

# Rethinking Depression Prediction from a Fine-Grained Subscore Modeling Perspective via Multi-Task Learning

Anonymous ACL submission

## Abstract

Standard depression assessment relies on instruments such as the clinician-rated Hamilton Depression Rating Scale (HAMD) and the patient-reported Patient Health Questionnaire (PHQ-8), but manual scoring is time-consuming and subject to inter-rater variability. Prior automated approaches typically regress a single total score or coarse severity category, lacking the fine-grained subscore-level supervision needed for precise clinical diagnosis. To address this, we propose MTSP (Multi-Task Subscore Prediction), a fine-grained model for subscore prediction via multi-task learning. MTSP achieves state-of-the-art performance on the public E-DAIC dataset (MAE 3.48, RMSE 4.57) and generalizes well to the public PDCH and a large-scale private clinical dataset (CIDH), outperforming total score regression baselines and Qwen3-14B direct scoring. We further show that multi-task learning is essential, subscore-level supervision improves assessment by better capturing symptom-cluster structure, and prior constraints plus task-level self-paced learning enhance robustness to subscore difficulty and annotation noise. Our code is available at <https://anonymous.4open.science/r/MTSP-1D51/>.

## 1 Introduction

Depression is a prevalent mental health disorder affecting over 280 million people worldwide (World Health Organization, 2025). Accurate and timely assessment of depression severity is essential for personalized treatment and longitudinal monitoring. In clinical practice, severity is typically quantified with standardized rating scales. Clinical instruments such as the Hamilton Depression Rating Scale (HAMD-13) (Hamilton, 1960) and brief self-report questionnaires such as the 8-item Patient Health Questionnaire (PHQ-8) (Kroenke et al., 2009) provide clinically validated measures of depressive symptomatology and are widely used to




Figure 1: **Illustration of individual-level heterogeneity.** Two representative patients from the CIDH dataset, NN\_00002178\_S001 (Patient A) and NN\_0000994\_S001 (Patient B), share the same total score (15) but exhibit distinct symptom patterns: Patient A shows a “Somatic–Sleep” pattern characterized by high scores in *Insomnia* (2, 2, 2), *Work & Interests* (3), *GI Symptoms* (2), and *Somatic Symptoms* (2), whereas Patient B presents a “Depressive–Cognitive” profile dominated by *Guilt* (3), *Suicide* (3), and *Psychic Anxiety* (2).

evaluate symptom severity. However, their repeated use is time- and resource-intensive, relies on trained psychologists or patient engagement, and remains vulnerable to subjectivity and inter-rater or self-report variability, demanding the development of *automated, reliable and exact* depression assessment systems.

Existing methods predominantly formulate depression assessment as a *coarse-grained* mapping  $X \rightarrow y$ , where  $y$  is a single global target such as the total score or a binary severity category (Ji et al., 2022; Zhao et al., 2025). We argue that this formulation suffers from intrinsic limitations because it overlooks the complex heterogeneity of depressive symptomatology (Buch and Liston, 2021). Patients with identical total scores can exhibit drastically different symptom profiles. For example, Figure 1 shows two representative patients with the same total score but qualitatively distinct patterns of task-level severity. When a model is trained




Figure 2: **Statistical discrepancy on the CIDH dataset** The x-axis represents each total score  $s$  (with sample size  $n_s$ ). The y-axis shows the *diversity across subscores*, defined as  $D_{s,k} = \max_{i \in \mathcal{P}_s} y_{i,k} - \min_{i \in \mathcal{P}_s} y_{i,k}$ , where  $y_{i,k}$  denotes the score of the  $k$ -th sub-score for the  $i$ -th sample, and  $\mathcal{P}_s$  is the sample set with the same total score  $s$ . Based on this, we compute  $\max_k D_{s,k}$ ,  $\text{mean}_k D_{s,k}$ , and  $\min_k D_{s,k}$ , which respectively indicates the max, mean, and min variation across subscores with the same total score.

solely to predict the total score, such individualized patterns inevitably make it difficult to capture clinically meaningful differences between patients. On the CIDH dataset, our statistical analysis (Figure 2) further demonstrates that such discrepancies are not anecdotal but systemic. The resulting curves reveal that, across different levels of depressive severity, even patients who share the same total score exhibit substantial task-level variability, with discrepancies being especially pronounced for those with total scores between 12 and 20.

To bridge this gap, we propose **MTSP (Multi-Task Subscore Prediction)**, a unified framework that shifts the paradigm from global coarse scoring to fine-grained subscore prediction by formulating it as a multi-task learning (MTL) problem. However, designing an effective MTL framework for depression assessment poses two key challenges.

**Challenge 1: How to build the correlation across subscore (task)** Depression assessment is inherently a multi-task prediction problem: subscores are not independent, but instead have correlation. Treating each subscore as an isolated task neglects these cross-task dependencies and can lead to clinically inconsistent predictions.

**Challenge 2: How to solve the inconsistent task difficulty.** Different subscores exhibit markedly distinct learning difficulties. Our preliminary experiments (in Fig.4) using Qwen3-14B with zero-shot prompting (non-reasoning mode) for direct PHQ-8 scoring reveal that even state-of-the-art LLMs show highly uneven performance across subscores. This suggests that depression assessment subscores possess inherently different difficulty levels, and that a uniform training schedule can cause

difficult tasks to interfere with the learning of easier ones. To address these challenges, MTSP employs a shared-specific task prediction framework. We construct a clinically informed *task correlation graph* that encodes prior knowledge about symptom correlations, and introduce prior constraint that penalize deviations among subscores within the same clinically defined clusters, encouraging predictions consistent with these graph-based priors. In addition, we further incorporate a **Task-level Self-Paced Learning (T-SPL)** strategy that operates on task rather than individual samples, adaptively reweighting task losses according to their learning difficulty. Our contributions are summarized as follows:

- **Fine-grained subscore modeling formulation.** We propose MTSP, a fine-grained framework that jointly predicts individual subscores of clinical scales to better describe patients’ profiles and predict symptoms.
- **Multi-task learning for subscores.** We capture the structured dependencies among clinical subscores using a MTL mechanism that explicitly models and regularizes the relationships among subscores. Meanwhile, we introduce a task-level self-paced learning strategy that explicitly accounts for subscore difficulty.
- **State-of-the-art performance.** On the public E-DAIC dataset, MTSP achieves state-of-the-art performance with an MAE of 3.48 and an RMSE of 4.57, outperforming both conventional total score regression baselines and direct scoring by the Qwen3-14B large language model. It also excels on the PDCH and CIDH datasets, demonstrating its robustness across different datasets.

## 2 Related Work

### 2.1 Natural Language Processing for Clinical Mental Health

Natural language processing (NLP) for mental health assessment has attracted sustained interest. Early work typically relied on feature engineering and rule- or lexicon-based pipelines to extract clinically meaningful signals from text (Gkotsis et al., 2017; Coppersmith et al., 2014; De Choudhury et al., 2013). With the rise of pre-trained language models (PLMs), recent approaches reduce reliance on manual features and improve generalization. Ji et al. (2022) introduce *MentalBERT*, a

domain-specific PLM for mental healthcare that consistently improves mental disorder detection benchmarks. In parallel, general-purpose LLMs such as ChatGPT have been explored in healthcare, and systematic reviews (e.g., Li et al. (2024)) summarize both their emerging use cases and their limitations in safety-critical clinical scenarios.

A growing line of work focuses on *clinical interviews* and explainable depression assessment Zhang et al. (2025); Zhao et al. (2025); Mandal et al. (2025); Bi et al. (2025). For example, Zhang et al. (2025) propose a personalized retrieval-augmented generation framework that retrieves evidence snippets from interview transcripts and generates natural-language explanations for depression decisions. Mandal et al. (2025) design a question-wise multimodal fusion model that predicts PHQ-8 item scores to enhance interpretability, and Bi et al. (2025) adopt a multi-agent LLM framework to operationalize structured psychiatric interviews and generate psychometric reasoning traces. Clinical rating instruments in routine practice require *fine-grained level scoring* over multiple correlated symptom dimensions. Most prior work either predicts a single total score or treats symptom-related targets independently, lacking explicit mechanisms to jointly model cross-task structure.

