

A Layered Multi-Expert Framework for Long-Context Mental Health Assessments

Authors

Tang, Jinwen; Guo, Qiming; Sun, Wenbo; Shang, Yi

DOI

[10.1109/CAI64502.2025.00080](https://doi.org/10.1109/CAI64502.2025.00080)

Licence

Dutch Copyright Act (Article 25fa)

Publication date

2025

Document Version

Final published version

Published in

Proceedings - 2025 IEEE Conference on Artificial Intelligence, CAI 2025

Citation (APA)

Tang, J., Guo, Q., Sun, W., & Shang, Y. (2025). A Layered Multi-Expert Framework for Long-Context Mental Health Assessments. In *Proceedings - 2025 IEEE Conference on Artificial Intelligence, CAI 2025* (pp. 435-440). (Proceedings - 2025 IEEE Conference on Artificial Intelligence, CAI 2025). IEEE. <https://doi.org/10.1109/CAI64502.2025.00080>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A Layered Multi-Expert Framework for Long-Context Mental Health Assessments

Jinwen Tang*, Qiming Guo[†], Wenbo Sun[‡], and Yi Shang*

*Electrical Engineering and Computer Science Department

University of Missouri, Columbia, Missouri, USA

{jt4cc, shangy}@umsystem.edu

[†]Department of Computer Science

Texas A&M University-Corpus Christi, Corpus Christi, Texas, USA

qguo2@islander.tamucc.edu

[‡]Department of Computer Science

Delft University of Technology, Delft, Netherlands

w.sun-2@tudelft.nl

Abstract—Long-form mental health assessments pose unique challenges for large language models (LLMs), which often exhibit hallucinations or inconsistent reasoning when handling extended, domain-specific contexts. We introduce *Stacked Multi-Model Reasoning (SMMR)*, a layered framework that leverages multiple LLMs and specialized smaller models as coequal "experts". Early layers isolate short, discrete subtasks, while later layers integrate and refine these partial outputs through more advanced long-context models. We evaluate SMMR on the DAIC-WOZ depression-screening dataset and 48 curated case studies with psychiatric diagnoses, demonstrating consistent improvements over single-model baselines in terms of accuracy, F1-score, and PHQ-8 error reduction. By harnessing diverse "second opinions", SMMR mitigates hallucinations, captures subtle clinical nuances, and enhances reliability in high-stakes mental health assessments. Our findings underscore the value of multi-expert frameworks for more trustworthy AI-driven screening.

Index Terms—NLP, LLM, Artificial Intelligence, Mental Health, Psychiatry, Explainable AI, Multi-Model Reasoning

I. INTRODUCTION

Mental health remains a pressing public health concern, complicated by stigma, limited clinical services, and socioeconomic barriers [1], [2]. Advances in large language models (LLMs) such as GPT and Mistral offer potential for discreet, scalable mental health screenings [3], helping individuals who might otherwise be reluctant to seek professional help due to judgment or resource constraints.

However, using LLMs in high-stakes mental health assessments poses distinct challenges. First, **mental health data are inherently subjective**: while instruments like Patient Health Questionnaire-8 (PHQ-8) provide structured guidelines, the interpretation of symptoms can vary significantly from one conversation or individual to another [4], [5]. Subtle linguistic nuances and personal expressions often shape how distress is perceived. Second, **long and complex transcripts**—such as multi-turn interviews or narrative case studies—can overwhelm advanced models, leading to hallucinations or inconsistencies

[6], [7]. Third, **specialized tools** like PHQ-8 offer only partial mitigation; a single LLM "expert" might still fail to capture nuanced details without additional context or repeated checks [8]. Although iterative reassessments and third-party evaluations can refine accuracy, they often rely on structured prompts or assume a known "best" configuration.

In response, we propose *Stacked Multi-Model Reasoning (SMMR)*, a framework in which diverse "expert" models collectively assess the same input without any being deemed categorically superior. SMMR leverages multiple reasoning steps, consolidating "second opinions" in a layered manner:

- 1) **Multiple Experts, One Pipeline**: Early layers employ smaller or specialized LLMs to generate preliminary assessments, treating each as an independent perspective.
- 2) **Iterative Refinement**: Subsequent layers integrate and reconcile these varying outputs, using long-context models to produce a cohesive final judgment.
- 3) **Robust Mental Health Screening**: By adopting a multi-expert approach, SMMR reduces hallucinations and inconsistencies common in extended conversations, improving diagnostic reliability for clinical interviews and narrative case studies.

