

AgentMental: An Interactive Multi-Agent Framework for Explainable and Adaptive Mental Health Assessment

Jinpeng Hu¹, Ao Wang¹, Qianqian Xie³, Zhuo Li², Hui Ma¹, Dan Guo^{1,*}

¹Hefei University of Technology

²The Chinese University of Hong Kong, Shenzhen

³Wuhan University

{jinpenghu, huima, guodan}@hfut.edu.cn, 2024170834@mail.hfut.edu.cn, 221019088@link.cuhk.edu.cn, xqq.sincere@gmail.com

Abstract

Mental health assessment is crucial for early intervention and effective treatment, yet traditional clinician-based approaches are limited by the shortage of qualified professionals. Recent advances in artificial intelligence have sparked growing interest in automated psychological assessment, yet most existing approaches are constrained by their reliance on static text analysis, limiting their ability to capture deeper and more informative insights that emerge through dynamic interaction and iterative questioning. Therefore, in this paper, we propose a multi-agent framework for mental health evaluation that simulates clinical doctor-patient dialogues, with specialized agents assigned to questioning, adequacy evaluation, scoring, and updating. We introduce an adaptive questioning mechanism in which an evaluation agent assesses the adequacy of user responses to determine the necessity of generating targeted follow-up queries to address ambiguity and missing information. Additionally, we employ a tree-structured memory, where the root node encodes the user’s basic information, while child nodes (e.g., topic and statement) organize key information according to distinct symptom categories and interaction turns. This memory is dynamically updated throughout the interaction to reduce redundant questioning and further enhance the information extraction and contextual tracking capabilities. Experimental results on the DAIC-WOZ dataset illustrate the effectiveness of our proposed method, which achieves better performance than existing approaches. Our code is released at <https://github.com/MindIntLab-HFUT/AgentMental>.

Introduction

Mental health issues significantly impact global well-being and represent a growing public health concern. According to the World Health Organization (WHO), over 300 million people worldwide suffer from depression, and the prevalence of anxiety disorders reaches up to 4% of the global population. Mental health assessment plays a vital role in facilitating early detection and timely intervention for mental health disorders. However, traditional mental health assessment methods rely on clinicians conducting structured interviews, while such methods usually require professional personnel, which makes it difficult to access assessment

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*Corresponding author.

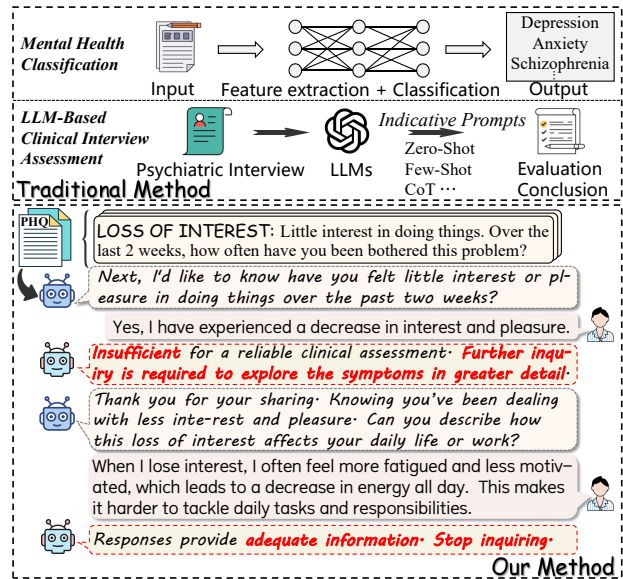


Figure 1: Comparison between our method and existing methods.

services in resource-scarce regions. Thus, many automated methods based on machine learning have been proposed to evaluate psychological states.

Prior efforts have used text classification techniques in detecting a range of psychological disorders, including depression, suicidal ideation, bipolar disorder, and anxiety (Ji et al. 2022; Khowaja et al. 2025; Agarwal, Dias, and Dollfus 2024). For example, Khowaja et al. (2025) introduced an emotion-aware contrastive network for mental disorder detection from social media posts. Agarwal, Dias, and Dollfus (2024) designed multi-view architectures that separately model therapist and patient utterances to enhance automated depression estimation. Recently, large language models (LLMs) have shown significant success in text comprehension and generation, leading to the development of LLM-based approaches for mental health (Yang et al. 2024b; Xu et al. 2024; Xiao et al. 2024; Hu et al. 2024). For example, Yang et al. (2024b) constructed a mental health instruction dataset and trained MentaLLaMA for interpretable mental health analysis. Ohse et al. (2024) investigated the perfor-

