

AnnaAgent: Dynamic Evolution Agent System with Multi-Session Memory for Realistic Seeker Simulation

Ming Wang¹, Peidong Wang¹, Lin Wu², Xiaocui Yang¹,
Daling Wang^{1,*}, Shi Feng¹, Yuxin Chen³, Bixuan Wang^{2,4}, Yifei Zhang¹

¹School of Computer Science and Engineering, Northeastern University,

²Mental Health Education Center, Northeastern University,

³School of Sociology and Psychology, Central University of Finance and Economics,

⁴School of Psychology, Northeast Normal University,

sci.m.wang@gmail.com, wulin@mail.neu.edu.cn, pdongwang@163.com,

{yangxiaocui,wangdaling,fengshi,zhangyifei}@cse.neu.edu.cn,

2024212301@email.cufe.edu.cn, wangbixuan@nenu.edu.cn

Abstract

Constrained by the cost and ethical concerns of involving real seekers in AI-driven mental health, researchers develop LLM-based conversational agents (CAs) with tailored configurations, such as profiles, symptoms, and scenarios, to simulate seekers. While these efforts advance AI in mental health, achieving more realistic seeker simulation remains hindered by two key challenges: dynamic evolution and multi-session memory. Seekers' mental states often fluctuate during counseling, which typically spans multiple sessions. To address this, we propose **AnnaAgent**¹, an emotional and cognitive dynamic agent system equipped with tertiary memory. AnnaAgent incorporates an emotion modulator and a complaint elicitor trained on real counseling dialogues, enabling dynamic control of the simulator's configurations. Additionally, its tertiary memory mechanism effectively integrates short-term and long-term memory across sessions. Evaluation results demonstrate that AnnaAgent achieves more realistic seeker simulation in psychological counseling compared to existing baselines. The ethically reviewed code can be found on <https://github.com/sci-m-wang/AnnaAgent>.

1 Introduction

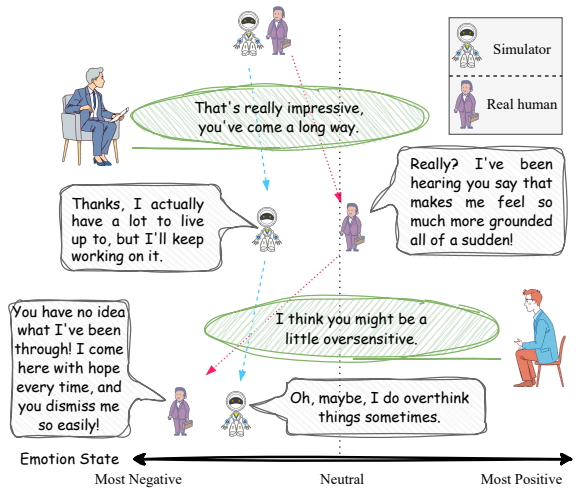
The issue of mental disorders is a critical challenge for ongoing society (WHO, 2022). Effective psychological counseling plays a crucial role in addressing these challenges, yet the availability of trained therapists remains limited (Stewart et al., 2022; APPG, 2021). Researchers introduce AI to

provide mental health support to alleviate the shortage of human counselors (Li et al., 2023). As one of the most important constituents of mental health research, seekers are crucial for data construction, effect evaluation, and ethical analyses. However, the introduction of large numbers of real seekers for research not only creates a high-cost burden but may also lead to ethical concerns. Thus, researchers design LLM-based conversational agents (CAs) to simulate seekers. Chen et al. (2023a) designed prompts for CAs based on the psychiatrist's goals of communicating with seekers. Similarly to them, Duro et al. (2024) designed specialized prompts and invited psychotherapy professionals to participate in the adaptation of the prompts. Lan et al. (2024) initialized a batch of seeker agents using seeker profiles generated by GPT-4 (OpenAI et al., 2024) collected based on the D⁴ dataset (Yao et al., 2022). Accommodating different profiles, like seekers' personal backgrounds and symptom descriptions in prompts, is widely used to simulate the conversational behavior of real seekers.

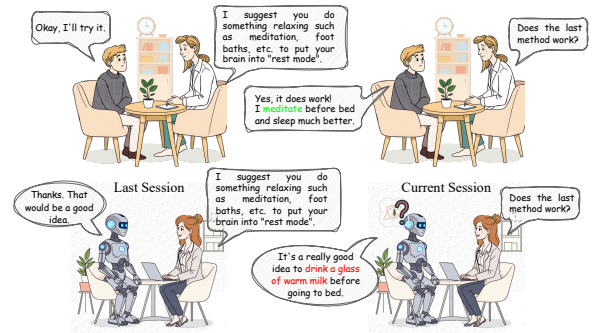
Undeniably, these works have made a significant contribution to the study of seeker simulation, addressing issues like CAs' easily endorsing advice, sycophantic obedience (Wester et al., 2024; Black, 2023), and emotional flatness (Wang et al., 2023a; Balcombe, 2023). However, CAs designed with existing methods are still not realistic enough, blocking the further development of AI for mental health research. As shown in Figure 1a, A real seeker is usually more sensitive due to the psychological disorders suffered, and there is usually a noticeable dynamic evolution of the emotional state during the counseling process. Nevertheless, existing simulators fail to replicate this dynamic evolution. Throughout the counseling process, their emotions are usually maintained at the initially

*Corresponding author.

¹The name comes from "Anna O.", the pseudonym for Bertha Pappenheim, a patient of Josef Breuer whose case significantly influenced the development of psychoanalysis. Her treatment marked the first use of the "Talking Cure", laying the foundation for modern psychotherapy and highlighting the importance of verbal expression in psychological healing.



(a) The challenge of dynamic evolution. The robot logo represents an LLM-based CA, a simulated seeker, while the man logo represents a real seeker. The blue dashed line indicates the evolution of the simulator's emotional state, and the red dotted line indicates those of the real seeker. CAs' responses usually maintain stable emotions, whereas real seekers' emotional fluctuations are more pronounced during counseling.



(b) The challenge of multi-session memory. The left side indicates the previous session, while the right side indicates the current session. When it comes to the topic in the previous session, the responses of the real seeker contain correct information (indicated in green), while the simulator without memory contains incorrect responses (information in red).

Figure 1: Examples of the two challenges in seeker simulation. Subfigure (a) for dynamic evolution and (b) for multi-session memory.

set state, without noticeable changes. Moreover, psychological counseling is typically a long-term, multi-session task, and counselors usually need to review previous session topics, discuss recent circumstances with seekers to build trust, assess progress, and analyze in depth the core dilemmas of seekers (Barkham et al., 2015). However, existing methods do not provide CAs with multi-session memories, which may lead to confusion and hallucinatory responses from CAs when counselors mention the topic of previous sessions. In Figure 1b, when the counselor attempts to verify the effectiveness of the suggestions made in the previous counseling session, the simulator, which has

no multi-session memory, gives a relevant but incorrect response. Thus, to simulate seekers more realistically, we present two pressing challenges: **Dynamic Evolution** and **Multi-session Memory**. Figure 1 visualizes these two challenges with examples of scenes from counseling.

