

Automating PTSD Diagnostics in Clinical Interviews: Leveraging Large Language Models for Trauma Assessments

Sichang Tu,¹ Abigail Powers,¹ Natalie Merrill,¹ Negar Fani,¹

Sierra Carter,² Stephen Doogan,³ Jinho D. Choi¹

¹Emory University, Atlanta, GA, USA

²Georgia State University, Atlanta, GA, USA

³Doogood Foundation, New York, NY, USA

{sichang.tu, abigail.d.powers, natalie.merrill, nfani,jinho.choi}@emory.edu
scarter66@gsu.edu, sdoogan@rlsciences.com

Abstract

The shortage of clinical workforce presents significant challenges in mental healthcare, limiting access to formal diagnostics and services. We aim to tackle this shortage by integrating a customized large language model (LLM) into the workflow, thus promoting equity in mental healthcare for the general population. Although LLMs have showcased their capability in clinical decision-making, their adaptation to severe conditions like Post-traumatic Stress Disorder (PTSD) remains largely unexplored. Therefore, we collect 411 clinician-administered diagnostic interviews and devise a novel approach to obtain high-quality data. Moreover, we build a comprehensive framework to automate PTSD diagnostic assessments based on interview contents by leveraging two state-of-the-art LLMs, GPT-4 and Llama-2, with potential for broader clinical diagnoses. Our results illustrate strong promise for LLMs, tested on our dataset, to aid clinicians in diagnostic validation. To the best of our knowledge, this is the first AI system that fully automates assessments for mental illness based on clinician-administered interviews.

1 Introduction

Mental health has become a vital element of overall well-being. The prevalence of mental illness poses, however, a critical challenge to healthcare, underscoring the urgent need for an increased capacity of mental health services. Only 29% of people with psychosis receive formal care, leaving a significant portion completely untreated (WHO: [World Health Organization \(2021\)](#)). Aside from obstacles such as high costs, limited awareness, and stigma surrounding mental health, the shortage of the mental health workforce has been a major factor exacerbating this gap. According to WHO, the average ratio of mental health workers per 100,000 population was 13, making it difficult for people to access reliable and readily administered mental health diagnostics, as well as subsequent support and interventions.

The emergence of Large Language Models (LLMs) has suggested innovative solutions to this challenge. Several studies have explored LLM applications in mental health for condition detection ([Zhang et al., 2022](#)), support and counseling ([Ma et al., 2023b](#)) as well as clinical decision-making ([Fu et al., 2023](#)), and shown the feasibility for LLMs to enhance the workforce of mental healthcare ([Hua et al., 2024](#)). By harnessing LLMs' ability to interpret languages that involve high expertise, it is possible to mitigate the service gap in the healthcare ecosystem through the automation of condition detection and diagnosis without the need of training so many professionals, which is both costly and time-consuming.

Despite these advancements, notable limitations persist in the current research on automatic diagnosis for mental health. Most studies have focused on prevalent conditions like stress ([Lamichhane, 2023](#)) and depression ([Qin et al., 2023](#)), with scant attention to less common but more severe conditions like Post-traumatic Stress Disorder (PTSD). Moreover, while prior studies have leveraged data from social media, clinical notes, and electronic health records, very few have utilized clinical interviews, and even in those cases, they rely on basic self-administered scales estimated in dialogues between computers and patients ([Galatzer-Levy et al., 2023](#)). No work has employed diagnostic interviews between real clinicians and patients that are systematically conducted, resulting in a dearth of practical research on the automatic diagnosis of mental illness.

In this paper, we present an LLM-based system that listens to hour-long conversations between clinicians and patients and performs diagnostic assessments for PTSD. Our final model is evaluated by clinicians specialized in PTSD, suggesting a great potential for LLMs while highlighting certain limitations (Section 6). Our primary contributions are:¹

¹Our final model is publicly available through our open-source project at <https://github.com/emorynlp/TraumaNLP>.

- A new dataset comprising over 700 hours of interviews between clinicians and patients is created. Every interview consists of multiple diagnostic sections, featuring a series of questions and corresponding assessments from clinicians based on the interview contents (Section 3).
- A novel and comprehensive pipeline is developed to process the interview dataset, so it can be used to build automatic assessment models on PTSD, which can be easily adapted to a broad range of diagnostic interviews (Section 4).
- Assessment models achieving promising results are developed using two state-of-the-art LLMs, showcasing LLMs' ability to answer diagnostic questions through information extraction and text summarization on the interviews (Section 5).

To the best of our knowledge, this is the inaugural system designed to conduct diagnostic assessments on mental health while interpreting real-world interviews administered by clinicians. We believe that this work will foster clinical collaboration between human experts and Artificial Intelligence, thus promoting equitable access to appropriate care for all populations affected by mental illness.

2 Related Work

Pre-trained language models have been widely applied in many healthcare tasks (Engelhardt et al., 2023; Hu et al., 2023; Peng et al., 2023; Ma et al., 2023a; Liu et al., 2023a). The emergence of LLMs has introduced new capabilities and innovations in healthcare to this domain (Nori et al., 2023; Cascella et al., 2023). This section introduces the related research of LLMs and their applications in healthcare, particularly in mental health.

2.1 LLMs in Mental Health

The advent of LLMs like GPT (OpenAI, 2023), Llama (Touvron et al., 2023), and PaLM (Chowdhery et al., 2022) has sparked research into their applications in mental health (Ji et al., 2023). One key area is using conversational agents for mental health support and counseling, where LLMs excel at generating empathetic responses (Lai et al., 2023; Ma et al., 2023b; Loh and Raamkumar, 2023), highlighting their potential as digital companions or on-demand service providers. Additionally, the research on decision-support systems for novice counselors underscores their potential to enhance mental healthcare provision (Fu et al., 2023).

Research has also explored LLMs in disease detection and diagnosis (Zhang et al., 2022), focusing on issues like depression (Qin et al., 2023), stress (Lamichhane, 2023), and suicidality (Bhaumik et al., 2023). Closer to our work, Bartal et al. (2023) use text-based narratives from new mothers to assess childbirth-related PTSD with GPT and neural network models. Although GPT showed moderate performance, it holds promise for clinical diagnosis with further refinement. These studies typically use zero/few-shot prompting for binary or multi-label classification, demonstrating LLMs' capabilities in detecting mental health issues without fine-tuning, despite challenges like unstable responses, potential bias, and interpretation inaccuracies.

Some research has pivoted towards fine-tuning LLMs for domain-specific performance enhancement. Xu et al. (2023) present two fine-tuned models, Mental-Alpaca and Mental-FLAN-T5, outperforming GPT-3.5 and GPT-4 in multiple mental health prediction tasks. Based on Llama-2, Yang et al. (2023) train MentaLLaMA on 105K social media data enhanced by GPT. The model performance is on par with other state-of-the-art methods, while providing interpretable analysis.

2.2 LLMs in Clinical Interview and Diagnosis

Research on using LLMs on clinical interview data and diagnosis is limited. Wu et al. (2023) utilize GPT to augment the Extended Distress Analysis Interview Corpus by generating a new dataset from provided profile and rephrasing existing data. The augmented data outperforms the original imbalanced data in PTSD diagnosis. Galatzer-Levy et al. (2023) adopt Med-PaLM-2 to predict Major Depression Disorder (MDD) and PTSD on eight item Patient Health Questionnaire and PTSD Checklist-Civilian version ratings.

3 PTSD Interview Data

This study utilizes data from diagnostic interviews administered as part of a larger study on risk and resiliency to the PTSD development in a population seeking medical care (Gluck et al., 2021). Participants were recruited from waiting rooms in primary care, gynecology and obstetrics, and diabetes medical clinics at a publicly funded, safety-net hospital. Data were collected from 2012 to 2023, and inclusion criteria were ages between 18 and 65 with the capacity to provide informed consent. The parent study was conducted according to the latest version

Section	Questions	Variables	Example Question	Example Variable
LBI	31	15	What has been your primary source of income over the past month?	lbi_a1
THH	39	20	In the past, have you been treated for any emotional or mental health problems with therapy or hospitalization?	thh_tx_yesno
CRA	17	20	What would you say is the one that has been most impactful where you are still noticing it affecting you?	critaprobenotes
CAP	241	92	In the past month, have you had any unwanted memories of the [Event] while you were awake, so not counting dreams?	dsm5capscrib01 trauma1_distress

Table 1: Statistics and examples for each of the four sections employed in this study.

of the Declaration of Helsinki (World Medical Association, 2013), and consent from the participants was obtained after explaining the procedures. The informed consent was approved by our Institutional Review Board and Research Oversight Committee.

3.1 Participants

Participants were paid \$60.00 for this interview and underwent semi-structured diagnostic interviews conducted by doctoral-level clinicians or doctoral students supervised by a licensed clinical psychologist on staff. A total of 411 interviews were conducted with 336 unique participants, some of whom had follow-up interviews after >1 month. 93.4% of the participants were women and 79.5% were Black or African American ($M_{age} = 31.4$), where 38.7% had a high school education or less and 57.9% reported a monthly household income of < \$1,000.

3.2 Interview Procedures

The diagnostic interview begins with a section of the Longitudinal Interval Follow-Up Evaluation to assess global adaptive functioning across various psychosocial domains, including work, household, relationship as well as general functioning, and life satisfaction in the past month (Keller et al., 1987). Videos of the interviews are recorded using online conferencing software such as Zoom and Microsoft Teams. Each interview lasts 1.5 hours on average, involving the participant and 1-2 interviewers.

