

Dialogue Language Model with Large-Scale Persona Data Engineering

Mengze Hong^{1,2} Chen Jason Zhang¹ Chaotao Chen² Rongzhong Lian² Di Jiang^{2*}

¹Hong Kong Polytechnic University ²AI Group, WeBank Co., Ltd

Abstract

Maintaining persona consistency is paramount in the application of open-domain dialogue systems, as exemplified by models like ChatGPT. Despite significant advancements, the limited scale and diversity of current persona dialogue datasets remain challenges to achieving robust persona-consistent dialogue models. In this study, drawing inspiration from the success of large-scale pre-training, we introduce PPDS, an open-domain persona dialogue system that employs extensive generative pre-training on a persona dialogue dataset to enhance persona consistency. Specifically, we present a persona extraction model designed to autonomously and precisely generate vast persona dialogue datasets. Additionally, we unveil a pioneering persona augmentation technique to address the invalid persona bias inherent in the constructed dataset. Both quantitative and human evaluations consistently highlight the superior response quality and persona consistency of our proposed model, underscoring its effectiveness.

1 Introduction

The open-domain dialogue systems have gained significant interest due to their wide industrial applications, such as customer service support (Song et al., 2021b; Hong et al., 2024b), virtual assistance (Mandamadiotis et al., 2021), and social chatbots (Zhou et al., 2024; Ng et al., 2025). Inspired by the recent success of large-scale pre-training in natural language processing, many neural dialogue models resort to pre-training on large-scale dialogue datasets (Xu and Zhao, 2021) and demonstrate substantial progress in open-domain dialogue. Notable examples include DialoGPT (Zhang et al., 2020b), SPACE (He et al., 2022), and Blender (Shuster et al., 2022) for English dialogue, as well as CDialGPT (Wang et al., 2020), PLATO-2 (Bao et al., 2021), and EVA (Gu et al., 2023) for Chinese.

*Corresponding author: dijiang@webank.com

Persona:	My husband is a middle school teacher.
Context:	Are you married?
DialoGPT:	Not yet, I am still single.

Figure 1: Example of a persona-inconsistent response generated by DialoGPT.

While these methods can generate fluent and coherent responses, maintaining context consistency - particularly persona consistency - remains a common challenge that can lead to negative experiences in real-world human-bot interactions (Welleck et al., 2019; An et al., 2024). As illustrated in Figure 1, the well-trained DialoGPT model struggles with maintaining persona consistency, i.e., revealing contradicted marriage status and generating out-of-character responses. Similar issues appear more frequently with the emerging LLM-based chatbots such as ChatGPT and Claude, where although the chatbot can comprehend user intention correctly (Hong et al., 2024a), the response behavior often deviates from the instruction prompt and results in user dissatisfaction (Tseng et al., 2024; Dam et al., 2024; Song et al., 2024).

To enhance user experience and preserve persona consistency, many research efforts propose introducing explicit personas into the dialogue model (Qian et al., 2018; Wu et al., 2021; Chen and Wei, 2025; Takayama et al., 2025). For example, the PERSONA-CHAT dataset (Zhang et al., 2018) represents persona through personality sentences. While such crowd-sourced datasets capture a variety of persona features, their small scale, which is limited by the high cost of annotation, prevents them from fully unlocking the potential of large-scale neural dialogue models. On the other hand, the Personality Assignment Dataset (Qian et al., 2018) leverages persona attributes from users' social media profiles to automatically create a significantly larger persona dataset. However, the persona diversity is still limited by the attribute set of user

profiles. Some studies impose persona consistency using Natural Language Inference (NLI), but their effectiveness is still limited by domain mismatches with general NLI datasets or the scale of dialogue-specific NLI datasets (Welleck et al., 2019).

In this paper, we present an efficient solution for constructing large-scale and diverse persona dialogue data, based on which we further pre-train an open-domain persona dialogue model called **PPDS** (*Pre-trained Persona Dialogue System*) to achieve persona consistency. To construct the dataset, we propose a persona extraction model based on the existing dialogue NLI dataset (Welleck et al., 2019), using a summarization approach to automatically and accurately extract personas from large-scale dialogue datasets, such as Reddit comments (Baumgartner et al., 2020). Strict filtering rules have been implemented to ensure the quality of the persona dataset. Next, we train a large-scale Transformer-based model on the persona dialogue dataset, aiming to enhance its persona consistency through extensive pre-training. Finally, we conduct extensive quantitative and human evaluations to verify the superiority of our model. The contributions of this paper are summarized as follows:

1. We propose a persona extraction model to automatically construct large-scale persona dialogue datasets from existing dialogue corpora. Experiments on Reddit comments present a significantly larger and more diverse dataset than current public datasets built from user profiles or human annotations.
2. We develop a novel open-domain dialogue system pre-trained on the constructed large-scale persona dialogue dataset to enhance persona consistency. A new persona augmentation technique is introduced to address the persona bias issue in the dataset.
3. Extensive experiments involving both quantitative and human evaluations demonstrate the effectiveness of our model compared to various baselines. We analyze the roles of pre-training, persona augmentation, and fine-tuning as key components, providing insights and best practices for industrial deployment.