## 2.2 Multi-Task Learning in Clinical Assessment

Beyond predicting a single total score, a few recent studies have begun to operate at the subscore level. van Steijn et al. (2022) focus on text-based depression severity prediction on the DAIC-WOZ corpus and use a multi-target regression framework to predict PHQ-8 item scores as interpretable symptom-level outputs that can be aggregated into overall severity. Mandal et al. (2025) propose QuestMF, a question-wise multimodal fusion framework that performs ordinal classification for each PHQ-8 item and derives the global PHQ-8 score by summing these predictions. In parallel, Zhang et al. (2025) formulate EMDRC as an explainable multimodal depression recognition task and introduce PHQ-aware multi-task methods that use PHQ-8 item labels as auxiliary supervision for utterance-level classification and dialogue-level symptom summarization. However, these approaches either treat PHQ subscores independently and do not consider the inherent task relations.

## 3 Method

### 3.1 Problem Formulation

Given a clinical interview transcript  $\mathbf{x}_i$ , our goal is to predict depression subscores  $\mathbf{y}_i = (y_{i1}, \dots, y_{iK})$  in a multi-task regression setting (Fig.3(a)), where  $K$  is the number of subscores and  $y_{ik} \in \mathbb{R}$  is the score for subscore  $k$  of sample  $i$ . We support two depression scales: **PHQ-8** is a self-report instrument containing  $K=8$  subscores (each 0–3), and **HAMD-13** is a clinician-rated scale with  $K=13$  text-assessable subscores (excluding 4 items requiring behavioral observation). See Appendix B for detailed subscore definitions.

### 3.2 MTSP

#### 3.2.1 Text Preprocessing

Clinical interview transcripts are first segmented into utterances using a hierarchical strategy: (1) split by newlines, (2) split by speaker tags (e.g., “doctor”, “patient”), (3) split by sentence delimiters, or (4) split by fixed-length chunks if no natural boundaries exist. Each utterance is then cleaned by removing non-text tags, filler words, and repeated characters. This preprocessing step is crucial for handling long transcripts.

#### 3.2.2 Base Language Model Encoding

The framework architecture is shown in Fig.3(b). First, we use domain-specific pre-trained BERT models as our base encoder to process the segmented utterances. For the English E-DAIC dataset, we employ MentalBERT (Ji et al., 2022), pre-trained on mental health-related text. For the Chinese HAMD-13 datasets (CIDH and PDCH), we use MedBERT-Chinese (Yang et al., 2021), pre-trained on Chinese medical corpora. For each utterance  $u_j$  in the transcript, we obtain a contextualized embedding  $\mathbf{h}_j = \text{BERT}(u_j)[\text{CLS}]$ , where  $\mathbf{h}_j \in \mathbb{R}^{d_{\text{bert}}}$  is the [CLS] token embedding that summarizes utterance  $u_j$ , and  $d_{\text{bert}} = 768$  is the hidden dimension. To improve efficiency, we encode all utterances across samples in a mini-batch, rather than encoding them individually. The sequence of utterance embeddings  $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_{N_i}] \in \mathbb{R}^{N_i \times d_{\text{bert}}}$  is fed into the Transformer encoder, where  $N_i$  is the number of utterances in transcript  $x_i$  (padded to the maximum sequence length in each batch).

#### 3.2.3 Shared Transformer Encoder

We apply a shared Transformer encoder to the sequence of utterance embeddings to capture tem-

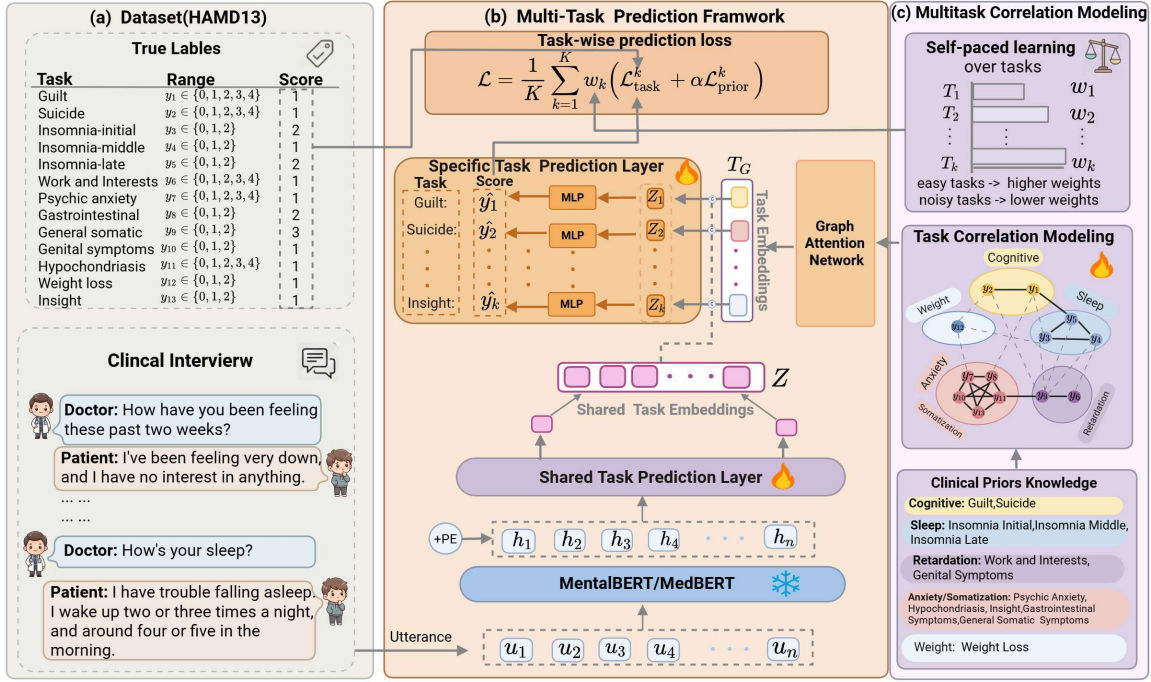


Figure 3: Overview of the proposed MTSP framework (illustrated with the HAMD-13 scale). (a) The illustration of a clinical interview in our dataset. (b) **Multi-task subscore prediction framework**. Each clinical interview is segmented into utterances, which are then encoded by a BERT-based encoder (MentalBERT / MedBERT) to obtain contextualized embeddings. Finally, these embeddings are passed through a shared task encoder and task-specific layers to jointly predict all subscores in a multi-task fashion. (c) **Multitask Correlation Modeling: 1. task correlation modeling** constructs a clinically informed task correlation graph over symptom subscores and incorporates clinical prior knowledge about symptom clusters as constraints to explicitly regularize relationships among subscores(tasks), while **2. task-level self-paced learning mechanism** adaptively reweights subscore losses according to their difficulty for making our model aware of the task difficulty.

poral dependencies and long-range interactions:  $\mathbf{Z} = \text{TransformerEncoder}(\mathbf{H}\mathbf{W}_{\text{proj}} + \mathbf{b}_{\text{proj}})$ , where  $\text{TransformerEncoder}(\cdot)$  consists of  $L_{\text{enc}}$  Transformer encoder layers,  $\mathbf{H} \in \mathbb{R}^{N \times d_{\text{bert}}}$  is the sequence of utterance embeddings,  $\mathbf{W}_{\text{proj}} \in \mathbb{R}^{d_{\text{bert}} \times d_{\text{enc}}}$  is a learnable projection matrix that maps BERT embeddings to the Transformer hidden dimension  $d_{\text{enc}}$ , and  $\mathbf{b}_{\text{proj}} \in \mathbb{R}^{d_{\text{enc}}}$  is the corresponding bias term that capture both local utterance semantics and global dialogue context. We then apply global average pooling over the sequence dimension to obtain a multi-task shared representation  $\mathbf{z} = \frac{1}{N} \sum_{n=1}^N \mathbf{Z}_n$ , where  $\mathbf{Z}_n \in \mathbb{R}^{d_{\text{enc}}}$  denotes the  $n$ -th row of  $\mathbf{Z}$  (i.e., the representation of the  $n$ -th utterance).

