

MORPHEUS*: Modeling Role from Personalized Dialogue History by Exploring and Utilizing Latent Space

Yihong Tang¹, Bo Wang^{2,3†}, Dongming Zhao⁴
Xiaoja Jin⁴, Jijun Zhang⁴, Ruifang He², Yuexian Hou²

¹School of New Media and Communication, Tianjin University, Tianjin, China

²College of Intelligence and Computing, Tianjin University, Tianjin, China

³Institute of Applied Psychology, Tianjin University, Tianjin, China

⁴AI Lab, China Mobile Communication Group Tianjin Co., Ltd.

{toyhom, bo_wang}@tju.edu.cn

Abstract

Personalized Dialogue Generation (PDG) aims to create coherent responses according to roles or personas. Traditional PDG relies on external role data, which can be scarce and raise privacy concerns. Approaches address these issues by extracting role information from dialogue history, which often fail to generically model roles in continuous space. To overcome these limitations, we introduce a novel framework **MO**del **R**oles from **P**ersonalized **D**ialogue **H**istory by **E**xploring and **U**tilizing **L**atent **S**pace (MORPHEUS) through a three-stage training process. Specifically, we create a persona codebook to represent roles in latent space compactly, and this codebook is used to construct a posterior distribution of role information. This method enables the model to generalize across roles, allowing the generation of personalized dialogues even for unseen roles. Experiments on both Chinese and English datasets demonstrate that MORPHEUS enhances the extraction of role information, and improves response generation without external role data. Additionally, MORPHEUS can be considered an efficient fine-tuning for large language models.

1 Introduction

Personalized Dialogue Generation (Chen et al., 2024b) (PDG), a critical method for capturing and responding to users' emotions (Wang et al., 2024) and preferences (Zhou et al., 2022) in dialogues, aims to model role bases according to given roles or personas, thus generating role-relevant responses to meet users' individual needs. Many methods have achieved commendable results by explicitly

* Morpheus excels at imitating others' forms, effortlessly embodying their physicality, posture, facial expressions, and vocal tone. He flawlessly embodies every intricate detail, including attire and unique mannerisms. — *Metamorphoses* by Ovid

† Corresponding author.

incorporating external role data into prompts and capturing the direct relationship between dialogue and roles during inference.

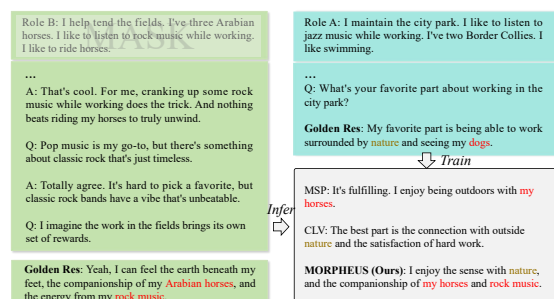


Figure 1: In the task of PDG when role data is masked, an example generated by two state-of-the-art (SOTA) models, MSP and CLV, alongside our model MORPHEUS, is presented. Roles A and B are similar, and we expect models to generalize from the data about A to learn how to generate B's response.

Despite their successes, existing PDG methods overly rely on extracting information from sparse external role knowledge bases (Zhang et al., 2018a; Song et al., 2019a; Wolf et al., 2019; Liu et al., 2020; Song et al., 2021); or brief textual persona descriptions (Qian et al., 2018; Zheng et al., 2020; Song et al., 2021). However, such external role data is often lacking and may raise privacy concerns.

Therefore, researchers attempt to leverage the role information presented in dialogue history (Li et al., 2016b; Chan et al., 2019; Ma et al., 2021). The MSP (Zhong et al., 2022) introduces a token-level retrieval-generation system that retrieves relevant personalized dialogues from a corpus and extracts tokens as prompts for generating responses. While MSP employs a retrieval strategy for the implicit modeling of roles, it treats different roles in isolation and overlooks the inherent associations across roles. As illustrated in Figure 1, these models must grasp the concept that, similar to Role A,

Role B derives comfort and enjoyment from “nature” and certain aspects of personal work and life, such as listening to music and spending time with pets, and utilize this understanding to craft appropriate responses, even when the persona data of Role B is masked.

To address this issue, CLV (Tang et al., 2023) attempts to model roles in latent space, dealing with similar characteristics through implicit clustering. Furthermore, PersonaPKL (Han et al., 2023) trains the model to perceive roles and then transfers this role perception ability to specific roles. However, these methods focus only on the associations between local role features, which leads to the model’s difficulty in generating personalized responses for unseen roles. In fact, by learning dialogues based on roles, models could potentially acquire a universal ability to model unseen roles.

To achieve this insight, we introduce MORPHEUS, a framework that **MO**del **R**oles from **P**ersonalized **D**ialogue **H**istory by **E**xploring and **U**tilizing latent **S**pace. First, we train the model to perceive roles, where it is supervised to generate personalized responses based on personas data and history, thereby establishing the relationship between persona information and dialogue history. Next, we create a persona codebook to represent roles compactly in latent space, initializing it with the encoded representations of existing personas. Finally, we construct a posterior distribution of persona information in the codebook using explicit personas and predict these codes using the dialogue history. Ultimately, our three-stage training method results in a codebook that generalizes the representation of all roles. During inference, MORPHEUS can automatically infer codes from the dialogue history, extract the code representations from the codebook, and use them to generate dialogue.

In summary, our contributions are as follows:

- (1) To our knowledge, this is the first systematic work considering role generalization and the generation of personalized dialogues for unseen roles;
- (2) We propose a novel framework MORPHEUS that improves the ability to generalize roles in the face of unseen roles or no external role data by modeling roles globally in the latent space;
- (3) By applying MORPHEUS, experimental results show significant improvements in the personalization of dialogue generation. Additionally, MORPHEUS can serve as a parameters-efficient

fine-tuning approach for large language models.