2 Related Work

2.1 Large-Scale Pre-Training

Large-scale pre-training has been a popular paradigm in natural language processing. With

large Transformer (Vaswani et al., 2017) model pre-training in massive plain texts and fine-tuning in downstream tasks, it has demonstrated substantial improvement and generality (Devlin et al., 2019; Liu et al., 2019). Recent attempts for larger models and data sizes further reveal the increasing potential of large-scale pre-training. Particularly, the GPT-3 (Brown et al., 2020) model with 175 billion parameters demonstrates strong zero-shot and few-shot learning capacities without task-specific fine-tuning on downstream tasks.

Motivated by the efficacy of large-scale pre-trained language models such as GPT-3 (Brown et al., 2020), UniLM (Dong et al., 2019) and T5 (Raffel et al., 2020), many recent efforts in dialogue try to build open-domain dialogue systems through large-scale pre-training on human-like dialogue. Equipped with large amounts of dialogue data collected from social media such as Reddit, Twitter, Weibo, etc, these models can generate human-like responses and enhance the engagingness of human-AI conversations. Although these methods have achieved substantial enhancements in open-domain dialogue, they still suffer from the consistency problem, especially persona consistency (Roller et al., 2021; Nie et al., 2021).

2.2 Persona Dialogue Model

To solve the problem of persona consistency, recent works focus on a data-driven approach where a persona dialogue dataset is introduced to capture the persona-related features. The persona include user identity (Li et al., 2016b), user profiles (Qian et al., 2018) and persona facts (Zhang et al., 2018; Mazaré et al., 2018). To leverage the persona information, many well-designed neural models are proposed, such as modeling mutual-persona (Liu et al., 2020) and multi-stage persona-based dialogue generation (Song et al., 2020a). Besides, there also exist many works (Wolf et al., 2019; Golovanov et al., 2019; Zheng et al., 2019; Roller et al., 2021; Lin et al., 2021) demonstrating that fine-tuning pre-trained models on persona dataset can obtain substantial improvement on persona consistency. However, due to the limitation of scale and diversity of the public persona dialogue dataset, these methods are still far from achieving satisfactory persona consistency.

In addition to capturing persona consistency implicitly, some works turn to explicitly imposing persona consistency by natural language inference (NLI). With an NLI model to judge whether a re-

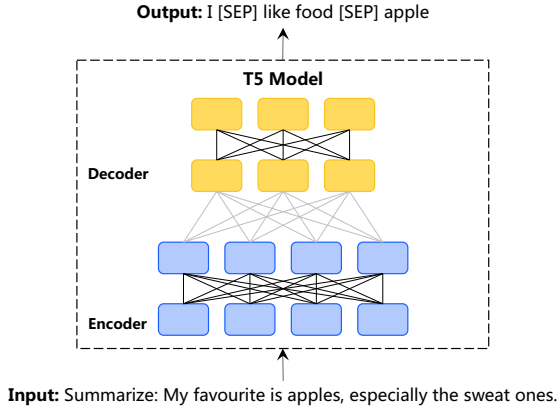


Figure 2: Overview of Persona Extraction Model.

sponse contradicts the personas, the dialogue models are able to improve their persona consistency by reranking (Welleck et al., 2019), unlikelihood training (Li et al., 2020; Song et al., 2021a) or reinforcement learning (Song et al., 2020b).

3 Large-scale Persona Dialogue Dataset

3.1 Persona Extraction Model

To construct the large-scale persona dialogue dataset, we first build a persona extraction model. Following (Welleck et al., 2019), we represent a persona as a triple, i.e. $p = \{e_1, r, e_2\}$, where e_1 , e_2 and r denote the subject, object and persona attribute respectively, e.g. (*i, like, swimming*). In particular, we propose to model the persona extraction problem as a summarization task, where the persona triple can be "summarized" from the utterance. Formally, given an utterance R , the persona extraction model outputs the corresponding persona p in a manner of generative summarization by considering the persona triple as a text $e_1 [SEP] r [SEP] e_2$, where the delimiter $[SEP]$ is used to distinguish each element in the persona triple. For utterances that are irrelevant to persona, we use a special token $[None]$ as their summarization, following the setting in (Welleck et al., 2019).