3.3 Psychiatric Diagnoses and Treatment

A total of 10 sections are applied during the interview. Among them, 4 sections are administered to the majority of participants; thus, this study focuses on those 4 sections. The first two sections, the Life Base Interview (LBI) and the Treatment History & Health (THH), are internally designed to assess the history of psychiatric diagnoses and treatment, as well as the presence of suicidality. The other two sections, the Criterion A (CRA) and the Clinician-

Administered PTSD Scale for DSM-5 (CAP), follow the standard diagnostic criteria for PTSD outlined in the Diagnostic and Statistical Manual of Mental Disorders (DSM-5; Weathers et al. (2018)). Every section is accompanied by a set of questions, linked to variables that store pertinent values derived from the corresponding answers. Table 1 shows statistics and examples for each of the 4 sections.²

LBI It assesses the participant’s functioning over the past month, addressing topics such as daily life, work, relationships with friends and family, and overall life satisfaction.

THH It covers the participant’s treatment/health history, including past physical and mental conditions as well as treatments received, such as medication and therapeutic services.

CRA It assesses whether the participant has been exposed to (threatened) death, serious injury, or sexual violence, with a focus on potential traumatic experiences the participant might have endured.

CAP It centers on issues the participant may have encountered due to traumatic events, including distress, avoidance of trauma-related stimuli, negative thoughts and feelings, and trauma-related arousal.

4 Data Processing

Every video is converted into an MP3 audio file and *transcribed* by two automatic speech recognizers, whose results are *aligned* to produce a high-quality transcript. The transcript is *segmented* into multiple sections based on the relevant questions, and each question is *paired* with its assessment result.

4.1 Transcription

Two commercial tools, Rev AI³ and Azure Speech-to-Text⁴, and an open-source tool, OpenAI Whis-

²Descriptions of all 10 sections are provided in Appendix A.

³Rev AI: <https://www.rev.ai>

⁴Azure Speech-to-Text: <https://bit.ly/42r24pA>

per (Radford et al., 2023), are tested for automatic speech recognition (ASR) on our dataset. Whisper gives the lowest Word Error Rate (WER; Klakow and Peters (2002)) of 0.13, compared to 0.21 and 0.16 from Rev AI and Azure, respectively. Whisper also exhibits better performance in handling noisy environments and numbers that Azure often misses or inaccurately transcribes (Table 2). Despite its superior ASR performance, Whisper does not identify speakers, a feature found in the others. Thus, both Azure and Whisper are run on all audios and their results are combined to obtain the best outcomes.

Tool	Examples
Azure	(1) I got 2020 on the 24 with three. Three will be 3 is turning 2116, one 15211. (2) They happened in 2017 and I'll be 60 next month, so 5556 something like that.
Whisper	(1) I got two to be 20 on the 24th, well, three, three is turning 20, one 16, one 15, two 11. (2) That happened in 2017 and I'll be 60 next month, so. 55, 56, something like that.

Table 2: Comparisons between Azure and Whisper transcripts, with equivalent tokens coded in matching colors.

4.2 Alignment

To map the speaker diarization (SD) output from Azure to the Whisper output, Align4D⁵ is used such that the first and last words of every utterance in the Azure output are aligned to their corresponding words in the Whisper transcript with speaker info, and form a speaker turn spanning all words between those words. Some words in the Whisper transcript may get left out from this mapping, which are combined with either preceding or following adjacent utterances using heuristics.

Text-based Diarization Error Rate (TDER; Gong et al. (2023)) is used, more suitable than traditional metrics like WER or Diarization Error Rate (DER; Fiscus et al. (2006)), for evaluating text-based SD. Transcripts from 29 audios produced by Microsoft Teams are used as the gold-standard, where Teams identifies speakers via different audio channels with near-perfect SD. Our aligned method achieves a TDER of 0.56, a significant improvement over the TDER of 0.62 achieved by Azure alone.

4.3 Segmentation

Each interview is conducted through multiple sections comprising a series of questions (Section 3.3), yet recorded as one continuous video. It is crucial to segment the video into sections, each of which is

⁵Align4D: <https://github.com/emorynlp/align4d>

split into sessions, where a session contains content relevant to a specific question. Here, a session is defined as a list of utterances where the first utterance includes the corresponding question, and it is followed by another session whose first utterance includes the next question (if it exists). Algorithm 1 describes how a section is matched in the transcript.

Algorithm 1: section_match(U, Q^c)

Input: U : a list of utterances, Q : a list of questions.

Output: An ordered list of tuples comprising utterance IDs and their matching scores.

```

1  $S \leftarrow \text{similarity\_matrix}(U, Q^c)$ ;
2  $T \leftarrow [\max(S_{*,i}) : 1 \leq i \leq |Q^c|]$ ;
3 if average( $T$ ) > 0.6 and
4   ( $|\text{select}(T, 0.8)| \geq 3$  or  $|\text{select}(T, 0.9)| \geq 2$ )
5   then return sequence_alignment( $S$ );
6 return  $\emptyset$ 
```

Let U be a list of utterances, and Q^c a list of core questions for a specific section.⁶ $S \in \mathbb{R}^{|U| \times |Q^c|}$ is created, where $S_{i,j}$ is a similarity score between $u_i \in U$ and $q_j \in Q^c$ (L1). $T \in \mathbb{R}^{|Q^c|}$ is then created by selecting the maximum similarity score for every question (L2). Given a function $\text{select}(T, s)$ that returns a list of scores in T greater than s , the section is matched if T 's average score is > 0.6 (L3) and if there exist at least 3 or 2 questions whose matching scores are > 0.8 or 0.9 , respectively (L4). If the section is matched, Gong et al. (2023)'s sequence alignment algorithm is applied to S , which returns an ordered list of utterance IDs and their matching scores for questions in Q^c ; otherwise, it returns an empty list (L5). In our case, Sentence Transformer is used to create embeddings for utterances & questions (Reimers and Gurevych, 2019), and cosine similarity is used to estimate the scores. Overlap between spans of two sections may occur due to incorrect matching. Algorithm 2 shows how to remove such overlaps. Let Q_i^c be a list of core questions for the i 'th section, and $R_i = \text{sm}(U, Q_i^c)$ (sm: section_match). Given (R_1, R_2) , R'_1 is created by taking a subset of R_1 whose utterance IDs exist in R_2 (L1), and R'_2 is created similarly (L2). If R'_1 contains more questions with scores > 0.6 than R'_2 , implying Q_1^c is more likely matched to the overlapped span than Q_2^c , R'_2 is removed from R_2 (L4); otherwise, R'_1 is removed from R_1 (L5).

Finally, Algorithm 3 shows how session spans are found for a specific section. C^e is a list of tuples

⁶Core questions are required for retrieving essential information, while optional questions depend on the answers to the core questions, so are often skipped during the interview.

Algorithm 2: `remove_overlap(R_1, R_2)`

Input: R_1, R_2 : ordered lists of tuples comprising utterance IDs and their matching scores for the first and second sections, respectively.

Output: (R_1, R_2): updated lists without overlaps.

```
1  $R'_1 \leftarrow [(i, s) : \forall(i, s) \in R_1 \wedge (i, *) \in R_2]$ ;
2  $R'_2 \leftarrow [(i, s) : \forall(i, s) \in R_2 \wedge (i, *) \in R_1]$ ;
3 if  $|\text{select}(R'_1, 0.6)| > |\text{select}(R'_2, 0.6)|$  then
4    $\lfloor$  return ( $R_1, R_2 \setminus R'_2$ )
5 return ( $R_1 \setminus R'_1, R_2$ )
```

comprising utterance IDs and their scores for the k 'th section created by Algorithms 1 and 2 (L1) (`ro`: `remove_overlap`). C^ℓ is created in the same manner, except adapting the Levenshtein Distance (LD) as the similarity metric (L2) (Levenshtein, 1966). `sel(C, s)` returns a list of tuples comprising utterance IDs and their matched question IDs, where the scores $> s$. `last(U, Q_k^c)` returns the first utterance ID of the $(k + 1)$ 'th section if exists; otherwise, it returns the last utterance ID of U . C is created by taking the intersection of C^e and C^ℓ whose scores > 0.8 and 0.7 , and the last utterance ID (L3).⁷

For each span U' of utterances between C_i and C_{i+1} (exclusive for both ends), a list Q' of optional questions related to C_i is created (L5-7). T^e is a list of tuples comprising utterance IDs in U' and their matched question IDs in Q' with scores > 0.8 , and T^ℓ is created using LD (L8-9). The intersection of T^e and T^ℓ is appended to a list O (L10), which is then merged with C and sorted to produce V (L11).

For each span U'' between V_i and V_{i+1} , a list Q'' of any questions have not been matched in that span is created (L14). Bipartite matching bw. U'' and Q'' are performed to find matches optimizing several criteria in Appendix B.1 (L15), accumulated, merged, and sorted to produce the final list (L16-17).

4.4 Assessment Pairing

Answers to the questions are used to determine the values of the variables (Table 1), resulting in many-to-many relations between questions and variables (many-questions to one-variable is the most common case). Our data comprises five variable types. (1) *Scale* assesses on an ordinal scale with ratings for intensity, severity, or likeness. (2) *Category* selects among binary choices or distinct class labels. (3) *Measure* captures various units such as duration, frequencies, and ages. (4) *Notes* are summarized

⁷Any section not matched by Algorithm 1 is considered absent.

Algorithm 3: `session_match($U, Q_{1..4}^c, Q_{1..4}^o, k$)`

Input: U : a list of utterances, $Q_{1..4}^{c/o}$: lists of core|optional questions for the 1..4'th sections, k : the section index to segment sessions in.