2 Related Work

Personalized Dialogue Generation Previous research has demonstrated that incorporating personality traits into dialogue generation is crucial (Chen et al., 2024a) for personalizing dialogue systems. The question of how to maximize the utilization of explicit persona data when available is a concern for most studies. These data include factual role descriptions (Qian et al., 2018; Zheng et al., 2020; Song et al., 2021) (e.g., "I enjoy swimming"), personality traits (e.g., introverted), knowledge bases (Zhang et al., 2018a; Song et al., 2019a; Wolf et al., 2019; Liu et al., 2020; Song et al., 2021) (e.g., gender: male), and speaking styles (Lu et al., 2023) (e.g., excited). These data contain rich persona information, enabling dialogue systems to efficiently and reliably extract information for use in dialogues.

In a word, the goal of all research is to model a role profile with personality, engaging in personalized, continuous dialogue with users. However, compared to explicit data, which is scarce and difficult to extract from the real world, dialogue histories can be collected on a large scale and also contain rich persona information. Modeling personality solely from dialogue history is challenging but offers significant advantages, including avoiding privacy issues associated with explicit data.

The essence of modeling personalities from dialogue history is to extract and utilize generic persona information. Some past studies, such as DHAP (Ma et al., 2021) and MSP (Zhong et al., 2022), enhance current dialogue generation by explicitly retrieving relevant dialogues from the entire corpus. These approaches achieve the goal of modeling roles from dialogue history but suffer from the limitations of corpus size and the inefficiency of retrieval-based methods. Unlike retrieve-based studies, such as PersonaWAE (Chan et al., 2019), model personalization in the latent space for the user side without considering the personality of the respondent.

PersonalPKT (Han et al., 2023), achieve the goal of inferring dialogue for target roles without explicit data by efficiently fine-tuning the model in two stages with explicit persona data, but this still requires additional training for each target. CLV (Tang et al., 2023) argues that explicit persona descriptions hide certain specific angle information.

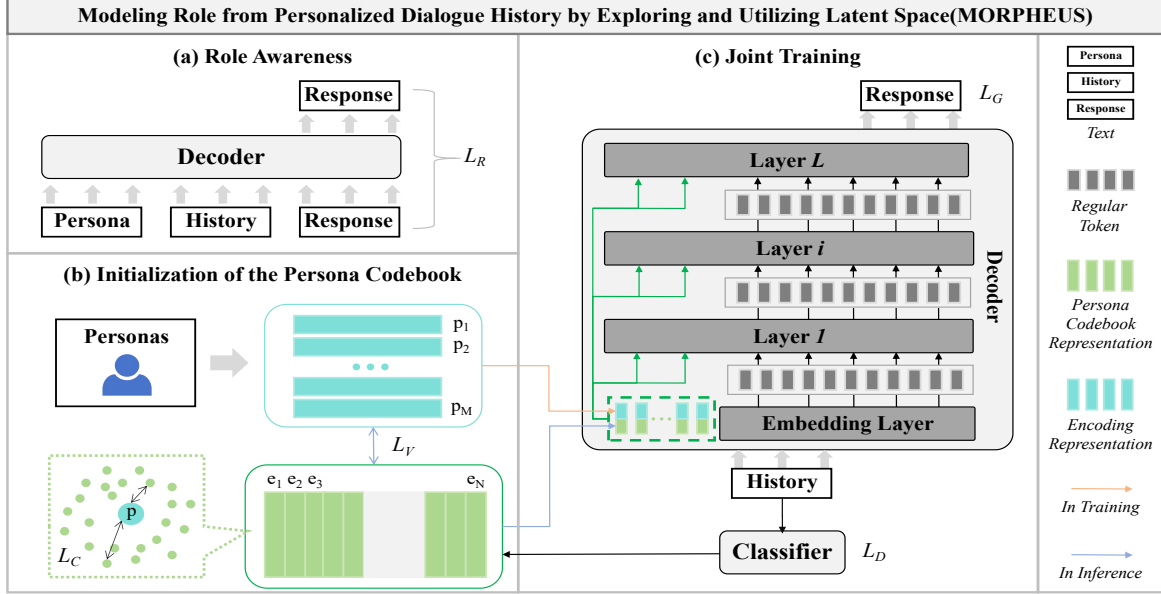


Figure 2: The overview structure of the proposed model.

It clusters persona information in a latent variable space and obtains information solely from dialogue history through a CAVE (Zhao et al., 2017). Furthermore, we believe that all persona information can actually be compressed in a larger latent space for generalization, which can then be utilized during dialogue generation.