The overview of the persona extraction model is illustrated in Figure 2. Specifically, we leverage the Text-to-Text Transfer Transformer (T5) (Raffel et al., 2020) pre-training model as the backbone of our persona extraction model. T5 combines many language problems into a text-to-text format for multi-task learning, achieving superior performance in summarization tasks. It is also regarded as a cost-efficient model in the current landscape, making it suitable for industrial deployment.

We employ the Dialogue NLI (DNLI) dataset

#Session	#Utterance	#Persona	#Token	#Token/Utterance
189M	470M	36M	12B	25.5

Table 1: Statistics of the constructed large-scale persona dialogue dataset.

(Welleck et al., 2019) as the training corpus for persona extraction. The DNLI dataset is built upon the PERSONA-CHAT (Zhang et al., 2018) dataset by manually annotating persona triple for each utterance. The details of the persona attributes can be referred to in the original paper (Welleck et al., 2019). Compared with the dataset presented in (Qian et al., 2018), whose personas are limited by the attribute set of user profiles, the proposed method can capture more diversified personas. We fine-tune the T5-large model on the DNLI dataset. Ultimately, the persona extraction model achieves a ROUGE-L score of 80.0% on the DNLI test set, demonstrating its effectiveness in summarizing personas from utterances.

3.2 Data Construction

To build the large-scale persona dialogue dataset, we employ the well-trained persona extraction model to automatically extract the persona from the utterances in Reddit comments (Baumgartner et al., 2020), which consists of 5,601,331,385 comments. After extracting the persona of each utterance, the following summarized personas are removed to ensure persona quality:

- Personas that do not follow the format " $e_1 [SEP] r [SEP] e_2$ ";
- Personas with attributes outside of the predefined set of persona attributes;
- Personas whose subject exceeds 5 tokens;
- Personas with semantic cosine similarity to the original utterance below 0.1, as measured by the sentence-transformer library (Reimers and Gurevych, 2019).

Finally, we merge the personas from the same character in a dialogue session as a persona profile. Table 1 shows the statistics of the constructed large-scale persona dialogue dataset. To the best of our knowledge, this dataset is the largest of its kind, featuring a diverse range of personas beyond the scale of any existing datasets. It is also worth noting that the scale can be further expanded by leveraging a large dataset of utterances.

4 Large-Scale Pre-Training

4.1 Model

Based on the constructed large-scale persona dataset, we pre-train a Transformer-based dialogue model PPDS. Formally, let $C = \{c_1, c_2, \dots, c_N\}$ denotes the dialogue context which consists of N utterances, R denotes the target response, and $P = \{p_1, p_2, \dots, p_M\}$ denotes the personas which consists of M triples of persona. The target of the proposed model M is to generate a persona consistent response \hat{R} based on both persona P and dialogue context C , i.e., $\hat{R} = M(C, P)$.




Figure 3: Network architecture of PPDS.

The network architecture of PPDS is illustrated in Figure 3. Similar to the existing pre-training dialogue model, it employs Transformer blocks as the backbone. In order for efficient training on large-scale datasets, PPDS adopts the unified Transformer (also known as UniLM (Dong et al., 2019)) instead of the typical encoder-decoder architecture for dialogue generation. By concatenating the persona, dialogue context, and response as a single input, the UniLM architecture can significantly reduce unnecessary computation of padding. The flexible mechanism of the self-attention mask can also simultaneously model the two tasks of dialogue context understanding and response generation with sharing parameters. Therefore, the UniLM architecture is more parameter-efficient than the encoder-decoder network (Bao et al., 2021). Additionally, UniLM has demonstrated promising performance across various downstream tasks (Huang et al., 2022; Bao et al., 2020), highlighting its superiority and suitability.

