

MentalSeek-Dx: Towards Progressive Hypothetico-Deductive Reasoning for Real-world Psychiatric Diagnosis

Xiao Sun^{1*}, Yuming Yang^{1*}, Xinyi Jiang², Yu Tian³, Junnan Zhu⁴, Jiang Zhong^{1†}
Qin Lei⁵, Jingwang Huang¹, Haoyang Zeng¹, Xinyu Zhou⁵, Xin Xiao¹, Kaiwen Wei^{1†}

¹ School of Computer Science, Chongqing University,

² School of Computer Science and Engineering, UNSW Sydney,

³ Department of Computer Science and Technology, Tsinghua University,

⁴ MAIS, Institute of Automation, Chinese Academy of Sciences,

⁵ The First Affiliated Hospital of Chongqing Medical University

sunx@stu.cqu.edu.cn, ymyang@cqu.edu.cn, jiangzhong@cqu.edu.cn, weikaiwen@cqu.edu.cn

Abstract

Mental health disorders represent a burgeoning global public health challenge. While Large Language Models (LLMs) have demonstrated potential in psychiatric assessment, their clinical utility is severely constrained by benchmarks that lack ecological validity and fine-grained diagnostic supervision. To bridge this gap, we introduce **MentalDx Bench**, the first benchmark dedicated to disorder-level psychiatric diagnosis within real-world clinical settings. Comprising 712 de-identified electronic health records annotated by board-certified psychiatrists under ICD-11 guidelines, the benchmark covers 76 disorders across 16 diagnostic categories. Evaluation of 18 LLMs reveals a critical *paradigm misalignment*: strong performance at coarse diagnostic categorization contrasts with systematic failure at disorder-level diagnosis, underscoring a gap between pattern-based modeling and clinical hypothetico-deductive reasoning. In response, we propose **MentalSeek-Dx**, a medical-specialized LLM trained to internalize this clinical reasoning process through supervised trajectory construction and curriculum-based reinforcement learning. Experiments on MentalDx Bench demonstrate that MentalSeek-Dx achieves state-of-the-art (SOTA) performance with only 14B parameters, establishing a clinically grounded framework for reliable psychiatric diagnosis. The dataset and code are available at [MentalSeek-Dx](#).

1 Introduction

Mental health disorders affect over one billion individuals worldwide, constituting a major contributor to the global burden of disease¹. However, the acute shortage of mental health professionals and clinical infrastructure often leads to delayed diagnoses and inadequate treatment, underscoring the

*These authors contributed equally to this work.

†Corresponding authors.

¹WHO : World Mental Health Today

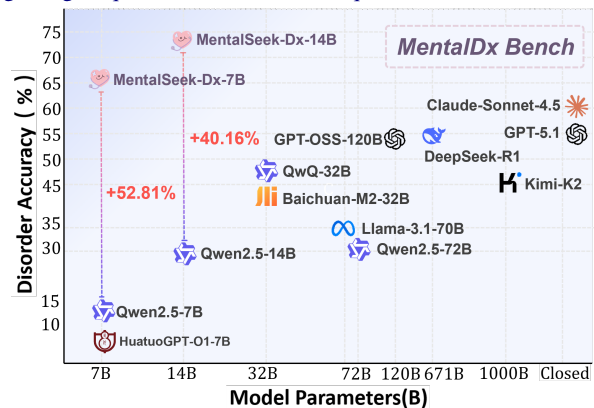


Figure 1: Diagnostic accuracy results on the **MentalDx** benchmark, evaluated across mainstream LLMs. The proposed model MENTALSEEK-DX achieves state-of-the-art performance while operating at smaller scales.

urgent need for scalable, automated diagnostic alternatives. In this context, LLMs have emerged as promising tools for psychiatric assessment (Ge et al., 2025). While early applications primarily focused on population-level surveillance or detecting surface-level psychological distress (Mangalik et al., 2024; Zhuang et al., 2024), recent efforts have advanced toward identifying specific clinical risks, such as suicidal ideation and depression symptoms (Levkovich et al., 2024; Kim, 2025; Lho et al., 2025; Wang et al., 2025). Despite this progress, a significant leap remains from symptom detection to the rigor of *disorder-level diagnosis* required in clinical decision support.

A primary bottleneck impeding this advancement is the lack of clinically realistic benchmarks. Existing evaluation frameworks are limited in two fundamental respects. (1) *Data Unrealism*: Most benchmarks rely on non-clinical proxies, such as social media posts (Wu et al., 2025) or synthetic narratives (Hu et al., 2024), which diverge substantially from the structured, jargon-rich, and information-dense documentation used in routine psychiatry. (2) *Annotation Coarseness*: Prevailing

datasets often capture only surface-level symptom mentions or assign broad disorder labels without grounding in formal diagnostic criteria (e.g., ICD-11) (Hengle et al., 2024). Consequently, models optimized on these benchmarks may learn superficial pattern matching (Liu et al., 2025; Ignashina et al., 2025; Lho et al., 2025) rather than the nuanced differential diagnosis essential for real-world clinical application.

To address these limitations, we introduce **MentalDx Bench**, the first benchmark specifically designed to evaluate fine-grained, disorder-level psychiatric diagnosis in authentic clinical settings. Unlike prior datasets (see Table 1), MentalDx Bench comprises 712 de-identified medical records directly sourced from medical centers. Each case is rigorously annotated by board-certified psychiatrists following strict ICD-11 guidelines, providing a dual-layer label structure that spans 16 broad diagnostic categories and 76 specific disorders. This clinically grounded design allows for a high-fidelity assessment of an LLM’s ability to navigate the complexity of real-world psychiatric cases.

We evaluate 18 mainstream and specialized LLMs on MentalDx Bench. The results reveal a striking phenomenon: while models achieve respectable accuracy in coarse diagnostic categorization (e.g., identifying a *Mood Disorder*), their performance degrades sharply on fine-grained disorder-level diagnosis (e.g., distinguishing *Bodily Distress Disorder*). Through psychiatrist-led diagnostic reviews, we identify the root cause as a fundamental **paradigm misalignment**. Mainstream LLMs predominantly rely on *associative shortcuts*—linking symptom keywords directly to disease labels. In contrast, clinical psychiatry demands a structured, *hypothetico-deductive reasoning* process: synthesizing symptoms, generating differential hypotheses, and rigorously excluding alternatives based on diagnostic criteria. This cognitive gap renders standard LLMs opaque and unreliable for precise clinical inference.

To resolve this misalignment, we propose **MentalSeek-Dx**, a domain-specialized LLM designed to align with the cognitive process of human psychiatrists. Inspired by clinical decision-making theories (Collins, 2017), MentalSeek-Dx adopts a progressive reasoning framework trained in two stages. First, we employ *Hypothetico-Deductive Trajectory Construction*, a supervised strategy decomposing diagnosis into explicit steps: symptom categorization, hypothesis generation, and differen-

Benchmark	Input	Scope	Hierarchy	Types
MindSET (Mankarious et al., 2025)	Narrative	Mental Health	Stratification	7
CRADLE Bench (Byun et al., 2025)	Narrative	Mental Health	Stratification	7
ANGST (Hengle et al., 2024)	Narrative	Mood	Stratification	2
IMHI Benchmark (Yang et al., 2024)	Narrative	Mood	Stratification	9
CLPsych (Kim, 2025)	Narrative	Suicide	Risk	1
D4 (Yao et al., 2022)	Narrative	Depression	Symptom	2
ReDSM5 (Bao et al., 2025)	Narrative	Depression	Symptom	1
MDPP (Fu et al., 2025)	Narrative	Depression	Symptom	1
EmoDep (Oh et al., 2026)	Clinical Note	Depression	Symptom	1
SMDH (Cohan et al., 2018)	Narrative	Psychiatry	Disorder	9
PsychiatryBench(Fouda et al., 2025)	Book	Psychiatry	Category+Disorder	50
MentalDx Bench (Ours)	EHR	Psychiatry	Category+Disorder	76

Table 1: Comparison of representative benchmarks for mental health diagnosis, where our **MentalDx Bench** features realistic EHR inputs and dual-level diagnostic labels at the category and disorder levels.

tial diagnosis. Second, we introduce *Reward-Based Curriculum Reinforcement Learning*, which progressively guides the model from mastering coarse categorization to refining fine-grained differential logic. As shown in Fig. 1, experiments demonstrate this approach instills clinically coherent reasoning, enabling our 14B-parameter model to achieve SOTA diagnostic accuracy, surpassing significantly larger general-purpose LLMs. In summary, our contributions are as follows:

(1) We introduce **MentalDx Bench**, the first benchmark dedicated to fine-grained, disorder-level diagnosis grounded in real-world clinical records. Expert-annotated under ICD-11 guidelines, it spans 76 disorders across 16 categories, bridging the gap between computational metrics and ecological clinical validity.

(2) Our systematic evaluation of 18 LLMs quantifies a critical **paradigm misalignment**: current models rely on superficial pattern associations rather than the structured clinical reasoning essential for psychiatry, resulting in significant performance collapse on fine-grained inference.

(3) We propose **MentalSeek-Dx**, a domain-specialized framework that aligns LLMs with clinical logic via progressive hypothetico-deductive reasoning and curriculum reinforcement learning. It achieves SOTA performance with only 14B parameters, demonstrating that reasoning alignment is the key to reliable psychiatric diagnosis.

2 Related Work

Psychiatric Benchmarks. Psychiatric NLP benchmarks have progressed from disorder-level classification of social media texts (Cohan et al., 2018) to symptom-level evaluation with clinically grounded labels (Yao et al., 2022; Bao et al., 2025; Fu et al., 2025). With LLMs, benchmarks increasingly adopt

structured diagnostic setups—first in narrow domains (Hengle et al., 2024; Yang et al., 2024), then across broader taxonomies and reasoning chains (Mankarious et al., 2025; Byun et al., 2025; Fouda et al., 2025)—expanding into suicide risk detection (Kim, 2025) and large-scale user-level corpora (Zheng et al., 2025). Some ground evaluation in real-world clinical decision-making (Liu et al., 2025), or simulate multi-step psychiatric workflows (Johri et al., 2025). Input modalities have shifted from narratives to clinical notes (Oh et al., 2026), enabling EHR-based modeling for outcome prediction (Sao and Lim, 2024), phenotype extraction (Frydman-Gani et al., 2025), and clinical reasoning (Zhu et al., 2025). However, most benchmarks remain based on non-clinical data, limiting clinical representativeness.

Psychiatric Methods. Large language models (LLMs) have emerged as promising tools to support psychiatric assessment (Ge et al., 2025). Early applications focused on detecting subtle indicators of psychological distress (Lho et al., 2025), enabling risk screening (Shin et al., 2024; Chim et al., 2024), subclinical symptom identification, and population-level stratification (Obradovich et al., 2024; Orrù et al., 2025). More recently, research has shifted toward diagnostic prediction, with two main approaches: one uses LLM-derived features combined with downstream classifiers (Radwan et al., 2024; Shewcraft et al., 2025; Lan et al., 2025b), and the other trains domain-specific LLMs in an end-to-end fashion (Xu et al., 2024; Weber et al., 2025; Dai et al., 2025). This line of work reflects growing interest in deploying LLMs to assist Psychiatric clinical decision-making. However, disorder-level, general-purpose psychiatric diagnostic modeling remains underexplored.

3 MentalDx Bench

We introduce MentalDx Bench, the first benchmark specifically designed for disorder-level psychiatric diagnosis across the full spectrum of psychiatric conditions, featuring real-world clinical data and expert annotations strictly aligned with guidelines:

(1) Data Source. The *MentalDx Bench* consists of 712 de-identified psychiatric diagnostic records sourced from partner medical centers. To ensure data quality and diagnostic integrity, we exclude records with incomplete workups, ambiguous conclusions, comorbidities, or sensitive personal content. All data processing is conducted entirely

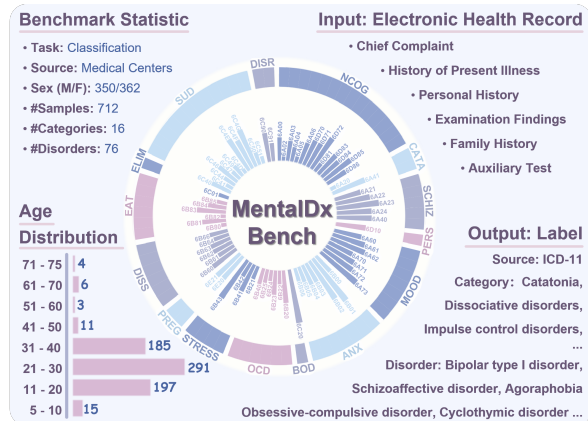


Figure 2: Overview of the *MentalDx Bench*, based on real-world medical centers, showing benchmark statistics, inputs, outputs, and data distributions.

within the hospital’s secure infrastructure.

(2) Ethical Compliance. The study protocol and annotation procedures were reviewed and approved by the institutional ethics boards of all participating clinical centers, ensuring the legal and ethical integrity. All data collection and labeling were conducted in accordance with the ethical principles of the *Declaration of Helsinki*. As all records were fully de-identified prior to processing, the requirement for individual informed consent was waived.

(3) Annotation Schema. To ensure standardized labeling, all cases in *MentalDx Bench* are annotated according to the ICD-11 taxonomy for *Mental, Behavioural, and Neurodevelopmental Disorders* (Chapter 6A), using both official diagnostic names and ICD codes. A team of 12 board-certified psychiatrists conducted the labeling process following ICD-11 guidelines and internal consensus protocols. Each case receives two levels of diagnostic labels: a Category-level label denoting the broader group (e.g., Mood Disorders) and a Disorder-level label specifying the exact condition (e.g., Schizophrenia). All annotations are independently assigned by two senior psychiatrists, with inter-rater agreement measured by Cohen’s κ coefficient ($\kappa > 0.8$). Disagreements are resolved through expert consensus meetings. Cases involving comorbidity are excluded to ensure label clarity.

(4) Dataset Characteristics. *MentalDx Bench* is designed to support structured, fine-grained diagnostic evaluation while capturing real-world psychiatric diversity across age groups and genders. To ensure balanced label distribution and mitigate long-tail effects, we sample a uniform number of cases, as shown in Fig 2. Each case contains a struc-

Model	ANX	CATA	SUD	BOD	STRESS	DISR	DISS	ELIM	EAT	PREG	MOOD	NCOG	NDEV	OCD	PERS	SCHIZ	CA	DA	JA
Small Language Models																			
Qwen2.5-7B	11.11	30.00	14.12	10.00	40.00	0.00	4.26	25.00	17.14	21.74	12.37	13.79	18.92	12.50	0.00	12.50	47.47	14.75	10.96
Qwen2.5-14B	18.06	63.33	29.41	0.00	52.50	44.44	21.28	25.00	62.86	82.61	22.68	13.79	48.65	40.00	0.00	10.94	61.66	29.92	24.72
Qwen3-32B	30.56	20.00	22.35	25.00	32.50	50.00	31.91	75.00	48.57	52.17	22.68	35.63	51.35	62.50	15.38	1.56	58.71	31.04	26.26
QWQ-32B	48.61	70.00	35.29	25.00	35.00	66.67	42.55	75.00	91.43	91.30	41.24	35.63	75.68	75.00	0.00	29.69	76.26	47.89	46.63
Large Language Models																			
LLaMA-3.1-70B	47.22	30.00	28.24	5.00	35.00	88.89	27.66	75.00	62.86	65.22	25.77	25.29	64.86	47.50	7.69	17.19	70.65	35.53	33.15
Qwen2.5-72B	34.72	40.00	32.94	0.00	35.00	38.89	34.04	75.00	74.29	78.26	23.71	16.09	62.16	55.00	30.77	6.25	70.51	33.57	29.78
Kimi-K2	55.56	33.33	44.71	0.00	57.50	72.22	48.94	100.00	91.43	65.22	24.74	32.18	89.19	70.00	23.08	10.94	75.70	45.08	43.40
Qwen3-Max	50.00	73.33	47.06	20.00	52.50	94.44	48.94	100.00	91.43	86.96	32.99	48.28	89.19	82.50	23.08	20.31	83.15	52.67	52.11
GPT-5.1	66.67	53.33	47.06	5.00	42.50	77.78	53.19	100.00	91.43	82.61	38.14	48.28	78.38	85.00	30.77	40.62	81.18	54.49	53.79
DeepSeek-R1	50.00	63.33	48.24	15.00	60.00	77.78	61.70	100.00	94.29	95.65	45.36	40.23	78.38	85.00	15.38	37.50	78.37	55.20	54.49
Claude-Sonnet-4.5	88.89	83.33	47.06	15.00	52.50	88.89	61.70	100.00	85.71	95.65	40.21	49.43	78.38	87.50	23.08	34.38	80.76	59.69	58.15
Medical-specific Large Language Models																			
ClinicalGPT-R1	9.72	23.33	10.59	5.00	27.50	5.56	2.13	0.00	14.29	17.39	16.49	5.75	13.51	5.00	0.00	1.56	36.10	10.53	6.46
HuatuogPT-o1-7B	9.72	20.00	12.94	5.00	25.00	5.56	0.00	50.00	5.71	8.70	15.46	6.90	5.41	10.00	0.00	17.19	43.26	11.24	8.01
GPT-OSS-20B	52.78	63.33	23.53	15.00	35.00	83.33	44.68	100.00	77.14	82.61	34.02	32.18	70.27	72.50	0.00	45.31	74.30	45.65	43.96
MedGemma-27B	54.17	50.00	28.24	25.00	42.50	55.56	29.79	75.00	74.29	65.22	25.77	32.18	72.97	70.00	23.08	40.62	66.99	42.84	39.89
Baichuan-M2	38.89	66.67	35.29	5.00	37.50	72.22	42.55	75.00	88.57	91.30	28.87	37.93	62.16	67.50	0.00	35.94	64.61	44.38	41.71
HuatuogPT-o1-72B	45.83	50.00	35.29	10.00	35.00	66.67	40.43	75.00	80.00	69.57	26.80	19.54	70.27	65.00	30.77	18.75	74.02	39.75	36.94
GPT-OSS-120B	54.17	66.67	38.82	20.00	50.00	88.89	51.06	50.00	94.29	65.22	41.24	43.68	70.27	75.00	38.46	45.31	75.84	52.53	50.00

Table 2: Performance of mainstream LLMs on MentalDx Bench (%). Models are grouped by parameter scale to examine how size affects performance across diagnostic categories. Each cell reports a model’s score on the corresponding task. For clarity, the best result in each category is shown in dark, the second-best in light.