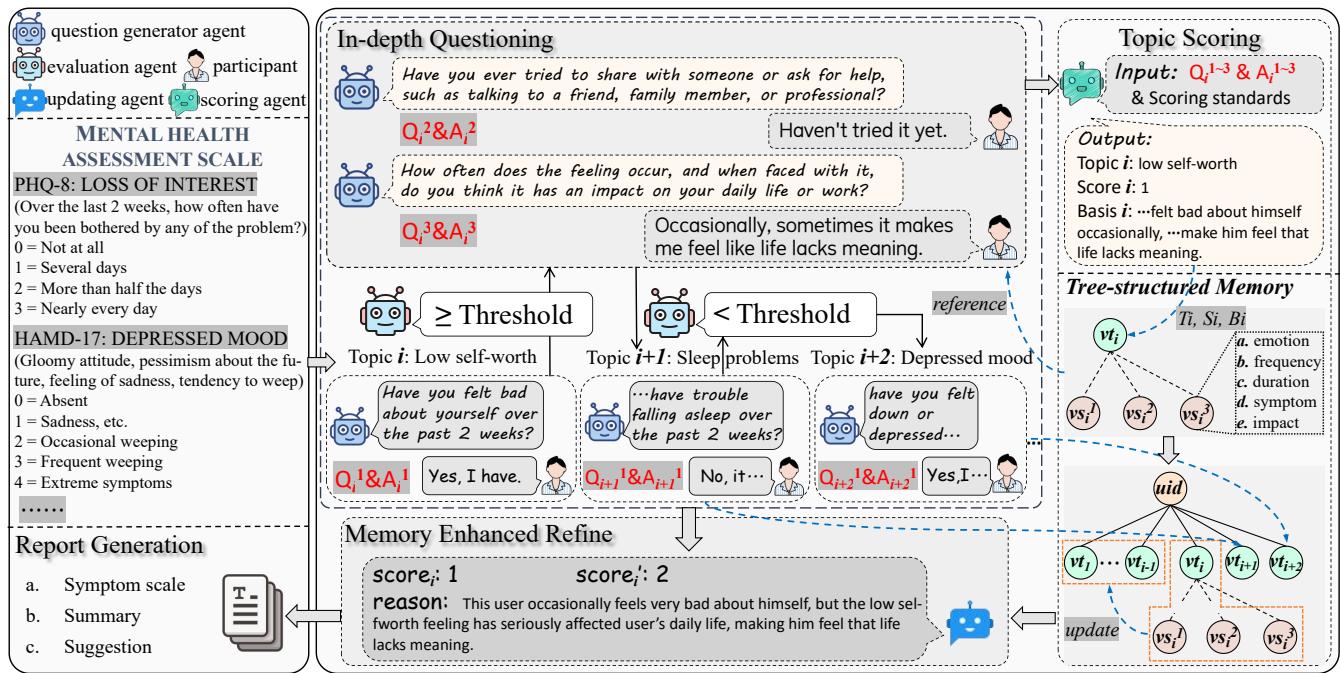


Figure 2: The overall framework of AgentMental for mental health assessment scales.

mance of GPT-4 in identifying social anxiety from semi-structured clinical interview data. Chen et al. (2024b) transformed the clinical interview into an expertise-inspired directed acyclic graph to improve depression detection.

Although these studies have achieved considerable progress, there are still two challenges. First, most existing approaches focus on passive analysis of static data, limiting their ability to handle scenarios where user-provided information is incomplete or insufficient due to factors such as user resistance or other contextual constraints. Second, although standardized diagnostic scales are essential for depression assessment and clinical interpretability, the associated counseling process often involves heterogeneous and multi-topic data, making it challenging for models to derive a holistic and coherent understanding of the user’s mental state. Recent developments in LLM-based role-playing have demonstrated robust performance across diverse tasks, providing a technological foundation for effectively simulating interactive dialogues and thereby enabling the acquisition of more comprehensive information (Yu et al. 2025).

Therefore, in this paper, we propose a multi-agent-based interactive framework to simulate clinical interview scenarios for mental health assessment (AgentMental), wherein dedicated agents are responsible for distinct roles such as question generation, response evaluation, score assignment, and updating. Prior to score assignment, we use an adaptive questioning mechanism to alleviate the issue of insufficient information in user responses, wherein an evaluation agent assesses the completeness of user responses to determine whether the question generation agent should produce targeted follow-up queries aimed at eliciting more comprehensive information. Furthermore, given the contextual dependencies inherent in scale-based assessments, we introduce a

tree-based memory to systematically record critical information, in which the root node contains fundamental user information, and topic nodes store item-level assessment scores and dialogue summaries, and statement nodes preserve turn-level evidence, such as symptom frequency, duration, etc. By dynamically updating this memory structure throughout the interactive process, the proposed method effectively reduces redundant questioning, facilitates coherent information flow across different topics, and consequently enhances the overall accuracy of the psychological assessment. In addition, we conduct comprehensive evaluations on the DAIC-WOZ dataset, and the results demonstrate that AgentMental consistently outperforms existing studies.

Methodology

In this section, we introduce the details of the proposed multi-agent-based framework designed for automated mental health assessment, including tree-structured memory, adaptive question generation, and memory-augmented scoring. The overall framework is shown in Figure 2.

Overall Workflow

Motivation. Based on discussions with psychologists, real-world psychological scale evaluations typically involve clinicians asking patients topic-related questions and subsequently assigning a rating score based on their responses (Sommers-Flanagan, Zeleke, and Hood 2015). This process allows psychologists to assess various aspects of a patient’s mental state in a structured yet flexible manner. Following this, to simulate this real-world evaluation process, we generate an initial question for each topic to ensure comprehensive coverage of key assessment criteria. Based on the

user response, we dynamically determine whether additional follow-up questions are needed for the current topic, facilitating a more in-depth evaluation of the individual’s mental state. Upon completion of the conversation, the system adaptively updates the scores along with explanatory reasoning and generates a psychological report that includes a dialogue summary and personalized recommendations.

Preparation. AgentMental incorporates four core agents, each fulfilling distinct roles within the system, including question generator agent AG_q , evaluation agent AG_{ev} , scoring agent AG_s , updating agent AG_u . The question generator agent is responsible for generating core assessment questions to ensure the structured progression of the dialogue. The evaluation agent performs necessity analysis and demand judgment, evaluating the completeness and depth of user responses to determine whether further questioning is required. The scoring agent focuses on analyzing user responses and generating quantitative scores for symptom severity based on clinical standards of the scale. The updating agent is responsible for information integration and semantic compression, as well as adjusting assessment scores based on the aggregated interaction history.

Tree-Structured Memory

Given the interdependencies between scale topics, we use a tree-structured memory that evolves throughout the conversation to record critical information, enhance information retention and reduce redundancy of questions. The memory consists of three types of nodes: a user node representing basic user attributes (e.g., occupation, sex, age and so on), topic nodes vt corresponding to individual assessment topics (e.g., score, summary and so on), and statement nodes vs encoding salient information extracted from user responses (e.g., emotions, frequency, duration, symptoms, impact). At each dialogue turn, several key dimensions of information, including emotion, frequency, duration, symptoms, and impact, are extracted by a LLM and a corresponding statement node is constructed to retain this information. Let i denote the index of the current topic. Upon completion of a i -th topic T_i , the following operations are performed:

- Instantiates a topic node vt_i recording the topic-specific assessment score S_i and behavioral summary B_i generated by the agent AG_s .
- Updates summary sections of previously instantiated topic nodes vt_1, \dots, vt_{i-1} to incorporate new relevant information.
- Establishes edges to all statement nodes associated with topic T_i .