3 Methodology

3.1 Overview

In this section, we first formalize the personalized dialogue generation with masked explicit person data and provide an overview of the framework we propose. Given a set of roles U , for each role u_i in U , \mathcal{P} represents the textual persona data to that role. Subsequently, there exists a multi-turn dialogue history \mathcal{C} between u_i and u_j , with the current turn awaiting a response from u_i . Our objective is to generate a personalized response \mathcal{R} that aligns closely with the u_i in response to the dialogue history \mathcal{C} . As previously mentioned, modeling persona or role profiles directly from dialogue history signifies that, during inference, the model does not require external data \mathcal{P}_{u_i} . Instead, it leverages the persona information \mathcal{P} implicit within the \mathcal{C} :

$$P(\mathcal{R}|\mathcal{C}) = \sum_{\mathcal{P}} P(\mathcal{R}|\mathcal{C}, \mathcal{P}) \cdot P(\mathcal{P}|\mathcal{C}). \quad (1)$$

As depicted in Figure 2, our proposed framework is divided into three primary training stages: (1) Role Awareness. Typically, an unconditional decoder is trained on a broad corpus, lacking specific

perception capabilities for dialogues and persona information. The goal of this stage is to fine-tune the model to enhance its perception of persona and the relationship between persona and dialogue, thereby enabling the generation of dialogues that are more aligned with role traits. (2) Persona Codebook(PC) Initialization. Considering our model will directly model roles from dialogue history and PC, the existing codebook might not align with the encoding representations of personas. Hence, we explored various initialization methods for PC, finding that PC generalized from specific persona encoding representations significantly outperforms random initialization in effectiveness. (3) Joint Training. This crucial training stage involves the model taking dialogue history as input, sampling PC representations from PC based on the implicit persona information within the dialogue, thereby producing personalized responses.

3.2 Role Awareness

In decoder-only architecture models, it is common practice to concatenate the persona with the dialogue history and then generate responses in an auto-regressive manner. Similarly, the hidden representation of role data can also be directly concatenated at the beginning of the model in the form of past keys and values. Drawing on the work of Han et al. (2023), we found that forcing the model to generate responses based on dialogue history, conditioned on the hidden state form of role data, can enhance the model’s joint perceptual ability

of persona information and dialogue history. Following this approach, a persona segment about the role is encoded, called encoding representation p_i , and then only the keys and values are concatenated at the forefront of the model. Subsequently, we fine-tune the model:

$$L_R = \sum_j^l \log P(\hat{\mathcal{R}}_i | \mathcal{P}, \mathcal{C}, \mathcal{R}_{<j}). \quad (2)$$

3.3 Initialization of the Persona Codebook

Upon creating a Persona Codebook(PC) $e \in \mathbb{R}^{N \times d}$ where N is the size of the discrete latent space and d is the hidden size. The e contains N PC representation e_k . Our aim is to compactly utilize the PC to represent a wide variety of persona information within this latent space. With the model already having a strong perception of roles, a critical issue arises: whether the e_k in the PC can align with the representation p_i of textual persona data encoded:

$$\mathcal{P}_0, \mathcal{P}_2, \dots, \mathcal{P}_{M-1} = \text{Split}(\mathcal{P}), \quad (3)$$

$$p_i = \text{Encoder}(\mathcal{P}_i), \quad (4)$$

where $p_i \in \mathbb{R}^d$. *Split* represents an explicit split at the string level, with a period as the separator, and M comes from the theory in CLV (Tang et al., 2023) that the description angle in the persona data can be divided into M categories.

This alignment is crucial for the e_k to be effective during decoding. Indeed, we find that the initialization of the PC is paramount in addressing this issue. Utilizing random uniform distribution for initializing often results in difficulty enhancing PC during subsequent parameter updates, potentially leading to the failure of joint training. Conversely, a well-initialized PC significantly smoothens the joint training process described in Section 3.4. To this end, we design three novel initialization methods to tackle this challenge: sequential persona initialization, average persona initialization, and EM (Dempster et al., 1977) (Expectation-Maximization) initialization.

During training, sequential initialization populates the uninitialized portions of the PC with the encoding representations p_i of persona data from the current batch in sequence, while average initialization performs an additional averaging operation on the batch's p_i . The EM algorithm-based initialization emerged as the most effective, operating independently from the joint training process. In EM initialization, we assume that the p_i of persona

information is sampled from one of several normal distributions, each centered at a e_k . The means and variances of these distributions are treated as parameters to be estimated by the EM algorithm.

The Expectation Step For a PC of size N , we establish N normal distributions, using all persona data encoding representations from the training set as samples, totaling $|D|$. For each encoding representation p_i , we calculate its posterior probability $P(z_i = k | p_i)$ of belonging to each normal distribution N_k (with mean μ_k and variance σ_k^2), where $z_i = k$ denotes the latent variable indicating that p_i belongs to the k -th distribution:

$$P(z_i = k | p_i) = \frac{P(p_i | z_i = k)P(z_i = k)}{\sum_{j=1}^N P(p_i | z_i = j)P(z_i = j)}, \quad (5)$$

where $P(z_i = k)$ is the prior probability of the latent variable $z_i = k$, assumed to be uniform distribution.

The Maximization Step The goal during the M-step is to update the parameters (mean μ_k and variance σ_k^2) of each distribution to maximize the expected log-likelihood of the posterior probabilities calculated in the current step:

$$\mu_k = \frac{\sum_{i=1}^{|D|} P(z_i = k | p_i) p_i}{\sum_{i=1}^{|D|} P(z_i = k | p_i)} \quad (6)$$

$$\sigma_k^2 = \frac{\sum_{i=1}^{|D|} P(z_i = k | p_i) (p_i - \mu_k)^2}{\sum_{i=1}^{|D|} P(z_i = k | p_i)} \quad (7)$$

The denominator $\sum_{i=1}^{|D|} P(z_i = k | p_i)$ represents the weighted count of all encoding representations p_i belonging to distribution k , essentially the "effective" sample size considering the uncertainty of each sample belonging to each code z . Through this method, we obtain accurate parameters for the normal distributions, using the derived mean μ_k as the initial value e_k for the z_k .

3.4 Joint Training

Following the initial setup of both components, our model embarks on learning how to predict codes based on the implicit persona information contained within the dialogue history. It then extracts PC representation e_k from the Persona Codebook (PC), thereby generating personalized responses. To train this mechanism, our joint training regimen is bifurcated into two concurrent steps: (1) Code Index Prediction and (2) Training of the PC.