Figure 3: Statistics of expert-annotated errors across 360 cases. Results highlight a paradigm misalignment between LLM reasoning and clinical diagnostic logic.

tured clinical narrative composed of multiple sections—chief complaint, history of present illness, personal and family history, examination findings, and auxiliary test results—providing rich contextual input for diagnostic modeling. Each case is labeled with a single disorder and its parent category, forming a two-level labeling structure used in a single-label classification setup. Details and the ethics statement can be found in Appendix 8.

4 Data Corpus and Curation

Each record was transformed into a bench case through a structured pipeline designed to preserve clinical realism while protecting privacy. Specifically, an LLM was used to de-identify the original clinical records by generalizing or masking sensitive attributes such as exact dates, locations, and organizations into non-identifiable clinical descrip-

tors like temporal ranges or generalized life events. All processed record profiles were subsequently reviewed and refined by licensed psychiatrists to ensure psychopathological consistency and clinical plausibility. We conducted a multi-stage privacy audit combining automated detection tools and expert verification. As shown in Table 3, all detected sensitive attributes, including 214 dates, 132 locations, and 138 organizations, were fully abstracted, with no confirmed cases of personally identifiable information leakage. All diagnostic conclusions and structured diagnosis fields were removed. The leakage audit in Table 4 shows that no exact matches are observed across all diagnostic entities, indicating that explicit label leakage was effectively eliminated. The remaining fuzzy matches, such as 65 for diagnostic categories and 180 for specific disorders, correspond to symptom descriptions or historical conditions rather than explicit diagnoses, as further summarized in Table 5.

5 Experiments on MentalDx Bench

5.1 Experiment Settings

We present the task definition, evaluation metrics, and evaluated models for MentalDx Bench.

Task Definition. After removing comorbid cases during preprocessing, the task is formulated as a multi-stage, single-label classification problem under a constrained candidate set. Given a clinical case x and a predefined list of ICD-11 categories and disorders, the model M is required to predict

Table 3: Privacy Audit Results

Entity Category	Presidio Flags	Anonymized	Breaches
MentalDx-Bench (n=712)			
Patient/Physician Name	0	–	0
Specific Date	214	214 (100%)	0
Detailed Location	132	132 (100%)	0
Organization/Workplace	138	138 (100%)	0
Contact Identifiers	0	–	0
Training Set (5% Sample, n=750)			
Patient/Physician Name	0	–	0
Specific Date	229	229 (100%)	0
Detailed Location	164	164 (100%)	0
Organization/Workplace	115	115 (100%)	0
Contact Identifiers	0	–	0

Table 4: Quantitative Audit Results

Target Entities	Count	Exact Match	Fuzzy Match
MentalDx-Bench (n=712)			
Diagnostic Categories	16	0.00% (0/712)	9.1% (65/712)
Specific Disorders	76	0.00% (0/712)	25.3% (180/712)
ICD-11 Codes	76	0.00% (0/712)	0.00% (0/712)
Diagnosis Headers*	5	0.00% (0/712)	0.00% (0/712)

* Diagnosis headers include terms such as “Final Diagnosis” or “Clinical Impression.”

the most likely category c and the corresponding disorder d based on the case description.

Evaluation Metrics. We evaluate model performance using three metrics: *Category Accuracy (CA)*, *Disorder Accuracy (DA)*, and *Joint Accuracy (JA)*. *CA* and *DA* measure whether the predicted category or disorder matches the ground truth, respectively, while *JA* requires both to be correct simultaneously. Let (c_i, d_i) denote the *gt* category and disorder for the i -th sample, and (\hat{c}_i, \hat{d}_i) be the predictions. The metrics are defined as:

$$CA, DA = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{y}_i = y_i], \quad (1)$$

$$JA = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{c}_i = c_i \wedge \hat{d}_i = d_i].$$

where N is the total number of test samples and $\mathbb{1}[\cdot]$ denotes the indicator function. With candidates provided in the input, we extract predictions through exact string matching.

Baselines. We evaluate 18 language models spanning 3 categories as follows: 1) Small Language Models, including: Qwen2.5-7B and 14B (Team et al., 2024) Qwen3-32B (Yang et al., 2025), and QWQ-32B (Team et al., 2024). 2) LLMs, including: LLaMA-3.1-70B (Patterson et al., 2022), Qwen2.5-72B (Team et al., 2024), Qwen3-Max (2025-10-30) (Yang et al., 2025), Kimi-

Table 5: Contextual breakdown of fuzzy matches.

Source Category	Proportion	Example
Past Medical History	44%	“History of anxiety and depression.”
Clinical Observations	56%	“Persistent depressive symptoms observed.”

K2 (2025-09-05) (Team et al., 2025), DeepSeek-R1 (2025-05-28) (Guo et al., 2025), GPT-5.1 (2025-11-13) (Bubeck et al., 2025), Claude-Sonnet-4.5 (2025-09-29) (Appel et al., 2025). 3) Medical-specific LLMs, including: ClinicalGPT-R1 (Lan et al., 2025a), HuatuoGPT-O1-7B and 72B (Chen et al., 2024), GPT-OSS-20B and 120B (Agarwal et al., 2025), Medgemma-27B (Sellergren et al., 2025), Baichuan-M2 (Dou et al., 2025).

5.2 Experiment Results

We aim to address the following research questions:

- **RQ1:** How well do existing models perform on psychiatric diagnosis?
- **RQ2:** What are the underlying causes of model limitations?

5.2.1 RQ1: How well do existing models perform on psychiatric diagnosis?

We evaluate 18 LLMs on MentalDx Bench as shown in Table 6 (The full category list and disorder explanations are provided in Appendix 9.1.), and find that most perform well on category prediction—for example, QWQ-32B achieves 76.26%, Qwen3-Max 83.15%, and GPT-OSS-120B 75.84%. However, some medical-specific models lag behind their general counterparts of similar scale; ClinicalGPT-R1 reaches only 36.10%, which is 11.37 percentage points lower than its base model Qwen2.5-7B (47.47%). Even with medical pre-training, gaps in coverage and distributional mismatch hinder generalization to psychiatric tasks.

While existing LLMs perform well on category-level diagnosis, their performance drops markedly at the disorder level. Qwen3-Max, despite achieving the highest *CA*, reaches only 52.67% on disorder prediction; ClinicalGPT-R1 performs even worse, with a *JA* of just 6.46%. Coarse-grained evaluation can obscure fine-grained diagnostic failures, motivating the need for fine-grained benchmarks. See Appendix 9.3 for examples.

5.2.2 RQ2: What are the underlying causes of model limitations?

We conducted expert validation to investigate diagnostic errors. For each of the 18 models, we sampled 20 incorrect predictions (360 total) and

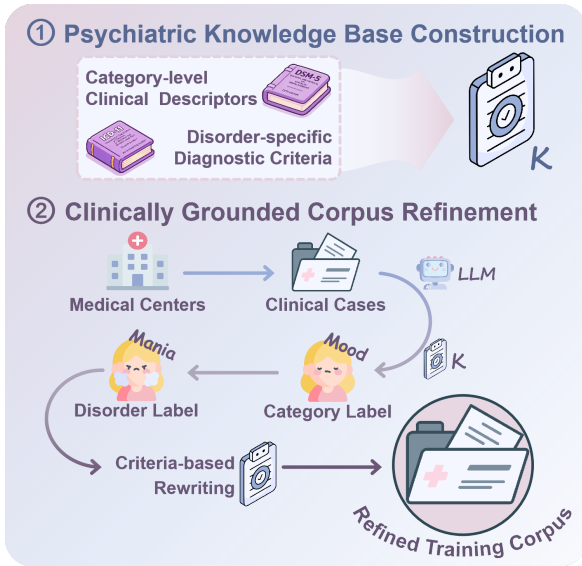


Figure 4: Overview of psychiatric knowledge base construction and clinically grounded corpus refinement.

assigned 12 psychiatrists to independently assess a random subset of 30 cases. As shown in Figure 3, most errors fall into three categories—*wrong inference* (25%), *identification failure* (22%), and *symptom omission* (21%)—whereas superficial issues like typos (7%) or refusals (4%) were far less common. This distribution suggests that the dominant failure modes stem not from surface-level mistakes, but from deeper deficits in reasoning and symptom understanding. In particular, many models exhibited fragmented diagnostic logic, including unclear symptom delineation, incomplete analysis of candidate conditions, and a lack of differential diagnostic ability. Error definitions and procedures are described in Appendix 9.5.

6 MentalSeek-Dx

Motivated by the above findings, we introduce MENTALSEEK-DX, a medical-specific LLM for psychiatric reasoning diagnosis.

6.1 Pretraining Setup

Structured Psychiatric Knowledge Base. We construct a structured knowledge base \mathcal{K} in collaboration with psychiatrists to support diagnostic reasoning and model training. \mathcal{K} follows a two-level hierarchy of diagnostic categories c and fine-grained disorders d , with each entry comprising clinical definitions, symptom clusters, functional impairments, diagnostic criteria, and differential considerations. Its design follows ICD-11, clinical guidelines, and expert consensus, ensuring cover-

age and clinical validity. Given a c or d , \mathcal{K} enables automatic retrieval of relevant knowledge to support both training and reasoning trajectory building. A detailed flow is shown in Figure 4. Excerpts from the Knowledge Base are provided in Appendix 9.2.

Clinically Grounded Corpus. We extract 15,000 clinical cases from a medical center, independently sourced and strictly disjoint from the MentalDx Bench. These cases are processed through a three-stage refinement pipeline, executed by DeepSeek-R1 (Guo et al., 2025), which performs labeling and rewriting under the guidance of \mathcal{K} .

Step 1. Given a case x , we first retrieve all category definitions $\{d_i\}_{i=1}^{16}$ from \mathcal{K} and prompt the LLM to select the most compatible category:

$$c = \text{LLM}(x, \{d_i\}) \quad \text{where } d_i \in \mathcal{K}_{\text{category}}. \quad (2)$$

Step 2. Conditioned on x and the predicted category c , we then extract all candidate disorders and their diagnostic criteria from $\mathcal{K}(c)$, and prompt the LLM to identify the most plausible disorder:

$$d = \text{LLM}(x, \mathcal{K}(c)). \quad (3)$$

Step 3. To improve alignment with (c, d) , we prompt the LLM to rewrite the original case x by referencing the diagnostic definitions and criteria of c and d , ensuring clinical consistency as:

$$x' = \text{LLM}_{\text{rewrite}}(x; \mathcal{K}(c, d)). \quad (4)$$

The entire automatic procedure is executed on-site within the medical center using an on-premise deployment of DeepSeek-R1, ensuring that all data remains local and under institutional governance. Prompt details for Step 1, Step 2, and Step 3 are provided in Appendix 9.7.

6.2 Hypothetico-Deductive Reasoning Trajectory Building

Hypothetico-deductive Reasoning. A foundational model of clinical decision-making, it characterizes how clinicians systematically generate, evaluate, and iteratively refine diagnostic hypotheses based on observed evidence (Jefford et al., 2011; Langridge et al., 2016; Collins, 2017). Inspired by this paradigm, we formulate psychiatric diagnosis as a structured three-step reasoning process: *symptom categorization*, *hypothesis generation*, and *differential diagnosis*. Together, these steps trace a closed diagnostic trajectory—identifying dysfunctions, proposing candidate disorders, and progressively refining hypotheses toward a coherent and well-supported final diagnosis.

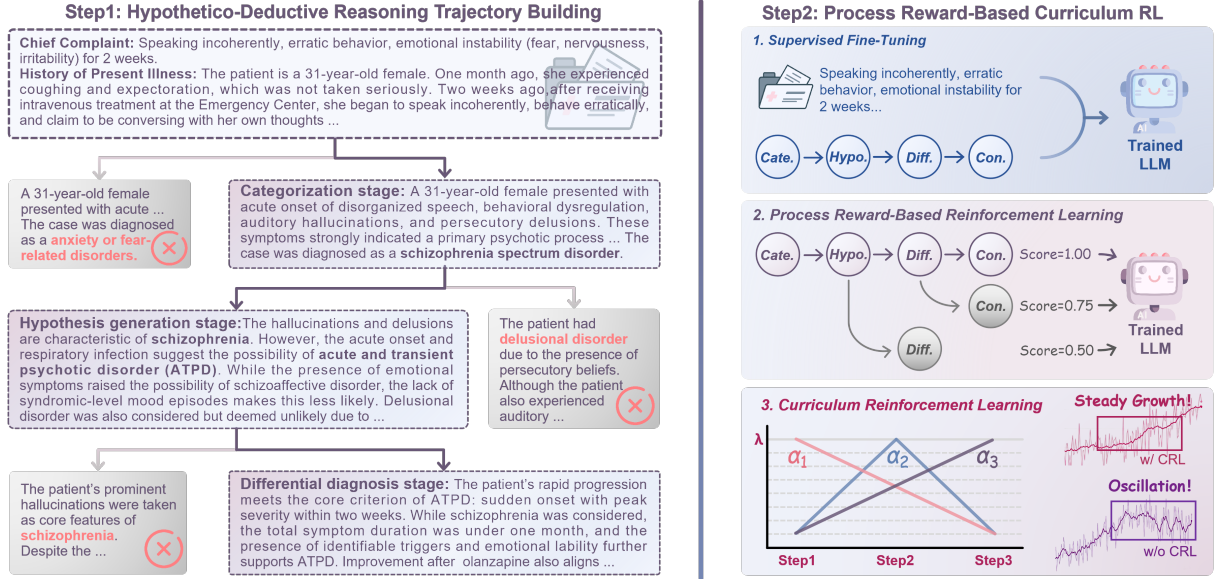


Figure 5: Training pipeline of MENTALSEEK-DX to address paradigm misalignment by modeling clinically grounded diagnostic reasoning, consisting of two stages: 1) Hypothetico-Deductive Reasoning Trajectory Building using \mathcal{K} as SFT corpus; 2) Reward-Based Curriculum Reinforcement Learning to guide the diagnostic process.