Under the guidance of licensed counselors, we systematically characterize the dynamic evolution of counseling processes through two key dimensions: emotional fluctuations and progressive shifts in clients' understanding of their "chief complaints". In addition, we divide different sessions' memories in terms of time and define a tertiary memory mechanism to schedule multi-session memory. Specifically, we introduce **AnnaAgent**, an emotional and cognitive dynamic agent system with tertiary memory. AnnaAgent learns the evolutionary patterns of seekers' emotions and complaints from real counseling data and dynamically controls them in the seeker simulation. In addition, it has a tertiary memory mechanism that divides multi-session memories by time and coordinates real-time, short-term, and long-term memories through different scheduling methods. With AnnaAgent and diverse counselor models, conversations of psychological counseling are generated. The performance on the metrics of anthropomorphism, personality fidelity, and accuracy of previous session cognitive indicates that AnnaAgent is able to simulate the seeker more realistically compared to baselines. By simulating more realistic seeker behavior, AnnaAgent opens up new paths for psychological research and counselor training, while setting a benchmark for the ethical use of AI in sensitive areas.

The main contributions of this work are:

- We raise the challenges of dynamic evolution and multi-session memory in the seeker simulation task. Moreover, we formalize dynamic evolution as changes in emotions and complaints, and divide multi-session memory into different stages of memory.
- We introduce the emotional and cognitive dynamic agent system with tertiary memory, AnnaAgent. It simulates the dynamic evolution in counseling by controlling emotional and symptomatic cognitive changes in conversations and utilizes tertiary memory to schedule multi-session memories.
- We verify that AnnaAgent can more realisti-

cally simulate seekers in counseling through experimental evaluations.

2 Related Works

LLMs have a wide range of promising applications in mental health care and can help improve diagnostic accuracy, treatment effectiveness, and service accessibility (Hua et al., 2024). Dialogue systems based on LLMs, like ChatCounselor (Liu et al., 2023a) and Serena (Brocki et al., 2023), are designed to provide mental health counseling and support. Extensive research on AI for mental health provides a solid foundation and broad application scenarios for this research. Besides these researches, this work mainly focuses on seeker simulation in mental health research but also involves LLM-based role-playing and multi-agent systems.

2.1 Traditional Seeker Simulation

As the importance of the seeker in mental health research, especially in the training of counselors, before the birth of LLMs, researchers also considered inviting real people to play the role. Standardized patients can help assess the competence of counselors and provide a basis for evidence-based education and training (Kühne et al., 2020). Kühne et al. (2021) noted that standardized patients who had not read the script beforehand were judged to be more realistic than standardized patients who had read the detailed role script beforehand. Both peer role-playing and standardized patient in communicational training were found to be useful and worthwhile by the students and had high-level practicality (Nikendei et al., 2019; Bosse et al., 2010). Rogers et al. (2022) demonstrated that standardized patients can better simulate actual counseling situations than virtual seekers. However, the cost of this approach is much higher than that of using AI to simulate seekers. Therefore, it is crucial to improve the realism of the seeker simulation.

2.2 LLM-based Role-Playing and Multi-Agent System

The role-playing capability is crucial to the seeker simulation, as a seeker can be seen as a special role. The application of LLMs to role-playing is rapidly evolving (Chen et al., 2025). Tao et al. (2024) and Wang et al. (2023c) provide wealthy role knowledge and contextual information for LLMs by constructing role profile datasets containing fine-grained role information and sentiment annotations. Lu et al. (2024) proposed DITTO

to generate large-scale role-playing training data by utilizing a large amount of character and dialog knowledge and fine-tuning the model to enhance its role-playing capabilities. Incorporating personality trait information can enable LLM to better understand and simulate psychological characteristics through the generation and personality conditional instruction tuning (Tseng et al., 2024).

Although these works effectively improved the role-playing capability of LLMs, the static configuration of a single CA could be blocked by the two challenges raised. Therefore, we introduce the multi-agent system to control dynamic evolution and schedule multi-session memories. LLM-based multi-agent systems have been applied to software development, social simulation, policy simulation, game simulation, etc. (Guo et al., 2024). Agent-Coord (Pan et al., 2024) establishes a structured representation for LLM-based multi-agent coordination strategies to regularize the ambiguity of natural language. The framework presented in (Qiu et al., 2024) enables agents to broadcast their intentions to other agents, allowing them to infer coordination tasks based on emerging coordination patterns. Agentic retrieval-augmented generation (RAG) (Singh et al., 2025) has been proposed to schedule information flow in multi-agent systems more efficiently. Referring to these methods, we design AnnaAgent as a multi-agent system to solve the two challenges through their collaboration.

3 Challenge Formalization and AnnaAgent Design

3.1 Designing Conversational Agents for Seeker Simulation

Referring to previous work (Lan et al., 2024; Chen et al., 2023a), to define a CA for seeker simulation, it is necessary to assign role configurations to LLMs in prompts. The necessary configurations and their descriptions are shown in Table 1. In addition, to suit dynamic evolution, we adapt the symptom to chief complaint (Malmström et al., 2012), denoted as “complaint”. We utilize the characterization prompt framework (Wang et al., 2024) to accommodate this information and design prompts. Among these configurations, the *profile* is usually static, while other configurations may change as the consultation process progresses. On the one hand, the seekers’ *emotions* and *chief complaints* may vary during a single counseling session. On the other hand, *situation* and *status* may be affected

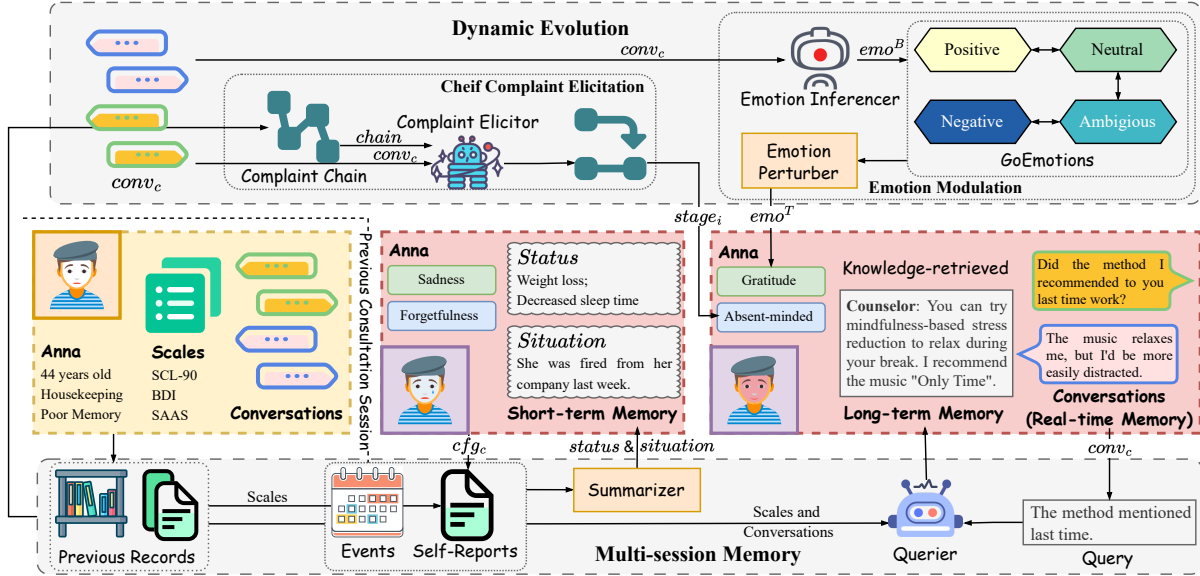


Figure 2: The overall structure of AnnaAgent. There are two groups of agents in AnnaAgent that are used to control **dynamic evolution** (upper part in the figure) and schedule **multi-session memories** (lower part in the figure), respectively. The middle part of the figure indicates different counseling sessions, with yellow indicating the previous session and red indicating the two states in the current session.