During the training phase of the PC, following the approach of VQ-VAE (van den Oord et al., 2017), we identify the e_k that most closely aligns with the persona’s encoding representation p_i :

$$k = \arg \min_k \|p_i - e_k\|_2. \quad (8)$$

We optimize the proximity of these PC representations by minimizing their distance, employing the following loss function:

$$L_V = \|sg[p_i] - e_k\|_2^2 + \beta \|sg[e_k] - p_i\|_2^2, \quad (9)$$

where the sg denotes the stop gradient operator, defined as an identity operation during forward computation but possesses a zero partial derivative, effectively rendering its operands as un-updated constants. The encoder optimizes the last loss term, while the embeddings are optimized by the first loss term. Our findings indicate that the algorithm exhibits sensitivity towards β because the encoder’s parameter volume significantly exceeds that of the intermediary processes. Across all experiments, we utilize $\beta = 0.05$.

Subsequently, the encoded representations participate in autoregressive dialogue generation as past keys and values, with its loss function as:

$$p_G = \text{concat}(p_1, p_2, \dots, p_M), \quad (10)$$

$$L_G = - \sum_j^l \log P(\mathcal{R}_j | p_G, \mathcal{C}, \mathcal{R}_{<j}), \quad (11)$$

In the realm of Code Index Prediction, the code index k that is nearest to the p_i is employed as the label. The dialogue history’s hidden states c are inputs, utilizing a multilayer MLP as the classifier:

$$c = \text{Decoder}(\mathcal{C}), \quad (12)$$

$$L_D = - \log P(y = k | c), \quad (13)$$

Furthermore, to ensure the PC possesses sufficient diversity and generalization capability, we introduce additional contrastive learning (Gao et al., 2021) loss pertaining to each latent variable within the codebook. This measure is aimed at enhancing the representational breadth and applicability of the PC across varied dialogic scenarios:

$$L_C = - \log \frac{e^{\text{sim}(p_i, e_k)/\tau}}{\sum_{j=1}^N e^{\text{sim}(p_i, e_j)/\tau}}, \quad (14)$$

where τ is a temperature hyperparameter and $\text{sim}(p_i, e_i)$ is the cosine similarity.

Dataset	# Train	# Valid	# Test
ConvAI2	43,410	4,213	2,138
Baidu PersonaChat	376,016	19,923	4,456

Table 1: Statistics of persona dialogue datasets.

4 Experiments

4.1 Datasets

ConvAI2 (Dinan et al., 2019) is an enriched English dataset centered on personalized dialogues that stem from specific role profiles. This dataset, an extension of the original PersonaChat (Zhang et al., 2018b), was meticulously curated by crowd workers to enhance its quality and relevance.

In a parallel vein, **Baidu PersonaChat**¹, a Chinese counterpart to ConvAI2, mirrors its structure by emphasizing personalization and leveraging individual details to drive engaging dialogues.

Table 1 presents statistics for the two datasets. We utilize these persona descriptions only during the training stage. We use the dataset curated by (Tang et al., 2023) but ensure no role leaks for the training, validation, and test data.

4.2 Baselines

In the exploration of personalized dialogue generation models without external role data, research in this domain can generally be categorized into two main types: First, the approach based on explicit retrieval, which relies on fetching the most relevant response or tokens from a predefined or dynamically generated response pool according to the current dialogue history. Second, the approach is based on implicit modeling, where the personalization or roles are learned and generated directly within the model itself, without the necessity for an explicitly defined response or tokens. Furthermore, when considering non-personalized dialogue generation models, we can categorize the existing model baselines into three distinct classes:

Non-Personalized Approaches The **Seq2Seq** (Sutskever et al., 2014) enhanced with an attention mechanism (Luong et al., 2015) exemplifies a key sequence modeling with context-focusing capabilities. The pre-trained **GPT-2** (Radford et al., 2019) is a cornerstone for dialogue creation upon fine-tuning with dialogue-specific datasets.

Approaches based on Explicit Retrieval **DHAP** (Ma et al., 2021) capitalizes on historical

¹<https://www.luge.ai/#/luge/dataDetail?id=38>

		Coherence			Diversity			Consistency	
		BLEU-1	ROUGE-L	Coh.Score	Dist-1	Dist-2	sBLEU ↓	Coh-Con.Score	P-Co
ConvAI2	Seq2Seq	3.41	5.48	35.89	2.04	3.98	13.22	10.85	3.13
	GPT-2	6.50	10.91	58.79	4.71	25.28	10.51	13.29	4.59
	DHAP	7.37	9.97	63.21	5.60	29.81	9.85	16.04	9.27
	MSP	8.55	10.96	64.19	5.11	28.60	9.93	15.45	8.92
	PersonaPKT	8.70	11.08	60.58	6.30	26.72	9.34	24.87	9.26
	CLV	11.16	15.04	70.83	4.31	26.17	10.14	23.01	9.38
	MORPHEUS(Ours)	12.67	16.18	73.19	5.89	28.74	8.97	31.57	11.64
Baidu PersonaChat	Seq2Seq	7.98	8.24	40.11	0.97	5.19	16.79	8.96	3.11
	GPT-2	10.16	12.29	49.72	3.08	20.98	13.32	12.14	5.30
	DHAP	11.23	11.58	53.89	3.11	22.10	13.42	12.30	7.95
	MSP	14.44	13.24	58.59	3.37	22.41	13.95	14.37	8.23
	PersonaPKT	13.82	15.57	53.95	2.98	21.83	13.10	19.86	8.17
	CLV	18.77	21.82	59.74	2.42	21.17	14.57	18.15	8.39
	MORPHEUS(Ours)	19.70	24.64	62.45	3.07	23.05	12.56	29.93	10.96

Table 2: Automatic evaluation on two datasets. The best results are in **bold**.

interaction data to dynamically craft query-sensitive user profiles, employing a specialized decoder for tailored responses. In a similar vein, **MSP** (Zhong et al., 2022) propels personalized dialogue generation forward by identifying conversations from users with similar profiles using User and Topic Refiners, and refining tokens for training via a Token Refiner, thus achieving superior performance in generating personalized dialogues without relying on predefined personas.