As shown in Figure 4, the model input is the concatenation of persona, dialogue context, and response. Its representation is calculated as the sum of the corresponding token, position, turn, and

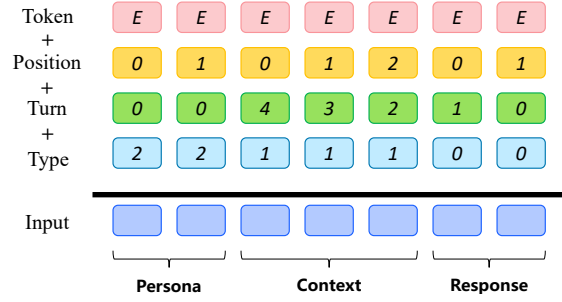


Figure 4: Input representation of PPDS.

type embeddings. The *token* is a BPE token in English input or a character in Chinese input. The *position* is the index of the token in an utterance. The *turn* is the turn distance of the utterance in dialogue context toward the target response. We assume that the closer utterances would be more relevant to the target response. Specifically, the turn index of target response and persona is 0. The *type* is used to distinguish the characters in the dialogue, where 0 refers to the responder (i.e., bot), 1 refers to the respondent (i.e., human), and 2 refers to the persona profiles.

In detail, the utterances in dialogue context are separated by a special token $[SEP]$. The persona is represented by the concatenation of persona triple (i.e., " $\{head\} \{relation\} \{tail\}$ "), and each persona is also separated by a special token $[SEP]$. Once we get the input representations, the UniLM Transformer will perform multi-head attention (Vaswani et al., 2017) on the input to transform the embeddings into a sequence of hidden representations H . Finally, we leverage *Softmax* to transform the hidden representations into the predictive probability of the target response (Jiang et al., 2023; Li et al., 2023). The details of the Transformer structure can refer to (Dong et al., 2019).

4.2 Training Objectives

In PPDS, the pre-training objective is to minimize the widely adopted negative log-likelihood (NLL) loss as follows:

$$\mathcal{L}_{NLL} = -E_{(P,C,R)}[\log p_{\theta}(\mathbf{R}|\mathbf{P}, \mathbf{C})] \quad (1)$$

$$= -E_{(P,C,R)}[\log p_{\theta}(r_t|\mathbf{P}, \mathbf{C}, r_{<t})] \quad (2)$$

where θ refers to the trainable parameters, T is the length of the target response and $r_{<t}$ denotes previously generated words.

4.3 Persona Augmentation

Since the large-scale persona dataset is built by extracting personas from the responses, there exists a persona bias in the raw persona dataset. Characters with extracted personas are often linked to persona-related responses, potentially misleading the model to generate such responses whenever a persona profile is present, regardless of its relevance to the dialogue context.

To mitigate this bias, we propose augmenting the dialogue with unrelated personas, compelling the model to identify the relevant persona based on the dialogue context. Specifically, we collect all extracted personas and sample some to supplement each dialogue’s persona profiles. If the sampled persona is of the same type as the existing one, we remove it to avoid contradictions. Finally, we merge this augmented dataset of unrelated personas with the raw persona dataset to eliminate the bias.

4.4 Pre-Training Details

We employ the augmented large-scale persona dataset as the pre-training corpora. To guarantee the data quality, we follow the elaborate cleaning process as PLATO-2 (Bao et al., 2021). After filtering, the data is split into training and validation sets. The training set contains 211M samples, and the validation set contains 0.2M samples. We reuse the BERT-base-uncased vocabulary (Devlin et al., 2019). The maximum sequence length for the persona, dialogue context, and target response are all set to 128. We use Adam as the optimizer with a learning rate scheduler of linear warmup and cosine decay. The warmup stage covers the first 30000 steps, and the peak learning rate is $5e-5$. The training of the model was done on 8 Nvidia Tesla V100 32G GPUs with a batch size of 256.

5 Experiments

5.1 Experiment Setup

We evaluate our models on persona dialogue generation with PERSONA-CHAT (Zhang et al., 2018) which is a crowd-sourced dataset covering rich persona features. The training and test sets are used for fine-tuning and evaluation, respectively.

Baselines. To evaluate the performance of PPDS, the following dialogue generation models including non-fine-tuned and fine-tuned ones are compared in the experiments.

- **Baseline:** Our vanilla PPDS, trained from scratch on PERSONA-CHAT without pre-training on the large-scale personal dialogue dataset.
- **DialoGPT:** Pre-trained on GPT-2 (Radford et al., 2019) using Reddit comments. We compare its medium version, which reports the best performance (Zhang et al., 2020b).
- **DialoGPT-finetuned:** Fine-tuned DialoGPT on PERSONA-CHAT by concatenating the persona with the dialogue context.
- **PPDS:** Our proposed model pre-trained on the large-scale persona dialogue dataset with persona augmentation.
- **PPDS-woP:** Pre-trained without a persona.
- **PPDS-woA:** Pre-trained without persona augmentation.
- **PPDS-finetuned:** Our PPDS pre-trained on the large-scale persona dialogue dataset with persona augmentation and fine-tuned on the PERSONA-CHAT.
- **PPDS-woP-finetuned:** PPDS fine-tuned on PERSONA-CHAT.
- **PPDS-woA-finetuned:** Pre-trained without persona augmentation and fine-tuned on PERSONA-CHAT.

