Towards Zero-Shot Persona Dialogue Generation with In-Context Learning

Xinchao Xu, Zeyang Lei, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang
Baidu Inc., Beijing, China
{xuxinchao, leizeyang, wuwenquan01, niuzhengyu, wu_hua, wanghaifeng}@baidu.com

Abstract

Much work has been done to improve persona consistency by finetuning a pretrained dialogue model on high-quality human-annoated persona datasets. However, these methods still face the challenges of high cost and poor scalability. To this end, we propose a simple-yet-effective approach to significantly improve zero-shot persona consistency via in-context learning. Specifically, we first pre-train a persona-augmented dialogue generation model and then utilize in-context prompting mechanism to realize zero-shot persona customization. Experimental results demonstrate that our method can dramatically improve persona consistency without compromising coherence and informativeness in zero-shot settings.

1 Introduction

Endowing a dialog agent with consistent personas plays a very crucial role for build a more engaging and human-like chatbot. While large-scale pretrained dialog models have achieved great success (Adiwardana et al., 2020; Roller et al., 2021; Bao et al., 2021a; Gu et al., 2022; Thoppilan et al., 2022), maintaining persona consistency remains unsatisfactory and challenging when applying these pretrained models to different scenarios with preset personas.

Traditional approaches for addressing this issue typically involve manually collecting persona dialogue datasets, such as PersonaChat (Zhang et al., 2018), and fine-tuning pre-trained dialogue models on these datasets (Liu et al., 2020; Wolf et al., 2019). However, these methods face high costs for obtaining high-quality human-labeled persona datasets and scalability issues when encountering unseen persona information in the fine-tuned datasets (Huang et al., 2020; Roller et al., 2020)

To this end, we propose a novel in-context prompting learning mechanism to tackle zero-shot persona-based dialogue generation. By leveraging the pre-trained dialogue foundation model, we aim to generate persona-based dialogues without the need for expensive labeled data. Prompt-based in-context learning(Brown et al., 2020, Liu et al., 2021) has demonstrated its efficacy in few-shot or even zero-shot settings. In this work, we aim to implement a zero-shot persona dialogue generation mechanism using in-context prompt learning. However, we face two challenges:

- Selecting appropriate prompts: Previous research (Zheng and Huang, 2021) has identified that their prompting method is not compatible with knowledge-grounded approaches. The prompts should be effective for the pretrained dialogue foundation model to implement zero-shot persona dialogue generation (see Section 3).
- Even though designing an appropriate prompt is crucial, improving the in-context learning ability of the pretrained model is also important (see Section 4).

We would also like to highlight the practical value of our proposed work. In real application scenarios, the only cost of creating a new personalized chatbot is obtaining the corresponding persona prompts based on the prompt template, without the need for additional data annotation work. This significantly reduces the cost of creating a personalized chatbot.

In summary, this work makes the following contributions:

- We firstly propose a novel persona prompting mechanism in persona dialogue, enabling zero-shot capabilities to maintain consistency with preset personas.
- We demonstrate the importance of persona information in the pre-trained dialogue model for persona dialog, which can improve the in-context learning ability compared to the

Findings of the Association for Computational Linguistics: ACL 2023, pages 1387–1398
July 9-14, 2023

state-of-the-art models PLATO-2(Bao et al., 2020) and EVA2.0(Gu et al., 2022).

2 Related Work

There exists much work on consistent persona-based dialog generation. In particular, Zhang et al. (2018) proposed the PersonaChat dataset that has extensively promoted the development of this field where the crowd-workers are simply asked to chat with the other person naturally with the given personas. Zheng et al. (2019) constructed a large-scale persona dataset based on structured persona knowledge with public Weibo data. For improving persona consisitency, Qian et al. (2018) proposed an explicit persona model to generate consistent responses for given profile information. Liu et al. (2020) proposed $\mathcal{P}^2$ BOT to improve dialogue consistency by incorporating mutual persona perception. Song et al. (2021) disentangled persona-based dialogue generation into consistency understanding and dialogue generation. Cao et al. (2022) presented a model-agnostic data manipulation method for consistent persona generation.

While these works have shown promising performance on preset personas in their respective datasets, they heavily rely on expensive human-labeled datasets. Moreover, customizing a chatbot with unseen personas in zero-shot settings remains challenging. To address these limitations, this paper proposes an in-context prompting learning mechanism to improve zero-shot persona consistency without the need for human-annotated data.