Besides, each topic node is also connected to the user node, thereby enabling context-aware filtering of inappropriate questions. For instance, querying a primary or secondary school student about work-related stress would be considered unsuitable, as individuals in this demographic are typically not engaged in formal employment.

Adaptive Question Generation

We first extract the topic for each scale, such as PHQ-8 (Kroenke et al. 2009) which contains 8 topics. Then the

question generator agent AG_q generates topic-specific questions to simulate interactions between psychologists and help-seekers, aiming to guide the user to give responses that can reflect their status on the given topic. Besides, in clinical settings, it is common for patients to face difficulties in clearly expressing their thoughts and emotions due to cognitive impairments, resistance, or communication barriers. For instance, a patient with severe depression may respond with minimal or vague answers, such as simply saying “I do have insomnia” when asked about their symptoms, without providing key details such as the frequency or duration of insomnia. Therefore, more specific inquiries are needed to elicit more comprehensive information from the patient, such as “Does this happen frequently?” or “How long does it take you to fall asleep?”. To replicate this dynamic interaction, we incorporate a context-aware follow-up mechanism, which can detect vague, ambiguous, or incomplete responses and generate appropriate clarification questions to guide the conversation. In detail, we use the evaluation agent AG_{ev} to systematically assess patient responses in terms of information completeness and symptom severity. It produces a necessity score on a scale from 0 to 2, representing the degree to which further questioning is warranted for the current topic. If the agent determines that the current information is insufficient for an accurate assessment (i.e., the score exceeds a threshold θ), AG_q further dynamically formulates targeted follow-up questions to encourage the patient to provide additional details. For follow-up questions, we incorporate several additional constraints:

- The in-depth questions require actively guiding the patient to elaborate on the severity, frequency, duration, and impact of their symptoms.
- The agent is encouraged to generate easy-to-answer questions, aiming to reduce the user’s cognitive burden and psychological resistance.

Let j represent the index of the follow-up question within i -th topic. Then in-depth question generation process can be formally defined as:

$$Q_i^{j+1} = AG_q(T_i, H_i^j, M_i^j), \quad (1)$$

H_i^j represents the contextual history of preceding question-answer pairs. The memory state M_i^j maintains a structured record of the accumulated dialogue history up to the $(i-1)$ -th topic, encompassing topic-level summaries and scores from T_1 to T_{i-1} , as well as salient information extracted from all prior conversational turns within i -th topic, spanning from (Q_i^1, A_i^1) to (Q_i^j, A_i^j) .

Once sufficient information is collected or the predefined upper limit of follow-up questions d is reached, the system automatically transitions from the in-depth questioning phase to the next topic, ensuring an efficient and structured evaluation. Therefore, the question generation process can be formulated as:

$$Q_{next} = \begin{cases} Q_i^{j+1}, & AG_{ev}(Q_i^{1 \sim j}, A_i^{1 \sim j}) > \theta \wedge j < d \\ Q_{i+1}^1, & \text{otherwise} \end{cases} \quad (2)$$

Memory-Augmented Scoring

After the complete conversation about a topic, the scoring agent is applied to assess the patient state.

$$[S_i, B_i] = AG_s(Q_i^{1 \sim N_i}, A_i^{1 \sim N_i}, SR) \quad (3)$$

where S_i is the score and B_i denotes the summary and supporting evidence, and N_i is the number of questions for topic T_i . SR refers to the scale rating standards. For instance, regarding the “depressed mood”, the scoring agent assigns a score of 1, with the summary and supporting evidence as: “The user described feeling down in the last two weeks, but only occasionally.” The procedure for scoring and memory updating is presented in Algorithm 1.

After all topics in the scale are evaluated with their respective scores and memory, a comprehensive update is performed to account for inter-topic correlations. The updating process can be formulated as:

$$[S', R] = AG_u(H, M_{topic}), \quad (4)$$

where H denotes the overall dialogue context, including all questions and responses, M_{topic} indicates aggregated information in topic nodes. S' represents the final assessment scores assigned to each topic and R represents the reasoning trace generated during the update process, as well as a comprehensive summary and professional suggestions for the current participant, which contributes to model explainability. To derive the overall outcome, scores from all topics in S' are aggregated into a total score, which is then mapped to a diagnostic category based on the corresponding psychological scale’s criteria.

Experiments

Experimental Setting

Dataset. We conduct experiments on the Distress Analysis Interview Corpus-Wizard of Oz (DAIC-WOZ) (Gratch et al. 2014), a widely adopted English-language dataset for clinical depression detection. This corpus includes interviews with 189 participants, captured through transcripts, audio recordings, and videos. In this work, we focus exclusively on the text modality derived from the interview transcripts. Following each interview, participants complete the PHQ-8, which evaluates eight depression-related symptoms. Participants with a PHQ-8 score of 10 or higher are classified as depressed, while those scoring below 10 are labeled as control. Following prior studies (Chen et al. 2024c), we adhere to the fixed data split of DAIC-WOZ and use the development set for evaluation. See Appendix A for details of the dataset.