We evaluate SMMR on two complex mental health datasets: the DAIC-WOZ dataset [9], designed for studying psychological distress (anxiety, depression, PTSD), and a curated set of narrative case studies drawn from real-world clinical scenarios. Our experiments demonstrate that SMMR offers a practical way to provide "second opinions" within a unified workflow—*without* requiring a single preselected model or intricate prompt engineering [10], [11]. By treating all models initially as independent experts, then effectively aggregating imperfect prompts and diverse outputs, SMMR navigates the subjective and nuanced landscape of mental health assessments more reliably, pointing toward safer AI-assisted screenings.

Note: The first two authors contributed equally to this work.

II. BACKGROUND

Artificial intelligence (AI) is reshaping mental health care by providing increasingly sophisticated analytical tools and treatment supports. Early systems relied on simple word-counting and text analysis [12], but the advent of large language models (LLMs) has radically expanded these capabilities. Modern AI-driven applications—such as semi-automated screening platforms, interactive self-reporting interfaces, and responsive mental health chatbots—now enable more nuanced interpretation of human language and broader access to mental health services [13].

Despite these advances, many solutions remain constrained by short-context analyses or single-model pipelines. For instance, Ohse et al. [14] investigate NLP models like BERT and GPT-4 for depression detection, yet their approach primarily relies on individual model outputs. Agrawal [15] highlights the potential of enhanced prompt engineering to improve explainability and intervention planning in mental health AI, though it too centers on a single-model framework. Similarly, Tang and Shang [8] propose a GPT-based system for pre-screening mental health disorders, demonstrating that domain-specific fine-tuning can improve early detection. However, these methods do not fully address the extended, subjective nature of clinical transcripts or case studies where long-format reasoning and multiple perspectives may be crucial.

In the following section, we introduce *Stacked Multi-Model Reasoning (SMMR)*, a layered framework designed to overcome the limitations of single-model or short-context approaches. By integrating multiple “expert” LLMs and specialized models within a unified pipeline, SMMR offers more robust error checks, interpretive consistency, and the capacity to reconcile diverse viewpoints in subjective mental health data.

III. METHOD

A. Stacked Multi-Model Reasoning (SMMR)

Motivation and Overview: Long-form mental health assessments often require careful synthesis of complex, multi-turn data. Even advanced Large Language Models (LLMs) can become prone to hallucinations or inconsistent reasoning in these extended contexts. To address this problem, we propose *Stacked Multi-Model Reasoning (SMMR)*, a framework that treats each LLM (or smaller specialized model) as an independent “expert,” without *a priori* knowledge of which model might be “best.” By layering multiple models—each offering a second opinion—SMMR leverages collective insights to mitigate the weaknesses of any single model. Figure 1 illustrates the conceptual architecture of SMMR, while Algorithm 1 provides pseudocode for the overall process.

a) Layer 1: Multiple Independent Experts.: In this initial step, SMMR applies multiple *single-step* models, each independently processing the input \mathcal{X} . These models could be smaller LLMs or specialized classifiers that excel at short-context tasks. Since we do not assume any model to be inherently superior, all outputs are aggregated on equal footing, forming R_1 for further refinement.

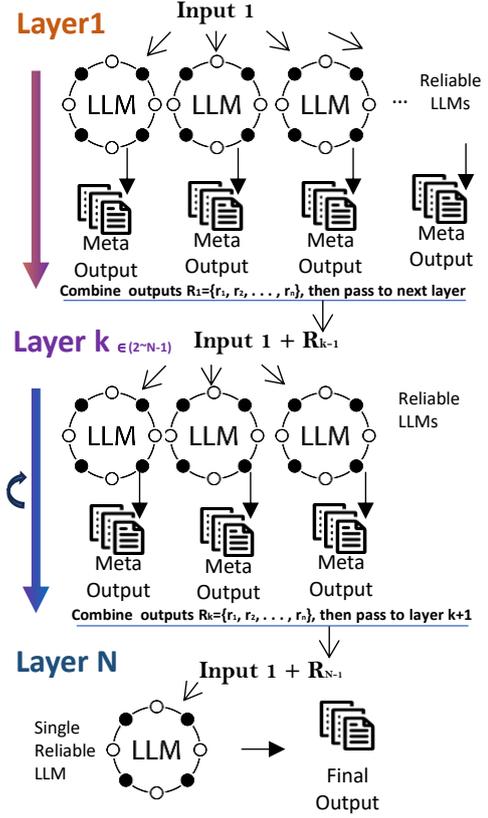


Fig. 1. Overview of the Stacked Multi-Model Reasoning (SMMR) framework. Each layer refines or integrates outputs from the previous layer to form the next layer’s input. No prior ranking of model quality is assumed.