Approaches based on Implicit Modeling PersonaPKT (Han et al., 2023) utilizes minimal additional parameters to embed persona-specific traits, leveraging lightweight tuning methods to transitively model roles in two phases via continuous vectors. **CLV** (Tang et al., 2023) categorize persona information in latent space to enhance personalized response generation with dialogue history.

Implementation Details are in Appendix A.1.

4.3 Evaluation Metrics

In order to obtain accurate comparisons, we use both automatic and human evaluations.

Automatic Evaluation We divide the automatic evaluation methods into three categories in order to evaluate and model the diversity, consistency, and coherence of the generated dialogues.

(1) **Coherence** BLEU-1/2 (Papineni et al., 2002) and ROUGE-L (Lin and Och, 2004) are classical words overlap-based metrics for measuring the similarity between generated responses and factual responses, which we believe can indirectly measure the coherence of dialogues. According to Tang et al. (2023), BLEU-3/4 imposes overly strict coverage requirements, making it unable to reflect the coherence between dialogue history and the response.

Therefore, we do not consider it as part of the coherence assessment.

(2) **Diversity**: The Dist-1/2 (Li et al., 2016a) assesses the presence of unique uni- or bi-grams within the generated responses, commonly employed for evaluating diversity at the corpus level. The self-BLEU (sBLEU) (Liu et al., 2022) computes the diversity by BLEU from the similarity between responses.

(3) **Consistency**: In dialogue generation, maintaining consistency between the generated responses and the role is crucial. Following Tang et al. (2023), we introduce two scoring models, Con.Score and Coh-Con.Score, to measure the consistency between responses and persona information. Furthermore, we employ the P-Co (Persona Cosine Comparison) (Song et al., 2019b) as a supplementary IDF-weighted words overlap-based metric to assess the semantic similarity between model responses and personas.

Human Evaluation Considering the auto evaluation criteria’s uncertainty, we conduct human assessments for all models. Specifically, We extract 100 data points (histories, responses, and role data) and engage three well-educated annotators to rank the responses from various models and normalize them into specific scores ranging from 0 to 1 to save costs. Our evaluation centers on readability, diversity, consistency, and coherence, prompting evaluators to rank 6 options for the five excellent model-generated responses and the factual responses. Detailed evaluation criteria can be found in Figure 4 in Appendix A.2.

Model	Readability	Diversity	Consistency	Coherence
DHAP	0.71	0.72	0.69	0.71
MSP	0.70	0.73	0.69	0.8
PersonaPKT	0.63	0.65	0.61	0.69
CLV	0.79	0.76	0.73	0.76
MORPHEUS(Ours)	0.82	0.84	0.77	0.85
Ground-Truth	0.92	0.89	0.93	0.91

Table 3: The result of human evaluation on ConvAI2.

4.4 Experimental Results

Automatic Evaluation Table 2 presents the performance of all models on automatic metrics in both Chinese and English datasets for the average value obtained by three random seeds, allowing for a clear observation of the enhancements in key metrics by our MORPHEUS. Specifically:

(1) **Coherence**: Whether through coverage metrics such as BLEU-1, Rouge-L, or learning indices like Coh-Con.Score, the MORPHEUS’s performance in consistency stands out.

(2) **Diversity**: MORPHEUS also excels at generating more diverse responses compared to other models. Additionally, we overcome the drawback commonly associated with latent variable models like CLV (Tang et al., 2023), which sacrifice diversity to achieve improvements in other performance metrics. This also indicates that MORPHEUS’s efforts are closer to the intended modeling target, leading to the convergence in the meanings of the responses. We will analyze this reason further in Section 5.1.

(3) **Consistency**: The guidance from Con.Score indicates the model’s ability to integrate persona information into generation, particularly evident in scenarios without external persona data, further highlighting the superiority of MORPHEUS. Experimental results demonstrate that the proposed model can generate more personalized responses than all baselines.

Human Evaluation The human evaluation outcomes for ConvAI2 are depicted in Table 3. We computed the Fleiss Kappa coefficient (Landis and Koch, 1977) among the three annotators by a novel approach wherein the ordinal rankings are conceptualized as discrete categories. The Kappa value is 0.65, indicating *substantial agreement* among the annotators. Overall, the human annotations align with the automatic evaluation findings, showcasing the strengths of MORPHEUS in personalized dialogue generation and fundamental readability.

Models	BLEU-1	Dist-1/2	Coh-Con.Score	P-Co
MORPHEUS(Ours)	12.67	5.89/28.74	31.57	11.64
w/o PC	6.41	4.83/25.97	13.20	4.46
w/o CL	9.24	4.19/24.93	23.02	8.96
w/o RA	8.56	5.03/27.39	29.05	9.56
<i>i/w</i> Random	10.99	5.12/27.95	28.53	10.32
<i>i/w</i> Se.	11.14	5.55/28.14	28.15	10.94
<i>i/w</i> Mean	11.69	5.41/28.30	30.12	11.19

Table 4: Ablation experiments results on ConvAI2.




