

# Long Time No See! Open-Domain Conversation with Long-Term Persona Memory

Xinchao Xu<sup>1\*</sup>, Zhibin Gou<sup>1,2\*</sup>, Wenquan Wu<sup>1</sup>, Zheng-Yu Niu<sup>1</sup>,  
Hua Wu<sup>1</sup>, Haifeng Wang<sup>1</sup> and Shihang Wang<sup>3</sup>

<sup>1</sup>Baidu Inc., China

<sup>2</sup>School of Computer Science, Beijing University of Posts and Telecommunications

<sup>3</sup>Columbia University

{xinchaoxu, wuwenquan01, niuzhengyu, wu\_hua, wanghaifeng}@baidu.com  
zebgou@gmail.com, sw3275@columbia.edu

## Abstract

Most of the open-domain dialogue models tend to perform poorly in the setting of long-term human-bot conversations. The possible reason is that they lack the capability of understanding and memorizing long-term dialogue history information. To address this issue, we present a novel task of Long-term Memory Conversation (LeMon) and then build a new dialogue dataset DuLeMon and a dialogue generation framework PLATO-LTM with a Long-Term Memory (LTM) mechanism. This LTM mechanism enables our system to accurately extract and continuously update long-term persona memory without requiring multiple-session dialogue datasets for model training. To our knowledge, this is the first attempt to conduct real-time dynamic management of persona information of both parties, including the user and the bot. Results on DuLeMon indicate that PLATO-LTM can significantly outperform baselines in terms of long-term dialogue consistency, leading to better dialogue engagingness<sup>1</sup>.

## 1 Introduction

Persona is crucial for open-domain dialogue systems to establish long-term intimacy with users (Huang et al., 2020). Existing persona dialogue datasets such as PersonaChat (Zhang et al., 2018; Dinan et al., 2019) and models (Li et al., 2016a; Zhang et al., 2017; Qian et al., 2018) have greatly facilitated the chatbot with configurable and persistent personalities.

Nevertheless, current open-domain dialogue systems still cannot build a long-term connection with humans. The possible reason is that they lack the capability of understanding and memorizing long-term dialogue history information, which we called

\* Equal contribution. The work was done when Zhibin Gou and Shihang Wang were doing internship at Baidu.

<sup>1</sup>Our data and codes are released at <https://github.com/PaddlePaddle/Research/tree/master/NLP/ACL2022-DuLeMon>

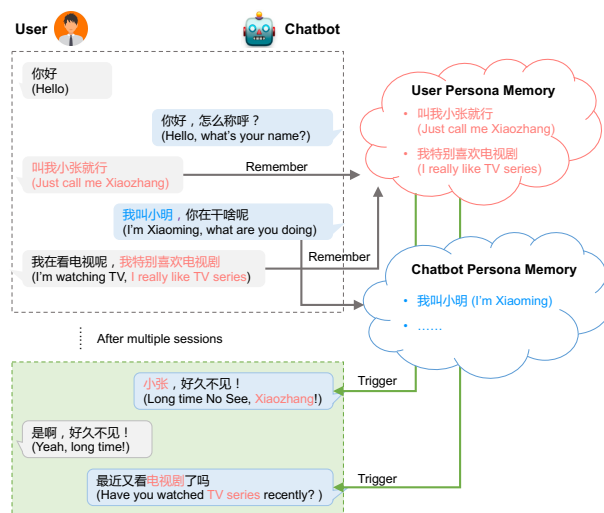


Figure 1: A sample of long-term conversation with memory. At first, the chat partner is not familiar with each other, so the goal is to get to know each other; Then, after multiple sessions, the chatbot already has a certain understanding and memory of the user's persona and its own persona, making the deep chat possible.

**long-term persona ability.** Remembering and actively utilizing the user's persona increases engagingness and contributes to long-term friendships between chatbot and user (Campos et al., 2018). Without this ability, the current state-of-the-art models, such as Meena (Adiwardana et al., 2020), Blender (Roller et al., 2021), and PLATO (Bao et al., 2020), tend to talk to people like strangers in long-term conversations.

Despite the importance and challenge of utilizing long-term persona in open-domain dialogue, as far as we know, the long-term persona ability of large-scale models is less studied due to a lack of both task design and corresponding dataset. Previous long-term persona dialogue systems (Kim et al., 2014; Bang et al., 2015) are mainly rule-based systems without large-scale pre-training models, in which researchers proposed various episodic memory architectures to extract, store and manage rel-

evant facts in prior interactions for use in future dialogs (Campos et al., 2018).