b) Layers 2 to $N - 1$: Iterative Refinement with Long-Context Models.: Subsequent layers introduce *long-context* LLMs capable of handling extended or detailed inputs. Each model in layer k takes the aggregated output R_{k-1} from the previous layer and refines it, generating new outputs r_k^j . These outputs are then aggregated into R_k . The iterative process allows each layer to provide a “second opinion,” reconciling discrepancies and minimizing hallucinations or inconsistencies that could arise from any single model.

c) Layer N : Final Consolidation by a Reliable Model.: In the last layer, a single, long-context *Reliable Model* M_N^* synthesizes the refined outputs from R_{N-1} into the ultimate result R_{final} . This top-tier model is selected for its stable performance on nuanced, long-context data—making it especially well-suited for high-stakes mental health evaluations.

d) Dynamic Stopping Based on Performance Optimization.: To ensure optimal performance and computational efficiency, the SMMR framework incorporates a dynamic stopping mechanism. After each layer processes and refines the aggregated outputs from the preceding layer, the framework evaluates the current performance using predefined metrics (e.g., Accuracy, F1-score, MAE). If the performance metrics improve

Algorithm 1: Stacked Multi-Model Reasoning (SMMR)

Input: Long-context data \mathcal{X} ; set of L_1 -type Single-Step Models $\{M_1^1, \dots, M_1^n\}$; set of L_k -type Long-Context Models $\{M_k^1, \dots, M_k^m\}$ for each layer k ; number of layers N .

Output: Final aggregated output R_{final} .

```
1 Layer 1: Initial Opinions
2 foreach model  $M_1^i \in \{\text{Single-Step Models}\}$  do
3   |  $r_1^i \leftarrow M_1^i(\mathcal{X})$ ;
4 end
5 Aggregate outputs:  $R_1 \leftarrow \{r_1^1, r_1^2, \dots, r_1^m\}$ ;
6 Layers 2 to  $N - 1$ : Iterative Refinement
7 for  $k = 2$  to  $N - 1$  do
8   | foreach model
9     |  $M_k^j \in \{\text{Long-Context Models at layer } k\}$  do
10    | |  $r_k^j \leftarrow M_k^j(R_{k-1})$ ;
11    | end
12   | Aggregate outputs:  $R_k \leftarrow \{r_k^1, r_k^2, \dots, r_k^m\}$ ;
13 end
14 Layer  $N$ : Final Consolidation
15 Use a single, reliable Long-Context Model  $M_N^*$ :
16  $R_{\text{final}} \leftarrow M_N^*(R_{N-1})$ ;
17 return  $R_{\text{final}}$ 
```

compared to the previous layer by a significant threshold δ , SMMR proceeds to the next layer for further refinement. This iterative process continues until adding another layer does not result in performance gains beyond the threshold. The final output used for evaluation is thus the result from the layer that achieved the highest performance metrics, optimizing both accuracy and resource utilization.

By dividing the reasoning process into distinct layers and integrating diverse model outputs at each stage, SMMR effectively mitigates the risk of hallucinations and maintains stronger consistency for complex mental health tasks.

B. Datasets and Task Setup

In order to demonstrate the effectiveness of our Stacked Multi-Model Reasoning (SMMR) framework, we evaluate on two complementary mental health datasets: (1) an externally sourced dataset (DAIC-WOZ) [9], and (2) a curated collection of narrative case studies. Our primary task involves predicting mental health risk and severity, operationalized through PHQ-8 scores or binary labels, based on extended transcripts or descriptive case data. By adopting a layered approach that draws on multiple “expert” models, SMMR aims to reduce hallucinations, enhance diagnostic fidelity, and improve the reliability of long-context LLM-based assessments.