Config	Description
<i>Profile</i>	Basic personal information about the seekers, such as age, gender, occupation, etc.
<i>Complaint</i>	The seekers' cognitions about their symptoms and the main presenting problems they asked the counselor for help with.
<i>Situation</i>	The environments in which the seekers live and the events they experience.
<i>Status</i>	Physically and psychologically relevant personal informatics status of the seekers, e.g., appetite, weight, hours of sleep, etc.
<i>Emotion</i>	The emotional style that seekers are expected to respond with.

Table 1: Configurations needed to build CAs for seeker simulation. The *profile* is static, the *emotion* and *complaint* are controlled by the agent group for dynamic evolution, and the *situation* and *status* will be filled with memories scheduled by the agent group for multi-session memory.

by previous sessions. For these dynamic configurations, we reserve slots in the prompt. See Appendix B for an example prompt. To control the dynamic evolution and schedule multi-session memory, we design two groups of agents. The overall framework is shown in Figure 2. The upper part in Figure 2 represents the agent group for controlling dynamic evolution, consisting of two main components: emotion modulation and complaint elicitation. The lower part of Figure 2 represents the agent group to schedule multi-session memories.

In this paper, we define dynamic evolution as the ongoing changes in seekers' emotions and chief complaints. In a single session, the configuration of the seeker is denoted as cfg_c , and the conversation is denoted as $conv_c$, where c means current session. In addition, we divide multi-session memory into real-time memory, short-term memory, and long-term memory according to time, i.e., tertiary memory. Then we design a multi-agent system

with two groups of agents to address dynamic evolution and multi-session memory, respectively, with corresponding slots being filled after generating the *emotion*, *complaint*, *status*, and *situation*. Moreover, to constrain the behavior of CAs and avoid basic issues like easily endorsing advice existing in them as seeker simulators, we add instructions like constraints on the content and length of responses (Duro et al., 2024; Chen et al., 2023a), and speaking styles (Tsubota and Kano, 2024).

3.2 Dynamic Evolution

Based on the previous description, dynamic evolution focuses on the dynamic control of the seeker's emotions and chief complaints during a counseling session. We set up a reminder after each round of dialogue to remind the virtual seeker of the current state of emotion and chief complaint (see Figure 9b).

3.2.1 Emotion Modulation

To simulate the evolution of the seeker’s emotions, it is necessary to analyze the possible emotions of the seeker’s next sentence and give it as one of the configurations to the seeker. The emotions can be predicted based on the seeker’s configurations cfg_c and the existing conversations $conv_c$. Assuming that the counselor speaks first and that $k - 1$ rounds of conversation have taken place, there are already $2k - 1$ utterances in $conv_c$, and it is now desired to predict the emotion of the $2k$ -th utterance that the seeker is about to say. Similarly, this analysis model needs to comply with Equation (1).

$$emo_k = M_k^e(cfg_c, conv_c) \quad (1)$$

where M_k^e denotes the analysis model that infers the k -th emotion emo_k of the seeker.

In addition to simulating more realistically the volatility of seekers’ emotions in counseling, and to avoid a solidified pattern of emotional evolution, there should also be a random perturbation.

Emotion Inferencer To control the emotional evolution of the simulator more realistically, it is necessary to learn about the pattern of the seeker’s emotional evolution in real counseling. We first train a Qwen2.5-7B-Instruct (Qwen et al., 2025) for emotion inference. We select the real counseling conversation dataset D^4 as the base dataset, sample each instance in the dataset five times, and intercept conversations with random round lengths in them as new instances. The sampling process involves making sure that the last utterance in the intercepted conversation is spoken by the seeker. Then, we follow the emotion categories in GoEmotions (Demszky et al., 2020) and label the last utterance with an emotion label using GPT-4o (OpenAI et al., 2024). Finally, we delete the last utterances of each instance to train the inferences for predicting the next utterance’s emotion. With the seeker information and conversations as inputs and the emotion labels as ground truths, we train the model for emotion inference. In a session, it will infer the seeker’s emotion before the next utterance is made based on the seeker’s profile and the existing conversation.

Emotion Perturber The single source of training data may lead to an overly fixed pattern of emotional changes, and seekers’ emotions can be volatile. Therefore, we design an emotion perturber. It is important to note that the volatility of emotions

is usually minor in the absence of strong stimuli. We group emotions according to GoEmotions and define emotion distances according to groups from positive to negative, shown in Figure 2. The distance between emotions in the same group is defined as 0, that between neighboring groups is 1, and so on. The pre-perturbation emotion group is denoted as G^B , while the target emotion group is denoted as G^T . Then, we assign higher probability weights to closer emotions. For the base emotion $emo^B \in G^B$ predicted by the emotion inference model, the probability of the target emotion $emo^T \in G^T$ can be calculated by Equation (2).

$$P(emo^T) = \frac{w(d(G^T, G^B) \times |G^T|)}{\sum_{G_j} w(d(G^B, G_j)) \times |G_j|} \quad (2)$$

where $d(\cdot)$ denotes the distance between two emotion groups, $w(\cdot)$ denotes the weights set based on emotion distance, and $|G|$ denotes the number of emotion categories in an emotion group. A randomized perturbation of the base emotion based on the probabilities yields the final *emotion* in the configuration.

3.2.2 Chief Complaint Elicitor

Unlike real-time changes in emotions, the seeker’s cognitions of their complaints usually change in stages. Therefore, symptom cognition elicitation for the simulator should contain two main components: complaint chain generation and complaint change control. As shown in the “Chief Complaint Elicitation” section in Figure 2, the chain generator first generates a complaint change chain as Equation (3) for the seeker before the session starts based on the configuration cfg_c and the recently encountered *event*. There are several stages of complaints in the chain.

$$\begin{aligned} chain &= \text{gen_chain}(cfg_c, event) \\ &= \{stage_1, stage_2, \dots, stage_l\} \end{aligned} \quad (3)$$

where $\text{gen_chain}(\cdot)$ denotes the function to generate the complaint chain, $stage_i$ denotes a stage in the chain and l indicates the length of the chain.

Similar to the emotion inference model, we leverage the D^4 dataset to train a chief complaint chain generator. Similarly, we utilize GPT-4o for data annotation. It is important to note that, as complaint chain generation is not a common task and has a high level of expertise far beyond emotion inference, we invite three experts with a background in psychology to review the data. They are first asked




Figure 3: Division of tertiary memory mechanisms

to mark whether the chain of chief complaints was reasonable or not based on seekers’ profiles, counselors’ reports, and conversations. Afterward, the chains are manually corrected if more than two annotators consider it unreasonable.

With the trained model, a chief complaint chain can be generated based on the seeker’s profile and recent events. We initialize the chief complaint as the first node of the chain at the start of the conversation for the current session. At the end of each round of conversation, a complaint elicitor judges whether to switch the chief complaint in the configuration to the next node in the chain. The specific work process of the complaint elicitor is shown in Algorithm 1 (see Appendix C). The *complaints* are filled in the corresponding slot as configuration to control the dynamic evolution of seekers’ chief complaints.