In addition, existing persona conversation datasets (Zhang et al., 2018; Dinan et al., 2019; Zheng et al., 2019) focus only on the consistency of the chatbot’s own persona and ignore the memory and utilization of the user’s persona. And they all set fixed persona that cannot be updated during the chat. Recently, Xu et al. (2021) proposed MSC dataset as a multi-session extension of PersonaChat, and its sessions are additionally annotated with summaries of important personal points. Similar to the previous episodic memory architecture, Xu et al. (2021) summarize and recall previous conversations for future dialogue generation. The stored documents in MSC will not be dynamically modified and will increase infinitely as the conversation progresses. Furthermore, the retrieval-augmented generative models rely on a long-session conversation dataset for training, which is expensive and difficult to annotate.

To address the limitations of existing models and the above issues, we define the **LeMon** (Long-term Memory Conversation) task and propose a new dataset named **DuLeMon**, which focuses not only on the consistency of the bot’s own persona but also on the active construction and utilization of the user’s persona in a long-term interaction (ie. mutual persona). We demonstrate an example dialogue in DuLeMon in Figure 1. In DuLeMon, we assume that the two speakers have previously interacted with each other and that the chatbot remembers part of the user’s persona. Besides, both the user and chatbot grounding persona are annotated in each utterance.

Based on our collected dataset, we carefully design a novel **PLATO-LTM** framework for the long-term persona dialogue setting by adding a plug-and-play long-term memory (LTM) to the state-of-the-art open-domain dialogue model (Bao et al., 2020). It enables us to study long-term persona conversations without relying on the long-session dataset. PLATO-LTM can extract both parties’ persona information from the conversation in real time, write it to persona memory respectively, and retrieve both parties’ persona information from memory to generate responses. The PLATO-LTM framework consists of three modules: (1) Persona Extractor (PE): The memory is updated by filtering irrelevant information and extracting persona sentences through a classifier. (2) Long-Term Memory (LTM): Two

separated long-term memories store the explicit persona information of interlocutors. (3) Generation Module: We use the large-scale model and the retrieved persona sentences of the user and chatbot are directly concatenated with dialogue context as model input.

Our major contributions are as follows:

- (1) We firstly propose the long-term persona chat task LeMon for Chinese long-term conversations. Our proposed DuLeMon dataset is also the largest multi-turn Chinese mutual persona chat dataset currently available.
- (2) We proposed a PLATO-LTM framework that extracts and remembers both user’s and the chatbot’s persona in real time, enabling the chatbot to have long-term persona dialogue without training on long-session data.
- (3) Automatic and human evaluation show that our method significantly improves the consistency of the state-of-the-art in long conversations, making the response more engaging while ensuring coherency.

## 2 Related Work

**Persona Dialogue:** As described in Huang et al. (2020), there is much work related to persona dialogue. Generally speaking, these works can be divided into implicit persona models and explicit persona models. In the implicit model, the persona is represented in the form of the semantic persona vector. Kim et al. (2014) proposed a retrieval-based method to integrate persona and user interests into the dialogue system. Because these models are implicit methods, they are not easy to interpret and control in target response generation. In Qian et al. (2018), an explicit persona model is proposed to generate consistent responses for given persona information. The persona information of the machine includes name, gender, hobbies, and so on. In this way, the given persona information can be better used for model generation. There are also many persona chat datasets that have been constructed to develop models, as shown in Table 1. In particular, the introduction of the PersonaChat (Zhang et al., 2018; Dinan et al., 2019) dataset has extensively promoted the development of this field where the crowd-workers are simply asked to "chat with the other person naturally and try to get to know each other." However, the user’s persona was unknown

Dataset	Persona	Mutual	# Dialogues	Language	Multi-turn
PersonaChat (Zhang et al., 2018)	Text	✗	10,907	English	Yes
PersonalDialog (Zheng et al., 2019)	Structure	✗	20,830,000	Chinese	part
XPersona (Lin et al., 2020)	Text	✗	16,878	Multilingual	Yes
PEC (Zhong et al., 2020)	Text	✗	355,000	English	Yes
PCR (Mazaré et al., 2018)	Text	✗	700,000,000	English	Yes
MSC (Xu et al., 2021)	Text	✓	5,001	English	Yes
DuLeMon (Ours)	Text	✓	27,501	Chinese	Yes

Table 1: Comparison of our dataset DuLeMon with other datasets.

to the bot, so the dialogue was like strangers exchanging information. In contrast, our proposed DuLeMon dataset requires the chatbot to actively remember and use the user’s persona to improve conversational engagements and increase the intimacy between interlocutors in long-term interactions.