5.2 Evaluation Metrics

We evaluate the response quality and persona consistency of the personal dialogue generation through both quantitative and human evaluations. For dialogue quality, we follow common practice (Zhang et al., 2018) to employ the following quantitative metrics: (1) Perplexity (**PPL**). Lower perplexity means better language modeling. (2) Distinct 1/2 (**Dist-1/2**) (Li et al., 2016a) denotes the ratio of distinct uni-grams/bi-grams, where higher distinct means better response diversity. (3) BertScore (**BS**) (Zhang et al., 2020a) measures the coherence similarity between predicted response and target response measured through the BERT model. For persona consistency, we employ the ratios of responses that are entailed (**E**), neutral (**N**), and contradicted (**C**) to the personas, which are measured by an NLI model. We also calculate Consistency Score (**CS**) (Madotto et al., 2019) to

Method	PPL ↓	Dist-1/2 ↑	BS ↑	E ↑	N	C ↓	CS ↑	Flu. ↑	Cohe. ↑	Info. ↑	P.C. ↑
Baseline	43.48	1.24/7.41	85.99	11.0	81.4	7.6	5.1	1.74	1.02	0.24	-0.20
DialoGPT	-	5.00/21.61	85.23	8.1	86.3	5.6	4.5	1.72	1.16	0.30	-0.12
PPDS-woP	20.19	3.76/19.42	86.00	11.6	83.0	5.3	10.7	1.76	1.26	0.44	0.02
PPDS-woA	18.19	3.19/16.67	86.16	33.8	58.8	7.4	41.6	1.74	1.40	0.88	0.14
PPDS	18.24	3.33/17.65	86.23	42.7	51.5	5.8	49.5	1.92	1.54	1.00	0.42
DialoGPT-finetuned	-	4.04/20.61	86.58	31.0	61.9	7.0	30.2	2.00	1.54	0.76	0.16
PPDS-woP-finetuned	15.93	3.10/14.86	86.38	19.2	75.8	5.0	18.1	1.98	1.56	0.56	0.04
PPDS-woA-finetuned	15.41	3.00/15.42	86.56	37.0	57.9	5.0	40.6	1.98	1.66	1.02	0.32
PPDS-finetuned	15.21	3.02/15.83	86.61	39.1	56.8	4.1	44.3	2.00	1.80	1.14	0.44

Table 2: Quantitative and human evaluation results. The best results are highlighted in **bold**.

measure persona consistency, which summarizes the result of NLI as follows:

$$NLI(\mathbf{R}, \mathbf{P}_i) = \begin{cases} -1, & \text{if } \mathbf{R} \text{ contradicts } \mathbf{P}_i, \\ 0, & \text{if } \mathbf{R} \text{ is neutral to } \mathbf{P}_i, \\ 1, & \text{if } \mathbf{R} \text{ entails } \mathbf{P}_i. \end{cases} \quad (3)$$

$$CS(\mathbf{R}) = \sum_{i=1}^M NLI(\mathbf{R}, \mathbf{P}_i) \quad (4)$$

The NLI model is fine-tuned on the DNLI dataset (Welleck et al., 2019) based on the pre-trained large RoBERTa model (Liu et al., 2019), achieving test set accuracy of 93.3% on DNLI.

As suggested in the empirical study (Liu et al., 2016), the correlation between quantitative metrics and human judgments may be weak in open-domain dialogue generation. Therefore, we also conduct human evaluations in the experiments. Thirty crowd-sourced workers were recruited based on criteria of diversity (Zhang et al., 2024) and strong proficiency in dialogue tasks. They were tasked with evaluating responses based on the following metrics: (1) Fluency (**Flu.**) measuring whether the response is fluent and grammatically correct. (2) Coherence (**Cohe.**) measuring whether the response is relevant and consistent with the context. (3) Informativeness (**Info.**) evaluating whether the response is informative or not given the context. (4) Persona Consistency (**P.C.**) checking whether the response has conflicts with the persona, where 1 means persona-related and consistent, 0 means neutral, and -1 means contradicted. Fluency, Coherence, and Informativeness are all rated on a scale of [0, 1, 2], with higher scores indicating better quality.