The DAIC-WOZ Database: A private dataset consists of 187 labeled interviews designed to assess psychological distress such as anxiety, depression, and PTSD. Each interview includes a PHQ-8 score and its corresponding binary label. Following the standard protocol in [9], we split the dataset into training and testing subsets. To create a more realistic, long-context

input for our models, each conversation was concatenated into a single data stream by aligning segments according to the speaker’s starting time. The consolidated dataset includes references to the speaker, the content of each segment, and punctuation marks consisting of a period followed by a slash (./) to denote the end of each speaking turn. We evaluate model outputs using:

- **PHQ-8 Score Estimation (0–24):** We compare the predicted PHQ-8 score to the ground truth, measuring accuracy with metrics such as Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE).
- **Binary Classification (PHQ-8 ≥ 10):** The model’s output is thresholded at 10 to indicate the presence or absence of clinically significant depressive symptoms.

Case Study Dataset: To complement the interview-style data from the DAIC-WOZ database, we collected 48 narrative case studies featuring professionals’ concluded psychiatric diagnoses. These cases, sourced from academic texts and clinical literature, include detailed demographic, behavioral, and contextual information. We manually extracted binary conclusions and types of disorders from the diagnoses provided by clinical professionals to verify the presence or absence of mental health concerns. Of these case studies, six are formatted as conversational transcripts, while the remaining 42 are presented in a descriptive format. An example is shown in Table I.

Because these cases do not provide PHQ-8 labels, we adopt an alternative evaluation scheme:

- **Mental Concern (0, 1, 2):** Determines whether the case indicates no mental health issue (0), presence of a mental health issue (1), or if the conclusion is indeterminate from the data (2).
- **Disorder-Type Identification:** Extracts the specific mental health disorder(s) (if any) from the text. We measure accuracy by comparing the model-identified disorders to a reference list of ground-truth labels, considering minor variations in naming as valid matches.

TABLE I
CASE STUDY EXAMPLE

Case: A 37-year-old white male infantryman stationed in Iraq arrived at a field hospital complaining that his superior officer placed poisonous ants in his helmet. His face is covered with excoriations from persistent scratching. On further examination, he is stuporous and has mildly slurred speech, tremor, and mint odor to his breath. Later his troop leader mentioned that his Humvee was littered with empty bottles of mouthwash and that the man has been reprimanded for falling asleep at his post. After a night of rest, he discussed his excessive use of mouthwash in place of alcohol, which is the only available form of alcohol in Iraq.

Conclusion: The individual in the provided case study exhibits symptoms consistent with delusions of persecution (believing his superior officer placed poisonous ants in his helmet), excoriations from persistent scratching indicative of possible hallucinations or delusions, stupor, slurred speech, tremors, and the smell of mint on his breath. Additionally, his excessive use of mouthwash as a substitute for alcohol suggests a coping mechanism or self-medication behavior. While the text does not explicitly provide a diagnosis, the symptoms described align with features of psychosis and substance abuse, possibly indicating comorbid conditions such as schizophrenia spectrum disorder and alcohol use disorder.

```

TEMPLATE = """
    You are an AI assistant tasked with simulating a mental health
    assessment based on the provided text. This is not a real clinical diagnosis,
    but a demonstration of how such an assessment might be structured. Using the
    information given, create a hypothetical assessment of the individual's mental
    health over the last two weeks using an 8-question survey. The survey evaluates:
        1. Lack of interest in activities
        2. Feelings of depression or hopelessness
        3. Sleep issues
        4. Low energy
        5. Changes in appetite
        6. Negative self-perception
        7. Concentration difficulties
        8. Unusual movement or speech patterns
    For each question, assign a hypothetical score from 0 (Not at all)
    to 3 (Nearly every day) based on the information provided in the text. Sum up
    these scores for a total ranging from 0 to 24.
    Here is the text to analyze:
    {input}
    Remember, this is a simulated assessment for demonstration purposes
    only. It's crucial to provide the total score as specified and keep your
    explanation concise.
    Below are the assessments from three other AI models. First model's
    result: {model1}; second model's result: {model2}; third model's result:
    {model3}. Please analyze these outcomes, consider the reasons for any
    variations, and provide your own assessment in response.
    Your response MUST follow this exact format, with each section
    separated by a newline:
        1. Total score: X (where X is the calculated score between 0 and 24)
        2. Explanation: A brief explanation of your reasoning for these
        scores in under 600 tokens.
    Your entire response, including the total score, and explanation,
    must not exceed 100 tokens.
    """

```

Fig. 2. SMMR Prompt for DAIC-WOZ Dataset

Model Selection: We initially conducted a pilot test with the smaller, local Mistral model to explore on-premise feasibility. However, the majority of outputs were invalid or incomplete, likely due to context window constraints and the domain-specific nature of the data. Consequently, we omitted local models and go with GPT models but note that future work could revisit them if they become more robust.

