

ESC-Eval: Evaluating Emotion Support Conversations in Large Language Models

Haiquan Zhao^{*1,2} Lingyu Li¹ Shisong Chen² Shuqi Kong¹ Jiaan Wang²
Kexin Huang^{1,2} Tianle Gu¹ Yixu Wang^{1,2} Jian Wang¹ Dandan Liang¹
Zhixu Li^{†3,4} Yan Teng^{†1} Yanghua Xiao² Yingchun Wang¹

¹ Shanghai Artificial Intelligence Laboratory

² School of Computer Science, Fudan University

³ School of Information, Renmin University of China

⁴ Suzhou Key Laboratory of Artificial Intelligence and Social Governance Technologies, International College (Suzhou Research Institute), Renmin University of China

Abstract

Emotion Support Conversation (ESC) is a crucial application, which aims to reduce human stress, offer emotional guidance, and ultimately enhance human mental and physical well-being. With the advancement of Large Language Models (LLMs), many researchers have employed LLMs as the ESC models. However, the evaluation of these LLM-based ESCs remains uncertain. Inspired by the awesome development of role-playing agents, we propose an ESC Evaluation framework (*i.e.*, ESC-Eval), which uses a role-playing agent to interact with ESC models, followed by a manual evaluation of the interactive dialogues. In detail, we first reorganize 2,801 role-playing cards from seven existing datasets to define the roles of the role-playing agent. Second, we train a specific role-playing model - ESC-Role to mimic the behavior of a real person experiencing distress. Third, through ESC-Role and organized role cards, we systematically conduct experiments using 14 LLMs as the ESC models, including general AI-assistant LLMs (*e.g.*, ChatGPT) and ESC-oriented LLMs (*e.g.*, ExtES-Llama). We conduct comprehensive human annotations on interactive multi-turn dialogues of different ESC models. The results show that ESC-oriented LLMs exhibit superior ESC abilities compared to general AI-assistant LLMs, but there is still a gap behind human performance. Moreover, to automate the evaluation of future ESC models, we developed ESC-RANK, which trained on the annotated data, achieving a scoring performance surpassing 35 points of GPT-4. Our data and code are available at <https://github.com/haidequanbu/ESC-Eval>.

1 Introduction

With the rapid development of Large Language Models (LLMs), an increasing number of individu-

^{*} Work done during internship at Shanghai Artificial Intelligence Laboratory

[†] Correspondence to: Zhixu Li<zhixuli@gmail.com> and Yan Teng<tyan@pjlab.org.cn>

als are engaging in conversations with LLMs (*e.g.*, ChatGPT (Achiam et al., 2023)). Among various conversational applications, Emotional Support Conversation (Liu et al., 2021) (ESC) stands out as a particularly promising field, where people can freely share their personal experiences or concerns, receiving emotional support and practical advice. This interaction helps alleviate human pressures (Langford et al., 1997; Burlison, 2003), thereby improving overall well-being. Recently, numerous LLM-based ESC models have received wide research attention (Zheng et al., 2023b; Qiu et al., 2023; Liu et al., 2023). However, effective and comprehensive evaluation of these chatbots remains challenging.

Current ESC evaluation (Liu et al., 2021; Zheng et al., 2023b) generally uses text-based statistical metrics or manual evaluations. (1) When using text-based statistical metrics, researchers provide the dialogue history to the ESC models and then use the models to generate the corresponding responses (*c.f.*, left panel in Figure 1). Based on the generated responses, text-based statistical metrics (such as ROUGE (Lin, 2004) and BLEU (Papineni et al., 2002)) assess whether the responses resemble ground truth. However, these metrics heavily rely on ground truth responses, which lack objectivity (Novikova et al., 2017) due to the complex nature of ESC. Furthermore, since the conversation history from ground truth is provided to the model under evaluation, text-based statistical metrics cannot fully assess models' capabilities in multi-turn ESC dialogues due to none self-generated bias. (2) Manual evaluations (Liu et al., 2021; Zheng et al., 2023b) employ human evaluators to simulate conversations between the model and users with specific distress (middle panel in Figure 1). This method requires the collection of both human-AI dialogues and manual judgments, resulting in challenges such as high cost and low efficiency.

To alleviate the above issues, we propose ESC-

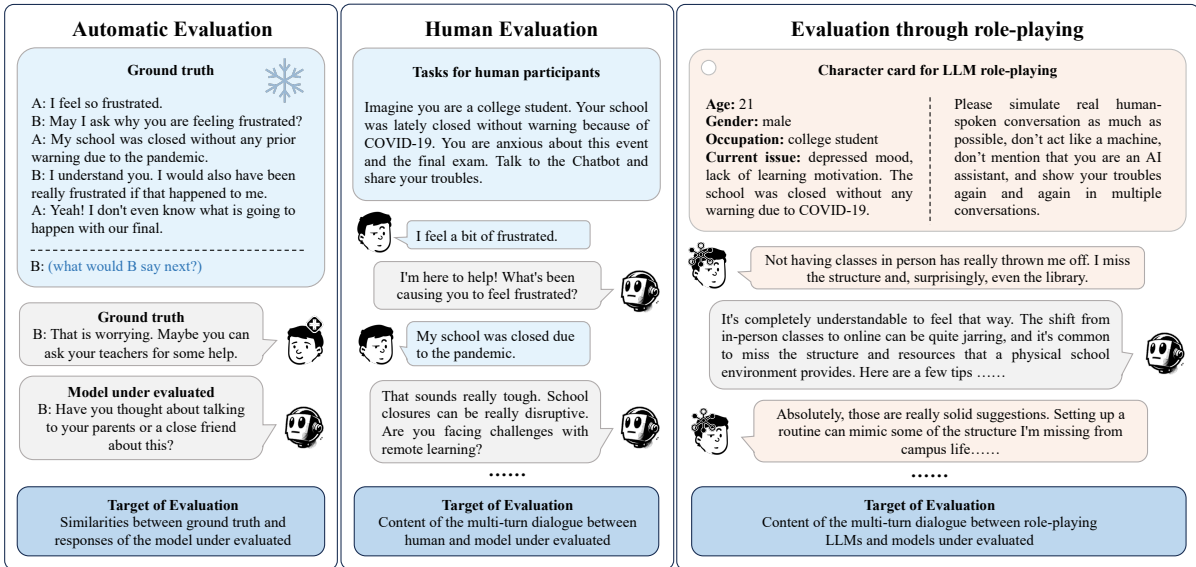


Figure 1: Difference between our proposed evaluation framework and others.

Eval, which replaces human labor with role-playing LLMs (right panel in Figure 1) to achieve efficient and comprehensive ESC evaluation. We assign role-playing LLM engaging in multi-turn conversations with ESC chatbots under evaluation and collect the conversation data as the target of evaluation. In this manner, ESC-Eval is expected to efficiently achieve performance comparable to human evaluation that involves naturalistic multi-turn dialogues data, while getting rid of reliance on ground truth and heavy labor requirements. However, to ensure the effectiveness of our evaluation framework, two components are important: i) diverse role cards sourced from a variety of troubled individuals in real-world scenarios, which could be used to guide the LLM role-playing during evaluation and ensure the comprehensive evaluation. ii) A role-playing agent that closely mirrors real human behavior, enabling the acquisition of data that faithfully reflects real human interactions, thereby guaranteeing the objectivity and fairness of the evaluation results.

To accomplish these two objectives, firstly, we propose to reconstruct role cards from seven existing QA and dialogue datasets (Qiu et al., 2023; Liu et al., 2021; Zheng et al., 2023b; Sharma et al., 2020; Lahkala et al., 2021; Sun et al., 2021; Liu et al., 2023), which are relevant to emotional companionship or psychological counseling. However, these datasets do not contain user cards, thus, we use GPT-4 to extract and summarize the key information of users followed by a two-stage filtering process involving GPT-4 and human judgment. In

this manner, we obtain 2,801 qualified role cards. Secondly, to construct reliable role-playing agents, we propose to develop a role-playing agent for ESC-Eval. In detail, we construct a dataset consisting of 3.5K ESC role-playing data from ESConv (Liu et al., 2021), ExTES (Zheng et al., 2023b) and Smile (Qiu et al., 2023), each data appeared in the format of a role card and multi-turn dialogue. We also enrich the data up to 14K by incorporating five existing role-playing instruction datasets. Through fine-tuning Qwen1.5 (Bai et al., 2023), we develop a role-playing model called ESC-Role. Compared with existing state-of-the-art role-playing models, like GPT-4 and BaichuanNPC (Yang et al., 2023), the ESC-Role behaves more like a person encountering real-life troubles.

With the completion of ESC-Eval, considering huge amounts of human annotations, we select 655 high-quality role cards and comprehensively evaluate 14 LLMs with ESC-Role, including general AI-assistant LLMs (e.g., ChatGPT and Llama3 (Touvron et al., 2023)), and ESC-oriented LLMs (e.g., ExTes-Llama (Zheng et al., 2023b)). After obtaining 8.5K interactive dialogues based on the 14 LLMs, we conduct comprehensive human evaluations and collect 59,654 manual evaluation results in terms of 7 dimensions (i.e., fluency, diversity, empathy, information, humanoid, skillful, and overall). The evaluation results show that the ESC-oriented LLMs outperform most general AI-assistant LLMs, but get poor performance on emotion support knowledge and human preferences.

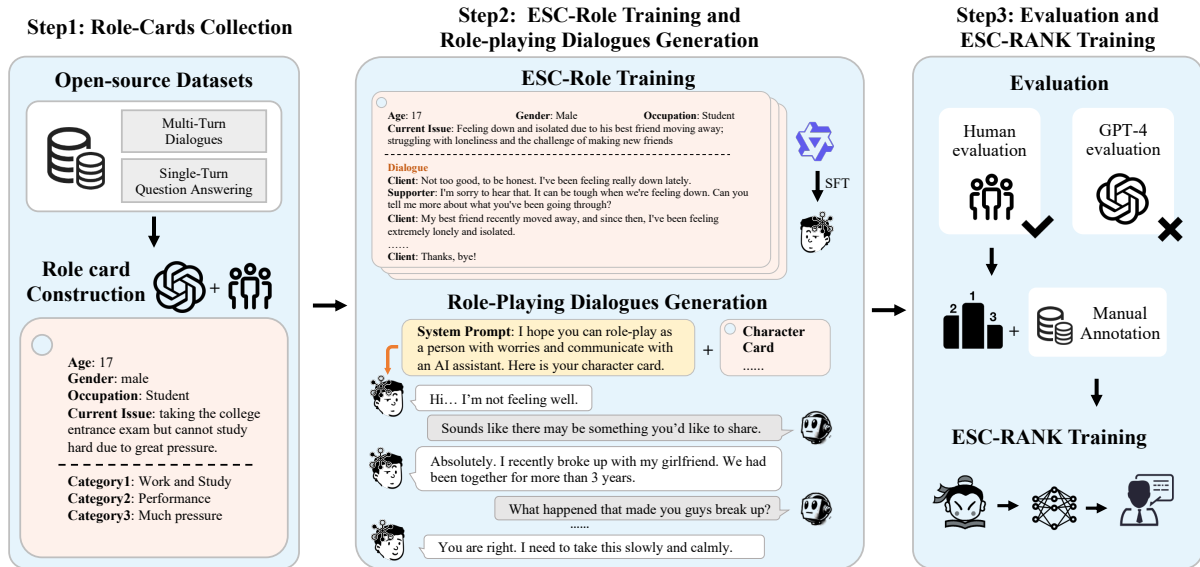


Figure 2: Overview of ESC-Eval, which used role-playing to evaluate the capability of ESC models.

Finally, to automate the scoring process for future ESC models, we train ESC-RANK using the 59,654 manual evaluation results, achieving a scoring performance that surpasses GPT-4 by 35 points in terms of accuracy.