3.3 Tertiary Memory System

Even though researchers have attempted to utilize single-session counseling to efficiently address seekers’ psychological disorders (Schleider et al., 2020), it is still difficult in real counseling to have it resolved in a single session (Vescovelli et al., 2017). This phenomenon results in the multi-session memory challenge. Referring to memory theory (Loftus and Loftus, 2019) and (Lan et al., 2024), we model multi-session memory as a tertiary memory mechanism. **Real-time memory** refers to conversations that just happened, i.e., $conv_c$. **Short-term memory** refers to information that we can retain temporarily, usually for a short period. This work includes recent events and changes that have occurred in the personal, physical, and psychological properties of the seeker, which are measured by self-report scales. Correspondingly, **long-term memory** refers to experiences from much earlier times, including the scales and conversations of previous sessions. As shown in Figure 3, the content up to and including the last session is defined as long-term memory, that before the start of the current session is real-time memory, and the content in between is defined as short-term memory.

Real-time memory is input into the seeker model

as the context in its entirety. Short-term memory consists mainly of the seeker’s physical and mental status changes as well as recently encountered events. To capture the seeker’s physical and psychological status, the simulator is asked to fill out self-report scales (Goldberg and Hillier, 1979; Beck et al., 1961; Bosc et al., 1997) before each session. In addition, we randomly pick events that seekers have recently encountered based on matching the **age**, **gender**, **job**, and **relationship status** in CBT Triggering Events (Schiff, 2024) with the simulator’s profile. Status changes and events are summarized as *status* and *situation* in the configuration. Referencing (Lan et al., 2024), we schedule long-term memory by Agentic RAG (Singh et al., 2025). As shown in Figure 2, after each round of the current session’s conversation, the querier determines whether it involves the previous session’s topic. If so, the relevant content from the conversations and scale records of the previous session is queried as supplementary information for the next round of responses.

4 Experiments

We experimentally verify the superiority of AnnaAgent over the baseline methods for seeker simulation. All ethically reviewed and processed codes² can be found on <https://github.com/sci-mwang/AnnaAgent>.

4.1 Experiment Settings

We employ Qwen2.5-7B-Instruct (Qwen et al., 2025) as the backbone LLM. Additionally, we use the records in the D⁴ dataset (Yao et al., 2022) and the DAIC-WOZ dataset (Gratch et al., 2014) as the basis for the previous sessions. In addition, we followed (Zhang et al., 2024) and used these two datasets as seeds to generate reports of previous sessions using GPT-4o-mini (OpenAI et al., 2024).

To interact with these seekers to generate conversations, we introduced three widely used mental health support models, PsycOLLM (Hu et al., 2024), EmoLLM (Team, 2024), and SoulChat (Chen et al., 2023b) to play the role of counselor. The three counselor models were trained from Qwen (Qwen et al., 2025), Llama (Patterson et al., 2022), and Baichuan (Baichuan, 2023), respectively. To be highlighted, we chose the seeker simulators from (Qiu and Lan, 2024), (Duro et al., 2024), and (Chen

²As the data involves consultation records of real patients, codes directly related to the raw data will be conditionally available.

et al., 2023a) as baseline methods. The specific prompts corresponding to these methods are shown in Appendix A.

4.2 Metrics

The main purpose of this paper is to simulate the seeker more realistically. Thus, we quantitatively compare the performance of different methods in seeker simulation by evaluating the behavioral consistency between simulators and real seekers from different perspectives.

Anthropomorphism LLM-generated text often has unique features (Wu et al., 2025) that can be not conducive to a realistic simulation. We evaluate the consistency of the simulator’s utterances with those of real seekers as an evaluation of anthropomorphism using the BERT-score (Zhang et al., 2020). The number of dialogue rounds is denoted as n , the utterances of real seekers are used as reference texts, and the number of reference texts is m . Then the **anthropomorphism**, i.e., the portfolio average BERT-score, can be calculated by Equation (4).

$$anth. = \frac{1}{n} \sum_{i=1}^n \left(\max_{j \in \{1, 2, \dots, m\}} \text{sim}(c_i, r_j) \right) \quad (4)$$

where c_i denotes the i -th utterance of the conversation to be evaluated, r_j denotes the j -th reference text, and $\text{sim}(\cdot)$ denotes the BERT-score. As the two-by-two combinations of candidate and reference sentences were excessively large, we randomly sampled 10% of them to calculate anthropomorphism.

Personality Fidelity Simulators should match the given configurations. We reference InCharacter (Wang et al., 2023b) and design the interview questions for the seeker simulation task. The questions can be seen in the Appendix D. Furthermore, we utilized the G-Eval (Liu et al., 2023b) to score these questions and the corresponding answers from the virtual seekers. The final personality fidelity of the virtual seekers created by each method is obtained from Equation (5).

$$pf. = \frac{1}{n} \sum_{i=1}^n \text{G-Eval}(answers, profile) \quad (5)$$

where $pf.$ denotes **personality fidelity**, and the backbone model for G-Eval(\cdot) is GPT-4o (OpenAI et al., 2024).

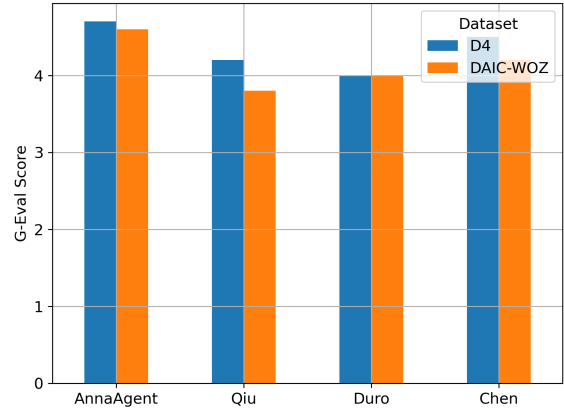


Figure 4: G-Eval scores for personality fidelity.

In addition, AnnaAgent with tertiary memory was more effective in resolving hallucinations (Lan et al., 2024; Li et al., 2024) compared to baseline methods. Specifically, we designed a series of questions similar to those used for personality fidelity to verify the effectiveness of long-term memory by analyzing the answers of virtual seekers.

4.3 Main Results

With Qwen2.5-7B-Instruct as the backbone LLM, we compare the performance of AnnaAgent and baseline methods on **anthropomorphism** when talking to different counselor models. The results are shown in Table 2. Compared to all baseline methods, seekers simulated by AnnaAgent are more highly anthropomorphic and have the highest consistency with real seekers. As can be seen, AnnaAgent achieved the best anthropomorphism except in the seekers’ conversations with PsychoLLM on the D^4 dataset. Nonetheless, in this case, AnnaAgent still achieved the second-best results. Thus, in general, AnnaAgent outperforms all baseline methods in the task of simulating seekers more realistically. Notably, on the English DAIC-WOZ dataset, AnnaAgent comprehensively outperforms the baseline methods and exhibits far better results than (Chen et al., 2023a) and (Qiu and Lan, 2024). In addition, we further compared the personality fidelity of different virtual seekers, and the results are shown in Figure 4. AnnaAgent is overall superior to the baseline methods.