Output: (R_1, R_2): updated lists without overlaps.

```
1  $C^e \leftarrow \text{ro}(\text{sm}^e(U, Q_k^c), \text{sm}^e(U, Q_{\forall j \neq k}^c))$ ;
2  $C^\ell \leftarrow \text{ro}(\text{sm}^\ell(U, Q_k^c), \text{sm}^\ell(U, Q_{\forall j \neq k}^c))$ ;
3  $C \leftarrow (\text{sel}(C^e, 0.8) \cap \text{sel}(C^\ell, 0.7)) \cup \text{last}(U, Q_k^c)$ ;
4  $O \leftarrow \emptyset$ ;
5 for  $i \leftarrow 1$  to  $(|C| - 1)$  do
6    $U' \leftarrow$  a list of utterances between  $C_i$  and  $C_{i+1}$ ;
7    $Q' \leftarrow$  a list of questions in  $Q_k^o$  related to  $C_i$ ;
8    $T^e = \text{sel}(\text{sm}^e(U', Q'), 0.8)$ ;
9    $T^\ell = \text{sel}(\text{sm}^\ell(U', Q'), 0.7)$ ;
10   $O \leftarrow O \cup (T^e \cap T^\ell)$ ;
11  $(V, W) = (\text{sorted}(C \cup O), \emptyset)$ ;
12 for  $i \leftarrow 1$  to  $(|V| - 1)$  do
13   $U'' \leftarrow$  a list of utterances between  $V_i$  and  $V_{i+1}$ ;
14   $Q'' \leftarrow$  a list of questions in  $Q_k^c \cup Q_k^o$  that are
    between  $V_i$  and  $V_{i+1}$ ;
15   $T \leftarrow$  the best bipartite matching results between
     $U''$  and  $Q''$  optimizing several criteria in B.1;
16   $W \leftarrow W \cup T$ 
17 return  $\text{sorted}(V \cup W)$ 
```

texts documented by the interviewers. (5) *Rule* is calculated based on predefined rules derived from the other variable types. Table 3 shows the statistics of all variables for each section in our dataset.

Type	Variables					Count
	LBI	THH	CRA	CAP	Total	
Scale	7	1	0	40	48	9,722
Category	4	9	15	3	31	4,258
Measure	2	0	1	24	27	3,482
Notes	1	10	3	0	14	1,146
Rule	1	0	1	25	27	6,326

Table 3: Statistics of the five types of variables. Examples of these variables are provided in Appendix B.2.

5 Experiments

5.1 Dataset

The original data contains 411 interviews (Sec. 3). Whisper tends to generate irrelevant or repetitive sequences when prolonged silences occur, rendering about $\approx 20\%$ of the resulting transcripts unusable. To address this issue, silence removal and noise cancellation techniques are applied, recovering $\approx 80\%$ of them. Among the 393 successful transcripts, 322

VT	Template
S&C	[INTRO]. Based on the patient’s interview history, please determine <i>{keywords}</i> that the patient <i>{symptom}</i> . [RETURN]. [REASON]. The "answer" should be in the range <i>{range}</i> . <i>{attributes}</i>
M	[INTRO]. Based on the patient’s interview history, please calculate <i>{keywords}</i> that the patient have <i>{symptom}</i> . [RETURN]. [REASON]. The "answer" should be <i>{type}</i> .
N	[INTRO]. Based on the formatted data from patient’s interview, please determine whether or not the formatted data includes this specified information <i>{single_slot}</i> . [RETURN]. The "reason" gives a brief explanation on whether the formatted data includes or omits the information. The "answer" should be either "yes" or "no", indicating the presence or absence of the information in formatted data.

Table 4: Instruction templates for **Scale**, **Category**, **Measure**, and **Notes** variables. VT: Variable type, [INTRO]: Imagine you are a professional clinician, [RETURN]: Return the answer as a JSON object with "reason" and "answer" as the keys, [REASON]: The "reason" should provide a brief justification or explanation for the answer.

of them have human assessments (§4.4), which are used to evaluate our approach (Table 5).

	Audios	Hours	Turns	Tokens
Original	411	703	116,501	6,035,027
Transcribe	393	651	90,174	5,499,662
Evaluation	322	515	71,412	4,335,977

Table 5: Statistics of our PTSD interview dataset.

Compared to other interview datasets⁸, our dataset is the largest in the mental health domain. While existing datasets often involve human-machine dialogues or crowdworker simulations, ours consists of formal diagnostic interviews conducted entirely by clinicians, making it the first clinician-administered interview dataset. Additionally, our dataset aims to generate comprehensive diagnostic reports rather than just single scores, providing more detailed resource for clinical practice.

5.2 Large Language Models (LLMs)

The state-of-the-art commercial and open-source large language models, GPT-4 and Llama-2 (Touvron et al., 2023), are adapted for our experiments.⁹ For each question, a model takes all sessions related to the variable to which the question pertains (§4.4), and an instruction to provide the answer and explanation. Table 4 shows our templates including replaceable patterns to generate the instruction for each variable type. For **Scale**, *{keywords}* can be replaced with "how severe", and *{symptom}* with "have unwanted dreams in the past month". For **Category**, *{keywords}* can be replaced with "which of the following categories best describes", and *{symptom}* with "usual employment status". To constrain the answer generated by the model, details such as the answer *{range}* for **S&C**, and the

value *{type}* for **Measure** are incorporated. **S** has a special pattern *{attributes}*, directing the model to return a particular score under certain conditions.

Assessing model performance for **Notes** poses a challenge as they must be compared against text summarized by interviewers. Given the complexity of this task, it is decomposed into multiple subtasks of binary classifications, information extraction, and categorization by adapting Chain-of-Thought (Wei et al., 2023). First, GPT is asked to generate a list of slots for each **N** variable, based on a batch of summary notes from interviewers. Because many of these slots have similar meanings, albeit varying in naming, GPT is again asked to cluster them. The clusters generated by GPT are manually refined, resulting in final grouped slots that cover 95+% of the initial generation. For each of these slots, an LLM is tasked with determining if relevant content for the slot is present in the provided sessions.¹⁰

5.3 Zero-shot V.S. Few-shot Settings

Zero-shot and few-shot settings are tested across all variable types¹¹. For **Scale**, two few-shot settings are explored: one including an example for a single scale point, and the other covering examples for all scale points. For the GPT model, few-shot settings mostly outperform zero-shot settings in predicting **Category**, **Measure**, and **Notes** variables. For **Scale**, the few-shot setting with a single example results in the lowest performance. On the other hand, the few-shot setting including examples for all scale points shows a slight improvement in model performance. Thus, few-shot settings are used for all experiments with GPT. In contrast, the Llama model consistently yields inferior outcomes with few-shot settings compared to zero-shot settings, leading us to adopt zero-shot settings for all Llama experiments.

⁸Statistics of the comparison is provided in Appendix C.1.

⁹Specific versions, parameters, and costs for these large language models are provided in Appendix C.3 and C.4.

¹⁰Appendix C.5 gives slot examples for **Notes** variables.

¹¹Appendix C.2 gives details on zero/few-shot settings.

Type	Count	Accuracy		RMSE		Bias		Recall	
		GPT-4	Llama-2	GPT-4	Llama-2	GPT-4	Llama-2	GPT-4	Llama-2
Scale	9,722	58.9	46.7	1.10	1.63	-0.04	0.51	-	-
Scale _g	9,722	67.3	59.0	0.85	1.01	-0.04	0.51	-	-
Category	4,258	77.2	63.6	-	-	-	-	-	-
Measure	3,482	64.4	56.5	-	-	-0.34	-0.004	-	-
Notes	1,146	-	-	-	-	-	-	48.1	52.7
Rule	6,326	68.4	59.8	0.80	0.92	-0.15	0.44	-	-

Table 6: Model performance on all variable types (§4.4) using four evaluation metrics (§5.4).

5.4 Evaluation Metrics

Since each variable type is uniquely defined, different evaluation metrics are employed accordingly. Accuracy is computed for all types except *Notes*. For *Notes*, since the model identifies the presence of information in the provided sessions based on predefined slots, Recall is used as the primary metric to gauge the coverage of relevant information detected by the model. For *Scale*, the Root Mean Square Error (RMSE) and Bias evaluation are used. RMSE quantifies the magnitude of errors, whereas Bias evaluation calculates the proportion of positive and negative residuals, thereby revealing any directional bias in the model predictions.

5.5 Results

Table 6 gives the results for each variable type. For *Scale*, additional evaluation is conducted for CAP whose original scaling ranges from 0 to 4 where 0 indicates the absence of symptoms, 1 denotes minimal symptoms, and 2+ are considered symptoms that meet or exceed the threshold for clinical significance. To reflect this clinical demarcation, scale points are categorized into three scale groups, 0, 1, and 2+, and evaluated as *Scale_g*.¹²

GPT consistently shows significantly higher accuracy, averaging 10.5% more across all types than Llama, and reaches an accuracy of 68.4% for *Rule* accumulating outcomes of other types. Regarding RMSE, GPT exhibits an error rate of 0.8 for *Rule* using results of *Scale*, implying that it is less than one scale off from human judgment on average. In terms of Bias, ranging from -1: *completely biased to negative* to 1: *completely biased to positive*, GPT displays a marginal bias toward negative for *Scale*, while Llama shows a strong positive bias, implying that GPT is a bit conservative in predicting a higher scale, whereas Llama tends to overestimate. GPT underestimates more than Llama for *Measure*, how-

ever, showing a slight negative bias of 0.15 for *Rule*. For *Notes*, Llama exhibits better performance with a recall of 52.7% than GPT, suggesting that Llama is more effective in retrieving relevant information. Considering that these models are not fine-tuned on our data, this level of performance is very promising, as we can achieve a robust model for practical use with further training.