Figure 3: Experiments with the different N on the ConvAI2 dataset. For ease of viewing, BLEU-1, rouge-L, P-Co are magnified by a factor of three and Dist-1 by a factor of ten.

5 Further Analysis

Further analyses are based on the ConvAI2 dataset, and similar phenomena can be observed on Baidu PersonaChat.

5.1 Ablation Study

We conducted ablation studies by removing specific modules of the MORPHEUS model, as presented in Table 4. Firstly, we examined the influence of the model’s core mechanism—the persona codebook (PC)—on its performance; upon the removal of the PC, the model regresses to a basic GPT-2, and its performance aligns with that of GPT-2. If we merely eliminate the contrastive learning from the joint training while retaining the VQ-VAE loss function, there is a noticeable decline in the model’s performance, especially the diversity. This underscores the effectiveness of employing contrastive learning to bring the encoding representations closer to the PC representations. Then, regarding various PC initialization methods, the EM(Ours) technique emerged as superior to random initial, sequential filling, or mean filling, suggesting that a global perspective on roles better optimizes the PC optimally. Finally, the removal of role-aware training results in a performance drop in certain metrics, indicating that aligning roles with

Role Data	I enjoy hiking . I've never been abroad. I'm a bank teller. I've never been to college. My favorite phone is the iPhone.	I like watching cooking shows. I'm not good at swimming at all. I enjoy roller coasters and skydiving. I'm a vegetarian, and I like Hummers.	I help tend the fields. I've three Arabian horses. I like to listen to rock music while working. I like to ride horses .
Dialogue History	Q: Hi, what are you up to? A: I'm getting ready to go out. How about you? What are you doing? Q: I'm traveling with my girlfriend right now.	...Q: I'm a carpenter, I started my own business. A: That's great! Q: How old are you? I feel you are so young. A: I just turned 17 . Q: Oh, you're really young, what do you do?	Q: I spent my most of the time in study. A: What are you studying? Q: I am studying horses. A: That's interesting. Q: What is your favorite band? A: I like rock. Q: Me too.
Golden-Response	Do you like traveling everywhere? I do too.	I don't have a job yet. I just moved here and haven't found one yet.	Do you like horses? I have three.
MSP	Could you tell me about some of your travel experiences ?	I'm a farmer.	Do you like music? I like pop music.
CLV	I often go out with my girlfriend too.	I'm out running.	I don't like jazz.
MORPHEUS(Ours)	Do you like traveling everywhere? I do too.	I'm still a student .	Do you like animals ?

Table 5: A case study. Keywords are marked in red.

method	BLEU-1	ROUGE-L	Dist-1	Dist-2	sBLEU ↓	P-Co	TPS(%)
Prompt-Tuning	8.15	10.02	4.24	11.27	7.62	8.81	0.03
Prefix-Tuning	9.23	11.08	4.12	13.19	9.13	8.14	0.17
P-Tuning	11.07	14.15	5.12	17.05	8.93	10.02	1.40
LoRA	14.06	18.91	5.97	20.73	9.79	11.83	0.23
MORPHEUS	16.98	20.03	7.34	32.28	7.92	12.30	0.16
LLaMA(Full)	28.11	23.51	7.53	33.59	7.88	18.96	100

Table 6: Comparative experimental results of the PEFT methods based on LLaMA2-7B in ConvAI2. The TPS means trainable parameter size.

dialogue through text is indispensable.

5.2 Effects of size of the persona codebook

In MORPHEUS, persona information is compressed into the persona codebook (PC), where the size of the PC, denoted as N , has a significant impact on the performance as reported in Figure 3. When N is small, the information compressed into the PC is more closely aligned with the generic information pertinent to the task of personalized dialogue generation rather than a specific role. Conversely, too large N can diminish the effectiveness of this compression. This reduction in effectiveness is attributed to an increase in noise resulting from an insufficient number of distinct personas, which leads to the dispersion of persona information and makes it challenging for the PC to be optimized effectively. On balance, the best performance is observed when N is 100.

5.3 PEFT

Our approach of concatenating the persona codebook(PC) representation vectors to the forefront of the model is akin to prompt-tuning and p-tuning. Thus, our technique can be considered an efficient fine-tuning method. Specifically, we freeze the parameters of the decoder(LLaMA-2 (Touvron et al., 2023)) responsible for text generation and only optimize the parameters within the PC during training,

Method	BLEU-1	ROUGE-L	Dist-1	Dist-2	sBLEU ↓	P-Co
LoRA	15.17	17.88	6.99	26.28	9.12	12.39
MORPHEUS	17.60	21.55	7.17	33.81	7.54	12.36
Mistral (Full)	29.24	24.49	7.48	36.74	7.17	19.11

Table 7: Performance comparison between LoRA, MORPHEUS and full fine-tuning based on Mistral-7B-v0.1 in ConvAI2.

with results presented in Table 6. We compare our method against LoRA, prompt-tuning, p-tuning, and full-parameter fine-tuning under a similar parameter budget. Findings reveal that our method exhibits distinct superiority and yields results that are closer to those obtained from full parameter fine-tuning.

Beyond LLaMA-2, we also conduct experiments with the Mistral-7B-v0.1 model, a newly popular LLM. As presented in Table 7, the performance results from Mistral further demonstrate the robustness of our proposed MORPHEUS. The BLEU, ROUGE, and Distinct metrics suggest that Mistral's architecture enhances performance in some areas of dialogue personalization compared to LLaMA. However, despite Mistral's improvements, the integration of persona information into LLMs remains a significant challenge. This observation highlights the complexity of generating truly personalized and contextually relevant dialogues, even with state-of-the-art models.