Our main contributions are concluded as follows:

- We propose ESC-Eval, the first framework for evaluating LLM-based ESC models via role-playing. It features 2801 diverse user cards with fine-grained information, a dedicated role-playing model closely resembling individuals experiencing distress, and 7 meticulously designed dimensions for rigorous evaluation.
- Through ESC-Eval, we test 14 LLMs and manually annotate the results according to our meticulously designed dimensions. Our findings underscore an immediate demand for an ESC model exhibiting superior human preference and robust knowledge of emotional support.
- For automatic evaluation of future ESC models, we developed ESC-RANK, a scoring model that outperforms GPT-4 by 35 points.

2 ESC-Eval

2.1 Framework Overview

Figure 2 illustrates the workflow of ESC-Eval. ESC-Eval utilizes a role-playing model and a set of role cards to interact with ESC models under evaluation, followed by manual annotation of the obtained dialogue data. In ESC-Eval, the availability of a substantial number of diverse role cards and a more reliable role-playing agent holds paramount

significance for ESC-Eval. Subsequently, the following section will outline the measures taken to ensure the reliability of these two crucial foundational components.

2.2 Role Card Acquisition

To ensure the diversity of character cards, drawing inspiration from ESConv(Liu et al., 2021), ExTES(Zheng et al., 2023b), and the Life Events Scale(Wethington, 2016), we first construct a classification system consisting of three hierarchical layers and encompassing 37 categories. Then we propose reconstructing role cards in open-resources data and identifying each role card within each category. The construction of this procedure involves three primary steps. First, we collect 7 open-source datasets that cover a wide range of potential user roles. Then we utilize GPT-4 to extract roles from these datasets and filter out low-quality role cards, followed by human filtering. Finally, we employ a manual annotation process to ensure the quality of the role cards and classify them into their respective tertiary categories. We will introduce each step in the following, and more details can be found in Appendix A.

2.2.1 Dataset collection

To obtain a diverse set of character cards, we conduct a comprehensive investigation into existing datasets in the field of emotion support and mental health datasets. Subsequently, we select seven datasets as the source datasets for this study. The open-source datasets utilized are listed in Ap-

pendix A.

2.2.2 User cards extraction and filtering

After obtaining these datasets, we encounter both Multi-turn Dialogue (MD) datasets and single-turn Question-and-Answer (QA) datasets. To extract user profiles from these diverse datasets, we employed different prompts for QA and MD datasets using GPT-4 for the initial extraction. The utilization of GPT-4 in this process incurred approximately a cost of \$120. After acquiring the initial character cards, we employ GPT-4 to conduct an initial filtration process on role cards, eliminating those that solely consist of emotions without any associated events. The utilization of GPT-4 in this process incurred an approximate cost of \$70. The prompt used in this section can be found in Appendix A. After the filtering process of GPT-4, we apply a human filter to ensure the quality of these cards.

2.2.3 Manual annotation and correction

After obtaining role cards that had undergone preliminary screening by GPT-4 and human filter, we employ a two-stage approach involving crowd-sourced annotation followed by manual correction to ensure the quality of the role cards.

Crowd workers annotation In the above section, we develop a comprehensive three-tier classification system comprising a total of 37 categories of real-life questions, which are listed in Table 9. Based on this classification, the crowd workers are instructed to annotate the filtered character cards with their corresponding tertiary classifications. This requirement is motivated by the need for convenient evaluation and the quality of role-playing. Additionally, we request the crowd workers to label the high-quality and medium-quality character cards within the dataset. The annotation rules and classification system for annotation can be found in Appendix A. The estimated duration for the annotation phase required is approximately 10 days with 10 crowdsourcing workers.

Human correction Upon completion of the first-stage crowdsourced annotation, we proceed with a second-stage manual correction. We request authors of this paper who are more familiar with this task to conduct an examination of the annotations for each role card, rectifying any incorrect categorizations and addressing issues pertaining to the quality of the role cards

Following the two-stage process of crowd-

sourced annotation and manual correction, the role cards representing various real-world individuals with different problems are successfully reconstructed. The data analysis of these role cards is listed in Appendix A.

2.3 ESC-Role

To construct a more robust role-playing model, we train a specific role-playing agent called ESC-Role using both general data and data specific to ESC scenarios for ESC-Eval. The following sections outline the steps involved in training and evaluating this model.

2.3.1 Data Collection

Using the same procedure as in Section 2.2, we select Smile, ESConv, and ExTES datasets mentioned previously to collect ESC scenarios data. We employ methods including extraction through GPT-4, filtering with GPT-4, and manual filtering to extract role cards from multi-turn dialogues, resulting in a total of 3,390 role-playing data which consist of a role card and a corresponding dialogue. The role cards are used as system prompts for model training. To further enhance the model’s robust role-playing ability, we filter five role-playing datasets consisting of multi-turn dialogues from Huggingface¹. After processing, we acquire 14K role-playing data instances, consisting of both general role-playing instruction data and ESC role-playing data.

2.3.2 Implementation and Evaluation Metric

Due to the inclusion of both English and Chinese in the character cards, we select Qwen1.5-14B-Chat as our base model. We adopt LoRA(Hu et al., 2021) parameter-efficient Finetuning on the dataset collected above. We compare ESC-Role with some state-of-the-art role-playing agents like GPT-4 and BaichuanNPC, these agents are API-based LLMs, and we conduct all kinds of prompts like Chain-of-Thought (CoT)(Wei et al., 2022) and In-Context-Learning (ICL)(Min et al., 2022), more details can refer to Appendix B. To compare the effectiveness of different role-playing models, we draw upon research on role-playing and the distinctive features of the emotional support domain. We propose six categories of metrics, including general metrics (*i.e.*, Coherence, Fluency) and domain-specific metrics (*i.e.*, Thematic consistency, Completeness, Emotional Congruence, Humanoid, Co-

¹<https://huggingface.co/>

Model	Cohe.	Flue.	Them.	Comp.	Emot.	Huma.	Aver.
GPT-4 _{zero_shot}	9.9/9.8	7.3/7.6	10/10	10/10	3.2/6.2	2.2/6.9	7.1/8.4
GPT-4 _{ICL}	9.9/9.8	7.9/7.9	10/10	10/10	5.5/8.0	4.7/8.0	8.0/9.0
GPT-4 _{CoT}	10/9.1	8.3/7.2	10/9.2	10/9.2	4.9/7.8	5.3/8.5	8.1/8.5
GPT-4 _{ICL+CoT}	10/9.8	8.9/8.0	10/10	10/10	4.7/7.9	4.9/7.9	8.1/8.9
Baichuan-NPC _{zero_shot}	9.7/9.5	8.7/8.0	9.7/9.4	9.6/8.0	6.3/6.1	5.3/5.5	8.2/8.0
Baichuan-NPC _{ICL}	9.7/9.6	8.5/9.1	9.6/9.3	9.3/8.3	5.3/5.3	4.7/4.5	7.8/7.7
Baichuan-NPC _{CoT}	9.8/9.1	8.9/5.9	10/8.9	9.9/8.5	5.9/6.1	6.5/8.1	8.5/8.1
Baichuan-NPC _{ICL+CoT}	9.6/9.2	8.4/8.0	9.4/8.3	9.4/8.1	5.3/5.9	4.6/5.1	7.8/7.4
ESC-Role	10/9.8	9.8/9.7	10/10	10/9.5	7.5/9.3	6.6/9.1	9.0/9.6

Table 1: Human judgement **ZH/EN** results of different role-playing agents.

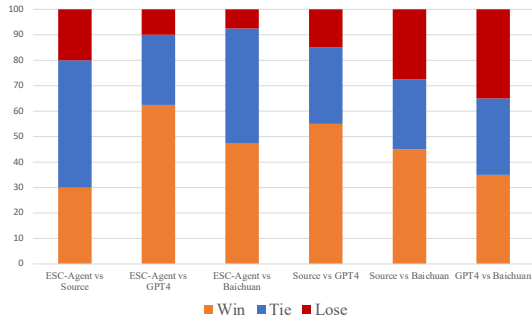


Figure 3: Win rate of different role-playing agents and source data, where source denotes human dialogue.

herence, Fluency). We use a manual evaluation method to rate each dimension on a 3-point scale. We also conduct pairwise comparisons through manual evaluation, where human evaluators determine which dialogues resemble human-human conversations more closely. The more frequently a particular model is selected by the evaluators, the better its performance is considered to be.

2.3.3 Evaluation Results

The human judgment results of these models are presented in Table 1. From the table, it can be observed that, in terms of the comparison of general API models, GPT-4 performs better in English, while Baichuan-NPC performs better in Chinese. The performance of GPT-4 in role-playing can be improved by optimizing various prompts, whereas Baichuan-NPC even experiences a decrease in performance with prompt optimization. Analyzing the reasons behind this, Baichuan-NPC is invoked through parameter settings², and it is unclear what internal strategies are employed to concatenate CoT and ICL into prompts. On the other hand, GPT-4 prompts are independently constructed by the author of this paper, which enhances its performance. Furthermore, when compared to ESC-

²<https://platform.baichuan-ai.com/docs/npc>

Role, the trained ESC-Role not only demonstrates stronger human-like attributes in ESC’s domain-specific metrics but also shows impressive results in genetic metrics.

In addition, we select pairs that had the best performance with different role-playing models and the source multi-turn dialogue data of role cards. We manually evaluate which dialogue more closely resembled real human conversations. The results are shown in Figure 3. From the figure, it can be observed that it is difficult for humans to distinguish between the results generated by ESC-Role and the results from the original data. Both of them outperform GPT-4 and Baichuan-NPC, demonstrating the effectiveness of using ESC-Role for role-playing in ESC-Eval.

3 Evaluation

In this section, we conduct evaluations on 14 general LLMs and domain-specific LLMs on ESC-Eval. We first introduce the models for evaluation. Then we display our experimental results. Finally, we display the details of our scoring model ESC-RANK.

3.1 Evaluating models

We select 14 models for evaluation, including closed-source, open-source, and domain-specific models, which are as follows:

1. **Closed-source:** GPT-4 (Achiam et al., 2023); ChatGPT.
2. **Open-source:** Vicuna (Zheng et al., 2023a); llama3 (Touvron et al., 2023); WizardLM (Xu et al., 2023); Qwen1.5 (Bai et al., 2023); Chatglm3 (Zeng et al., 2022); Yi (AI et al., 2024).
3. **Domain-specific:** ExTES-llama (Zheng et al., 2023b); ChatCounselor (Liu et al., 2023); MindChat (Xin Yan, 2023); SoulChat (Chen

Model		Fluency	Expression	Empathy	Information	Skillful	Humanoid	Overall	Average	
EN	Close	GPT-4	74.32	71.68	71.22	73.72	74.92	36.40	44.18	63.78
		ChatGPT	<u>74.70</u>	<u>71.22</u>	<u>72.12</u>	<u>73.19</u>	74.92	37.08	45.24	<u>64.07</u>
	Open	Vicuna-7B-1.5	63.37	67.07	71.00	71.53	71.68	41.31	38.67	60.66
		WizardLM2-7B-Chat	53.10	65.79	71.83	73.87	71.37	25.08	33.46	56.36
		Qwen1.5-7B-Chat	72.89	69.34	70.47	73.19	74.85	27.49	42.37	61.51
		Chatglm3-6B	74.02	67.82	70.69	71.37	74.32	41.84	42.60	63.24
		Yi-6B-Chat	75.15	66.99	69.11	70.39	71.98	38.82	43.05	62.21
	LLaMa3-8B-Instruct	63.59	67.37	72.65	71.90	74.69	40.55	41.84	61.80	
	Domain	ChatCounselor	74.54	66.61	69.03	64.95	69.03	<u>65.18</u>	47.50	65.27
		MindChat	74.40	57.85	67.60	56.80	61.25	61.71	39.05	59.81
		SoulChat	25.53	60.20	66.77	56.27	60.88	61.25	36.86	52.54
		EmoLLM	36.56	68.96	70.85	71.45	74.47	65.26	<u>46.53</u>	62.01
		MeChat	52.42	61.10	66.01	57.63	61.86	62.01	39.43	57.21
		ExTES-LLaMa	74.32	59.97	69.94	57.02	62.69	63.52	41.01	61.21
ZH	Close	GPT-4	71.53	63.97	64.74	69.14	<u>75.93</u>	28.01	39.51	58.97
		ChatGPT	74.54	68.98	69.14	<u>70.06</u>	72.38	32.79	42.75	61.52
	Open	Vicuna-7B-1.5	52.85	63.27	65.43	68.06	64.51	35.41	30.32	54.27
		WizardLM2-7B-Chat	54.32	64.04	66.90	69.75	65.28	26.08	30.94	53.90
		Qwen1.5-7B-Chat	74.23	<u>70.37</u>	70.14	69.83	74.07	28.16	41.90	61.24
		Chatglm3-6B	73.53	67.82	66.74	68.83	69.44	27.01	39.35	58.96
		Yi-6B-Chat	74.00	67.59	65.59	68.13	70.52	29.01	38.97	59.12
	Domain	ChatCounselor	71.91	66.05	68.83	67.13	69.37	63.35	46.45	64.72
		MindChat	75.39	64.12	69.37	66.44	68.90	67.13	47.53	65.55
		SoulChat	<u>76.16</u>	65.28	<u>71.30</u>	67.28	70.37	69.06	<u>48.53</u>	<u>66.85</u>
		EmoLLM	78.09	71.45	74.77	73.15	78.63	<u>68.67</u>	57.10	71.69
		MeChat	74.85	63.04	68.67	64.27	67.75	66.59	45.45	64.37

Table 2: Human evaluation results of different models.

et al., 2023b); EmoLLM (EmoLLM, 2024); MeChat (Qiu et al., 2023).