**Dialogue Model with External Memory:** As described in Lim (2012), there are various memory models used by the rule-based dialogue systems. In Bang et al. (2015), user-related information is memorized and used to rewrite the response. In Elvir et al. (2017), a unified episodic memory architecture for Embodied Conversational Agents (ECAs) is proposed. They describe a process that determines the prevalent contexts in the conversations obtained from the interactions. In Campos et al. (2018), the authors introduce an agent that uses its conversational memory to revisit shared history with users to maintain a coherent social relationship over time. However, they find it challenging to leverage the shared history with individual users and hard to accommodate expected conversational coordination patterns. Apart from studies in rule-based dialogue systems mentioned above, Xu et al. (2021) shows how large-scale pre-training generative dialogue models trained on existing datasets perform poorly in the long-term conversation setting and proposes a new extended English conversation dataset, entitled Multi-Session Chat (MSC). Different from them, our novel dataset DuLeMon does not rely on long sessions with high collection costs to study long-term memory problems in the persona chat, with significant differences in task design and data collection.

### 3 Data Collection

**Task Definition.** Given dialogue context  $c = \{u_1, s_1, u_2, s_2, \dots, u_{t-1}, s_{t-1}, u_t\}$ , where  $u$  and  $s$  represent the user and the chatbot respectively. Each speaker has its corresponding persona descrip-

tion that consists of a set of sentences, we define the user persona as  $\rho^u = \{\rho_1^u, \rho_2^u, \dots, \rho_m^u\}$ , and the chatbot persona as  $\rho^s = \{\rho_1^s, \rho_2^s, \dots, \rho_n^s\}$ . Given the dialogue context  $c$ , user persona  $\rho^u$  and chatbot persona  $\rho^s$ , we are interested in finding the corresponding persona and predicting the chatbot response  $s_t$ .

To support our task, we collect and release a new dataset, entitled DuLeMon. In DuLeMon, the chatbot actively remembers and reasonably uses what the user has said about their persona while maintaining consistency in its persona, allowing the conversation to proceed more deeply. In a nutshell, our DuLeMon dataset has two essential features: During the conversation, the chatbot can see the persona of both parties; the other is that the persona associated with the response is explicitly annotated in our dataset. Unlike the PersonaChat dataset, the setting in DuLeMon is that one speaker plays the role of a chatbot, and the other plays the user’s role. We elaborate on the construction process of the dataset as the following.

**(1) Persona collection:** The persona is mainly from the translation and rewriting of persona in PersonaChat. The chatbot’s persona is only visible to itself, and the chatbot can use its persona information to chat with the user, as shown in Figure 2. The user’s persona contains two parts: persona that the chatbot already knows and persona that the chatbot does not know. The first part is the user’s persona that the chatbot has learned through historical conversations. This part is randomly selected from multiple personas of each user. The chatbot needs to use this information to guide the conversation during the chat process. It should be noted that in order to simulate the situation at the beginning of the chat, this part may be empty.

**(2) Dialogue collection:** For each dialogue, two crowd-workers (one plays the chatbot, the other plays the user) are randomly paired and given random persona. They are required to organize a di-





