Personalized Topic Selection Model for Topic-Grounded Dialogue

Shixuan Fan[1,2], Wei Wei*[1,2], Xiaofei Wen[1,2],
Xianling Mao[3], Jixiong Chen[4], Dangyang Chen* [2,5]
[1]Cognitive Computing and Intelligent Information Processing (CCIIP) Laboratory,
School of Computer Science and Technology, Huazhong University of Science and Technology
[2]Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL)
[3]Department of Computer Science and Technology, Beijing Institute of Technology
[4]Brilliance Technology Co. Ltd.
[5]Ping An Property & Casualty Insurance company of China
fanshixuan, weiw, xfwen}@hust.edu.cn,
maoxl@bit.edu.cn, chenjixiong@brilliance.com.cn,
chendangyang273@pingan.com.cn

Abstract

Recently, the topic-grounded dialogue (TGD) system has become increasingly popular as its powerful capability to actively guide users to accomplish specific tasks through topic-guided conversations. Most existing works utilize side information (*e.g.,* topics or personas) in isolation to enhance the topic selection ability. However, due to disregarding the noise within these auxiliary information sources and their mutual influence, current models tend to predict user-uninteresting and contextually irrelevant topics. To build user-engaging and coherent dialogue agent, we propose a **P**ersonalized topic s**E**lection model for **T**opic-grounded **D**ialogue, named **PETD**, which takes account of the interaction of side information to selectively aggregate such information for more accurately predicting subsequent topics. Specifically, we evaluate the correlation between global topics and personas and selectively incorporate the global topics aligned with user personas. Furthermore, we propose a contrastive learning based persona selector to filter out irrelevant personas under the constraint of lacking pertinent persona annotations. Throughout the selection and generation, diverse relevant side information is considered. Extensive experiments demonstrate that our proposed method can generate engaging and diverse responses, outperforming state-of-the-art baselines across various evaluation metrics.

1 Introduction

The dialogue systems have been widely used for task-specific interactions like customer service (Zhao et al., 2023; Li et al., 2023b) or emotional interaction (Wei et al., 2019; Lai et al., 2023). Early works adopt a passive stance in response to user queries, yielding generic or irrelevant responses (Liu et al., 2022a,b; Lu et al., 2023). However, real-world conversational scenarios often manifest heightened complexity, necessitating dialogue agents to adeptly manage topics and actively guide conversations. Although large language models, such as ChatGPT, exhibit closely resemble abilities of humans, they fall short in topic management and exhibit suboptimal initiative (Cao, 2023; Hudeček and Dušek, 2023). Consequently, topic-grounded dialogue systems (TGDs), which can proactively predict appropriate future topics and generate diverse and informative responses around new topics, have recently attracted considerable attention.

Indeed, the core of topic-grounded dialogue systems lies in the effective exploitation of diverse side information (*i.e.,* global topics or user personas) to precisely predict subsequent topics. The former (Xu et al., 2021; Zou et al., 2021; Wen et al., 2022) fully models the topic transitions overall topic sequences and infers the subsequent topic via the co-occurrences of adjacent topics over such topic transferring information, and the latter (Zhou et al., 2020; Ren et al., 2022) takes account of all user personas to model user preferences for topic selection.

Despite achieving promising results, previous studies still face challenges in adequately modeling side information. This limitation arises from inadequate consideration of the interconnections between different side information, leading to the indiscriminate integration of both relevant and irrelevant information for topic selection. This issue can be delineated from two perspectives. ***Firstly***, users with different personas may select different topics based on the same history topics. As shown in Figure 1 (a), global co-occurrence topic frequency dis-

* Corresponding author.

Findings of the Association for Computational Linguistics: ACL 2024, pages 7188–7202
August 11-16, 2024 ©2024 Association for Computational Linguistics

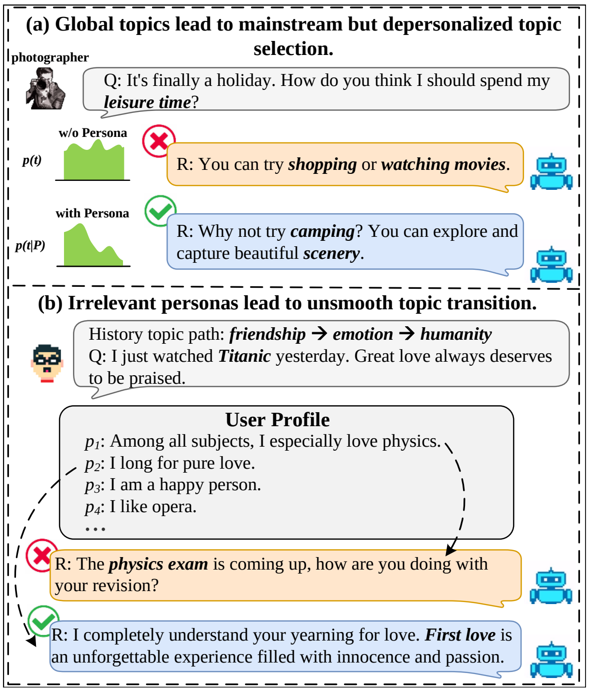

Figure 1: The irrelevant information in global topics and user personas misleads the model to choose depersonalized and unsmooth topics. The topics are bolded and italicized.

tribution exhibits a uniform profile due to its amalgamation of diverse persona users choosing topics. Existing works indiscriminately aggregate global topics resulting in biasing models towards mainstream but depersonalized topics. ***Secondly***, each selected topic may exclusively align with specific user personas, rather than all of them. As shown in Figure 1 (b), existing methods indiscriminately encode multiple personas in user profiles, resulting in models that may be misled by contextually irrelevant personas and choose unsmooth topics. Furthermore, the absence of annotated labels indicating the correspondence between personas and topics introduces additional challenges for modeling the interaction among side information and filtering out irrelevant information.

To address these problems, we propose a novel method called **P**ersonalized topic s**E**lection model for **T**opic-grounded **D**ialogue (**PETD**), which effectively mitigates the impact of irrelevant side information by considering the interaction among topics and personas for more accurately predicting subsequent topics. To disentangle the global topics that emerge from overlapping numerous personas, we establish the n-to-n correspondence between topics and personas at global level. We further selectively aggregate global topics aligned with user personas for efficiently filtering out irrelevant global topics to user personas and more accurate user preference modeling. To prevent the intrusion of context-irrelevant personas during topic selection, we employ a persona selector to predict personas likely to be manifested by the user in the next turn based on the topic path and dialogue context. In light of existing research, which indicates challenges in ensuring optimal training for each submodule through end-to-end single-task supervised training alone (Zhang and Yang, 2021), we design a contrastive learning based auxiliary task to address the impediment of lacking relevant persona annotations for targeted training of the persona selector submodule. This task aims to specifically optimize the persona selector and amplify the representation distinction between different personas. Finally, we utilize side information after filtering out irrelevant information to assist topic selection and dialogue generation.

To summarize, the contributions of this paper are listed as follows:

- We identify the problem in current paradigms, which is the isolated and indiscriminate integration of side information, resulting in depersonalized and unsmooth topic selection.
- We propose a persona-specific global topic expansion method to selectively aggregate global topics under different user personas.
- We exploit fine-grained personas to guide topic selection and a contrastive learning based auxiliary task is proposed to optimize the persona selector and enhance the distinction of different personality representations.