Baselines. To explore the effectiveness of our proposed method, we use the following prompt-based studies as our baselines, including Zero-Shot, Few-Shot, Chain-of-Thought (CoT) (Wei et al. 2022), and Chain-of-Logic (Zhao et al. 2024). In addition to prompt-based methods, we extended our comparative analysis to include several basic multi-agent-based approaches, which include MDAgents (Kim et al. 2024), Debate (Tran et al. 2025), and Majority Voting (Wang et al. 2023). See Appendix D for a detailed introduction. For the aforementioned LLM-based methods,

Algorithm 1: Topic Scoring

Input: Scale topics T , Scale rating standards SR , User ID uid

Output: Memory M

```
1:  $M \leftarrow InitializeMemory(uid)$ 
2:  $d \leftarrow 3, i \leftarrow 1$  {Initialize time step}
3: for each theme  $T_i$  in  $T$  do
4:   Initialize  $j \leftarrow 0$ 
5:   while  $j < d$  do
6:      $Q_i^{j+1} \leftarrow AG_q(T_i, H_i^j, M_i^j)$ 
7:      $A_i^{j+1} \leftarrow User(Q_i^{j+1})$ 
8:     {emotion, frequency, etc}  $\leftarrow LLM(A_i^{j+1})$ 
9:      $M.add(vs_i^{j+1})$  {Create the statement node}
10:    necessity  $\leftarrow AG_{ev}(T_i, Q_i^{1 \sim j+1}, A_i^{1 \sim j+1})$ 
11:    if necessity  $< \theta$  then
12:      break
13:    end if
14:     $j \leftarrow j + 1$ 
15:  end while
16:   $(S_i, B_i) \leftarrow AG_s(Q_i^{1 \sim j}, A_i^{1 \sim j}, SR)$ 
17:   $M.add(vt_i)\{T_i, S_i, B_i\}$ 
18:   $M.update(vt_1, vt_2, \dots, vt_{i-1})$ 
19:   $i \leftarrow i + 1$ 
20: end for
21: return  $M$ 
```

we implement two distinct experimental settings for effective comparison. In the first setting, the model performs binary classification using the full interview transcript as input, denoted as bc . In the second setting, the model predicts item-wise scores based on the interview text and relevant psychological scale topics. These item-level predictions are then aggregated into a total score, which is mapped to a categorical label according to the criteria defined by the scale, denoted as sl . There are also some task-specific learning-based approaches, including Multi-View (Agarwal, Dias, and Dollfus 2024), and SEGA (Chen et al. 2024b).

Evaluation. In our experiments, we leverage transcribed text of dialogues from the DAIC-WOZ dataset as prior knowledge and prompt LLMs to simulate individuals with or without mental disorders. Subsequently, our proposed interactive framework engages with the simulated individuals, implementing multi-turn dialogues that reflect realistic mental health scenarios. In validating the efficacy of prediction scores, we use mean absolute error (MAE) to evaluate precision between the predicted scores and the ground truth. Additionally, based on the classification outputs, model performance is evaluated using Cohen’s Kappa coefficient (Kappa), the F1 scores for the Control and Depressed classes (denoted as F1[C] and F1[D]), and the Macro F1 score. See Appendix B.1 for a detailed introduction for these metrics.

Implementation Details. We adopt the AutoGen framework (Wu et al. 2024) to implement our multi-agent system. We employ the Qwen2.5 series (e.g., Qwen2.5-14B-Instruct and Qwen2.5-72B-Instruct) models as the core foundation models for comparative experiments, with the

LLM	Methods	MAE↓	Kappa↑	F1[C]↑	F1[D]↑	Macro F1↑
-	Multi-View	-	-	-	-	77.0
	SEGA	-	-	88.4	81.5	84.9
	SEGA++	-	-	90.9	84.6	87.8
Qwen2.5-14b	Zero-Shot (<i>bc/sl</i>)	- / 3.314	51.2 / 56.8	81.8 / 88.5	69.2 / 66.7	75.5 / 77.6
	5-Shot (<i>bc/sl</i>)	- / 3.086	63.4 / 66.3	86.4 / 89.8	76.9 / 76.2	81.6 / 83.0
	CoT (<i>bc/sl</i>)	- / 3.143	56.5 / 64.8	84.4 / 90.2	72.0 / 73.7	78.2 / 81.9
	Chain-of-Logic (<i>bc/sl</i>)	- / 3.000	62.0 / 66.3	87.0 / 89.8	75.0 / 76.2	81.0 / 83.0
	Majority Voting (<i>bc/sl</i>)	- / 2.857	56.5 / 62.0	84.4 / 87.0	72.0 / 75.0	78.2 / 81.0
	Debate (<i>bc/sl</i>)	- / 2.857	51.2 / 67.7	81.8 / 89.4	69.2 / 78.3	75.5 / 83.8
	MDAgents (<i>bc/sl</i>)	- / 2.629	62.0 / 67.7	87.0 / 89.4	75.0 / 78.3	81.0 / 83.8
	AgentMental	2.543	72.4	92.0	80.0	86.0
Qwen2.5-72b	Zero-Shot (<i>bc/sl</i>)	- / 3.286	53.0 / 56.8	81.0 / 88.5	71.4 / 66.7	76.2 / 77.6
	5-Shot (<i>bc/sl</i>)	- / 3.029	63.4 / 66.3	86.4 / 89.8	76.9 / 76.2	81.6 / 83.0
	CoT (<i>bc/sl</i>)	- / 3.029	58.1 / 64.8	83.7 / 90.2	74.1 / 73.7	78.9 / 81.9
	Chain-of-Logic (<i>bc/sl</i>)	- / 3.114	63.4 / 64.8	86.4 / 90.2	76.9 / 73.7	81.6 / 81.9
	Majority Voting (<i>bc/sl</i>)	- / 2.771	63.4 / 67.7	86.4 / 89.4	76.9 / 78.3	81.6 / 83.8
	Debate (<i>bc/sl</i>)	- / 2.800	54.7 / 68.9	80.0 / 88.9	73.3 / 80.0	76.7 / 84.4
	MDAgents (<i>bc/sl</i>)	- / 2.800	63.4 / 73.6	86.4 / 91.7	76.9 / 81.8	81.6 / 86.7
	AgentMental	2.514	79.8	93.9*	85.7*	89.8*

Table 1: Comparison of the evaluation results of our method with baseline methods on the dataset DAIC-WOZ. **Bold** indicates best performance. *bc* represents binary classification, and *sl* represents item-level predictions. Results with * are better than the baseline ($p < 0.05$) based on a one-tailed unpaired t-test.