Persona	I like to drink wine
Context	... A: Which do you prefer white or red? B: Fermented. I prefer well-aged and fermented. I could bathe in it! A: Good choice. I always like a nice dry white wine.
PPDS-woP	My secret is that I don't drink . I don't know why.
PPDS	Me too. My body would be so strong if I had a dry wine !

Table 3: A qualitative example of a persona-consistent response generated by our model (PPDS) compared to a model pre-trained without a persona (PPDS-woP).

5.3 Evaluation Result

The evaluation results are summarized in Table 2. The baseline model struggles to perform well due to the limited scale of the PERSONA-CHAT dataset, which is inadequate for its large parameters. In contrast, all methods with large-scale pre-training generate fluent and coherent responses, demonstrating the benefits of such pre-training. However, DialoGPT and PPDS-woP exhibit poor persona consistency with low *CS* and *P.S.* scores. Although fine-tuning improves the persona consistency of DialoGPT-finetuned and PPDS-woP-finetuned, the performance remains unsatisfactory.

Our PPDS-woA and PPDS models, pre-trained on the large-scale persona dialogue dataset, achieve significantly better persona consistency scores, surpassing both PPDS-woP-finetuned and DialoGPT-finetuned by a large margin. This indicates that large-scale pre-training on persona dialogue data can greatly enhance the persona consistency of dialogue models. Additionally, with our proposed persona augmentation, PPDS demonstrates superior persona consistency and response quality in both quantitative and human evaluations compared to PPDS-woA, confirming the effectiveness of persona augmentation in mitigating bias in the constructed dataset. The improvements in reducing

contradictions and enhancing coherence and informativeness are particularly notable. Ultimately, through pre-training on the large-scale persona dialogue dataset with persona augmentation and subsequent fine-tuning on the PERSONA-CHAT dataset, our PPDS-finetuned achieves the highest scores in most quantitative and human evaluations, showcasing its superior understanding of persona consistency. A qualitative example is presented in Table 3 to further illustrate the effectiveness of our model in maintaining persona consistency.

6 Conclusion

In this work, we introduce a summarization-based persona extraction model to construct a large-scale persona dialogue dataset. Based on the dataset, we propose PPDS, an open-domain persona dialogue system that leverages large-scale pre-training for achieving persona consistency in dialogue generation. Both quantitative and qualitative evaluations demonstrate the effectiveness of our approach. Given that the experiments were conducted with relatively cost-efficient models and still yielded promising results, this work encourages future research to apply these techniques in building large-scale dialogue models and enhancing dialogue generation systems for industrial applications.

Beyond the discussed techniques, we also encourage exploration towards the following directions in constructing better persona datasets and training persona-consistent dialogue models for different application scenarios. First, the recent emergence of LLM-in-the-loop methodologies (Hong et al., 2025) offers a promising path by incorporating the natural language understanding capabilities of LLMs to enhance the persona extraction process for complex, compositional personas. Second, the source of extraction can extend from textual data to multimodal datasets, particularly conversational speech that contains rich persona information (Song et al., 2020c, 2022). Lastly, extending the persona from individual behaviors to larger entities, such as brand personality (Aaker, 1997), would further enhance the practical value of the proposed methods in various downstream domains, such as the hospitality and service sectors (Ng et al., 2024).

References

Jennifer L Aaker. 1997. Dimensions of brand personality. *Journal of marketing research*, 34(3):347–356.

Ruichuan An, Sihan Yang, Ming Lu, Kai Zeng, Yulin Luo, Ying Chen, Jiajun Cao, Hao Liang, Qi She, Shanghang Zhang, and Wentao Zhang. 2024. [McIlava: Multi-concept personalized vision-language model](#). *Preprint*, arXiv:2411.11706.

Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Songhao Piao, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2020. Unilmv2: pseudo-masked language models for unified language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning*, ICML’20. JMLR.org.

Siqi Bao, Huang He, Fan Wang, Hua Wu, Haifeng Wang, Wenquan Wu, Zhen Guo, Zhibin Liu, and Xinchao Xu. 2021. [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.

Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. In *ICWSM*, pages 830–839.

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. In *NeurIPS 2020*.

Kuiyun Chen and Yanbin Wei. 2025. [Upes: Unbiased persona construction for dialogue generation](#). In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.

Sumit Kumar Dam, Choong Seon Hong, Yu Qiao, and Chaoning Zhang. 2024. [A complete survey on llm-based ai chatbots](#). *Preprint*, arXiv:2406.16937.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186.