2 Related Work

Nowadays, existing methods in TGD often share a paradigm that decomposes the task into two related sub-task, namely topic selection and response generation (Qin et al., 2020; Xu et al., 2020a, 2021; Zou et al., 2021). In this work, we mainly focus on the topic selection task, which can be broadly categorized into knowledge-based, global topic-based and user-based methods.

Knowledge-Based Methods. Tang et al. (2019) first propose the target-guided dialogue topic selection task and develop a rule-based model based on the similarity of the next topic and target. DKRN (Qin et al., 2020) further utilizes semantic correlation to improve the smoothness of selected topics. Xu et al. (2020b) and CKC (Zhong et al., 2021) model topic transition as a continuous walk on the commonsense graph (CKG), which effectively reduces the neighborhood candidate topic space. Considering topic selection based on neighborhood entities on CKG does not conform to real dialogues, ECCF (Li et al., 2023a) mines high-frequency topic transition from real dialogues to expand commonsense graph. However, strict neighborhood constraints in topic selection limit the generalization of these methods to real-world conversations (Li et al., 2022).

Global Topic-Based Methods. This kind of method aims to leverage relevant topic transitions in other paths to aid subsequent topic selection in the local topic path. Xu et al. (2020a) construct a global topic transition graph utilizing all topic paths in dialogue corpora to capture utterance-level correlations for topic selection. On this basis, Xu et al. (2021) extend the method by employing Discrete Variational Auto-Encoder (VAE) with Graph Neural Network (GNN) to aggregate global neighborhood topics into local topic paths for modeling global level topic correlations. CG-nAR (Zou et al., 2021) considers the joint influence of multiple turns of historical topics on the next topic selection and uses a dynamic graph attention mechanism to select subsequent topics. SGTA (Wen et al., 2022) represented the global co-occurrence frequency of topics as a multivariate skew-normal distribution with hybrid kernel functions to assist in selecting relevant topics with high global co-occurrence frequencies.

However, it is crucial to note that users with different personas tend to choose different topics associated with the same historical topic sequence. The aforementioned works indiscriminately fuse the topic paths of users with different personas at the global level, leading to the erroneous modeling of user preferences and depersonalized topic selection. In contrast, PETD exclusively integrates global topics related to the user personas, avoiding interference from global topics corresponding to irrelevant personas.

User-Based Methods. Recently, user-based works attempt to model user characteristics for the enhancement of user satisfaction in predicted topics. TG-ReDial (Zhou et al., 2020) proposes a topic-grounded dialogue dataset with user personas and employs pre-trained language models to independently encode historical context and all user personas for topic prediction. UPCR (Ren et al., 2022) uses user embedding instead of text-described personas to model user characteristics and combines with dialogue history for topic selection. However, these methods statically model user personas and inject all user personas into each turn of the conversation, resulting in the risk of generating context-independent topics and responses. In contrast, PETD selectively considers relevant personas based on the historical conversation and topic path, taking into account the dynamic transition of the personas exhibited by the user during the conversation.

3 Method

3.1 Problem Formulation

Let $\mathcal{C} = \{c_1, c_2, ..., c_{|\mathcal{C}|}\}$ represent a multi-turn dialogue context and $tp = \{t_1, t_2, ..., t_{|tp|}\}$ is topic path of the dialog $\mathcal{C}$, where t_j refers to the topics discussed at the j-th turn. We assume that user u is taken from a set $\mathcal{U}$ with a set of predefined personas $\mathcal{P}_u = \{p_1, p_2, ..., p_{|\mathcal{P}_u|}\}$, where each persona is described as a sentence $p_i = \{w_j\}_{j=1}^{|p_i|}$. Given conversation $\mathcal{C}_{[1:j]}$, user u, user personas $\mathcal{P}_u$ and topic path $tp_{[1:j]}$, the goal of topic-grounded dialogue is to predict topics of the next turn t_{j+1} and generate the corresponding response $\mathcal{R}_{j+1}$. We split the target function of the model through the Bayesian formula to correspond to the three subtasks, persona selection $\mathcal{P}^+_{j+1}$, topic selection t_{j+1}, and response generation $\mathcal{R}_{j+1}$, separately. We formulate target function y^* as follows:

$$\begin{aligned} y^* \triangleq & \prod_{j=1}^{|\mathcal{C}|-1} P(\mathcal{P}^+_{j+1}|u, \mathcal{C}_{[1:j]}, tp_{[1:j]}, \mathcal{P}_u) \cdot \\ & \prod_{j=1}^{|\mathcal{C}|-1} P(t_{j+1}|u, \mathcal{C}_{[1:j]}, tp_{[1:j]}, \mathcal{P}^+_{j+1}) \cdot \\ & \prod_{j=1}^{|\mathcal{C}|-1} P(\mathcal{R}_{j+1}|u, \mathcal{C}_{[1:j]}, \mathcal{P}^+_{j+1}, t_{j+1}). \end{aligned} \tag{1}$$

For the sake of brevity, we omit all temporal subscripts below. Table 5 in the Appendix lists the notations used in this paper.

3.2 Input Representation

We use two multiple multi-layer transformer encoders as the backbone to encode the dialogue context $\mathcal{C}$ and topic path tp respectively. We concat

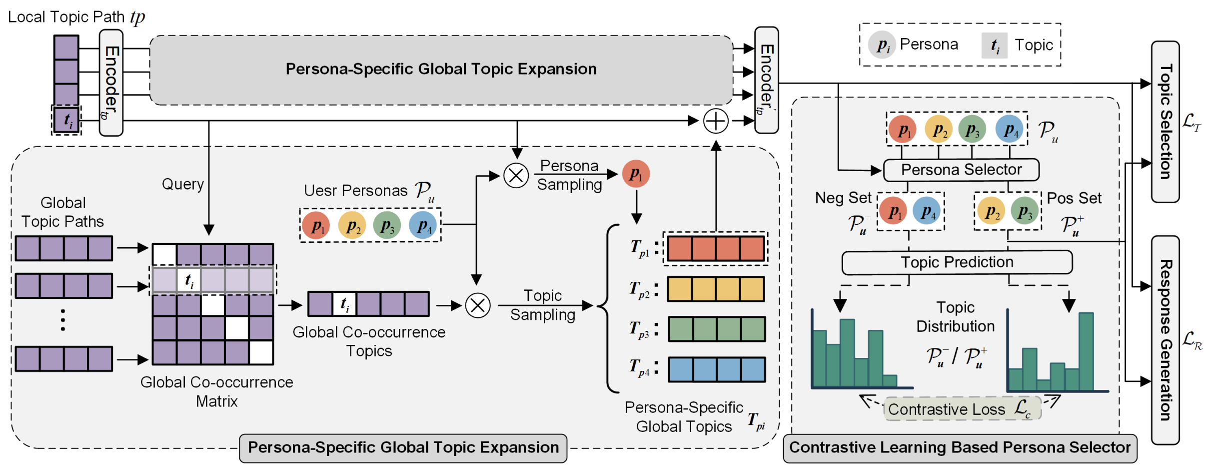

Figure 2: The structure of PETD. We use different colors to represent different personas and their corresponding global topics.

all sentences into a single paragraph before encoding. $H_{\mathcal{C}} \in R^{|\mathcal{C}|\times d}$ and $H_{tp} \in R^{|tp|\times d}$ are used to represent the encoded sentence level dialogue context representation and topic representation, respectively.

$$\begin{aligned} H_{\mathcal{C}} &= [h_{c_1}, h_{c_2}, ..., h_{c_{|\mathcal{C}|}}] = \text{encoder}_{\mathcal{C}}(\mathcal{C}), \\ H_{tp} &= [h_{t_1}, h_{t_2}, ..., h_{t_{|tp|}}] = \text{encoder}_{tp}(tp). \end{aligned} \tag{2}$$

Users and personas are parameterized as $E_{\mathcal{U}}$ and $E_{\mathcal{P}}$ respectively by indexing corresponding values from the embedding matrix. To utilize the semantic information in persona description, we use BERT (Kenton and Toutanova, 2019) to encode personas at the sentence level as the initialization embedding. For users, we use random initialization embedding.