### 3.2.4 Multitask Correlation Modeling

To explicitly model the structured relationships among clinical subscores, we construct a **task correlation graph** (shown in Fig.3(c)) where each node represents a subscore and edges encode symptom correlations based on clinical prior knowledge. We employ a Graph Attention Network

(GAT) (Velickovic et al., 2018) to learn task embeddings that capture these relationships.

The task correlation graph is constructed based on clinically validated symptom clusters (3 clusters for PHQ-8, 5 for HAMD-13; see Appendix C for details). The graph includes intra-cluster edges (weight  $\omega_{\text{intra}}$ , initialized to 1.0) and inter-cluster edges (weight  $\omega_{\text{inter}}$ , initialized to 0.6), which are learned during training to adapt to data while respecting clinical priors. The graph structure includes two types of edges with learnable weights parameterized from clinical priors.

- **Intra-cluster edges:** the connected edges within each symptom cluster, reflecting higher correlations among related symptoms, with weight  $\omega_{\text{intra}}$  (initialized to 1.0 to reflect strong within-cluster correlations).
- **Inter-cluster edges:** the edges between clusters based on clinical correlations (e.g., Core Depression  $\leftrightarrow$  Cognitive Function, Sleep  $\leftrightarrow$  Concentration), with weight  $\omega_{\text{inter}}$  (initialized to 0.6 based on empirical analysis).

**Task Embedding Learning.** We maintain learnable task embeddings  $\mathbf{E} \in \mathbb{R}^{K \times d_e}$ , where  $K$  is the number of subscores and  $d_e$  is the task-embedding dimension. After projecting  $\mathbf{E}$  and the pooled text representation  $\mathbf{z} \in \mathbb{R}^d$  to the GAT hidden size  $d_g$ , we propagate them over the clinically informed task correlation graph using a series of GAT layers,  $\mathbf{T}_G = \text{GAT}(\mathbf{E}, \mathbf{z}, \mathcal{G}; \omega_{\text{intra}}, \omega_{\text{inter}})$ , where  $\mathcal{G}$  is the task correlation graph and  $\omega_{\text{intra}}, \omega_{\text{inter}}$  control intra- and inter-group edge weights. The resulting  $\mathbf{T}_G \in \mathbb{R}^{K \times d_g}$  encodes both task-specific information and cross-task dependencies.

### 3.2.5 Multi-Task Prediction

For each subscore  $k \in \{1, \dots, K\}$ , we first obtain a task-specific feature by concatenating the pooled text representation and the projected GAT-enhanced task embedding as  $\mathbf{z}_k = [\mathbf{z}; \text{Proj}_t(\mathbf{t}_k)]$ , where  $\mathbf{z} \in \mathbb{R}^{d_z}$  is the shared task embedding,  $\mathbf{t}_k = \mathbf{T}_G[k] \in \mathbb{R}^{d_g}$  is the embedding for subscore  $k$  obtained from the task correlation graph, and  $\text{Proj}_t(\cdot)$  is a learnable linear projection applied only to  $\mathbf{t}_k$ . For continuous score prediction, we normalize predictions to  $[0, 1]$  during training and denormalize to the valid range  $[0, C_k]$  in inference, where  $C_k$  is the maximum score for subscore  $k$ . To this end, the subscore is then predicted as  $\hat{y}_k = C_k \cdot \sigma(\mathbf{w}_k^\top \mathbf{z}_k + b_k)$ , where  $\mathbf{w}_k$  and  $b_k$  are the parameters of the task-specific head.

## 3.3 Objective Function

### 3.3.1 Subscore Prediction Loss

The primary supervision signal is the loss on predictions from each subscore. For sample  $i$  and subscore  $k$ , let  $y_k^{(i)}$  and  $\hat{y}_k^{(i)}$  denote the ground-truth and predicted scores, respectively. We use the mean squared error between  $\hat{y}_k^{(i)}$  and  $y_k^{(i)}$  as the prediction loss, and at the end of each epoch  $e$  we compute the average loss for each subscore  $k \in \{1, \dots, K\}$  over the training set as  $\bar{\mathcal{L}}_k^{(e)} = \frac{1}{M} \sum_{i=1}^M (\hat{y}_k^{(i)} - y_k^{(i)})^2$ , where  $M$  is the number of training samples.

### 3.3.2 Prior Loss

Given the task correlation graph, we incorporate a constraint loss that encourages predictions of highly correlated subscores within each cluster to be similar across all training examples. Given a training set with  $M$  examples, the prior loss is de-

defined as

$$\mathcal{L}_{\text{prior}} = \sum_{i=1}^M \sum_{c \in \mathcal{C}} \gamma_c \sum_{(k, k') \in \mathcal{P}_c} w_{kk'}^{(c)} (\hat{y}_{ik} - \hat{y}_{ik'})^2, \quad (1)$$

where  $\mathcal{C}$  is the set of symptom clusters,  $\mathcal{P}_c$  is the set of subscore pairs  $(k, k')$  within cluster  $c$ ,  $\hat{y}_{ik}$  and  $\hat{y}_{ik'}$  are the predicted scores for subscores  $k$  and  $k'$  of example  $i$ ,  $w_{kk'}^{(c)} \in [0, 1]$  is the correlation weight between subscores  $k$  and  $k'$  within cluster  $c$  (derived from clinical literature and empirical analysis), and  $\gamma_c > 0$  controls the relative importance of cluster.

### 3.3.3 Task-Level Self-Paced Learning

In practice, different subscores exhibit substantially different learning difficulty: items such as *Depressed* and *Concentrating* are relatively easy to optimize because patients usually respond to these questions explicitly and they are mentioned frequently in the dialogue, whereas *Appetite* is harder since eating-related information is rarely discussed directly, and *Suicide* is even more challenging because cues about suicidal ideation are often expressed implicitly or euphemistically, leading to noisier supervision. To address this, following self-paced learning (Kumar et al., 2010), we propose a **task-level** curriculum strategy that dynamically adjusts the contribution of each task based on its learning difficulty.

**Task Weights.** Given the average task-level loss  $\bar{\mathcal{L}}_k^{(e)}$ , we assign a self-paced weight  $w_k^{(e)}$  to each task. Tasks with loss below the threshold receive positive weights proportional to their easiness, while tasks with loss exceeding the threshold are temporarily excluded from training.

$$w_k^{(e)} = \begin{cases} 1 - \frac{\bar{\mathcal{L}}_k^{(e)}}{\lambda_{\text{spl}}^{(e)}}, & \text{if } \bar{\mathcal{L}}_k^{(e)} < \lambda_{\text{spl}}^{(e)}, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where  $\lambda_{\text{spl}}^{(e)}$  is the pace parameter at epoch  $e$ .

**Pace Schedule.** The pace parameter  $\lambda_{\text{spl}}$  controls the curriculum progression and increases linearly over epochs:  $\lambda_{\text{spl}}^{(e)} = \lambda_{\text{min}} + \frac{e}{E} \cdot (\lambda_{\text{max}} - \lambda_{\text{min}})$ , where  $E$  is the total number of epochs. We set  $\lambda_{\text{min}} = 0.5$  and  $\lambda_{\text{max}} = 2.0$ . Early in training, only easy tasks (with small loss) are learned, and as training progresses, harder tasks are gradually incorporated. The task weights can be further smoothed using a logarithmic pace function.

