

Learning to Memorize Entailment and Discourse Relations for Persona-Consistent Dialogues

Ruijun Chen¹, Jin Wang^{1*}, Liang-Chih Yu² and Xuejie Zhang¹

¹School of Information Science and Engineering, Yunnan University, Yunnan, China

²Department of Information Management, Yuan Ze University, Taiwan

chenrj@mail.ynu.edu.cn, wangjin@ynu.edu.cn, lcyu@saturn.yzu.edu.tw, xjzhang@ynu.edu.cn

Abstract

Maintaining engagement and consistency is particularly important in dialogue systems. Existing works have improved the performance of dialogue systems by intentionally learning interlocutor personas with sophisticated network structures. One issue with this approach is that it requires more personal corpora with annotations. Additionally, these models typically perform the next utterance prediction to generate a response but neglect the discourse coherence in the entire conversation. To address these issues, this study proposes a method of learning to memorize entailment and discourse relations for persona-consistent dialogue tasks. Entailment text pairs in natural language inference dataset were applied to learn latent entailment relations as external memories by premise-to-hypothesis generation task. Furthermore, an internal memory with a similar architecture was applied to the discourse information in the dialogue. Placing orthogonality restrictions on these two memory spaces ensures that the latent entailment relations remain dialogue-independent. Both memories collaborate to obtain entailment and discourse representation for the generation, allowing a deeper understanding of both consistency and coherence. Experiments on two large public datasets, PersonaChat and DSTC7-AVSD, demonstrated the effectiveness of the proposed method. Both automatic and human evaluations indicate that the proposed model outperforms several strong baselines in terms of both persona consistency and response coherence. Our source code is available at <https://github.com/Chenrj233/LMEDR>.

Introduction

Traditional chat models lack specificity and personality consistency. Only when they access a sufficiently large dataset will they have the opportunity to generate piecemeal and uninformative responses in a chat setting. For two consecutive questions with similar meanings in a two-round dialogue, that is, *what is your job* and *what do you do*, the model replied to the former: *I am a lawyer*, while the latter: *I am a doctor* (Welleck et al. 2020). This issue arises because of the lack of a consistent personality as well as an explicit memory towards plausibility as they are typically trained to produce a response given only recent dialogue history (Shum, He, and Li 2018).

*Corresponding author

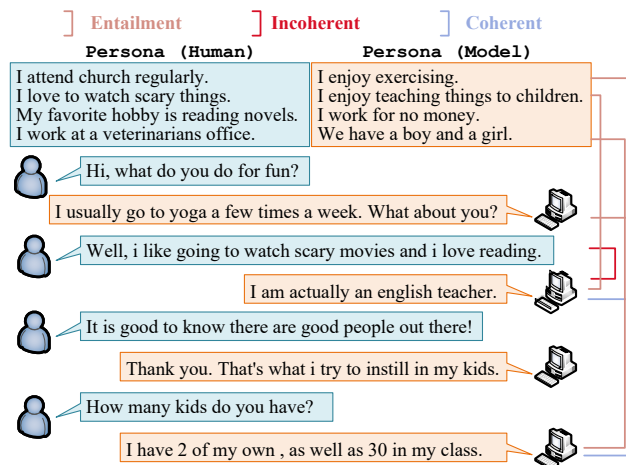


Figure 1: The conceptual diagram of introducing natural language inference in persona-based dialogue.

One solution that maintains consistency in a dialogue system is to provide a set of persona profiles that describe the character and then generate responses according to the persona. **Persona** can be defined as the composition of identity elements, such as profiles and background personal facts.

The expected outcome is that dialogue models generate a response consistent with the given persona. The PersonaChat dataset (Zhang et al. 2018), widely adopted to support the training of persona-consistent dialogues, was manually annotated by two annotators to act as part of a predefined persona and chat naturally to know each other during the conversation. However, given the time and effort needed to annotate more persona corpora to cover all possibilities, it is difficult to extend the application of such persona-related information to the daily usage of dialogue.

As humans, our knowledge of the concepts and the semantic relationship behind the language can allow us to rearrange unstructured data so that we can understand and analyze it. Essentially, we can robustly learn novel concepts with minimal supervision, benefitting from the well-known ability of natural language inference (NLI). Figure 1 shows an example of introducing NLI in a persona-based dialogue. Given a persona as a *premise*, we can determine whether