6 Error Analysis

A thorough error analysis is conducted by proportionally sampling 100+ examples per variable type. Six types of major errors are identified (Table 7), with only two attributed to LLMs and the remainder caused by external factors, implying that the true LLM performance may be even higher.

Misaligned Reasoning One predominant error type occurs when models deviate from instructions of the rating scheme, presenting seemingly logical reasoning, although it ultimately leads to incorrect conclusions. In Table 7, both models fail to align the key term provided by the participant, *extremely*, with the definition of score 4 - “Extreme, dramatic physical reactivity”. Llama tends to deviate further than GPT, resulting in a higher RMSE.

False Negatives is a major error type caused by:

1. Inaccurate assessments by clinicians. In Table 7, the participant reports *five times a week*, yet the clinician incorrectly records the frequency of monthly basis as 5, which should have been 20 times a month.
2. Ambiguity in *Scale* where answers may fall between two scales, resulting in potentially valid model predictions being marked incorrect.
3. The model’s inability to recognize paraphrased information in *Notes*, mistakenly indicating the absence of slot information. This issue particularly affects GPT’s performance due to its strict interpretation of wording variations.

¹²Appendix C.6/C.7 presents results for each section/variable.

Type	History	Gold	Auto		Ext
			GPT	LM	
MR	Have you had any physical reactions when something reminded you of what happened? ... I had a horrible headache. ... How many times in the past month has that happened? ... Those two times. ... How long did it take you to sort of feel back to normal? I swear. It took me a minute. I got up. I got a glass of water. It took me about. I say two to three hours. ... So how bad was that Headache? Do you think there are any other symptoms? It was extremely. I never had. I had it like that.	4	3	2	
FN	... can you think about like how often that might happen in the last month about? I feel like about like five times a week.	5	20	20	✓
EI	... when did those start for you? ... So, since around age 12, at least yeah yeah because it took me a long time to really trust my stepfather.	480	NA	108	✓
TE	... how satisfied and fulfilled have you felt about your life, with zero being like not at all, couldn't have a worse life, and 10 being perfect, couldn't have a better life? I would say a C, because it's a lot more things that I want to do to be at a 10.	2	3	3	✓
SM	So how many times in the past month would you say some things made you upset that reminded you of it? Rarely, maybe like two, three times? Very rarely.	2	1	1	✓
CR	... thinking about your work in the past month, how have you been doing? ... It's a normal, consistent, um, it's a normal, consistent routine where I do the same thing, do the same thing every day.	40	NA	40	

Table 7: Examples of the six error types. MR: Misaligned Reasoning, FN: False Negative, EI: External Information, TE: Transcription Error, SM: Session Mismatching, CR: Commonsense Reasoning. Gold: clinician’s answers, Auto: model-predicted answers. Ext: errors caused by external factors, not LLMs. NA: the model predicts None. Clinician’s questions are highlighted in blue. Patient’s key information to the questions are highlighted in red.

External Information One common issue is the absence of external information, such as the prior knowledge about the patient (e.g., medical histories, demographics) or the content of previous interview questions. In Table 7, although both models see the onset of symptoms at age 12, they fail to provide an accurate response of the total symptom duration in months because the patient’s current age (that is 52) is not provided in the transcript. In this case, GPT tends to generate a None answer, while Llama tends to hallucinate the patient’s age, and thus produces an answer based on an arbitrary assumption.

Transcription Error Transcription errors from automatic speech recognizers often cause LLMs to incorrectly interpret the answers, especially with short responses (e.g., *yes, no*, single digits like 6), medical terminologies, or non-verbal cues such as nodding. In Table 7, the number ‘6’ is incorrectly transcribed as ‘C’ in the participant’s response.

Session Mismatching A question can be mismatched with the transcript, especially when the clinician extensively paraphrases it. In such cases, the segmented session may or may not contain all the necessary information to answer the question. In Table 7, both models correctly answer based on the patient’s response (1: Minimal). However, due to the mismatch, the session is missing a part where the patient also indicates 2 (clearly present but still manageable), which is recorded as gold.

Commonsense Reasoning The models’ limitations extend to inferring basic human experiences. Unable to deduce standard working hours from *a normal, consistent routine* in Table 7, the models fall short of clinician-like assumptions of a typical 40-hour workweek, showcasing a gap in applying commonsense logic to the assessment.

7 Conclusion

In this study, we undertook the task of automating PTSD diagnostics using 411 clinician-administered interviews. To ensure the data quality, we develop an end-to-end pipeline streamlining transcription, alignment, segmentation, and assessment pairing. We also construct a pioneering framework for this task by leveraging two state-of-the-art LLMs. Our findings reveal the substantial potential of LLMs in assisting clinicians with diagnostic validation and decision-making processes. Our error analysis suggests future directions for improvement, such as incorporating external information or commonsense knowledge to engineer more comprehensive instructions. We envision that this framework holds promise for addressing a broader spectrum of mental health conditions and offers novel insights into LLM applications within the mental health domain. We plan to collect more data and train a custom LLM to better preserve patients’ privacy, and develop a dialogue system to conduct the interviews.

Limitations

Although the experiment results prove the capability of LLMs to automate PTSD diagnosis, their applications in real-world unsupervised clinical settings are premature. To avoid the possible negative influence of model errors on the patients, we recommend using this framework as a supportive tool for clinicians in diagnostics and decision-making.

It should be noted that the clinician annotated gold assessment data is not perfect, which may affect evaluation accuracy. However, this framework makes it easier to identify and refine the inaccuracies in the gold assessment data and thus improve its overall validity. We leave the data augmentation as the next step of our future work.

In addition, the experiments in this paper utilize LLMs without fine-tuning. One limitation is that we have little control over the model predictions. The models, especially Llama-2, generate unexpected outputs that violate the instructions. Furthermore, data privacy concerns restrict the use of models like GPT for clinical data. To address these issues and enhance framework adaptability, future work will focus on developing more controllable, open-sourced models that guarantee data protection in line with clinical domain restrictions.

Due to strict Institutional Review Board (IRB) regulations concerning the confidentiality of real patient information, we are unable to release the dataset, even in an anonymized format. However, recognizing the importance of contributing to the research community, we are pleased to announce that we will release the framework utilized in our study. This, we believe, will facilitate further research and innovation, as our methodology is versatile and can be adapted to a wide array of mental health conditions, provided the requisite interview question sets and video/transcripts are available.

Ethical Considerations

The diagnostic interview data used in this paper was collected with informed consent approved by the Institutional Review Board (IRB) and Research Oversight Committee. The authors and clinicians involved in the research have passed Research, Ethics, Compliance, and Safety Training through Collaborative Institutional Training Initiative¹³ (CITI Program). For the use of LLMs, this study exclusively employs anonymized interviews,

¹³<https://about.citiprogram.org>

ensuring the confidentiality and privacy of all participants. All practices in this research adhere to the ACL Code of Ethics.

Acknowledgements

We gratefully acknowledge the support of the DooGood Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DooGood Foundation.

References

- Alon Bartal, Kathleen Jagodnik, Sabrina Chan, and Sharon Dekel. 2023. [Chatgpt Demonstrates Potential for Identifying Psychiatric Disorders: Application to Childbirth-Related Post-Traumatic Stress Disorder](#).
- Runa Bhaumik, Vineet Srivastava, Arash Jalali, Shanta Ghosh, and Ranganathan Chandrasekaran. 2023. Mindwatch: A smart cloud-based ai solution for suicide ideation detection leveraging large language models. *medRxiv*, pages 2023–09.
- Marco Cascella, Jonathan Montomoli, Valentina Bellini, and Elena Bignami. 2023. [Evaluating the Feasibility of Chatgpt in Healthcare: An Analysis of Multiple Clinical and Research Scenarios](#). *Journal of Medical Systems*, 47(1).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling Language Modeling with Pathways](#).
- Zachary Englhardt, Chengqian Ma, Margaret E. Morris, Xuhai "Orson" Xu, Chun-Cheng Chang, Lianhui Qin, Daniel McDuff, Xin Liu, Shwetak Patel, and Vikram Iyer. 2023. [From Classification to Clinical Insights: Towards Analyzing and Reasoning About Mobile and Behavioral Health Data With Large Language Models](#).