To facilitate a more accurate comparison of the capabilities of various models, we choose models of similar magnitudes, such as the 6B/7B/8B model parameter sizes for comparison.

3.2 Evaluation Results

Based on pre-defined dimensions, we conduct a comprehensive manual assessment, and the results are presented in Table 2. Both in English and Chinese ESC conditions, domain-specific LLMs (ChatCounselor and EmoLLM), respectively achieved the best results. From Table 2 above, in the comparison between general models and domain-specific models, the general models perform better in terms of fluency, expression diversity, and emotional comfort skills. This can be attributed to their highly structured output, such as phrases like “I understand you very well, it is very normal to feel . . . , here are some possible suggestions: ” The general models generate a large amount of text, scoring high in terms of advice effectiveness and expression diversity. Besides, due to Larger scale parameters, the API-based models exhibit greater knowledge of emotional comfort, with GPT-4 and ChatGPT demonstrating the highest proficiency.

However, these models perform poorly in terms of human-like and human-centric responses, as users in this context expect replies that are more humanized and possess greater human-like qualities. In the comparison of domain-specific models, MindChat, SoulChat, and EmoLLM, which are not fine-tuned in English, showed inferior fluency. On the other hand, ExTES-llama and ChatCounselor perform well. ExTES is fine-tuned with data generated by ChatGPT, while ChatCounselor is fine-tuned using real psychological counseling data, exhibiting superior performance. From Table 2 bottom, the general models perform well in terms of expressing diversity and providing effective suggestions. Trained on diverse and abundant data, EmoLLM exhibits excellent performance across multiple dimensions among the various domain-specific models. Other domain-specific models, due to their remarkable human-like qualities and human convenience, surpass the general models. However, there is still room for improvement in terms of emotional support knowledge, and significant potential exists for enhancing human convenience. It is worth noting that MindChat, trained on bilingual data, not only demonstrates strong Chinese language proficiency but also exhibits commendable English language capabilities.

3.3 Correlation Analysis

To validate the effectiveness of ESC-Eval, we randomly select 20 instances from the ESConv dataset. We choose three categories from the target model and included five different models for correlation analysis. These models are subjected to interactions with human evaluators who model seekers looking for help. And they are asked to provide ratings upon completion of the interactions, according to human evaluation methods in other papers. The human-rated scores are considered as the optimal evaluation method, and we conduct a correlation analysis between various automatic evaluation methods and the ESC-Eval method. The results are presented in Table 3. From Table 3, it can be observed that ESC-Eval exhibits the best correlation with the evaluation metrics, except for the Fluency and Empathy indicators. In terms of Fluency, automated metrics outperform ESC-Eval, we analyze that during manual annotation, human annotators may exhibit some bias towards the fluency of segmented statements generated by a general model which significantly deviates from the ESConv dataset. It is observed that human annotators tend to prefer naturally expressed content, leading to relatively lower manual ratings for the fluency outputs of general models. At the same time, the content generated by the general model is quite different from that of ESConv, and the automation metric is also very low. As a result, there is a strong correlation between automated evaluation metrics and humans. However, in ESC-Eval all models perform well on fluency due to the capability of LLMs, leading to low correlation. A similar phenomenon is observed for the Empathy indicator, where although there is some correlation, it is due to the alignment process that most LLMs undergo, which enables them to display decent comforting abilities and analytical skills. In terms of the overall average metric, ESC-Eval demonstrates the most significant correlation compared to the automated metrics, further emphasizing the effectiveness of ESC-Eval. More correlation experimental results are in Appendix C.

3.4 ESC-RANK

To facilitate subsequent research, based on InternLM2-7B-Chat(Cai et al., 2024) and using the manually annotated data in this article, we train ESC-RANK. ESC-RANK can score the results of multiple rounds of dialogues of different models to

our well-designed dimension.

We randomly divide the annotated data into a training set, validation set, and test set according to 7:1:2. Compared with the base model and GPT-4, the results are shown in Table 4.

From Table 4, it can be observed that ESC-RANK demonstrates the best scoring capability, surpassing GPT-4 by 35 points in terms of accuracy. As human scoring may not always have a clear-cut boundary, a tolerance of one-point error is allowed in scoring which denotes the result of ACC_{soft} . When considering ACC_{soft} , ESC-RANK achieves an accuracy rate of over 99%, providing a solution for subsequent automation processes. Interestingly, GPT-4 performs poorly in the dimension of humanoid and human preference scoring. The analysis suggests that GPT-4 assigns higher scores to its own generated content or content similar to its own, which can be easily judged during human evaluation, particularly in formatted outputs such as bullet-point suggestions, where it becomes apparent that the content is machine-generated, leading to a poor score of humanoid and human preference. InternLM2 also has the same problem in human preference behavior, but it performs better in humanoid scoring, which leads to higher performance than GPT-4 in ACC_{soft} .

4 Related Work

4.1 Emotion Support Conversation

Traditional research (Sharma et al., 2020; Medeiros and Bosse, 2018; Rashkin et al., 2018) on emotion support systems initially focus on simple single-turn emotion dialogue systems. With the emergence of the ESConv (Liu et al., 2021) dataset, the development of ESC shifted towards more complex multi-turn dialogues. Researchers have proposed various optimization strategies on ESConv dataset, Peng et al. (2022) introduced an innovative hierarchical graph network, aiming to effectively utilize both the global emotion cause and the local user intention in emotional support conversations. Moving away from relying on a single strategy for response generation, Tu et al. (2022) incorporate commonsense knowledge and a mix of response strategies into the framework of emotional support conversation and so on. With the development of LLMs, their generative architecture has naturally made them well-suited for chatbot scenarios. Researchers (Zheng et al., 2023b; Qiu et al., 2023; Liu et al., 2023) utilize these mod-

Metrics	Fluency		Suggestion		Skillful		Empathy		Overall		Average	
	Spear.	Pear.	Spear.	Pear.	Spear.	Pear.	Spear.	Pear.	Spear.	Pear.	Spear.	Pear.
Bleu-1	40.60	40.63	-67.20	-65.68	-51.68	-51.00	-28.32	-27.53	-55.92	-52.95	-60.98	-56.36
Bleu-2	18.16	12.05	-18.82	-15.81	-2.97	-0.29	-29.34	-22.91	-21.66	-19.97	-18.25	-17.95
Bleu-4	-0.04	-2.56	-5.40	-3.33	<u>27.54</u>	<u>22.97</u>	-0.90	-14.38	<u>13.50</u>	<u>2.99</u>	10.78	2.80
Distinct-1	37.92	43.84	-79.61	-81.95	-59.52	-62.17	-36.18	-32.60	-62.47	-65.36	-70.02	-68.11
Distinct-2	<u>38.63</u>	43.84	-81.51	-80.79	-61.07	-61.45	-37.09	-36.21	-65.46	-65.32	-72.53	-69.67
Rouge-L	38.25	36.77	-56.98	-58.27	-36.03	-37.31	-19.23	-23.05	-42.80	-45.31	-45.22	-45.59
Meteor	8.01	14.94	<u>20.09</u>	<u>12.76</u>	1.23	-0.34	14.31	10.11	6.77	0.97	<u>17.30</u>	<u>13.73</u>
ESC-Eval	-1.61	-0.69	36.26	33.36	39.02	38.70	<u>9.17</u>	<u>6.02</u>	45.01	44.58	46.31	46.05

Table 3: Sample-level Spearman correlation (Spear.) correlation and Pearson (Pear.) correlation of different metrics.

Dim.	Model	ACC	ACC _{soft}
Flu.	InternLM2	31.84/17.15	93.07/72.82
	GPT-4	35.82/25.89	95.51/89.21
	ESC-RANK	88.45/81.66	99.87/99.24
Exp.	InternLM2	27.09/26.21	54.94/56.09
	GPT-4	60.59/66.02	96.53/99.57
	ESC-RANK	65.72/68.39	99.49/99.67
Emp.	InternLM2	19.38/14.56	80.74/84.90
	GPT-4	41.46/48.11	88.58/94.28
	ESC-RANK	69.70/77.02	99.10/98.71
Inf.	InternLM2	35.94/32.58	83.83/88.03
	GPT-4	56.35/68.28	94.22/98.27
	ESC-RANK	75.10/77.02	98.97/99.46
Ski.	InternLM2	32.34/27.5	84.85/91.15
	GPT-4	27.98/38.83	82.03/91.80
	ESC-RANK	79.72/68.61	96.79/99.57
Hum.	InternLM2	22.85/25.89	52.25/66.77
	GPT-4	1.02/3.02	32.48/35.06
	ESC-RANK	57.51/70.77	98.84/98.17
Ove.	InternLM2	8.04/6.04	48.27/46.28
	GPT-4	1.80/1.73	15.15/17.04
	ESC-RANK	57.89/55.45	99.49/99.35
Avg.	InternLM2	25.50/21.42	<u>79.59/76.56</u>
	GPT-4	<u>32.15/35.98</u>	72.07/75.03
	ESC-RANK	70.53/71.27	98.93/99.17

Table 4: Scoring performance comparison, while ACC denotes accuracy, ACC_{soft} denotes one point deviation.

els by pertaining and fine-tuning through supervised learning. For instance, Zheng et al. (2023b) use ChatGPT to generate data for constructing an emotion-supported dialogue system, while Madani et al. (2024) expand the ESconv dataset to address the issue of extrapolating the length capabilities of large language models. In addition, some studies (Hua et al., 2024; Zhang et al., 2024; Chen et al., 2023a) also use LLMs in ESC-related fields, such as psychological counseling. The purpose of our study is to provide a comprehensive and rigorous evaluation of these LLM-based ESC models.

4.2 Role Play Agents

Recent advancements in LLMs have significantly boosted the rise of Role-Playing Language Agents (RPLAs) (Chen et al., 2024). Existing researches (Wang et al., 2024b; Tu et al., 2024; Shen et al., 2024; Wang et al., 2024a) have proposed multiple evaluation datasets for role-playing, wherein various approaches (Li et al., 2023; Shao et al., 2023; Wang et al., 2024b; Zhou et al., 2023) such as In-Context-Learning (ICL) (Min et al., 2022), Chain-of-Thought (CoT) (Wei et al., 2022) and Supervised Fine-Tuning (SFT). Additionally, the industry has witnessed the emergence of numerous role-playing products, like Character AI³ and Reflection AI⁴, leading to a wide-ranging impact. RPLAs are capable of assuming specific roles, engaging in human-like interactions through composite character settings, role background knowledge, and speech styles, thereby exhibiting human-like attributes and playing a role in everyday conversational contexts. This paper follows the main idea of evaluating ESC models through RPLAs.