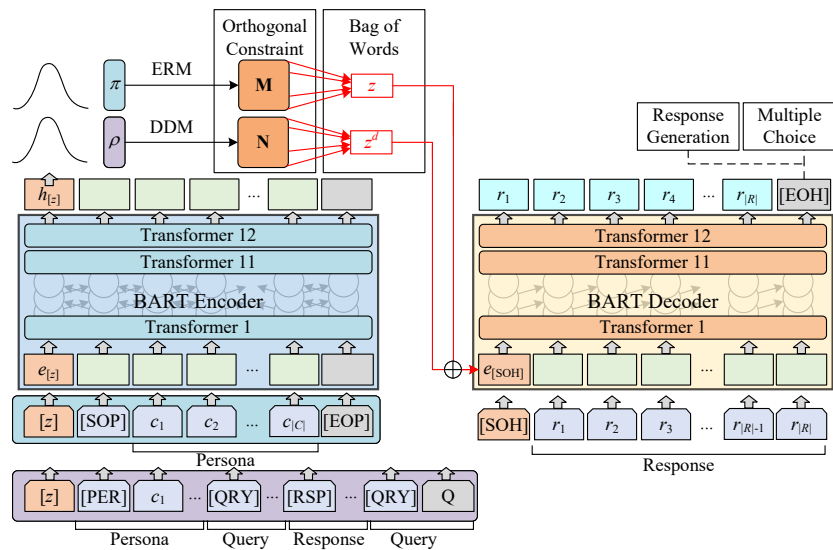


Figure 2: Overall architecture of the proposed method for persona-consistent dialogue generation.

the *hypothesis* of the response utterance is true (entailment), false (contradiction), or undetermined (neutral).

Recent studies have sought to improve the consistency of the dialogue system by modeling the understanding between interlocutors (Liu et al. 2020). Song et al. (2021) disentangled persona-based dialogue generation into two subtasks—response generation and consistency understanding—and used unlikelihood training to make the decoder generate contradictory dialogue responses as few as possible. However, multiple subtasks require multiple encoders, leading to a complex generation model structure. Nie et al. (2021) introduced a contradiction detection task to evaluate the consistency in dialogues.

Despite continuing efforts to improve the engagement and consistency of dialogue systems, understanding persona-response consistency is still difficult. The key challenges are twofold: 1) Existing methods apply sophisticated structures to learn persona consistency, which requires more annotated corpora for training. However, persona-based corpora are still insufficient and difficult to collect. 2) Dialogue-generating models typically neglect discourse information. Discourse coherence is a crucial component of the effectiveness of a conversation, encompassing how utterances are connected and how the entire dialogue is organized to convey information to the interlocutor. Existing models usually perform the next utterance prediction for response generation but ignore the dialogue discourse coherence. As indicated in Figure 1, *I am actually an English teacher* seems to be an appropriate and persona-consistent response to the query. However, this response is incoherent in the context of an entire conversation.

To address these issues, this study proposes a method of learning to memorize entailment and discourse relations for persona-consistent dialogue tasks. We applied an encoder-decoder architecture from BART (Lewis et al. 2020). To explicitly understand the consistency of personas, we designed

an external memory to store the latent entailment relations between premises and the entailment hypothesis, independent of dialogue itself. In addition, discourse relations were learned and stored in internal latent memory. The latent entailment relations are ensured to be dialogue-independent by imposing orthogonality constraints on the two memory spaces. Given personas and dialogue queries, both memories work jointly to obtain the entailment and discourse representation by the BART encoder. The generation was finally accomplished by the BART decoder with two extra training objectives, which further acquired the ability to understand both consistency and coherence.

Comparative experiments were conducted using the PersonaChat (Dinan et al. 2020) and DSTC7-AVSD (Alamri et al. 2019). Both automatic and human evaluations show that the proposed method generalizes well under different settings and outperforms several strong baselines on most metrics, especially persona consistency, indicating that the proposed method can produce better persona-consistent dialogue responses.

The remainder of this paper is organized as follows. Section 2 provides a brief review of the related work. Section 3 describes the proposed model, which learns to memorize entailment and discourse relations by using latent variables. Section 4 summarizes the specific experimental setup for the two public dialogue datasets and the corresponding analysis of the results. Finally, conclusions are drawn in Section 5.

Related Work

Persona-based Dialogues

Generation-based dialogue systems usually use the sequence-to-sequence (seq2seq) model (Sutskever, Vinyals, and Le 2014) as the backbone. After the persona is introduced into the dialogue, it is necessary to adopt an effective method to integrate role information into the dialogue, such as persona embedding (Li et al. 2016b). Subsequently, with

the development of large-scale pre-trained language models, an increasing number of methods (Wolf et al. 2019; Roller et al. 2021; Lin et al. 2021; Zheng et al. 2020) have leveraged pre-training and fine-tuning to improve persona-based dialogue, but the problem of dialogue consistency remains unsolved. Therefore, Liu et al. (2020) have attempted to model the understanding between interlocutors to improve the consistency of dialogue systems. A new perspective (Song et al. 2021) decomposes persona-based dialogue tasks into consistent understanding and dialogue generation significantly improves dialogue-consistent generation based on natural language inference.

Latent Modeling

In a dialogue scene, the factors that associate dialogue context with dialogue responses are often difficult to observe and explain; therefore, modeling the latent space of dialogue can help improve the performance of dialogue generation. Optimus (Li et al. 2020) combines the advantages of BERT (Devlin et al. 2019) and GPT-2 (Radford et al. 2020) for large-scale pre-training in the form of VAE (Kingma and Welling 2014) to model the latent variable space. PLATO (Bao et al. 2020) introduces discrete latent variables to solve the one-to-many relationship in response generation. DialogVED (Chen et al. 2022) introduces continuous latent variables into an enhanced encoder-decoder pre-training framework to improve the relevance and diversity of dialogue responses. All these methods show great promise for modeling dialogue-related features in latent space. This paper extends the idea by additionally memorizing NLI relations as latent dialogue-independent features.