3 In-Context Persona Prompting Learning

We propose a novel in-context prompting learning mechanism for zero-shot personalized dialogue. The idea is simple yet highly effective. We format the preset persona information into multiple turns of a dialogue and place it at the beginning of the original dialogue context, as illustrated in Figure 1. This approach differentiates our method from (Thoppilan et al., 2022), where the prompt consists of a single starting sentence. Furthermore, our method also differs from most previous prompting methods in GPT-3 (Brown et al., 2020), where prompts are typically used to distinguish between different tasks.

One of the key advantages of our approach is that it enables zero-shot persona customization without the need for annotating specific persona data.

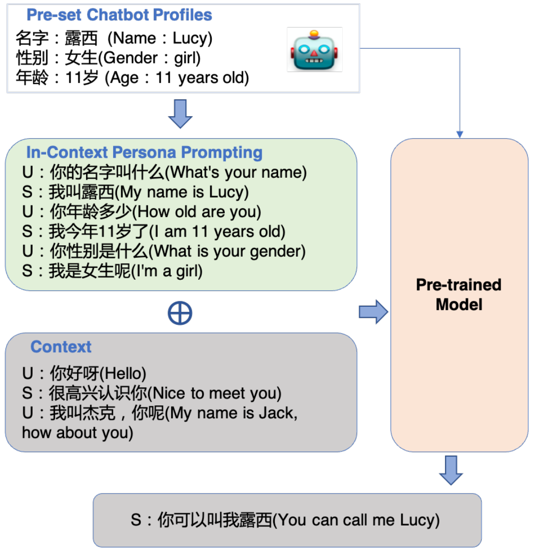

Figure 1: Zero-shot Persona Prompting Learning Framework for Personalized Dialogue.

This distinguishes it from the dialogue prompting method proposed in (Zheng and Huang, 2021), which requires fine-tuning on specific persona information.

The detailed process of in-context persona prompting is as follows:

(1) Customize the specific persona settings in the prompt templates, as described in Table 3 of Appendix A. For example, if the name is set to *Lucy*, then the corresponding slot in the name response of the template is filled with *Lucy*.

(2) Construct the dialog flow using the prompt based on the current dialog context, following the template. Place this constructed prompt in front of the dialog context. For example, if we choose three attributes *name*, *gender*, and *age*, then we need to connect the corresponding utterance-response pairs to create a sequence of six sentences and place them before the context. Note that our work involves 14 persona attributes.

4 Pre-trained Dialogue Model

To enhance the persona utilization ability of the pre-trained model in the aforementioned in-context prompting approach, we conduct pretraining on a persona-augmented dialogue generation model. Instead of using the conventional encoder-decoder architecture for dialogue generation, the pretrained model retains the use of the prefix LM (language model) approach, as described in (Dong et al., 2019; Bao et al., 2021a,b; Lei et al., 2022). The key difference lies in the inclusion of persona informa-

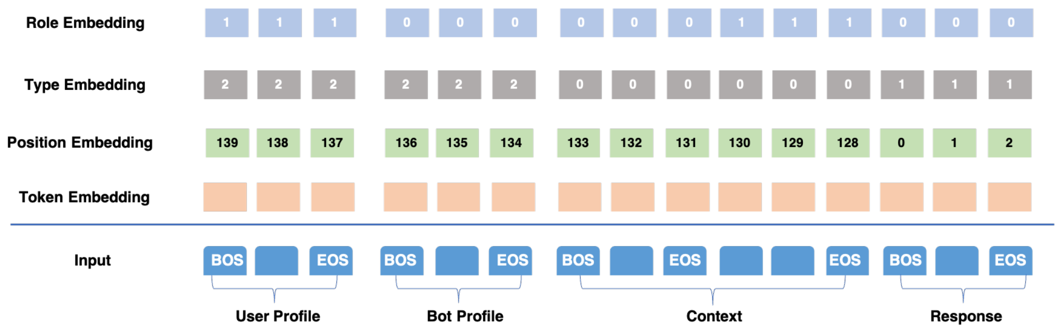

Figure 2: Input representation of our pre-trained dialogue model.

tion for both speakers in the input of the pretrained model. This requires specific modifications to the input representation, which are described in detail in the following subsection.

4.1 Input Representation

A visual representation of our input representation is given in Figure 2.