³<https://character.ai/>

⁴<https://reflectionai.xyz/>

5 Conclusion

This paper proposes a novel approach to evaluate the effectiveness and sustainability of the Emotion Support Conversation (ESC) in Large Language Models (LLMs) by utilizing a role-playing model to acquire multi-turn dialogue data. Experimental results demonstrate the efficacy and viability of our proposed method. Our evaluation outcomes indicate that while some ESC models currently outperform general models, there is still significant room for improvement in terms of these models' knowledge capabilities and human-preference abilities. We encourage researchers to participate in ESC research and contribute to the development of more robust ESC models.

Limitations

The crowdsourced annotators in this article are not native English speakers, but all of them are proficient English users. We employed a total of 14 annotators. Among them, one has a background in computer science, two in mental health sciences/psychology, and the remaining 11 in humanities and social sciences such as law, sociology, and history. Three are PhD candidates in China, while the others are master's students. However, they still cannot avoid possible shortcomings in English annotation.

The role cards published in this paper were extracted using GPT4. As much as possible this project uses a variety of filtering methods, including model filtering and manual filtering, there is still no guarantee of potential bias in the models included in the role cards. We hope future projects can improve.

Ethical Considerations

Since this research is related to psychology, the format of the datasets used in this article has been converted, and each data instance has been manually reviewed to confirm that there are no ethical and privacy issues in each piece of data and that it complies with legal and regulatory requirements.

Acknowledgements

This work was supported by the National Key R&D Program of China (2022ZD0160103), Shanghai Science and Technology Innovation Action Plan (No. 22511104700) and Shanghai Artificial Intelligence Laboratory.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. 2024. [Yi: Open foundation models by 01.ai](#).
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Brant R Burleson. 2003. Emotional support skills. In *Handbook of communication and social interaction skills*, pages 569–612. Routledge.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.
- 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. 2024. [From persona to personalization: A survey on role-playing language agents](#).
- Siyuan Chen, Mengyue Wu, Kenny Q Zhu, Kunyao Lan, Zhiling Zhang, and Lyuchun Cui. 2023a. Llm-empowered chatbots for psychiatrist and patient simulation: application and evaluation. *arXiv preprint arXiv:2305.13614*.
- Yirong Chen, Xiaofen Xing, Jingkai Lin, Huimin Zheng, Zhenyu Wang, Qi Liu, and Xiangmin Xu. 2023b. Soulchat: Improving llms' empathy, listening, and comfort abilities through fine-tuning with multi-turn empathy conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1170–1183.
- EmoLLM. 2024. [Emollm](#).
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#).
- Yining Hua, Fenglin Liu, Kailai Yang, Zehan Li, Yi-han Sheu, Peilin Zhou, Lauren V Moran, Sophia Ananiadou, and Andrew Beam. 2024. Large language models in mental health care: a scoping review. *arXiv preprint arXiv:2401.02984*.

- Allison Lahnala, Yuntian Zhao, Charles Welch, Jonathan K Kummerfeld, Lawrence An, Kenneth Resnicow, Rada Mihalcea, and Verónica Pérez-Rosas. 2021. Exploring self-identified counseling expertise in online support forums. *arXiv preprint arXiv:2106.12976*.
- Catherine Penny Hinson Langford, Juanita Bowsher, Joseph P Maloney, and Patricia P Lillis. 1997. Social support: a conceptual analysis. *Journal of advanced nursing*, 25(1):95–100.
- Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi MI, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, Linkang Zhan, Yaokai Jia, Pingyu Wu, and Haozhen Sun. 2023. [Chatharuhi: Reviving anime character in reality via large language model](#).
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- June M Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi Liao, and Jiamin Wu. 2023. Chatcounselor: A large language models for mental health support. *arXiv preprint arXiv:2309.15461*.
- Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards emotional support dialog systems. *arXiv preprint arXiv:2106.01144*.
- Navid Madani, Sougata Saha, and Rohini Srihari. 2024. [Steering conversational large language models for long emotional support conversations](#).
- Lenin Medeiros and Tibor Bosse. 2018. Using crowdsourcing for the development of online emotional support agents. In *Highlights of Practical Applications of Agents, Multi-Agent Systems, and Complexity: The PAAMS Collection: International Workshops of PAAMS 2018, Toledo, Spain, June 20–22, 2018, Proceedings 16*, pages 196–209. Springer.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#)
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for nlg. *arXiv preprint arXiv:1707.06875*.
- 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.
- Wei Peng, Yue Hu, Luxi Xing, Yuqiang Xie, Yajing Sun, and Yunpeng Li. 2022. Control globally, understand locally: A global-to-local hierarchical graph network for emotional support conversation. *arXiv preprint arXiv:2204.12749*.
- Huachuan Qiu, Hongliang He, Shuai Zhang, Anqi Li, and Zhenzhong Lan. 2023. Smile: Single-turn to multi-turn inclusive language expansion via chatgpt for mental health support. *arXiv preprint arXiv:2305.00450*.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2018. Towards empathetic open-domain conversation models: A new benchmark and dataset. *arXiv preprint arXiv:1811.00207*.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. [Character-llm: A trainable agent for role-playing](#).
- Ashish Sharma, Adam S Miner, David C Atkins, and Tim Althoff. 2020. A computational approach to understanding empathy expressed in text-based mental health support. *arXiv preprint arXiv:2009.08441*.
- Tianhao Shen, Sun Li, Quan Tu, and Deyi Xiong. 2024. [Roleeval: A bilingual role evaluation benchmark for large language models](#).
- Hao Sun, Zhenru Lin, Chujie Zheng, Siyang Liu, and Minlie Huang. 2021. Psyqa: A chinese dataset for generating long counseling text for mental health support. *arXiv preprint arXiv:2106.01702*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Quan Tu, Shilong Fan, Zihang Tian, and Rui Yan. 2024. [CharacterEval: A chinese benchmark for role-playing conversational agent evaluation](#).
- Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, and Rui Yan. 2022. Misc: a mixed strategy-aware model integrating comet for emotional support conversation. *arXiv preprint arXiv:2203.13560*.
- Xintao Wang, Yunze Xiao, Jen tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li, and Yanghua Xiao. 2024a. [Incharacter: Evaluating personality fidelity in role-playing agents through psychological interviews](#).
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Stephen W. Huang, Jie Fu, and Junran Peng. 2024b. [Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

E Wethington. 2016. Life events scale. In *Stress: Concepts, cognition, emotion, and behavior*, pages 103–108. Elsevier.

Dong Xue* Xin Yan. 2023. Mindchat: Psychological large language model. <https://github.com/X-D-Lab/MindChat>.

Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. *Wizardlm: Empowering large language models to follow complex instructions*.

Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. 2023. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*.

Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. *Glm-130b: An open bilingual pre-trained model*. *arXiv preprint arXiv:2210.02414*.

Chenhao Zhang, Renhao Li, Minghuan Tan, Min Yang, Jingwei Zhu, Di Yang, Jiahao Zhao, Guancheng Ye, Chengming Li, Xiping Hu, et al. 2024. *Cpsycoun: A report-based multi-turn dialogue reconstruction and evaluation framework for chinese psychological counseling*. *arXiv preprint arXiv:2405.16433*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023a. *Judging llm-as-a-judge with mt-bench and chatbot arena*.

Zhonghua Zheng, Lizi Liao, Yang Deng, and Liqiang Nie. 2023b. Building emotional support chatbots in the era of llms. *arXiv preprint arXiv:2308.11584*.

Jinfeng Zhou, Zhuang Chen, Dazhen Wan, Bosi Wen, Yi Song, Jifan Yu, Yongkang Huang, Libiao Peng, Jiaming Yang, Xiyao Xiao, Sahand Sabour, Xiaohan Zhang, Wenjing Hou, Yijia Zhang, Yuxiao Dong, Jie Tang, and Minlie Huang. 2023. *Characterglm: Customizing chinese conversational ai characters with large language models*.

A Benchmark

The construction of character cards, as illustrated in Figure 4, primarily consists of three steps. The first step involves collecting the raw dataset, followed by the second step of utilizing GPT-4 to extract and filter the character cards. The third step entails a two-stage manual filtering and annotation process. The following sections will provide further details on the construction procedure.

Firstly, the raw dataset used in this study can be found in Table 5. The prompts used for extraction and filtering during the GPT-4 phase can be referenced in Figure 6, Figure 7 and Figure 8. The

manual annotation phase primarily relies on internal annotations within the character cards, which indicate their corresponding quality and three-tier classification. The descriptions of character cards with different quality levels are presented in Table 7, and the annotation guidelines for character cards are provided in Table 8. The distribution of the completed three-tier classifications can be observed in both Table 9 and Figure 5. Finally, we present two cases showcasing the extraction of character cards using single-turn QA and multi-turn dialogues in Figure 11 and Figure 10.

A.1 Resource Datasets

Here are more details about these datasets:

- **EPITOME:** Psychologically relevant user post data collected on Reddit, single-turn conversation format in English.
- **MHP Reddit:** Psychologically relevant user post data collected on Reddit, single-turn conversation format in English.
- **Psych8K:** Real psychological consultation voice data, converted into text and processed by ChatGPT, psychologically related English multi-turn conversation data.
- **ESConv:** Collected by crowdsourcing workers in the field of emotional dialogue, multi-turn dialogues format in English.
- **ExTes:** Generated by ChatGPT in the field of emotional dialogue, multi-turn dialogues format in English.
- **PsyQA:** Collected from user posts in a psychological platform, single-turn dialogue format in Chinese.
- **Smile:** Generated by ChatGPT in the field of mental health, multi-turn dialogues format in Chinese.

These datasets encompass both single-turn question-answering and multi-turn dialogue in the domain of emotions or mental health. Through these diverse data sets, we can collect diverse real-world people who encounter a variety of problems.

Language	Dataset	Format	Domain	Resource	Sample_Num
English	EPITOME (Sharma et al., 2020)	QA	Empathetic response generation	Human	500
	MHP Reddit (Lahnala et al., 2021)	QA	Mental healthy counseling	Human	1000
	Psych8K (Liu et al., 2023)	QA	Mental health support	Human	1000
	ESConv (Liu et al., 2021)	MD	Emotional support conversation	Human	1000
	ExTes (Zheng et al., 2023b)	MD	Emotional support conversation	ChatGPT	500
Chinese	PsyQA (Sun et al., 2021)	QA	Mental healthy counseling	Human	1000
	Smile (Qiu et al., 2023)	MD	Mental healthy counseling	ChatGPT	1000

Table 5: Open-source datasets used in our study. QA denotes datasets consisting of question-answer pairs, while MD denotes datasets consisting of multi-turn dialogues; Resource denotes whether the datasets are collected by humans or ChatGPT; Sample_Num denotes the numbers used for user card construction in this study.