### 3.3.4 Overall Objective.

The overall training objective combines: (1) sub-score prediction losses weighted by task-level self-paced weights  $w_k^{(e)}$ , and (2) the medical domain prior constraint loss  $\mathcal{L}_{\text{prior}}$ . The task correlation graph with learnable edge weights ( $\omega_{\text{intra}}$  and  $\omega_{\text{inter}}$ ) is learned jointly via the GAT, while the cluster constraint loss penalizes deviations among subscores within the same clinically defined clusters. The overall objective function is defined as:

$$\mathcal{L} = \frac{1}{M} \sum_{i=1}^M \frac{\sum_{k=1}^K w_k^{(e)} (\hat{y}_k^{(i)} - y_k^{(i)})^2}{\sum_{k=1}^K w_k^{(e)}} + \beta \mathcal{L}_{\text{prior}}, \quad (3)$$

where  $\beta > 0$  is a hyperparameter controlling the weight of the constraint loss. The normalization by  $\sum_{k=1}^K w_k^{(e)}$  ensures stable gradient magnitudes as the number of active subscores changes during curriculum learning.

## 4 Experiments

### 4.1 Datasets

We evaluate MTSP on three clinical interview datasets, i.e., public **E-DAIC** (Ringeval et al., 2019), **PDCH** (Cao et al., 2025) and private **CIDH (Clinical interview dialogue based on the Hamilton Depression Rating Scale)**, covering both PHQ-8 and HAMD-13 scales. Dataset statistics and experimental settings are provided in Appendix A and E.

### 4.2 Experimental Results

Table 1 reports results on the E-DAIC dataset, including both development and test performance. All results are averaged over multiple random seeds to ensure statistical robustness.

**Comparison with the state of the art.** On E-DAIC (PHQ-8), MTSP (Full) achieves a test MAE of 3.48, outperforming all prior methods. Compared to the best prior work PTTSD-seq2seq (Schmidt et al., 2025) (MAE 3.85), MTSP achieves a 9.6% relative improvement. MTSP also outperforms the Qwen3-14B zero-shot baseline (MAE 5.39, 35.4% improvement), demonstrating that our compact model, which employs only 12 transformer layers (with merely 2 being trainable) under the multi-task subscore supervision, is more effective than relying solely on large language model prompting. **Ablation analysis.** We conduct ablation experiments across all three datasets (Table 2) to assess each component’s contribution.

| Model  | Dev         |             | Test        |             |
|--|-------------|-------------|-------------|-------------|
|  | MAE↓        | RMSE↓       | MAE↓        | RMSE↓       |
| <i>LLM Baseline:</i>                         |             |             |             |             |
| Qwen3-14B (Zero-Shot)                        | 5.43        | 5.49        | 5.39        | 5.47        |
| <i>Other Work on E-DAIC:</i>                 |             |             |             |             |
| Ray et al. (2019)                            | –           | 4.37        | 4.02        | 4.73        |
| Sadeghi et al. (2023)                        | 3.65        | 5.27        | 4.26        | 5.37        |
| Sadeghi et al. (2024) (Pr3 + Whisper)        | 3.17        | 4.51        | 4.22        | 5.07        |
| Schmidt et al. (2025) (s2o-MentalBERT)       | 3.55        | 4.58        | 4.18        | 5.23        |
| Schmidt et al. (2025) (s2s-all-MiniLM-L6-v2) | 3.47        | 4.57        | 3.85        | 4.52        |
| <i>Our Methods:</i>                          |             |             |             |             |
| <b>MTSP (Full)</b>                           | <b>2.94</b> | <b>4.18</b> | <b>3.48</b> | <b>4.57</b> |

Table 1: Main results on the E-DAIC (PHQ-8) dataset. We report mean absolute error (MAE) and root mean squared error (RMSE) on the development and test sets.

Comparing individual contributions, the medical domain prior constraint (Prior) provides the largest improvement: on E-DAIC, adding Prior alone reduces MAE from 3.93 to 3.58 (8.9% improvement), while adding T-SPL alone reduces MAE to 3.76 (4.3% improvement). Notably, combining T-SPL and Prior Loss yields additional gains (MAE 3.48 vs. 3.58 / 3.76), demonstrating that these two modules are complementary because they address different aspects of the learning problem: Prior Loss constrains the prediction space based on medical knowledge, while T-SPL optimizes the learning dynamics by scheduling tasks according to difficulty. Moreover, we observe that Prior alone achieves MAE 3.65 (1.6% improvement) and MAE 3.66 (3.9% improvement) on the CIDH and PDCH dataset, respectively. The larger improvements on the PDCH suggest that medical domain priors are especially valuable when training data is limited (PDCH has only 99 samples).

### 4.3 Analysis

We conduct detailed analyses to investigate (1) task difficulty variation and the effectiveness of self-paced learning, and (2) the ability to capture fine-grained symptom.

#### 4.3.1 Task Difficulty Analysis

To understand why task-level self-paced learning is effective, we analyze per-task performance on PHQ-8. Figure 4 compares the task-level MAE of Qwen3-14B (zero-shot), MTSP without T-SPL, and MTSP with T-SPL on the E-DAIC test set.

The results reveal a significant difficulty dispar-

| Model Variant                        | E-DAIC (PHQ-8) |             | CIDH (HAMD-13) |             | PDCH (HAMD-13) |             |
|--------------------------------------|----------------|-------------|----------------|-------------|----------------|-------------|
|                                      | MAE↓           | RMSE↓       | MAE↓           | RMSE↓       | MAE↓           | RMSE↓       |
| <i>Our Ablation (MTSP Variants):</i> |                |             |                |             |                |             |
| MTSP (Total Score)                   | 4.02           | 4.76        | 3.72           | 4.65        | 4.58           | 5.47        |
| MTSP (Subscore)                      | 3.93           | 4.91        | 3.71           | 4.72        | 3.81           | 4.79        |
| + Task-Level SPL                     | 3.76           | 4.82        | 3.66           | 4.64        | 3.74           | 4.73        |
| + Prior                              | 3.58           | 4.55        | 3.65           | 4.57        | 3.66           | 4.68        |
| <b>MTSP (Full)</b>                   | <b>3.48</b>    | <b>4.57</b> | <b>3.63</b>    | <b>4.55</b> | <b>3.54</b>    | <b>4.56</b> |

Table 2: Ablation study of MTSP variants across all three datasets. We report mean absolute error (MAE) and root mean squared error (RMSE) of the total score on E-DAIC (PHQ-8), CIDH (HAMD-13), and PDCH (HAMD-13).




Figure 4: **The Task-level MAE comparison on PHQ-8 (E-DAIC test set).** Each subscore shows three bars: LLM (Qwen3-14B zero-shot, red), MTSP w/o SPL (blue), and MTSP w/ SPL (green). Tasks are annotated with difficulty levels based on LLM performance: *Easy* (MAE < 0.65), *Medium* (0.65–0.80), and *Hard* (> 0.80). T-SPL consistently improves performance across most subscores, especially on hard tasks like *Appetite* and *Concentrating*.

ity across tasks. Depressed is classified as “Easy” (MAE 0.57) because depressed mood is explicitly discussed in interviews. In contrast, Appetite (MAE 1.05) and Concentrating (MAE 0.89) are “Hard” tasks because their symptoms are often expressed implicitly and in ambiguous ways rather than being mentioned directly. By investigating the learning procedure of T-SPL, we find that it learns tasks in the order of increasing task difficulty. In early epochs (pace  $\lambda = 0.2$ ), only easy tasks receive non-zero weights: *Depressed* ( $w = 0.85$ ), *NoInterest* ( $w = 0.78$ ), and *Moving* ( $w = 0.72$ ) are prioritized, while hard tasks like *Appetite* ( $w = 0$ ) and *Concentrating* ( $w = 0$ ) are excluded. As training progresses and  $\lambda$  increases linearly, medium-difficulty tasks (*Tired*, *Failure*, *Sleep*) are incorporated around epoch 5–10, and hard tasks enter the curriculum after epoch 10. By epoch 15, all tasks have comparable weights ( $w \approx 0.7$ – $0.9$ ), allowing the model to refine predictions across all subscores. This progressive curriculum prevents gradient interference from difficult tasks during early training, reducing average MAE from 0.70 to 0.68. The improvement is es-

| Subscore       | MTSP MAE↓   | Qwen3-14B MAE↓ | Improv.      |
|----------------|-------------|----------------|--------------|
| NoInterest     | 0.61        | 0.91           | 32.5%        |
| Depressed      | 0.62        | 0.77           | 19.5%        |
| Sleep          | 0.86        | 1.08           | 20.0%        |
| Tired          | 0.77        | 1.04           | 25.7%        |
| Appetite       | 0.79        | 1.42           | 44.4%        |
| Failure        | 0.77        | 0.93           | 16.6%        |
| Concentrating  | 0.78        | 1.09           | 27.8%        |
| Moving         | 0.54        | 0.95           | 43.8%        |
| <b>Average</b> | <b>0.72</b> | <b>1.04</b>    | <b>30.5%</b> |

Table 3: The comparison between MTSP and Qwen3-14B (zero-shot) for each subscore on E-DAIC (PHQ-8) test set. MTSP outperforms the LLM on all 8 subscores, with an average improvement of 30.5%.

pecially obvious on hard tasks, i.e., *Appetite*: 0.65  $\rightarrow$  0.60, *Concentrating* : 0.67  $\rightarrow$  0.62.