5.4 Case Study

For the case study, we select examples, including some that are clear failures, yet MORPHEUS consistently performs as well as or better than baselines. The results in Table 5 indicate that MORPHEUS effectively aligns dialogue history and role.

In case 1, the questioner is traveling, and while MSP and CLV both mentioned travel-related con-

tent, they failed to infer the respondent’s hobby from the dialogue history. MORPHEUS accurately captured the fact that the respondent was preparing to go out, creating a response closely aligned with the role setting of "enjoy hiking". In case 2, when the questioner inquires about the profession, MSP and CLV both fail to notice the respondent’s age mentioned in the history, whereas MORPHEUS generates a response that is consistent with the history and close to the ground-truth response. The same is true for case 3, due to the confusing history, MSP and CLV struggle to make relevant responses, while MORPHEUS’s response is closer to the role.

Overall, facing previously unseen roles in the test set and without the external role data, MSP and CLV are easily misled by confusing history and unable to produce coherent responses. MORPHEUS overcomes this by accurately utilizing dialogue history to make reasonable inferences about the roles.

6 Conclusion

In this work, we introduce the MORPHEUS framework for generating personalized responses. Unlike existing efforts, we model roles in a generalized manner from the dialogue history of roles and align roles with dialogues through a multi-stage training process to enhance the personalization of dialogue generation. Experimental results demonstrate that our model is capable of generating highly personalized responses even when faced with previously unseen roles, efficiently utilizing persona information while mitigating the risk of privacy breaches.

Limitations

The MORPHEUS framework introduces a significant advancement in the field of Personalized Dialogue Generation (PDG), particularly through its innovative use of latent space to model roles from dialogue history. However, like any study, there are areas that were beyond the scope of our current investigation. For instance, while MORPHEUS has been rigorously tested on Chinese and English datasets, its adaptability to other languages and dialects remains an area for future exploration. This limitation is not a reflection of the model’s effectiveness but rather an opportunity to further validate its versatility across diverse linguistic contexts. Additionally, the current implementation of the persona codebook, although effective in compactly representing roles, offers potential for further refinement to enhance its representation capabilities.

It’s important to note that these considerations do not detract from the model’s current achievements but highlight pathways for continued research and development. Furthermore, our study’s focus on the technical development and application of MORPHEUS meant that a deeper dive into the broader societal implications of PDG technologies was outside our immediate scope. As the field progresses, these considerations will become increasingly important to address.

Ethics Statement

This study acknowledges the ethical considerations inherent in the development and application of personalized dialogue generation technologies. While MORPHEUS aims to enhance the personalization of dialogue generation without relying on external role data, thus mitigating privacy concerns, it is imperative to consider the potential for misuse, such as the creation of misleading or harmful content under the guise of personalization. To address these concerns, the research strictly adheres to ethical guidelines, ensuring that all data used for training and evaluation is sourced from publicly available datasets that are free from unethical content. Moreover, the study employs anonymization techniques during human evaluation to protect the privacy of individuals. Despite these precautions, we recognize the importance of ongoing ethical scrutiny and the development of robust frameworks to prevent the misuse of personalized dialogue technologies, especially as they become more sophisticated and widely adopted.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (62376188, 62272340, 62276187, 62376192, 62166022).

References

- Zhangming Chan, Juntao Li, Xiaopeng Yang, Xiuying Chen, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. [Modeling personalization in continuous space for response generation via augmented wasserstein autoencoders](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024a. [From persona to personalization: A survey on role-playing language agents](#).

- Yi-Pei Chen, Noriki Nishida, Hideki Nakayama, and Yuji Matsumoto. 2024b. [Recent trends in personalized dialogue generation: A review of datasets, methodologies, and evaluations](#). In *International Conference on Language Resources and Evaluation*.
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. [Maximum likelihood from incomplete data via the em - algorithm plus discussions on the paper](#).
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2019. The second conversational intelligence challenge (convai2). *arXiv preprint arXiv:1902.00098*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [Simcse: Simple contrastive learning of sentence embeddings](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Xu Han, Bin Guo, Yoon Jung, Benjamin Yao, Yu Zhang, Xiaohu Liu, and Chenlei Guo. 2023. [Personapkt: Building personalized dialogue agents via parameter-efficient knowledge transfer](#). In *SUSTAINLP*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Diederik Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *International Conference on Learning Representations*.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 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*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016b. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- Chin-Yew Lin and Franz Josef Och. 2004. [Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics](#). In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics, ACL '04*, page 605–es, USA. Association for Computational Linguistics.
- Guangyi Liu, Zeyu Feng, Yuan Gao, Zichao Yang, Xiaodan Liang, Junwei Bao, Xiaodong He, Shuguang Cui, Zhen Li, and Zhiting Hu. 2022. [Composable text controls in latent space with odes](#). In *Conference on Empirical Methods in Natural Language Processing*.
- 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 *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1417–1427, Online. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Zhenyi Lu, Wei Wei, Xiaoye Qu, Xian ling Mao, Dangyang Chen, and Jixiong Chen. 2023. [Miracle: Towards personalized dialogue generation with latent-space multiple personal attribute control](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). *ArXiv*, abs/1508.04025.
- Zhengyi Ma, Zhicheng Dou, Yutao Zhu, Hanxun Zhong, and Ji-Rong Wen. 2021. [One chatbot per person: Creating personalized chatbots based on implicit user profiles](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 555–564, New York, NY, USA. Association for Computing Machinery.
- 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*, pages 311–318.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, and Ting Liu. 2021. [BoB: BERT over BERT for training persona-based dialogue models from limited personalized data](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1:*