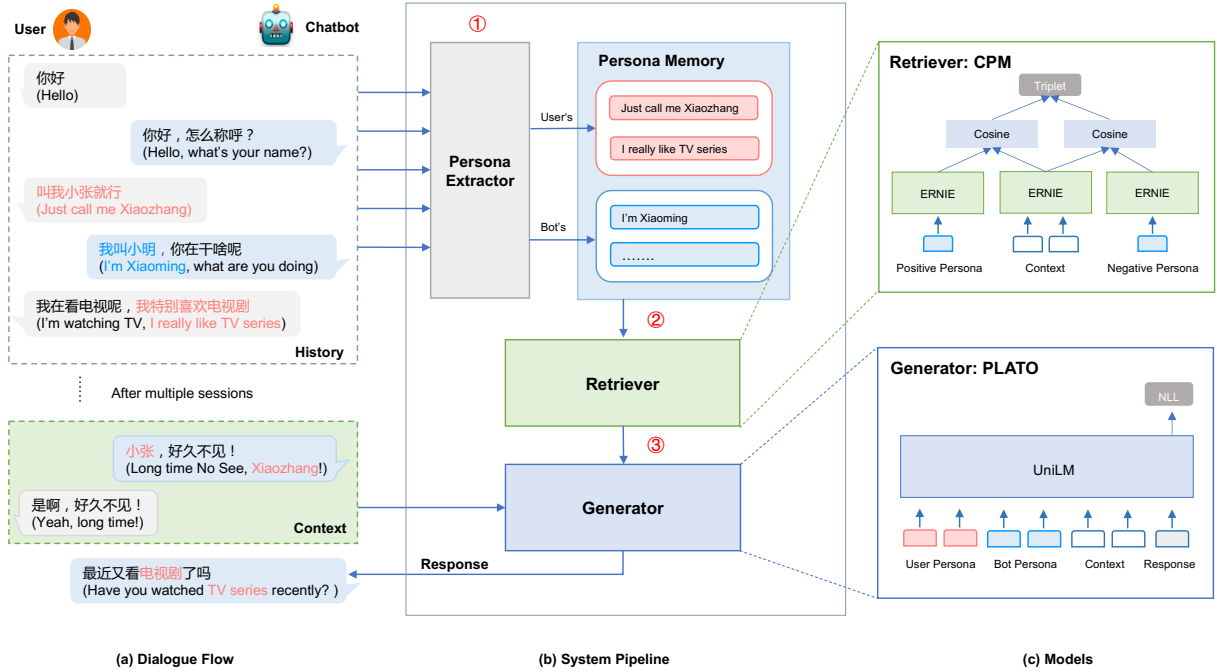


Figure 3: Illustration of our system PLATO-LTM. (a) shows the dialogue flow. (b) describes the modules and pipeline of our system. It consists of a persona extractor (PE), a long-term persona memory, a retriever, and a generator. ① The long-term memory contains both **user persona** and **chatbot persona** extracted from the dialogue history by PE. ② The retriever uses context as query to retrieve related personas in memory ③ concatenates the retrieved text to the context and use the generator to produce the generated response. (c) details our generator PLATO-2 and ranker CPM (Context Persona Matching).

for classification.

**Training procedure.** First, we collect the first-version training dataset, in which there are 6k utterances (from the DuLeMon corpus and Chinese social forum corpus) being human-annotated with positive or negative class labels. Second, using the aforementioned dataset, we train five ERNIE-CNN network (with different pre-training parameter versions) based models (called pc-stage1). Third, we employ these five models to automatically annotate 1.4 million utterances with labels, where these utterances are collected from the DuLeMon and the online Chinese social forum. We then refine this augmented dataset as the final-version dataset with the following steps: (a) Given an utterance, if there are at least two of the above five models identifying it as a positive sample, then it is attached with a positive label, (b) otherwise it is attached with a negative label. Finally, we train the five models on the final-version dataset and select the one with the best performance as our persona extractor (named pc-stage2).

**Inference procedure.** First, given an utterance, we segment it into clauses with the use of punctuation marks. Second, we use the persona extractor

mentioned above to classify each clause with a label and then collect the clause with a positive label as persona sentences.

## 4.2 Long-Term Memory

The long-term memory (LTM) module maintains memories to store the historical persona information from the user and the chatbot, respectively. The most critical operations are reading and writing based on the context persona matching (CPM) model. We use context encoder  $E_c(\cdot)$  to encode the current context  $c$ , and use persona encoder  $E_\rho(\cdot)$  to encode the persona  $\rho_i$ .  $E(\cdot)$  is the encoder's output on the first input token ( $[CLS]$ ), corresponding to the input's pooled representation.

The encoder  $E_c$  and  $E_\rho$  is initialized with the ERNIE model and then trained on our DuLeMon corpus. For each training sample, we define the positive persona as the persona used in the current user's utterance and the bot's response (including bot persona and user persona seen by bot), and the negative persona as the remaining persona of the current session. Given context  $c$ , a positive persona  $\rho^+$ , and a negative persona  $\rho^-$ , we use triplet loss

to tune the network as:

$$\max(\text{sim}(c, \rho^+) - \text{sim}(c, \rho^-) + \alpha, 0)$$

We set the margin  $\alpha = 0.2$  in our experiments. Below we describe the specific read and write process of the long-term memory module.