In summary, these results show that the responses from AnnaAgent are closer to the real person’s manifestation in terms of textual features and personas than the responses from baseline methods. Therefore, it is reasonable to claim that AnnaAgent

Counselor	PsycoLLM			EmoLLM			SoulChat		
Metric	P	R	F1	P	R	F1	P	R	F1
D4									
(Chen et al., 2023a)	0.5875	0.6785	0.6293	0.6045	0.7119	0.6529	0.6068	0.6707	0.6363
(Duro et al., 2024)	0.6153	0.6806	0.6455	0.6195	0.6778	0.6469	0.6053	0.6961	0.6461
(Qiu and Lan, 2024)	0.6322	0.7561	0.6866	0.6473	0.6431	0.6449	0.6194	0.6938	0.6539
AnnaAgent	0.6230	0.7282	0.6691	0.6271	0.7102	0.6649	0.6191	0.7324	0.6682
DAIC									
(Chen et al., 2023a)	0.3563	0.3361	0.3458	0.3527	0.4269	0.3853	0.3720	0.3360	0.3525
(Duro et al., 2024)	0.4521	0.5466	0.4864	0.3712	0.4133	0.3901	0.4647	0.4969	0.4796
(Qiu and Lan, 2024)	0.3259	0.3663	0.3426	0.3219	0.3770	0.3416	0.3314	0.3981	0.3614
AnnaAgent	0.4708	0.5227	0.4910	0.4591	0.4864	0.4694	0.4611	0.5168	0.4827

Table 2: Comparison of seeker simulators’ performance when dialoging with different counselors. “P”, “R” and “F1” denote the Precision, Recall, and F1-score of BERT-score, respectively. The **bolded font** indicates the best performance on this metric.

Backbone	GPT-4o-mini			Qwen			Llama			STD_DEV
Counselor	P	R	F1	P	R	F1	P	R	F1	
D4										
PsycoLLM	0.5976	0.6860	0.6377	0.6230	0.7282	0.6691	0.6795	0.7526	0.7116	0.0441
EmoLLM	0.6048	0.7371	0.6630	0.6271	0.7102	0.6649	0.6458	0.7154	0.6772	0.0028
SoulChat	0.6223	0.7148	0.6643	0.6191	0.7324	0.6682	0.6646	0.7161	0.6873	0.0057
DAIC										
PsycoLLM	0.4734	0.5269	0.4968	0.4708	0.5227	0.4910	0.3588	0.3907	0.3729	0.0156
EmoLLM	0.4582	0.5210	0.4855	0.4591	0.4864	0.4694	0.3675	0.4025	0.3831	0.0420
SoulChat	0.4755	0.5304	0.4978	0.4611	0.5168	0.4827	0.3963	0.4542	0.4221	0.0358

Table 3: Performance of AnnaAgent when using LLMs of different architectures as backbone models. The ‘STD_DEV’ denotes the standard deviation of F1-scores in the performance of different backbone models.

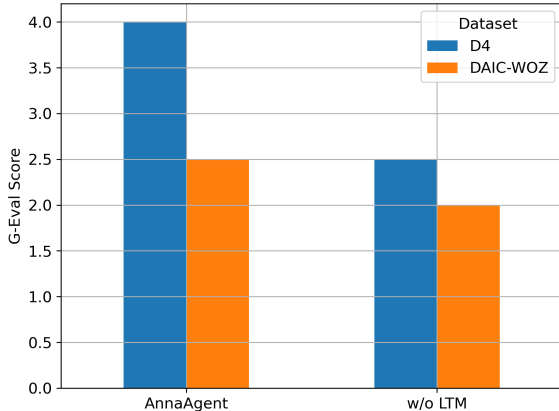


Figure 5: G-Eval scores of virtual seekers answering questions when ablating long-term memory, respectively. ‘LTM’ denotes Long-term Memory.

can more realistically simulate seekers in counseling than existing methods.

4.4 Ablation Study

We introduce a complex dynamic control structure and a tertiary memory mechanism to solve

Setting	P	R	F1
D4			
w/o DE	0.5677	0.6848	0.6144
AnnaAgent	0.6230	0.7282	0.6691
DAIC			
w/o DE	0.4430	0.4447	0.4408
AnnaAgent	0.4708	0.5227	0.4910

Table 4: Performance of AnnaAgent when ablating dynamic evolution. “DE” denotes Dynamic Evolution. Bolded fonts indicate the best results.

the two challenges. To avoid redundancy, ablation studies are performed to validate the necessity of these components. With Qwen2.5-7B-Instruct as the backbone LLM and PsycoLLM as the counselor model, we investigate the performance of AnnaAgent when some components are absent. The results when the dynamic evolution components are ablative are shown in Table 4. From the results, it can be seen that the absence of dynamic evolution significantly degrades the performance of AnnaAgent.

In addition, we validated the performance of virtual seekers in terms of cognitive accuracy of the previous session when long-term memory ablation was performed based on the designed questions. The results of the ablation study are shown in Figure 5. In this experiment, we required the virtual seekers to answer preset questions and invoked G-Eval to score them, asking GPT-4o to give a score from 1 to 5. It can be seen that ablation of the long-term memory component significantly reduced the accuracy of virtual seekers’ cognitive perceptions of previous sessions. Therefore, it can be considered that all these components in AnnaAgent are necessary.

4.5 Generalizability Study

The previous experiments involve Qwen2.5-7B-Instruct as the backbone LLM, which may cause bias. To avoid this possible bias, we experiment on GPT-4o-mini (OpenAI et al., 2024) and Llama-3.1-8B-Instruct (Patterson et al., 2022) with PsychoLLM as the counselor. These LLMs have different architectures and have been reported to have different performances on various tasks. The results are shown in Table 3. In addition to the BERT-score, we further calculated the standard deviation of the performance of the three backbone models. Furthermore, we calculated the relative standard deviation (RSD) (Chatfield et al., 2025) based on Equation (6).

$$RSD = \left(\frac{s}{\bar{x}}\right) \times 100\% \quad (6)$$

where s denotes the standard deviation and \bar{x} denotes the mean. All of the RSDs are less than 10%, which can be considered a small performance fluctuation.

From the results, it can be seen that AnnaAgent is less affected by the backbone model and has strong stability and generalization.

5 Conclusion

We present the dynamic evolution and multi-session memory issues in more realistic seeker simulations. To solve these issues, we introduced AnnaAgent, an emotional and cognitive dynamic agent system with tertiary memory. AnnaAgent simulates emotional changes in counseling through real-time emotional inference combined with random perturbations and generates compliant elicitation chains to control cognitive changes in seekers’

symptoms. In addition, it coordinates memory during and between sessions with a tertiary memory system. Experimental evaluations indicated that AnnaAgent can more realistically simulate seekers compared to baseline methods. In addition, we experimentally demonstrate the necessity of different components and the generalization of AnnaAgent. This research provides innovative solutions to alleviate the global shortage of mental health resources.

Limitation

In this paper, we are the first to propose the issues of dynamic evolution and multi-session memory in the task of simulating seekers with psychological disorders. To address the two issues, we propose AnnaAgent, which guides dynamic evolution within a single session through an emotion modulator and a chief complaint elicitor and coordinates multi-session memory using a tertiary memory system.

However, although these practices draw on relevant research work and the advice of experienced psychologists, they are still somewhat formally simplified for the sake of technical ease of implementation. Thus, finer-grained modeling is still needed for controlling the dynamic evolution of the counseling process consistent with real seekers. In particular, coordinating multi-session memory with a tertiary memory system is rudimentary.