To Evaluate: For each dataset instance—whether an interview transcript from DAIC-WOZ or a narrative case study—SMMR processes the *entire* text in a layered manner (the prompts shown in Figure 2 and Figure 3, respectively):

- 1) *Layer 1* applies multiple smaller or specialized models in parallel to gather initial assessments.
- 2) *Layers 2 through N – 1* employ advanced long-context LLMs to refine, reconcile, and aggregate these preliminary outputs.
- 3) *Layer N* employs a single, reliable long-context model to finalize the assessment and generate the ultimate decision, whether it be a binary classification or a PHQ-8 score.
- 4) *Dynamic Stopping:* The SMMR framework automatically

terminates additional layering when further layers do not yield performance improvements beyond a predefined threshold.

This pipeline effectively creates multiple “checkpoints” for error correction and multi-expert verification, ultimately aiming to improve the reliability of mental health evaluations in lengthy and complex conversations.

IV. RESULTS

We evaluated the SMMR framework on both DAIC-WOZ and our curated case studies, comparing its performance against single-model baselines. Tables II and III present the key metrics, reflecting a consistent advantage for SMMR-enhanced approaches.

A. DAIC-WOZ Performance

Table II highlights the performance of GPT-3.5-turbo and GPT-4-turbo across training, testing, and validation splits. Alongside classification measures such as Accuracy (Acc.), F1, Macro F1 (MF1.), Macro Precision (MPrec.), Macro Recall

```

TEMPLATES = {
  # Determine mental health concern from a case
  'mental_concern': """
    Read the given case study or conversation and determine if the person has a
    mental health issue. Choose the most appropriate number (0, 1, or 2) from the
    following options:
    - 0: No, the conclusion does not indicate a mental health issue.
    - 1: Yes, the conclusion indicates a mental health issue.
    - 2: It is indeterminate from the given conclusion.
    Content: {input}
    Below are the assessments from three other AI models. First model's result:
    {model1}; second model's result: {model2}; third model's result: {model3}.
    Please analyze these outcomes, consider the reasons for any variations, and
    provide your own assessment in response.
    Please respond with a SINGLE number (0, 1, or 2) only."""
  # Identify types of mental disorders from a case
  'mental_disorder_type': """
    Read the given case study or conversation and determine if the person has a
    mental health issue, and if so, what type or types of disorders are mentioned?
    If no specific disorders are identified in the text, please type "NA".
    Content: {input}
    Below are the assessments from three other AI models. First model's result:
    {model1}; second model's result: {model2}; third model's result: {model3}.
    Please analyze these outcomes, consider the reasons for any variations, and
    provide your own assessment in response.
    The response should be a list of words or phrases separated by commas."""
}

```

Fig. 3. SMMR Prompt for Case Study Dataset

TABLE II
PERFORMANCE ENHANCEMENT FOR SMMR ACROSS DAIC-WOZ SUBSETS

Dataset	Model	Method	Acc.	F1	MF1	MPrec.	MRec.	ROC AUC	MAE	RMSE
Training	GPT-3.5-turbo	Baseline	0.51	0.49	0.51	0.60	0.61	0.61	5.23	6.17
		SMMR Enhanced	0.69	0.59	0.67	0.68	0.72	0.72	4.19	5.18
	GPT-4-turbo	Baseline	0.76	0.64	0.73	0.72	0.75	0.75	3.32	4.07
		SMMR Enhanced	0.79	0.69	0.76	0.75	0.79	0.79	3.33	4.24
Testing	GPT-3.5-turbo	Baseline	0.55	0.57	0.55	0.70	0.68	0.68	6.04	6.81
		SMMR Enhanced	0.76	0.70	0.75	0.76	0.81	0.81	4.22	5.54
	GPT-4-turbo	Baseline	0.77	0.67	0.74	0.74	0.77	0.77	3.21	4.16
		SMMR Enhanced	0.80	0.74	0.79	0.79	0.84	0.84	3.76	4.78
Validation	GPT-3.5-turbo	Baseline	0.60	0.63	0.60	0.73	0.70	0.70	4.86	5.54
		SMMR Enhanced	0.80	0.76	0.79	0.80	0.83	0.83	3.29	3.96
	GPT-4-turbo	Baseline	0.83	0.75	0.81	0.81	0.81	0.81	2.54	3.08
		SMMR Enhanced	0.83	0.79	0.82	0.82	0.85	0.85	2.43	3.37