### 4.3.2 Subscore analysis

To evaluate whether MTSP enable fine-grained subscore prediction by capturing individual symptom patterns, especially for patients with identical total scores, we conduct the following analyses on the E-DAIC and CIDH test sets.

**Subscore prediction performance.** Table 3 reports the comparison between MTSP and Qwen3-14B (zero-shot) for each subscore on the E-DAIC test set. MTSP achieves an average subscore MAE of 0.72, significantly outperforming the LLM (1.04). MTSP shows substantial improvements across all 8 PHQ-8 subscores, with particularly large gains on difficult items like *Appetite* (0.60 vs. 1.05, 43% improvement) and *Concentrating* (0.62 vs. 0.89, 30% improvement). This shows that subscore prediction under the guidance of multi-task correlation is more effective in fine-grained symptom assessment.

**Subscore distribution analysis.** To validate that MTSP capture fine-grained subscores, Figure 5 shows the subscore distributions for each

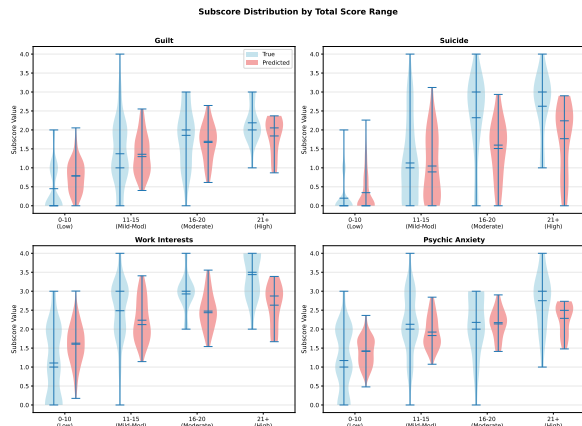




Figure 5: **Subscore distribution comparison by total score range.** Violin plots showing true (blue) vs. predicted (red) subscore distributions for four representative subscores across different total score ranges. The similar distribution shapes indicate that MTSP preserves symptom heterogeneity within each severity group.

total score ranges using violin plots. For each total score range, we compare the true subscore distributions (blue) with predicted distributions (red). The similarity in distribution shapes indicate that MTSP captures the underlying symptom patterns rather than merely predicting average values.

To further illustrate how MTSP learns clinically meaningful representations, we visualize the learned features using t-SNE (Van der Maaten and Hinton, 2008). We extract the final-layer features (before the output heads) for all test samples and project them to 2D space. Fig. 6 shows the data distributions obtained by MTSP (Total Score) and MTSP (Subscore). Patients A ( $\blacktriangle$ ) and B ( $\blacktriangle$ ) mentioned in the introduction section have the same total score. In Fig. 6a, for MTSP (Total Score), A and B are mapped close together because they share the same total score. In contrast, MTSP (Subscore) correctly separates them, recognizing their clinically distinct symptom profiles. This validates that MTSP (Subscore) learns representations that separate patients with different symptom patterns even when total scores are the same—a capability absent in total score supervision. Figure 6b provides complementary evidence by coloring points according to subscore pattern similarity. We project each patient’s 13-dim subscore vector to a 1D score via PCA, where similar colors indicate similar symptom profiles. In MTSP (Total Score), the points within the same clusters exhibit varying colors. In contrast, MTSP (Subscore) shows that spatially close points share similar colors, indicating that pa-



(a) Colored by total score. Two patients with a total score of 15 are highlighted: A ( $\blacktriangle$ ) and B ( $\blacktriangle$ ) have different symptom patterns.



(b) Colored by subscore pattern similarity: similar colors indicate similar symptom profiles.

Figure 6: **Data distribution visualization via t-SNE.**

tients with similar symptom profiles are mapped to nearby regions in the feature space. This confirms that multi-task subscore supervision enables the model to learn symptom-aware and discriminative representations that preserve patient heterogeneity, enabling fine-grained clinical assessment that captures individual patient differences.

## 5 Conclusion

We have proposed MTSP, a multi-task subscore prediction framework for automated depression assessment that shifts the paradigm from coarse total score regression to fine-grained subscore modeling. MTSP comprises three key components: (1) a multi-task prediction framework that jointly predicts all subscores, providing richer supervision signals than single total score prediction; (2) a task correlation graph with prior constraints that explicitly models clinically validated symptom relationships; and (3) task-level self-paced learning that adaptively handles varying subscore difficulty during training. Experiments on E-DAIC (PHQ-8), CIDH (HAMD-13), and PDCH (HAMD-13) datasets demonstrate that MTSP achieves state-of-the-art performance, outperforming both total score regression baselines and LLM scoring. These results suggest that fine-grained subscore modeling via multi-task learning is a promising direction for building clinically meaningful automated depression assessment systems.

## 537 Limitations

538 Our evaluation is limited to three datasets (E-  
539 DAIC, CIDH, PDCH) from specific clinical con-  
540 texts. While CIDH provides a relatively large  
541 sample (1,689 interviews), PDCH contains only  
542 99 samples, and E-DAIC focuses on virtual inter-  
543 viewer interactions rather than real doctor-patient  
544 consultations. Future work should validate on  
545 larger, more diverse populations across different  
546 healthcare systems.

547 We assess only 13 text-based HAMD subscores,  
548 excluding 4 items (Depressed Mood, Retardation,  
549 Agitation, Depersonalization) that require direct  
550 behavioral observation. Extending MTSP to mul-  
551 timodal inputs would enable complete HAMD-  
552 17 assessment but requires careful handling of  
553 video/audio data.

554 The model has not undergone prospective clin-  
555 ical validation. Real-world deployment would re-  
556 quire collaboration with psychiatrists to evaluate  
557 failure modes, examine potential biases across pa-  
558 tient subgroups, and establish appropriate clinical  
559 workflows. Additionally, HAMD annotations them-  
560 selves have known inter-rater variability, and our  
561 model may inherit these annotation inconsistencies.

562 Finally, we emphasize that MTSP is designed as  
563 a clinical decision support tool, not a replacement  
564 for professional psychiatric evaluation. Automated  
565 depression assessment raises important questions  
566 about privacy, accountability, and the appropriate  
567 role of AI in mental healthcare.

## 568 References

569 R Michael Bagby, Andrew G Ryder, Deborah R  
570 Schuller, and Margarita B Marshall. 2004. The  
571 Hamilton Depression Rating Scale: has the gold stan-  
572 dard become a lead weight? *American Journal of*  
573 *Psychiatry*, 161(12):2163–2177.

574 Baichuan-M2 Team, Chengfeng Dou, Chong Liu, Fan  
575 Yang, Fei Li, Jiyuan Jia, Mingyang Chen, Qiang Ju,  
576 Shuai Wang, Shunya Dang, Tianpeng Li, Xiangrong  
577 Zeng, Yijie Zhou, Chenzheng Zhu, Da Pan, Fei Deng,  
578 Guangwei Ai, Guosheng Dong, Hongda Zhang, and  
579 15 others. 2025. *Baichuan-m2: Scaling medical ca-*  
580 *patibility with large verifier system. arXiv preprint*  
581 *arXiv:2509.02208*.