Item	CoT	MDAgents	Ours
Loss of Interest	37.8	45.8	38.7
Depressed Mood	55.8	57.1	55.7
Sleep Problems	49.6	60.7	63.6
Fat or LowEn	45.7	56.6	67.0
Appetite or Weight	13.3	19.1	28.1
Low Self-Worth	36.8	35.2	55.2
Concentration	25.6	28.5	38.1
Psychomotor	20.5	27.3	28.9

Table 2: Item-level fine-grained assessment results of F1 score on DAIC-WOZ. ‘Fat or LowEn’, ‘Appetite or Weight’, ‘Concentration’ and ‘Psychomotor’ refer to ‘Fatigue or Low Energy’, ‘Appetite or Weight Changes’, ‘Concentration Difficulties’ and ‘Psychomotor Changes’.

temperature parameter set to 0. Besides, we use Deepseek-R1 (i.e., Deepseek-R1-Distill-Qwen-32B) to simulate users (Guo et al. 2025). The experiments are conducted on an NVIDIA A6000 GPU platform, utilizing the vLLM framework (Kwon et al. 2023) for acceleration. The hyperparameters d and θ are set to 3 and 1, respectively.

Overall Performance

To explore the effectiveness of our proposed method, we compare it with the aforementioned studies, with the results reported in Table 1. There are several observations. First, given the same model architecture, models with more parameters tend to have better performance, as demonstrated by the comparison between Qwen2.5-14B-Instruct and Qwen2.5-72B-Instruct. This observation aligns with

prior findings that larger models tend to demonstrate enhanced capabilities in language understanding and reasoning (Team 2024). Second, methods evaluated under the *sl* configuration consistently demonstrate superior performance compared to those under the *bc* configuration. This performance disparity can be attributed to the structured guidance offered by the psychological scale in the *sl* setting, which enhances the extraction of critical information and reduces semantic ambiguity present in long-form psychological interviews. Third, models augmented with prompt engineering outperform their non-augmented counterparts. This improvement may stem from the model’s ability to better interpret user input, infer underlying mental states, and produce more accurate and reliable assessment results. Fourth, multi-agent-based approaches (i.e., Majority Voting, Debate, MDAgents and AgentMental) further exhibit superior performance than most prompt-based methods, which underscores the benefits of collaborative agent coordination. AgentMental achieves the best performance in most evaluation metrics, confirming the effectiveness of integrating an in-depth questioning mechanism with a tree-based memory structure. The in-depth questioning mechanism is capable of generating easy-to-answer yet informative questions, guiding users to provide more valuable and comprehensive information for better understanding of their mental state. Meanwhile, the memory module helps reduce information redundancy and strengthens the coherence and interdependence of the collected information.

Item-level Fine-Grained Analysis

To further validate the framework’s capability to capture subtle differences in psychological states, we conducted a

Module		MAE↓	Kappa↑	F1-score↑
In-depth	Memory			
✓	✓	2.514	79.8	89.8
✓	✗	3.000	58.7	79.0
✗	✓	3.314	60.4	80.1
✗	✗	3.400	47.2	73.5

Table 3: The results of ablation study.

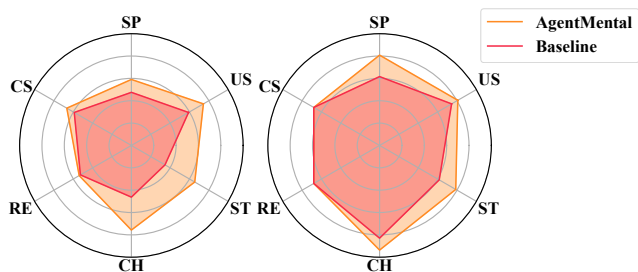
fine-grained analysis of 8 topics within the PHQ-8 scale, with the results shown in Table 2. First, AgentMental demonstrates superior performance across the majority of scale topics compared to the CoT and MDAgents, demonstrating its enhanced effectiveness. Second, substantial performance gains are observed for symptoms such as “Fatigue or Low Energy” and “Sleep Problems”, which often involve observable physiological or behavioral indicators. These improvements suggest that our method benefits from its interactive design, where the ability to pose follow-up questions enables more precise and contextually grounded assessment. Third, items characterized by high subjectivity and contextual dependency, such as “psychomotor changes”, remain challenging to assess reliably within automated systems. This difficulty likely stems from their reliance on subtle linguistic cues, including emotional tendencies throughout the interaction and signs of cognitive dissonance, which place higher demands on the model’s reasoning and understanding capabilities.

Ablation Study

To explore the effectiveness of different modules in our proposed method, we conduct an ablation study comparing the full model to three variants: I without in-depth questioning mechanism, II without memory to store information, III without both of them. The results are reported in Table 3. It is observed that removing any component from our proposed method results in a noticeable decline in performance, highlighting the contribution of each component in enhancing the framework’s ability to assess mental health. Furthermore, the integration of in-depth questioning and memory allows the full model to outperform the variants, validating our design choice of combining these modules into a multi-agent framework. This may be attributed to the complementary benefits of the two components: in-depth questioning facilitates the elicitation of richer user information, while the memory enhances the model’s ability to capture inter-topic relationships and retain relevant contextual details.

GPT/Human Evaluation

To evaluate clinical utility and interaction quality, we conduct a double-blind evaluation experiment involving human and GPT assessments across 5 participants cases with varying symptom severities, comparing its performance against baseline methods (i.e., model III in **Ablation Study**). Each case comprises the full dialogue text, assessment scale scores, a summary of the participant’s symptoms, and corresponding suggestions. We adopt the multi-dimensional em-



(a) Human Evaluation (b) GPT Evaluation

Figure 3: Comparison of our method with the baseline method through automatic and human evaluations, where US, SP, CS, RE, CH, ST refer to Understanding, Sympathy, Consistency, Realism, Coherence, and Satisfaction.

pathy framework and use Understanding and Sympathy to assess perceived empathy during psychological interviews (Xu and Jiang 2024). In addition, we follow Yang et al. (2024c) to measure simulation effectiveness and user experience using four metrics: Consistency, Realism, Coherence, and Satisfaction.