**Write:** We use the PE module to identify the persona in the dialogue history as the candidate information to be written. It needs to eliminate duplicates before writing. Specifically, calculate the cosine similarity with the persona in memory to get the most approximate persona  $\rho_j$ . When the similarity between  $\rho_i$  and  $\rho_j$  exceeds the given duplication threshold  $s_{dup}$ , replace  $\rho_j$  in memory with  $\rho_i$ ; otherwise, write  $\rho_i$  directly into the memory. When writing to memory, save  $\{\rho_i, E_\rho(\rho_i)\}$  pair for the subsequent reading. We measure the distance with the cosine similarity as:

$$\text{sim}(\rho_i, \rho_j) = \cos(E_\rho(\rho_i), E_\rho(\rho_j)) \quad (1)$$

**Read:** The reading process can be regarded as the retrieval process from memory. First, we use the efficient similarity search of dense vectors to select candidates. Then a matching model is utilized to score the relevance of the candidates to the current context. The similarity between the context and the persona using cosine similarity:

$$\text{sim}(c, \rho_i) = \cos(E_c(c), E_\rho(\rho_i)) \quad (2)$$

The top  $k$  persona candidates  $\rho^u$  in the user memory and top  $k$  candidates  $\rho^s$  in the chatbot memory are used for response generation. To model persona sparsity in dialogue, we filter out the persona, whose similarity score is lower than the similarity threshold  $s_c$ .

### 4.3 Generation Module

We trained our model on the basis of the PLATO-2 (Bao et al., 2020) architecture which adopts the generic transformer language model (Vaswani et al., 2017) and leverages a stack of masked multi-head self-attention layers to train on massive dialogue data<sup>3</sup>.

Given the conversation context  $c = \{u_1, s_1, u_2, s_2, \dots, u_{t-1}, s_{t-1}, u_t\}$ , the corresponding user persona  $\rho^u$  and chatbot persona  $\rho^s$ , the ground truth response as

<sup>3</sup>There are two stages within the PLATO-2 model, the first stage conduct candidate responses generation and the second stage conduct responses selection. We only implement our work on the first stage of PLATO-2.

$r = \{x_{m+1}, x_{m+2}, \dots, x_N\}$ , the conditional probability of  $p(r|c, \rho^u, \rho^s)$  can be written as the product of a series of conditional probabilities:

$$p(r|c, \rho^u, \rho^s) = \prod_t^N p(r_t|c, \rho^u, \rho^s, r_{<t}) \quad (3)$$

Therefore, we need to minimize the following negative log-likelihood (NLL) loss:

$$\begin{aligned} \mathcal{L}_{NLL} &= -\mathbb{E} \log p(r|c, \rho^u, \rho^s) \\ &= -\mathbb{E} \sum_{t=1}^T \log p(r_t|c, \rho^u, \rho^s, r_{<t}) \end{aligned} \quad (4)$$

where  $T$  is the length of the target response  $r$  and  $r_{<t}$  denotes previously generated words. Since the response generation is a uni-directional decoding process, each token in the response only attends to those before it. As for the context, bi-directional attention is enabled for better natural language understanding.

We added two strategies to distinguish different roles in the dialogue and prevent the confusing use of persona information.

- Role Embedding (Bao et al., 2021): different role embedding is used to distinguish the persona of different chat parties, abbreviated role\_embed.
- Role Token: splicing "system persona" before the chatbot persona and "user persona" before the user persona, abbreviated role\_token.

## 5 Experiments

In this section, we present the baselines, experiment settings, model comparisons, and results of experiments.

### 5.1 Compared Methods

As baselines, we select state-of-the-art methods to compare with our method.

- PLATO-2 (Bao et al., 2020): The SOTA open-domain dialogue model.
- PLATO-FT: The PLATO-2 model fine-tuned on our proposed DuLeMon dataset.
- PLATO-LTM: The PLATO-FT model with our proposed long-term memory (LTM).

- PLATO-LTM w/o PE: PLATO-LTM without the persona extractor (PE) module, which stores all history utterances (user and bot separately) into memory without persona extraction.

## 5.2 Experiment Settings

**Automatic Evaluation Metrics.** We use Precision, Recall and F1 to evaluate the persona classification model. For the long-term memory module, we use the AUC and recall@k to evaluate the ranking model. We evaluate responses generated by the models using PPL, BLEU (Papineni et al., 2002), and F1 with reference to the human-annotated responses and DISTINCT-1/2 (Zhao et al., 2017). More recently, Adiwardana et al. (2020) has shown the correlation between perplexity and human judgment in open-domain chat models.