3.3 Persona-Specific Global Topic Expansion

In this section, we use personas to selectively fuse global topics to alleviate the detrimental impact of persona-irrelevant global topics on topic selection. For each topic in the topic path, we first calculate the correlation between global co-occurrence topics and user personas to decouple the global topics under different personas. The correspondence score s_{ij} between persona p_i and topic t_j is calculated as follows:

$$s_{ij} = e_{p_i} W e_{t_j}, \tag{3}$$

where $W \in R^{d\times d}$ is a learnable matrix, e_{p_i} and e_{t_j} are the embeddings of persona p_i and topic t_j. We sample the most relevant k topics for each persona p_i, named persona-specific global topics T_{p_i}.

Subsequently, we select the relevant user personas for each turn of the historical conversation. Context representation h_{c_i}, user embedding e_u, and topic representation h_{t_i} are used together for correspondence score s'_{ij} calculation:

$$s'_{ij} = f([h_{c_i}; e_u; h_{t_i}], e_{p_j}), \tag{4}$$

$$f(h_i, h_j) = \text{MLP}(h_i) \cdot h_j^{\mathrm{T}}, \tag{5}$$

where [;] represents the concat operation on multiple elements and T means matrix transpose. Here, we build a mask matrix $M' \in R^{|tp|\times|\mathcal{P}_u}$ to sample relevant personas. A simple threshold sampling strategy is used to populate the mask matrix according to the correlation score, as follows:

$$m'_{ij} = 1 \text{ if } \sigma(s'_{ij}) \geq 0.5 \text{ else } 0, \tag{6}$$

where σ is the sigmoid activation function.

For each turn, we aggregate persona-specific global topics corresponding to the selected relevant personas into local topic to obtain global-enhanced topic representation h'_{t_i}:

$$h'_{t_i} = \text{FFN}(h_{t_i} + \sum_{j=1}^{|\mathcal{P}_u|}(s'_{ij} \cdot m'_{ij} \cdot \sum_{t_k \in T_{p_i}} (s_{jk} \cdot e_{t_k}))). \tag{7}$$

Another multi-layer transformer encoder is used to encode the topic path after global topic aggregation.

$$\begin{aligned} \widetilde{H}'_{tp} &= \text{encoder}'_{tp}(H'_{tp}) \\ &= \text{encoder}'_{tp}([h'_{t_1}, h'_{t_2}, ..., h'_{t_{|tp|}}]). \end{aligned} \tag{8}$$

3.4 Contrastive Learning Based Persona Selector

In this section, we use a persona selector to select personas that are relevant to the next target topics.

To optimize the persona selector without supervisory signals and amplify the difference between relevant and irrelevant persona representations, we design a topic prediction auxiliary task based on contrastive learning.

We use the scoring functions corresponding to formulas 5 and 6 to evaluate the relevance $S'' \in R^{|\mathcal{P}_u|}$ of each persona and construct a mask vector $M'' \in R^{|\mathcal{P}_u|}$.

$$s''_{p_i} = f([H_{\mathcal{C}}; e_u; \widetilde{H}'_{tp}], e_{p_i}). \tag{9}$$

According to the mask vector M'', we divide the user persona set into positive persona set $\mathcal{P}_u^+$ (where m''_{p_i}=1) and negative persona set $\mathcal{P}_u^-$ (where m''_{p_i}=0). We further use contrastive learning to reinforce the difference between positive and negative persona sets. We first aggregate the positive and negative persona sets separately to obtain the representations $h_{\mathcal{P}_u^+}$ and $h_{\mathcal{P}_u^-}$:

$$\begin{aligned} h_{\mathcal{P}_u^+} &= \sum_{p_i \in \mathcal{P}_u^+} \frac{\exp(s''_{p_i})}{\sum_{p_i \in \mathcal{P}_u^+} \exp(s''_{p_i})} \cdot e_{p_i}, \\ h_{\mathcal{P}_u^-} &= \sum_{p_i \in \mathcal{P}_u^-} \frac{\exp(-s''_{p_i})}{\sum_{p_i \in \mathcal{P}_u^-} \exp(-s''_{p_i})} \cdot e_{p_i}, \end{aligned} \tag{10}$$

where each weight in the negative set is transformed into a negative value, thereby amplifying the presence of less relevant personas within the negative set representation. We only use personas to predict the next topic, as follows:

$$g(h_{\mathcal{P}_u^+}, e_{t_i}) = \frac{\exp(h_{\mathcal{P}_u^+} W e_{t_i}^{\mathrm{T}})}{\sum_{t_j \in \mathcal{T}} \exp(h_{\mathcal{P}_u^+} W e_{t_j}^{\mathrm{T}})}. \tag{11}$$

The objective function for contrastive learning is defined as follows:

$$\mathcal{L}_c = -\log(g(h_{\mathcal{P}_u^+}, e_{t_i}) - g(h_{\mathcal{P}_u^-}, e_{t_i})), \tag{12}$$

where e_{t_i} is the embedding vector of the target topic. The contrastive learning objective function aims to augment the divergence in prediction scores between positive and negative persona sets for subsequent correct topics. Consistent with our intuition, relevant personas tend to predict correct topics while irrelevant personas tend to predict wrong ones. Through this task, we cleverly transform topic annotations into supervision signals of relevant personas to optimize the persona selector.

3.5 Topic Selection

We use the same dot product similarity as formula 5 for topic selection, and the probability of each topic being selected is defined as $p(t_i)$:

$$p(t_i) = \frac{\exp(f([H_{\mathcal{C}}; e_u; \widetilde{H}'_{tp}; h_{\mathcal{P}_u^+}], e_{t_i}))}{\sum_{t_j \in \mathcal{T}} \exp(f([H_{\mathcal{C}}; e_u; \widetilde{H}'_{tp}; h_{\mathcal{P}_u^+}], e_{t_j}))}, \tag{13}$$

where $H_{\mathcal{C}}$ is the dialogue context representation, e_u is the user embedding, $\widetilde{H}'_{tp}$ is the topic path represention after global topic aggregation, $h_{\mathcal{P}_u^+}$ is the positive persona set representation. We set a cross-entropy loss to optimize the parameter:

$$\mathcal{L}_{\mathcal{T}} = -log(p(t_i)). \tag{14}$$

3.6 Response Generation

We use the transformer decoder with copy mechanism (Gu et al., 2016) to generate responses. At the ζ-step, previously generated response embedding $e_{w_{[1:\zeta-1]}}$ and other auxiliary information are fed into the decoder together to generate the current word.

$$h_\zeta = \text{decoder}([H_{\mathcal{C}}; e_u; \widetilde{H}'_{tp}; h_{\mathcal{P}_u^+}; e_{w_{[1:\zeta-1]}}]). \tag{15}$$

The probability of generating the word w_ζ is the sum of both the generation probability $p_g(w_\zeta)$ and the copy probability $p_c(w_\zeta)$:

$$\begin{aligned} p(w_\zeta) &= p_g(w_\zeta) + p_c(w_\zeta), \\ p_g(w_\zeta) &= f(h_t, e_{w_\zeta}), \\ p_c(w_\zeta) &= f(h_t, h_{w_\zeta}),\ w_\zeta \in \mathcal{C}, t_i, \end{aligned} \tag{16}$$

where we use the encoder output h_{w_ζ} of dialogue context and selected topics for copynet generation.

