | 52.94(0.32) |
| HRED | 3.96(0.71) | 2.73(1.05) | 2.60(1.16) | 64.71(0.32) |
| Profile Memory | 4.04(0.68) | 3.08(1.01) | 3.10(1.10) | 58.82(0.40) |
| Per.-CVAE | 3.61(1.02) | 2.63(1.09) | 2.78(1.29) | 85.29(0.34) |
| PEE | 4.13(0.76) | 3.46(1.07) | 3.44(1.13) | 76.47(0.36) |

Table 4. Automatic evaluation results of PEE with different hops in mutual-reinforcement multi-hop retrieval mechanism.

| Hops | BLEU1 | BLEU2 | BLEU3 | BLEU4 | F1 | Average | Extrema | Greedy |
|-------|----------------|----------------|---------------|---------------|--------------|---------------|---------------|---------------|
| PEE-1 | 22.5956 | 11.2877 | 6.0405 | 3.4315 | 18.32 | 0.8631 | 0.5009 | 0.6902 |
| PEE-2 | 22.9758 | 11.4999 | 6.2383 | 3.6327 | 18.68 | 0.8654 | 0.4979 | 0.6861 |
| PEE-3 | 23.1926 | 11.5166 | 6.1248 | 3.4977 | 18.41 | 0.8691 | 0.5010 | 0.6906 |
| PEE-4 | 22.3422 | 11.0628 | 5.8804 | 3.3678 | 18.25 | 0.8618 | 0.4985 | 0.6824 |
| PEE-5 | 22.2892 | 11.1789 | 5.9878 | 3.4148 | 18.55 | 0.8591 | 0.4993 | 0.6811 |

4.4 Evaluation Metrics

We use different evaluation metrics (automated and human) to demonstrate the effectiveness of our model. In this subsection, we will give a brief introduction to those metrics.

Automatic Metrics. We report three different automatic metrics: **BLEU@N**: BLEU is an algorithm which has been widely used in machine translation and dialogue system to evaluate the quality of the generated text. It measures the N-gram overlap between the generated response and ground truth.

F1-Measure: It measures the accuracy of the generated response considering both the precision and the recall. We treat the predicted and target response as bags of tokens, and compute their F1 score.

Embedding-based similarity: Embedding Average (Average), Embedding Extrema (Extrema), and Embedding Greedy (Greedy) [11]. These embedding-based metrics measure semantic similarity between the generated response and the ground truth.

Human Metrics. It is not enough to only automatically evaluate dialogue systems, so we randomly sample about 100 dialogues from test data and hire 5 volunteers to evaluate. We use four metrics: fluency, engagingness, consistency, and persona detection.

Fluency: It measures the quality of generated sentence, e.g., whether grammar is correct.

Engagingness: It measures whether the generated sentence is appropriate and interesting.

Consistency: It measures whether the generated sentence has some relationships with the history and persona description.

Persona detection: For each dialogue, given generated responses and two set of persona sentences (one is real and another is fake),

we ask students to choose which one is the real description of chatbot.

The first three metrics are scored between 1-5. For persona detection, Scoring 1 means the choice is correct and 0 means the choice is wrong and 0.5 means people can not judge.

5 Results and Analysis

5.1 Experimental Results and Ablation Study

Automatic evaluation. Comparative automatic evaluation results are presented in Table 2. Our model outperforms baselines on all automatic metrics. This demonstrates that our model generates more appropriate responses by persona exploration and exploitation. Especially, our model improves approximately 15.17% over seq2seq on BLEU1. Comparing with PED, PED+PE has better scores on most metrics. This is because explored persona information contributes to generate more informative responses. Comparing with PED+PE, both PED+PE+p-BoWs and PED+PE+P-Match perform better, which is because P-BoWs loss and P-Match loss supervise model to exploit persona information more precisely.

According to automatic evaluation results of PEE with different hops in mutual-reinforcement multi-hop retrieval mechanism in Table4, PEE-2 outperforms PEE-1 on most metrics. This demonstrates the interaction between two types persona information improves the performance of our model. Analyzing various indicators, the PEE works best when hop is 3. When the number of hops exceeds 3, the effect drops, that may because the query vector contains little information of current decoder state s_t after several update operations.

Human evaluation. The results of human evaluation are listed in Table 3. Our model significantly outperforms most of the baselines in terms of all the metrics. Particularly, our model increases approximately 30.01% over profile memory on persona detection. This demonstrates that persona exploration and exploitation is beneficial to improve the usage of persona information and enrich the responses. Per.-CVAE has the highest persona detection metric, but it pays too much attention to persona, resulting in very poor grammar, relevance, and fluency of the generated responses.