Reasoning Trajectory Building. Given an input case x , we construct a reasoning trajectory (s, h, d) (i.e., symptom, hypothesis, and differential) by decomposing the diagnostic process into the three stages above. At each stage, the model generates the corresponding reasoning component conditioned on the input case, previous outputs, and structured knowledge retrieved from \mathcal{K} . Specifically, in the categorization stage, the model identifies relevant diagnostic categories c by matching features in x with characteristic symptoms in $\mathcal{K}(c)$, yielding s . In the hypothesis generation stage, the model produces a candidate set $h = \{h_1, h_2, \dots\}$ based on (x, s) and the disorder-specific features under category c . In the differential diagnosis stage, the model compares each $d \in h$ using their full feature sets from $\mathcal{K}(d)$, and selects the most plausible diagnosis. We retain only trajectories that successfully complete all trajectory as SFT corpus. For detailed procedures, see Appendix 9.8.

6.3 Process Reward-Based Curriculum Reinforcement Learning

Reinforcement Learning. After supervised training on HDR, we further optimize MENTALSEEK-DX using *Process Reward-Based Curriculum Reinforcement Learning (PRCRL)*, which builds upon Group Relative Policy Optimization (GRPO) (Guo et al., 2025). For each case x , we sample candidate trajectories y_1, \dots, y_G from the current policy $\pi_{\theta_{\text{old}}}$

and assign scalar rewards R_i :

$$\hat{A}_i = \frac{R_i - \mu_R}{\sigma_R + \epsilon}, \quad (5)$$

where μ_R and σ_R denote the mean and standard deviation of rewards within the group. The policy is then updated using a PPO-style clipped objective:

$$\mathcal{J}(\theta) = \mathbb{E}_{x, y_i} \left[\min(r_i \hat{A}_i, \text{clip}(r_i, 1 - \epsilon, 1 + \epsilon) \hat{A}_i) \right]. \quad (6)$$

where r_i is the importance sampling ratio.

Reward Composition. The effectiveness of PRCRL hinges on reward design. Instead of a single outcome-level reward, we define stage-aware, process-level rewards that reflect structured clinical reasoning. Each reward R_i is computed as a weighted combination of multiple components:

$$R_{\text{diagnosis}} = \alpha_1 R_{\text{cat.}} + \alpha_2 R_{\text{hypo.}} + \alpha_3 R_{\text{diff.}}, \quad (7)$$

where $R_{\text{cat.}}$, $R_{\text{hypo.}}$, and $R_{\text{diff.}}$ correspond to category identification, hypothesis generation, and differential diagnosis, respectively, and α_k are fixed coefficients controlling their relative contribution.

1) *Symptom categorization reward.* At the first stage, MENTALSEEK-DX is required to identify the correct high-level diagnostic category based on presenting symptoms. We formulate this stage as a coarse-grained classification task and define:

$$R_{\text{cat.}} = \mathbb{I}[C_{\text{pred}} = C_{\text{gt}}], \quad (8)$$

where C_{pred} and C_{gt} denote the predicted and ground-truth diagnostic categories, and $\mathbb{I}[\cdot]$ is the

Model	CA	DA	JA
Qwen3-32B	58.71	31.04	26.26
QWQ-32B	76.26	47.89	46.63
Baichuan-M2	64.61	44.38	41.71
GPT-OSS-120B	75.84	52.53	50.00
DeepSeek-R1	78.37	55.20	54.49
Claude-Sonnet-4.5	80.76	59.69	58.15
MENTALSEEK-DX-7B	82.58	67.56	67.13
MENTALSEEK-DX-14B	83.99	70.08	69.38

Table 6: Comparison of MENTALSEEK-DX with other models, achieving state-of-the-art on all metrics.

indicator function. This reward enforces category fidelity and limits diagnostic error propagation.

2) *Hypothesis generation reward.* The second stage evaluates the quality of hypothesis construction, focusing on whether the correct disorder is included and appropriately prioritized among candidate diagnoses. Let $\mathcal{H}_{\text{pred}}$ denote the predicted hypothesis list and i be the rank position of the ground-truth disorder d_{gt} if present. We define:

$$R_{\text{hypo.}} = \mathbb{I}[i \leq 4] \cdot \left(1 - \frac{i-1}{4}\right). \quad (9)$$

This formulation rewards both coverage and ranking quality, encouraging concise yet focused hypotheses and penalizing overinclusive reasoning.

3) *Differential diagnosis reward.* At the final stage, the model is expected to converge to a single, clinically correct diagnosis. We define:

$$R_{\text{diff.}} = \mathbb{I}[d_{\text{pred}} = d_{\text{gt}}], \quad (10)$$

where d_{pred} and d_{gt} denote the predicted and ground-truth diagnoses. This term directly measures diagnostic correctness and anchors the entire reasoning process toward clinically valid outcomes.

Curriculum Learning. To guide the model through the progressive structure of clinical reasoning, inspired by curriculum learning (Wang et al., 2021; Liu et al., 2021), we adopt a stage-aware reward scheduling strategy. Specifically, we divide training into five epochs and adjust the stage-wise reward weights $(\alpha_1, \alpha_2, \alpha_3)$ over time. Early epochs prioritize category-level accuracy by setting higher α_1 , while later epochs gradually increase α_2 or α_3 to emphasize hypothesis generation or differential diagnosis. This training schedule helps the model internalize diagnostic reasoning in a structured manner, supporting stable optimization and improved fine-grained diagnostic performance. For a detailed overview of the training flow, see Fig 5.

	CA	DA	JA
<i>Baseline LLMs</i>			
Qwen2.5-7B-Instruct	47.47	14.75	10.96
Qwen2.5-14B-Instruct	61.66	29.92	24.72
<i>Hypothetico-Deductive Reasoning (HDR)</i>			
7B-SFT w/o HDR	68.14 (+20.67)	42.56 (+27.81)	40.14 (+29.18)
7B-SFT w/ HDR	82.02 (+34.55)	62.78 (+48.03)	62.50 (+51.54)
14B-SFT w/o HDR	69.21 (+7.55)	47.44 (+17.52)	46.88 (+22.16)
14B-SFT w/ HDR	83.71 (+22.05)	64.47 (+34.55)	64.47 (+39.75)
<i>Reinforcement Learning (RL)</i>			
7B-HDR w/o RL	82.02 (+34.55)	62.78 (+48.03)	62.50 (+51.54)
7B-HDR w/ RL	82.16 (+34.69)	66.01 (+51.26)	65.87 (+54.91)
7B-HDR w/ CRL	82.58 (+35.11)	67.56 (+52.81)	67.13 (+56.17)
14B-HDR w/o RL	81.59 (+19.93)	64.47 (+34.55)	64.47 (+39.75)
14B-HDR w/ RL	83.71 (+22.05)	67.28 (+37.36)	67.28 (+42.56)
14B-HDR w/ CRL	83.99 (+22.33)	70.08 (+40.16)	69.38 (+44.66)

Table 7: The ablation results on MentalDx Bench.




Figure 6: Expert evaluation of diagnostic reasoning, showing superior performance of MENTALSEEK-DX.

7 MENTALSEEK-DX Evaluation

We train MENTALSEEK-DX based on Qwen2.5-7B and 14B (Team et al., 2024); training details and hyperparameter settings can see in Appendix 9.10.

Evaluation Results. We compare MentalSeek-Dx with other LLMs on MentalDx Bench. As shown in Table 6, our model achieves state-of-the-art performance on all metrics among mainstream models, despite having only 7B and 14B parameters. Detailed results are provided in Appendix 9.4. Complementary to these quantitative results, we conduct a blinded evaluation (Figure 6) with 12 medical experts using a Likert scale to assess the reasoning process along four dimensions: Clinical Logical Rigor, Evidence Strength, Critical Completeness, and Differential Diagnosis Adequacy. Results indicate that MentalSeek-Dx, guided by progressive hypothetico-deductive reasoning, not only improves diagnostic accuracy but also demonstrates structured clinical reasoning toward transparent, accurate, and clinically trustworthy diagnosis. Annotation guidelines can see in Appendix 9.9.

Ablation Study. We evaluate the contribution



INPUT	
Basic Information: Female patient, 31 years old.	
Chief Complaint: The patient experienced coughing and sputum production one month ago, which was not taken seriously. Two weeks ago, after receiving intravenous treatment at "City Emergency Center, she began to exhibit nonsensical speech, disorganized behavior, and reported that she was conversing with her own thoughts. She also felt nervous, fearful, and experienced episodes of agitation...	
MentalSeek-Dx	
First, I need to classify the patient's syndrome based on the combination of symptoms and disease course. Fellow clinical definitions, ... which leads to the classification under the Schizophrenia and Primary Psychotic Disorders...	
Finally, I need to apply diagnostic criteria and key differential points to complete the differential diagnosis... Her symptoms progressed ... within just two weeks, meeting the core criterion of Acute and Transient Psychotic Disorder for "acute onset of symptoms with maximum severity reached within two weeks." Regarding the disease course, from the onset of the first psychiatric and behavioral abnormalities to hospitalization, it was just over two weeks, far short of the one-month duration required by Schizophrenia ...	
Claude-Sonnet-4.5	
Key Information Summary: Female, 31 years old, acute onset....	
Diagnostic Analysis:	
1)Schizophrenia (Most Likely Diagnosis)	
Supporting Factors: Persistent positive symptoms (hallucinations, delusions); Good response to antipsychotic medication, although...	
2)Acute and Transient Psychotic Disorder (Less Likely)	
Against This Diagnosis: Symptoms are persistent and there is a tendency ...	
Final Diagnosis: Persistentmost consistent with Schizophrenia, classified under: Schizophrenia and Primary Psychotic Disorders.	

Figure 7: Comparison between Claude-Sonnet-4.5 and MENTALSEEK-DX, showing a paradigm misalignment.

of each training stage in MENTALSEEK-DX on Qwen2.5-7B and 14B (Table 7). All components consistently improve diagnostic performance. Incorporating hypothetico-deductive reasoning yields substantial gains over direct training, increasing *CA*, *DA*, and *JA* by 13.88%, 20.22%, and 22.36% on Qwen2.5-14B, respectively, underscoring the importance of clinical reasoning in resolving the paradigm misalignment between pattern-based associations and diagnostic reasoning. Beyond SFT, reinforcement learning further improves diagnostic precision. On Qwen2.5-7B, RL adds 3.23% in *DA* and 3.37% in *JA*, indicating that RL helps refine model decisions in challenging cases. Notably, while CRL brings modest gains, it effectively stabilizes the RL training process (see Appendix 9.11).

Case Study. In Fig. 7, we present a case comparison. Claude-Sonnet-4.5 shows paradigm misalignment by mapping symptoms directly to a disorder and then re-inferring its category. This reverse process reflects unstructured knowledge and leads to diagnostic error. In contrast, MentalSeek-Dx follows a clinically grounded path: it first constrains category-level reasoning, then formulates disorder hypotheses, and finally applies diagnostic criteria. More details are provided in Appendix 9.12.

8 Conclusion

This work bridges the critical gap between computational psychiatry and clinical reality. By introducing **MentalDx Bench**, we established the first ecologically valid framework for evaluating

fine-grained disorder diagnosis under strict ICD-11 standards. Our investigation uncovers a fundamental *paradigm misalignment*: while general-purpose LLMs rely on associative shortcuts, they fail to replicate the rigorous *hypothetico-deductive reasoning* essential for clinical practice. We address this by proposing **MentalSeek-Dx**, a specialized framework that aligns model cognition with psychiatrist-level logic through trajectory-based supervision and curriculum reinforcement learning. Achieving state-of-the-art results with only 14B parameters, our work demonstrates that *cognitive alignment*, rather than mere parameter scale, is the decisive factor for reliable automated diagnosis.

Limitations

Despite the promising results, this study has limitations inherent to the complexity of computational psychiatry. First, while **MentalDx Bench** offers high ecological validity, the dataset primarily comprises textual Electronic Health Records. In actual clinical practice, diagnosis is multimodal, incorporating patient demeanor, speech prosody, and non-verbal cues. Our current focus on textual modalities serves as a controlled foundational step, but it may not fully capture the complete spectrum of diagnostic signals available in face-to-face consultations. Second, although **MentalSeek-Dx** demonstrates advanced reasoning capabilities, it is not immune to the broader challenges of interpretability and safety in Large Language Models. In high-stakes medical settings, our model is designed strictly as a clinical decision support tool. It cannot, and should not, replace the final judgment of qualified human psychiatrists. The human-in-the-loop paradigm remains essential to mitigate potential risks associated with algorithmic hallucinations or biases in long-tail rare disorders.

Ethics Statement

This study was conducted in strict adherence to ethical guidelines, ensuring the protection of participants' rights and the integrity of the research process. The study protocol, including all data collection, labeling, and annotation procedures, was thoroughly reviewed and approved by the institutional ethics boards of all participating clinical centers. This approval process ensured that all aspects of the study met the highest standards for legal and ethical conduct in clinical research.

Data collection and annotation activities fol-

lowed the ethical principles outlined in the *Declaration of Helsinki*, which emphasizes respect for individuals, the protection of privacy, and the commitment to research integrity. In particular, all processes were designed to minimize any potential risks to participants and to uphold their dignity throughout the study.

In order to protect participant confidentiality and privacy, all records used in the study were fully de-identified prior to any processing or analysis. This de-identification process ensured that no personally identifiable information was retained, further safeguarding participants' anonymity. As a result of this de-identification, the requirement for individual informed consent was waived by the institutional ethics boards. This decision was based on the fact that the study utilized fully anonymized data, with no identifiable link to individual participants, and posed no risk to personal privacy.

Additionally, the ethical standards for handling and analyzing sensitive data were rigorously followed throughout the study. The research was conducted solely for scientific purposes, with all data being used in a manner that was consistent with the ethical guidelines governing clinical research.

For transparency, we note that the dataset used in this study was reviewed and approved by the **relevant institutional ethics review committees under formal protocols**. Due to anonymization requirements, detailed ethics approval identifiers and data governance specifics are not disclosed at this stage. These details will be updated in the final published version of the paper where permitted.

Acknowledgements

We gratefully acknowledge the support of the National Key R&D Program of China (No. 2024YFC2707805), the National Natural Science Foundation of China (No. 62506050, No. 62176029), the China Postdoctoral Science Foundation Funded Project (No. 2024M763867), the Chongqing Science and Technology Innovation and Application Development projects (No. CSTB2023TIAD-KPX0064), and the Open Competition Program of Chongqing Municipal Commission of Economy and Information Technology (No. YJX-202500100200X).

The experimental and computational work in this research runs on the Huawei Cloud AI Compute Service. We appreciate the stable compute supply from this platform. We sincerely thank the anonymous

reviewers for their insightful comments and constructive suggestions.