**Human Evaluation Metrics.** In human evaluation, we employ three utterance-level metrics, including coherence, consistency, engagingness. Three crowd-sourcing workers are asked to score the response/dialogue quality on a scale of [0, 1, 2]. The higher score, the better. These criteria are discussed as follows:

- Coherence: an utterance-level metric, measuring whether the response is relevant and consistent with the context.
- Consistency: an utterance-level metric, evaluating whether the response is consistent with the persona in the dialogue history.
- Engagingness: an utterance-level metric, assessing whether the annotator would like to talk with the speaker for each response in the long-term conversation.

## 5.3 Results

In this part, we first analyze the effects of each module and then analyze the results of the manual evaluation of our entire system, PLATO-LTM.

### 5.3.1 Results of Persona Extractor

We measure the performance of the persona extractor. To measure the performance of different models, we manually annotated the test set (the number of test sets is 200). We select the best of the first and second stages. The result is shown in Table 3. The pc-stage2 model is better than that of the pc-stage1 model. The F1 of the model exceeds

Model	ACC	Precision	Recall	F1
pc-stage1	0.91	0.96	0.84	0.90
pc-stage2	0.92	0.95	0.87	0.91

Table 3: Comparison of two-stage models of our persona classifier.

0.9, which shows that our model can effectively recognize the persona information from the dialogue history and ensure that the persona information can be correctly stored in the long-term memory. Therefore, the pc-stage2 model is adopted in our system to recognize the persona in the dialogue history.

### 5.3.2 Selection of Generative Models

The generative model utilizes the current context and persona information retrieved from long-term memory to generate the response. We first evaluate the effect of the CPM model on retrieval persona information. The AUC on the automatic test set is 0.76, recall@5 is 0.83, which shows that our model can efficiently retrieve relevant persona from the long-term memory.

The effect of the generative model reflects the model’s ability to use the content of long-term memory to generate the response. Therefore, we select the best generative model to utilize better the retrieved persona information to generate. The result is shown in Table 4. We use the 12L model to conduct experiments to compare different models. The experiment results show that PLATO-FT + role\_embed + role\_token is the best. Compared to PLATO-FT, the PPL can decrease to 13.377, showing that both strategies are effective. In order to further improve the model, we increased the model size and further trained with the 32L model. Experiment results have shown that the PPL of the 32L model is lower than the 12L model by 4.4 and F1 increased by 2.5, which can further improve the generative model. Therefore, PLATO-FT 32L + role\_embed + role\_token model is adopted in our system.

### 5.3.3 Human Evaluation

Self-chat has been widely used in the evaluation of dialogue systems (Li et al., 2016b; Roller et al., 2021; Bao et al., 2020), where the model plays the roles of both parties in the dialogue. To better control variables, we use our proposed PLATO-LTM as a user simulator in our experiments and ask all chatbots (including PLATO-LTM) to chat sepa-

Model	PPL	BLUE-1/2	DISTINT-1/2	F1
PLATO-FT 12L	13.641	0.190/0.081	0.061/0.277	21.02
PLATO-FT 12L + role_embed	13.387	0.180/0.080	0.062/0.274	20.98
PLATO-FT 12L + role_token	13.553	0.193/0.081	0.060/0.272	21.28
PLATO-FT 12L + role_embed + role_token	<b>13.377</b>	<b>0.194/0.081</b>	0.060/0.267	<b>21.59</b>
PLATO-FT 32L + role_embed + role_token	9.380	0.194/0.087	0.068/0.296	22.61

Table 4: Comparison of automatic evaluation metric results among different generative models.

Model	Coherence	Consistency	Engagingness
PLATO-2	<b>1.70</b>	0.13	1.46
PLATO-FT	1.59	0.40	1.40
PLATO-LTM	1.67	<b>0.87</b>	<b>1.54</b>
PLATO-LTM w/o PE	1.57	0.49	1.43

Table 5: Comparison of human evaluation metric results on self-chat dialogues among our model and baselines. All the above generation models are 32L. The PLATO-FT is with role embedding and role token strategies.

rately with the user simulator. After that, the crowdsourcing workers evaluate only the responses generated by the chatbots other than the simulator. The details are as follows.

Each chatbot chats with the user simulator for 10 episodes, each containing 4 long sessions, and each session contains 16 rounds. As in Bao et al. (2020), we do not impose any restrictions on the chats except for specifying session openings. We pre-select some session openings from the DuLeMon test set, start the interactive conversation with these openings, and ask the two bots to perform chats given the context.