- Jonathan G. Fiscus, Jerome Ajot, Martial Michel, and John S. Garofolo. 2006. The Rich Transcription 2006 Spring Meeting Recognition Evaluation. In *Proceedings of International Workshop on Machine Learning and Multimodal Interaction*, pages 309–322.
- Guanghui Fu, Qing Zhao, Jianqiang Li, Dan Luo, Changwei Song, Wei Zhai, Shuo Liu, Fan Wang, Yan Wang, Lijuan Cheng, Juan Zhang, and Bing Xiang Yang. 2023. [Enhancing Psychological Counseling with Large Language Model: A Multifaceted Decision-Support System for Non-Professionals.](#)
- Isaac R. Galatzer-Levy, Daniel McDuff, Vivek Nataraajan, Alan Karthikesalingam, and Matteo Malgaroli. 2023. [The Capability of Large Language Models to Measure Psychiatric Functioning.](#)
- Rachel L. Gluck, Georgina E. Hartzell, Hayley D. Dixon, Vasiliki Michopoulos, Abigail Powers, Jennifer S. Stevens, Negar Fani, Sierra Carter, Ann C. Schwartz, Tanja Jovanovic, Kerry J. Ressler, Bekh Bradley, and Charles F. Gillespie. 2021. [Trauma exposure and stress-related disorders in a large, urban, predominantly african-american, female sample.](#) *Archives of Women’s Mental Health*, 24(6):893–901.
- Chen Gong, Peilin Wu, and Jinho D. Choi. 2023. [Aligning Speakers: Evaluating and Visualizing Text-based Speaker Diarization Using Efficient Multiple Sequence Alignment.](#) In *Proceedings of the 35th IEEE International Conference on Tools with Artificial Intelligence, ICTAI’23.*
- Jonathan Gratch, Ron Artstein, Gale Lucas, Giota Stroutou, 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).
- Yan Hu, Qingyu Chen, Jingcheng Du, Xueqing Peng, Vipina Kuttichi Keloth, Xu Zuo, Yujia Zhou, Zehan Li, Xiaoqian Jiang, Zhiyong Lu, Kirk Roberts, and Hua Xu. 2023. [Improving Large Language Models for Clinical Named Entity Recognition via Prompt Engineering.](#)
- Yining Hua, Fenglin Liu, Kailai Yang, Zehan Li, Yi-han Sheu, Peilin Zhou, Lauren V. Moran, Sophia Ananiadou, and Andrew Beam. 2024. [Large Language Models in Mental Health Care: a Scoping Review.](#)
- Shaoxiong Ji, Tianlin Zhang, Kailai Yang, Sophia Ananiadou, and Erik Cambria. 2023. [Rethinking Large Language Models in Mental Health Applications.](#)
- Martin B. Keller, Philip W. Lavori, Barbara Friedman, Eileen Nielsen, Jean Endicott, Pat McDonald-Scott, and Nancy C. Andreasen. 1987. [The Longitudinal Interval Follow-up Evaluation. A comprehensive method for assessing outcome in prospective longitudinal studies.](#) *Archives Of General Psychiatry*, 44(6):540–548.
- Dietrich Klakow and Jochen Peters. 2002. [Testing the Correlation of Word Error Rate and Perplexity.](#) *Speech Communication*, 38(1):19–28.
- Tin Lai, Yukun Shi, Zicong Du, Jiajie Wu, Ken Fu, Yichao Dou, and Ziqi Wang. 2023. [Psy-Llm: Scaling up Global Mental Health Psychological Services with Ai-based Large Language Models.](#)
- Bishal Lamichhane. 2023. [Evaluation of Chatgpt for Nlp-based Mental Health Applications.](#)
- Vladimir I Levenshtein. 1966. [Binary Codes Capable of Correcting Deletions, Insertions, and Reversals.](#) *Soviet Physics Doklady*, 10(8):707–710.
- Jialin Liu, Changyu Wang, and Siru Liu. 2023a. [Utility of Chatgpt in Clinical Practice.](#) *Journal of Medical Internet Research*, 25:e48568.
- June M. Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi Liao, and Jiamin Wu. 2023b. [Chatcounselor: A Large Language Models for Mental Health Support.](#)
- Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. [Towards emotional support dialog systems.](#)
- Siyuan Brandon Loh and Aravind Sesagiri Raamkumar. 2023. [Harnessing Large Language Models’ Empathetic Response Generation Capabilities for Online Mental Health Counselling Support.](#) *arXiv preprint arXiv:2310.08017.*
- Chong Ma, Zihao Wu, Jiaqi Wang, Shaochen Xu, Yaonai Wei, Zhengliang Liu, Xi Jiang, Lei Guo, Xiaoyan Cai, Shu Zhang, Tuo Zhang, Dajiang Zhu, Dinggang Shen, Tianming Liu, and Xiang Li. 2023a. [Impressiongpt: An Iterative Optimizing Framework for Radiology Report Summarization with Chatgpt.](#)
- Zilin Ma, Yiyang Mei, and Zhaoyuan Su. 2023b. [Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support.](#) In *AMIA Annual Symposium Proceedings*, volume 2023, page 1105. American Medical Informatics Association.
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023. [Capabilities of Gpt-4 on Medical Challenge Problems.](#)
- OpenAI. 2023. [Gpt-4 Technical Report.](#)
- Cheng Peng, Xi Yang, Kaleb E Smith, Zehao Yu, Aokun Chen, Jiang Bian, and Yonghui Wu. 2023. [Model Tuning or Prompt Tuning? A Study of Large Language Models for Clinical Concept and Relation Extraction.](#)

- Wei Qin, Zetong Chen, Lei Wang, Yunshi Lan, Weijie Ren, and Richang Hong. 2023. [Read, diagnose and chat: Towards explainable and interactive llms-augmented depression detection in social media.](#)
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. [Robust Speech Recognition via Large-scale Weak Supervision.](#) In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*, pages 28492–28518.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-Bert: Sentence Embeddings using Siamese Bert-Networks.](#)
- David V Sheehan, Yves Lecrubier, K Harnett Sheehan, Patricia Amorim, Juris Janavs, Emmanuelle Weiller, Thierry Hergueta, Roxy Baker, Geoffrey C Dunbar, et al. 1998. The mini-international neuropsychiatric interview (mini): the development and validation of a structured diagnostic psychiatric interview for dsm-iv and icd-10. *Journal of clinical psychiatry*, 59(20):22–33.
- Ying Shen, Huiyu Yang, and Lin Lin. 2022. [Automatic depression detection: An emotional audio-textual corpus and a gru/bilstm-based model.](#)
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open Foundation and Fine-Tuned Chat Models.](#)
- Michel Valstar, Björn Schuller, Kirsty Smith, Timur Al-maev, Florian Eyben, Jarek Krajewski, Roddy Cowie, and Maja Pantic. 2014. [Avec 2014: 3d dimensional affect and depression recognition challenge.](#) In *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge, AVEC '14*, page 3–10, New York, NY, USA. Association for Computing Machinery.
- Frank W. Weathers, Michelle J. Bovin, Daniel J. Lee, Denise M. Sloan, Paula P. Schnurr, Danny G. Kaloupek, Terence M Keane, and Brian P. Marx. 2018. [The Clinician-Administered PTSD Scale for DSM-5 \(CAPS-5\): Development and initial psychometric evaluation in military veterans.](#) *Psychological Assessment*, 30(3):383–395.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.](#)
- World Health Organization. 2021. *Mental health atlas 2020.* World Health Organization.
- World Medical Association. 2013. [World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects.](#) *The Journal of the American Medical Association*, 310(20):2191–2194.
- Yuqi Wu, Jie Chen, Kaining Mao, and Yanbo Zhang. 2023. [Automatic Post-Traumatic Stress Disorder Diagnosis via Clinical Transcripts: A Novel Text Augmentation with Large Language Models.](#) In *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 1–5.
- Xuhai Xu, Bingsheng Yao, Yuanzhe Dong, Saadia Gabriel, Hong Yu, James Hendler, Marzyeh Ghassemi, Anind K. Dey, and Dakuo Wang. 2023. [Mental-Llm: Leveraging Large Language Models for Mental Health Prediction via Online Text Data.](#)
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Sophia Ananiadou, and Jimin Huang. 2023. [Mental-lama: Interpretable Mental Health Analysis on Social Media with Large Language Models.](#)
- 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.](#)
- Tianlin Zhang, Annika M. Schoene, Shaoxiong Ji, and Sophia Ananiadou. 2022. [Natural language processing applied to mental illness detection: a narrative review.](#) *npj Digital Medicine*, 5(1).

A Section Details

Table 8 - 11 give examples for 4 core sections. Each example includes the standard interview Question, the Variable that the question belongs to, and the example Sessions between the Clinician and the Participant.

The Mini International Neuropsychiatric Interview (MINI) is a brief, structured diagnostic interview for diagnosing 17 major psychiatric disorders (Sheehan et al., 1998). We adopt 6 modules from MINI to assess conditions such as Major Depressive Episode (MDE), Mania & Hypomania (MH), PTSD (past incidents), Psychosis Symptoms (PS), Substance Use Disorder (SUD), and Alcohol Use Disorder (AUD). Table 12 provides an example from the MDE module.

Q	What has been your primary source of income over the past month?
V	lbi_a1
S	C: You got to do it all over again. Are you working full time? P: Yes.
Q	How would you rate your overall satisfaction on a scale of 1 to 10, with 1 being the best and 10 being the worst?
V	lbi_e1
S	C: In the past month, like how satisfied have you felt with your life? If we were doing like a scale of one to 10, one is like, it's the worst. This is the worst I've ever had in my life. 10 being like, this is, I'm living my best life. Living my life like it's golden. P: I actually feel like that now. I actually do. Cause until January 1st of this year, I had been unemployed the last two years.

Table 8: Two examples of the LBI section.

Q	Do you have any current physical health conditions?
V	thh_medicalcond
S	C: OK, so now we're going to move on to talking about your health and treatment history. Do you currently have, do you have any current physical health conditions? Did you say no? OK, I couldn't hear what you were saying. Go ahead. P: I have a skin condition called eczema.
Q	In the past, have you been treated for any emotional/mental health problems with therapy or hospitalization?
V	thh_tx_yesno
S	C: In the past, have you been treated for any emotional or mental health problem with therapy or hospitalization? P: No. Yes.

Table 9: Two examples of the THH section.

Q	Tell me a little bit more about what happened.
V	trauma1whathappened
S	C: OK, and what would that be? P: My mom worked at the airport here in xxx. It was the food catering place. They put the food, made the food for the planes. When I was a child, every year, they would sponsor a day at xxx. They would go out there and barbecue. We took over the whole picnic area. You had free entrance to the park, plus tickets to do all the little fair games and all that good stuff. Having a good time. My mom asked my stepfather to go with us because he had a car. He said he didn't wanna go and he wasn't going nowhere. So my mom put us all on the bus. We drove the bus out there. When we came home, it was like 11 o'clock. Of course, we living in xxx. You know that bus ride was long. It was dark, dark when we got home and she had all three of her children with her. My mom unlocked the door, closed that door, the house was pitch black. That man shot down them steps at my mama and all three of her children five times.