References

- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, and 1 others. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Ruth Appel, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon, and Michael Stern. 2025. Anthropic economic index report: Uneven geographic and enterprise ai adoption. *arXiv preprint arXiv:2511.15080*.
- Eliseo Bao, Anxo Pérez, and Javier Parapar. 2025. Redsm5: A reddit dataset for dsm-5 depression detection. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 6323–6327.
- Sébastien Bubeck, Christian Coester, Ronen Eldan, Timothy Gowers, Yin Tat Lee, Alexandru Lupsasca, Mehtaab Sawhney, Robert Scherrer, Mark Sellke, Brian K Spears, and 1 others. 2025. Early science acceleration experiments with gpt-5. *arXiv preprint arXiv:2511.16072*.
- Grace Byun, Rebecca Lipschutz, Sean T Minton, Abigail Lott, and Jinho D Choi. 2025. Cradle bench: A clinician-annotated benchmark for multi-faceted mental health crisis and safety risk detection. *arXiv preprint arXiv:2510.23845*.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. 2024. [Huatuoqpt-o1](#), towards medical complex reasoning with llms. *Preprint*, arXiv:2412.18925.
- Jenny Chim, Adam Tsakalidis, Dimitris Gkoumas, Dana Atzil-Slonim, Yaakov Ophir, Ayah Zirikly, Philip Resnik, and Maria Liakata. 2024. Overview of the clpsych 2024 shared task: Leveraging large language models to identify evidence of suicidality risk in online posts. In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 177–190.
- Arman Cohan, Bart Desmet, Andrew Yates, Luca Soldaini, Sean MacAvaney, and Nazli Goharian. 2018. Smhd: a large-scale resource for exploring online language usage for multiple mental health conditions. *arXiv preprint arXiv:1806.05258*.
- Mark Collins. 2017. *Clinical reasoning in image guided radiotherapy: a multimethod study*. Sheffield Hallam University (United Kingdom).
- Chongyuan Dai, Jinpeng Hu, Hongchang Shi, Zhuo Li, Xun Yang, and Meng Wang. 2025. Psyche-r1: Towards reliable psychological llms through unified

- empathy, expertise, and reasoning. *arXiv preprint arXiv:2508.10848*.
- Chengfeng Dou, Chong Liu, Fan Yang, Fei Li, Jiyuan Jia, Mingyang Chen, Qiang Ju, Shuai Wang, Shunya Dang, Tianpeng Li, and 1 others. 2025. Baichuanm2: Scaling medical capability with large verifier system. *arXiv preprint arXiv:2509.02208*.
- Aya E Fouda, Abdelrahman A Hassan, Radwa J Hanafy, and Mohammed E Fouda. 2025. Psychiatrybench: A multi-task benchmark for llms in psychiatry. *arXiv preprint arXiv:2509.09711*.
- Clara Frydman-Gani, Alejandro Arias, Maria Perez Vallejo, John Daniel Londoño Martínez, Johanna Valencia-Echeverry, Mauricio Castaño, Alex AT Bui, Nelson B Freimer, Carlos Lopez-Jaramillo, and Loes M Olde Loohuis. 2025. Large language models for psychiatric phenotype extraction from electronic health records. *medRxiv*.
- Changzeng Fu, Zelin Fu, Qi Zhang, Xinhe Kuang, Jiacheng Dong, Kaifeng Su, Yikai Su, Wenbo Shi, Junfeng Yao, Yuliang Zhao, and 1 others. 2025. The first mpdd challenge: Multimodal personality-aware depression detection. *arXiv preprint arXiv:2505.10034*.
- Zhuohan Ge, Nicole Hu, Darian Li, Yubo Wang, Shihao Qi, Yuming Xu, Han Shi, and Jason Zhang. 2025. A survey of large language models in mental health disorder detection on social media. In *2025 IEEE 41st International Conference on Data Engineering Workshops (ICDEW)*, pages 164–176. IEEE.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Amey Hengle, Atharva Kulkarni, Shantanu Deepak Patankar, Madhumitha Chandrasekaran, Sneha D’silva, Jemima S Jacob, and Rashmi Gupta. 2024. Still not quite there! evaluating large language models for comorbid mental health diagnosis. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16698–16721.
- Jinpeng Hu, Tengpeng Dong, Gang Luo, Hui Ma, Peng Zou, Xiao Sun, Dan Guo, Xun Yang, and Meng Wang. 2024. Pyscollm: Enhancing llm for psychological understanding and evaluation. *IEEE Transactions on Computational Social Systems*, 12(2):539–551.
- Mariia Ignashina, Paulina Bondaronek, Dan Santel, John Pestian, and Julia Ive. 2025. Llm assistance for pediatric depression. *arXiv preprint arXiv:2501.17510*.
- Elaine Jefford, Kathleen Fahy, and Deborah Sundin. 2011. Decision-making theories and their usefulness to the midwifery profession both in terms of midwifery practice and the education of midwives. *International Journal of Nursing Practice*, 17(3):246–253.
- Shreya Johri, Jaehwan Jeong, Benjamin A Tran, Daniel I Schlessinger, Shannon Wongvibulsin, Leandra A Barnes, Hong-Yu Zhou, Zhuo Ran Cai, Eliezer M Van Allen, David Kim, and 1 others. 2025. An evaluation framework for clinical use of large language models in patient interaction tasks. *Nature medicine*, 31(1):77–86.
- Laerdon Kim. 2025. A baseline for self-state identification and classification in mental health data: Clpsych 2025 task. *arXiv preprint arXiv:2504.14066*.
- Wuyang Lan, Wenzheng Wang, Changwei Ji, Guoxing Yang, Yongbo Zhang, Xiaohong Liu, Song Wu, and Guangyu Wang. 2025a. Clinicalgpt-r1: Pushing reasoning capability of generalist disease diagnosis with large language model. *arXiv preprint arXiv:2504.09421*.
- Xiaochong Lan, Zhiguang Han, Yiming Cheng, Li Sheng, Jie Feng, Chen Gao, and Yong Li. 2025b. Depression detection on social media with large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 2155–2171.
- Neil Langridge, Lisa Roberts, and Catherine Pope. 2016. The role of clinician emotion in clinical reasoning: Balancing the analytical process. *Manual therapy*, 21:277–281.
- Inbar Levkovich, Eyal Rabin, Michal Brann, and Zohar Elyoseph. 2024. Large language models outperform general practitioners in identifying complex cases of childhood anxiety. *Digital Health*, 10:20552076241294182.
- Silvia Kyungjin Lho, Sang-Cheol Park, Hahyun Lee, Da Young Oh, Hyeonjin Kim, Soomin Jang, Hee Yeon Jung, So Young Yoo, Su Mi Park, and Jun-Young Lee. 2025. Large language models and text embeddings for detecting depression and suicide in patient narratives. *JAMA Network Open*, 8(5):e2511922–e2511922.
- Fenglin Liu, Shen Ge, and Xian Wu. 2021. Competence-based multimodal curriculum learning for medical report generation. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*, pages 3001–3012.
- Shuyu Liu, Ruoxi Wang, Ling Zhang, Xuequan Zhu, Rui Yang, Xinzhu Zhou, Fei Wu, Zhi Yang, Cheng Jin, and Gang Wang. 2025. Psychbench: A comprehensive and professional benchmark for evaluating the performance of llm-assisted psychiatric clinical practice. *arXiv preprint arXiv:2503.01903*.
- Siddharth Mangalik, Johannes C Eichstaedt, Salvatore Giorgi, Jihu Mun, Farhan Ahmed, Gilvir Gill, Adithya V. Ganesan, Shashanka Subrahmanya, Nikita Soni, Sean AP Clouston, and 1 others. 2024. Robust language-based mental health assessments in time and space through social media. *NPJ Digital Medicine*, 7(1):109.

- Saad Mankarious, Ayah Ziriky, Daniel Wiechmann, Elma Kerz, Edward Kempa, and Yu Qiao. 2025. Mindset: Advancing mental health benchmarking through large-scale social media data. *arXiv preprint arXiv:2511.20672*.
- Nick Obradovich, Sahib S Khalsa, Waqas U Khan, Jina Suh, Roy H Perlis, Olusola Ajilore, and Martin P Paulus. 2024. Opportunities and risks of large language models in psychiatry. *NPP—Digital Psychiatry and Neuroscience*, 2(1):8.
- Jaedong Oh, Jooyoung Lim, and Hayoung Oh. 2026. Clinically validated depression dataset aligned with dsm-5 criteria for major depressive disorder (mdd). *Expert Systems with Applications*, 296.
- Graziella Orrù, Giulia Melis, and Giuseppe Sartori. 2025. Large language models and psychiatry. *International Journal of Law and Psychiatry*, 101:102086.
- David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R So, Maud Texier, and Jeff Dean. 2022. The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7):18–28.
- Ahmad Radwan, Mohannad Amarneh, Hussam Alawneh, Huthaifa I Ashqar, Anas AlSobeh, and Aws Abed Al Raheem Magableh. 2024. Predictive analytics in mental health leveraging llm embeddings and machine learning models for social media analysis. *International Journal of Web Services Research (IJWSR)*, 21(1):1–22.
- Mavin Sao and Hoi-Jeong Lim. 2024. Miroberta: Mental illness text classification with transfer learning on subreddits. *IEEE Access*.
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cian Hughes, Charles Lau, and 1 others. 2025. Medgemma technical report. *arXiv preprint arXiv:2507.05201*.
- Ryan Allen Shewcraft, John Schwarz, and Mariann Micsinai Balan. 2025. Algorithmic classification of psychiatric disorder-related spontaneous communication using large language model embeddings: Algorithm development and validation. *JMIR AI*, 4:e67369.
- Daun Shin, Hyoseung Kim, Seunghwan Lee, Younhee Cho, and Whanbo Jung. 2024. Using large language models to detect depression from user-generated diary text data as a novel approach in digital mental health screening: instrument validation study. *Journal of Medical Internet Research*, 26:e54617.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, and 1 others. 2025. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*.
- Qwen Team and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).
- Xin Wang, Yudong Chen, and Wenwu Zhu. 2021. A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):4555–4576.
- Yi-Fan Wang, Ming-Da Li, Su-Hong Wang, Yin Fang, Jie Sun, Lin Lu, and Wei Yan. 2025. Large language models in clinical psychiatry: Applications and optimization strategies. *World Journal of Psychiatry*, 15(11):108199.
- Samantha Weber, Nicolas Deperrois, Robert Heun, Laura Frühschütz, Anna Monn, Stephanie Homan, Andrea Häfliger, Erich Seifritz, Tobias Kowatsch, MULTICAST consortium Jäger Lena 8 Schulte-brauks Katharina 9 Gershov Sapir 9 Mocellin Jacopo 1 4, and 1 others. 2025. Using a fine-tuned large language model for symptom-based depression evaluation. *npj Digital Medicine*, 8(1):598.
- Shurui Wu, Xinyi Huang, and Dingxin Lu. 2025. Psychological health knowledge-enhanced llm-based social network crisis intervention text transfer recognition method. In *Proceedings of the 2025 International Conference on Health Big Data*, pages 156–161.
- Xuhai Xu, Bingsheng Yao, Yuanzhe Dong, Saadia Gabriel, Hong Yu, James Hendler, Marzyeh Ghassemi, Anind K Dey, and Dakuo Wang. 2024. Mental-llm: Leveraging large language models for mental health prediction via online text data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1):1–32.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024. Mental-lama: interpretable mental health analysis on social media with large language models. In *Proceedings of the ACM Web Conference 2024*, pages 4489–4500.
- 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. *arXiv preprint arXiv:2205.11764*.
- Shouwen Zheng, Yingzhi Tao, and Taiqi Zhou. 2025. Rsd-15k: A large-scale user-level annotated dataset for suicide risk detection on social media. In *2025 IEEE 41st International Conference on Data Engineering Workshops (ICDEW)*, pages 190–196. IEEE.
- Mingcheng Zhu, Zhiyao Luo, Yu Liu, and Tingting Zhu. 2025. Medtpe: Compressing long ehr sequence for llm-based clinical prediction with token-pair encoding. In *Proceedings of the 11th Mining and Learning from Time Series Workshop@ KDD*, volume 2025.

Muni Zhuang, Dongsheng Cheng, Xin Lu, and Xu Tan. 2024. Postgraduate psychological stress detection from social media using bert-fused model. *PLoS one*, 19(10):e0312264.

9 Appendix

9.1 Category and Disorder List

In this section, we provide the full forms and brief explanations of the category abbreviations used in the main text. Fig 8 presents the names and corresponding codes for the 76 specific disorders included in the study.

The full form of **ANX** is Anxiety or fear-related disorders. Anxiety or fear-related disorders are a group of mental health conditions characterized by excessive fear and anxiety, along with associated behavioral disturbances. These symptoms are severe enough to cause significant distress or impairment in an individual's personal, social, educational, occupational, or other important areas of functioning. According to the ICD-11, specific types of disorders under this category include generalized anxiety disorder, panic disorder, agoraphobia, specific phobias, social anxiety disorder, and separation anxiety disorder. Patients may exhibit specific cognitive features, such as excessive worry about particular situations or objects, which help distinguish between different types of anxiety or fear-related disorders.

The full form of **CATA** is Catatonia. Catatonia is a clinical syndrome characterized by psychomotor disturbances, which manifest as a combination of reduced, increased, or abnormal psychomotor activity. Typical symptoms include stupor (maintaining a fixed posture for extended periods), rigidity (abnormal increase in muscle tension), waxy flexibility (limbs can be manipulated and maintained in a posture), mutism, negativism (active or passive resistance to instructions), echolalia (repetition of others' speech), or echopraxia (repetition of others' actions). This syndrome may arise secondary to psychiatric disorders such as schizophrenia, bipolar disorder, and depressive disorders, and can also be triggered by drug reactions, neurological diseases, or physical illnesses.

The full form of **SUD** is Disorders due to substance use or addictive behaviors. These disorders refer to functional impairment or distress caused by the repeated use of psychoactive substances (including drugs) or specific repetitive rewarding behaviors, manifesting in cognitive, behavioral, and physiological symptoms. The core features include

a persistent craving for the substance or behavior, impaired control, prioritizing use despite harm, and neuroadaptive changes. Key characteristics also include difficulty controlling the frequency, duration, and intensity of use, withdrawal symptoms after reducing or stopping use (such as tremors, anxiety, insomnia), tolerance (requiring increased doses or intensity for the same effect), functional impairment in personal, family, occupational, or social roles, and continued use despite knowing the harm to physical and mental health.

The full form of **BOD** is Disorders of bodily distress or bodily experience. These disorders are mental health conditions primarily characterized by bodily symptoms. Patients often show excessive worry about their health, accompanied by persistent physical symptoms such as pain, fatigue, and digestive issues, among others. These symptoms typically cannot be fully explained by routine medical examinations.

The full form of **STRESS** is Disorders specifically associated with stress. These disorders are a group of mental health conditions triggered by clearly identifiable stressors or traumatic events. Patients typically present with acute stress disorder, which occurs within days of exposure to extreme stressors and is characterized by intense fear, helplessness, and flashbacks. Post-traumatic stress disorder (PTSD) lasts longer and is marked by the recurrent re-experiencing of traumatic events, avoidance behaviors, and heightened arousal. Adjustment disorder is another form, where emotional or behavioral problems arise from identifiable life events that exceed normal coping mechanisms, typically occurring within one month of the event and affecting daily functioning.

The full form of **DISR** is Disruptive behavior or dissociative disorders. Impulse control disorders are a group of mental health conditions characterized by the difficulty of patients to resist strong, inappropriate desires or impulses, leading to the performance of certain specific behaviors. These behaviors often cause long-term harm to the individual or others, or result in significant impairment in important functional areas such as social, family, or occupational domains. Although patients typically recognize the negative impact of these behaviors, they are unable to resist the impulse in the short term. Common ICDs include pyromania, kleptomania, intermittent explosive disorder, and compulsive sexual behavior disorder.

The full form of **DISS** is Dissociative disorders.

Dissociative disorders are complex mental health conditions primarily characterized by the involuntary partial or complete disintegration of psychological functions such as identity, memory, consciousness, perception, or behavior. These symptoms are not directly caused by drugs, substance abuse, or other psychiatric, behavioral, or neurodevelopmental disorders, and must be distinguished from normal manifestations in cultural or religious practices. The main types include, but are not limited to, dissociative amnesia (selective memory loss), depersonalization-derealization disorder (feeling detached from reality), dissociative fugue (manifesting as a confused state), and dissociative identity disorder (formerly known as multiple personality disorder).

The full form of **ELIM** is Elimination disorders. Elimination disorders refer to the persistent or recurrent inability to control urination or defecation after reaching the normal developmental age. The main manifestations include enuresis, which is the involuntary urination during the night or day, typically in bed or clothing, and encopresis, which is the inappropriate defecation in places such as clothing. This diagnosis is applicable when children fail to achieve typical bladder and bowel control by the expected developmental age (enuresis by age 5 and encopresis by age 4). If the incontinence symptoms can be fully attributed to other health conditions, congenital or acquired abnormalities in the urinary or gastrointestinal systems, or medication side effects, it should not be diagnosed as an elimination disorder.

The full form of **EAT** is Feeding or eating disorders. Feeding or eating disorders are a group of mental health conditions characterized by abnormal eating behaviors that cannot be explained by other health conditions and do not conform to normal developmental patterns in the cultural context. These include, but are not limited to, rumination/regurgitation, avoidant/restrictive food intake, anorexia nervosa, bulimia, binge eating, and pica. These disorders often involve issues with controlling food intake, concerns about weight and body shape, or inappropriate dietary choices.

The full form of **PREG** is Mental or behavioral disorders associated with pregnancy, childbirth, or the puerperium. These disorders refer to mental health issues that occur during pregnancy, during childbirth, and within the first six weeks postpartum. Such disorders include, but are not limited to, depressive disorders, anxiety disorders, and psy-

chotic disorders. Based on their severity and specific manifestations, they can be further categorized into non-psychotic disorders (e.g., postpartum depression) and psychotic disorders (e.g., postpartum psychosis). These mental health conditions significantly affect a woman's psychological state and social functioning, requiring professional medical intervention.