582 Guanqun Bi, Zhuang Chen, Zhoufu Liu, Hongkai Wang,  
583 Xiyao Xiao, Yuqiang Xie, Wen Zhang, Yongkang  
584 Huang, Yuxuan Chen, Libiao Peng, and Minlie  
585 Huang. 2025. *MAGI: Multi-agent guided interview*  
586 *for psychiatric assessment. In Findings of the Asso-*  
587 *ciation for Computational Linguistics: ACL 2025,*

pages 24898–24921, Vienna, Austria. Association  
for Computational Linguistics.

Amanda M. Buch and Conor Liston. 2021. *Dissecting*  
*diagnostic heterogeneity in depression by integrating*  
*neuroimaging and genetics. Neuropsychopharmacol-*  
*ogy*, 46(1):156–175.

Iain M Cameron, John R Crawford, Kenny Lawton, and  
Ian C Reid. 2008. Psychometric comparison of PHQ-  
9 and HADS for measuring depression severity in  
primary care. *British Journal of General Practice*,  
58(546):32–36.

Peng Cao, Yifan Zhang, Lei Wang, and 1 others. 2025.  
*A multimodal depression consultation dataset of*  
*speech and text with HAMD-17 assessments. Scien-*  
*tific Data*, 12:89.

Glen Coppersmith, Mark Dredze, and Craig Harman.  
2014. Quantifying mental health signals in twitter.  
*In Proceedings of the Workshop on Computational*  
*Linguistics and Clinical Psychology: From Linguistic*  
*Signal to Clinical Reality*, pages 51–60.

Munmun De Choudhury, Michael Gamon, Scott Counts,  
and Eric Horvitz. 2013. Predicting depression via so-  
cial media. *In Proceedings of the International AAAI*  
*Conference on Web and Social Media*, volume 7,  
pages 128–137.

Robert D Gibbons, Duncan C Clark, and David J Kupfer.  
1993. Exactly what does the Hamilton Depression  
Rating Scale measure? *Journal of Psychiatric Re-*  
*search*, 27(3):259–273.

George Gkotsis, Anika Oellrich, Sumithra Velupillai,  
Maria Liakata, Tim JP Hubbard, Richard JB Dob-  
son, and Rina Dutta. 2017. Characterisation of men-  
tal health conditions in social media using informed  
deep learning. *Scientific Reports*, 7(1):45141.

Max Hamilton. 1960. A rating scale for depression.  
*Journal of Neurology, Neurosurgery, and Psychiatry*,  
23(1):56–62.

Shaoxiong Ji, Tianlin Zhang, Luna Ansari, Jie Fu,  
Prayag Tiwari, and Erik Cambria. 2022. Mental-  
BERT: Publicly available pretrained language mod-  
els for mental healthcare. *In Proceedings of the Lan-*  
*guage Resources and Evaluation Conference*, pages  
7184–7190.

Kurt Kroenke, Robert L Spitzer, and Janet BW Williams.  
2001. The PHQ-9: validity of a brief depres-  
sion severity measure. *Journal of General Internal*  
*Medicine*, 16(9):606–613.

Kurt Kroenke, Tara W Strine, Robert L Spitzer,  
Janet BW Williams, Joyce T Berry, and Ali H Mok-  
dad. 2009. The PHQ-8 as a measure of current de-  
pression in the general population. *Journal of Affec-*  
*tive Disorders*, 114(1-3):163–173.

|     |   |   |     |
|-----|---|---|-----|
| 640 | M Pawan Kumar, Benjamin Packer, and Daphne Koller.                  | Petar Velickovic, Guillem Cucurull, Arantxa Casanova,   | 696 |
| 641 | 2010. Self-paced learning for latent variable mod-                  | Adriana Romero, Pietro Lio, and Yoshua Bengio.  | 697 |
| 642 | els. In <i>Advances in Neural Information Processing</i>            | 2018. Graph attention networks. In <i>International</i>   | 698 |
| 643 | <i>Systems</i> , volume 23.   | <i>Conference on Learning Representations</i> .   | 699 |
| 644 | Jianning Li, Amin Dada, Jens Kleesiek, and Jan Egger.               | World Health Organization. 2025. Depressive disorder  | 700 |
| 645 | 2024. <i>Chatgpt in healthcare: A taxonomy and sys-</i>             | (depression). <a href="https://www.who.int/news-room/fact-sheets/detail/depression">https://www.who.int/news-room/fact-sheets/detail/depression</a> . Accessed: | 701 |
| 646 | <i>tematic review</i> . <i>Computer Methods and Programs in</i>     | 2025-12-15.   | 702 |
| 647 | <i>Biomedicine</i> , 245:108013.                                    |   | 703 |
| 648 | Aishik Mandal, Dana Atzil-Slonim, Tamar Solorio,                    | Feihong Yang, Xuwen Wang, and Jiao Li. 2021. <i>Ex-</i>   | 704 |
| 649 | and Iryna Gurevych. 2025. <i>Enhancing depression</i>               | <i>ploration and research on applying BERT to chinese</i>   | 705 |
| 650 | <i>detection via question-wise modality fusion</i> . In <i>Pro-</i> | <i>clinical natural language processing</i> . GitHub repos-   | 706 |
| 651 | <i>ceedings of the 10th Workshop on Computational</i>               | <i>itory</i> . Original title in Chinese: “BERT”. Accessed:   | 707 |
| 652 | <i>Linguistics and Clinical Psychology (CLPsych 2025)</i> ,         | 2026-01-04.   | 708 |
| 653 | pages 44–61, Albuquerque, New Mexico. Associa-                      |   |     |
| 654 | tion for Computational Linguistics.                                 | Linhai Zhang, Ziyang Gao, Deyu Zhou, and Yulan He.  | 709 |
| 655 | Anupama Ray, Siddharth Kumar, Rutvik Reddy, Pre-                    | 2025. <i>Explainable depression detection in clinical in-</i>   | 710 |
| 656 | rana Mukherjee, and Ritu Garg. 2019. <i>Multi-level</i>             | <i>terviews with personalized retrieval-augmented gen-</i>  | 711 |
| 657 | <i>attention network using text, audio and video for</i>            | <i>eration</i> . In <i>Findings of the Association for Compu-</i>   | 712 |
| 658 | <i>depression prediction</i> . In <i>Proceedings of the 9th In-</i> | <i>tational Linguistics: ACL 2025</i> , pages 9927–9944,  | 713 |
| 659 | <i>ternational on Audio/Visual Emotion Challenge and</i>            | Vienna, Austria. Association for Computational Lin-   | 714 |
| 660 | <i>Workshop (AVEC '19)</i> , pages 81–88, New York, NY,             | guistics.   | 715 |
| 661 | USA. Association for Computing Machinery.                           |   |     |
| 662 | Fabien Ringeval, Björn Schuller, Michel Valstar,                    | Taolin Zhang, Zerui Cai, Chengyu Wang, Minghui Qiu,   | 716 |
| 663 | Nicholas Cummins, Roddy Cowie, Leili Tavabi, Max-                   | Bite Yang, and Xiaofeng He. 2021. <i>SMedBERT:</i>  | 717 |
| 664 | imilian Schmitt, and 1 others. 2019. <i>AVEC 2019</i>               | <i>A knowledge-enhanced pre-trained language model</i>  | 718 |
| 665 | <i>workshop and challenge: State-of-mind, detecting</i>             | <i>with structured semantics for medical text mining.</i>   | 719 |
| 666 | <i>depression with AI, and cross-cultural affect recogni-</i>       | <i>arXiv preprint arXiv:2108.08983</i> .  | 720 |
| 667 | <i>tion</i> . In <i>Proceedings of the 9th International on Au-</i> |   |     |
| 668 | <i>dio/Visual Emotion Challenge and Workshop (AVEC</i>              | Xianbing Zhao, Yiqing Lyu, Di Wang, and Buzhou  | 721 |
| 669 | <i>'19)</i> , pages 3–12. ACM.                                      | Tang. 2025. <i>Predicting depression in screening in-</i>   | 722 |
| 670 | Misha Sadeghi, Bernhard Egger, Reza Agahi, Robert                   | <i>terviews from interactive multi-theme collaboration</i> . In   | 723 |
| 671 | Richer, Klara Capito, Lydia Helene Rupp, Lena                       | <i>Findings of the Association for Computational Lin-</i>   | 724 |
| 672 | Schindler-Gmelch, Matthias Berking, and Bjoern M.                   | <i>guistics: ACL 2025</i> , pages 23025–23035, Vienna,  | 725 |
| 673 | Eskofier. 2023. <i>Exploring the capabilities of a lan-</i>         | Austria. Association for Computational Linguistics.   | 726 |
| 674 | <i>guage model-only approach for depression detection</i>           |   |     |
| 675 | <i>in text data</i> . In <i>2023 IEEE EMBS International Con-</i>   | <b>A Dataset Details</b>  | 727 |
| 676 | <i>ference on Biomedical and Health Informatics (BHI)</i> ,         | <b>E-DAIC (Extended Distress Analysis Interview</b>   | 728 |
| 677 | pages 1–5, Pittsburgh, PA, USA. IEEE.                               | <b>Corpus)</b> . E-DAIC (Ringeval et al., 2019) is a pub-   | 729 |
| 678 | Misha Sadeghi, Robert Richer, Bernhard Egger, Lena                  | licly available dataset from the AVEC 2019 chal-  | 730 |
| 679 | Schindler-Gmelch, Lydia Helene Rupp, Farnaz                         | lenge, containing clinical interviews with PHQ-8  | 731 |
| 680 | Rahimi, Matthias Berking, and Bjoern M Eskofier.                    | annotations. The dataset includes 275 audio record-   | 732 |
| 681 | 2024. <i>Harnessing multimodal approaches for depres-</i>           | ings and transcripts of structured psychiatric in-  | 733 |
| 682 | <i>sion detection using large language models and facial</i>        | terviews conducted by a virtual interviewer (Ellie).  | 734 |
| 683 | <i>expressions</i> . <i>npj Mental Health Research</i> , 3(1):66.   | Each interview is annotated with 8 PHQ-8 subscore   | 735 |
| 684 | Fabian Schmidt, Seyedehmoniba Ravan, and Vladimir                   | scores (each 0–3) and a total score (0–24). We use  | 736 |
| 685 | Vlassov. 2025. Probabilistic textual time series de-                | the standard train/validation/test split (163/56/56   | 737 |
| 686 | pression detection. <i>arXiv preprint arXiv:2511.04476</i> .        | samples) for evaluation.  | 738 |
| 687 | ArXiv:2511.04476 [cs.CL].   | <b>CIDH (Clinical interview dialogue based on</b>   | 739 |
| 688 | Laurens Van der Maaten and Geoffrey Hinton. 2008.                   | <b>the Hamilton Depression Rating Scale)</b> . This   | 740 |
| 689 | Visualizing data using t-SNE. <i>Journal of Machine</i>             | dataset contains 1,689 clinical interviews with   | 741 |
| 690 | <i>Learning Research</i> , 9(86):2579–2605.                         | HAMD-13 annotations, collected in collaboration   | 742 |
| 691 | Floris van Steijn, Gizem Sogancioglu, and Heysem                    | with Nanjing Brain Hospital (Nanjing Medical Uni-   | 743 |
| 692 | Kaya. 2022. <i>Text-based interpretable depression</i>              | versity Affiliated Brain Hospital). The dataset in-   | 744 |
| 693 | <i>severity modeling via symptom predictions</i> . In <i>Pro-</i>   | cludes 271 male and 621 female participants, with   | 745 |
| 694 | <i>ceedings of the 24th ACM International Conference</i>            | ages ranging from 11 to 54 years. Each inter-   | 746 |
| 695 | <i>on Multimodal Interaction (ICMI)</i> , pages 139–147.            | view is a structured psychiatric evaluation tran-   | 747 |
|     |   | script conducted by trained clinicians. Due to the  | 748 |