Note: Acc, MF1, MPrec, and MRec stand for Accuracy, Macro F1, Macro Precision, and Macro Recall, respectively.

(MRec.), and ROC AUC (for PHQ-8 ≥ 10 classification), we also track Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) for PHQ-8 scoring. Overall, the SMMR-enhanced models outperform baselines on every split, with particularly notable improvements in binary classification accuracy and F1. For example, GPT-3.5-turbo sees its testing accuracy jump from 0.55 to 0.76, while MAE decreases significantly (6.04 to 4.22), suggesting better handling of nuanced long-context data.

Importantly, SMMR also boosts consistency across different subsets. The validation accuracy for GPT-3.5-turbo, for instance, jumps from 0.60 to 0.80, showing that leveraging multiple “experts” and iterative refinement reduces hallucinations and more accurately captures PHQ-8 severity signals.

B. Case Study Dataset Performance

Table III summarizes results on our 48 narrative case studies, where we evaluate binary detection of mental health

TABLE III
MENTAL HEALTH DETECTION ON CASE STUDY DATASET

Method	Valid (%)	Acc.	F1	Ave.	SD
GPT-4o Baseline	92	0.95	0.98	6.97	3.20
GPT-4o+SMMR	100	0.93	0.97	6.85	2.90
GPT-3.5 Baseline	98	0.91	0.95	6.66	3.17
GPT-3.5+SMMR	100	0.93	0.97	7.03	2.57
GPT-4 Baseline	100	0.92	0.96	7.02	2.95
GPT-4+SMMR	100	0.91	0.95	7.40	2.67

Note: This table shows the results for both binary mental health detection (accuracy and F1 scores) and the correctness of disorder type identification (Ave, SD).

concerns and the correctness of disorder-type identification. Although the baseline metrics here are relatively high for GPT-4 variants (exceeding 0.90 accuracy), SMMR demonstrates small but meaningful gains and, in some configurations, yields a 100% valid output rate. Notably, GPT-3.5 sees its accuracy rise from 0.91 to 0.93 under SMMR, while its F1 measure improves from 0.95 to 0.97. Beyond classification, SMMR also stabilizes disorder-type identification; for instance, the average correctness (Ave.) for GPT-4 climbs from 7.02 to 7.40, indicating richer extraction of relevant diagnostic details.

In summary, these findings confirm that layering multiple models curbs inconsistencies and incomplete responses, leading to more robust and thorough assessments in both interview transcripts and case narratives.

V. CONCLUSION

In this study, we introduced Stacked Multi-Model Reasoning (SMMR) as a layered solution for long-context mental health assessments. By progressively refining the outputs of multiple “expert” models, SMMR demonstrated notable improvements in diagnostic accuracy, F1-scores, and PHQ-8 estimation on both the DAIC-WOZ dataset and a curated set of case studies. These gains highlight how a structured, multi-expert approach can mitigate hallucinations, capture subtle clinical cues, and enhance consistency in sensitive, high-stakes domains.

Despite these encouraging results, several limitations warrant attention. First, our model diversity was limited to commercial GPT variants; although these models are powerful, our findings may not generalize to local or open-source LLMs without further experimentation. Second, mental health assessments impose strict confidentiality and ethical constraints. While SMMR can reduce errors and inconsistencies, thorough clinical validation and prospective studies remain necessary to confirm its real-world safety and efficacy. Additionally, the data scarcity inherent in mental health research restricts the breadth of our evaluation, and the computational overhead of running sequential layers can pose scalability challenges. Finally, conflict resolution among layered “experts” raises questions about decision transparency, bias, and the need for IRB-reviewed protocols when dealing with vulnerable populations.