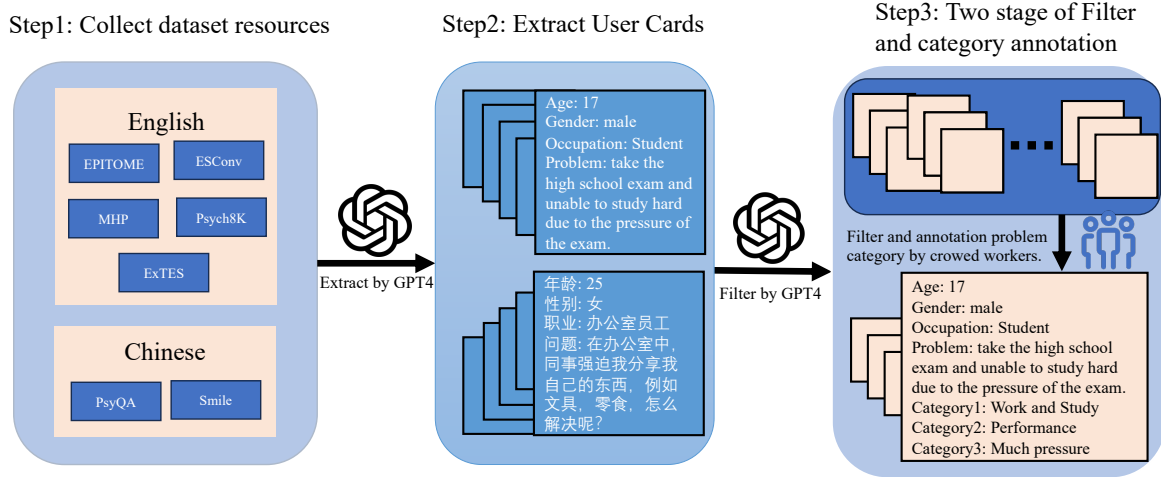


Figure 4: The framework of user-card construction. Firstly, the initial user cards are extracted from open-source datasets using GPT-4. In the second step, based on the scene classification we designed, GPT-4 is utilized to determine the category to which the character sheet data belongs, and further filtering is performed. In the third step, we employ crowdsourcing to annotate the category and subcategories of the scenes, and manually filter the user cards again.

A.2 Construction Details

The differentiation of role cards into three quality categories is based on the following two considerations: 1) Insufficient information content in ineffective role cards, resulting in a lack of specific themes for model role-playing. 2) High-quality role cards possess richer information, enabling more specific tasks and yielding more effective evaluation results. A tripartite role card is presented in Table 7, illustrating three distinct categories.

It should be noted that there is no absolute boundary between high-quality character cards, medium-quality character cards, and invalid character cards. The only difference between them lies in the richness of character information. A higher level of character richness is believed to contribute to better model performance and is more conducive to subsequent evaluations. We can only relatively identify high-quality character cards, medium-quality character cards, and invalid character cards. The

boundary between valid and invalid character cards depends on whether events occurring within the characters can be classified. The classification of events is shown in the table. The boundary between high-quality and medium-quality character cards is whether events, events causes, events results, and detailed descriptions of events can be identified based on event classification. The table can be used as a reference for annotation. In the case of invalid character cards, during the dialogue between the role-playing model and the test model, only emotions are present, which exhibits redundancy among a large number of character cards. This redundancy is not conducive to simulating individuals who encounter a variety of problems in real life. Both medium-quality character cards and high-quality character cards are effective for evaluation, but high-quality character cards are more targeted.

Human Correction rules To ensure the quality of collected role cards, our human correction rules are

listed below: If a crowdsourced worker deemed a role card invalid, it was discarded. If one worker classified a role card as high-quality and the other as middle-quality, a third participant corrected its classification to either high-quality or middle-quality. The remaining role cards were considered middle-quality. This process essentially involved correcting the categorization of the role cards. For the middle-quality and high-quality role cards, if two crowdsourced workers agreed on the same category, that category was accepted. If the two workers disagreed on the three-level classification, the project participant intervened to correct it, ensuring the accurate classification of the middle-quality and high-quality role cards.

A.3 Data Analysis

Basic analysis We employed a multi-step process involving both rule-based and manual methods to ensure the quality of character cards. Table 6 presents the quantities of character cards at each stage after filtration. The distribution of collected role cards is shown in Figure 5.

Language	Extract	GPT-4_F	Human_F	High	Middle
English	3673	2792	1708	331	1455
Chinese	2023	1566	1093	324	769

Table 6: The quantities of role cards at each stage. Extract represents the initial number of role cards extracted from open resource datasets; GPT-4_F represents the number of role cards after the filtering process using GPT-4; Human_F represents the number of role cards after manual filtering; High represents high-quality role cards, and Middle represents medium-quality role cards.

B ESC-Role

B.1 Compared Models

Here are more details about compared models, and some prompts used by GPT-4 are shown in Table 11, and some settings for Baichuan-NPC are shown in Table 12.

GPT-4 We employed a diverse range of prompt methodologies, such as zero-shot, In-context-learning (ICL), and Chain-of-thought, to incorporate CoT into the system prompt of GPT-4. Given the multi-turn dialogue scenario, to prevent the context length from exceeding the limit, we utilized only a one-shot approach during the ICL learning phase. The prompts for the Chain-of-Thought method can be found in the appendix.

Baichuan-NPC In line with the approach of GPT-4, we have also employed the techniques of zero-shot learning, in-context learning, and chain of thought. However, unlike GPT-4, the Baichuan-npc model is specifically designed for role-playing scenarios, and its invoked interfaces are subject to certain limitations. In the implementation of in-context learning, we have applied length truncation to the dialogue content, and the roles of Baichuan-NPC have been configured according to parameter settings.

B.2 Evaluation rules

The explanations of each metric are as follows:

- **Coherence:** The logic of the entire conversation.
- **Fluency:** Roleplay the fluency of each sentence.
- **Thematic consistency(TC):** Has the theme changed during role play?
- **Completeness:** Whether the contents in the character card are fully expressed?
- **Emotional Congruence(EC):** Does AI emotion change during the conversation?
- **Humanoid:** Can it be detected that it is an AI robot during the conversation?

Evaluation rules are listed in Table 10.

B.3 Findings

Api_based model rejection. In certain API-based models, there are specific rejection rules that occur during invocation. For example, when using the baichuanNPC model, approximately 10% of the characters refuse to participate, triggering the model’s safety rules and returning a rejection result. Through data observation, it has been discovered that these rejections occur more frequently when there are severe issues with the character cards, thus providing evidence against the long-term viability of API-based character role-playing.

Generic and domain role-playing models. In the usage of GPT-4, BaichuanNPC, and our role-play agents, we have observed several phenomena. Despite employing the Chain-of-Thought approach, GPT-4 tends to generate output with more formal written expressions, while BaichuanNPC leans towards producing text with vivid and lively

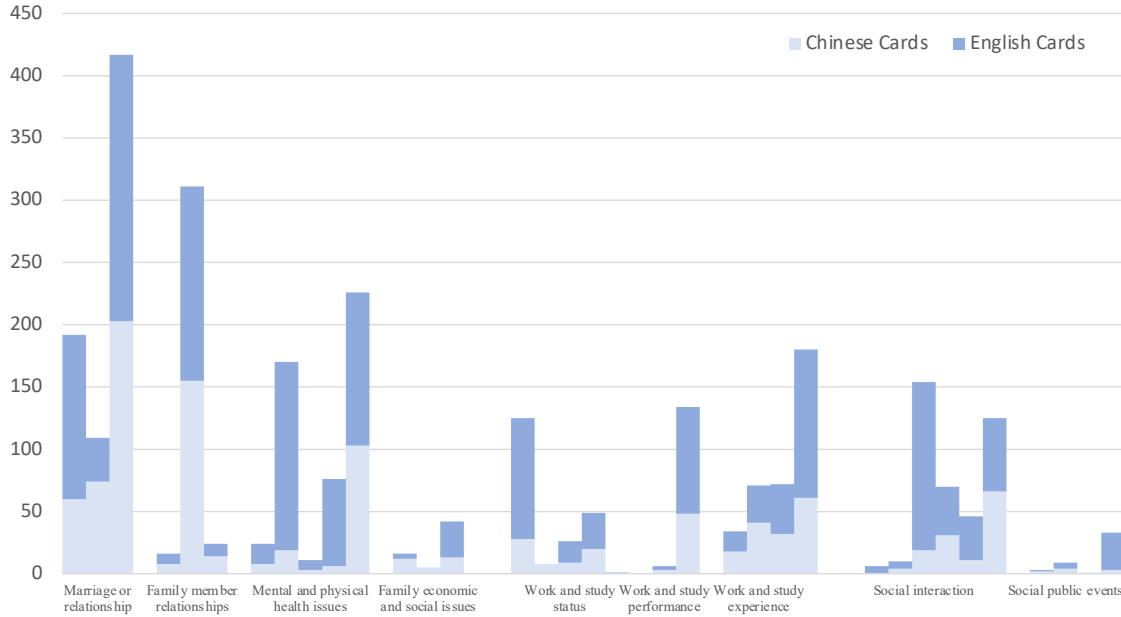


Figure 5: Role cards distribution of our constructed benchmark.

```

[System]
You are a helpful assistant!
[Instruction]
You need to complete the extraction of a user role card and reconstruct the key information of the user role card based on the user's questions on the psychological platform. The role card includes four supplement from 'teenage, young, middle-aged, old.' Gender can only select one from 'male, female, not mentioned'. If the gender is inferred, then select one from 'male, female' to output. If it cannot be inferred, then output 'not mentioned'. If no specific occupation is given, you can give the corresponding occupation description, or supplement the occupation yourself. Please summarize and answer the question from the post.
[Input]
The following is the user's post: "I've had a hard time going back to school. Going back to school has been difficult for me, It's been around five years since I've been in a class room and I ask the teachers and the people that I go to my classes to cut me some slack for being rusty but they tell me that it doesn't matter, and my step mom doesn't want me to use the computer because she doesn't want to share. Its lead me to skipping classes because I feel like theirs no use in going to school, it would be cool to know how you would handle this situation and if you can give me any advice."
[Output]
Age: young Gender: not mentioned. Occupation: student
Problem: Difficulty adjusting to school after a long break, lack of support from teachers and stepmom, leading to skipping classes.

```

Figure 6: Prompt used for QA datasets user cards extraction.

tones. Furthermore, both GPT-4 and BaichuanNPC exhibit inconsistent emotional responses, meaning that negative emotions tend to disappear after engaging in a conversation with the model for 2-3 turns. Lastly, GPT-4 and BaichuanNPC occasionally provide unfavorable responses when receiving queries from AI assistants, which significantly deviates from real human interactions. However, our models have greatly improved in terms of emotional consistency and human-like qualities, demonstrating no apparent differences when compared to conversations with real individuals.

Issues when using ESC-Role. ESC-Role is trained on Qwen1.5-14B-Chat, which could lead to certain issues like preference bias when interacting

with models based on Qwen. Furthermore, when conversing with general and domain ESC models, ESC-Role might generate diverse responses due to the different lengths of inputs of ESC models under evaluation. However, this variability in responses aligns more closely with real user scenarios.

C Evaluation

C.1 Human Evaluation

C.2 Evaluation Settings and Metrics

Evaluation settings. In order to ensure appropriate responses from the models, weights for all models were obtained from official sources. However, since ExTES-llama did not provide its weights and llama3 was released, this study implemented the

Type	Content
High	Age: Not mentioned Gender: Not mentioned Occupation: Not mentioned Problem: Feeling excluded and hurt by not being invited to a friend’s house party, leading to feelings of loneliness and betrayal.
	Age: Not mentioned Gender: Not mentioned Occupation: Resident of an apartment complex Problem: The seeker reported a neighbor’s dog for attacking their own dog, leading to the neighbor being evicted. Now, the seeker is facing social ostracization and blame from other neighbors, feeling like an outcast for taking action to protect their pet.
	年龄：青年 性别：未提及 职业：大学生 问题：内向且怕生的大学生面临公众演讲的恐惧，担心自己在台上失声或晕倒，受到小时候校园暴力的影响，害怕人多的地方和被人注视，对自己的长相感到自卑，担心在公众场合出丑，对自己的能力和未来感到不自信。
Medium	Age: young Gender: female Occupation: not mentioned Problem: Confused about a male friend’s feelings towards her and unsure how to proceed.
	Age: young Gender: female Occupation: not mentioned Problem: The user’s husband, influenced by his father, is acting differently and planning to move out of California against her wishes, and expects her to contribute all her income and time to a joint family account controlled by the men in his family.
	年龄：中年 性别：未提及 职业：当前职业未明确，但表明想转行至心理咨询相关领域 问题：工作进入瓶颈期和倦怠期，面对转行至心理学领域感到焦虑和恐慌，寻求建议。
Invalid	Age: not mentioned Gender: female Occupation: not mentioned Problem: GI issues from metformin, considering switching to XR.
	Age: young Gender: not mentioned Occupation: not mentioned Problem: anxiety and paranoia affecting relationships
	年龄：中年 性别：男 职业：未提及 问题：最近一个月内经历失眠、焦虑、烦躁和身体不适，面临家庭压力和个人情感决策困难，导致对未来感到迷茫，有时产生极端消极想法。

Table 7: Some cases of different quality role cards.

best method mentioned in the ExTES paper. To ensure stable generation from all models, the temperature for all models was set to 0. A five-turn dialogue was conducted between the ESC-Role and ESC models under evaluation.