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 *NeurIPS*, pages 13042–13054.

Sergey Golovanov, Rauf Kurbanov, Sergey I. Nikolenko, Kyryl Truskovskiy, Alexander Tselousov, and Thomas Wolf. 2019. Large-scale transfer learning for natural language generation. In *ACL*, pages 6053–6058.

- Yuxian Gu, Jiaxin Wen, Hao Sun, Yi Song, Pei Ke, Chujie Zheng, Zheng Zhang, Jianzhu Yao, Lei Liu, Xiaoyan Zhu, et al. 2023. Eva2. 0: Investigating open-domain chinese dialogue systems with large-scale pre-training. *Machine Intelligence Research*, 20(2):207–219.
- Wanwei He, Yinpei Dai, Min Yang, Jian Sun, Fei Huang, Luo Si, and Yongbin Li. 2022. [Unified dialog model pre-training for task-oriented dialog understanding and generation](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 187–200, New York, NY, USA. Association for Computing Machinery.
- Mengze Hong, Wailing Ng, Yifei Wang, Di Jiang, Yuanfeng Song, Chen Jason Zhang, and Lei Chen. 2025. Position: Towards llm-in-the-loop machine learning for future applications.
- Mengze Hong, Yuanfeng Song, Di Jiang, Wailing Ng, Yanjie Sun, and Chen Jason Zhang. 2024a. [Dial-in llm: Human-aligned dialogue intent clustering with llm-in-the-loop](#). *Preprint*, arXiv:2412.09049.
- Mengze Hong, Yuanfeng Song, Di Jiang, Lu Wang, Zichang Guo, and Chen Jason Zhang. 2024b. [Expanding chatbot knowledge in customer service: Context-aware similar question generation using large language models](#). *Preprint*, arXiv:2410.12444.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. [Layoutlmv3: Pre-training for document ai with unified text and image masking](#). In *Proceedings of the 30th ACM International Conference on Multimedia, MM '22*, page 4083–4091, New York, NY, USA. Association for Computing Machinery.
- Di Jiang, Chen Zhang, and Yuanfeng Song. 2023. *Probabilistic topic models: Foundation and application*. Springer.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL HLT*, pages 110–119.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016b. A persona-based neural conversation model. In *ACL*.
- Margaret Li, Stephen Roller, Ilia Kulikov, Sean Welleck, Y-Lan Boureau, Kyunghyun Cho, and Jason Weston. 2020. Don't say that! making inconsistent dialogue unlikely with unlikelihood training. In *ACL*, pages 4715–4728.
- Yawen Li, Di Jiang, Rongzhong Lian, Xueyang Wu, Conghui Tan, Yi Xu, and Zhiyang Su. 2023. [Heterogeneous latent topic discovery for semantic text mining](#). *IEEE Transactions on Knowledge and Data Engineering*, 35(1):533–544.
- Zhaojiang Lin, Andrea Madotto, Yejin Bang, and Pascale Fung. 2021. The adapter-bot: All-in-one controllable conversational model. In *IAAI*, pages 16081–16083.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *EMNLP*, pages 2122–2132.
- 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. In *ACL*, pages 1417–1427.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In *ACL*, pages 5454–5459.
- Antonis Mandamadiotis, Georgia Koutrika, Stavroula Eleftherakis, Apostolis Glenis, Dimitrios Skoutas, and Yannis Stavrakas. 2021. Datagent: The imminent age of intelligent data assistants. In *VLDB*, pages 2815–2818.
- Pierre-Emmanuel Mazaré, Samuel Humeau, Martin Raïson, and Antoine Bordes. 2018. Training millions of personalized dialogue agents. In *EMNLP*, pages 2775–2779.
- Wailing Ng, Fei Hao, and Chen Zhang. 2024. From function to relation: Exploring the dual influences of warmth and competence on generative artificial intelligence services in the hospitality industry. *Journal of Hospitality & Tourism Research*, page 10963480241292016.
- Wailing Ng, Fei Hao, and Chen Zhang. 2025. [Avatar for hotels green training](#). *International Journal of Hospitality Management*, 126:104068.
- Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, and Jason Weston. 2021. I like fish, especially dolphins: Addressing contradictions in dialogue modeling. In *ACL/IJCNLP*, pages 1699–1713.
- Qiao Qian, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Assigning personality/profile to a chatting machine for coherent conversation generation. In *IJCAI*, pages 4279–4285.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits

- of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP-IJCNLP*, pages 3980–3990.
- 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 *EACL*, pages 300–325.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, Morteza Behrooz, William Ngan, Spencer Poff, Naman Goyal, Arthur Szlam, Y-Lan Boureau, Melanie Kambadur, and Jason Weston. 2022. [Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage](#). *Preprint*, arXiv:2208.03188.
- Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, and Ting Liu. 2021a. Bob: BERT over BERT for training persona-based dialogue models from limited personalized data. In *ACL/IJCNLP*, pages 167–177.
- Haoyu Song, Yan Wang, Weinan Zhang, Xiaojiang Liu, and Ting Liu. 2020a. Generate, delete and rewrite: A three-stage framework for improving persona consistency of dialogue generation. In *ACL*, pages 5821–5831.
- Haoyu Song, Wei-Nan Zhang, Jingwen Hu, and Ting Liu. 2020b. Generating persona consistent dialogues by exploiting natural language inference. In *AAAI*, pages 8878–8885.
- Yuan-Feng Song, Yuan-Qin He, Xue-Fang Zhao, Han-Lin Gu, Di Jiang, Hai-Jun Yang, and Li-Xin Fan. 2024. A communication theory perspective on prompting engineering methods for large language models. *Journal of Computer Science and Technology*, 39(4):984–1004.
- Yuanfeng Song, Di Jiang, Xiaoling Huang, Yawen Li, Qian Xu, Raymond Chi-Wing Wong, and Qiang Yang. 2020c. [Goldenretriever: A speech recognition system powered by modern information retrieval](#). In *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, page 4500–4502, New York, NY, USA. Association for Computing Machinery.
- Yuanfeng Song, Raymond Chi-Wing Wong, Xuefang Zhao, and Di Jiang. 2022. [Voicequerysystem: A voice-driven database querying system using natural language questions](#). In *Proceedings of the 2022 International Conference on Management of Data*, SIGMOD '22, page 2385–2388, New York, NY, USA. Association for Computing Machinery.
- Yuanfeng Song, Xuefang Zhao, Di Jiang, Xiaoling Huang, Weiwei Zhao, Qian Xu, Raymond Chi-Wing Wong, and Qiang Yang. 2021b. Smartsales: An ai-powered telemarketing coaching system in fintech. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2774–2776.
- Junya Takayama, Masaya Ohagi, Tomoya Mizumoto, and Katsumasa Yoshikawa. 2025. [Persona-consistent dialogue generation via pseudo preference tuning](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5507–5514, Abu Dhabi, UAE. Association for Computational Linguistics.
- Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. [Two tales of persona in LLMs: A survey of role-playing and personalization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16612–16631, Miami, Florida, USA. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, pages 5998–6008.
- Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong Jiang, Xiaoyan Zhu, and Minlie Huang. 2020. A large-scale chinese short-text conversation dataset. In *NLPCC*, volume 12430, pages 91–103.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. In *ACL*, pages 3731–3741.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Yuwei Wu, Xuezhe Ma, and Diyi Yang. 2021. Personalized response generation via generative split memory network. In *NAACL-HLT*, pages 1956–1970.
- Yi Xu and Hai Zhao. 2021. [Dialogue-oriented pre-training](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2663–2673, Online. Association for Computational Linguistics.
- Chen Jason Zhang, Yunrui Liu, Pengcheng Zeng, Ting Wu, Lei Chen, Pan Hui, and Fei Hao. 2024. Similarity-driven and task-driven models for diversity of opinion in crowdsourcing markets. *The VLDB Journal*, 33(5):1377–1398.
- 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 *ACL*, pages 2204–2213.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. Bertscore: Evaluating text generation with BERT. In *ICLR*.

- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. DIALOGPT : Large-scale generative pre-training for conversational response generation. In *ACL*, pages 270–278.
- Yinhe Zheng, Guanyi Chen, Minlie Huang, Song Liu, and Xuan Zhu. 2019. Personalized dialogue generation with diversified traits. *arXiv preprint arXiv:1901.09672*.
- Jinfeng Zhou, Zhuang Chen, Dazhen Wan, Bosi Wen, Yi Song, Jifan Yu, Yongkang Huang, Pei Ke, Guanqun Bi, Libiao Peng, JiaMing Yang, Xiyao Xiao, Sahand Sabour, Xiaohan Zhang, Wenjing Hou, Yijia Zhang, Yuxiao Dong, Hongning Wang, Jie Tang, and Minlie Huang. 2024. [CharacterGLM: Customizing social characters with large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1457–1476, Miami, Florida, US. Association for Computational Linguistics.