- For human evaluation, we employ five graduate students independently to review each case, and the average value is taken as the final result to reduce individual bias and ensure reliability.
- For GPT-4o automated evaluation, each case is scored independently three times, and the average value is taken as the final result.

The evaluation results are presented in Figure 3. It is observed that our method outperforms the baseline across most aforementioned metrics, further demonstrating the effectiveness of our design for psychological evaluation. Besides, the model demonstrates significantly greater improvements in Coherence, and Satisfaction relative to other measures. This improvement may be attributed to the interactive design of AgentMental, which promotes more engaging, contextually coherent, and user-centered dialogue. Such an interaction paradigm further enhances users’ perceived satisfaction.

Case Study

To further investigate the effectiveness of our method, we conduct a qualitative analysis of two representative cases. As shown in Figure 4, we present a comparison of the interaction processes and rating outcomes between our proposed method and the baseline model (i.e., Qwen2.5-72B-Instruct) in evaluating the topic of “Sleep Problems” and “Loss of Interest” on the PHQ-8. In these cases, the baseline employs a single-question approach that captures only surface-level responses, whereas AgentMental utilizes multi-turn interactions to elicit deeper and more informative user input. Triggered by the evaluation agent, follow-up inquiries aim to explore the underlying causes of symptoms and uncover critical details, thereby refining the assessment and demonstrating the framework’s capacity for dynamic adaptation to user-provided input. Through the collection of the key information, e.g., (“leaves him tired and less focused...work and

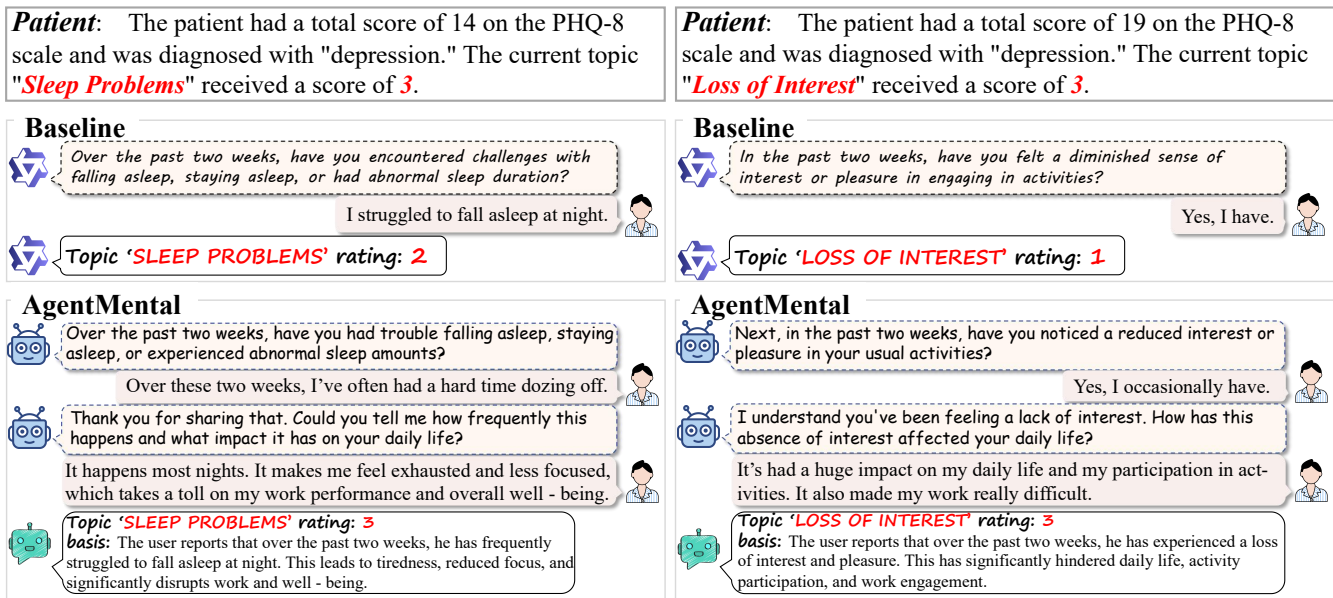


Figure 4: Examples of PHQ-8 topics “SLEEP PROBLEMS” and “LOSS OF INTEREST” from Baseline and AgentMental.

well-being” in the left example and “significantly impacting daily life...activities and work” in the right case), AgentMental can give a more precise score.

Related Work

Multi-Agent Systems

Multi-agent systems (Dorri, Kanhere, and Jurdak 2018) involve multiple autonomous agents that collaborate or compete to solve complex problems, adapting dynamically to changing environments (Li et al. 2023; Chen et al. 2024a; Chan et al. 2024; Guo et al. 2024; Wang et al. 2024; Hu et al. 2025). Recently, multi-agent systems have been increasingly applied in the medical domain to enhance assessment and intervention by simulating interactions between patients and virtual agents (Fan et al. 2025; Almansoori, Kumar, and Cholakkal 2025). For example, the MDAgents (Kim et al. 2024) leverages a dynamic team of collaborative agents to develop highly personalized, comprehensive treatment plans for patients with multiple chronic conditions. Similarly, the PsychoGAT framework (Yang et al. 2024c) employs agents to transform traditional questionnaires into interactive, game-based assessments, thereby boosting engagement and accuracy.