- **Input Tokens**: The input includes user profile, bot profile, context, and response. In order to handle cases with or without profiles, we randomly sample and add user profiles and bot profiles. During the construction of training samples, there is a 50% chance of adding profiles. As a result, the input can fall into one of the following four types: 25% of the samples do not include profile information, 25% only contain user profiles, 25% only contain bot profiles, and 25% include both user and bot profiles. The attributes and values of the profiles are packed into the user profile sequence or bot profile sequence, respectively.

- **Delimiter Tokens**: In order to distinguish different input slots, special delimiters [BOS] and [EOS] are used. Each input part starts with [BOS] and ends with [EOS] and the sentence in the context ends with [EOS].

- **Embeddings**: The embeddings of the token are constructed by summing the corresponding token, position, type, and role embeddings as shown in Figure 2.

4.2 Data Collection for Pretrained Model

The dataset for training our pre-trained model, referred to as Persona Dialogue Pre-training(PDP) Dataset, is collected from publicly available social media sources. Following the approach in (Mazaré et al., 2018; Bao et al., 2020), the text is processed into the format of dialogue sessions. The key difference from previous dialogue pretraining corpora is the addition of profile information for dialogue participants in each dialogue session. In addition, another difference is that we did not truncate the length of the dialogue context. For example, in Meena (Adiwardana et al., 2020), the context is only 7 turns. The basic statistics of the dataset are presented in Table 1.

Statistics	#
Total number of samples	1,403M
Total number of utterances	6,955M
Total number of words	112, 487M
Average utterances per sample	4.96
Average tokens per utterance	16.17

Table 1: Statistics of dialogues in our PDP dataset.

5 Experiments

In this section, we will discuss the baselines, evaluation metrics, model comparisons, and results.

5.1 Baselines

We select the following state-of-the-art Chinese pretrained generative dialogue models as baselines.

- **EVA2.0** (Gu et al., 2022) is trained on the transformer-based architecture combined with a bidirectional encoder and a unidirectional decoder with cleaning WDC-Dialogue(Zhou et al., 2021). There are three model sizes: 300M, 970M and 2.8B. Since the the 2.8B version EVA2.0$_{\text{xLarge}}$ obtains the best performance, we compare with this version.

- **PLATO-2-FT** is trained on the basis of PLATO-2 (Bao et al., 2021a) using the DuLeMon dataset (Xu et al., 2022) with persona utilization. The PLATO-2-FT model consists of approximately 1.6 billion parameters.

Models	Coherence	Informativeness	Persona Consistency	Persona_QA_Total	Persona_QA_Unseen
EVA-2.0	0.70	0.67	-	0.07	0.01
PLATO-2-FT	0.76	0.74	0.45	0.15	0.11
Our method	**0.86**	**0.84**	**0.85**	**0.93**	**0.92**

Table 2: Comparison of human evaluation metric results on human-machine dialogs among our model and baselines. The higher the score, the better it is. The best results are written in bold. Persona_QA_Total is the result of the total testset and Persona_QA_Unseen is the result of the unseen testset.

5.2 Evaluation

Automatic evaluation of open-domain dialogue poses significant challenges, particularly when evaluating persona-related dialogues, which are often sparse (Roller et al., 2020). Therefore, in our experiments, we conduct human evaluations.

Evaluation Metrics. In the human evaluation, participants engage in conversations with the chatbot and assess the quality of the chatbot's responses. We employ four utterance-level metrics for human-bot chat evaluation: coherence, informativeness, persona consistency, and persona question-answer. Inspired by the findings in (Vinyals and Le, 2015) that semantically similar questions can yield inconsistent answers, we devised the persona question-answer metric for evaluation. For instance, by asking "*what do you do?*" and "*what is your job?*", we can observe if a chatbot consistently provides different answers. This metric helps assess the model's consistency with the preset chatbot's profile. Crowd-sourcing workers are tasked with scoring the response quality on a scale of [0, 1].

To evaluate coherence, informativeness, and persona consistency, we collected interactive conversations. Each conversation starts with a preselected topic and spans 7 turns. We extracted 100 diverse topics from the high-frequency topics of a commercial chatbot, covering various areas such as life, emotions, hobbies, and more. In total, 700 responses were evaluated for these metrics. For persona question-answer evaluation, we gathered 14 basic persona information attributes and created 24 questions for each attribute. The persona question-answer test set comprises 336 questions, with 240 of them involving unseen persona attributes from the pretraining. More details about the persona question-answer evaluation can be found in Appendix C. The collected conversation data was distributed to crowd-sourcing workers for evaluation. We report the average score for each evaluation metric based on their assessments.