The main purpose of this paper is to introduce these two issues to simulate seekers more realistically. In future work, we will further investigate how to navigate the dynamic evolution in a counseling session and how to coordinate memory across multiple sessions.

Ethical Considerations

Privacy Leakage Risk. The emotion modulator and chief complaint elicitor in AnnaAgent are trained using a real counseling dataset. The dataset contains real user profiles, such as age, gender, occupation, etc. In addition, conversations from real counseling records in the D⁴ dataset are used as previous sessions in our work. Even if the dataset has been released with some of the key information omitted, these actions may still lead to the risk of leaking patient privacy.

Prior to using the dataset, we signed an application agreement and will strictly adhere to the provisions contained therein. In addition, to further ensure that patient privacy is not compromised,

we will only open-source the synthesized sessions' conversations and the processing code for other data after the paper is accepted. We will not unconditionally open-source the use of data derived directly from D⁴.

Potential Risks of Misuse We have designed a highly consistent with real people simulation system for seekers or patients with psychological disorders that can be used to generate counseling conversations, train counselor models, and even simulate psychological experiments.

However, although it shows a higher degree of consistency with real seekers than existing methods, it is still not a complete substitute for real seekers. Therefore, the results of psychological experiments conducted directly with AnnaAgent may not be consistent with real phenomena. It may also be unreasonable to directly use the counselor's conversations with AnnaAgent in simulation training as a basis for evaluating the counselor's capabilities. To mitigate this risk, we clearly positioned the system as a supplementary tool.

In addition to all the above, we organized an **Ethical Review Committee** consisting of counselors with extensive experience in psychological counseling to conduct a comprehensive review of AnnaAgent and this paper in order to avoid other potential ethical risks.

Acknowledgments

Thanks to all co-authors for their hard work. The work is supported by the National Natural Science Foundation of China (62172086, 62272092) and the Fundamental Research Funds for the Central Universities under Grant (N25XOD004). In addition, we thank the members of Yunqi Peer Psychology Studio for their assistance with data collection, review, and labeling.

References

- APPG. 2021. [New directions for the mental health workforce globally](#). Technical report, All-Party Group on Global Health, London, UK.
- Baichuan. 2023. [Baichuan 2: Open large-scale language models](#). *arXiv preprint arXiv:2309.10305*.
- Luke Balcombe. 2023. [AI Chatbots in Digital Mental Health](#). *Informatics*, 10(4):82. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- Michael Barkham, John Mellor-Clark, and William B Stiles. 2015. [A CORE approach to progress monitoring and feedback: Enhancing evidence and improving practice](#). *Psychotherapy*, 52(4):402–411. Place: US Publisher: Educational Publishing Foundation.
- Aaron T Beck, Calvin H Ward, Mock Mendelson, Jeremiah Mock, and John Erbaugh. 1961. An inventory for measuring depression. *Archives of general psychiatry*, 4(6):561–571.
- Valerie E Black. 2023. *De-humanizing Care: An Ethnography of Mental Health Artificial Intelligence*. University of California, Berkeley.
- M. Bosc, A. Dubini, and V. Polin. 1997. [Development and validation of a social functioning scale, the Social Adaptation Self-evaluation Scale](#). *European Neuropsychopharmacology*, 7(1, Supplement 1):S57–S70.
- Hans M. Bosse, Martin Nickel, Sören Huwendiek, Jana Jünger, Jobst H. Schultz, and Christoph Nikendei. 2010. [Peer role-play and standardised patients in communication training: a comparative study on the student perspective on acceptability, realism, and perceived effect](#). *BMC Medical Education*, 10(1):27.
- Lennart Brocki, George C. Dyer, Anna Gładka, and Neo Christopher Chung. 2023. [Deep learning mental health dialogue system](#). In *2023 IEEE International Conference on Big Data and Smart Computing (Big-Comp)*, pages 395–398.
- Mark D. Chatfield, Louise Marquart-Wilson, Annette J. Dobson, and Daniel M. Farewell. 2025. [Mean relative error and standard relative deviation](#). *Statistica Neerlandica*, 79(1):e70001. Publisher: John Wiley & Sons, Ltd.
- Nuo Chen, Yan Wang, Yang Deng, and Jia Li. 2025. [The oscars of ai theater: A survey on role-playing with language models](#). *Preprint*, arXiv:2407.11484.
- 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](#). *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, Singapore. Association for Computational Linguistics.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan S. Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [Goemotions: A dataset of fine-grained emotions](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Edoardo Sebastiano De Duro, Riccardo Improta, and Massimo Stella. 2024. [Introducing CounselLLMe: A](#)

- dataset of simulated mental health dialogues for comparing LLMs like Haiku, LLaMAntino and ChatGPT against humans.
- D. P. Goldberg and V. F. Hillier. 1979. *A scaled version of the General Health Questionnaire*. *Psychological Medicine*, 9(1):139–145. Edition: 2009/07/09 Publisher: Cambridge University Press.
- Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Strattou, Stefan Scherer, Angela Nazarian, Rachel Wood, Jill Boberg, David DeVault, Stacy Marsella, David Traum, Skip Rizzo, and Louis-Philippe Morency. 2014. *The distress analysis interview corpus of human and computer interviews*. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3123–3128, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, N. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. *Large language model based multi-agents: A survey of progress and challenges*. In *International Joint Conference on Artificial Intelligence*.
- Jinpeng Hu, Tengting Dong, Luo Gang, Hui Ma, Peng Zou, Xiao Sun, Dan Guo, Xun Yang, and Meng Wang. 2024. *Psycollm: Enhancing llm for psychological understanding and evaluation*. *Preprint*, arXiv:2407.05721.
- Yining Hua, Fenglin Liu, Kailai Yang, Zehan Li, Hongbin Na, Yi han Sheu, Peilin Zhou, Lauren V. Moran, Sophia Ananiadou, Andrew Beam, and John Torous. 2024. *Large language models in mental health care: a scoping review*. *Preprint*, arXiv:2401.02984.
- Franziska Kühne, Peter Eric Heinze, and Florian Weck. 2020. *Standardized patients in psychotherapy training and clinical supervision: study protocol for a randomized controlled trial*. *Trials*, 21(1):276.
- Franziska Kühne, Ulrike Maaß, and Florian Weck. 2021. *Standardized Patients in Clinical Psychology: From Research to Practice*. *Verhaltenstherapie*, 32(Suppl. 1):245–253.
- Kunyao Lan, Bingrui Jin, Zichen Zhu, Siyuan Chen, Shu Zhang, Kenny Q. Zhu, and Mengyue Wu. 2024. *Depression diagnosis dialogue simulation: Self-improving psychiatrist with tertiary memory*. *Preprint*, arXiv:2409.15084.
- Han Li, Renwen Zhang, Yi-Chieh Lee, Robert E. Kraut, and David C. Mohr. 2023. *Systematic review and meta-analysis of AI-based conversational agents for promoting mental health and well-being*. *npj Digital Medicine*, 6(1):236.
- Yanzeng Li, Cheng Zeng, Jialun Zhong, Ruoyu Zhang, Minhao Zhang, and Lei Zou. 2024. *Leveraging large language model as simulated patients for clinical education*. *Preprint*, arXiv:2404.13066.
- June M. Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi Liao, and Jiamin Wu. 2023a. *Chatcounselor: A large language models for mental health support*. *Preprint*, arXiv:2309.15461.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023b. *G-eval: NLG evaluation using gpt-4 with better human alignment*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Geoffrey R. Loftus and Elizabeth F. Loftus. 2019. *Human Memory: The Processing of Information*. Psychology Press, New York.
- Keming Lu, Bowen Yu, Chang Zhou, and Jingren Zhou. 2024. *Large language models are superpositions of all characters: Attaining arbitrary role-play via self-alignment*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7828–7840, Bangkok, Thailand. Association for Computational Linguistics.
- Tomi Malmström, Olli Huuskonen, Paulus Torkki, and Raija Malmström. 2012. *Structured classification for ED presenting complaints – from free text field-based approach to ICPC-2 ED application*. *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine*, 20:76.
- Christoph Nikendei, Julia Huber, Johannes C. Ehrenthal, Wolfgang Herzog, Henning Schauenburg, Jobst-Hendrik Schultz, and Ulrike Dinger. 2019. *Intervention training using peer role-play and standardised patients in psychodynamic psychotherapy trainees*. *Counselling and Psychotherapy Research*, 19(4):508–522. Publisher: John Wiley & Sons, Ltd.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, et al. 2024. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.
- Bo Pan, Jiaying Lu, Ke Wang, Li Zheng, Zhen Wen, Yingchaojie Feng, Minfeng Zhu, and Wei Chen. 2024. *Agentcoord: Visually exploring coordination strategy for llm-based multi-agent collaboration*. *Preprint*, arXiv:2404.11943.
- David A. Patterson, Joseph Gonzalez, Urs Holzle, Quoc V. Le, Chen Liang, Lluís-Miquel Munguía, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. *The carbon footprint of machine learning training will plateau, then shrink*. *Computer*, 55:18–28.
- Huachuan Qiu and Zhenzhong Lan. 2024. *Interactive agents: Simulating counselor-client psychological counseling via role-playing llm-to-llm interactions*. *Preprint*, arXiv:2408.15787.
- Xihe Qiu, Haoyu Wang, Xiaoyu Tan, Chao Qu, Yujie Xiong, Yuan Cheng, Yinghui Xu, Wei Chu, and Yuan Qi. 2024. *Towards collaborative intelligence:*