LLM for Mental Health

In recent years, LLMs have shown great potential in natural language processing, leveraging extensive datasets and advanced architectures to comprehend and generate human-like text. Models such as GPT-4 (Achiam et al. 2023), and Qwen (Yang et al. 2024a) exhibit near-human capabilities in diverse tasks, including text generation, text classification, sentiment analysis, and contextual understanding. In the realm of mental health, LLMs also offer significant potential in powering mental healthcare domain, such as dialogue generation, depression and anxiety detection (Yang

et al. 2024b; Xu et al. 2024; Dai et al. 2025). For example, MeChat (Qiu et al. 2024), which simulates psychotherapy dialogues to provide emotional support. Furthermore, the LLAMADRS framework (Kebe et al. 2025) utilizes LLMs to analyze clinical interview transcription data, achieving depression symptom scoring close to human assessment. Chen et al. (2024b) transforms clinical interviews into a directed acyclic graph and improve it with principle-guided data using LLMs. Although these approaches have brought significant improvements, they often lack explicit explainability and fail to guide patients to gradually disclose more valuable information during evaluation. Compared to these methods, our method offers an alternative solution to introduce an in-depth questioning mechanism to actively elicit richer information from patients, mitigating cognitive barriers or resistance, while simultaneously providing a transparent process to enhance the explainability of the evaluation.

Conclusion

In this paper, we propose a multi-agent framework designed to simulate clinical doctor-patient dialogues for psychological evaluation. Specifically, we integrate an adaptive in-depth questioning mechanism designed to elicit richer user information, dynamically adjusting its queries based on the adequacy evaluation of user responses. Besides, we introduce a dynamic tree-based memory to structurally record critical information from previous dialogue turns, which is used to reduce question duplication and further improve performance. Experimental results on the DAIC-WOZ dataset demonstrate that our proposed method outperforms existing baselines, showcasing its significant potential for advancing automated mental health assessment.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China under Grant 62402158, Grant 62502145 and Grant 62272144; by the Anhui Provincial Natural Science Foundation Grant 2408085QF188, and Grant 2408085J040; by the Fundamental Research Funds for the Central Universities Grant JZ2025HGTA0162, and Grant JZ2025HGQA0134; and by the Major Project of Anhui Provincial Science and Technology Breakthrough Program (202423k09020001).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Agarwal, N.; Dias, G.; and Dollfus, S. 2024. Analysing relevance of discourse structure for improved mental health estimation. In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, 127–132.
- Almansoori, M.; Kumar, K.; and Cholakkal, H. 2025. Self-Evolving Multi-Agent Simulations for Realistic Clinical Interactions. *arXiv preprint arXiv:2503.22678*.
- Chan, C.-M.; Chen, W.; Su, Y.; Yu, J.; Xue, W.; Zhang, S.; Fu, J.; and Liu, Z. 2024. ChatEval: Towards Better LLM-based Evaluators through Multi-Agent Debate. In *The Twelfth International Conference on Learning Representations*.
- Chen, W.; Su, Y.; Zuo, J.; Yang, C.; Yuan, C.; Chan, C.-M.; Yu, H.; Lu, Y.; Hung, Y.-H.; Qian, C.; Qin, Y.; Cong, X.; Xie, R.; Liu, Z.; Sun, M.; and Zhou, J. 2024a. AgentVerse: Facilitating Multi-Agent Collaboration and Exploring Emergent Behaviors. In *The Twelfth International Conference on Learning Representations*.
- Chen, Z.; Deng, J.; Zhou, J.; Wu, J.; Qian, T.; and Huang, M. 2024b. Depression detection in clinical interviews with LLM-empowered structural element graph. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8181–8194.
- Chen, Z.; Deng, J.; Zhou, J.; Wu, J.; Qian, T.; and Huang, M. 2024c. Depression Detection in Clinical Interviews with LLM-Empowered Structural Element Graph. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 8181–8194. Mexico City, Mexico: Association for Computational Linguistics.
- Dai, C.; Hu, J.; Shi, H.; Li, Z.; Yang, X.; and Wang, M. 2025. Psyche-r1: Towards reliable psychological llms through unified empathy, expertise, and reasoning. *arXiv preprint arXiv:2508.10848*.
- Dorri, A.; Kanhere, S. S.; and Jurdak, R. 2018. Multi-agent systems: A survey. *Ieee Access*, 6: 28573–28593.
- Fan, Z.; Wei, L.; Tang, J.; Chen, W.; Siyuan, W.; Wei, Z.; and Huang, F. 2025. AI Hospital: Benchmarking Large Language Models in a Multi-agent Medical Interaction Simulator. In Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B. D.; and Schockaert, S., eds., *Proceedings of the 31st International Conference on Computational Linguistics*, 10183–10213. Abu Dhabi, UAE: Association for Computational Linguistics.
- Gratch, J.; Artstein, R.; Lucas, G.; Stratou, G.; Scherer, S.; Nazarian, A.; Wood, R.; Boberg, J.; DeVault, D.; Marsella, S.; Traum, D.; Rizzo, S.; and Morency, L.-P. 2014. The Distress Analysis Interview Corpus of human and computer interviews. In Calzolari, N.; Choukri, K.; Declerck, T.; Loftsson, H.; Maegaard, B.; Mariani, J.; Moreno, A.; Odijk, J.; and Piperidis, S., eds., *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, 3123–3128. Reykjavik, Iceland: European Language Resources Association (ELRA).
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Guo, T.; Chen, X.; Wang, Y.; Chang, R.; Pei, S.; Chawla, N. V.; Wiest, O.; and Zhang, X. 2024. Large Language Model Based Multi-agents: A Survey of Progress and Challenges. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 8048–8057. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Hu, J.; Dong, T.; Gang, L.; Ma, H.; Zou, P.; Sun, X.; Guo, D.; Yang, X.; and Wang, M. 2024. Psycollm: Enhancing llm for psychological understanding and evaluation. *IEEE Transactions on Computational Social Systems*.
- Hu, J.; Shi, H.; Dai, C.; Li, Z.; Song, P.; and Wang, M. 2025. Beyond emotion recognition: A multi-turn multimodal emotion understanding and reasoning benchmark. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 5814–5823.
- Ji, S.; Zhang, T.; Ansari, L.; Fu, J.; Tiwari, P.; and Cambria, E. 2022. MentalBERT: Publicly Available Pretrained Language Models for Mental Healthcare. In Calzolari, N.; Béchet, F.; Blache, P.; Choukri, K.; Cieri, C.; Declerck, T.; Goggi, S.; Isahara, H.; Maegaard, B.; Mariani, J.; Mazo, H.; Odijk, J.; and Piperidis, S., eds., *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, 7184–7190. Marseille, France: European Language Resources Association.
- Kebe, G. Y.; Girard, J. M.; Liebenthal, E.; Baker, J.; De la Torre, F.; and Morency, L.-P. 2025. LlaMADRS: Prompting Large Language Models for Interview-Based Depression Assessment. *arXiv preprint arXiv:2501.03624*.
- Khowaja, S. A.; Nkenyereye, L.; Khuwaja, P.; Al Hamadi, H.; and Dev, K. 2025. Depression Detection From Social Media Posts Using Emotion Aware Encoders and Fuzzy Based Contrastive Networks. *IEEE Transactions on Fuzzy Systems*, 33(1): 43–53.