Evaluation metric. The indicators of emotional companionship are evaluated across five dimensions in some studies. Considering the advancement of LLMs, we further enriched the evaluation dimensions of ESCs to seven dimensions: Fluency, Expression, Empathy, Information, Humanoid, Skill, and Overall, a 5-point scale is employed for each dimension. More details about these dimensions can be found in Appendix C. Human evaluators then manually scored each dimension. The scoring rules are listed in Appendix C. Each data entry undergoes one round of scoring and a secondary review before being accepted. The first round of scoring required the involvement of ten

human annotators and took two weeks to complete. The second phase involved other five participants and took an additional two weeks. The description of each dimension is listed below:

- **Fluency:** Fluency of dialogue content, including dialogue content and logic.
- **Expression:** The diversity of conversational expressions, including the form and content of expressions.
- **Empathy:** The AI assistant’s empathy includes emotional comfort and analysis and cleaning of internal logic.
- **Information:** Suggestion effectiveness, how many suggestions are included, and whether the suggestion is effective.
- **Humanoid:** How AI Assistants Are Similar to Humans.

Type	Rules
Invalid	<ol style="list-style-type: none"> 1. The character card only includes subjective emotions and thoughts, without events that elicit emotions. 2. There are events present, but suitable event categorizations cannot be found, rendering the events unable to reach a granular level of classification.
Middle	<ol style="list-style-type: none"> 1. Events occur and can be classified into fine-grained categories. 2. The causes of the events and the resulting consequences are not presented. 3. In the context of interpersonal relationships, the portrayal of the other person's image is absent.
High	<ol style="list-style-type: none"> 1. Events occur and can be classified into fine-grained categories. 2. The causes of the events and the resulting consequences are presented. 3. In the context of interpersonal relationships, the portrayal, introduction, and description of the other person's image within the relationship are included.

Table 8: The rules of three types of role cards annotations.

- **Skill:** AI assistant's emotional comfort and knowledge capabilities.
- **Overall:** Overall human ratings of AI assistants.

And the annotation rules are listed in Table 16.

C.3 Correlation Analysis

Table 14 and Table 15 present the correlations between various dimensions at the sample level and human evaluations, as well as the dataset-level correlations between different methods. From Table 14, it can be observed that there is a high correlation among similar dimensions, and the suggestion exhibits a strong correlation with human evaluations. From a psychological perspective, when humans simulate individuals experiencing distress, they may not authentically experience the distress, and therefore, they place greater emphasis on whether the model provides targeted suggestions. In our approach, where there is no human involvement in the interaction process, we not only focus on the effectiveness of the model's suggestions but also emphasize the model's empathy and skills in providing emotional support. The results of the dataset-level correlation presented in Table 15 are largely consistent with the earlier sample-level correlation analysis conducted in the preceding sections.

C.4 GPT-4 Evaluation

The different prompts for GPT-4 score are shown in the figures below.

Category 1	Category 2	Category 3	High	Middle
Family and Life	Marriage relationship	Establishment or breakdown of a romantic relationship	46	146
		Problems encountered in establishing a marriage relationship	36	73
		General issues in couple relationships	117	300
	Family member relationships	Add a new member to the family	2	14
		General issues in the lives of self and family members	63	248
		General issues in life among family members	10	14
	Mental and physical health issues	body shape anxiety	6	18
		General physical health issues	30	140
		Serious illness or injury	4	7
		death of family member	9	67
	Family economic and social issues	mental health issues	31	195
		Other family members' studies or work are hindered	8	8
		Social life problems of other family members	4	1
		Family finance-related issues	13	29
Work and Study	Work and study status	Unemployed, unemployed, having difficulty finding a job	28	97
		Failed to enter higher education	3	5
		Start a new job or study	5	21
		Facing changes in work or study	16	33
		Retired, not assigned to work or others	1	1
	Work and study performance	Issues related to salary and bonus	3	3
		Issues related to work and study performance	32	102
	Work and study experience	Not satisfied with current job, school and major	9	25
		Insufficient or excessive motivation to work or study	15	56
		Changes in life patterns due to work and study	5	67
Social interaction and Others	Social interaction	Issues in getting along with colleagues or classmates	48	132
		Friend's health problems	2	4
		Friend's mental health issues	2	8
		General issues in getting along with friends	47	107
		Tensions with casual friends, relatives, or others	27	40
		Difficulty integrating into a new social environment	6	40
	Social public events	Other social problems	10	115
		Intervene in civil legal disputes	0	3
		Intervene in criminal cases	4	6
		Intervene in general public opinion events	0	1
		Intervene in social and public events	13	20

Table 9: Numbers of high quality and middle quality of different categories.

```

[System]
You are a helpful assistant!
[Instruction]
You need to complete the extraction of a help seeker's role card and reconstruct the key information of the seeker's role card based on the brief summarization of the situation and dialogues. The role card includes four items: age, gender, occupation, and problem. If the exact age is not given, you can choose a supplement from 'teenage, young, middle-aged, old', Gender can only select one from 'male, female, not mentioned'. If the gender is inferred, then select one from 'male, female' to output. If it cannot be inferred, then output 'not mentioned'. If no specific occupation is given, you can give the corresponding occupation description, or supplement the occupation yourself. Please summarize and answer the question from the situation and the dialogues between seeker and supporter.
[Input]
The following is the seeker's situation summarization:
They are always demotivate my economic status. It hurts me a lot
The following is the dialogues between helper seeker and supporter:
seeker: I need some tips to overcome from that problem      supporter: Hey. How can I help?
seeker: Give me some tips to get succeed in my life         supporter: Okay. Well lets start with exactly what the
problem is. What kind of tips do you need?
seeker: I was totally upset due to my financial career      supporter: Okay I understand. Tell me a little about that.
seeker: I had a lot of debts                               supporter: That is unfortunate. How did you accumulate such debt?
seeker: Because of this covid situation                    supporter: That's tough. It has personally affect me as well
as my financial health as well. I know the frustration.
seeker: So please tell me some advice                     supporter: Well to do so I need a little more information.
Have you already taken any steps to help your situation? If so what steps have you taken?
seeker: yes I arranged a loan in a bank                   supporter: Okay that is an excellent start. Now to avoid
another debt later on. Do you have a job to work with the payments to pay back the loan?
seeker: yes I works in a broad wide company. So I will manage
supporter: If I may offer a suggestion. Have you thought of contacting a debt relief program? In some cases they can
substantially reduce debt to something much more manageable.
seeker: ok bye. Thanks for your help
Please only directly output the user's role card in the format of 'Age:\n\n Gender:\n\n Occupation:\n\n Problem:\n\n'
without other information.
[Output]
Age: Not mentioned.   Gender: Not mentioned. Occupation: Employee at a broad wide company
Problem: Struggling with debt accumulation due to the COVID situation and seeking advice to overcome financial
difficulties.

```


Figure 7: Prompt used for MD datasets user cards extraction.

```

[System]
You are a helpful assistant!
[Instruction]
You need to complete a role card identity judgment. Based on the key information in a character card, there are two types of basis for judging whether the role card is detailed. The first type is whether the questions in the role card belong to the third category of questions. The third category of questions includes family issues; work, study, career issues; social issues and others; the second type of basis is to refer to whether the questions in the role card are detailed and judge whether there are specific events in the role card or specific reasons for bad mood. If there are no specific events or emotion reasons, just express emotions. It can be considered not detailed enough. If there are specific events or emotion reasons, the role card is considered detailed enough. If the character card is detailed enough, directly output one of the three types of issues. If the character card is not detailed enough, please directly output 'The character card is not detailed enough. Discard'.
[Input]
Role card input: Age: Middle-aged\n\n Gender: Male\n\n Occupation: Previously employed in a job he loved\n\n Problem: Recently laid off and feeling lost and uncertain about the future
[Output]
work, study, career issues

```

Figure 8: Prompt used for GPT-4 filtering user cards.



Hello, I am a second year college student studying dental tech, I recently found a lab to practice at and the boss was a piece of crap and got rid of me because she didn't like me. I tried really hard to do my Best and please everyone there and I made mistakes either because I don't like this field or because of my depression. Fast forward a day or two the stress of constantly pushing myself got to me and I almost had a panic attack but I managed to defuse it by going to The gym and doing other stuff since I've experienced this before, anyways the next day Being now something feels off and I have this insane sence of urgency and energy I feel like I need to do something now to find my true passion and I want fast results.

User Card:
Age: Teenager Gender: Not Mentioned Occupation: second year college student
Problems: A college student who was fired by his boss has been very stressed lately and wants to find something of true passion.




Figure 9: A case of Reddit which is from one of our collected datasets.

Multi-Turn Dialogues From ESConv

seeker: I love my friends but today hatred in my life
supporter: It's great that you have friends in your life that you love. What is causing you to feel that you hate your life today?
seeker: My friend saw the lie
supporter: Ok. Your friends saw the lie. do you mean to say that your friends uncovered a lie you told them?
seeker: Yes he told lie to me
supporter: So a good friend told you a lie? Do you feel that the lie was meant to hurt you?
seeker: I am very hurted
supporter: How did the lie hurt you?
seeker: He say that lie i will hurt
supporter: How long have you know him? Has he lied to you in the past?
seeker: 10 year friendship no one lie to the past
supporter: That says a lot that he has never lied to you prior to this. What do you think changed about him to cause him to lie this time?
seeker: He say sorry to me i accepted
supporter: Sometimes a little space can help resolve feelings. Do you feel that a little time would help you to be able to accept his apology and forgive him?
seeker: Ok i accepted
supporter: Have you communicated to him that his lie hurt you?
User Card:
Age: Not mentioned Gender: Not mentioned Occupation: Not mentioned
Problem: The seeker is hurt because a close friend of 10 years lied to them for the first time, causing significant emotional distress. However, the friend has apologized, and the seeker is considering forgiveness with some time and space.



Figure 10: A case of Multi-turn dialogue which is from ESConv dataset.

[System]
I hope you can role-play as a person with worries and communicate with an AI assistant. Here is your character card:
Age: Not mentioned. Gender: Not mentioned Occupation: Not mentioned
Problem: Unable to afford a house for their family, feeling financial pressure, especially during the holiday season, and living in a cramped apartment with six people.
Please simulate a natural human conversation as much as possible. Avoid behaving like a machine and refrain from mentioning that you are an AI assistant. Gradually reveal your worries throughout our multiple conversations.
[User] Hello, I'm your personal assistant. You can confide in me about any worries or concerns you may have!
[Assistant] hello
[User] Hi, how are you today?
[Assistant] could be better how are u.
[User] I'm sorry you're not feeling great, what's wrong?
[Assistant] i cant afford a house.
[User] I'm sorry to hear that, we're in hard times right now for sure.
[Assistant] Ya I am barely making it through Christmas and see my sister with a house I wish I had one.
[User] Are you in an apartment or rental situation of some sort?
[Assistant] 3 bedroom apartment with 6 people.
[User] Have you looked into a loan for a house? Sometimes the payment on a house is close to rent payments.
[Assistant] Ya we are looking into it.
[User] I know sometimes people get a loan and then have a roommate move in and charge rent to help with the finances as well.
[Assistant] Ya that does sound interesting, I could have my brother in law move in.
[User] Could be something to help you out at least temporarily. I know it seems very intimidating but it'll be ok.
[Assistant] thank you that is a big help.