The results are shown in Table 5, from which we can get the following key results:

**(1) The long-Term Memory mechanism can significantly improve dialogue consistency.** As shown in Table 5, in terms of dialogue consistency, our two models, PLATO-LTM and PLATO-FT, can achieve scores of 0.87 and 0.40, respectively, which is significantly better than the baseline model PLATO-2. Furthermore, when we compare the performance of PLATO-LTM with PLATO-FT, it can be seen that the use of Long-Term Memory and persona extractor can boost the performance of PLATO-FT with a relative improvement of 118%. Moreover, the model of PLATO-LTM w/o PE can achieve a score of 0.49, which is still better than the PLATO-FT model. It indicates that long-term memory without a persona extractor is still effective in improving persona consistency.

**(2) With the long-term memory mechanism, the use of persona extractor can significantly improve persona consistency and dialogue engagingness.** As shown in Table 5, in terms of dia-

logue consistency, the two models, PLATO-LTM (using PE) and PLATO-LTM w/o PE, can achieve scores of 0.87 and 0.49 respectively, indicating that the use of persona extractor can significantly improve dialogue consistency. In terms of dialogue engagingness, PLATO-LTM can obtain a score of 1.54, outperforming the baseline model PLATO-2. In addition, when we remove PE from PLATO-LTM, its performance drops from 1.54 (the score of PLATO-LTM) to 1.43 (that of PLATO-LTM w/o PE), indicating that the use of persona extractor can improve the performance of PLATO-FT.

**(3) Fine-tuning on the small-scale dataset will slightly hurt the performance of pre-trained dialogue models in dialogue coherence.** In terms of dialogue coherence, the PLATO-FT model (fine-tuned on our dataset) achieve a score of 1.59, which is lower than that of the baseline model PLATO (not finetuned on our dataset). The possible reason is that during the self-play procedure for system evaluation, their dialogs usually cover a wide range of topics, and then it is challenging to generate appropriate or coherent responses when given these open-domain topics in contexts. The finetuning procedure might hurt the capability of the pre-trained dialogue model in terms of response appropriateness or dialogue coherence, leading to the inferior performance of PLATO-LTM and its variants.

## 6 Conclusion

In this paper, We present a novel LeMon (Long-term Memory Conversation) task and then build the corresponding dataset DuLeMon, introducing long-term persona modelling into large-scale generative



dialogue models. We further propose a Long-Term Memory (LTM) as a plug-in component of state-of-the-art large-scale generative dialogue models. LTM consists of user memory and chatbot memory, where the user memory is for understanding and memorizing persona information mentioned by the user, and the chatbot memory attempts to keep its persona information to be continuously updated over time. Experiment results show that our system PLATO-LTM can make effective use of both parties' persona information from dialogue history to enhance dialogue consistency and engagingness when conducting a long-term conversation. In the future, we will further study the possibility of using reinforcement learning with human feedback signals to help long-term conversation.

## 7 Ethical Considerations

We are sure that DuLeMon 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 DuLeMon and steps taken to ensure the potential problems with the quality of the dataset do not create additional risks.

## 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.
- Jeesoo Bang, Hyungjong Noh, Yonghee Kim, and Gary Geunbae Lee. 2015. [Example-based chat-oriented dialogue system with personalized long-term memory](#). In *2015 International Conference on Big Data and Smart Computing (BIGCOMP)*, pages 238–243.
- 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, Zhihua Wu, Zhen Guo, Hua Lu, Xinxian Huang, Xin Tian, Xinchao Xu, Yingzhan Lin, and Zhengyu Niu. 2021. [Plato-xl: Exploring the large-scale pre-training of dialogue generation](#).
- Joana Campos, James Kennedy, and Jill F. Lehman. 2018. Challenges in exploiting conversational memory in human-agent interaction. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, page 1649–1657, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander H. Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhunoye, Alan W. Black, Alexander I. Rudnicky, Jason Williams, Joelle Pineau, Mikhail S. Burtsev, and Jason Weston. 2019. [The second conversational intelligence challenge \(convai2\)](#). *CoRR*, abs/1902.00098.
- Miguel Elvir, Avelino J. Gonzalez, Christopher Walls, and Bryan Wilder. 2017. [Remembering a conversation – a conversational memory architecture for embodied conversational agents](#). *Journal of Intelligent Systems*, 26(1):1–21.
- Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. [Challenges in building intelligent open-domain dialog systems](#). *ACM Trans. Inf. Syst.*, 38(3).
- Yonghee Kim, Jeesoo Bang, Junhwi Choi, Seonghan Ryu, Sangjun Koo, and Gary Geunbae Lee. 2014. [Acquisition and use of long-term memory for personalized dialog systems](#). In *Multimodal Analyses enabling Artificial Agents in Human-Machine Interaction - Second International Workshop, MA3HMI 2014, Held in Conjunction with INTERSPEECH 2014, Singapore, Singapore, September 14, 2014, Revised Selected Papers*, volume 8757 of *Lecture Notes in Computer Science*, pages 78–87. Springer.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar; A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1746–1751. ACL.
- 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*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P. Spithourakis, Jianfeng Gao, and William B. Dolan. 2016a. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016b. [Deep reinforcement learning for dialogue generation](#).
- Mei Yii Lim. 2012. [Memory Models for Intelligent Social Companions](#), pages 241–262. Springer Berlin Heidelberg, Berlin, Heidelberg.