As same as topic selection, we also use the cross-entropy function to optimize model parameters for a response utterance of length $|\mathcal{R}|$:

$$\mathcal{L}_{\mathcal{R}} = -\frac{1}{|\mathcal{R}|} \sum_{\zeta=1}^{|\mathcal{R}|} log(p(w_\zeta)). \tag{17}$$

4 Experiments

In this section, we demonstrate that our proposed method achieves the state-of-the-art and the importance of each proposed component through exhaustive experiments. We also give detailed hyperparameter analysis (A.3), persona selection analysis (A.4), and case study (A.5) in the Appendix.

Dataset	TG-ReDial	Persona-Chat
dialogue	10,000	9,935
utterance	129,392	147,039
topic	2,571	2,409
persona	2,433	6,737
avg personas per user	10	5

Table 1: Statistics of the datasets.

4.1 Datasets

To evaluate the effectiveness of PETD, following previous works (Wen et al., 2022; Ren et al., 2022; Zou et al., 2021; Kishinami et al., 2022), we conduct experiments on two widely used benchmark datasets, **TG-ReDial** (Zhou et al., 2020) and **Persona-Chat** (Zhang et al., 2018), for target-guided topic-grounded dialogue. Table 1 presents the statistics of the datasets. We give a more detailed description of datasets in the Appendix A.1.

4.2 Baselines

In order to demonstrate the effectiveness of PETD, we compare it with three category baselines.

(1) Knowledge-based Methods. **DKRN** (Qin et al., 2020) proposes an explicit knowledge-routed method for topic selection. **CKC** (Zhong et al., 2021) models topic transition on the commonsense graph and utilizes graph neural network to model the correlation between topics. **ECCF** (Li et al., 2023a) augments high-frequency topic transitions into the commonsense graph and aggregates neighborhood information using relation-aware attention.

(2) Global topic-based Methods. **CG-nAR** (Zou et al., 2021) constructs a topic global transition graph for topic selection, and designs a non-autoregressive topic-grounded insertion transformer decoder for response generation. **SGTA** (Wen et al., 2022) introduces a latent space for flexibly integrating global topic transition probabilities with sequence topic prediction probabilities.

(3)User-based Methods. **Profile-Bert, TG-CRS** (Zhou et al., 2020) introduces profile information for dialogue topic selection for the first time. **Profile-Bert** uses a pre-trained language model to encode semantic information in all user personas for topic selection. **TG-CRS** uses user profile, dialogue context, and topic sequence for topic selection. **UPCR** (Ren et al., 2022) uses user embedding and dialogue history to model long-term and short-term user interests respectively for topic selection.

4.3 Implementation Details

All of the baselines and our method are implemented in PyTorch and trained on RTX 4090 24GB. We keep a maximum of 7 turns of historical dialogue for all methods and allocate $k = 10$ global topics for each persona. The embedding size is set to 768 and the L2 regularization weight is 1e-6. Throughout the experiments, we use Adam optimizer (Kingma and Ba, 2015). Its initial learning rate is 1e-4 and the batch size is set to 80. In order to prevent overfitting, the dropout rate is fixed at 0.1. For all datasets, we split the dataset into training/validation/testing sets. We train the model for up to 100 epochs and early stop the training in advance when the hit@3 and bleu-1 don't improve for 10 consecutive epochs on the validation set.

4.4 Evaluation Metric

Automatic Evaluation. We jointly evaluate the abilities of topic selection and response generation from several different perspectives. (1) For topic selection accuracy, following previous works (Tang et al., 2019; Zhou et al., 2020; Wen et al., 2022), we adopt *Hit@k* (k=1,3,5) as evaluation metrics. (2) We report *perplexity* (*PPL*) (Horgan, 1995) and *BLEU-n* (n=1, 2) (Papineni et al., 2002) to evaluate the coherence and word overlap of generated utterances. (3) Following previous works (Zhou et al., 2020; Wen et al., 2022), we employ *Distinct-n* (n=1, 2) (Li et al., 2016) to evaluate the diversity of the generated response.

Human Evaluation. Following previous works (Zhou et al., 2020; Wen et al., 2022), we adopt *Relevance*, *Fluency* and *Informativeness* of the generated utterances with the rating range of [0, 2]. We recruit three experienced annotators to evaluate 100 randomly selected dialogues. The Fleiss Kappa is 0.68, indicating consistency in the estimates of annotators. The evaluation details are shown in the Appendix A.2.

4.5 Main Result

The evaluation results are shown in Table 2.

4.5.1 Performance on Topic Selection

Our methods consistently outperform all baselines, achieving an average performance increase of 4.90%, 3.64%, and 3.17% for Hit@1/3/5 respectively compared to the best-performing baseline. This improvement is attributed to PETD, which filters out the irrelevant information from global topics and user persona sets, enabling more precise

Dataset	Method	Topic Selection			Dialogue Generation							
		Hit@1	Hit@3	Hit@5	PPL	BLEU-1	BLEU-2	Distinct-1	Distinct-2	Relevance	Fluency	Informativeness
TG-ReDial	DKRN	0.402	0.482	0.507	-	-	-	-	-	-	-	-
	CKC	0.591	0.786	0.827	-	-	-	-	-	-	-	-
	ECCF	0.601	0.839	0.852	34.131	0.263	0.161	0.017	0.088	1.37	1.26	1.51
	CG-nAR	0.566	0.764	0.829	52.417	0.161	0.103	0.015	0.047	1.45	1.08	1.42
	SGTA	0.621	0.852	0.867	21.616	0.301	0.191	<u>0.023</u>	0.124	1.54	**<u>1.52</u>**	1.53
	Profile-BERT	0.499	0.821	0.834	23.552	0.287	0.117	0.019	0.090	1.14	1.48	1.19
	TG-CRS	0.613	0.816	0.830	<u>19.223</u>	0.280	0.173	0.021	0.094	1.55	1.51	1.37
	UPCR	<u>0.808</u>	<u>0.883</u>	<u>0.907</u>	41.234	<u>0.316</u>	<u>0.200</u>	0.022	<u>0.132</u>	<u>1.57</u>	1.47	<u>1.58</u>
	PETD	**0.837***	**0.899**	**0.920***	**17.076***	**0.351***	**0.224***	**0.031***	**0.176***	**1.75**	1.51	**1.65**
Persona-Chat	DKRN	0.468	0.515	0.533	-	-	-	-	-	-	-	-
	CKC	0.583	0.733	0.773	-	-	-	-	-	-	-	-
	ECCF	0.634	0.775	0.839	35.499	0.257	0.166	0.038	0.238	1.35	1.36	1.43
	CG-nAR	0.542	0.592	0.613	18.237	0.184	0.134	0.041	0.237	1.34	1.27	1.39
	SGTA	0.664	<u>0.857</u>	<u>0.907</u>	<u>15.648</u>	<u>0.327</u>	0.207	<u>0.047</u>	<u>0.241</u>	1.52	1.43	<u>1.57</u>
	Profile-BERT	0.502	0.820	0.836	19.938	0.269	0.170	0.031	0.228	1.19	1.47	1.28
	TG-CRS	0.657	0.846	0.863	16.269	0.285	0.181	0.037	0.237	1.49	<u>1.54</u>	1.37
	UPCR	<u>0.685</u>	0.855	0.865	26.735	0.311	<u>0.216</u>	0.016	0.192	<u>1.62</u>	1.45	1.52
	PETD	**0.727***	**0.904***	**0.951***	**13.462***	**0.346***	**0.252***	**0.061***	**0.254**	**1.67**	**1.58**	**1.63**