Learning to Memorize for Persona-consistent Dialogue

The task of dialogue generation can be defined as the next utterance prediction, where a target response utterance $\mathcal{R} = [r_1, r_2, \dots, r_{|\mathcal{R}|}]$ is predicted given a conversation query $\mathcal{Q} = [q_1, q_2, \dots, q_{|\mathcal{Q}|}]$ according to given persona constraints $\mathcal{C} = [c_1, c_2, \dots, c_{|\mathcal{C}|}]$. For convenience, the sentences $(\mathcal{R}, \mathcal{Q}, \mathcal{C})$ are mapped to the vector representation $x = \{R, Q, C\}$. Further, natural language inference data (Welleck et al. 2020; Williams, Nangia, and Bowman 2018) $\mathcal{N} = \{P^{(n)}, H^{(n)}\}_{n=1}^N$, which consists of the entailed text pairs of premise and hypothesis, was used to learn the entailment relation to preserve consistency in dialogue generation.

Figure 2 shows the overall architecture of the proposed learning to memorize the entailment and discourse relations model for a persona-consistent dialogue. The backbone model is based on BART (Lewis et al. 2020), which performs repeated two-stage training, i.e., learning to memorize and persona-consistent dialogue generation. The key insight of the proposed model is that it maps both the entailment relation and discourse information to latent spaces. Based on this information, an external memory module enforces premise-to-hypothesis generation to map the textual entailed pair to the *Dialogue-Independent* latent space, which can be memorized and stored in a memory structure \mathbf{M} . Similarly, the discourse information was mapped using an inter-

nal memory module \mathbf{N} to learn the *Dialogue-Related* features. For generation, both entailment and discourse representation can be obtained from memory and enhance persona consistency in dialogue generation with additional entailment and discourse information.

Learning to Memorize

Entailment Relation Memory (ERM). ERM is an external memory which is used to learn and store entailment relations for persona consistency. If a given hypothesis H can be inferred from the premise P , the relationship of the pair is entailment. For persona-based dialogue, such an entailment relationship can be introduced to generate consistent responses.

Given a dataset of textual entailed pairs $\mathcal{N} = \{P^{(n)}, H^{(n)}\}_{n=1}^N$, textual entailment generation was adopted to learn a latent variable z , which represents the latent form of entailment relations in natural language inference, defined as

$$p(H, z | P) = p(z | P)p(H | z, P) \quad (1)$$

Based on the BART encoder, we introduce a special latent token $[z]$, a start-of-premise token [SOP], and an end-of-premise token [EOP] to the premise for latent entailment relation learning. By using the tokenizer and adding position embeddings, the input of the premise is transformed as

$$\mathbf{E}_{\text{ERM}} = [e_{[z]}, e_{[\text{SOP}]}, e_{p_1}, e_{p_2}, \dots, e_{p_{|P|}}, e_{[\text{EOP}]}] \quad (2)$$

We introduce a latent entailment relation memory structure \mathbf{M} parameterized by θ , as shown in Figure 3, where each element represents a certain latent factor, defined as

$$\mathbf{M} = [M_1, \dots, M_k] \in \mathbb{R}^{k \times d} \quad (3)$$

where k is the number of latent factors in entailment relations, and d is the dimension of the memory element. The hidden state of the last layer of the BART encoder, that is, $h_{[z]}$ corresponding to $e_{[z]}$, was applied to learn the distribution of the latent entailment relations $z \sim p(z | P)$ by

$$\pi = \text{softmax}(W_{\pi} h_{[z]} + b_{\pi}) \quad (4)$$

where π represents the probability of each element in \mathbf{M} . Then, the latent entailment representation z can be easily obtained from \mathbf{M} :

$$z = \sum_{i=1}^k \pi_i M_i \quad (5)$$

To memorize the latent entailment relations, we use entailment representation z from the memory \mathbf{M} with the weights π , along with the premise to generate the corresponding hypothesis. The obtained entailment representation z was added to the special start-of-hypothesis token [SOH] of the decoder, denote as

$$\hat{e}_{[\text{SOH}]} = e_{[\text{SOH}]} + z \quad (6)$$

Latent memory can keep track of the entailment relation with the representation of the source premise by both reading and writing during generation. Notably, it can be updated by backpropagation of the premise-to-hypothesis generation.

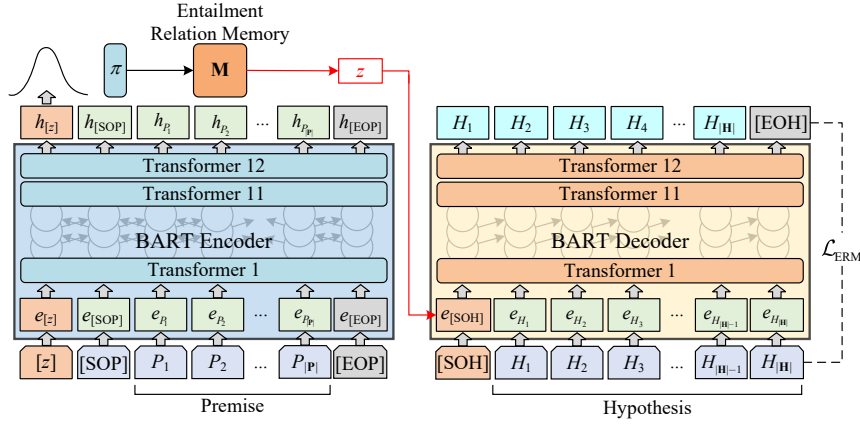


Figure 3: Learning to memorize the entailment relations in latent variables.