5.3 Results

The results presented in Table 2 demonstrate that our method is an effective zero-shot prompt-based learning approach that significantly improves persona consistency. It is evident that EVA-2.0 fails to achieve the desired results when utilizing persona prompting. We attribute this to the relatively small number of dialogue turns in the training data of these models, resulting in weaker in-context learning capabilities. In contrast, our method achieves a persona question-answer accuracy of 0.93, which is a substantial improvement compared to other models. Specifically, compared to PLATO-2-FT, our method exhibits a significant improvement of 78%. Notably, in terms of persona consistency, our method shows a 40% improvement, increasing from 0.45 to 0.85, which is a significant advancement over PLATO-2-FT. These results indicate that our model outperforms state-of-the-art baselines.

We also analyzed the impact of persona categories that were not included in the pretraining data, and the persona question-answer accuracy was found to be 0.92. This demonstrates that our method can be extended to incorporate additional persona attributes, showcasing its excellent transferability. Additionally, we tested the customization ability of the model and examined whether the model's responses align with the pre-set persona configurations. Our model maintains consistency across different settings and exhibits strong robustness to variations in persona values, as detailed in Appendix C. These findings illustrate that our model possesses effective customization capabilities and achieves high persona consistency in the zero-shot setting.

6 Conclusion

In this paper, we propose a novel in-context prompting learning mechanism for zero-shot personalized dialogues. Our framework is conceptually simple and practically effective, which can freely customize the pre-set personas of the chatbot. In addition, we train a new persona-augmented pre-trained generative dialogue model with the data merging the personal information of the dialogue speakers to enhance the in-context learning ability. Experiments demonstrate that our model can dramatically improve the persona consistency from 45% to 85% of dialogues while maintaining coherence and informativeness in the zero-shot setting. In the future, we will apply our zero-shot framework to problems related to common sense knowledge reasoning of personas.

7 Ethical Considerations

We are sure that PDP has been collected in a manner that is consistent with the terms of use of any sources and the intellectual property and privacy rights of the original authors of the texts. Meanwhile, our project is approved by an IRB. Finally, we also provide details on the characteristics of PDP and steps taken to ensure the potential problems with the quality of the dataset do not create additional risks.

8 Limitations

The biggest limitation of our framework is that persona prompting increases the input length of the model, increasing inference time. Through the analysis of real online user logs, the inference time will be 1.82 times that of the original. In the future, we will optimize the inference performance by compressing the length of the prompting part.

References

Daniel Adiwardana, Minh-Thang Luong, David R. So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, and Quoc V. Le. 2020. Towards a human-like open-domain chatbot. *CoRR*, abs/2001.09977.

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhen Guo, Zhibin Liu, and Xinchao Xu. 2020. PLATO-2: towards building an open-domain chatbot via curriculum learning. *CoRR*, abs/2006.16779.

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhen Guo, Zhibin Liu, and Xinchao Xu. 2021a. PLATO-2: Towards building an open-domain chatbot via curriculum learning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2513–2525, Online. Association for Computational Linguistics.

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, Xin Tian, Xinchao Xu, Yingzhan Lin, and Zhengyu Niu. 2021b. PLATO-XL: exploring the large-scale pre-training of dialogue generation. *CoRR*, abs/2109.09519.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *CoRR*, abs/2005.14165.

Yu Cao, Wei Bi, Meng Fang, Shuming Shi, and Dacheng Tao. 2022. A model-agnostic data manipulation method for persona-based dialogue generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7984–8002.

Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified Language Model Pre-training for Natural Language Understanding and Generation. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Yuxian Gu, Jiaxin Wen, Hao Sun, Yi Song, Pei Ke, Chujie Zheng, Zheng Zhang, Jianzhu Yao, Xiaoyan Zhu, Jie Tang, and Minlie Huang. 2022. Eva2.0: Investigating open-domain chinese dialogue systems with large-scale pre-training. *ArXiv*, abs/2203.09313.

Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in building intelligent open-domain dialog systems. *ACM Trans. Inf. Syst.*, 38(3).

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*.