Table 10: An example of the CRA section.

Q	Tell me a little bit more about what happened.
V	dsm5capscribt_01trauma1_distress
S	C: To this day, let's say over the past month. So since like the beginning of April, end of March, have you had unwanted memories of this event? Does it randomly pop into your mind at all? Like while you're awake? P: Well, actually my daughter's in an abusive relationship. So yes, I do think about it a lot. Every time I see her, all I think about is my mom. How she endured it.
Q	How often in the past month?
V	dsm5capscritc02trauma1_num
S	C: So in the last month, thinking about the things that you have tried to avoid, how often would you say you've done that? P: I guess every day. I don't know. I just, the most I've done is just, and me avoiding stuff is me just sitting here smoking and playing my video game. That avoids me from thinking about anything negative in my life. And I just try to avoid that.

Table 11: Two examples of the CAP section.

Q	For the past two weeks, were you depressed or down, or felt sad, empty or hopeless most of the day, nearly every day?
V	miniv7_mde_c_a1
S	C: I'm going to ask you some different questions. We're going to focus on the past two weeks right now. So for the past two weeks, did you feel depressed, down, sad, empty or hopeless for most of the day, almost every day the past two weeks? P: Um, no.

Table 12: An example of the MINI section.

B Data Preprocessing Details

B.1 Final Matching Criteria

The best bipartite result should follow the criteria:

- All matching IDs need to be ascending.
- Only edges whose embedding cosine similarity > 0.4 are kept.
- Maximize: $y = \sum_{i=1}^n a_i \cdot x_i$, subject to $x_i \geq 0$, for $i = 1, \dots, n$.

In our case, let $n = 9$, with the following variables:

- x_1 : the sum of Sentence Transformer (ST) cosine similarity scores of all edges
- x_2 : the sum of Levenshtein Distance (LD) similarity scores of all edges
- x_3 : the average ST cosine similarity scores of all matched questions
- x_4 : the average LD similarity scores of all matched questions
- x_5 : the total number of matched core questions
- x_6 : the total number of matched questions that take the maximum ST cosine similarity result
- x_7 : the total number of matched questions that take the maximum LD similarity result
- x_8 : the total number of matched core questions that take the maximum ST cosine similarity result
- x_9 : the total number of matched core questions that take the maximum LD similarity result

And the coefficients are set as:

- $a_1 = a_2 = 1$
- $a_3 = a_4 = 1$
- $a_5 = a_6 = a_7 = 0.1$
- $a_8 = a_9 = 0.2$

B.2 Variable Examples

Table 13 - 17 show examples for each variable type. Every example includes the Variable name, replaceable Patterns for prompt generation (Section 5), answer Range, and covered Questions. Note that *Measure*, *Notes*, and *Rule* variables do not have a predefined range. And *Rule* variables are calculated from the results of their Related Variables.

V	dsm5capscrib01trauma1_distress
P	{ <i>keywords</i> }: how intense in the past month { <i>symptom</i> }: unwanted memories of the traumatic event while awake { <i>attributes</i> }: - If the symptom only exists in dreams, the answer should be 0. - If the symptom is not perceived as involuntary and intrusive, the answer should be 0.
R	0: None, 1: Minimal, minimal distress or disruption of activities 2: Clearly Present, distress clearly presented but still manageable, some disruption of activities 3: Pronounced, considerable distress, difficulty dismissing memories, marked disruption of activities 4: Extreme, incapacitating distress, cannot dismiss memories, unable to continue activities
Q	In the past month, have you had any unwanted memories of it while you were awake, so not counting dreams? - How does it happen that you start remembering it? - Are these unwanted memories, or are you thinking about it on purpose? - How much do these memories bother you? - Are you about to put them out of your mind and think about something else? - Overall, how much of a problem is this for you? - How so?

Table 13: An example of the *Scale* variable. Questions start with - are optional questions that might be skipped based on the participant's response.

V	lbi_a1
P	{ <i>keywords</i> }: which of the following categories best describes, { <i>symptom</i> }: usual employment status
R	1: Full-Time Gainful Employment 2: Part-Time Gainful Employment (30 hours or less/week) 3: Unemployed But Expected by Self or Others 4: Unemployed But Not Expected by Self or Others (e.g., physically disabled) 5: Retired 6: Homemaker 7: Student (Includes Part-Time) 8: Leave of Absence Due to Medical Reasons (e.g., holding job; plans to return) 9: Volunteer Work - Full Time 10: Volunteer Work - Part Time 11: Other 888: N/A
Q	What has been your primary source of income over the past month?

Table 14: An example of the *Category* variable.

V	dsm5capscribtb01trauma1_num
P	{ <i>keywords</i> }: how intense in the past month { <i>symptom</i> }: unwanted memories of the traumatic event while awake { <i>type</i> }: an integer representing the frequency of the symptom in the past month
Q	- How often have you had these memories in the past month?

Table 15: An example of the *Measure* variable. The corresponding question for this question is optional, which might be skipped if the participant denies the presence of the symptom.

V	critaprobnotes
P	{ <i>slots</i> }: - trauma_reactions - trauma_details - life_changes - coping_and_changes - worldview_changes - health_concerns - family_and_social_context - nightmare_details - intrusive_experiences - trauma_cognition - trust_and_safety - impact_assessment - age_and_time_factors - substance_use - therapy_and_progress - eating_disorders
Q	You discussed a number of traumas in the last visit with our team members. What would you say is the one that has been most impactful where you are still noticing it affecting you? -* How much do you think about what happened to this day? -* How often do you have nightmares about what happen? -* How much did it change the way you think about yourself and the world? - In the past month, which of these have you thought about more often or had nightmares about or find yourself purposely avoiding thinking about? - Are there any other stressors that you find yourself thinking about when you don't want to or avoiding?

Table 16: An example of the *Notes* variable. Questions start with - are optional questions which might be skipped based on the participant’s response. Questions start with * are recurrent questions which might be asked multiple times during the interview.

V	dsm5capscribtb01trauma1
R	0: Absent 1: Mild/subthreshold 2: Moderate/threshold 3: Severe/markedly elevated 4: Extreme/incapacitating
RV	dsm5capscribtb01trauma1_distress dsm5capscribtb01trauma1_num

Table 17: An example of the *Rule* variable.

C Experiments Details

C.1 Dataset Comparison

Table 18 gives the comparison with related datasets in the metal health domain.

Dataset	A	H	Turns	Utters
DAIC (Gratch et al., 2014)	189	51	-	-
AViD (Valstar et al., 2014)	300	240	-	-
EATD (Shen et al., 2022)	162	2.26	-	-
Psych8k (Liu et al., 2023b)	260	260	-	-
D4 (Yao et al., 2022)	-	-	28,855	81,559
ESConv (Liu et al., 2021)	-	-	-	31,410
Ours	322	515	71,412	142,824

Table 18: Comparisons with existing mental health interview/dialogue datasets in terms of **Audio** counts, total **Hours**, total and utterances.

C.2 Details on Zero-shot/Few-shot Settings

We randomly sampled 30 instances for each variable type and asked both models to predict under zero-shot and few-shot settings. For the GPT model, few-shot settings generally yield better performance. However, the Llama model consistently fails to follow instructions as the context length grows, leading to significant degradation with few-shot prompting. Additionally, we observed a 28% increase in the likelihood of generating an unexpected response format, such as deviating from the requested JSON format, when using few-shot settings.

C.3 Experiment Costs

GPT-4 The pricing of the GPT-4 Turbo model is \$0.01/1K tokens for input and \$0.03/1K tokens for output. We spend approximately \$300 (upper bound) to complete GPT experiments in this paper.

Llama-2 We use a single NVIDIA H100 GPU for Llama inferences with a batch size of 1, taking

Type	Zero-shot		Few-shot	
	GPT-4	Llama-2	GPT-4	Llama-2
Scale	60.0	50.0	63.3	36.7
Scale ₁	-	-	56.7	40.0
Category	43.3	40.0	46.7	33.3
Measure	56.7	56.7	60.0	50.0
Notes	41.0	42.7	43.6	34.9

Table 19: Model performance on zero-shot and few-shot settings. Scale₁ refers to the few-shot setting that only include one example for a single scale point. Accuracy is the metric used for all variable types except Notes variables, which are evaluated using Recall.

roughly 10 seconds per request. Completing a full set of experiments on all samples requires ~3 days.

C.4 LLM Configurations

We utilize gpt-4-1106-preview, the latest GPT-4 Turbo model, and llama-2-70b-chat-hf, the largest Llama-2 model. For GPT, to enhance the stability and consistency of the model output, we configure the *temperature* parameter to 0. This adjustment makes the model’s response more deterministic. Besides, we also employ parameters exclusive to GPT-4 Turbo and GPT-3.5 Turbo, namely *response_format* and *seed*. Setting *response_format* to "json_object" constrains the model to generate parsable JSON strings, facilitating easier data handling and analysis. Despite ChatGPT’s non-deterministic nature, *seed* parameter enables users to obtain consistent outputs across multiple requests, as long as there are no changes at the system level.

As for the Llama, we conduct experiments involving different *temperature*, *top_p*, and *repetition_penalty* separately. The results indicate that the model gives better performance with a *temperature* setting of 0.3, a *top_p* of 0.9, and a *repetition_penalty* of 1.

C.5 Slot Examples for Notes Variable

Table 20 outlines the process for generating, merging, and formatting the slots in *Notes* variables (§5.2). Initially, we compile all clinician-summarized notes for each *Notes* variable and input them into the GPT model using the NSG prompt to produce a list of slots. Due to potential overlaps, the NSM prompt directs the model to consolidate these slots into clusters, ensuring both conciseness and comprehensiveness. Subsequently, the NSF prompt is used to format both the gold-standard summaries and the corresponding interview ses-

sions, facilitating a straightforward comparison of the structured slot arrangements.