Algorithm 1: Latent memory learning

Input: A set of entailed text pair $\mathcal{N} = \{P^{(n)}, H^{(n)}\}_{n=1}^N$ and persona dialogue dataset $\mathbf{x} = \{R, Q, C\}$.

Parameter: Memory \mathbf{M} , \mathbf{N} parameterized by θ , ϕ , a pre-trained BART parameterized by φ .

- 1: **repeat**
 - 2: **Stage 1:**
 - 3: Train the BART model and memory \mathbf{M} with input \mathcal{N} .
 - 4: Minimizing Eq. (7) and optimize θ , φ .
 - 5: **Stage 2:**
 - 6: Fixed parameter θ .
 - 7: Train the BART model and memory \mathbf{N} with input x .
 - 8: Minimizing Eq. (19) and optimize ϕ , φ .
 - 9: **until** convergence
-

The objective of the pre-training is to optimize memory \mathbf{M} and model parameters φ by minimizing the language modeling loss:

$$\begin{aligned} \mathcal{L}_{ERM} &= -\mathbb{E}_{z \sim p_{\theta, \varphi}(z|P)} \log p_{\theta, \varphi}(H|z, P) \\ &= -\mathbb{E}_{z \sim p_{\theta, \varphi}(z|P)} \sum_{t=1}^{|H|} \log p_{\theta, \varphi}(H_t | z, P, H_{<t}) \end{aligned} \quad (7)$$

Dialogue Discourse Memory (DDM). Discourse coherence is an important aspect of dialogue text quality. It encompasses how utterances are connected, as well as how the entire dialogue is organized to convey information to the interlocutor (Bao et al. 2020). Similarly, we memorized the discourse information into an internal memory \mathbf{N} parameterized by ϕ . Let $\mathcal{C} = [c_1, c_2, \dots, c_{|C|}]$ denote the persona of the agent, $Q^{(1)}, \dots, Q^{(m)}$ represent m dialogue queries, and $\mathcal{R}^{(1)}, \dots, \mathcal{R}^{(m)}$ represent the target responses. Here, Q and \mathcal{R} are consecutive context-response pairs from the same dialogue session, which are used to capture the correct discourse relation in a dialogue. The input of the BART encoder is the concatenation of the persona and dialogue content, de-

noted as

$$\mathbf{E}_{DDM} = [e_{[z]}, e_{[PER]}, \mathcal{C}, e_{[QRY]}, Q^{(1)}, e_{[RSP]}, \mathcal{R}^{(1)}, \dots, e_{[QRY]}, Q^{(m)}] \quad (8)$$

where [QRY] and [RSP] are two special tokens that indicate the beginning of the query and the response, respectively.

Latent memory $\mathbf{N} \in \mathbb{R}^{l \times d}$ was introduced, where l represents the kind of latent dialogue discourse information. It learns the distribution of the latent dialogue discourse using

$$\rho = \text{softmax}(W_{\rho} h_{[z]} + b_{\rho}) \quad (9)$$

Then, the latent dialogue discourse representation was calculated by

$$z^d = \sum_{j=1}^l \rho_j N_j \quad (10)$$

In addition, latent entailment relations should be independent of the dialogue context. Therefore, we propose imposing an orthogonal constraint on these two memory spaces to reduce the correlation between the different latent spaces. The orthogonal constraint can encourage the latent memory to learn more features and reduce redundant features. To accomplish this goal, we used cosine similarity to represent the relationship between them, described as

$$\cos(\mathbf{M}, \mathbf{N}) = \frac{\mathbf{M}\mathbf{N}^{\top}}{\|\mathbf{M}\|_2 \|\mathbf{N}\|_2} \quad (11)$$

where $\|\cdot\|_2$ represents L_2 normalization. Once the vectors of different latent-memory spaces are orthogonal to each other, the correlation between them is the lowest. Thus, the training objective of the DDM can be defined to minimize the cosine similarity between \mathbf{M} and \mathbf{N} :

$$\mathcal{L}_{DDM} = \sum_{i \leq k, j \leq l} \left(\frac{M_i N_j^{\top}}{\|M_i\|_2 \|N_j\|_2} \right)^2 \quad (12)$$

Persona-consistent Responses Generation

For generation, the model leverages latent memories of both entailment relations and dialogue discourse to preserve consistency for generating a persona-based response, denoted

Model	Original			Revised		
	Hits@1	PPL	F1	Hits@1	PPL	F1
KV Profile Memory	54.8	-	14.25	38.1	-	13.65
Generative Profile Memory	10.2	35.01	16.29	9.9	34.94	15.71
LIC	17.3	-	17.79	16.2	-	16.83
Transfertransfo	82.1	17.51	19.09	-	-	-
P ² BOT	81.9	15.12	19.77	68.6	18.89	19.08
Our	89.5	10.99	21.99	85.0	13.42	19.54