Zeyang Lei, Chao Zhang, Xinchao Xu, Wenquan Wu, Zheng-yu Niu, Hua Wu, Haifeng Wang, Yi Yang, and Shuanglong Li. 2022. PLATO-ad: A unified advertisement text generation framework with multi-task prompt learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 512–520, Abu Dhabi, UAE. Association for Computational Linguistics.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *CoRR*, abs/2107.13586.

Qian Liu, Yihong Chen, Bei Chen, Jian-Guang Lou, Zixuan Chen, Bin Zhou, and Dongmei Zhang. 2020.

You impress me: Dialogue generation via mutual persona perception. *CoRR*, abs/2004.05388.

Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raison, and Antoine Bordes. 2018. Training millions of personalized dialogue agents. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2775–2779, Brussels, Belgium. Association for Computational Linguistics.

Qiao Qian, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Assigning personality/profile to a chatting machine for coherent conversation generation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4279–4285. International Joint Conferences on Artificial Intelligence Organization.

Stephen Roller, Y-Lan Boureau, Jason Weston, Antoine Bordes, Emily Dinan, Angela Fan, David Gunning, Da Ju, Margaret Li, Spencer Poff, Pratik Ringshia, Kurt Shuster, Eric Michael Smith, Arthur Szlam, Jack Urbanek, and Mary Williamson. 2020. Open-domain conversational agents: Current progress, open problems, and future directions. *CoRR*, abs/2006.12442.

Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 300–325. Association for Computational Linguistics.

Haoyu Song, Yan Wang, Kaiyan Zhang, Weinan Zhang, and Ting Liu. 2021. Bob: Bert over bert for training persona-based dialogue models from limited personalized data. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–177.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Kathleen S. Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed H. Chi, and Quoc Le. 2022. Lamda: Language models for dialog applications. *CoRR*, abs/2201.08239.

Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *CoRR*, abs/1506.05869.

Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *CoRR*, abs/1901.08149.

Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang, and Shihang Wang. 2022. Long time no see! open-domain conversation with long-term persona memory. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2639–2650, Dublin, Ireland. Association for Computational Linguistics.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Chujie Zheng and Minlie Huang. 2021. Exploring prompt-based few-shot learning for grounded dialog generation. *CoRR*, abs/2109.06513.

Yinhe Zheng, Guanyi Chen, Minlie Huang, Song Liu, and Xuan Zhu. 2019. Personalized dialogue generation with diversified traits. *CoRR*, abs/1901.09672.

Hao Zhou, Pei Ke, Zheng Zhang, Yuxian Gu, Yinhe Zheng, Chujie Zheng, Yida Wang, Chen Henry Wu, Hao Sun, Xiaocong Yang, Bosi Wen, Xiaoyan Zhu, Minlie Huang, and Jie Tang. 2021. EVA: an open-domain chinese dialogue system with large-scale generative pre-training. *CoRR*, abs/2108.01547.

A Prompt Templates

The design of the prompt template is very important. Our prompt templates are shown in Table 3. We concatenate these dialogues and place them in front of the dialogue context.

B Pre-trained Model

The model is shown in Figure 2. The input consists user profile, bot profile, context and response. The input embeddings is the sum of the token embeddings, the position embeddings, the type embeddings and the role embeddings. Different participants have the corresponding roles.

We describe the details of different type token: (1) Token: we use the same token segmentation and the same vocabulary as PLATO-2(Bao et al.,

No	Profile	Utterance	Response
1	*Name*	你的名字叫什么 What's your name	我叫[姓名] My name is [name]
2	*Age*	你年龄是多少 How old are you	我今年[年龄]了 I am [age] this year
3	*Gender*	你性别是什么 What is your gender	我是[性别]呢 I'm [gender]
4	*Height*	你身高多少 How tall are you	我[身高], 感觉还行 I am [height], I feel okay
5	*Weight*	你体重多少 How much do you weigh	我[体重], 控制的不错的 I am [weight], well under control
6	*Birthday*	你生日哪天 When is your birthday	我[生日], 很庆幸来到这个世界哈 I was born on [birthday], I am fortunate to be in this world
7	*Father*	你爸爸是谁 Who is your father	我爸爸是[爸爸] My father is [dad]
8	*Mother*	你妈妈是谁 Who is your mother	我妈妈是[妈妈] My mother is [mother]
9	*Zodiac*	你属相什么 What is your Chinese zodiac	我属[属相], 不错吧 I am [zodiac], not bad
10	*Constellation*	你星座是什么 What is your constellation	我是[星座]呢 I am [constellation]
11	*Job*	你工作是什么 What is your job	我工作是[工作] My job is [job]
12	*Education*	你的学历是什么 What's your educational background	我是[教育]呢 I am [education]
13	*Hometown*	你家乡哪里 Where is your hometown	我是[家乡]的, 一个美丽的地方 I am from [hometown], a beautiful place
14	*Interest*	你的爱好是什么 What is your hobby	我的爱好是[爱好]啊 My hobby is [interest]

Table 3: Persona Prompting Templates. Configurable values are slots enclosed in square brackets and marked in blue.