Table 2: The performance of PETD and all baselines. The results of the best baseline and best performance in each column are underlined and in boldface respectively. We do not report dialogue generation results for DKRN and CKC because their methods rank candidate sentences instead of generation. Significant improvements compared to the best baseline are marked with * (t-test, p $\leq$ 0.05).

Dataset	Method	Topic Selection			Dialogue Generation							
		Hit@1	Hit@3	Hit@5	PPL	BLEU-1	BLEU-2	Distinct-1	Distinct-2	Relevance	Fluency	Informativeness
TG-ReDial	PETD (0.2B)	0.837	0.899	0.920	17.076	0.351	0.224	0.031	0.176	1.75	1.51	1.65
	Llama2 (7B)	0.639	-	-	12.710	0.314	0.226	0.044	0.197	1.71	1.75	1.81
	Llama2-COT (7B)	0.658	-	-	11.492	0.349	0.239	0.042	0.193	1.78	1.74	1.81
	PETD† (7.1B)	0.837	0.899	0.920	10.505	0.394	0.247	0.057	0.215	1.81	1.75	1.83
Persona-Chat	PETD (0.2B)	0.727	0.904	0.951	13.462	0.346	0.252	0.061	0.254	1.67	1.58	1.63
	Llama2 (7B)	0.603	-	-	11.929	0.317	0.231	0.054	0.257	1.72	1.79	1.74
	Llama2-COT (7B)	0.620	-	-	9.876	0.320	0.248	0.067	0.271	1.74	1.81	1.77
	PETD† (7.1B)	0.727	0.904	0.951	9.082	0.379	0.273	0.068	0.287	1.77	1.78	1.81

Table 3: The performance of PETD and LLMs. For Llama2 and Llama2-COT, We evaluate the accuracy of topic prediction by detecting whether the target topic is included in the generated responses. We report the number of parameters of each model in symbol ().

utilization of side information for dialogue topic selection.

Overall, global topic-based methods tend to outperform knowledge-based methods due to their capacity to incorporate other topic paths, capturing correlations between different dialogues. Nonetheless, global topic-based methods disregard the influence of persona on topic transition and treat the global information under different personas equally. Conversely, our method differentially considers the global topics under different personas, resulting in improved performance in comparison to global topic-based methods.

The user-based methods almost achieve the best results overall baselines due to the incorporation of user persona information. It is worth noting that the poor performance of Profile-BERT arises from its failure to dynamically model the personas exhibited by the user during dialogue and using the full personas resulting in noise that heavily interferes with the topic selection. UPCR achieves the best experimental performance among all baselines as it employs user embeddings and topic path to model both long-term and short-term user preferences. However, these user-based methods either consider full personas or solely rely on user embeddings, without adequately tackling the problem of irrelevant information within user persona sets. We design a person selector and optimize it using a contrastive learning-based topic selection auxiliary task to mitigate irrelevant information in user persona sets. In summary, PETD, through its consideration of fine-grained persona, exhibits the ability to provide a more personalized and coherent topic selection compared to state-of-the-art baselines.

4.5.2 Performance on Response Generation

Overall, SGTA and UPCR achieve better experimental results in all baselines. The performance of SGTA comes from considering multiple potential topics simultaneously during dialogue generation. The enhanced performance of UPCR results from incorporating user embedding as user characteristics in the generation process. Our method achieves the best performance, with an average improvement of 9.4%, and 25.9% in Bleu and Distinct metrics compared to the current state-of-the-art baseline, respectively. The significant improvement in PETD performance arises from a more accurate selection of topics and user personas. The utterances generated by PETD exhibit an average enhancement of 7.28% and 4.13% on the relevance and informativeness metrics, respectively, which indicates that our method can better align with the user personas and dialogue topics.

4.6 Compaered with LLMs

To demonstrate the importance of topic selection for the Large Language Models (LLMs), we choose Llama2-7b-chat (Touvron et al., 2023) after instruction tuning as the backbone for experiments. We use two methods to generate responses for LLM, (1) using the prompt to generate responses for the model at once (Llama2), and (2) using Chain-of-Thought (Wei et al., 2022) to select relevant personas before generating responses (Llama2-COT). We also design an enhanced variant of our method (PETD†) by using Llama2-7b-chat as the decoder to generate responses. For all experiments, we fine-tune three epochs on the corresponding dataset using lora (Hu et al., 2021) for Llama-7b-chat. The prompts are shown in the Table 7 in Appendix.

The experimental results are shown in Table 3. We find that although LLM can generate more fluent and diverse responses, the accuracy of topic selection is lower than most baselines of topic-grounded dialogue. This is consistent with previous research results (Cao, 2023; Hudeček and Dušek, 2023). Without an external topic management module, Llama2, like most conversation models, tends to discuss the current topic rather than expand on new ones. Compared to Llama2, the improvement of Llama2-COT indicates that LLM is still subject to interference from irrelevant personas in the prompt. PETD† achieve optimal performance by introducing predicted topics and relevant personas into LLM, indicating that we can cleverly utilize small models (topic selection) to improve the initiative and information of LLM.

4.7 Ablation Study

To investigate the effectiveness of PETD, we conduct detailed ablation experiments around two key components of PETD. The experimental results are shown in Table 4.

Persona-Specific Global Topic Expansion. We first eliminate the global topic aggregation for w/o global topic. The dramatic drop in this variant performance demonstrates the important role of global information in modeling user preferences. To demonstrate the effectiveness of personas in global topic aggregation, we contrast two variants: one employing topic similarity (w topic similar) and the other utilizing global co-occurrence (w co-occurrence) to select and aggregate global topics. We observe that PETD w co-occurrence and PETD w topic similar achieve similar performance, indicating that the statistical frequency of global co-occurrence can be well fitted by the topic representation similarity. This similarity serves as the basis for our successful decoupling of global topics through personas and topics similarity. However, the performance of these two variants remains notably lower than that of PETD, highlighting the necessity of distinctively accounting for global topics aligned with different personas.