Table 1: Automatic evaluation results of different methods for persona-based dialogue generation.

as

$$p(R, z, z^d | Q, C) = p(R | z, z^d, Q, C) p(z | C) p(z^d | Q, C) \quad (13)$$

We applied the persona as a premise and fed it into the model and obtained the latent entailment relations representation z under the condition of persona from latent entailment memory \mathbf{M} . The latent entailment relation representation z and the latent dialogue discourse representation z^d are then added to the special start-of-hypothesis token $e_{[\text{SOH}]}$ of the decoder input, denoted as

$$\hat{e}_{[\text{SOH}]} = e_{[\text{SOH}]} + z + z^d \quad (14)$$

Additionally, we use a bag-of-words loss (Zhao, Zhao, and Eskenazi 2017) to facilitate the latent variable,

$$\begin{aligned} \mathcal{L}_{\text{BOW}} &= -\mathbb{E}_{z \sim p_\varphi(z|C), z^d \sim p_{\phi, \varphi}(z^d|C, Q)} \sum_{t=1}^{|R|} \log p_{\phi, \varphi}(R_t | C, Q, z, z^d) \\ &= -\mathbb{E}_{z \sim p_\varphi(z|C), z^d \sim p_{\phi, \varphi}(z^d|C, Q)} \sum_{t=1}^{|R|} \log \frac{e^{f(R_t)}}{\sum_{v \in V} e^{f(v)}} \end{aligned} \quad (15)$$

where V is the whole vocabulary and $f(v)$ represents the predicted probability of token v appearing in the target response.

For response generation, we also trained the parameters of the model using a language modeling loss function:

$$\begin{aligned} \mathcal{L}_{\text{LM}} &= -\mathbb{E}_{z \sim p_\varphi(z|C), z^d \sim p_{\phi, \varphi}(z^d|C, Q)} \log p_{\phi, \varphi}(R | C, Q, z, z^d) \\ &= -\mathbb{E}_{z \sim p_\varphi(z|C), z^d \sim p_{\phi, \varphi}(z^d|C, Q)} \sum_{t=1}^{|R|} \log p_{\phi, \varphi}(R_t | R_{<t}, C, Q, z, z^d) \end{aligned} \quad (16)$$

Similar to (Wolf et al. 2019), we randomly sampled t interference responses to train the model to select the correct response, i.e., as a multiple-choice task. Specifically, we use the hidden state of the last token output by the decoder to predict the scores \hat{y}_i of each candidate response and calculate the cross-entropy loss with the ground-truth label y_i , denoted as

$$\hat{y} = \text{softmax}(W_h h_{eos} + b_h) \quad (17)$$

$$\mathcal{L}_{\text{CLS}} = -\sum_{i=1}^{t+1} \hat{y}_i \log(y_i) \quad (18)$$

In summary, the objective of our model is to minimize the loss:

$$\mathcal{L}(\phi, \varphi) = \mathcal{L}_{\text{DDM}} + \mathcal{L}_{\text{BOW}} + \mathcal{L}_{\text{LM}} + \mathcal{L}_{\text{CLS}} \quad (19)$$

Model	Hits@1	PPL	F1
PE-Trans	89.4	-	-
ImageS2S	-	11.19	21.30
BART	86.9	11.85	20.72
Our	89.5(90.1)	10.99 (11.00)	21.99 (21.96)

Table 2: Comparison with pre-trained language models fine-tuned on PersonaChat original mode.

Model	Dist-1	Dist-2	C.Score
LIC	2.31	10.71	19.13
P ² BOT	1.87	10.08	23.84
BOB	2.59	13.90	22.45
Our	2.47	13.82	25.31

Table 3: Automatic evaluation results of different methods for persona consistency.

While training persona-consistent responses generation, we fixed the parameters of \mathbf{M} ; the specific process is described in Algorithm 1.

Experiments

Dataset

Dialogue Dataset. We conducted experiments on two publicly available dialogue datasets to evaluate the performance of the proposed method.

- **ConvAI2 PersonaChat** (Dinan et al. 2020) is a chit-chat dataset based on PersonaChat (Zhang et al. 2018). It provides the personas of the interlocutor and is designed to facilitate dialogue models to generate more consistent and engaging responses.
- **DSTC7-AVSD** (Alamri et al. 2019) provides a conversational question-answering dataset, which is similar to that of PersonaChat. We used this dataset to explore the performance of our proposed method on a contextual knowledge-based dialogue task. The system must generate answers based on the given context and background knowledge of the conversation. We used only text information for the experiments.

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr
CMU	0.718	0.584	0.478	0.394	0.267	0.563	1.094
PLATO	0.784	0.637	0.525	0.435	0.286	0.596	1.209
ProphetNet	0.824	0.691	0.582	0.487	0.313	0.635	1.382
DialogVED	0.822	0.692	0.582	0.489	0.312	0.636	1.391
Our	0.801	0.680	0.576	0.488	0.316	0.631	1.403

Table 4: Automatic evaluation results of different methods for DSTC7-AVSD.