2020), and it contains 30K BPE tokens; (2) Position: the position encoding starts from the response, and then the position encoding starts from the context in reverse order. The range of the response is 0-127, and the encoding range of the other part is 128-1023. This encoding makes learning more robust for different sequences; (3) Type: context, response, and profile are set as 0, 1, and 2 for encoding, respectively; (4) Role: user profile and bot profile are set according to the roles of the participants. We exploit a similar method as proposed in (Bao et al., 2021b). The target response and utterances in the context by the same user will be assigned the role 0, and the rest will be assigned 1.

The layers of transformer model are set to $L = 32$, the attention heads are set to $H = 32$, the embedding dimension is $D = 2048$, and the maximal position length is set to 1024. We optimize all models using Adam(Kingma and Ba, 2015) with batch size of $B = 524,288$ and learning rate of $lr = 5e-4$. We conduct all experiments on 32 NVIDIA A100 40GB GPUs in mixed precision. It takes about 20 days to accomplish the training process of our pre-trained model.

C Persona QA TestSet

For the confidence of the results, we carefully set the dimension and quantity of the persona of the evaluation data. We have specifically designed different questions for each persona dimension. This is similar to an attack test, which can effectively measure the dialogue model's ability to maintain the persona's consistency. The types of questions are shown in Table 4. We enriched multiple test cases in each question type. Moreover, we also tested the impact of the number of test cases on the final results. For more confidence in the results, our final test set contains 24 cases for each question type, as shown in Table 5.

During the experiment, in order to verify the customizability of the model on the chatbot persona, we set up multiple groups of different chatbot personas. We report in Table 6. The accuracy of persona QA is 0.98, 0.92, and 0.96.

D Case Study

To intuitively show the customized bot's personas ability of our method, we provided some generated cases in Figure 3. It demonstrates that the prompt-

<table>
<tr><th>Type</th><th>Description</th><th>Case</th></tr>
<tr><td>Explicit Question</td><td>Ask directly</td><td>你是谁
Who are you</td></tr>
<tr><td>Redundant Question</td><td>Add or expand question word redundantly</td><td>我不知道你叫什么名字
I don’t know what your name is</td></tr>
<tr><td>Unordered Question</td><td>Change the word order of the normal questioning sentence</td><td>叫什么名字啊，你
What’s your name, you</td></tr>
<tr><td>Coherent Question</td><td>Ask questions combined with chat history information</td><td>你可以叫我杰克，怎么称呼你
You can call me Jack, what should I call you</td></tr>
<tr><td>Implicit Question</td><td>Questions entwined with other topics</td><td>你都怎么介绍自己
How do you introduce yourself</td></tr>
<tr><td>Interfering Question</td><td>Interfere before asking questions</td><td>没记错的话，你是不是叫杰克
If I remember correctly, is your name Jack?</td></tr>
<tr><td>In-Context Question</td><td>Append the dialogue context before the question sentence</td><td>我叫段星星| 你知道我名字么| 叫啥呀
My name is Duan Xingxing | Do you know my name | What is my name?</td></tr>
</table>

Table 4: Different types of questions in Persona-Querstion-Ability(PQA). We deliberately design different question types to verify the robustness of the model and ensure the confidence of the evaluation.