- Zhaojiang Lin, Zihan Liu, Genta Indra Winata, Samuel Cahyawijaya, Andrea Madotto, Yejin Bang, Etsuko Ishii, and Pascale Fung. 2020. [Xpersona: Evaluating multilingual personalized chatbot](#). *CoRR*, abs/2003.07568.
- 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.
- 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.
- 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, 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.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. 2019. [Ernie 2.0: A continual pre-training framework for language understanding](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Jing Xu, Arthur Szlam, and Jason Weston. 2021. [Beyond goldfish memory: Long-term open-domain conversation](#).
- 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.
- Weinan Zhang, Ting Liu, Yifa Wang, and Qingfu Zhu. 2017. [Neural personalized response generation as domain adaptation](#). *CoRR*, abs/1701.02073.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 654–664. Association for Computational Linguistics.
- Yinhe Zheng, Guanyi Chen, Minlie Huang, Song Liu, and Xuan Zhu. 2019. [Personalized dialogue generation with diversified traits](#). *CoRR*, abs/1901.09672.
- Peixiang Zhong, Chen Zhang, Hao Wang, Yong Liu, and Chunyan Miao. 2020. [Towards persona-based empathetic conversational models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6556–6566, Online. Association for Computational Linguistics.

## A Details of Data Collection

The collection processes of DuLeMon are as follows.

- The crowdworkers enter the chat interface in pairs, and role 1 initiates a conversation;
- The chat content can include opening greetings, self-introduction, chatting content that conforms to the persona information, asking the other party’s questions, answering the other’s questions, and so on. The information used in the chat must be consistent with the given personal information;
- The dialogue contains at least 8 turns (each person speaks at least 8 utterances);

At the same time, we also let the crowdworkers pay attention to the follows: 1. Use as many words as possible, and do not repeat them. The overall dialogue strives to be natural, smooth, and not embarrassing. 2. Do not simply copy and paste the sentences in the personal information and express them as richly as possible. If it is found that 50% of the fragments of any given sentence appear in the conversation, it is a non-compliant conversation. 3. When using persona information, do not copy it entirely, and talk about relevant content around the persona. For example, if the persona setting contains the sentence "I am a painter", the response can be that "I have painted many beautiful paintings and held several exhibitions"; 4. If the question raised by the other speaker is not covered in the given personal information, the reply can be freely used; if there is any reference or related information in the given personal information, reply according it.

## B Details of Models

**Generation Model** For the Generation model, We follow PLATO-2 (Bao et al., 2020). The maximum length of context, user persona, and chatbot persona are set to 384, 76, and 52, respectively. The vocabulary contains 30K Chinese BPE tokens. We optimize all models using Adam (Kingma and Ba, 2015) with every batch of  $B = 16384$  tokens and learning rate of  $lr = 5e - 5$ . We conduct all experiments on NVIDIA V100 32GB and A100 48GB GPUs.

**Long-term Memory** For both user memory and chatbot memory, we set duplication threshold  $s_{dup} = 0.95$ , number of candidates  $K = 5$ , and similarity threshold  $s_c = 0.7$ . Due to the persona sparsity of dialogue and the efficiency of our persona storage, we do not limit the memory capacity.

## C Cases of PLATO-LTM

To concretely demonstrate the long-term persona ability in a long-term conversation, we further provide a cherry-picked example of one episode conversation (between PLATO-LTM and PLATO-2) in Figure 4.



Figure 4: A cherry-picked example of one episode conversation between PLATO-LTM and PLATO-2.