- Kim, Y. H.; Park, C.; Jeong, H.; Chan, Y. S.; Xu, X.; McDuff, D.; Breazeal, C.; and Park, H. W. 2024. MDAgents: An Adaptive Collaboration of LLMs for Medical Decision-Making. In *Neural Information Processing Systems*.
- Kroenke, K.; Strine, T. W.; Spitzer, R. L.; Williams, J. B.; Berry, J. T.; and Mokdad, A. H. 2009. The PHQ-8 as a measure of current depression in the general population. *Journal of affective disorders*, 114(1-3): 163–173.
- Kwon, W.; Li, Z.; Zhuang, S.; Sheng, Y.; Zheng, L.; Yu, C. H.; Gonzalez, J.; Zhang, H.; and Stoica, I. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, 611–626.
- Li, G.; Hammoud, H.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023. Camel: Communicative agents for” mind” exploration of large language model society. *Advances in Neural Information Processing Systems*, 36: 51991–52008.
- Ohse, J.; Hadžić, B.; Mohammed, P.; Peperkorn, N.; Fox, J.; Krutzki, J.; Lyko, A.; Mingyu, F.; Zheng, X.; Rättsch, M.; et al. 2024. GPT-4 shows potential for identifying social anxiety from clinical interview data. *Scientific Reports*, 14(1): 1–12.
- Qiu, H.; He, H.; Zhang, S.; Li, A.; and Lan, Z. 2024. SMILE: Single-turn to Multi-turn Inclusive Language Expansion via ChatGPT for Mental Health Support. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 615–636. Miami, Florida, USA: Association for Computational Linguistics.
- Sommers-Flanagan, J.; Zeleke, W. A.; and Hood, M. 2015. The clinical interview. *The encyclopedia of clinical psychology*. John Wiley & Sons, Inc, Hoboken, 1–9.
- Team, Q. 2024. Qwen2. 5: A party of foundation models. *Qwen (Sept. 2024)*. url: <https://qwenlm.github.io/blog/qwen2,5>.
- Tran, K.-T.; Dao, D.; Nguyen, M.-D.; Pham, Q.-V.; O’Sullivan, B.; and Nguyen, H. D. 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. arXiv:2501.06322.
- Wang, Q.; Wang, Z.; Su, Y.; Tong, H.; and Song, Y. 2024. Rethinking the Bounds of LLM Reasoning: Are Multi-Agent Discussions the Key? In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6106–6131. Bangkok, Thailand: Association for Computational Linguistics.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q. V.; Chi, E. H.; Narang, S.; Chowdhery, A.; and Zhou, D. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; et al. 2024. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Xiao, M.; Xie, Q.; Kuang, Z.; Liu, Z.; Yang, K.; Peng, M.; Han, W.; and Huang, J. 2024. HealMe: Harnessing Cognitive Reframing in Large Language Models for Psychotherapy. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1707–1725.
- Xu, X.; Yao, B.; Dong, Y.; Gabriel, S.; Yu, H.; Hendler, J.; Ghassemi, M.; Dey, A. K.; and Wang, D. 2024. Mental-llm: Leveraging large language models for mental health prediction via online text data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1): 1–32.
- Xu, Z.; and Jiang, J. 2024. Multi-dimensional Evaluation of Empathetic Dialogue Responses. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 2066–2087. Miami, Florida, USA: Association for Computational Linguistics.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; et al. 2024a. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Yang, K.; Zhang, T.; Kuang, Z.; Xie, Q.; Huang, J.; and Ananiadou, S. 2024b. MentaLLaMA: interpretable mental health analysis on social media with large language models. In *Proceedings of the ACM Web Conference 2024*, 4489–4500.
- Yang, Q.; Wang, Z.; Chen, H.; Wang, S.; Pu, Y.; Gao, X.; Huang, W.; Song, S.; and Huang, G. 2024c. PsychoGAT: A Novel Psychological Measurement Paradigm through Interactive Fiction Games with LLM Agents. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 14470–14505. Bangkok, Thailand: Association for Computational Linguistics.
- Yu, Y.; Yu, R.; Wei, H.; Zhang, Z.; and Qian, Q. 2025. Beyond Dialogue: A Profile-Dialogue Alignment Framework Towards General Role-Playing Language Model. In Che, W.; Nabende, J.; Shutova, E.; and Pilehvar, M. T., eds., *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 11992–12022. Vienna, Austria: Association for Computational Linguistics. ISBN 979-8-89176-251-0.
- Zhao, X.; Li, M.; Lu, W.; Weber, C.; Lee, J. H.; Chu, K.; and Wermter, S. 2024. Enhancing Zero-Shot Chain-of-Thought Reasoning in Large Language Models through Logic. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 6144–6166. Torino, Italia: ELRA and ICCL.