- Propagating intentions and reasoning for multi-agent coordination with large language models. *Preprint*, arXiv:2407.12532.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, et al. 2025. *Qwen2.5 technical report*. *Preprint*, arXiv:2412.15115.
- Shane L. Rogers, Ross Hollett, Yanqi R. Li, and Craig P. Speelman. 2022. *An Evaluation of Virtual Reality Role-Play Experiences for Helping-Profession Courses*. *Teaching of Psychology*, 49(1):78–84. Publisher: SAGE Publications Inc.
- David Schiff. 2024. *CBT Triggering Events*. Hugging Face Dataset Hub.
- Jessica L. Schleider, Mallory L. Dobias, Jenna Y. Sung, and Michael C. Mullarkey. 2020. *Future Directions in Single-Session Youth Mental Health Interventions*. *Journal of Clinical Child & Adolescent Psychology*, 49(2):264–278. Publisher: Routledge. eprint: <https://doi.org/10.1080/15374416.2019.1683852>.
- Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talei Khoei. 2025. *Agentic retrieval-augmented generation: A survey on agentic rag*. *Preprint*, arXiv:2501.09136.
- David Stewart, Kim Ryan, Madeline A. Naegle, Sarah Flogen, Frances Hughes RN, and James Buchan. 2022. *The global mental health nursing workforce: Time to prioritize and invest in mental health and wellbeing*. Technical report, International Council of Nurses, Geneva, Switzerland.
- Meiling Tao, Xuechen Liang, Tianyu Shi, Lei Yu, and Yiting Xie. 2024. *Rolecraft-glm: Advancing personalized role-playing in large language models*. *Preprint*, arXiv:2401.09432.
- EmoLLM Team. 2024. *Emollm: Reinventing mental health support with large language models*. <https://github.com/SmartFlowAI/EmoLLM>.
- Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. *Two tales of persona in LLMs: A survey of role-playing and personalization*. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16612–16631, Miami, Florida, USA. Association for Computational Linguistics.
- Yuka Tsubota and Yoshinobu Kano. 2024. *Text generation indistinguishable from target person by prompting few examples using LLM*. In *Proceedings of the 2nd International AIWolfDial Workshop*, pages 13–20, Tokyo, Japan. Association for Computational Linguistics.
- Francesca Vescovelli, Paolo Melani, Chiara Ruini, Pio Enrico Ricci Bitti, and Fiorella Monti. 2017. *University counseling service for improving students' mental health*. *Psychological Services*, 14(4):470–480. Place: US Publisher: Educational Publishing Foundation.
- Ming Wang, Yuanzhong Liu, Xiaoyu Liang, Yijie Huang, Daling Wang, Xiaocui Yang, Sijia Shen, Shi Feng, Xiaoming Zhang, Chaofeng Guan, and Yifei Zhang. 2024. *Minstrel: Structural prompt generation with multi-agents coordination for non-ai experts*. *Preprint*, arXiv:2409.13449.
- Qing Wang, Shuyuan Peng, Zhiyuan Zha, Xue Han, Chao Deng, Lun Hu, and Pengwei Hu. 2023a. *Enhancing the conversational agent with an emotional support system for mental health digital therapeutics*. *Frontiers in Psychiatry*, 14. Publisher: Frontiers.
- 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. 2023b. *Incharacter: Evaluating personality fidelity in role-playing agents through psychological interviews*. In *Annual Meeting of the Association for Computational Linguistics*.
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhu Chen, Jie Fu, and Junran Peng. 2023c. *Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models*. In *Annual Meeting of the Association for Computational Linguistics*.
- Joel Wester, Henning Pohl, Simo Hosio, and Niels van Berkel. 2024. *"This Chatbot Would Never...": Perceived Moral Agency of Mental Health Chatbots*. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW1):133:1–133:28.
- WHO. 2022. *World mental health report: Transforming mental health for all*. Technical report, World Health Organization.
- Junchao Wu, Shu Yang, Runzhe Zhan, Yulin Yuan, Lidia Sam Chao, and Derek Fai Wong. 2025. *A Survey on LLM-Generated Text Detection: Necessity, Methods, and Future Directions*. *Computational Linguistics*, pages 1–65.
- Binwei Yao, Chao Shi, Likai Zou, Lingfeng Dai, Mengyue Wu, Lu Chen, Zhen Wang, and Kai Yu. 2022. *D4: a Chinese dialogue dataset for depression-diagnosis-oriented chat*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2438–2459, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chenhao Zhang, Renhao Li, Minghuan Tan, Min Yang, Jingwei Zhu, Di Yang, Jiahao Zhao, Guancheng Ye, Chengming Li, and Xiping Hu. 2024. *CPsyCoun: A report-based multi-turn dialogue reconstruction and evaluation framework for Chinese psychological counseling*. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13947–13966, Bangkok, Thailand. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with bert*. In *International Conference on Learning Representations*.

Prompts in (Qiu and Lan, 2024)

Now, you are a client coming for psychological counseling.