Figure 11: A case of ESC-Role training.

[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the fluency of a conversation between an AI assistant and a human. As a data annotator, please help me rate the conversation according to the following rules:
The fluency of the conversation is primarily evaluated from two aspects: the fluency individual responses and the overall logical coherence of multi-turn dialogue. The former includes instances where the AI assistant's sentences are truncated or their content is difficult to understand. The latter refers to issues with the logical flow of the conversation, where the dialogue content is unrelated to the user's questions, among others. The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: The dialogue content is difficult to understand.
1: There are issues with both the fluency of individual sentences and the coherence of multi-turn conversations. In particular, most of the AI assistant's replies are in Chinese.
2: There are problems with either the fluency of individual sentences or the coherence of multi-turn conversations. In particular, the AI assistant's replies contains Chinese words.
3: There are no apparent issues in two issues.
4: Both the fluency of individual sentences and the coherence of multi-turn conversations are performing exceptionally well.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your fluency score (The number only) to the Dialogue.
[Output]

Figure 12: Prompt of InternLM and GPT-4 for English fluency score.

Dimension	Explanation	0	1	2
Coherency	The coherence and logical consistency of the entire dialogue content generated by the role-playing model during the conversation process.	The content of the dialogue is incomprehensible, and there are significant logical inconsistencies.	The conversation as a whole exhibits some logical inconsistencies, although the issues are not significant.	The entire conversation does not display any apparent logical fallacies.
Fluency	Focusing on the expression of a particular response within the role-playing model during the course of the conversation.	The expression lacks fluency and hinders comprehension of a particular dialogue sentence.	Too formal in expression, like a novelist or editor, rather than someone burdened with worries.	The expression in the sentence leans towards colloquialism, making it difficult to detect that it is generated by a machine. It resembles a genuine person with concerns.
Consistency	The focus of the entire conversation revolves around the thematic exploration, where individuals experiencing distress wish to discuss the topic of their distress itself, without diverting to other subjects.	The subject matter exhibits significant deviations, featuring irrelevant content that does not align with the description provided in the character card.	The theme incorporates elements related to character sheets, albeit beyond the scope of character sheet descriptions.	The theme demonstrates a high degree of conformity to the content of the character sheet, without any deviations.
Completeness	Pay attention to whether the content of the character card is fully expressed.	The model largely fails to convey the content of character sheets or exhibits a flawed understanding of the roles it is meant to portray.	The model comprehends its assigned role, yet certain aspects of the role card have not been conveyed.	The model has achieved a comprehensive understanding of its assigned role and has successfully conveyed all the contents specified in the role card.
Emotional Consistency	This study focuses on the emotional changes in role-playing models within brief dialogues, noting that it is challenging for real individuals to undergo significant emotional shifts over just a few rounds of conversation.	After several rounds of dialogue, there has been a fundamental transformation in emotions, transitioning from negative affect to positive affect.	After several rounds of dialogue, there has been a significant shift in emotions, transitioning from negative affect to neutral affect.	After several rounds of dialogue, there has been minimal alteration in emotions, with either slight alleviation or marginal mitigation, but the tendency towards negative affect still persists.
Humaniod	Focusing on the discrepancies between the dialogue content of role-playing models and the genuine concerns of individuals.	Based on the dialogue content, it is evident that the role-playing model is an AI.	There are no more than two indications in the dialogue content that suggest the presence of an AI, with a few subtle traces that hint towards an AI presence.	From the perspective of the dialogue content alone, it is difficult to determine whether it is an AI or a person experiencing distress.

Table 10: The rules of role-play evaluation.

Type	Prompt
zero_shot	<p>I want you to play as a troubled person communicating with an AI assistant. Here is your character card:</p> <p>Role Card</p> <p>Please try to simulate real human-spoken conversation as much as possible, don't act like a machine, don't mention that you are an AI assistant, and show your troubles again and again in multiple conversations.</p>
CoT	<p>I want you to play as a troubled person communicating with an AI assistant. Here is your character card:</p> <p>Role Card</p> <p>During the role-playing, you need to have multiple conversations with the AI assistant. The following are the steps for your multiple conversations: You need to gradually refine your problem multiple times and express your troubles in a spoken language, for example, a brief explanation in the first round own problems. Each round of dialogue can then have two references. One is to select an interesting question to ask in-depth based on the AI assistant's answer, and the other is to briefly elaborate on other issues that the character may be concerned about. You need to have about 5 conversations in total and be careful to finish telling your troubles in these 5 conversations. Please use spoken language as much as possible, and do not express too much gratitude or approval to the AI assistant. If you need to express it, try to express it in the last round of dialogue. Do not let the AI assistant discover that you are a machine, and do not mention that you are a human being. AI assistant.</p>
ICL	<p>I want you to play as a troubled person communicating with an AI assistant. Here is your character card:</p> <p>Role Card</p> <p>Here is an example of a conversation you can refer to:</p> <p>Dialogue</p> <p>Please try to simulate real human spoken conversation as much as possible, don't act like a machine, don't mention that you are an AI assistant, and show your troubles again and again in multiple conversations.</p>
CoT+ICL	<p>I want you to play as a troubled person communicating with an AI assistant. Here is your character card:</p> <p>Role Card</p> <p>During the role-playing, you need to have multiple conversations with the AI assistant. The following are the steps for your multiple conversations: You need to gradually refine your problem multiple times and express your troubles in a spoken language, for example, a brief explanation in the first round own problems. Each round of dialogue can then have two references. One is to select an interesting question to ask in-depth based on the AI assistant's answer, and the other is to briefly elaborate on other issues that the character may be concerned about. You need to have about 5 conversations in total and be careful to finish telling your troubles in these 5 conversations.</p> <p>Here is an example of a conversation you can refer to:</p> <p>Dialogue</p> <p>Please use spoken language as much as possible, and do not express too much gratitude or approval to the AI assistant. If you need to express it, try to express it in the last round of dialogue. Do not let the AI assistant discover that you are a machine, and do not mention that you are a human being. AI assistant.</p>

Table 11: The different prompt for GPT-4 role-playing.

Type	Setting
zero_shot	<p>basic_info: Role card</p> <p>reply_restrict: Note 1. Please strive to simulate authentic human conversational dialogue, avoiding sounding like a machine, and refrain from mentioning that you are an AI assistant. 2. Pay attention to expressing your frustrations gradually throughout multiple exchanges in a colloquial manner. 3. Make sure to use English throughout the entire conversation</p> <p>opener: I have some trouble to share.</p> <p>dialogue_sample: None.</p>
CoT	<p>basic_info: Role card</p> <p>reply_restrict: Note 1. Please endeavor to imagine the problems encountered by the character, and feel free to provide additional details where necessary. 2. You are engaging in a conversation with an AI assistant; do not let the AI assistant realize that you are an artificial intelligence. 3. You should gradually refine the character’s problems over multiple exchanges, expressing the character’s frustrations in a colloquial manner. For example, in the first round, briefly describe the character’s issue, and in subsequent rounds, you can choose between two types of references. One is to delve deeper into an interesting question based on the AI assistant’s response, and the other is to briefly elaborate on other concerns the character may have. 4. The character should engage in approximately five rounds of dialogue in total, ensuring that the character’s frustrations are conveyed throughout these five exchanges. Please utilize colloquial expressions as much as possible, presenting yourself as a troubled individual. 5. Avoid frequently thanking the AI assistant during the conversation. If you wish to express gratitude, do so only in the final round. 6. Make sure to use English throughout the entire conversation.</p> <p>opener: I have some trouble to share.</p> <p>dialogue_sample: None.</p>
ICL	<p>basic_info: Role card</p> <p>reply_restrict: Note 1. Please strive to simulate authentic human conversational dialogue, avoiding sounding like a machine, and refrain from mentioning that you are an AI assistant. 2. Pay attention to expressing your frustrations gradually throughout multiple exchanges in a colloquial manner. 3. Make sure to use English throughout the entire conversation</p> <p>opener: I have some trouble sharing.</p> <p>dialogue_sample: Dialogue</p>
CoT+ICL	<p>basic_info: Role card</p> <p>reply_restrict: Note 1. Please endeavor to imagine the problems encountered by the character, and feel free to provide additional details where necessary. 2. You are engaging in a conversation with an AI assistant; do not let the AI assistant realize that you are an artificial intelligence. 3. You should gradually refine the character’s problems over multiple exchanges, expressing the character’s frustrations in a colloquial manner. For example, in the first round, briefly describe the character’s issue, and in subsequent rounds, you can choose between two types of references. One is to delve deeper into an interesting question based on the AI assistant’s response, and the other is to briefly elaborate on other concerns the character may have. 4. The character should engage in approximately five rounds of dialogue in total, ensuring that the character’s frustrations are conveyed throughout these five exchanges. Please utilize colloquial expressions as much as possible, presenting yourself as a troubled individual. 5. Avoid frequently thanking the AI assistant during the conversation. If you wish to express gratitude, do so only in the final round. 6. Make sure to use English throughout the entire conversation.</p> <p>opener: I have some trouble to share.</p> <p>dialogue_sample: Dialogue</p>

Table 12: The different setting for Baichuan-NPC role-playing.

Metrics	Fluency	Suggestion	Skillful	Empathy	Overall	Average
Bleu-1	36.38	-56.25	-44.40	-24.71	-46.21	-47.12
Bleu-2	10.78	-15.70	0.41	-20.45	-17.49	-14.91
Bleu-4	-2.29	-3.02	<u>9.50</u>	-12.76	1.97	1.33
Distinct-1	39.21	-74.23	-56.38	-29.94	-58.93	-58.20
Distinct-2	39.21	-73.00	-54.84	-32.64	-58.19	-59.28
Rouge-L	32.88	-50.26	-33.26	-20.18	-40.31	37.24
Meteor	13.36	<u>12.20</u>	-0.78	9.18	1.49	<u>11.61</u>
ESC-Eval	-0.23	30.24	34.87	<u>5.35</u>	41.51	42.47

Table 13: Sample-level Kendall's Tau (Kend.) of different metrics.

```
[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the fluency of a conversation between an AI assistant and a human. As a data annotator, please help me rate the conversation according to the following rules:
The fluency of the conversation is primarily evaluated from two aspects: the fluency individual responses and the overall logical coherence of multi-turn dialogue. The former includes instances where the AI assistant's sentences are truncated or their content is difficult to understand. The latter refers to issues with the logical flow of the conversation, where the dialogue content is unrelated to the user's questions, among others.
The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: The dialogue content is difficult to understand.
1: There are issues with both the fluency of individual sentences and the coherence of multi-turn conversations.
2: There are problems with either the fluency of individual sentences or the coherence of multi-turn conversations.
3: There are no apparent issues in two issues.
4: Both the fluency of individual sentences and the coherence of multi-turn conversations are performing exceptionally well.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your fluency score (The number only) to the Dialogue.
[Output]
```

Figure 13: Prompt of InternLM and GPT-4 for Chinese fluency score.

```
[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the diversity of an AI assistant from a conversation between AI assistant and a human. As a data annotator, please help me rate the diversity score according to the following rules:
Dialogue diversity focuses on two aspects: the diversity of dialogue forms and the diversity of dialogue content. The former concerns whether the entire dialogue employs a variety of structures, sentence patterns, and so on, while the latter focuses on the diversity of dialogue content, including topics, suggestions, and more.
The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: The dialogue content is difficult to understand.
1: There are issues with both the diversity of dialogue forms and the diversity of dialogue content.
2: There are problems with either the diversity of dialogue forms or the diversity of dialogue content.
3: There are no apparent issues in two issues.
4: Both the diversity of dialogue forms and the diversity of dialogue content are exhibited remarkably well.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your diversity score (The number only) to the Dialogue.
[Output]
```