- Long Papers*), pages 167–177, Online. Association for Computational Linguistics.
- Haoyu Song, Weinan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019a. [Exploiting persona information for diverse generation of conversational responses](#). In *IJCAI*.
- Haoyu Song, Weinan Zhang, Yiming Cui, Dong Wang, and Ting Liu. 2019b. [Exploiting persona information for diverse generation of conversational responses](#). In *International Joint Conference on Artificial Intelligence*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Yihong Tang, Bo Wang, Miao Fang, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2023. [Enhancing personalized dialogue generation with contrastive latent variables: Combining sparse and dense persona](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv*, abs/2302.13971.
- Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. [Neural discrete representation learning](#). *ArXiv*, abs/1711.00937.
- Yan Wang, Bo Wang, Yachao Zhao, Dongming Zhao, Xiaojia Jin, Jijun Zhang, Ruifang He, and Yuexian Hou. 2024. [Emotion recognition in conversation via dynamic personality](#). In *International Conference on Language Resources and Evaluation*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. [Transfertransfo: A transfer learning approach for neural network based conversational agents](#).
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. [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.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018b. [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.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Yinhe Zheng, Rongsheng Zhang, Minlie Huang, and Mao Xiaoxi. 2020. [A pre-training based personalized dialogue generation model with persona-sparse data](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:9693–9700.
- Hanxun Zhong, Zhicheng Dou, Yutao Zhu, Hongjin Qian, and Ji-Rong Wen. 2022. [Less is more: Learning to refine dialogue history for personalized dialogue generation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5808–5820, Seattle, United States. Association for Computational Linguistics.
- Jinfeng Zhou, Bo Wang, Minlie Huang, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2022. [Aligning recommendation and conversation via dual imitation](#). In *Conference on Empirical Methods in Natural Language Processing*.

A Appendix

A.1 Default Parameter Settings

Our experiments are done based on pre-trained GPT-2, and we tried various model structures and hyperparameters, and the final hyperparameters are as follows: the size of GPT-2 embedding and GPT-2 hidden vector is 768. All word embedding dimensions are set to 768, and we use word2vec to initialize word embedding. The number of layers of Transformer is 12. The size of the persona codebook N is set from 20 to 500 (default is 100), the MLP input dimension and output dimension in the model are kept the same as the hidden vector, and the number of batches was set to 16. The maximum learning rate is $1e-4$. The training of the proposed model was done on an Nvidia Tesla V100 16G GPU. The total training time takes approximately 10 hours. The temperature hyperparameter τ is 0.5. The pre-trained models used in these experiments of this paper include gpt2², gpt2-chinese-cluecorpussmall³, xlm-roberta-base⁴, and chinese-roberta-wwm-ext⁵.

We use kernel sampling (Holtzman et al., 2020) as our decoding strategy, use the Adam (Kingma

²<https://huggingface.co/gpt2>

³<https://huggingface.co/uer/gpt2-chinese-cluecorpussmall>

⁴<https://huggingface.co/xlm-roberta-base>

⁵<https://huggingface.co/hfl/chinese-roberta-wwm-ext>

and Ba, 2014) optimizer to train the model and use AdamW (Loshchilov and Hutter, 2019) to warm up the generator. Please refer to the published project for additional details, which is publicly available⁶.

When calculating BLEU, ROUGE, and Distinct, we use the TikToken-based word segmentation method `cl100k_base`⁷ for all models to ensure cross-language data fairness.

A.2 Evaluation Criteria of Human

Upon the acquisition of human evaluators' rankings for the responses generated by six distinct models, the algorithm assigns scores in a manner that reflects both the quality and the relative positioning of these responses. Specifically, the scoring scheme is bifurcated based on whether a model's response precedes or follows the gold standard response in the ranking hierarchy.

For responses that are ranked higher than the gold standard, a uniform score of 1 is allocated. This decision is rooted in the assumption that surpassing the gold standard in human evaluators' rankings is indicative of exceptional quality, warranting the highest possible score within the algorithm's framework.

Moreover, recognizing the inherent challenge of applying the Fleiss Kappa coefficient directly to ranking-based assessments, a novel approach is adopted wherein the ordinal rankings provided by the evaluators are conceptualized as discrete categories. This conceptualization facilitates the application of the Fleiss Kappa coefficient to the task at hand, thereby enabling a quantifiable measure of inter-rater reliability in the context of a ranking-based evaluation framework.

we recruited 3 annotators with graduate degrees for \$20 an hour. Detailed evaluation criteria can be found in Figure 4.

⁶<https://github.com/Toyhom/MORPHEUS>

⁷<https://github.com/openai/tiktoken>

Dialogue Response Evaluation

Please read the following instructions carefully before proceeding with the evaluation.

Instructions

You will be presented with a series of dialogue contexts and six corresponding responses. Your task is to sort these responses for each of the following criteria: readability, diversity, consistency, and coherence.

Drag and drop the responses to sort them from best to worst according to the criterion being evaluated. The top should be the response you consider best in terms of the criterion, and the bottom should be the worst.

Evaluation Criteria

Readability:

Evaluate the naturalness and fluency of the response. Consider the correctness of grammar, spelling, and sentence structure, as well as its ease of understanding.

Diversity:

Assess the variety of expressions and ideas in the response. Look for special vocabularies, ways of expression, and viewpoints.

Consistency:

Evaluate how well the response aligns with the given role or background information. Ensure the response is consistent with the role set.

Coherence:

Assess the logical connection of the response to the dialogue context. The answer should logically follow from the previous conversation, forming a meaningful dialogue.

Continue

Figure 4: The evaluation criteria of human.