The full form of **MOOD** is Mood disorders. Mood disorders are a group of mental health conditions characterized by significant and persistent changes in mood or emotional states. These disorders include two main types: bipolar disorder and depressive disorder. They are typically manifested by several major types of mood episodes: depressive episodes, which involve persistent low mood, loss of interest, and reduced energy; manic episodes, which involve abnormally elevated mood, increased activity, and reduced need for sleep; mixed episodes, where both depressive and manic symptoms are present simultaneously; and hypomanic episodes, which are similar to manic episodes but less severe, with minimal impairment of social functioning. These mood episodes are not independent diagnostic entities but are the primary components of mood disorders.

The full form of **NCOG** is Neurocognitive disorders. Neurocognitive disorders refer to major clinical impairments in cognitive function that are acquired, rather than developmental or congenital. Specifically, these disorders do not include cognitive issues that are present from birth, nor do they cover problems that typically emerge during the developmental stage (which are classified as neurodevelopmental disorders). Instead, neurocognitive disorders refer to a decline in cognitive function from a previously attained level.

The full form of **NDEV** is Neurodevelopmental disorders. Neurodevelopmental disorders are a group of mental or behavioral disorders that appear during an individual's developmental stage, typically in early childhood, especially before school age, and are characterized by developmental deficits that cause impairment in social, academic, or occupational functioning. The core features of these disorders include significant difficulties in acquiring and performing specific intellectual, motor, language, or social functions. While many mental and behavioral disorders can emerge during the developmental period, only those disorders that are primarily characterized by neurodevelopmental impairments are classified under this group.

The full form of **OCD** is Obsessive-compulsive or related disorders. Obsessive-compulsive or related disorders are a group of mental health conditions characterized by repetitive thoughts and behaviors. The main symptoms include: obsessive thoughts, which are unwanted, recurring ideas, doubts, or impulses, and compulsive behaviors, which are repetitive actions or mental rituals. These symptoms persist despite the patient's awareness of their meaninglessness, leading to significant personal distress and impairment in social functioning.

The full form of **PERS** is Personality disorders and related traits. Personality disorders refer to a persistent and stable pattern of emotional experiences, cognitive patterns, and behaviors that significantly deviate from the cultural expectations of the individual's background, leading to widespread functional impairment or subjective distress. This pattern typically emerges during adolescence or early adulthood, with its stability confirmed through observation in adulthood. Personality disorders and related personality traits cover two main aspects: prominent personality features or patterns, where individuals exhibit significant deviations in personality traits, but the severity or extent may not fully meet the diagnostic criteria for a personality disorder, and personality disorders, which include various subtypes (such as borderline, antisocial, schizoid, etc.), characterized by persistent functional impairment, and must exclude direct consequences of other psychiatric disorders or substance use.

The full form of **SCHIZ** is Schizophrenia or other primary psychotic disorders. Schizophrenia and other primary psychotic disorders are a group of mental health conditions characterized by significant impairment in reality testing. The core symptoms include: positive symptoms such as delusions, hallucinations, and disorganized speech; negative symptoms such as flat affect, reduced volition, and social withdrawal; and cognitive dysfunctions such as impairments in attention, memory, and executive function. These symptoms cause the individual's thoughts, emotions, and behaviors to deviate markedly from the cultural norms, and cannot be attributed to other psychiatric disorders, such as bipolar disorder or substance use disorders.

9.2 Knowledge Base

In our work, we have developed a comprehensive and structured psychiatric knowledge base, denoted as \mathcal{K} , to facilitate diagnostic reasoning and improve

model training. This knowledge base was designed through extensive collaboration with expert psychiatrists, ensuring it is grounded in clinical relevance and real-world psychiatric practice. The structure of \mathcal{K} is organized into a two-tiered hierarchy: the top level consists of broad diagnostic categories, such as mood disorders or psychotic disorders, while the lower level consists of more specific disorders, such as schizophrenia or bipolar disorder. Each entry in \mathcal{K} contains both general and specific clinical information. For each category, the knowledge base includes formal definitions, prototypical symptom clusters, common functional impairments, and the range of conditions that fall under the category. For example, in mood disorders, descriptors include definitions of typical mood disturbances like depressive or manic episodes and their impact on cognitive and social functioning, as well as the differential diagnostic boundaries with other conditions.

For each specific disorder, \mathcal{K} provides disorder-specific diagnostic criteria aligned with international standards, such as the ICD-11. These criteria include core symptoms, auxiliary features, and typical course. Additionally, the knowledge base outlines physical signs, relevant examination findings, and differential considerations that help distinguish it from other similar disorders. In the case of schizophrenia, for instance, the criteria include the required symptoms, the necessary duration for diagnosis, and the signs needed to differentiate it from other psychotic disorders.

The creation of \mathcal{K} was informed by a combination of authoritative sources, including the ICD-11, clinical guidelines from psychiatric associations, and expert consensus. This approach ensures that the knowledge base remains comprehensive, clinically valid, and directly relevant to psychiatric practice.

Once developed, \mathcal{K} serves as an essential resource for both model training and diagnostic reasoning. It allows for the automatic retrieval of relevant knowledge for a given diagnostic category or disorder, facilitating tasks such as training machine learning models and supporting hypothetico-deductive reasoning in psychiatric diagnosis. This ensures that both model-driven and human-driven diagnostic processes are informed by a reliable and expert-curated knowledge base.

Fig10 illustrates the broader category-level knowledge for schizophrenia. This category encompasses a range of disorders characterized by

Specific Disorders List

["6A20 Schizophrenia", "6B00 Generalised anxiety disorder", "6A72 Dysthymic disorder", "6E20 Mental or behavioural disorders associated with pregnancy, childbirth or the puerperium, without psychotic symptoms", "6A61 Bipolar type II disorder", "6A00 Disorders of intellectual development", "6D70 Delirium", "6E21 Mental or behavioural disorders associated with pregnancy, childbirth or the puerperium, with psychotic symptoms", "6B20 Obsessive-compulsive disorder", "6A40 Catatonia associated with another mental disorder", "6D72 Amnestic disorder", "6A62 Cyclothymic disorder", "6A23 Acute and transient psychotic disorder", "6B43 Adjustment disorder", "6A21 Schizoaffective disorder", "6A70 Single episode depressive disorder", "6A73 Mixed depressive and anxiety disorder", "6A71 Recurrent depressive disorder", "6A24 Delusional disorder", "6A60 Bipolar type I disorder", "6C20 Bodily distress disorder", "6A22 Schizotypal disorder", "6A41 Catatonia induced by substances or medications", "6D71 Mild neurocognitive disorder", "6D85 Dementia due to diseases classified elsewhere", "6C4G Disorders due to use of unknown or unspecified psychoactive substances", "6C4E Disorders due to use of other specified psychoactive substances, including medications", "6C46 Disorders due to use of stimulants including amphetamines, methamphetamine or methcathinone", "6B04 Social anxiety disorder", "6C40 Disorders due to use of alcohol", "6D10 Personality disorder", "6A02 Autism spectrum disorder", "6B02 Agoraphobia", "6B60 Dissociative neurological symptom disorder", "6C91 Conduct-dissocial disorder", "6C90 Oppositional defiant disorder", "6B83 Avoidant-restrictive food intake disorder", "6B40 Post traumatic stress disorder", "6C4F Disorders due to use of multiple specified psychoactive substances, including medications", "6D84 Dementia due to psychoactive substances including medications", "6B63 Possession trance disorder", "6B03 Specific phobia", "6B41 Complex post traumatic stress disorder", "6A03 Developmental learning disorder", "6C51 Gaming disorder", "6D86 Behavioural or psychological disturbances in dementia", "6A06 Stereotyped movement disorder", "6D83 Frontotemporal dementia", "6B80 Anorexia Nervosa", "6B23 Hypochondriasis", "6C49 Disorders due to use of hallucinogens", "6B81 Bulimia Nervosa", "6A05 Attention deficit hyperactivity disorder", "6B62 Trance disorder", "6B84 Pica", "6B21 Body dysmorphic disorder", "6B05 Separation anxiety disorder", "6B64 Dissociative identity disorder", "6B61 Dissociative amnesia", "6A04 Developmental motor coordination disorder", "6C4D Disorders due to use of dissociative drugs including ketamine and phencyclidine [PCP]", "6B66 Depersonalization-derealization disorder", "6B25 Body-focused repetitive behaviour disorders", "6B42 Prolonged grief disorder", "6B82 Binge eating disorder", "6C44 Disorders due to use of sedatives, hypnotics or anxiolytics", "6B22 Olfactory reference disorder", "6C01 Encopresis", "6D81 Dementia due to cerebrovascular disease", "6B06 Selective mutism", "6C50 Gambling disorder", "6B24 Hoarding disorder", "6C41 Disorders due to use of cannabis", "6B85 Rumination-regurgitation disorder", "6C47 Disorders due to use of synthetic cathinones"]

Figure 8: **Specific Disorders List.** A list of 76 disorders mentioned in the paper, including their corresponding codes and specific names.



Figure 9: Comparison of three reward strategies during reinforcement learning. The progressive process-based reward yields a smoother and more stable training trajectory compared to disorder-based and fixed process-based rewards.

significant impairments in reality testing. The figure highlights the diagnostic category itself, with specific ICD-11 codes (6A20 to 6A24), offering a detailed view of the various positive, negative, and cognitive dysfunction symptoms associated with schizophrenia. It presents an overview of the disorder's general features, helping contextualize the

condition in a larger diagnostic framework. This knowledge is essential for understanding the overarching classification of schizophrenia and its primary manifestations, providing the basis for accurate model training and reasoning during diagnostic processes.

Fig 11 delves deeper into a specific disorder within the category of Feeding or Eating Disorders (ICD-11 code: 6B82). It presents the detailed diagnostic criteria for Binge Eating Disorder (BED), which includes frequent binge episodes, the loss of control over eating, and the absence of compensatory behaviors seen in disorders like bulimia nervosa. The figure outlines both mandatory and supportive criteria for BED, with clinical manifestations such as emotional distress post-binge, emotional eating, and social avoidance. This knowledge representation zooms in on the specific disorder, providing precise diagnostic criteria and clinical

Diagnostic Category Knowledge Example

Diagnostic Category: Schizophrenia or Other Primary Psychotic Disorders

Code List:

- 6A20
- 6A21
- 6A22
- 6A23
- 6A24

Definition: Schizophrenia and other primary psychotic disorders are a group of mental disorders characterized by significant impairment in reality testing. These disorders are marked by:

• Positive Symptoms:

- Delusions
- Hallucinations
- Disorganized speech or behavior

• Negative Symptoms:

- Emotional blunting
- Reduced willpower or motivation
- Social withdrawal or isolation

• Cognitive Dysfunction:

- Impaired attention
- Memory deficits
- Executive function impairments

These symptoms cause the patient's thoughts, emotions, and behaviors to deviate significantly from cultural norms and cannot be attributed to other mental disorders, such as bipolar disorder or substance use disorders.

Figure 10: Diagnostic Category Knowledge Example for Schizophrenia and Other Primary Psychotic Disorders, detailing the ICD-11 codes, definition, and key symptoms, including positive, negative, and cognitive dysfunction.

cal features that are essential for accurate diagnosis and clinical decision-making.

9.3 Benchmark Example

Each example in the **MentalDx Bench** consists of two main components: the patient's Electronic Health Record (EHR) and the diagnostic conclusion provided by the doctor. The EHR includes a comprehensive set of details, such as the patient's Basic Information, Chief Complaint, Past Medical and Family History, Physical Examination results, and Mental Status Examination. The doctor's diagnostic conclusion includes a diagnostic category, a specific disorder from that category, and the corresponding disorder code according to the classification system. A specific example illustrating this structure is shown in the fig 12, which demonstrates how all these elements are integrated into a complete case representation.

Fig13 and 14 display the chain of thinking and

final diagnostic results of HuatuoGPT-O1-72B and GPT-OSS-120B, respectively, when diagnosing based on patient medical records. Both models correctly identified the diagnostic category, but there were discrepancies in the diagnosis of the specific disorder. Specifically, HuatuoGPT-O1-72B did not strictly adhere to the information in the medical records during its diagnostic process, instead making assumptions about potential conditions the patient might have, which led to an incorrect disorder diagnosis. Meanwhile, GPT-OSS-120B overlooked key information in the medical records, such as the patient's multiple hospitalizations due to depression, resulting in an inaccurate diagnosis. Although existing models perform reasonably well in certain areas, they still exhibit notable limitations during the diagnostic process. When handling complex cases, models often rely on hypothetical reasoning rather than comprehensively analyzing the patient's medical records or fail to consider important back-

ground information, such as hospitalization history, which can ultimately lead to incorrect diagnoses.

9.4 Benchmark Results

Tab 8 presents the performance of various language models on the **MentalDx Bench**, with results shown for individual diagnostic categories as well as for overall category accuracy (**CA**), disorder accuracy (**DA**), and joint accuracy (**JA**). Each cell represents the model’s score on a specific diagnostic task, expressed as a percentage. **CA** refers to the overall category diagnostic accuracy, **DA** refers to the accuracy in diagnosing disorders, and **JA** is the accuracy when both category and disorder diagnoses are correct simultaneously. The best result in each category is highlighted in dark, with the second-best in light.

Our models, **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B**, show outstanding performance across a number of key diagnostic tasks. Specifically, **MentalSeek-Dx-7B** achieves strong results in *ANX* (81.94%), *CATA* (70.00%), *SUD* (48.24%), and *BOD* (25.00%), demonstrating its capability in accurately diagnosing these critical mental health categories.

Furthermore, **MentalSeek-Dx-14B** shows further improvements, particularly in *CATA* (90.00%) and *DISS* (82.98%), surpassing other models in these areas. This demonstrates that scaling up the model enhances diagnostic accuracy, reinforcing the effectiveness of **MentalSeek-Dx** in handling complex mental health tasks.

In terms of overall category accuracy (**CA**), both **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B** perform excellently, with **MentalSeek-Dx-14B** achieving the highest accuracy at 67.56%, outpacing many other large language models. The disorder accuracy (**DA**) also shows strong results, particularly with **MentalSeek-Dx-14B** at 69.38%, further highlighting the model’s strength in diagnosing mental health disorders.

Notably, in terms of joint accuracy (**JA**), which considers both category and disorder diagnosis, **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B** again perform impressively, with **MentalSeek-Dx-14B** reaching 69.38%, showcasing its robust capability in simultaneously diagnosing both the overall category and the specific disorder.

Overall, the **MentalSeek-Dx** models excel across diagnostic categories, overall category accuracy, disorder accuracy, and joint accuracy, outperforming many other state-of-the-art and medical-

specific models. This demonstrates the innovative contribution of **MentalSeek-Dx** in the field of psychiatric diagnosis, especially with the breakthroughs achieved by **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B** in applying large-scale language models to mental health diagnosis.

9.5 Error Types Explanation

Tab 9 presents a breakdown of different types of errors observed in the diagnostic processes, each represented by an abbreviation. NR (No Reasoning) indicates cases where the model outputs an incorrect conclusion without providing any reasoning. WI (Wrong Inference) occurs when symptoms are correctly identified, but the reasoning leads to an incorrect diagnosis. IE (Identification Error) happens when the model lists potential diagnoses but fails to identify the correct one. OP (Overprediction) refers to situations where the model provides multiple possible diagnoses without a clear primary diagnosis. RA (Refusal to Answer) happens when the model refuses to answer, often citing safety concerns. Typo (Typographical Error) represents errors in characters that prevent the correct diagnosis from being made. Lastly, Other (Uncategorized Errors) covers any errors that do not fit into the above categories. These categories help to identify common pitfalls and guide improvements in diagnostic models.