excessive length of raw clinical transcripts (often exceeding 10,000 tokens), we use Baichuan-M2-32B (Baichuan-M2 Team et al., 2025) to generate structured summaries for each HAMD-13 subscore. Each summarized transcript contains concise descriptions relevant to all 13 HAMD subscores (see Appendix G for the summarization prompt). All data were collected with informed consent and anonymized. We use an approximate 70/15/15 split, resulting in 1,182/253/254 samples for train/validation/test. All models are trained on the CIDH training set, and hyperparameters are selected on the validation set.

**PDCH (Public Depression Consultation in Hospital).** PDCH (Cao et al., 2025) is a publicly available multimodal dataset containing clinical psychiatric interviews with HAMD-17 annotations. The dataset includes audio, video, and text transcripts of real doctor-patient consultations. We extract the 13 text-assessable tasks (excluding 4 items requiring direct behavioral observation: Depressed Mood, Retardation, Agitation, and Depersonalization/Derealization) for evaluation. We use a 70/15/15 split, resulting in 69/15/15 samples for train/validation/test.

| Statistic        | E-DAIC | CIDH    | PDCH    |
|------------------|--------|---------|---------|
| Scale            | PHQ-8  | HAMD-13 | HAMD-13 |
| Total samples    | 275    | 1,689   | 99      |
| Age range        | -      | 11–54   | 18–65   |
| Gender (M/F)     | -      | 271/621 | -       |
| Subscore         | 8      | 13      | 13      |
| Score range      | 0–24   | 0–36    | 0–36    |
| Avg. total score | -      | 10.2    | 11.4    |

Table 4: Summary of dataset statistics. E-DAIC uses PHQ-8 (8 subscales), while CIDH and PDCH use HAMD-13 (13 text-assessable subscales).

## B Subscore Definitions

**PHQ-8 Subscores** (each scored 0–3): (1) NoInterest, (2) Depressed, (3) Sleep, (4) Tired, (5) Appetite, (6) Failure, (7) Concentration, (8) Psychomotor.

**HAMD-13 Subscores:** (1) Guilt (0–4), (2) Suicide (0–4), (3) Insomnia-initial (0–2), (4) Insomnia-middle (0–2), (5) Insomnia-late (0–2), (6) Work and Interests (0–4), (7) Psychic Anxiety (0–4), (8) GI Symptoms (0–2), (9) Somatic Symptoms (0–2), (10) Genital Symptoms (0–2), (11) Hypochondriasis (0–4), (12) Weight Loss (0–2), (13) Insight (0–2).

## C Symptom Cluster Definitions

The task correlation graph is constructed based on clinically validated symptom clusters (Kroenke et al., 2001, 2009; Cameron et al., 2008; Bagby et al., 2004; Gibbons et al., 1993).

### PHQ-8 Clusters (3 clusters):

- *Core Depression:* Anhedonia, Depressed Mood, Worthlessness
- *Cognitive Function:* Concentration, Psychomotor
- *Somatic Symptoms:* Sleep Problems, Fatigue, Appetite Changes

### HAMD-13 Clusters (5 clusters):

- *Cognitive:* Guilt, Suicide
- *Sleep:* Insomnia-Initial, Insomnia-Middle, Insomnia-Late
- *Retardation:* Work & Interests, Genital Symptoms
- *Anxiety/Somatization:* Psychic Anxiety, GI Symptoms, Somatic Symptoms, Hypochondriasis, Insight
- *Weight:* Weight Loss

The graph structure includes two types of edges:

- **Intra-cluster edges:** Connect subscores within each symptom cluster, reflecting higher correlations among related symptoms. Initialized with weight  $\omega_{\text{intra}} = 1.0$ .
- **Inter-cluster edges:** Connect subscores between clusters based on clinical correlations (e.g., Core Depression  $\leftrightarrow$  Cognitive Function, Sleep  $\leftrightarrow$  Concentration). Initialized with weight  $\omega_{\text{inter}} = 0.6$ .

## D Evaluation Metrics

Let  $y_{i,k}$  and  $\hat{y}_{i,k}$  denote the ground-truth and predicted scores for sample  $i \in \{1, \dots, M\}$  and task  $k \in \{1, \dots, K\}$ .