5.2 Persona Analysis

In order to further evaluate the ability of the model to combine persona, for each multi-turn dialogue, we count the number of words that appear in both persona sentences and generated responses, and divide the number by the total number of words in persona sentence to get **Persona use ratio**. It measures the probability of persona

Table 5. Case studies.

| | Case 1 | Case 2 |
|------------------------|---|---|
| Personas for Speaker B | I worked at hollister in high school. I am a professional skater. I play bluegrass music. I do not like chips. | I write short stories for a living. I used to work at a carnival. I like to drink scotch to relax. I like dark superhero movies. |
| Historical Utterances | A: Hi. How is your night going? B: Good, just left a bluegrass concert. A: Ooh. Interesting. What else do you do for fun? B: Skate. I love it. What do you do? A: That is so exciting! I am currently a stay at home mom. B: cool, how many kids do you have? A: I have three kids and pregnant with my fourth. I love being a mom. | A: Hi, how are you doing tonight? I am good so far. B: I am good, relaxing with a glass of scotch to end the night. A: That is nice and relaxing I love to get a good workout. B: Yes. I have been busy working on a new short story to release. A: That sounds very interesting hope all is good for you. B: Yeah. What kind of movies are you interested in? A: I like to watch romance and some scary movies is okay for me. |
| Explored words | concert, band, piano, guitar, rap ... | ebook, thriller, horror, creepy, comic ... |
| Response | seq2seq: That is cool. Do you have any pets? HRED: That is cool. Do you have any pets? Profile Memory: I have a lot of kids, but i have a daughter that is so cool. Per.-CVAE: I am a coach I might play tomorrow. I do roofing. PEE: Wow that is a great thing. I like to play piano with my family. | seq2seq: I like movies too. I am a loner. HRED: That is cool. I am a fan of a movie. Profile Memory: I have been watching movies on TV. Per.-CVAE: I like movies too, I am not sure I like that. What are your hobbies? PEE: That is scary. I write a lot of horror stories. |




Figure 3. Average persona use ratio of all models in different turns.

words being used and punishes the repeated use of the same persona information in responses of different turns. We calculate the average persona use ratio of all models with different turns and show them in Figure 3. Per.-CVAE pays too much attention to persona, which seriously affects the quality of the generated responses. We do not consider Per.-CVAE here. We can see that our model outperforms all the baseline methods. There are three reasons: first, persona information retrieval module thinks about the influence of persona in history when selecting persona information; second, external persona words contribute to utilize persona description that has an indirect relationship with current topic. third, the P-BoWs loss and the P-Match loss encourage the model to generate more persona-related words.

5.3 Case Study

Table 5 depicts some cases generated by PEE, Seq2Seq, HRED, and Profile Memory. From the comparisons, we can see that PEE model can use explored persona information to generate more persona-oriented informative responses. For example, in case 1, one of persona descriptions for speaker B is “I play bluegrass music.”, speaker A mentioned “kids” and “mom” in the query utterance. The explored persona word “piano” is related to “music” and “family”, the word “family” has correlation with “kids” and “mom”. So, the response generated by PEE follows the clues above and implies persona information simultaneously. What’s more, it leads the topic to a new field that speakers are familiar with, making the next reply have more content to facilitate. When the previous topic “work” is drawing to an end, our model can use persona and its extended words to convert topics to “music”, however, responses of other baseline models do not reflect personality information.




Figure 4. Visualization of matching weights on external persona words memory in the last step of mutual-reinforcement multi-hop memory retrieval.

Similarly, in case 2, explored persona word “horror” is related with “scary movies” and “stories”, so our PEE model uses word “horror” to rich response. In order to show the contribution of explored persona words more clearly and directly, we show the matching weights on external persona words memory in the last step of mutual-reinforcement multi-hop memory retrieval mechanism in Figure 4.

6 CONCLUSION

In this work, we propose a neural topical expansion framework, namely Persona Exploration and Exploitation (PEE), for the unstructured persona-oriented dialogue systems. Different from previous work that trained purely on predefined persona description, our model extends external persona information by a VAE-based topic model. By fusing predefined persona descriptions and external persona information, the responses our model generates can more accurately and properly represent the user persona while maintaining the consistency of the dialogue. Experimental comparisons and analysis demonstrated that our approach outperforms a set of state-of-the-art baselines in terms of both automated metrics and human evaluations. For future work, we will extend persona information dynamically and jointly train persona exploration and exploitation.

7 ACKNOWLEDGEMENT

This work is supported by the Natural Science Foundation of China (61972234, 61902219, 61672324, 61672322), the Tencent AI Lab Rhino-Bird Focused Research Program (JR201932), the Foundation of State Key Laboratory of Cognitive Intelligence, iFLYTEK,

P.R.China (COGOSC- 20190003), the Fundamental Research Funds of Shandong University, Ahold Delhaize, the Association of Universities in the Netherlands (VSNU), and the Innovation Center for Artificial Intelligence (ICAI).