As shown in Figure 3, the most frequent error is **Wrong Inference (WI)**, with a score of 89, indicating that this error occurred the most often in the dataset. The least frequent error is **Refusal to Answer (RA)**, with a score of only 8, representing the fewest occurrences among the error types. The distribution of errors revealed that the majority of mistakes fell into three primary categories: *wrong inference* (25%), *identification failure* (22%), and *symptom omission* (21%). These categories reflect key areas where the models struggled, with *wrong inference* representing errors in the diagnostic reasoning process, *identification failure* reflecting difficulties in recognizing the correct diagnosis from the set of possible conditions, and *symptom omission* pointing to the failure of the models to consider or properly incorporate key symptoms in their diagnostic conclusions. On the other hand, superficial errors, such as *typos* (7%) and *refusal to answer* (4%), were considerably less frequent. This finding highlights that the majority of diagnostic errors were not merely due to trivial issues, like typographical mistakes or the model’s refusal to

Model	ANX	CATA	SUD	BOD	STRESS	DISR	DISS	ELIM	EAT	PREG	MOOD	NCOG	NDEV	OCD	PERS	SCHIZ	CA	DA	JA
Small Language Models																			
Qwen2.5-7B	11.11	30.00	14.12	10.00	40.00	0.00	4.26	25.00	17.14	21.74	12.37	13.79	18.92	12.50	0.00	12.50	47.47	14.75	10.96
Qwen2.5-14B	18.06	63.33	29.41	0.00	52.50	44.44	21.28	25.00	62.86	82.61	22.68	13.79	48.65	40.00	0.00	10.94	61.66	29.92	24.72
Qwen3-32B	30.56	20.00	22.35	25.00	32.50	50.00	31.91	75.00	48.57	52.17	22.68	35.63	51.35	62.50	15.38	1.56	58.71	31.04	26.26
QWQ-32B	48.61	70.00	35.29	25.00	35.00	66.67	42.55	75.00	91.43	91.30	41.24	35.63	75.68	75.00	0.00	29.69	76.26	47.89	46.63
Large Language Models																			
LLaMA-3.1-70B	47.22	30.00	28.24	5.00	35.00	88.89	27.66	75.00	62.86	65.22	25.77	25.29	64.86	47.50	7.69	17.19	70.65	35.53	33.15
Qwen2.5-72B	34.72	40.00	32.94	0.00	35.00	38.89	34.04	75.00	74.29	78.26	23.71	16.09	62.16	55.00	30.77	6.25	70.51	33.57	29.78
Kimi-K2	55.56	33.33	44.71	0.00	57.50	72.22	48.94	100.00	91.43	65.22	24.74	32.18	89.19	70.00	23.08	10.94	75.70	45.08	43.40
Qwen3-Max	50.00	73.33	47.06	20.00	52.50	94.44	48.94	100.00	91.43	86.96	32.99	48.28	89.19	82.50	23.08	20.31	83.15	52.67	52.11
GPT-5.1	66.67	53.33	47.06	5.00	42.50	77.78	53.19	100.00	91.43	82.61	38.14	48.28	78.38	85.00	30.77	40.62	81.18	54.49	53.79
DeepSeek-R1	50.00	63.33	48.24	15.00	60.00	77.78	61.70	100.00	94.29	95.65	45.36	40.23	78.38	85.00	15.38	37.50	78.37	55.20	54.49
Claude-Sonnet-4.5	88.89	83.33	47.06	15.00	52.50	88.89	61.70	100.00	85.71	95.65	40.21	49.43	78.38	87.50	23.08	34.38	80.76	59.69	58.15
Medical-specific Large Language Models																			
ClinicalGPT-R1	9.72	23.33	10.59	5.00	27.50	5.56	2.13	0.00	14.29	17.39	16.49	5.75	13.51	5.00	0.00	1.56	36.10	10.53	6.46
HuatuogPT-o1-7B	9.72	20.00	12.94	5.00	25.00	5.56	0.00	50.00	5.71	8.70	15.46	6.90	5.41	10.00	0.00	17.19	43.26	11.24	8.01
GPT-OSS-20B	52.78	63.33	23.53	15.00	35.00	83.33	44.68	100.00	77.14	82.61	34.02	32.18	70.27	72.50	0.00	45.31	74.30	45.65	43.96
MedGemma-27B	54.17	50.00	28.24	25.00	42.50	55.56	29.79	75.00	74.29	65.22	25.77	32.18	72.97	70.00	23.08	40.62	66.99	42.84	39.89
Baichuan-M2	38.89	66.67	35.29	5.00	37.50	72.22	42.55	75.00	88.57	91.30	28.87	37.93	62.16	67.50	0.00	35.94	64.61	44.38	41.71
HuatuogPT-o1-72B	45.83	50.00	35.29	10.00	35.00	66.67	40.43	75.00	80.00	69.57	26.80	19.54	70.27	65.00	30.77	18.75	74.02	39.75	36.94
GPT-OSS-120B	54.17	66.67	38.82	20.00	50.00	88.89	51.06	50.00	94.29	65.22	41.24	43.68	70.27	75.00	38.46	45.31	75.84	52.53	50.00
Ours																			
MentalSeek-Dx-7B	81.94	70.00	48.24	25.00	52.50	88.89	72.34	100.00	97.14	86.96	54.64	66.67	94.59	90.00	61.54	56.25	82.58	67.56	67.13
MentalSeek-Dx-14B	84.72	90.00	48.24	20.00	52.50	88.89	82.98	100.00	97.14	91.30	57.73	66.67	89.19	85.00	53.85	67.19	83.99	70.08	69.38

Table 8: Performance of mainstream LLMs on MentalDx Bench (%). Models are grouped by parameter scale to examine how size affects performance across diagnostic categories. Each cell reports a model’s score on the corresponding task. For clarity, the best result in each category is shown in dark, the second-best in light.

Abbreviation	Full Form	Explanation
NR	No Reasoning	Directly outputs an incorrect conclusion without reasoning.
WI	Wrong Inference	Symptoms are identified but the reasoning leads to an incorrect diagnosis.
IE	Identification Error	A list of potential diagnoses is given but the correct diagnosis is not identified.
OP	Overprediction	No clear primary diagnosis, but multiple possible diagnoses are listed.
RA	Refusal to Answer	Refuses to answer citing safety concerns.
Typo	Typographical Error	Character errors that prevent matching the correct diagnosis.
Other	Uncategorized Errors	Other errors that don’t fit into the above categories.

Table 9: Explanation of abbreviations for error types in diagnosis

respond, but instead stemmed from more profound problems related to the reasoning and understanding of symptoms.

The error distribution further suggests that the dominant sources of failure in the models lie in deeper, more structural flaws in their diagnostic logic. Specifically, many models exhibited fragmented or incomplete diagnostic reasoning. This included unclear or vague delineation of symptoms, meaning the models struggled to differentiate between overlapping or similar conditions. Additionally, the models often failed to conduct a comprehensive analysis of the potential conditions, leading to incomplete or erroneous diagnostic conclusions. Moreover, the inability to effectively perform differential diagnosis, a critical component of psychiatric diagnosis, was a common limitation. These findings emphasize the need for further refinement in model training, particularly in improving the depth

of reasoning, symptom analysis, and differential diagnostic capabilities, to reduce errors and enhance diagnostic accuracy.

9.6 Prompt Details

Given a detailed medical record and a list of 16 distinct diagnostic categories, each accompanied by a brief definition and explanation, the model is tasked with inferring the most relevant suspected diagnoses based on the provided medical information. The specific format for the prompt is shown in the fig 15. The model should perform a comprehensive analysis of various aspects of the medical record to assess which diagnostic categories best match the features of the case. Each diagnostic category has a clear and distinct definition with specific symptom characteristics. The model’s task is to identify which categories most closely align with the information provided in the medical record.

The classification process must carefully consider all relevant details in the record, particularly the form of symptom presentation, onset and duration, severity, and factors such as the patient's age and sex. Additionally, the model should not restrict itself to a single category; rather, it should consider multiple potentially relevant categories based on the medical history. Therefore, the model should be capable of identifying one or more plausible suspected categories, demonstrating a thorough understanding and accurate analysis of the medical record.

Based on the suspected categories identified, the system will activate the classification tasks for all specific diseases associated with these categories. A total of 76 different diseases will need to be analyzed and classified in detail. The classification tasks for each disease will be handled according to the features of its respective category, ensuring that the diagnostic results for each disease accurately reflect its unique clinical manifestations and pathological characteristics.

In practice, these classification tasks will be processed in parallel, meaning that for each disease classification requirement, the system will perform multiple disease analyses simultaneously, rather than processing them sequentially one by one. This parallel processing approach not only improves efficiency but also ensures the independence and accuracy of the diagnosis for each disease. The specific format for the prompt is shown in the fig 15.

By activating the corresponding classification tasks and processing them in parallel, the system can quickly and comprehensively analyze the clinical characteristics of each disease, combining the existing suspected category information to make accurate inferences about each disease's diagnosis.

9.7 Refinement Pipeline Details

Given a specified disorder, the original medical record, the diagnostic criteria for the specified disorder, clinical manifestations, and common scenarios, the model should perform a detailed refinement and modification of the original medical record based on this information, ensuring that the record aligns solely with the characteristics and criteria of the specified disorder. During this modification process, the model should thoroughly analyze and comprehend all the diagnostic criteria and typical clinical manifestations of the specified disorder, ensuring that every part of the medical record adheres

to the relevant requirements of the disorder. Any symptoms, signs, or medical history unrelated to the specified disorder should be removed or adjusted to avoid misleading or irrelevant descriptions.

The refined medical record should emphasize and retain symptoms, signs, medical history, and other relevant elements that are highly associated with the specified disorder, ensuring that the record perfectly aligns with the diagnostic criteria for the disorder. Special attention should be given to the typical symptoms in the clinical presentation, symptom duration, the patient's medical history, and other factors that may influence the diagnosis. Furthermore, the model should ensure that the modified medical record meets the standards of medical writing in terms of semantics, logic, and structure, ensuring clarity and ease of understanding, while avoiding any ambiguous or contradictory descriptions that could lead to misdiagnosis or inaccurate conclusions. The specific format for the prompt is shown in the Fig 17,20and 18.

Ultimately, the modified medical record should fully comply with the diagnostic requirements for the specified disorder, making it a reliable clinical reference for subsequent diagnosis and treatment decisions.

9.8 Trajectory Building Details

We construct a Hypothetico-Deductive Reasoning Trajectory that explicitly mirrors the clinical diagnostic process adopted by human psychiatrists. Specifically, the trajectory is decomposed into three sequential reasoning stages: category-level hypothesis formation, disorder-level hypothesis narrowing, and criteria-based differential confirmation. At the first stage, the model is guided to reason from the raw medical record toward a broad diagnostic category by grounding symptom patterns and illness course in formal clinical definitions. Conditioned on this category-level hypothesis, the second stage further refines the reasoning to identify a small set of plausible disorders within the category, leveraging detailed clinical manifestations to generate competing diagnostic hypotheses. Finally, the model performs strict differential reasoning by systematically aligning patient evidence with diagnostic criteria and key exclusion points, culminating in a single confirmed diagnosis. By structuring reasoning as a progressive hypothesis generation and elimination process, our trajectory enforces a transparent, stepwise diagnostic logic that reflects

real-world psychiatric decision-making, rather than allowing conclusions to emerge implicitly or heuristically.

9.9 Medical Reasoning Quality

Tab 10 presents a breakdown of the abbreviations used for assessing the quality of medical reasoning in large language models (LLMs) generation. The abbreviations include CLR (Clinical Logical Rigor), which measures the alignment of the CoT with the clinical diagnostic process in psychiatry; ES (Evidence Strength), which evaluates whether the CoT is based on authoritative guidelines, consensus, or reliable research; CC (Critical Completeness), which checks if the CoT adequately utilizes both positive and negative information; and DDA (Differential Diagnosis), which assesses whether the CoT systematically lists and excludes potential differential diagnoses. These abbreviations help to evaluate and ensure the quality of reasoning within LLMs' diagnostic processes.

To further validate our reasoning process, we conducted a blinded evaluation, inviting 12 experienced medical experts. The experts used a Likert scale to assess four core dimensions of the reasoning process: Clinical Logical Rigor (CLR), Evidence Strength (ES), Critical Completeness (CC), and Differential Diagnosis Adequacy (DDA). Each dimension evaluates the quality and reliability of the reasoning process from different perspectives. Specifically, CLR assesses whether the reasoning follows a standard clinical thought process and whether any logical inconsistencies are present; ES examines whether the evidence cited in the reasoning is authoritative and credible; CC evaluates whether all relevant information is considered and whether alternative diagnoses are reasonably excluded; and DDA focuses on whether all possible diagnoses have been reasonably ruled out. Each dimension has clear rating criteria, and the experts scored various aspects of the reasoning process based on these standards. The specific rating guidelines are shown in Fig 21, 22, 23, and 24. The results provided valuable feedback, helping to optimize the diagnostic system and enhance its transparency, accuracy, and clinical reliability.

9.10 Hyperparameter Settings

We trained **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B** by building upon the pre-existing models Qwen2.5-7B-Instruct and Qwen2.5-14B-Instruct, respectively. The training process was carried out

in two distinct stages, each designed to address different aspects of fine-tuning and optimization for the models. The experiments were conducted on 16 A800 GPUs.

In the first stage, we performed supervised fine-tuning using a dataset of 14,093 carefully refined medical records. These records were selected to ensure they represented a broad range of medical diagnoses, allowing the models to learn from high-quality data. This stage lasted for 3 epochs, with a learning rate of 5×10^{-5} and a batch size of 128. The primary goal of this stage was to expose the models to the specific language and terminology used in psychiatric and medical contexts, enhancing their ability to handle complex diagnostic tasks.

Following the supervised fine-tuning, the second stage focused on reinforcement learning using GPRO (Generalized Proximal Policy Optimization) applied to **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B**. In this stage, we aimed to improve the model's diagnostic accuracy by addressing specific errors identified during Stage 1. The errors were analyzed and selected to form a training corpus, which was used to guide the learning process in areas where the models exhibited weaknesses. By leveraging model generations that showed diagnostic errors, we were able to focus the training on correcting those mistakes.

The reinforcement learning process utilized an Actor-Critic architecture, with separate learning rates for the Actor and Critic components set at 3×10^{-6} and 1×10^{-5} , respectively. The training was performed with a batch size of 64 and 3 epochs, providing enough iterations for effective fine-tuning. Several additional hyperparameters were used to guide the learning process, such as a discount factor of 1.0, which ensures equal weighting between future and immediate rewards, a value loss coefficient of 0.5 to control the learning of state values, and a clipping range of 0.2 to avoid large, destabilizing updates. A KL divergence coefficient of 0.001 was applied to maintain a balance between the old and new policies, ensuring that the learned policy did not deviate too far from the baseline. For each diagnostic prompt, the model generated 8 candidate responses, which were then evaluated to select the most appropriate diagnostic output.

This two-stage training procedure, combining supervised fine-tuning and reinforcement learning, allowed **MentalSeek-Dx-7B** and **MentalSeek-Dx-14B** to effectively learn from both curated data and

Abbreviation	Full Form	Explanation
CLR	Clinical Logical Rigor	The CoT aligns with the clinical diagnostic process in psychiatry.
ES	Evidence Strength	The CoT is based on authoritative guidelines, consensus, or reliable research.
CC	Critical Completeness	The CoT adequately utilizes both positive and negative information.
DDA	Differential Diagnosis	The CoT’s differential diagnoses are systematically listed and excluded.

Table 10: Explanation of abbreviations for CoT quality assessment in LLMs generation

targeted error correction, enhancing their performance in the challenging task of psychiatric diagnosis.

9.11 Reinforcement Learning Setup

The results presented in Fig 9 show that although the improvement with CRL is relatively small, it significantly facilitates a smoother transition during the reinforcement learning process. In our experiment, we compared three different reward strategies to assess their effectiveness in handling reinforcement learning tasks. The first strategy, Disorder-based reward, awards points only for the last predicted disorder. If the predicted disorder matches the actual disorder, points are awarded; otherwise, no points are given. The advantage of this strategy is its simplicity and clarity, but its major limitation is its heavy reliance on the final disorder prediction, which may hinder the optimization of earlier reasoning stages. The second strategy, Process-based reward, provides rewards for three key components—category, hypothesis list, and disorder prediction—with fixed scoring proportions of 1:1:1 for each component. This strategy takes a more holistic approach by considering each element’s contribution, but its limitation lies in the fixed ratio, which may not adapt flexibly to the learning needs at different stages, potentially leading to insufficient optimization in certain areas. The third strategy, **Progressive Process-based reward**, dynamically adjusts the reward distribution based on the training step. In different steps, the reward distribution changes as follows: in step 1, the ratio of category:list:disorder is 2:1:1; in step 2, it is 1:2:1; and in step 3, it is 1:1:2. This approach’s advantage is its ability to gradually emphasize disorder prediction as training progresses, allowing the model to focus more on precise diagnosis in the later stages while placing greater importance on category and hypothesis list in the earlier stages.