Model	Fluency	Consistency	Avg
LIC	3.27	2.11	2.69
P ² BOT	3.51	2.20	2.86
Our	3.57	2.31	2.94

Table 5: Human evaluation results.

	Hits@1	F1	BLEU-4	C.Score
BART	86.9	20.72	0.01289	21.32
w/o ERM	89.3	21.70	0.01597	22.09
w/o DDM	88.4	21.72	0.01406	24.66
w/o OC	88.8	21.84	0.01485	25.01
Our	89.5	21.99	0.01561	25.31

Table 6: Analysis of ablation experiments on the PersonaChat original mode.

NLI Dataset. Additionally, two natural language inference datasets were used to learn the corresponding ERM.

- **DNLI** (Welleck et al. 2020) is a dialogue inference dataset based on PersonaChat. The dataset consisted of sentence pairs labeled as entailment, neutral, or contradictory.
- **MNLI** (Williams, Nangia, and Bowman 2018) is a multi-genre natural language inference corpus and one of the largest corpora available for recognizing textual entailments.

Compared Methods

Baseline. Both **KV profile memory** and **generative profile memory** (Zhang et al. 2018) are the official baselines for PersonaChat, whereas **LIC** (Golovanov et al. 2020) was the best performing model in ConvAI2 for human evaluation, and **Transfertransfo** (Wolf et al. 2019) achieved the best performance on automatic metrics in ConvAI2. For DSTC7-AVSD, **CMU** (Sanabria, Palaskar, and Metze 2018) achieved the best performance for all evaluation metrics in the competition.

State of the art. **P²BOT** (Liu et al. 2020) models the understanding between the interlocutors and obtains a new state-of-the-art on ConvAI2. **Poly-Encoder Transformer** (PE-Trans) (Humeau et al. 2019) was pre-trained on Reddit and fine-tuned on ConvAI2 to obtain the best performance in the response selection task. **Imagseq2Seq Dodecadi-**

alogue (ImageS2S) (Shuster et al. 2020) was trained on 12 tasks and fine-tuned on the ConvAI2 generation task. **BOB** (Song et al. 2021) disentangles persona-based dialogue generation into consistent understanding and dialogue generation tasks to ensure that the model has a better consistent understanding. For DSTC7-AVSD, **PLATO** (Bao et al. 2020) was the first large-scale pre-trained dialogue language model that introduced a discrete variable for one-to-many relationship modeling. **ProphetNets** (Qi et al. 2020) proposed a pre-training objective for predicting multiple future tokens to enhance the performance of pre-trained language models on natural language generation tasks. **DialogVED** (Chen et al. 2022) is a pre-trained model framework to enhance the encoder-decoder by introducing continuous latent variables and obtaining the state-of-the-art on the DSTC7-AVSD.

Implementation Details

The proposed model was initialized using BART-large. AdamW (Loshchilov and Hutter 2019) was applied to optimize the model, with an initial learning rate of $8e-6$. We used **DNLI** on **PersonaChat** and **MNLI** on **DSTC7-AVSD** for the ERM learning. The batch size was 64 for training stage 1, and we used a batch size of two with a gradient accumulation of eight for training stage 2. The types of ERM and DDM were both set to 10 for PersonaChat and set to 20 and 5 for DSTC7-AVSD. The proposed model trained on one NVIDIA RTX 3090 with PyTorch framework. For dialogue generation, we used a beam search, and the maximum sequence length was set to 50.

Automatic Evaluation

Following the official automatic evaluation, we used **Hits@1**, **Perplexity** (PPL), and **F1** for automatic evaluation on the **PersonaChat** dataset. **Hits@1** is the probability that the golden response ranks the highest among the 20 candidate responses. **Perplexity** was used to calculate the negative log-likelihood of the golden response from the model. **F1** is calculated from the precision and recall at the word level between the predicted and golden responses. For persona consistency, we apply **Consistency Score** (C.Score) (Madotto et al. 2020), which leverages a referee model to predict consistency between response and persona. **Dist-1/2** (Li et al. 2016a) is used to measure response diversity.

The **BLEU**, **METEOR**, **ROUGE-L**, and **CIDEr** metrics were reported for **DSTC7-AVSD** automatic evaluation, similar to DSTC7 reviews (Alamri et al. 2019).

Table 1 presents the automatic evaluation results of the different methods for persona-based dialogue generation. As

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr
BART	0.791	0.668	0.562	0.473	0.307	0.619	1.372
w/o ERM	0.795	0.671	0.565	0.477	0.313	0.625	1.395
w/o DDM	0.799	0.676	0.571	0.482	0.309	0.625	1.388
w/o OC	0.782	0.660	0.557	0.473	0.309	0.620	1.388
Our	0.801	0.680	0.576	0.488	0.316	0.631	1.403

Table 7: Analysis of ablation experiments on the DSTC7-AVSD.

indicated, the proposed model outperformed all the baselines on the PersonaChat dataset, especially on the revised dataset. The proposed model achieved significant improvements in response selection tasks because ERM can provide more entailed information about the persona while DDM can capture the appropriate discourse information in the dialogue, making it easier to distinguish the correct response from the candidate responses. The improvement in PPL and F1 also shows that ERM and DDM can further improve the consistency of the persona and the quality of the response.