Dateset	Method	Hit@1	Hit@3	Hit@5
TG-ReDial	PETD	0.837	0.899	0.920
	w/o global topic	0.801	0.871	0.890
	w topic similar	0.814	0.883	0.910
	w co-occurrence	0.813	0.885	0.913
	w/o persona	0.759	0.874	0.894
	w/o persona selection	0.814	0.872	0.902
	w random persona selection	0.729	0.862	0.857
	w/o auxiliary task	0.816	0.877	0.903
	w/o contrastive learning	0.823	0.884	0.912
Persona-Chat	PETD	0.727	0.904	0.951
	w/o global topic	0.703	0.834	0.901
	w topic similar	0.704	0.868	0.927
	w co-occurrence	0.718	0.875	0.895
	w/o persona	0.676	0.778	0.857
	w/o persona selection	0.718	0.872	0.890
	w random persona selection	0.689	0.798	0.876
	w/o auxiliary task	0.718	0.873	0.893
	w/o contrastive learning	0.720	0.887	0.916

Table 4: The Performance of Ablation Study.

Contrastive Learning Based Persona Selector. We design following variants: deletes complete personas (w/o persona), deletes persona selection module (w/o persona selection), instead random selection of persona selection module (w random persona selection), and deletes contrastive learning based auxiliary task (w/o auxiliary task), deletes contrastive learning (w/o contrastive learning), respectively. Generally speaking, richer side information will make the model perform better. When all personas are given (w/o persona selection), the performance of the model is lower than that of selecting personas (PETD), indicating that the model is poisoned by irrelevant persona, proving that persona selection is necessary. For PETD with random persona selection. The number of selected personas is set to 2, consistent with the number of relevant personas selected by our methods (PETD) in most scenarios. Since random selection has a high probability of selecting irrelevant personas, the performance of the PETD w random person selection variant decreases significantly, demonstrating the sensitivity of the model to irrelevant personas. We believe that persona selection is an essential component for intelligent agents as conversation scenarios become more complex and user personas increase significantly. We also find that the inclusion or exclusion persona selection module has little impact on model performance when the model is without the auxiliary task. This finding suggests that optimizing the persona selector using a specific end-to-end training method is challenging and the auxiliary task we proposed can subtly address this problem through targeted optimization. After removing contrastive learning, the model's performance drops significantly, although it is slightly higher than the variant that removes all the auxiliary tasks. The significant drop in performance demonstrates the necessity of using contrastive learning to increase the difference between irrelevant and relevant persona representations and to specifically optimize the persona selector.

5 Conclusion

In this work, we leverage the interplay between topics and personas to improve the accuracy of topic selection by removing redundant noise in side information. We notice that user under different personas selects different topics, and existing global topic-based methods ignore this difference. Simultaneously, the complicated persona information in the user profile contains plenty of noise, and only a few are related to the next topic. To tackle the above problems, we propose a novel model, named PETD. We construct a corresponding topic set for each persona at the global level and selectively amalgamate globally pertinent topic sets aligned with user personas to exclude persona-irrelevant global topics. We subsequently develop a persona selector, curbing the adverse influence of irrelevant personas on topic selection. A contrastive learning based auxiliary task is proposed to optimize the persona selector and increase the distance between different persona representations in unlabeled scenarios. Comprehensive experiments showcase the superior ability of our method to achieve more precise topic selection and produce captivating and varied responses, outperforming all benchmarks across various evaluation metrics.

Limitations

First, our method only uses the simplest top-k sampling and threshold sampling for topic and persona selection and does not experiment with more clever sampling methods. Second, due to the large experimental scale and fairness issues, we used a transformer decoder of a similar size to GPT2 in most of the experiments. Although we demonstrated the effectiveness of our method for LLMs in Section 4.6, this is still a limitation. Third, considering the simplicity of the method and the gap between the structure of the knowledge graph and real conversations, this study does not discuss how to introduce the knowledge graph into topic prediction. However, external knowledge represented by a knowledge graph can be considered similar to the global topic co-occurrence matrix, as both model topic correlation relationships outside the sequence. Therefore, our method can be extended to other side information such as knowledge graphs through similar decoupling, selection, and fusion methods.

Ethics Statement

In a broad sense, introducing personality information into topic-grounded conversations may indeed lead to user profile privacy leaks and false identity forgery. However, in this work, personality information and responses are limited to the scope of the experiment and are not enough to threaten real conversations. In addition, all models in this paper are trained and evaluated on datasets collected in the public corpus, and the dataset corpus is only used for experimental purposes. The dataset we use does not contain unethical language.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62276110, No. 62172039 and in part by the fund of Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL). The authors would also like to thank the anonymous reviewers for their comments on improving the quality of this paper.

References

Lang Cao. 2023. Diaggpt: An llm-based chatbot with automatic topic management for task-oriented dialogue. *arXiv preprint arXiv:2308.08043*.

Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. 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.

J. Horgan. 1995. From complexity to perplexity. *Scientific American*, 272(6):104–109.

Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Vojtěch Hudeček and Ondřej Dušek. 2023. Are large language models all you need for task-oriented dialogue? In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 216–228.

Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Yosuke Kishinami, Reina Akama, Shiki Sato, Ryoko Tokuhisa, Jun Suzuki, and Kentaro Inui. 2022. Target-guided open-domain conversation planning. In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, pages 660–668. International Committee on Computational Linguistics.

Yongquan Lai, Shixuan Fan, Zeliang Tong, Weiran Pan, and Wei Wei. 2023. Conversational aspect-based sentiment quadruple analysis with consecutive multi-view interaction. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 162–173. Springer.

Dawei Li, Yanran Li, Jiayi Zhang, Ke Li, Chen Wei, Jianwei Cui, and Bin Wang. 2022. C^3KG: A Chinese commonsense conversation knowledge graph. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1369–1383, Dublin, Ireland. Association for Computational Linguistics.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. 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.

Siheng Li, Wangjie Jiang, Pengda Si, Cheng Yang, Yao Qiu, Jinchao Zhang, Jie Zhou, and Yujiu Yang. 2023a. Enhancing dialogue generation with conversational concept flows. In *Findings of the Association for Computational Linguistics: EACL 2023, Dubrovnik, Croatia, May 2-6, 2023*, pages 1484–1495. Association for Computational Linguistics.

Wendi Li, Wei Wei, Xiaoye Qu, Xian-Ling Mao, Ye Yuan, Wenfeng Xie, and Dangyang Chen. 2023b. Trea: Tree-structure reasoning schema for conversational recommendation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2970–2982.

Jiayi Liu, Wei Wei, Zhixuan Chu, Xing Gao, Ji Zhang, Tan Yan, and Yulin Kang. 2022a. Incorporating causal analysis into diversified and logical response generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 378–388.

Yifan Liu, Wei Wei, Jiayi Liu, Xianling Mao, Rui Fang, and Dangyang Chen. 2022b. Improving personality consistency in conversation by persona extending. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 1350–1359.

Zhenyi Lu, Wei Wei, Xiaoye Qu, Xian-Ling Mao, Dangyang Chen, and Jixiong Chen. 2023. Miracle: Towards personalized dialogue generation with latent-space multiple personal attribute control. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5933–5957.

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.

Jinghui Qin, Zheng Ye, Jianheng Tang, and Xiaodan Liang. 2020. Dynamic knowledge routing network for target-guided open-domain conversation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8657–8664. AAAI Press.