Looking ahead, we plan to expand SMMR by refining conflict-resolution strategies, exploring multi-modal data (e.g.,

audio and video cues), and investigating cost-effective deployments that balance performance with resource usage. We also intend to collaborate with clinical professionals to ensure that SMMR’s layered outputs align with human judgment in ethically rigorous settings. Taken together, these developments could make SMMR a robust foundation for safer and more reliable AI-driven mental health screening, addressing the complexity and subjectivity that define this crucial application area.

ACKNOWLEDGMENT

ChatGPT was used to revise the writing to improve the spelling, grammar, and overall readability.

REFERENCES

- [1] A. Gulliver, K. M. Griffiths, and H. Christensen, “Perceived barriers and facilitators to mental health help-seeking in young people: a systematic review,” *BMC psychiatry*, vol. 10, no. 1, pp. 1–9, 2010.
- [2] B. A. Pescosolido and C. A. Boyer, “How do people come to use mental health services? current knowledge and changing perspectives,” 1999.
- [3] R. OpenAI, “Gpt-4 technical report. arxiv 2303.08774,” *View in Article*, vol. 2, p. 13, 2023.
- [4] A. Ng, R. Kornfield, S. M. Schueller, A. K. Zalta, M. Brennan, and M. Reddy, “Provider perspectives on integrating sensor-captured patient-generated data in mental health care,” *Proceedings of the ACM on human-computer interaction*, vol. 3, no. CSCW, pp. 1–25, 2019.
- [5] M. Marsolek, S. Emert, J. R. Dietch, E. Tucker, and D. Taylor, “0969 predicting differences between objective and subjective sleep parameters with mental health questionnaires,” *SLEEP*, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:269288613>
- [6] L. Huang, W. Yu, W. Ma, W. Zhong, Z. Feng, H. Wang, Q. Chen, W. Peng, X. Feng, B. Qin *et al.*, “A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions,” *ACM Transactions on Information Systems*, 2023.
- [7] M. Dahl, V. Magesh, M. Suzgun, and D. E. Ho, “Large legal fictions: Profiling legal hallucinations in large language models,” *Journal of Legal Analysis*, vol. 16, no. 1, pp. 64–93, 2024.
- [8] J. Tang and Y. Shang, “Advancing mental health pre-screening: A new custom gpt for psychological distress assessment,” *arXiv preprint arXiv:2408.01614*, 2024.
- [9] J. Gratch, R. Artstein, G. M. Lucas, G. Stratou, S. Scherer, A. Nazarian, R. Wood, J. Boberg, D. DeVault, S. Marsella *et al.*, “The distress analysis interview corpus of human and computer interviews,” in *LREC*. Reykjavik, 2014, pp. 3123–3128.
- [10] X. Xu, B. Yao, Y. Dong, S. Gabriel, H. Yu, J. Hendler, M. Ghassemi, A. K. Dey, and D. Wang, “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*, vol. 8, no. 1, pp. 1–32, 2024.
- [11] H. Yu and S. McGuinness, “An experimental study of integrating fine-tuned llms and prompts for enhancing mental health support chatbot system,” *Journal of Medical Artificial Intelligence*, pp. 1–16, 2024.
- [12] Y. R. Tausczik and J. W. Pennebaker, “The psychological meaning of words: Liwc and computerized text analysis methods,” *Journal of language and social psychology*, vol. 29, no. 1, pp. 24–54, 2010.
- [13] H. Li, R. Zhang, Y.-C. Lee, R. E. Kraut, and D. C. Mohr, “Systematic review and meta-analysis of ai-based conversational agents for promoting mental health and well-being,” *NPJ Digital Medicine*, vol. 6, no. 1, p. 236, 2023.
- [14] J. Ohse, B. Hadžić, P. Mohammed, N. Peperkorn, M. Danner, A. Yorita, N. Kubota, M. Rättsch, and Y. Shiban, “Zero-shot strike: Testing the generalisation capabilities of out-of-the-box llm models for depression detection,” *Computer Speech & Language*, vol. 88, p. 101663, 2024.
- [15] A. Agrawal, “Illuminate: A novel approach for depression detection with explainable analysis and proactive therapy using prompt engineering,” *arXiv preprint arXiv:2402.05127*, 2024.