Table 2 compares the results of the proposed method with some pre-trained language models fine-tuned on the PersonaChat. The results of training with 14 interference responses are in parentheses. The proposed model also achieves better results compared to models trained on larger corpora and more tasks. Additionally, adding more interference responses can improve the accuracy of the model to select the correct response. Table 3 shows the result of persona consistency. The proposed model still obtained the highest C.Score, which indicates that the responses generated by the proposed model perform the best on persona consistency.

Table 4 shows the experimental results of DSTC7-AVSD. The proposed model achieves the best results for METROR and CIDEr and is close to the best model for the other metrics. Better performance on CIDEr shows that in background-based dialogue question answering, the proposed model can use ERM to capture key information in the background and combine the dialogue history to generate high-quality responses.

Human Evaluation

We conducted a human evaluation of the state-of-the-art models (LIC, P²BOT) and the proposed method. We randomly sampled 100 responses generated by these models from the original PersonaChat dev set for the human evaluation. Both **fluency** and **consistency** were applied as criteria. Four human annotators were asked to rate **fluency** on a scale of 1 to 5 and **consistency** on a scale of 1 to 3, where the Fleiss’s kappa of fluency and consistency are 0.578 and 0.671. Here, **fluency** indicates the smoothness of responses and conversation, coherence is included, where 1 means terrible and 5 represents very satisfying. **Consistency** represents the consistency between persona and response, which reflects whether the model can maintain persona consistency, where 1 means it does not match persona, 2 means it is irrelevant, and 3 means it is consistent with persona. As shown in Table 5, the results of the human evaluation are consistent with the automatic evaluation, and our model

outperforms the previous best-performing model on human evaluation in terms of both fluency and consistency. Several examples of the generated responses are provided in the Appendix to help illustrate the effectiveness of our model.

Ablation Study

We conducted ablation experiments on the PersonaChat and DSTC7-AVSD to explore the impact of each module. The ablation results are presented in Tables 6 and 7, respectively.

Effect of ERM. After removing the ERM, the C.Score becomes lower on PersonaChat, which shows that the ERM can make the response generated by the model more consistent with its persona and improve the persona consistency in the dialogue response. In dialogue question answering, ERM can capture the key information in the background so that the generated answers can achieve better results for CIDEr.

Effect of DDM. The large decline in BLEU-4 indicates that the quality of the responses generated by the model deteriorates without DDM. The metric of METEOR has a high correlation with the results of human judgment. Without DDM, METEOR drops to a large extent, validating the important role that DDM plays in response generation and DDM can effectively capture the connections of the utterances in the dialogue, thereby making the responses more coherent and natural.

Effect of the Orthogonal Constraint (OC). Imposing orthogonality constraints on the two latent spaces effectively reduces redundant features and makes the features captured by the model easier to distinguish, thereby improving the quality of generation and performance in response selection.

Conclusion

In this paper, we propose a dialogue generation method for learning to memorize entailment and discourse relations with latent variables. Combining latent entailment relations and dialogue discourse relations makes generated responses more coherent and consistent. Experiments on the PersonaChat dataset demonstrate the effectiveness of the proposed method. The results on the DSTC7-AVSD dataset also show that learning entailment and discourse relations are beneficial for dialogue question-answering generation.

Future works will attempt to explore different latent relations in text pairs on different datasets and combine the discourse relations in the dialogue to make the dialogue generation in the desired direction.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant Nos. 61966038 and 62266051, the Ministry of Science and Technology, Taiwan, ROC, under Grant No. MOST 111-2628-E-155-001-MY2 and the Postgraduate Research and Innovation Foundation of Yunnan University under Grant No.2021Z076. The authors would like to thank the anonymous reviewers for their constructive comments.