Zhaochun Ren, Zhi Tian, Dongdong Li, Pengjie Ren, Liu Yang, Xin Xin, Huasheng Liang, Maarten de Rijke, and Zhumin Chen. 2022. Variational reasoning about user preferences for conversational recommendation. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pages 165–175. ACM.

Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.

Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019. Target-guided open-domain conversation. 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 5624–5634. Association for Computational Linguistics.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Wei Wei, Jiayi Liu, Xianling Mao, Guibing Guo, Feida Zhu, Pan Zhou, and Yuchong Hu. 2019. Emotion-aware chat machine: Automatic emotional response generation for human-like emotional interaction. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1401–1410.

Xiaofei Wen, Wei Wei, and Xian-Ling Mao. 2022. Sequential topic selection model with latent variable for topic-grounded dialogue. In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 1209–1219. Association for Computational Linguistics.

Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. 2021. Discovering dialog structure graph for coherent dialog generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 1726–1739. Association for Computational Linguistics.

Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020a. Conversational graph grounded policy learning for open-domain conversation generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1835–1845. Association for Computational Linguistics.

Jun Xu, Haifeng Wang, Zhengyu Niu, Hua Wu, and Wanxiang Che. 2020b. Knowledge graph grounded goal planning for open-domain conversation generation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9338–9345. AAAI Press.

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.

Yu Zhang and Qiang Yang. 2021. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609.

Sen Zhao, Wei Wei, Yifan Liu, Ziyang Wang, Wendi Li, Xian-Ling Mao, Shuai Zhu, Minghui Yang, and Zujie Wen. 2023. Towards hierarchical policy learning for conversational recommendation with hypergraph-based reinforcement learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2459–2467.

Peixiang Zhong, Yong Liu, Hao Wang, and Chunyan Miao. 2021. Keyword-guided neural conversational model. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 14568–14576. AAAI Press.

Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards topic-guided conversational recommender system. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 4128–4139. International Committee on Computational Linguistics.

Yicheng Zou, Zhihua Liu, Xingwu Hu, and Qi Zhang. 2021. Thinking clearly, talking fast: Concept-guided non-autoregressive generation for open-domain dialogue systems. 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 2215–2226. Association for Computational Linguistics.

Symbol	Description
$\mathcal{U}, u$	the set of all users and a user
$\mathcal{P}, p$	the set of all persona and a persona
$\mathcal{T}, t$	the set of all topics and a topic
$\mathcal{C}, c$	a conversation and a utterance in conversation
$\mathcal{W}, w$	the set of all words and a word
$\mathcal{R}$	a response
$\mathcal{P}_\mathcal{C}$	the persona set of a conversation
$\mathcal{P}_\mathcal{C}^+, \mathcal{P}_\mathcal{C}^-$	the position and negative persona set of a conversation
tp	topic path
e, E	a embedding vector and the embedding matrix
h, H	hidden state generated in encoder
S, s_{ij}	correspondence score matrix and the score between i and j
M, m_{ij}	mask matrix and the mask value between i and j
W	a learnable parameter matrix
d	dimension of embedding and hidden vector

Table 5: Glossary.

Criteria
Relevance
2: Fits the user's personality and is related to the current conversation
1: Relevant to one of the user's personality or current conversation, irrelevant or conflicting with the other
0: Does not resemble any user's personality and context history
Fluency
2: Fluent and easy to read
1: Grammatically formed
0: Not a complete sentence or hard to read
Informativeness
2: Have clear and specific meaning
1: Contain a few informative words
0: Meaningless sentence

Table 6: Criteria of human evaluation.

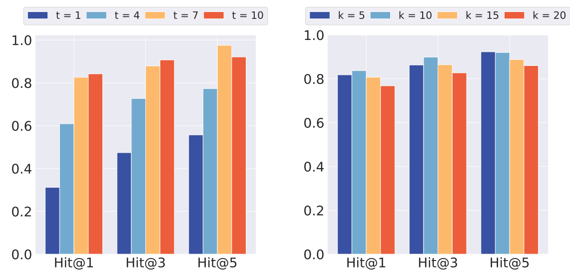

(a) history turns (b) topics per persona

Figure 3: Impact of the number of history turns (a) and topics per persona (b) on TG-ReDial dataset.

A More Experiments

A.1 Datasets

TG-ReDial[1] (Zhou et al., 2020) is a dialogue dataset in the movie domain, composed of 10,000 two-party dialogues between a seeker and a recommender. The dataset is structured in a topic-guided way, i.e. each dialogue in the TG-ReDial dataset includes a topic path to achieve target-guided topic-grounded dialogue. Each dialogue has 7.9 topics and each utterance contains 19.0 words. We use the same data preprocessing strategy as in (Wen et al., 2022; Ren et al., 2022).

Persona-Chat[2] (Zhang et al., 2018) is an open-domain dialogue dataset, which covers a broad range of topics. Following previous works (Qin et al., 2020; Zhong et al., 2021; Zou et al., 2021; Kishinami et al., 2022), we use TF-IDF and part-of-speech (POS) to extract topics from dialogue utterance. Inspired by (Zhong et al., 2021; Kishinami et al., 2022), we use the first topic in the follow-up dialogue that does not constant movement on the ConceptNet (Speer et al., 2017) as the target to construct the dataset into the form of target-guided topic-grounded dialogue task.

A.2 Human Evaluation Metrics

Relevance is used to evaluate the relevance of selected topics and generated sentences to historical conversations and the user's personality. *Fluency* is used to measure the fluency of generated utterances. *Informativeness* is used to evaluate whether the generated utterance revolves around the topics and user personas. The detailed scoring criteria are shown in Table 6.

Considering that we conduct evaluations on Chinese and English datasets, each evaluator we recruit is required to be a native Chinese speaker and have a high level of English proficiency. All human evaluations are conducted anonymously.

A.3 Hyperparameter Research

To explore the sensitivity of our proposed PETD model to hyperparameters, we conduct experiments on two hyperparameters, the number of history turns $t \in \{1, 4, 7, 10\}$, global topics corresponding to each persona $k \in \{5, 10, 15, 20\}$ on TG-ReDial dataset. The experimental results are shown in Figure 3.

The performance of the model increases with the number of history turns, which is consistent with previous works (Zou et al., 2021; Wen et al., 2022), indicating that the user's next topic is re-

[1]https://github.com/Lancelot39/TG-ReDial (Apache-2.0 license)

[2]https://parl.ai/projects/personachat/ (MIT license)

Model	Prompt
Llama2	The following is a conversation between an AI assistant called Assistant and a human user called User. The assistant needs to guide the conversation to target topic based on the user's personality and historical conversations. The users' personas are following: <persona 1>, <persona 2>, ...,<persona n> The topic history is following: <topic 1>, <topic 2>, ..., <topic n> The target topic is <target topic> The conversation is following: Assistant: <utterance 1>, User: <utterance 2>, Assistant: <utterance 3>, ..., User: <utterance n>, Assistant:
Llama2(COT)	**Turn 1:** The following is a conversation between an AI assistant called Assistant and a human user called User. The assistant needs to guide the conversation to target topic based on the user's personality and historical conversations. Please think step by step and first predict the personas that users may be interested in in the next turn. The users' personas are following: <persona 1>, <persona 2>, ...,<persona n> The topic history is following: <topic 1>, <topic 2>, ..., <topic n> The target topic is <target topic> The conversation is following: Assistant: <utterance 1>, User: <utterance 2>, Assistant: <utterance 3>, ..., User: <utterance n>, Relevant Personas: **Turn 2:** Please generate a response based on the selected personas and conversation history, and gradually guide the conversation to the target topics. Assistant:

Table 7: The prompt of Llama2 and Llama2(COT). The symbol <> is a placeholder that represents the corresponding data in the dataset.