C.6 Model Performance by Sections

Table 21 presents model performances by each section. Note that THH section lacks *Measure* and *Rule* variables, whereas CRA section does not contain *Scale* variables. The grouped scale_g is exclusively applied within the CAP section.

C.7 Model Performance by Variables

Table 22 lists results for each variable following the four evaluation metrics (Section 5.4).

Step	Template
NSG	<p>As a clinician who has conducted interviews with multiple patients, you are tasked with structuring the interview data into a more organized format. To achieve this, identify general "slots" from the interview question and answers. These slots should represent key themes or types of information that can be adapted to various responses from different patients.</p> <p>For each identified slot, provide a brief explanation of why it has been chosen, focusing on its relevance and utility in categorizing interview data.</p> <p>Your findings should be presented in a JSON format as a list, for example: [{"reason": "This slot captures the primary health concern of the patient, a common theme across all interviews", "slot": "primary_health_concern"}, {"reason": "This slot pertains to the patient's lifestyle habits, which is crucial for understanding health context", "slot": "lifestyle_habits"}].</p> <p>Remember to ensure that the slots are broad enough to be applicable across different patient responses yet specific enough to offer meaningful categorization.</p>
NSM	<p>Imagine you are a clinician who documents patient interviews in a structured, slot-filling manner. Sometimes, certain slots may have overlapping or similar content. Your task is to review a given list of slots and merge those that are similar. The merged results should be returned as a JSON object, where each key represents a merged slot, and the corresponding value is a list of the original slots that have been combined under this merged category.</p> <p>For instance, if the list of slots is: ["daily_routine", "work_events", "daily_activity", "daytime_activities", "work_routine"], a possible merged result could be: {"daily_routine": ["daily_routine", "daily_activity", "daytime_activities"], "work": ["work_events", "work_routine"]}.</p> <p>When you receive a list of slots, analyze and merge them accordingly, ensuring that the merged slots are logically grouped and accurately represent the original information categories.</p>
NSF	<p>Imagine you are a professional clinician. Based on the patient's interview history, please extract specific information and fill in the following slots: {slots}. If the interview history does not provide information for any of these slots, please enter an empty string ("") for that slot. Return the answer as a JSON object.</p>

Table 20: Prompts used for Notes Variable Slot structure Generation, Merging, and Formatting.

	Type	Count	Accuracy		RMSE		Bias		Recall	
			GPT-4	Llama-2	GPT-4	Llama-2	GPT-4	Llama-2	GPT-4	Llama-2
LBI	Scale	1,281	54.6	44.7	1.26	1.42	0.46	0.45	-	-
	Category	594	74.6	67.3	-	-	-	-	-	-
	Measure	99	68.7	66.7	-	-	-0.16	-0.09	-	-
	Notes	203	-	-	-	-	-	-	42.0	50.8
	Rule	215	43.3	37.7	0.94	0.98	0.44	0.43	-	-
THH	Scale	29	55.2	51.7	1.20	1.25	0.23	0.43	-	-
	Category	1,527	92.6	85.9	-	-	-	-	-	-
	Notes	633	-	-	-	-	-	-	52.5	59.8
CRA	Category	1,737	63.7	42.4	-	-	-	-	-	-
	Measure	143	63.6	55.9	-	-	-0.58	-0.36	-	-
	Notes	310	-	-	-	-	-	-	47.2	43.5
	Rule	146	91.8	71.9	0.38	0.97	0.83	-0.51	-	-
CAP	Scale	8,412	59.6	47.0	1.07	1.66	-0.14	0.52	-	-
	Scale _g	8,412	69.3	61.2	0.77	0.93	-0.15	0.53	-	-
	Category	400	81.0	64.5	-	-	-	-	-	-
	Measure	3,240	64.2	56.3	-	-	-0.33	0.01	-	-
	Rule	5,965	68.8	60.4	0.81	0.92	-0.19	0.46	-	-

Table 21: Model performances on 4 sections (§4.4) using four evaluation metrics (§5.4).

Variable	Count	Acc		RMSE		Bias		Recall	
		GPT	LM2	GPT	LM2	GPT	LM2	GPT	LM2
Scale Variable									
lbi_a2b	199	53.8	37.7	1.31	2.05	0.37	0.68	-	-
lbi_a3	201	56.7	52.7	0.96	0.99	0.31	0.47	-	-
lbi_a4	63	50.8	55.6	0.90	1.02	0.23	0.29	-	-
lbi_b1a_family	207	46.9	54.6	1.78	1.36	0.49	0.11	-	-
lbi_b2	212	60.8	44.8	1.19	1.02	0.47	0.49	-	-
lbi_d	194	52.1	38.7	1.29	1.29	0.53	0.53	-	-
lbi_e1	205	58.5	36.1	0.91	1.66	0.65	0.42	-	-
dx_understanding	29	55.2	51.7	1.20	1.25	0.23	0.43	-	-
dsm5capscritb01trauma1_distress	257	59.5	53.3	0.89	1.67	-0.44	0.48	-	-
dsm5capscritb02trauma1_distress	254	69.7	53.9	0.72	1.16	-0.56	0.32	-	-
dsm5capscritb03trauma1_distress	249	67.9	51.8	0.73	1.35	-0.05	0.68	-	-
dsm5capscritb04trauma1_distress	259	57.1	40.5	0.96	1.25	-0.35	0.47	-	-
dsm5capscritb05trauma1_distress	243	63.8	56.8	0.90	1.04	-0.11	0.37	-	-
dsm5capscritc01trauma1_distress	253	46.2	39.9	1.77	1.92	-0.34	0.53	-	-
dsm5capscritc02trauma1_distress	243	58.0	45.7	0.99	1.20	-0.04	0.64	-	-
dsm5capscritd01trauma1_distress	242	66.1	53.7	0.92	1.13	-0.10	0.30	-	-
dsm5capscritd02trauma1_distress	256	56.6	36.7	0.85	1.31	-0.06	0.83	-	-
caps5trauma1related_d02	164	57.9	55.5	0.97	0.85	-0.71	0.07	-	-
dsm5capscritd03trauma1_distress	248	61.7	58.9	0.94	0.92	-0.56	0.24	-	-
dsm5capscritd04trauma1_distress	252	56.0	49.2	0.93	1.13	-0.03	0.55	-	-
caps5trauma1related_d04	160	63.8	54.4	0.89	0.84	-0.28	0.10	-	-
dsm5capscritd05trauma1_distress	253	57.7	47.8	1.00	1.18	-0.08	0.53	-	-
caps5trauma1related_d05	138	53.6	44.9	1.06	0.96	-0.56	0.21	-	-
dsm5capscritd06trauma1_distress	255	53.5	47.5	1.01	1.23	0.09	0.66	-	-
caps5trauma1related_d06	156	51.3	41.0	0.98	0.90	-0.47	0.35	-	-
dsm5capscritd07trauma1_distress	257	59.5	45.5	0.88	1.22	0.04	0.67	-	-
caps5trauma1related_d07	128	55.5	44.5	0.96	0.94	-0.16	0.35	-	-
dsm5capscritc01trauma1_distress	257	60.3	46.7	0.79	1.13	0.33	0.78	-	-
caps5trauma1related_e01	148	52.7	33.8	3.54	3.44	-0.74	0.06	-	-
dsm5capscritc02trauma1_distress	251	67.3	61.0	0.71	1.11	0.02	0.31	-	-
caps5trauma1related_e02	50	74.0	58.0	1.09	1.26	-0.38	0.43	-	-
dsm5capscritc03trauma1_distress	255	51.4	47.1	1.09	1.20	0.32	0.54	-	-
caps5trauma1related_e03	155	50.3	51.6	0.93	0.86	-0.40	0.17	-	-
dsm5capscritc04trauma1_distress	252	63.1	52.8	0.85	1.05	-0.03	0.60	-	-
caps5trauma1related_e04	117	50.4	53.0	0.99	0.88	-0.55	0.13	-	-
dsm5capscritc05trauma1_distress	256	59.8	53.5	0.81	0.99	-0.13	0.58	-	-
caps5trauma1related_e05	161	57.8	41.6	1.09	0.99	-0.79	0.51	-	-
dsm5capscritc06trauma1_distress	256	53.5	52.7	1.02	1.06	0.09	0.37	-	-
caps5trauma1related_e06	181	63.0	38.7	1.0	10.2	-0.67	0.37	-	-
dsmiv_future_frequency_current	251	80.1	48.6	0.81	6.02	0.40	0.80	-	-
dsmiv_future_intens_current	246	69.1	40.2	0.93	1.77	0.61	0.90	-	-