References

- Alamri, H.; Cartillier, V.; Das, A.; Wang, J.; Cherian, A.; Essa, I.; Batra, D.; Marks, T. K.; Hori, C.; Anderson, P.; Lee, S.; and Parikh, D. 2019. Audio visual scene-aware dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, (CVPR-2019)*, 7558–7567.
- Bao, S.; He, H.; Wang, F.; Wu, H.; and Wang, H. 2020. PLATO: Pre-trained Dialogue Generation Model with Discrete Latent Variable. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, (ACL-2020)*, 85–96.
- Chen, W.; Gong, Y.; Wang, S.; Yao, B.; Qi, W.; Wei, Z.; Hu, X.; Zhou, B.; Mao, Y.; Chen, W.; Cheng, B.; and Duan, N. 2022. DialogVED: A Pre-trained Latent Variable Encoder-Decoder Model for Dialog Response Generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, (ACL-2022)*, 4852–4864.
- Devlin, J.; Chang, M. W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-HLT 2019)*, 4171–4186.
- Dinan, E.; Logacheva, V.; Malykh, V.; Miller, A.; Shuster, K.; Urbanek, J.; Kiela, D.; Szlam, A.; Serban, I.; Lowe, R.; Prabhumoye, S.; Black, A. W.; Rudnicky, A.; Williams, J.; Pineau, J.; Burtsev, M.; and Weston, J. 2020. The Second Conversational Intelligence Challenge (ConvAI2). *arXiv preprint arXiv:1902.00098*.
- Golovanov, S.; Kurbanov, R.; Nikolenko, S.; Truskovskiy, K.; Tselousov, A.; and Wolf, T. 2020. Large-scale transfer learning for natural language generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, (ACL-2019)*, 6053–6058.
- Humeau, S.; Shuster, K.; Lachaux, M.-A.; and Weston, J. 2019. Poly-encoders: Architectures and Pre training Strategies for Fast and Accurate Multi sentence Scoring. In *International Conference on Learning Representations, (ICLR-2020)*.
- Kingma, D. P.; and Welling, M. 2014. Auto-encoding variational bayes. In *Proceedings of 2nd International Conference on Learning Representations, (ICLR-2014)*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, (ACL-2020)*, 7871–7880.
- Li, C.; Gao, X.; Li, Y.; Peng, B.; Li, X.; Zhang, Y.; and Gao, J. 2020. OPTIMUS: Organizing sentences via pre-trained modeling of a latent space. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, (EMNLP-2020)*, 4678–4699.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-2016)*, 110–119.
- Li, J.; Galley, M.; Brockett, C.; Spithourakis, G. P.; Gao, J.; and Dolan, B. 2016b. A persona-based neural conversation model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, (ACL-2016)*, 994–1003.
- Lin, Z.; Madotto, A.; Bang, Y.; and Fung, P. 2021. The Adapter-Bot: All-In-One Controllable Conversational Model. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence, (AAAI-2021)*, 16081–16083.
- Liu, Q.; Chen, Y.; Chen, B.; LOU, J.-G.; Chen, Z.; Zhou, B.; and Zhang, D. 2020. You Impress Me: Dialogue Generation via Mutual Persona Perception. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, (ACL-2020)*, 1417–1427.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled weight decay regularization. In *Proceedings of the 7th International Conference on Learning Representations, (ICLR-2019)*.
- Madotto, A.; Lin, Z.; Wu, C. S.; and Fung, P. 2020. Personalizing dialogue agents via meta-learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, (ACL-2019)*, 5454–5459.
- Nie, Y.; Williamson, M.; Bansal, M.; Kiela, D.; and Weston, J. 2021. I like fish, especially dolphins: Addressing contradictions in dialogue modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, (ACL-2021)*, 1699–1713.
- Qi, W.; Yan, Y.; Gong, Y.; Liu, D.; Duan, N.; Chen, J.; Zhang, R.; and Zhou, M. 2020. ProphetNet: Predicting future n-gram for sequence-to-sequence pre-training. In *Findings of the Association for Computational Linguistics, (EMNLP-2020)*, 2401–2410.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2020. Language Models are Unsupervised Multitask Learners. *OpenAI Blog*, 1(May): 1–7.
- Roller, S.; Dinan, E.; Goyal, N.; Ju, D.; Williamson, M.; Liu, Y.; Xu, J.; Ott, M.; Shuster, K.; Smith, E. M.; Boureau, Y. L.; and Weston, J. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics, (EACL-2021)*, 300–325.
- Sanabria, R.; Palaskar, S.; and Metze, F. 2018. CMU Sinbad’s Submission for the DSTC7 AVSD Challenge. *Dstc7*.

- Shum, H.; He, X.; and Li, D. 2018. From Eliza to XiaoIce: challenges and opportunities with social chatbots. *Frontiers of Information Technology and Electronic Engineering*, 19(1): 10–26.
- Shuster, K.; JU, D.; Roller, S.; Dinan, E.; Boureau, Y.-L.; and Weston, J. 2020. The Dialogue Dodecathlon: Open-Domain Knowledge and Image Grounded Conversational Agents. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, (ACL-2020)*, 2453–2470.
- Song, H.; Wang, Y.; Zhang, K.; Zhang, W. N.; and Liu, T. 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, (ACL-2021)*, 167–177.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Proceedings of Advances in Neural Information Processing Systems, (NIPS-2014)*, 3104–3112.
- Welleck, S.; Weston, J.; Szlam, A.; and Cho, K. 2020. Dialogue natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, (ACL-2019)*, 3731–3741.
- Williams, A.; Nangia, N.; and Bowman, S. R. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-2018)*, 1112–1122.
- Wolf, T.; Sanh, V.; Chaumond, J.; and Delangue, C. 2019. TransferTransfo: A Transfer Learning Approach for Neural Network Based Conversational Agents. *arXiv preprint arXiv:1901.08149*.
- Zhang, S.; Dinan, E.; Urbanek, J.; Szlam, A.; Kiela, D.; and Weston, J. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, (ACL-2018)*, 2204–2213.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 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)*, 654–664.
- Zheng, Y.; Zhang, R.; Mao, X.; and Huang, M. 2020. A Pre-Training based personalized dialogue generation model with persona-sparse data. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, (AAAI-2020)*, 9693–9700.