Dateset	Method	Recall	Precision	F1
TG-ReDial	PETD	0.754	0.854	0.801
	PETD w/o auxiliary task	0.614	0.546	0.578
	Llama2-COT	0.415	0.472	0.442
Persona-Chat	PETD	0.790	0.833	0.811
	PETD w/o auxiliary task	0.674	0.508	0.579
	Llama2-COT	0.516	0.489	0.502

Table 8: The performance of persona selection.

lated to long historical dialogue, and it is necessary to model interest transition in historical dialogues. We observe that the improvement of model performance slows down significantly when the number of turns reaches 7 rounds. To balance efficiency and performance, we choose 7 turns as the hyperparameter of the main experiment. The performance of PETD is not sensitive to the number of topics corresponding to each persona, but the performance of the model still degrades when n is too small or large. When k exceeds 10, the performance of the model gradually decreases as k increases, indicating that forcibly assigning too many topics for each persona will introduce irrelevant information.

A.4 Persona Selection Analysis

We evaluated the persona prediction accuracy of our proposed method (PETD), a variant of PETD that deletes the contrastive learning based auxiliary task (PETD w/o auxiliary task), and Llama2-7b-chat using the thinking chain (Llama2-COT). To evaluate the accuracy of persona selection, we manually annotated 100 pieces of data with relevant personas for each dataset. We adopt *Recall*, *Precision* and *F1* as the evaluation metrics. The experimental results are shown in Table 8.

We find that our proposed method accurately predicts the relevant personas that users are likely to display in the next turn. The PETD w/o auxiliary task variant has a significant decrease in precision, indicating that without contrastive learning based auxiliary task, the model has difficulty distinguishing between relevant and irrelevant personas, and tends to predict more irrelevant personas. We found that the reason for the lower performance of Llama2-COT is that, through prompts, Llama2 prefers to select relevant personas shown in the conversation history rather than predicting relevant personalities in the future.

A.5 Case Study

We provide an example from the TG-ReDial dataset in Table 9. We notice that SGTA chooses globally relevant *friendship* as the topic because of its lack of consideration for user personas. Although TG-CRS considers the persona information, it brutally injects a whole range of personas, relevant and irrelevant, into the topic selection model and chooses

Personas	p_1: I like expressing emotions. p_2: I seem to fall in love lately. p_3: I like children very much. p_4: I yearn for love very much. p_5: I am a person who is easily moved.	p_6: I like music very much. p_7: I really like the impressive screenplay. p_8: I'm currently studying psychology. p_9: I like famous actors. p_{10}: I want good luck.
Context	A: How have you been lately? B: I've been fine lately. How long does it take for an ***actor*** to become popular since their debut? A: There is no specific standard for this, some are very fast and some are many years. Anyway, it is not easy for them to be well known by everyone, and some of them ***succeed*** by certain means. B: Such success is not worth mentioning, let's talk about something else, are you going to take the exam soon, ***good luck*** to you.	
Response	**Gold Response:** Haha, I hope this time it will pass smoothly. When you say I can reap ***love*** again, it will be even more perfect.	
	SGTA: Thank you, how lucky I am to have your ***friendship***.	
	TG-CRS: Let's not talk about exams, have you listened to ***music*** recently, Stephen Chow's new song?	
	UPCR: I know a way to improve ***luck***, you know?	
	Llama2: Thank you for your blessings. By the way, I have watched many movies recently. Can I recommend some movies about ***love*** to you?	
	PETD: Hey, the exam can pass smoothly. I just watched a movie about love last night, and I was moved to tears. I don't know when I will reap my own ***love*** again. **select personas:** p_2, p_4, p_5	

Table 9: One case extracted from TG-ReDial. The topics are bolded and italicized.

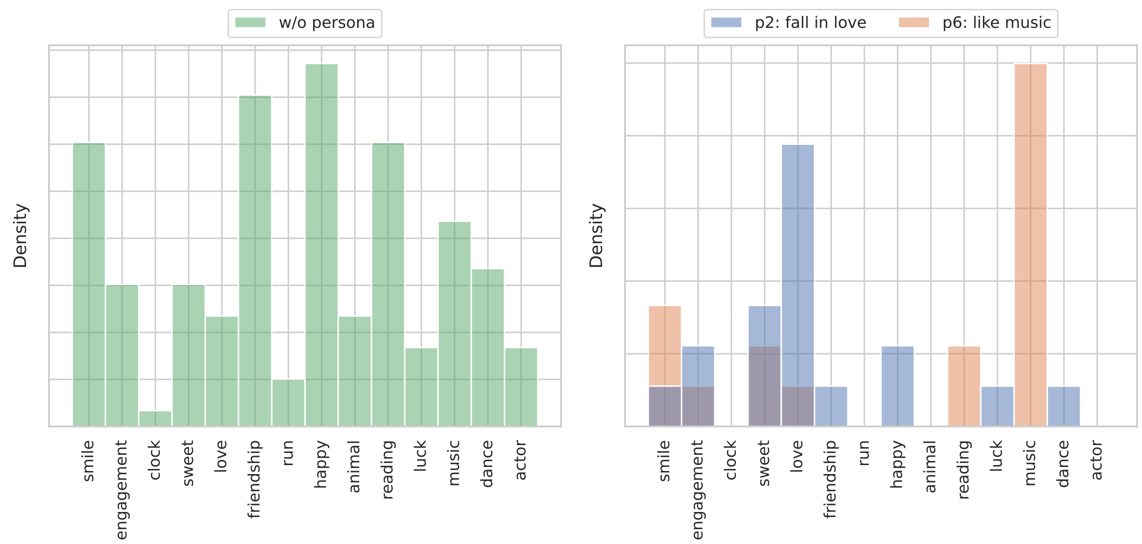

Figure 4: The global co-occurrence frequency density plot of the topic 'good luck' in the TG-ReDial dataset. For brevity, we do not draw all global topics.

persona-relevant but incorrect *music* as the topic. UPCR, which models long-term preference based on user embeddings, mistakenly chooses *luck* as the topic due to the lack of explicit modeling of each persona's fine-grained granularity. In contrast, PETD accurately models user preference, selects the corresponding personas, and further chooses the correct topics by taking into account the interaction of personas and topics to selectively aggregate relevant side information. Additionally, the PETD response is also affected by the persona of 'I am a person who is easily moved.', and generated utterance 'I was moved to tears', which effectively increased the user's interest in the dialogue. The improved performance of PETD is attributed to filtering out irrelevant information on global topics and user personas.

We further provide the global co-occurrence frequency density plot of the topic 'good luck' in the TG-ReDial dataset. As shown in Figure 4, the co-occurrence frequency density under the global topic and a certain persona has obvious distribution differences. We find that the global co-occurrence topic frequency is a multimodal distribution because it can be regarded as a superposition state of the global co-occurrence topic distribution under multiple personas. This explains that indiscriminately incorporating side information often leads the model to irrelevant or uninteresting topic choices.