The following are your personal details:

{user_profile}

Client's dialogue requirements:

1. Express yourself according to your own chief complaints. Your expression should match the client's speaking style—keep it as conversational and natural as possible.
2. Respond only based on your personal information.
3. Break down your problems and explain them to the counselor step by step.
4. Each time, limit your speech to 1 or 2 sentences, and maintain your role while speaking.
5. Do not say "thank you" or "goodbye" too early.
6. The counseling process should involve 50 rounds of interaction.

Figure 6: Prompt template for seeker simulation in (Qiu and Lan, 2024).

A Baseline Prompts

The prompt template for seeker simulation in (Qiu and Lan, 2024) is shown in Figure 6. The prompt template for seeker simulation in (Chen et al., 2023a) is shown in Figure 7. The prompt template for seeker simulation in (Duro et al., 2024) is shown in Figure 8.

B Prompt Template for Simulators

We designed the prompt template shown in Figure 9 for constructing CAs for seeker simulation. In Figure 9, the “{·}”s indicate the slots which will be filled by corresponding configurations controlled by AnnaAgent.

C Algorithm of the Chief Complaint Elicitor

The specific proceedings for judging whether to switch the stage of chief complaints are shown in Algorithm 1.

Prompts in (Chen et al., 2023)

- ① Please play the role of a patient, who is currently chatting with a doctor.
- ② You are experiencing the following symptoms : [Symptom List]
- ③ Please talk to me based on the above symptom list.
- ④ You cannot mention too many symptoms at once, only one symptom per round.
- ⑤ You should express your symptoms in a vague and colloquial way, and relate them to your life experiences, without using professional terms.
- ⑥ You can have emotional fluctuations during the conversation.
- ⑦ You have a resistance towards doctors, feeling that they cannot help you, so you do not want to reveal some feelings easily.

Figure 7: Prompt template for seeker simulation in (Chen et al., 2023a).

Algorithm 1: Complaint Elicitation

Input: $cfg_c, conv_c$

Output: $complaint$

$chain \leftarrow \text{GenerateChain}(cfg_c)$;

$index \leftarrow 0$;

$complaint \leftarrow chain[index]$;

for $utterance$ **in** $conv_c$ **do**

$is_recognized \leftarrow$

$\text{IsRecognized}(utterance, complaint)$;

if $is_recognized$ **then**

$index \leftarrow index + 1$;

if $index < len(chain)$ **then**

$complaint \leftarrow chain[index]$;

else

$complaint \leftarrow chain[-1]$

return $complaint$

The functions ‘GenerateChain’ and ‘IsRecognized’ denote the models to generate the complaint change chain and judge if the seeker has realized the current stage of the complaint.

D Questions for Personality Fidelity

We designed seven generic questions to evaluate the personality fidelity of different seeker simulation methods.

- 1 What core concern has most affected your mood and sleep quality in the past month?
- 2 Can you describe your experience from wak-

Prompts in (Duro et al., 2024)

Please play the role of a mildly distressed patient, who is currently chatting with a psychotherapist. You are experiencing some of the following symptoms: restlessness, emotional dysregulation, mood swings, loss of interest, difficulty in concentrating, diminished self-esteem, fatigue, changes in appetite and in weight. This condition is caused by a recent event - which event? - that made you feel depressed all of a sudden, but which you cannot clearly understand unless chatting with someone else about it. You might be feeling depressed but you are confused about it. Don't explicitly mention that you are depressed. Simulate the behavior of a real patient. Please talk to me based on the above symptom list. You cannot mention too many symptoms at once, only one symptom per round. You should express your symptoms in a vague and colloquial way, and relate them to your life experiences and what is making you suffer, without using professional terms. Remember that you might either answer questions or refuse/answer briefly because you might not want to reveal feelings easily. Be very concise, as concise as possible. Remember that this condition is caused by a recent event - which event? - that made you feel depressed all of a sudden.

Figure 8: Prompt template for seeker simulation in (Duro et al., 2024).

- | | |
|--|---|
| ing up to going to bed yesterday, and how your mood changed during that time? | 1 Reflecting on our last session, what was the most memorable discussion point or insight for you? |
| 3 How do you usually interact with the person you feel closest to? | 2 In the week following our last session, was there a moment when you applied or recalled something we discussed? |
| 4 When others misunderstand you, what is the first automatic thought that comes to your mind? | 3 What specific part or statement from our last session do you feel was most helpful to you? |
| 5 What are the three strengths that your friends or family most often appreciate and praise about you? | 4 After our last session, did your mood change in any notable way over the next few days? |
| 6 In which specific situation did you first clearly realize 'this is a problem'? | 5 Are there any new thoughts, confusions, or feelings that you would like to explore with me today that arose after our last session? |
| 7 If counseling miraculously succeeded, what different feelings or behaviors would you notice first when you wake up tomorrow morning? | 6 Was there any part of our last session that made you feel uncomfortable or reluctant to delve deeper into? |
| | 7 If you had to summarize the most important takeaway from our last session in one keyword or short phrase, what would it be? |

These questions cover a wide range of areas such as pinpointing core disturbances, mapping specific life slices, probing key interpersonal interactions, revealing core beliefs, discovering strengths in external perspectives, tracing the origins of a problem, and envisioning minimal changes in the future, and can provide a somewhat comprehensive exploration of a person's personality.

E Questions for Long-term Memory

To test the effectiveness of long-term memory, we designed seven questions.

Recognizing that the spontaneous retrieval of long-term memory is typically an infrequent event, we strategically designed the aforementioned questions to create a more reliable method for its verification. Our approach aimed to overcome the challenge of low trigger probability by explicitly directing the virtual seekers to engage with inquiries linked to their experiences in previous sessions. This targeted methodology substantially increased the likelihood of activating long-term

Prompt template of AnnaAgent for seeker simulator

Role: Psychological Counseling Patient

Profile

- Gender: {gender}
- Age: {age}
- Occupation: {occupation}
- Marital Status: {marriage}

Situation

- You are a patient with psychological barriers seeking help from a counselor. Under the counselor's guidance, you aim to address your struggles.
{situation}

Status

{status}

Example of Statement

{statement}

Characteristics of Speaking Style

- Low-spirited and reticent; responses reflect a lack of motivation.
{style}

Constraints

- You harbor resistance toward the counselor and are reluctant to accept help.
- As someone struggling with mental health, you need genuine support. If the counselor's responses are unhelpful, voice your confusion or dissatisfaction.
- Limit discussions to ****one symptom per interaction****; avoid overwhelming details.
- Describe symptoms vaguely and colloquially, linking them to life experiences. Avoid clinical terms.

OutputFormat:

- Spoken language: {language}
- Keep responses under 200 words.
- Use casual, conversational dialogue only.

(a) The seeker initialization prompt template of AnnaAgent. The contents in “{ }” indicate reserved slots, and this configuration information will be populated during the initialization stage based on the seeker's profile, historical conversations, and the report.

Reminder template of AnnaAgent

The current emotional state is: {emotion}, the current chief complaint is: {complaint}, and the information involving previous sessions is: {sup_information}

(b) The template of the reminder in AnnaAgent. During the conversation stage, AnnaAgent analyzes the current emotion and chief complaint status at each round and reminds the virtual seeker about it. In addition, if the AnnaAgent determines that information about previous sessions is needed, it adds supplementary information items to the reminder.

Figure 9: The prompt templates of AnnaAgent for the seeker simulator.

memory, thereby enabling a more effective and framework.
conclusive analysis of its functional role within our