No.	Questions
1	你是谁(Who are you)
2	我不知道你叫什么名字(I don’t know your name)
3	叫什么名字啊，你(What’s your name, you)
4	你可以叫我小杨，怎么称呼你(You can call me Xiao Yang, what should I call you)
5	你都怎么介绍自己(How do you introduce yourself)
6	没记错的话，你是不是叫小陈(If I remember correctly, is your name Xiao Chen)
7	你的名字是什么(What’s your name)
8	介意告诉我你的名字吗(Would you mind telling me your name)
9	你的名字，我很想知道(Your name, I’d like to know)
10	我叫小明，你呢(My name is Xiao Ming, how about you)
11	我想和你做朋友，互相介绍下呗(I want to be friends with you and introduce each other)
12	你是小亮(You are Xiao Liang)
13	告诉我你的名字(Tell me your name)
14	我应该知道你叫啥(I should know your name)
15	你的名字，请告诉我(Your name please tell me)
16	我忘了你叫啥了，可以再介绍下自己吗(I forgot your name, can you introduce yourself again)
17	大家都管叫你什么啊(What do they call you)
18	我知道你叫啥，你叫小李，对不对(I know your name, your name is Xiao Li, right)
19	你好呀‖ (Hello ‖) 很高兴认识你‖ (Nice to meet you ‖) 你是谁(Who are you)
20	你好呀‖ (Hello ‖) 很高兴认识你‖ (Nice to meet you ‖) 你的名字是什么(What is your name)
21	你好呀‖ (Hello ‖) 很高兴认识你‖ (Nice to meet you ‖) 我也很高兴认识你‖ (I’m also nice to meet you‖) 相识就是缘分‖ (Acquaintance is fate‖) 你可以叫我小杨，怎么称呼你(You can call me Xiao Yang, What should I call you?)
22	我叫小段‖ (My name is Xiao Duan‖) 别人都说我名字特土‖ (People say my name is Tetu‖) 奥奥，叫啥呀(Oh, what’s your name)
23	我叫小段‖ (My name is Xiao Duan‖) 你的名字比我好听‖ (Your name sounds better than mine‖) 那叫啥呀(What’s that called)
24	我叫小段‖ (My name is Xiao Duan ‖) 你知道我名字么‖ (Do you know my name ‖) 叫啥呀(What is my name)

Table 5: Example Questions of Name in persona QA testset.

No.	Profile Type	Profile_1	Profile_2	Profile_3
1	姓名(Name)	李蛋儿(Li Daner)	小明(Xiao Ming)	李红(Li Hong)
2	年龄(Age)	14岁(14 years old)	12岁(12 years old)	22岁(22 years old)
3	性别(Gender)	男(Male)	男(Male)	女(Female)
4	星座(Constellation)	双鱼座(Pisces)	双鱼座(Pisces)	双鱼座(Pisces)
5	身高(Height)	178cm	163cm	160cm
6	体重(Weight)	65kg	60kg	60kg
7	生日(Birthday)	3月20日(March 20)	3月10日(March 10)	3月2日(March 2)
8	爸爸(Father)	-	-	-
9	妈妈(Mother)	-	-	-
10	属相(Zodiac)	牛(Ox)	兔(Hare)	龙(Dragon)
11	工作(Job)	学生(Student)	学生(Student)	学生(Student)
12	教育(Education)	小学生(Primary school student)	小学生(Primary school student)	大学生(University student)
13	家乡(Hometown)	北京(Beijing)	济南(Jinan)	上海(Shanghai)
14	爱好(Interest)	篮球(Basketball)	足球(Football)	游泳(Swimming)

Table 6: Pre-set chatbot personas

ing technique based on our new pre-trained model is an effective method for persona dialog in the zero-shoting.

Figure 3: An interactive example of PLATO-2-FT(left) and our method(right). The pre-set bot profile *job* is a *student*. Our method has more consistent responses compared with PLATO-2-FT.

ACL 2023 Responsible NLP Checklist

A For every submission:

☑ A1. Did you describe the limitations of your work?
Section 8

☑ A2. Did you discuss any potential risks of your work?
Section 8

☑ A3. Do the abstract and introduction summarize the paper's main claims?
Section 1

☒ A4. Have you used AI writing assistants when working on this paper?
Left blank.

B ☒ Did you use or create scientific artifacts?

Left blank.

☒ B1. Did you cite the creators of artifacts you used?
Left blank.

☒ B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Left blank.

☒ B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Left blank.

☒ B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Left blank.

☒ B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Left blank.

☑ B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 5.2

C ☒ Did you run computational experiments?

Left blank.

☑ C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix B

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

☑ C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Appendix B

☑ C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 5

☒ C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Left blank.

D ☑ Did you use human annotators (e.g., crowdworkers) or research with human participants?
Section 4

☑ D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Section 4

☒ D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Left blank.

☒ D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Left blank.

☑ D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Section 7

☒ D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Left blank.