continued on next page

continued from previous page

Variable	Count	Acc		RMSE		Bias		Recall	
		GPT	LM2	GPT	LM2	GPT	LM2	GPT	LM2
dsm5capscritg_trauma1_distress	228	53.9	43.4	1.08	1.39	0.35	0.69	-	-
dsm5capscritg_trauma1_impair	226	51.8	42.0	0.93	1.27	-0.28	0.57	-	-
dsm5capscritg_trauma1_fx	205	54.1	29.3	1.10	1.56	-0.04	0.81	-	-
dsm5depersonalization_sev	255	67.5	52.2	0.80	1.25	-0.08	0.49	-	-
caps5trauma1related_diss01	76	53.9	31.6	1.16	1.26	0.31	0.19	-	-
dsm5derealization_sev	249	63.1	30.9	0.98	1.74	0.20	0.88	-	-
caps5trauma1related_diss02	70	55.7	27.1	1.11	1.25	-0.03	0.53	-	-
Category Variable									
lbi_a1	200	70.0	41.0	-	-	-	-	-	-
lbi_student	201	95.0	89.1	-	-	-	-	-	-
lbi_c1a	192	57.8	71.9	-	-	-	-	-	-
lbi_c2	1	100	100	-	-	-	-	-	-
thh_medicalcond	206	92.7	88.8	-	-	-	-	-	-
thh_tx_curr_yesno	215	94.9	80.9	-	-	-	-	-	-
thh_tx_yesno	233	89.7	87.6	-	-	-	-	-	-
feedback_helpful	79	94.9	89.9	-	-	-	-	-	-
thh_txneed_yesno	96	92.7	88.5	-	-	-	-	-	-
thh_psychmed_curr_yesno	194	92.3	88.7	-	-	-	-	-	-
thh_psychmed_yesno	198	95.5	93.4	-	-	-	-	-	-
thh_suicide_yesno	236	90.7	77.1	-	-	-	-	-	-
thh_suicide_pw_yesno	70	94.3	78.6	-	-	-	-	-	-
trauma1lifeeventscl	146	61.6	12.3	-	-	-	-	-	-
trauma1_exposure_type___1	146	77.4	67.1	-	-	-	-	-	-
trauma1_exposure_type___2	146	77.4	43.2	-	-	-	-	-	-
trauma1_exposure_type___3	146	67.8	28.1	-	-	-	-	-	-
trauma1_exposure_type___4	146	65.8	22.6	-	-	-	-	-	-
caps_e1_lt	145	62.1	45.5	-	-	-	-	-	-
caps_e1_ltsel	73	64.4	64.4	-	-	-	-	-	-
caps_e1_ltother	74	41.9	44.6	-	-	-	-	-	-
caps_e1_si	146	43.8	39.7	-	-	-	-	-	-
caps_e1_siself	61	54.1	65.6	-	-	-	-	-	-
caps_e1_siother	61	60.7	29.5	-	-	-	-	-	-
caps_e1_tpi	146	54.1	52.7	-	-	-	-	-	-
caps_e1_tpiself	79	84.8	75.9	-	-	-	-	-	-
caps_e1_tpiother	77	49.4	26.0	-	-	-	-	-	-
trauma1_nomemory	145	75.2	44.8	-	-	-	-	-	-
dsm5caps_critf_cur1_yesno	202	78.7	41.6	-	-	-	-	-	-
dsm5caps_critf_cur1_c	198	83.3	87.9	-	-	-	-	-	-
Measure Variable									
lbi_a2a	99	68.7	66.7	-	-	41.9	45.5	-	-
trauma1age	143	63.6	55.9	-	-	21.2	31.7	-	-
dsm5capscritb01trauma1_num	162	63.6	58.0	-	-	37.3	52.9	-	-

continued on next page

continued from previous page

Variable	Count	Acc		RMSE		Bias		Recall	
		GPT	LM2	GPT	LM2	GPT	LM2	GPT	LM2
dsm5capscrib02trauma1_num	98	74.5	63.3	-	-	28.0	52.8	-	-
dsm5capscrib03trauma1_num	84	72.6	59.5	-	-	47.8	76.5	-	-
dsm5capscrib04trauma1_num	177	62.1	58.8	-	-	17.9	42.5	-	-
dsm5capscrib05trauma1_num	137	59.1	57.7	-	-	28.6	50.0	-	-
dsm5capscritc01trauma1_num	170	59.4	53.5	-	-	31.9	54.4	-	-
dsm5capscritc02trauma1_num	140	63.6	54.3	-	-	27.5	54.7	-	-
dsm5capscritd01trauma1_num	87	50.6	48.3	-	-	32.6	82.2	-	-
dsm5capscritd02trauma1_num	168	76.2	69.0	-	-	47.5	59.6	-	-
dsm5capscritd03trauma1_num	120	65.0	57.5	-	-	23.8	43.1	-	-
dsm5capscritd04trauma1_num	166	72.3	68.1	-	-	39.1	45.3	-	-
dsm5capscritd05trauma1_num	138	65.9	59.4	-	-	42.6	46.4	-	-
dsm5capscritd06trauma1_num	155	69.7	63.2	-	-	27.7	40.4	-	-
dsm5capscritd07trauma1_num	140	61.4	60.0	-	-	40.7	53.6	-	-
dsm5capscrit01trauma1_num	135	65.9	62.2	-	-	21.7	54.9	-	-
dsm5capscrit02trauma1_num	61	83.6	68.9	-	-	80.0	89.5	-	-
dsm5capscrit03trauma1_num	159	73.0	67.3	-	-	37.2	34.6	-	-
dsm5capscrit04trauma1_num	131	68.7	60.3	-	-	24.4	50.0	-	-
dsm5capscrit05trauma1_num	168	69.0	66.1	-	-	21.2	35.1	-	-
dsm5capscrit06trauma1_num	184	72.8	61.4	-	-	40.0	31.0	-	-
dsmcaps_critf_cur1_nummonths	191	49.7	22.5	-	-	60.4	81.8	-	-
dsm5caps_critf_cur1_b	181	35.7	22.0	-	-	17.9	19.0	-	-
dsm5depersonalization_num	84	59.5	51.2	-	-	32.4	65.9	-	-
dsm5derealization_num	3	100	33.3	-	-	0.00	50.0	-	-
Notes Variable									
life_base_typicalday	203	-	-	-	-	-	-	42.0	50.8
thh_medicalcond_desc	100	-	-	-	-	-	-	56.8	80.1
thh_tx_curr_descr	59	-	-	-	-	-	-	53.6	73.8
thh_tx_descr	135	-	-	-	-	-	-	44.0	57.4
dx_knowledge	33	-	-	-	-	-	-	59.4	48.5
dx_lackknowledge	20	-	-	-	-	-	-	60.7	37.9
feedback_info	66	-	-	-	-	-	-	75.1	48.1
thh_txneed_desc	45	-	-	-	-	-	-	59.7	49.4
thh_psychmed_descr	89	-	-	-	-	-	-	40.4	59.0
thh_suicide_desc	73	-	-	-	-	-	-	56.9	67.0
thh_suicide_pw_desc	13	-	-	-	-	-	-	62.2	62.3
critaprobnotes	143	-	-	-	-	-	-	50.8	37.4
trauma1whathappened	143	-	-	-	-	-	-	42.7	51.4
trauma1describe	24	-	-	-	-	-	-	51.4	48.9
Rule Variable									
lbi_e2	215	43.3	37.7	0.94	0.98	0.44	0.43	-	-
caps_e1_crita	146	91.8	71.9	0.38	0.97	0.83	-0.51	-	-
dsm5capscrib01trauma1	253	62.8	63.6	0.81	0.90	-0.51	0.28	-	-

continued on next page

continued from previous page

Variable	Count	Acc		RMSE		Bias		Recall	
		GPT	LM2	GPT	LM2	GPT	LM2	GPT	LM2
dsm5capscrib02trauma1	250	88.0	70.0	0.53	0.80	-0.47	0.47	-	-
dsm5capscrib03trauma1	246	86.2	63.4	0.54	0.96	0.00	0.82	-	-
dsm5capscrib04trauma1	255	67.5	60.0	0.86	0.95	-0.57	0.25	-	-
dsm5capscrib05trauma1	241	74.7	69.7	0.73	0.75	-0.18	0.26	-	-
dsm5capscrit01trauma1	250	55.2	54.8	0.94	0.97	-0.64	0.36	-	-
dsm5capscrit02trauma1	242	71.9	61.2	0.83	0.93	-0.29	0.51	-	-
dsm5capscritd01trauma1	239	81.2	66.5	0.68	1.02	0.24	0.60	-	-
dsm5capscritd02trauma1	222	62.6	46.4	0.79	1.09	-0.16	0.83	-	-
dsm5capscritd03trauma1	246	72.0	72.0	0.85	0.74	-0.48	0.36	-	-
dsm5capscritd04trauma1	251	63.7	62.2	0.94	1.02	-0.08	0.35	-	-
dsm5capscritd05trauma1	252	59.9	53.6	0.98	1.00	-0.19	0.42	-	-
dsm5capscritd06trauma1	254	55.9	50.8	1.03	1.10	-0.07	0.57	-	-
dsm5capscritd07trauma1	255	63.1	60.0	0.85	0.95	-0.17	0.59	-	-
dsm5capscite01trauma1	255	72.5	51.4	0.69	0.90	0.03	0.66	-	-
dsm5capscite02trauma1	250	90.4	76.8	0.38	0.73	0.75	0.90	-	-
dsm5capscite03trauma1	220	57.3	58.6	0.91	0.96	0.09	0.43	-	-
dsm5capscite04trauma1	250	75.6	72.0	0.71	0.79	-0.08	0.63	-	-
dsm5capscite05trauma1	254	65.7	67.7	0.80	0.77	-0.36	0.54	-	-
dsm5capscite06trauma1	254	55.9	52.8	1.05	0.96	0.05	0.27	-	-
dsmcaps_critf_admin	28	75.0	100	0.50	0.00	-1.00	-1.00	-	-
dsm5depersonalization	246	85.4	64.2	0.61	0.78	-0.06	0.70	-	-
dsm5derealization	243	75.3	39.1	0.69	1.14	0.27	0.76	-	-
dsm5capsglobalvalidtrauma1	255	63.5	63.5	0.84	0.84	-1.00	-1.00	-	-
dsm5capsglobalsevtrauma1	254	44.1	42.9	0.91	0.97	0.21	0.45	-	-

Table 22: Model performances on all variable (§4.4) using four evaluation metrics (§5.4).