The reward updates for the process-based strategy are normalized to ensure comparability across different training phases. The experimental results show that while the Disorder-based reward

strategy provides relatively balanced feedback, its sparse scoring does not significantly enhance the model’s stability or accuracy. The Process-based reward strategy exhibits significant fluctuations during training, often interrupting the learning process and making stable improvement difficult. In contrast, the **Progressive Process-based reward** strategy demonstrates a more stable upward trend in the scores, indicating that this dynamic reward adjustment approach is more effective in gradually optimizing the model. In conclusion, the **Progressive Process-based reward** strategy clearly facilitates a smoother reinforcement learning process, particularly when handling complex diagnostic tasks, by reducing instability in reasoning and improving the accuracy of the final diagnosis.

9.12 Case Study

We selected a representative case for a comparative analysis, as shown in Fig 25, to illustrate the effects of different reasoning processes. Claude-Sonnet-4.5 employs a pattern-based association reasoning method, inferring specific disorders by identifying known patterns and retrospectively mapping them to diagnostic categories. However, this approach has significant flaws, manifesting as a paradigm misalignment. It relies on pre-set pattern associations while bypassing a structured psychiatric reasoning process, thereby neglecting a comprehensive evaluation of the patient’s condition. Specifically, Claude-Sonnet-4.5 directly links hallucinations to schizophrenia without considering alternative diagnoses or contextual factors. Its reasoning process fails to adequately explore differential diagnoses and does not effectively differentiate or validate various diagnostic possibilities. This reasoning model reflects a lack of a structured internal knowledge framework, ultimately leading to diagnostic errors and inaccurate predictions. The entire diagnostic process is shown in Fig 27.

In contrast, MentalSeek-Dx adopts a more structured, clinically grounded reasoning paradigm. This system strictly constrains the categorization reasoning process, ensuring that each step is based

on clinical evidence and diagnostic criteria. First, the system clearly defines the diagnostic domain and narrows the reasoning scope, then generates multiple hypotheses, which are further tested through scientific methods. Finally, MentalSeek-Dx applies standardized diagnostic criteria for differential diagnosis, ensuring that each hypothesis undergoes thorough clinical validation. This process not only follows a standard medical reasoning path but also makes the diagnostic process more transparent and accurate. Through this structured reasoning approach, MentalSeek-Dx provides more precise and interpretable diagnostic predictions, significantly reducing the risk of diagnostic errors and enhancing clinical decision-making reliability. The entire diagnostic process is shown in Fig 26.

Diagnostic Disorder Knowledge Example

Diagnostic Category: Feeding or Eating Disorders

ICD-11 Code: 6B82

Specific Disorder: Binge Eating Disorder

Definition:

Binge Eating Disorder (BED) is an eating behavior disorder characterized by frequent and repeated binge eating episodes (e.g., at least once a week over a few months). During each binge episode, the individual feels a lack of control over eating and consumes an amount of food significantly greater than usual or in a short period. This eating behavior is often associated with intense distress and negative emotions such as shame, guilt, or disgust. Unlike bulimia nervosa, individuals with BED do not regularly engage in compensatory behaviors such as self-induced vomiting, excessive exercise, or laxative abuse to prevent weight gain.

Clinical Manifestations:

Symptoms:

- **Uncontrollable Binge Eating:** Consuming large amounts of food in a short time, often accompanied by a sense of loss of control over eating (high: 70%-90%).
- **Post-Binge Distress:** Emotional distress after binge eating episodes, including feelings of guilt, shame, or self-loathing (high: 70%-90%).
- **Emotional Eating:** Using food to cope with negative emotions such as anxiety or loneliness (moderate: 50%-70%).
- **Night Eating:** Frequent eating at night, especially during emotional fluctuations (moderate: 30%-50%).
- **Increased Appetite:** Persistent hunger or desire to eat even when not physically hungry (moderate: 30%-50%).
- **Social Avoidance:** Avoiding social situations due to embarrassment or fear of eating in public (low: 10%-30%).

Diagnostic Criteria:

Mandatory Criteria (Diagnostic Basis):

1. Repeated binge episodes: Occurring at least once a week for at least 3 months. During each binge, the individual feels a lack of control and consumes a larger quantity of food than usual, typically within a short period.
2. Emotional distress: Significant negative emotions such as shame, guilt, or disgust follow the binge eating episode.
3. Lack of compensatory behaviors: Unlike bulimia nervosa, individuals with BED do not engage in regular compensatory behaviors (e.g., self-induced vomiting, excessive exercise, or laxative abuse) to prevent weight gain.

Supportive Criteria (Clinical and Epidemiological Basis):

- Frequent night eating: Eating frequently at night, especially when bored or under stress.
- Increased appetite: Persistent craving for food even when not hungry.
- Social avoidance: Avoiding social situations due to fear of eating in public.
- Weight concerns: Significant concern about body weight and shape, with a low self-evaluation.

Threshold Criteria:

1. Fulfillment of all conditions in the **Mandatory Criteria** confirms the diagnosis.
2. If only some of the **Mandatory Criteria** are met but multiple **Supportive Criteria** are present, further assessment of psychosocial factors and family history is necessary.

Figure 11: Binge Eating Disorder. Diagram illustrating the diagnostic criteria and clinical manifestations of Binge Eating Disorder, along with its ICD-11 code and specific symptoms.

EHR

- **Basic Information:** Female patient, 39 years old.
- **Chief Complaint:** Six days ago, the patient developed a poor mood, poor sleep, difficulty falling asleep, and reduced sleep duration, sleeping only about 5 hours per night. The patient's mother brought her to the hospital for admission this afternoon. This is the fourth hospitalization. During this episode, there were no hallucinations, auditory hallucinations, self-harm, suicidal ideation or behavior, chills, or fever.
- **Past Medical History and Family History:** The patient's great-grandfather had a mental illness, her great-grandfather's daughter had depression, her uncle's son had bipolar disorder, and her cousin has depression.
- **Physical Examination:** T 36.3°C, P 86 beats/min, R 20 breaths/min, BP 127/78 mmHg. The patient is awake and oriented, with steady breathing, coherent speech, and clear articulation. She cooperated during the examination. No jaundice on the skin or mucous membranes, no swollen superficial lymph nodes, neck soft with no resistance, no jugular vein distension, the trachea is centrally located, and the chest is normal. There is no widened intercostal space, percussion of the lungs reveals a clear sound, normal breath sounds with no rales or wheezing, heart borders are not enlarged on percussion, heart rate is within normal limits, rhythm is regular, no murmurs. The abdomen is flat, with no tenderness or rebound tenderness, no liver or spleen enlargement upon palpation, negative liver-jugular reflux test, and no pitting edema in both lower limbs.
- **Mental Status Examination:** The patient is awake, oriented, and focused with appropriate responses to questions, speaking calmly. She is appropriately dressed, has a flat affect, is partially self-sufficient in daily activities, and is actively engaging with others, being friendly toward medical staff and family members. In terms of self-awareness, there are no auditory hallucinations. Her speech rate is normal, with no formal thought disorder or content disorder. Common knowledge, judgment, comprehension, and calculation abilities are normal, as are both short-term and long-term memory.
- **Ground Truth:** [" diagnostic_category " : " Mood disorders " , "disorder_code": "6A71", "specific_disorder": " Recurrent depressive disorder "]

Figure 12: **Electronic Health Record (EHR):** Includes patient details, chief complaint, past medical and family history, physical examination, mental status examination, and ground truth.

HuatuogPT-01-72B

CoT: <think>The patient is a 39-year-old female who has recently been feeling down for the past six days, with sleep disturbances, sleeping only about 5 hours a night. This is her fourth hospitalization, suggesting a possible chronic issue. She has not experienced hallucinations, self-harm behavior, or fever, and there are no physical discomfort symptoms. Her family history includes mental illness, depression, and bipolar disorder, which makes me consider a mood disorder.

Mental examination shows that the patient is awake and oriented, with no hallucinations, and there are no obvious thought disturbances. Although her facial expression is somewhat flat, she is friendly toward both medical staff and family, indicating relatively good social functioning. Her self-awareness is preserved, suggesting that she has some insight into her condition.

Based on her symptoms and medical history, I suspect she may have a mood disorder, particularly Bipolar Disorder Type II, as Bipolar II typically presents with depressive episodes and hypomanic episodes, and she may have experienced hypomania in the past and is currently in a depressive phase. However, I must also consider other possibilities, such as Major Depressive Disorder.

Her family history of Bipolar Disorder increases the likelihood of her having this condition. However, I also need to rule out other conditions, such as anxiety disorders or adjustment disorders. Although she has no obvious anxiety symptoms, long-term stress could trigger an adjustment disorder.

In summary, I believe the most likely diagnosis is Bipolar II Disorder. Although further tests are needed for confirmation, the current evidence points in this direction. I will continue to monitor her symptoms to ensure the accuracy of the diagnosis. </think>

Conclusion: <answer> { "diagnostic_category": "Mood disorders", "candidate_disorders": ["6A61", "6A70", "6A71"], "disorder_code": "6A61", "specific_disorder": "Bipolar type II disorder" } </answer>

Figure 13: HuatuogPT-01-72B: Chain of thinking and diagnostic conclusion.

GPT-OSS-120B

CoT: <think> The patient is a 39-year-old woman who, for the past six days, has experienced a depressed mood, insomnia (≈5 h/night), and reduced overall functioning. There are no psychotic symptoms, suicidal ideation, or notable medical abnormalities. Mental status examination shows clear consciousness, appropriate orientation, intact attention, memory, and judgment, and no thought disorder. The presentation is consistent with a depressive episode. According to ICD-11, depressive episodes are classified under Mood disorders with the codes: 6A70 Single episode depressive disorder (first or isolated episode), 6A71 Recurrent depressive disorder (if there have been previous depressive episodes), 6A72 Dysthymic disorder (requires ≥2 years of chronic low mood, which is not present), and 6A73 Mixed depressive and anxiety disorder (requires prominent anxiety, not emphasized here). The most appropriate primary diagnosis is 6A70 Single episode depressive disorder. The diagnostic category is therefore "Mood disorders". Candidate codes are listed for completeness </think>

Conclusion: <answer>{"diagnostic_category": "Mood disorders","candidate_disorders": ["6A61", "6A70", "6A71"], "disorder_code": "6A61","specific_disorder": "Bipolar type II disorder"}</answer>

Figure 14: GPT-OSS-120B: Chain of thinking and diagnostic conclusion.

Prompt for the Category classification

Task

You are a psychiatric diagnostic expert. I will now provide the patient's medical history. Please analyze the provided medical information and give the corresponding diagnosis, i.e., the relevant category. I will also give you a list of all categories and their corresponding brief explanations.

Reference for Diagnostic Categories and Explanations

{ "Neurodevelopmental disorders": "Neurodevelopmental disorders are a group of mental or behavioral disorders that appear during an individual's developmental stage, typically in early childhood, especially before school age, and are characterized by developmental deficits that cause impairment in social, academic, or occupational functioning. The core features of these disorders include significant difficulties in acquiring and performing specific intellectual, motor, language, or social functions. While many mental and behavioral disorders can emerge during the developmental period, only those disorders that are primarily characterized by neurodevelopmental impairments are classified under this group",

"Schizophrenia or other primary psychotic disorders": "Schizophrenia and other primary psychotic disorders are a group of mental health conditions characterized by significant impairment in reality testing. The core symptoms include: positive symptoms such as delusions, hallucinations, and disorganized speech; negative symptoms such as flat affect, reduced volition, and social withdrawal; and cognitive dysfunctions such as impairments in attention, memory, and executive function. These symptoms cause the individual's thoughts, emotions, and behaviors to deviate markedly from the cultural norms, and cannot be attributed to other psychiatric disorders, such as bipolar disorder or substance use disorders.",

... other categories ... }

Output Format:

The diagnostic results should be output in JSON format, including:

- "diagnostic_category": Choose the most likely category or categories from the above list of disorders.

Note

- Do not output ''' json or similar content, strictly follow the output example format.

Output Example:

```
-----  
{  
  "id": "123456"  
  "diagnostic_category": [ "Disorders specifically associated with stress"  
]  
}
```

Input:

- Test medical record: <MEDICAL_RECORD>

Figure 15: Prompt for the Category classification

Prompt for the Disorder diagnosis

Task

You are a psychiatric diagnostic expert. I will provide the patient's medical record, along with the specific disorder and its diagnostic criteria. Your task is to determine whether the patient meets the criteria for the given disorder.

Disorder and Diagnostic Criteria

Disorder: "6A20"

Diagnostic Criteria:

- **Mandatory Criteria (for Diagnosis):**
 - **Symptom Duration:** Core symptoms (positive symptoms, negative symptoms, or cognitive impairment) must persist for at least one month.
 - **Functional Impairment:** Significant impairment in important areas such as work, education, and social interactions.
 - **Exclusion of Other Causes:** Exclude similar symptoms caused by other medical conditions (such as brain tumors or epilepsy) or substance use (including drug abuse).
- **Supporting Criteria (Clinical and Epidemiological Evidence):**
 - **Typical Clinical Presentation:**
 - * **Positive Symptoms:** Hallucinations (especially auditory hallucinations), delusions (paranoid delusions, delusions of reference, etc.), thought disorder (thought fragmentation, poverty of thought).
 - * **Negative Symptoms:** Blunted affect, social withdrawal, lack of motivation, poverty of speech.
 - * **Cognitive Impairment:** Attention deficits, memory problems, executive function impairments.
 - * **Behavioral Abnormalities:** Stereotyped movements, bizarre behaviors, catatonia (e.g., waxy flexibility, echopraxia).
 - **Neuropsychological Testing:** Impairment in multiple cognitive domains.
- **Threshold Criteria:**
 - The patient must meet all conditions of the "Mandatory Criteria" and exhibit at least two core symptoms (either positive, negative, or cognitive impairment).

The diagnosis process should also take into account the patient's detailed medical history, family history, and clinical presentation to ensure the exclusion of other potential causes.

Output Format:

- Provide the diagnosis result in JSON format.
- Do not include any extra formatting like " " json or similar. Strictly follow the output sample format.

Output Sample:

```
{
  "id": "123456",
  "has_disorder_6A20": "no"
}
```

Input:

- Patient's medical record: <MEDICAL_RECORD>

Figure 16: Prompt for the Category classification

Prompt for Selected Disorder

Task Definition
 You are a professional psychiatrist. You are given an anonymized clinical case that needs to be revised and completed to strictly align with the diagnostic criteria of <selected_disorder>. Any irrelevant or non-conforming symptoms should be removed.

Input
 You will receive the following information for ICD-11 code 6A20:

Diagnostic Criteria
 <disorder_category>

Clinical Manifestations
 <clinical_manifestation>

Record to Revise
 <record_input>

Core Requirements

- The revised case must fully meet the diagnostic criteria of the target disorder. If any core criteria are missing, you must complete them through generation.
- Remove unrelated symptoms or information to avoid introducing comorbid or confounding conditions.

Output Format
 Output in JSON format only, without any explanations or “ ‘ json tags, as follows:

```
{
  "revised_record": "...
}
```

Figure 17: Prompt for Selected Disorder

Prompt for Category Reasoning

You are an experienced psychiatric expert with extensive clinical diagnostic experience. You will first be provided with: the patient’s medical record, diagnostic category, and clinical definition, and need to supplement the reasoning process between the two.

Requirements:

- Act as a professional doctor, presenting your internal diagnostic monologue.
- Start with the patient’s medical record, gradually combining the clinical definition standards to pinpoint the category.
- Evidence-based reasoning should naturally link the reasoning process with the definition.
- Begin with: "Alright, let me go through the medical reasoning step by step. First, I need to combine the symptom combinations and the course of the illness to frame the broad syndrome category this patient belongs to..."