REFERENCES

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [2] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. In *Advances in Neural Information Processing Systems 14*, pages 601–608, 2002.
- [3] E. Chu, P. Vijayaraghavan, and D. Roy. Learning personas from dialogue with attentive memory networks. *arXiv preprint arXiv:1810.08717*, 2018.
- [4] J. Chung, Ç. Gülçehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [5] C. K. Joshi, F. Mi, and B. Faltings. Personalization in goal-oriented dialog. *arXiv preprint arXiv:1706.07503*, 2017.
- [6] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [7] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [8] J. Li, M. Galley, C. Brockett, G. Spithourakis, J. Gao, and B. Dolan. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, 2016.
- [9] Y. Li, H. Su, X. Shen, W. Li, Z. Cao, and S. Niu. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, 2017.
- [10] Z. Lin, A. Madotto, C.-S. Wu, and P. Fung. Personalizing dialogue agents via meta-learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5459, 2019.
- [11] C. Liu, R. Lowe, I. V. Serban, M. Noseworthy, L. Charlin, and J. Pineau. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*, 2016.
- [12] L. Luo, W. Huang, Q. Zeng, Z. Nie, and X. Sun. Learning personalized end-to-end goal-oriented dialog. *arXiv preprint arXiv:1811.04604*, 2018.
- [13] S. Ma, X. Sun, Y. Wang, and J. Lin. Bag-of-words as target for neural machine translation. *arXiv preprint arXiv:1805.04871*, 2018.
- [14] P. Mazaré, S. Humeau, M. Raison, and A. Bordes. Training millions of personalized dialogue agents. *arXiv preprint arXiv:1809.01984*, 2018.
- [15] C. Meng, P. Ren, Z. Chen, C. Monz, J. Ma, and M. de Rijke. Refnet: A reference-aware network for background based conversation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 2020.
- [16] K. Mo, Y. Zhang, S. Li, J. Li, and Q. Yang. Personalizing a dialogue system with transfer reinforcement learning. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence*, pages 5317–5324, 2018.
- [17] O. Olabiyi, A. Khazane, A. Salimov, and E. T. Mueller. An adversarial learning framework for A persona-based multi-turn dialogue model. *arXiv preprint arXiv:1905.01992*, 2019.
- [18] J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing*, pages 1532–1543, 2014.
- [19] Q. Qian, M. Huang, H. Zhao, J. Xu, and X. Zhu. Assigning personality/profile to a chatting machine for coherent conversation generation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4279–4285, 2018.
- [20] P. Ren, Z. Chen, C. Monz, J. Ma, and M. de Rijke. Thinking globally, acting locally: Distantly supervised global-to-local knowledge selection for background based conversation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 2020.
- [21] I. V. Serban, A. Sordoni, Y. Bengio, A. C. Courville, and J. Pineau. Hierarchical neural network generative models for movie dialogues. *arXiv preprint arXiv:1507.04808*, 2015.
- [22] I. V. Serban, A. Sordoni, R. Lowe, L. Charlin, J. Pineau, A. C. Courville, and Y. Bengio. A hierarchical latent variable encoder-decoder model for generating dialogues. *arXiv preprint arXiv:1605.06069*, 2016.
- [23] N. A. Smith, D. Card, and C. Tan. Neural models for documents with metadata. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018.
- [24] H. Song, W. Zhang, Y. Cui, D. Wang, and T. Liu. Exploiting persona information for diverse generation of conversational responses. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5190–5196, 2019.
- [25] M. Steyvers and T. Griffiths. Probabilistic topic models. In T. Landauer, D. McNamara, S. Dennis, and W. Kintsch, editors, *Latent Semantic Analysis: A Road to Meaning.*, 2006.
- [26] J. Urbanek, A. Fan, S. Karamcheti, S. Jain, S. Humeau, E. Dinan, T. Rocktäschel, D. Kiela, A. Szlam, and J. Weston. Learning to speak and act in a fantasy text adventure game. *arXiv preprint arXiv:1903.03094*, 2019.
- [27] O. Vinyals and Q. V. Le. A neural conversational model. *arXiv preprint arXiv:1506.05869*, 2015.
- [28] J. Wang, X. Wang, F. Li, Z. Xu, Z. Wang, and B. Wang. Group linguistic bias aware neural response generation. In *Proceedings of the 9th SIGHAN Workshop on Chinese Language Processing*, pages 1–10, 2017.
- [29] X. Yan, J. Guo, Y. Lan, and X. Cheng. A bitern topic model for short texts. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1445–1456, 2013.
- [30] M. Yang, Z. Zhao, W. Zhao, X. Chen, J. Zhu, L. Zhou, and Z. Cao. Personalized response generation via domain adaptation. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1021–1024, 2017.
- [31] S. Zhang, E. Dinan, J. Urbanek, A. Szlam, D. Kiela, and J. Weston. 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)*, pages 2204–2213, 2018.
- [32] Y. Zheng, G. Chen, M. Huang, S. Liu, and X. Zhu. Personalized dialogue generation with diversified traits. *arXiv preprint arXiv:1901.09672*, 2019.
- [33] Y. Zheng, R. Zhang, X. Mao, and M. Huang. A pre-training based personalized dialogue generation model with persona-sparse data. *arXiv preprint arXiv:1911.04700*, 2019.