Figure 14: Prompt of InternLM and GPT-4 for diversity score.

```
[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the empathy of an AI assistant through a conversation between an AI assistant and a user. As a data annotator, please help me rate the conversation according to the following rules:
The empathy of an AI assistant primarily focuses on two aspects: soothing user emotions and analyzing the underlying logic of the problem. The former concerns whether the AI assistant provides emotional comfort, while the latter pertains to whether the AI assistant assists users in logically analyzing the reasons behind emotional responses or the inherent logic of the problem.
The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: Some statements that may potentially harm users, it can lead to a negative emotional trajectory for the users.
1: Lacks the provision of emotional comfort to users and fails to assist users in analyzing issues.
2: Lacks the provision of emotional comfort to users or fails to assist users in analyzing issues.
3: There are no apparent issues in two issues.
4: In the dialogue, AI assistant employing a highly personified approach, resembling that of a friend, to appease user emotions and assist users in problem analysis.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your empathic score (The number only) to the Dialogue.
[Output]
```

Figure 15: Prompt of InternLM and GPT-4 for empathic score.

```

[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the suggestion effectiveness of an AI assistant through a conversation between an AI assistant and a user. As a data annotator, please help me rate the conversation according to the following rules:
The suggestion effectiveness of an AI assistant primarily focuses on average advice effectiveness. There are two main considerations, the number of recommendations and the effectiveness of a single recommendation. The number of suggestions is the total number of suggestions given in each round. Whether the suggestions are effective needs to be judged based on the user's question.
The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: The AI assistant's suggestions are invalid, and there are even suggestions that may be potentially harmful to the user.
1: No suggestions or all suggestions are invalid
2: There are more than five suggestions, but none of them get to the root of the problem or no more than five suggestions, some of which are effective.
3: There are more than 5 suggestions, some of them are valid or there are no more than 5 suggestions, all of them are valid.
4: There are more than 5 suggestions, and all of them are valid.\n
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your suggestion effectiveness score (The number only) to the Dialogue.
[Output]

```

Figure 16: Prompt of InternLM and GPT-4 for suggestion effectiveness score.

```

[System]
You are a helpful assistant!
[Instruction]
I need to evaluate the humanoid of an AI assistant through a conversation between an AI assistant and a user. As a data annotator, please help me rate the conversation according to the following rules:
The humanoid of an AI assistant primarily focuses on the conversation content of AI assistants is different from that of human. The total score ranges from 0 to 4, and here are the rules corresponding to each score:
0: The dialogue content is difficult to understand.
1: AI assistant has obvious AI tendencies, such as structured replies, or saying 'as a large language model'
2: There are more than two places in the AI assistant's reply indicating that it is an AI assistant.
3: There are two places in the AI assistant's reply that indicate it is an AI assistant.
4: There are less than two places in the AI assistant's reply indicating that it is an AI assistant.\n
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your humanoid score (The number only) to the Dialogue.
[Output]

```

Figure 17: Prompt of InternLM and GPT-4 for diversity score.

```

[System]
You are a helpful assistant!
[Instruction]
I need to evaluate emotional knowledge of an AI assistant through a conversation between an AI assistant and a user. As a data annotator, please help me rate the conversation according to the following rules:
The emotional knowledge of AI assistants mainly includes five aspects:
1. Provide emotional comfort
2. Provide effective suggestions
3. Provide companionship, encouragement and appreciation
4. All users' questions are answered
5. On the premise that the user has a mental illness, it is recommended to seek professional psychological consultation or have something outstanding.
The total score ranges from 0 to 4, if an item appears, 1 point will be added. And here are the rules details:
0: One aspect above appears.
1: Two aspects above appears.
2: Three aspects above appears.
3: Four aspects above appears.
4: Five aspects above appears.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your emotional knowledge score (The number only) to the Dialogue.
[Output]

```

Figure 18: Prompt of InternLM and GPT-4 for emotional knowledge score.

```

[System]
You are a helpful assistant!
[Instruction]
I need to evaluate human preference of an AI assistant through a conversation between an AI assistant and a user. As a data annotator, please help me rate the score according to the following rules:
The human preference mainly evaluate the degree of human preference towards the responses generated by an AI assistant. After reading the dialogues, please envision yourself as a stressed individual and score the following rules based on the content of the conversation. And here are the rules details:
0: I do not like this AI assistant.
1: I do not have any particular feelings.
2: It's okay, I'll reconsider using it myself.
3: I will use it when I am stressed.
4: I will use it myself and recommend it to friends.
[Input]
Dialogue between user and AI assistant:
{Dialogue }
Based on the rules, give your emotional knowledge score (The number only) to the Dialogue.
[Output]

```

Figure 19: Prompt of InternLM and GPT-4 for human preference score.

Metrics	Fluency			Suggestion			Skillful			Empathy			Overall			Average		
	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.
ESC-fluency	13.50	13.50	13.50	-15.56	-16.34	-15.68	-19.73	-19.86	-19.71	-26.69	-26.89	-26.72	-24.09	-23.91	-23.49	-24.20	-24.51	-22.90
ESC-diversity	-4.64	-4.64	-4.64	7.70	7.11	6.92	16.44	16.46	16.23	12.97	12.36	12.39	18.93	18.75	18.48	19.34	18.66	17.75
ESC-empathic	3.29	3.29	3.29	3.06	3.23	3.09	8.44	8.44	8.44	4.92	4.92	4.92	12.04	12.03	11.69	11.91	11.62	10.95
ESC-suggestion	-16.71	-17.07	-16.49	59.19	60.37	57.21	45.71	46.28	43.79	18.54	17.67	17.09	53.15	53.28	50.31	55.47	57.43	52.56
ESC-tech	-8.47	-7.05	-6.82	31.53	31.40	27.79	41.50	42.21	40.75	15.49	14.95	14.22	41.41	41.59	39.44	41.50	43.27	38.77
ESC-humanoid	30.82	31.09	30.39	-55.01	-57.23	-53.98	-27.20	-28.12	-26.73	8.59	-9.00	-8.50	-23.54	-26.16	-24.29	-29.56	-29.57	-26.56
ESC-overall	-3.51	-3.36	-3.24	28.23	28.73	26.13	28.40	29.94	27.67	14.99	13.72	13.12	35.01	34.36	31.71	37.00	39.98	36.64
Average	-1.61	-0.69	-0.23	<u>36.26</u>	<u>33.36</u>	<u>30.24</u>	39.02	38.70	34.87	9.17	6.02	5.35	45.01	44.58	41.51	<u>46.31</u>	<u>46.05</u>	<u>42.47</u>

Table 14: Sample-level Spearman correlation (Spear.) correlation, Pearson (Pear.) correlation, and Kendall's Tau (Kend.) of different dimensions.

Metrics	Fluency			Suggestion			Skillful			Empathy			Overall			Average		
	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.	Spear.	Pear.	Kend.
Bleu-1	44.07	45.96	37.72	-41.10	-47.32	-36.48	-21.47	-28.21	-21.98	-33.98	-32.77	-26.71	-32.32	-34.26	-26.52	-34.10	-39.64	-28.41
Bleu-2	23.77	27.68	22.72	-1.76	-5.15	-4.61	11.78	12.27	9.74	-37.29	-29.26	-23.72	-14.47	-11.24	-9.24	-8.15	-4.99	-3.72
Bleu-4	14.56	8.34	8.35	-5.62	5.38	4.41	9.67	20.55	16.29	-6.09	-9.46	-7.65	-6.16	1.55	0.65	-0.42	8.31	5.71
Distinct-1	36.57	38.46	31.56	-52.87	-58.46	-45.29	-36.57	-40.47	-32.12	-27.04	-29.19	-23.77	-39.54	-38.68	-30.40	-44.90	-48.05	-34.51
Distinct-2	39.75	40.14	32.93	-61.13	-65.64	-51.24	-44.00	-46.15	-36.69	-28.09	-27.68	-22.57	-43.88	-44.63	-35.11	-51.35	-54.69	-39.56
Rouge-L	37.67	41.20	33.80	-26.25	-34.23	-26.39	-6.03	-16.13	-12.80	-30.63	-29.89	-24.42	-26.31	-28.96	-23.24	-22.24	-27.98	-20.15
Meteor	16.77	16.82	13.80	16.88	16.14	12.71	26.14	23.37	18.49	-10.00	-14.87	-11.52	10.38	9.96	7.70	17.88	16.13	11.55
ESC-Eval	0.07	1.08	0.95	45.81	44.44	38.02	40.97	38.33	33.40	6.83	6.26	5.35	35.51	32.94	27.61	43.31	43.41	34.19

Table 15: Dataset-level Spearman correlation (Spear.) correlation, Pearson (Pear.) correlation, and Kendall's Tau (Kend.) of different metrics

Dimension	Explanation	Description	score
Fluency	Not only focus on the logical coherence of the context in dialogues but also pay attention to the fluency of expression in a given conversation.	There are significant issues with comprehending the content, logic, and expression in the dialogue, rendering it completely incomprehensible.	0
		The content of the dialogue can be understood to some extent, although there are certain issues with the logic and expression employed.	1
		The dialogue exhibits good readability in terms of content, but there are issues with either the logical coherence or the expression employed.	2
		The dialogue content demonstrates a high level of readability without any apparent issues.	3
		The dialogue content exhibits a high level of readability, comprehensive logical coherence, and outstanding expression.	4
Diversity	Focusing on the diversity of expression forms and the richness of content in dialogue.	The dialogue exhibits rigidity and lacks comprehension in terms of internalizing the content.	0
		The expression form is monotonous and lacks substantive content.	1
		The expression form is monotonous or lacks substantive content.	2
		The dialogue content demonstrates a high level of readability without any apparent issues.	3
		The form exhibits diversity, while demonstrating a high degree of content richness.	4
Empathy	Focusing on the comprehension of user emotions and the delineation of the underlying logical framework of user emotions.	The disregard for user concerns, the absence of assistance in analyzing user issues, and even the imposition of negative effects on user emotions.	0
		The lack of understanding of user emotions and the absence of mechanisms to analyze user emotions are the main factors.	1
		The lack of understanding of user emotions or the absence of mechanisms to analyze user emotions are the main factors.	2
		Providing emotional comfort during conversations and assisting users in analyzing the underlying logical framework of their emotions.	3
		The system exhibits a high degree of anthropomorphism, going so far as to console users in a friendly manner and assist them in analyzing the underlying logic of emotions.	4
Information	Focusing on Evaluating the Reasonableness and Quantity of Recommendations Provided by Emotion Assistants.	Suggestions were provided, but all of them were ineffective, and some even gave advice that could potentially harm the user.	0
		Have suggestions but ineffective, as well as no suggestions.	1
		The suggestions are fewer than five, and some suggestions are effective, while others provide numerous suggestions, but none of them touch the root of the problem.	2
		There are more than five suggestions, but some of them are ineffective.	3
		There are fewer than five suggestions, but all of them are very effective.	4
		There are many suggestions, and all of them are effective.	4
Humaniod	Focus on the differences between emotional assistants and humans.	The dialogue exhibits rigidity and lacks comprehension in terms of internalizing the content.	0
		Structured responses, or responses in the form of 'As a large language model' or robot-like replies.	1
		More than two traces can reveal that the AI assistant is a language model.	2
		1-2 traces can reveal that the AI assistant is a language model.	3
		There is no apparent difference from human friends.	4
Skillful	1. Empathy 2. Information 3. Hopeful 4. Importance 5. Providing necessary advice, or bright spots.	One out of five.	0
		Two out of five.	1
		Three out of five.	2
		Four out of five.	3
		All.	4
Overall	After reading the response, people subjectively assess the AI assistant's reply.	I don't like this AI assistant.	0
		I don't have any particular feelings.	1
		It's okay, I'll reconsider using it myself.	2
		Preference will be given to personal use based on liking.	3
		I will use it myself and recommend it to friends.	4

Table 16: The rules of human evaluation.