Input Information

- **Patient Medical Record:** <INPUT>
- **Diagnostic Category:** <DIAGNOSTIC_CATEGORY>
- **Clinical Definition:** <DEFINITION>

Output Example

```
{
  "category_reasoning": "Alright, let me go through the medical reasoning step by step.
    First, I need to combine the symptom combinations and the
    course of the illness to frame the broad syndrome category
    this patient belongs to...", // Here, place the reasoning
    process you generate.
  "diagnostic_category": "<DIAGNOSTIC_CATEGORY>"
}
```

Figure 18: Prompt for Category Reasoning

Prompt for Disorder Reasoning

You are an experienced psychiatric expert with extensive clinical diagnostic experience. You will first be provided with: the patient's medical record, the corresponding list of disorders, and clinical manifestations. Your task is to list the possible corresponding disorders the patient may have based on the provided information.

Requirements:

- Act as a professional doctor, presenting your internal diagnostic monologue.
- Start with the patient's medical record, gradually combining the clinical manifestations to pinpoint 2-3 suspected disorders.
- Use evidence-based reasoning to naturally and closely connect the reasoning process with the clinical manifestations.
- Begin with: "Next, I need to work within the syndrome range and, combining clinical symptoms, deduce the possible disorders..."

Input Information:

- **Patient's Medical Record:** <INPUT>
- **Candidate Disorder List:** <SPECIFIC_DISORDER>

Output Example:

```
{
  "disorder_reasoning": "Next, I need to work within the syndrome range and, combinin
                        clinical symptoms, deduce the possible disorders...",
                        Here, place the reasoning process you generate
  "candidate_disorders": [...] // "code" list
}
```

Figure 19: Prompt for Differential Reasoning

Prompt for Differential Reasoning

You are an experienced psychiatric expert with extensive clinical diagnostic experience. You will first be provided with: the patient's medical record, a list of suspected disorders, the corresponding diagnostic criteria, and the final confirmed diagnosis. Your task is to supplement the differential diagnostic reasoning process.

Requirements:

- Act as a professional doctor, presenting your internal diagnostic monologue.
- Strictly follow the diagnostic criteria for differential diagnosis, and conclude the diagnosis as <CODE>.
- Use evidence-based reasoning to naturally and closely link the differential reasoning process with the diagnostic criteria.
- Begin with: "Finally, I need to combine the diagnostic criteria and key differential points to complete the differentiation and confirm the final diagnosis..."

Input Information:

- **Patient's Medical Record:** <INPUT>
- **Candidate Disorder List:** <SPECIFIC_DISORDER>

Output Example:

```
{
  "differential_reasoning": "Finally, I need to combine the diagnostic criteria and key
                        differential points to complete the differentiation and confirm the final diagnosis...",
                        Here, place the reasoning process you generate
  "confirmed_disorder": "<CODE>"
}
```

Figure 20: Prompt for Differential Reasoning

Clinical Logical Rigor Rating Criteria

Dimension 1: Clinical Logical Rigor

Purpose:

Evaluate whether the reasoning process aligns with standard clinical thinking. This dimension assesses the internal consistency and logical progression of the diagnostic reasoning path.

Core Focus:

- Does the reasoning follow a coherent clinical flow (e.g., from symptom presentation to diagnostic hypothesis to investigation or confirmation)?
- Are transitions between reasoning steps medically justified?
- Is there evidence of logical gaps or contradictions?

Guidelines for Rating (1–5 Likert Scale):

- **Score 5 – Excellent:**
The reasoning process is rigorous and mirrors professional diagnostic workflows. Each step logically builds upon the previous one, demonstrating clear clinical justification and no inconsistencies.
- **Score 4 – Good:**
Mostly consistent reasoning with minor issues. Logical flow is preserved overall, but one or two steps may lack explicit justification or be slightly underdeveloped.
- **Score 3 – Fair:**
The reasoning includes some coherent transitions but also contains noticeable logical gaps, unclear assumptions, or weak linkages between major steps.
- **Score 2 – Poor:**
The reasoning shows significant inconsistencies or unjustified jumps between steps. Clinical logic is frequently broken or poorly structured.
- **Score 1 – Very Poor:**
The reasoning path is fragmented or incoherent, with no clear structure resembling clinical logic. Diagnostic conclusions appear arbitrary or unsupported by prior steps.

Figure 21: Clinical Logical Rigor Rating Criteria

Evidence Strength Rating Criteria

Dimension 2: Evidence Strength

Purpose:

Assess the strength, clarity, and credibility of the medical evidence used to support the reasoning. This dimension evaluates whether the knowledge cited is grounded in authoritative and up-to-date clinical resources.

Core Focus:

- Is the reasoning supported by evidence from recognized guidelines (e.g., ICD, DSM, WHO), expert consensus, or peer-reviewed clinical studies?
- Are the references to evidence explicit, precise, and contextually appropriate?

Guidelines for Rating (1–5 Likert Scale):

- **Score 5 – Excellent:**
The reasoning explicitly cites high-quality sources such as current clinical guidelines, landmark trials, or authoritative reviews. The evidence is specific, well-integrated, and directly supports the reasoning.
- **Score 4 – Good:**
The reasoning is generally well-supported by clinical knowledge, with some explicit references. However, some sources may be less current or less clearly stated.
- **Score 3 – Fair:**
The evidence is partially grounded in clinical understanding but lacks specificity. The reasoning may include general or vague references such as “studies show” or “it is commonly believed,” without citing concrete sources.
- **Score 2 – Poor:**
The reasoning relies on outdated, unclear, or weak sources. It may reflect misconceptions or obsolete diagnostic categories, reducing its credibility.
- **Score 1 – Very Poor:**
No discernible clinical evidence is used. The reasoning appears speculative, inaccurate, or contrary to established medical knowledge.

Figure 22: Evidence Strength Rating Criteria

Critical Completeness Rating Criteria

Dimension 3: Critical Completeness

Purpose:

Evaluate whether the reasoning thoroughly integrates all relevant clinical information, including both positive findings and critical negative evidence.

Core Focus:

- Does the reasoning account for all key clues in the case, including negative symptoms, lab findings, and contextual information?
- Are any important signs, contradictions, or alternative possibilities overlooked?

Guidelines for Rating (1–5 Likert Scale):

- **Score 5 – Excellent:**
The reasoning comprehensively incorporates all relevant case information, including less obvious symptoms or negative findings. Contradictory data is acknowledged and thoughtfully analyzed.
- **Score 4 – Good:**
Most critical information is considered, though a few minor details or counterpoints may be overlooked. Overall, the analysis remains balanced and well-supported.
- **Score 3 – Fair:**
The reasoning captures the main symptoms but partially overlooks important secondary findings or negative evidence. Some diagnostic bias may be present.
- **Score 2 – Poor:**
Key details—especially disconfirming or contradictory evidence—are neglected. The analysis appears narrow or overly focused on a single hypothesis.
- **Score 1 – Very Poor:**
The reasoning is based on incomplete information, ignoring significant findings that undermine the conclusion. The result is a superficial or misleading diagnostic interpretation.

Figure 23: Critical Completeness Rating Criteria

Differential Diagnosis Adequacy Rating Criteria

Dimension 4: Differential Diagnosis Adequacy

Purpose:

Assess whether the reasoning demonstrates a structured and comprehensive approach to differential diagnosis.

Core Focus:

- Has the response considered a broad but relevant range of possible diagnoses?
- Are common, high-probability, or life-threatening conditions addressed?
- Are exclusions logically and explicitly justified?

Guidelines for Rating (1–5 Likert Scale):

- **Score 5 – Excellent:**
A well-structured differential diagnosis is presented, including both common and critical conditions. Diagnoses are prioritized by likelihood, and exclusions are clearly justified using clinical evidence.
- **Score 4 – Good:**
Several plausible alternatives are considered, and key conditions are ruled out with reasonable justification. Some minor omissions may exist, but the diagnostic reasoning remains strong.
- **Score 3 – Fair:**
A limited differential is provided with partial justification. Some relevant conditions may be missing or excluded without clear rationale. The diagnostic scope is somewhat narrow.
- **Score 2 – Poor:**
Only one diagnosis is offered, or the differential includes mostly irrelevant or unsupported options. Little to no effort is made to rule out other possibilities.
- **Score 1 – Very Poor:**
No meaningful differential diagnosis is attempted. The reasoning is overly simplistic, with no recognition of alternative explanations.

Figure 24: Differential Diagnosis Adequacy Rating Criteria

EHR

- **Basic Information:** Female patient, 31 years old.
- **Chief Complaint:** The patient experienced coughing and sputum production one month ago, which was not taken seriously. Two weeks ago, after receiving intravenous treatment at ** City Emergency Center, she began to exhibit nonsensical speech, disorganized behavior, and reported that she was conversing with her own thoughts. She also felt nervous, fearful, and experienced episodes of agitation. She then visited the emergency department of our hospital, where the diagnosis was unclear, and medication was not specified. During this period, the patient did not take any medication, and her condition did not improve. Four days ago, her symptoms worsened. Two days ago, she visited our outpatient clinic, and one day ago, she started taking Olanzapine 5mg at night, with noticeable improvement in her condition. She was admitted to our department for treatment today.
- **Past Medical History and Family History:** The patient has a family history of hepatitis B, and the patient herself has a small positive hepatitis B surface antigen (HBsAg).
- **Physical Examination:** T 36.3°C, P 86 beats/min, R 20 breaths/min, BP 127/78 mmHg. The patient is fully alert, with stable respiration, coherent speech, and clear articulation. She cooperated during the examination. There is no jaundice in the skin and mucous membranes, no enlargement of superficial lymph nodes, no neck stiffness, no resistance, and no jugular vein distention. The trachea is midline, the chest is normal, and there is no widening of the intercostal spaces. Percussion of the lungs reveals clear resonance, with normal breath sounds and no rales or wheezing. Heart borders are not enlarged upon percussion, heart rate is 86 beats/min, rhythm is regular, and there are no murmurs. The abdomen is flat, without tenderness or rebound tenderness. The liver and spleen are not palpable, and there is no hepatosplenomegaly. The hepatojugular reflux test is negative, and there is no pitting edema in both lower limbs. Physiological reflexes are present, and pathological reflexes are absent.
- **Mental Status Examination:** The patient is alert, oriented, with normal attention, and provides relevant responses in a calm voice. She is dressed appropriately, with a natural expression, self-care abilities, and actively interacts with others. She is friendly toward medical staff and family. The patient lacks insight, experiences auditory hallucinations of a commentative type, has normal speech rate, and exhibits thought content disturbances, including paranoid delusions. There is no formal thought disorder, and basic common sense, judgment, understanding, calculation, short-term memory, and long-term memory are all normal.
- **Ground Truth:** ["diagnostic_category": "Schizophrenia or other primary psychotic disorders", "disorder_code": "6A20", "specific_disorder": "Schizophrenia "]

Figure 25: EHR inputs.

MentalSeek-Dx

- Let me break down the medical reasoning in stages.
- **First, I need to categorize the patient's syndrome based on the combination of symptoms and disease course.** The patient is a 31-year-old female who developed nonsensical speech, disorganized behavior, auditory hallucinations (commentative hallucinations), and persecutory delusions during the course of the illness, which strongly suggests a psychotic disorder. According to clinical definitions, the core features of schizophrenia and other primary psychotic disorders are significant impairment in reality testing, manifested by positive symptoms (such as delusions and hallucinations), negative symptoms, and cognitive dysfunction. In this case, the patient exhibits clear positive symptoms: auditory hallucinations (commentative in nature, involving conversations with her thoughts) and persecutory delusions, which align with the definition of distorted reality. Moreover, the disease course shows that the symptoms lasted more than two weeks and worsened during the untreated period, with partial relief after Olanzapine treatment. This supports the chronic nature of primary psychotic disorders. Additionally, the mental status exam reveals a lack of insight, with no significant thought formal disorder, but prominent content disorder, and no negative symptoms or significant cognitive impairment, which further rules out other disorders like substance-induced or affective psychosis. Finally, considering her medical history (no family history of psychiatric disorders) and physical examination findings that rule out organic causes, **the symptoms cannot be attributed to other conditions, leading to a diagnosis within the schizophrenia spectrum.**
- **Next, I need to deduce the possible disease entities within this syndrome range, based on the clinical symptoms.** The patient is a 31-year-old female with an acute onset, having started with coughing and sputum production over the past month, followed by psychiatric and behavioral abnormalities. The mental status examination shows the patient is fully alert, oriented, but with clear positive symptoms: commentative auditory hallucinations and persecutory delusions, accompanied by emotional symptoms such as nervousness, fear, and restlessness. It is worth noting that her symptoms significantly worsened after receiving intravenous treatment for a respiratory infection, and her condition continued to deteriorate during the untreated period, but improved after starting Olanzapine treatment. The core symptoms of the patient are psychotic symptoms, including hallucinations and delusions, which fit the characteristics of schizophrenia spectrum disorders. However, her acute onset and clear physical disease trigger (respiratory infection) raise the possibility of an acute transient psychotic disorder. Additionally, the patient exhibits significant emotional symptoms (nervousness, fear, agitation), but these do not meet the criteria for a distinct manic or depressive episode, making a schizoaffective disorder less likely. Delusional disorder is also not a good fit because, in addition to persecutory delusions, the patient has clear hallucinations. **Overall, the most likely diagnoses are schizophrenia and acute transient psychotic disorder, and further observation of the disease progression is required to exclude psychotic symptoms caused by an underlying physical condition.**
- **Finally, I need to apply diagnostic criteria and differential points to complete the differentiation and confirm the final diagnosis.** The patient is a 31-year-old female with an acute onset, with symptoms progressing from coughing and sputum production to nonsensical speech and disorganized behavior within just two weeks, **meeting the core criteria of 6A23 for "acute onset of symptoms with maximum severity reached within two weeks."** The patient exhibits clear commentative auditory hallucinations and persecutory delusions, fulfilling the major symptom requirements of 6A23. **Regarding the disease course, the time from the first psychiatric and behavioral abnormalities to hospitalization was just over two weeks, far from the one-month duration required by 6A20.** The patient had a clear history of physical illness prior to the onset (coughing and sputum production, likely a respiratory infection), and after the onset, she experienced emotional fluctuations such as nervousness and fear, which aligns with the psychosocial factors and physical illness criteria in the supporting conditions of 6A23. The improvement in her condition after Olanzapine treatment further supports the characteristics of an acute psychotic disorder. **While the patient has persecutory delusions and hallucinations, the duration of symptoms, acute onset pattern, and clear triggering factors are more consistent with 6A23 than 6A20.**

Figure 26: MentalSeek-Dx (ours) prediction.

Claude-Sonnet-4.5

• Key Information Summary

- **Basic Information:** Female, 31 years old, Cough and sputum production one month ago, followed by the onset of psychiatric symptoms two weeks after intravenous treatment.
- **Main Symptoms:** Nonsensical speech, Disorganized behavior, Auditory hallucinations (commentative), Persecutory delusions, Claims "I am talking to my own thoughts", Nervousness, fear, Occasional agitation

• Diagnostic Analysis

• Schizophrenia (Most Likely Diagnosis)

Supporting Factors: Persistent positive symptoms (hallucinations, delusions); Good response to antipsychotic medication, although symptoms have not fully resolved and there is a risk of relapse; Normal cognitive function, no significant mood symptoms, fitting the characteristics of acute onset and gradual progression

• Acute and Transient Psychotic Disorder (Less Likely)

Against This Diagnosis: Symptoms are persistent and there is a tendency for recurrence, whereas acute transient psychotic disorders typically have a short duration and rapid improvement

• Schizoaffective Disorder (Less Likely)

Against This Diagnosis: Mood symptoms are not prominent, and the diagnosis does not meet the criteria for schizoaffective disorder due to the predominance of psychotic symptoms

• Exclusion Diagnoses

- **Substance-Induced Psychotic Disorder:** No history of psychoactive substance use, and symptoms are unrelated to substance use.
- **Organic Psychotic Disorder:** Physical examination was normal, and cognitive function was intact.
- **Mood Disorder with Psychotic Features:** No symptoms of mania or depression that meet the criteria, and the symptoms have persisted beyond the typical duration for mood disorders.

• Final Diagnosis

Persistent psychotic symptoms (hallucinations, delusions, etc.) with a good response to antipsychotic medication, no significant mood symptoms, relatively preserved cognitive function, symptoms not fully resolved, and a possibility of recurrent relapse, most consistent with schizophrenia (Schizophrenia), classified under: Schizophrenia or other primary psychotic disorders.

Figure 27: Claude-Sonnet-4.5 prediction.