- **Mean Absolute Error (MAE):** Our primary metric, defined as

$$\text{MAE} = \frac{1}{MK} \sum_{i=1}^M \sum_{k=1}^K |\hat{y}_{i,k} - y_{i,k}|.$$

- **Root Mean Squared Error (RMSE):** Defined as

$$\text{RMSE} = \sqrt{\frac{1}{MK} \sum_{i=1}^M \sum_{k=1}^K (\hat{y}_{i,k} - y_{i,k})^2},$$

which penalizes larger errors more heavily than MAE.

- **Task-level MAE:** MAE computed separately for each task  $k$ ,

$$\text{MAE}_k = \frac{1}{M} \sum_{i=1}^M |\hat{y}_{i,k} - y_{i,k}|,$$

which allows us to analyze which symptoms are easier or harder to predict.

- **Total Score MAE and RMSE:** Error metrics on the summed score per interview,

$$\text{MAE}_{\text{total}} = \frac{1}{M} \sum_{i=1}^M \left| \sum_{k=1}^K \hat{y}_{i,k} - \sum_{k=1}^K y_{i,k} \right|,$$

$$\text{RMSE}_{\text{total}} = \sqrt{\frac{1}{M} \sum_{i=1}^M \left( \sum_{k=1}^K \hat{y}_{i,k} - \sum_{k=1}^K y_{i,k} \right)^2},$$

reflecting error at the overall severity level.

- **Task Accuracy:** Exact-match accuracy, i.e., the proportion of task predictions that match the ground-truth label, averaged over all tasks and samples:

$$\text{Acc} = \frac{1}{MK} \sum_{i=1}^M \sum_{k=1}^K \mathbb{I}[\hat{y}_{i,k} = y_{i,k}].$$

## E Detailed Experimental Setup

### E.1 Training Setup

All models are trained on the respective training splits (E-DAIC for PHQ-8, CIDH and PDCH for HAMD-13), and we select hyperparameters based on validation performance. We use Adam optimizer with learning rate  $2 \times 10^{-4}$  and cosine annealing learning rate schedule with minimum learning rate  $1 \times 10^{-4}$ . The mini-batch size is 50. We train for 80 epochs with early stopping (patience=15) based on validation total score MAE. Experiments are repeated with multiple random seeds, and we report mean and standard deviation.

### E.2 Model Configuration

**Encoder.** For the English E-DAIC dataset, we use MentalBERT (Ji et al., 2022), a BERT model pre-trained on mental health-related text. For the Chinese HAMD-13 datasets (CIDH and PDCH), we use MedBERT-Chinese (Zhang et al., 2021), a BERT model pre-trained on Chinese medical corpora. The BERT embedding dimension is  $d_{\text{bert}}=768$ . All experiments are conducted on a single NVIDIA A100 80GB GPU.

**Transformer encoder.** MTSP uses  $L_{\text{enc}}=2$  Transformer encoder layers with hidden dimension  $d=200$ , 4 attention heads, feedforward dimension  $4d=800$ , and dropout rate 0.3.

**Task correlation graph GAT.** The task correlation graph module is configured as follows:

- Task embedding dimension:  $d_e=64$
- GAT hidden dimension:  $d_g=128$
- Number of GAT layers:  $L_{\text{gat}}=2$
- Number of attention heads:  $H=4$
- GAT dropout rate: 0.1
- LeakyReLU negative slope:  $\alpha=0.2$
- Edge weight initialization:  $\omega_{\text{intra}}=1.0$ ,  $\omega_{\text{inter}}=0.6$
- Fusion type: gated fusion

**Loss weights.** For PHQ-8, we set medical domain prior constraint weights  $\gamma_1=1.0$  (Core Depression),  $\gamma_2=0.5$  (Cognitive),  $\gamma_3=0.8$  (Somatic) and overall constraint weight  $\beta=0.01$ . For HAMD-13, we use five clusters with weights  $\gamma_1=1.0$  (Cognitive),  $\gamma_2=0.8$  (Sleep),  $\gamma_3=0.6$  (Retardation),  $\gamma_4=0.5$  (Anxiety/Somatization),  $\gamma_5=0.3$  (Weight) and  $\beta=0.01$ .

**Self-paced learning.** The pace parameter  $\lambda_{\text{spl}}$  increases linearly over epochs, with task weights computed using a linear pace function. We set  $\lambda_{\text{min}}=0.5$  and  $\lambda_{\text{max}}=2.0$ .

## F Hyperparameter Sensitivity

### F.1 Implementation Details

We implement MTSP in PyTorch with Hugging Face Transformers and run all experiments on NVIDIA GPUs. For E-DAIC (English), we use MentalBERT (Ji et al., 2022) as the base encoder; for HAMD-13 datasets (Chinese), we use

MedBERT-Chinese (Zhang et al., 2021). The Transformer encoder uses  $L = 2$  layers with hidden dimension  $d_{\text{tfm}} = 200$ , 4 attention heads, and dropout rate 0.3. The task correlation graph module uses learnable task embeddings with dimension  $d_e = 64$ , GAT hidden dimension  $d_g = 128$ ,  $L_{\text{gat}} = 2$  GAT layers with  $H = 4$  attention heads, and dropout rate 0.1. We optimize with Adam (learning rate  $2 \times 10^{-4}$ ) using mini-batch size 50, and train for 80 epochs with early stopping (patience=15) based on validation performance. For PHQ-8, we apply prior constraint with weights  $\gamma_1 = 1.0$ ,  $\gamma_2 = 0.5$ ,  $\gamma_3 = 0.8$  and constraint weight  $\beta = 0.01$ . Task-level self-paced learning uses linear lambda growth, where  $\lambda_{\text{min}} = 0.5$  and  $\lambda_{\text{max}} = 2.0$ .

We conducted sensitivity analysis on key hyperparameters:

| Hyperparameter                          | Range             | Optimal |
|---|-------------------|---------|
| <i>Transformer Encoder:</i>             |                   |         |
| Transformer layers ( $L_{\text{enc}}$ ) | 1, 2, 3           | 2       |
| Hidden dimension ( $d$ )                | 100, 200, 300     | 200     |
| Attention heads                         | 2, 4, 8           | 4       |
| <i>Task Correlation Graph GAT:</i>      |                   |         |
| Task embed dim ( $d_e$ )                | 32, 64, 128       | 64      |
| GAT hidden dim ( $d_g$ )                | 64, 128, 200      | 128     |
| GAT layers ( $L_{\text{gat}}$ )         | 1, 2, 3           | 2       |
| GAT heads ( $H$ )                       | 2, 4, 8           | 4       |
| $\omega_{\text{cross}}$                 | 0.4, 0.6, 0.8     | 0.6     |
| Fusion type                             | concat, add, gate | gate    |
| <i>Training:</i>                        |                   |         |
| Learning rate                           | 1e-4, 2e-4, 5e-4  | 2e-4    |
| Batch size                              | 16, 32, 50        | 50      |

Table 5: Hyperparameter sensitivity analysis for MTSP

## G Prompts

### G.1 Summarization Prompt (HAMD-13)

The following prompt generates structured summaries from raw clinical transcripts using Baichuan-M2-32B:

You are a professional psychiatrist. Analyze the clinical interview and extract information for HAMD-13 assessment.

subscores:

1. Guilt (self-blame feelings)
2. Suicide (ideation or attempts)
3. Insomnia-Initial (falling asleep)
4. Insomnia-Middle (waking at night)
5. Insomnia-Late (early awakening)
6. Work/Interests (reduced interest)
7. Psychic Anxiety (anxiety symptoms)
8. GI Symptoms (appetite, digestion)
9. Somatic Symptoms (fatigue)
10. Genital Symptoms (libido loss)

11. Hypochondriasis (health concerns)
12. Weight Loss (weight changes)
13. Insight (illness awareness)

[Transcript]: {transcript}

Output a summary for each subscore.

### G.2 HAMD-13 Scoring Prompt

The prompt for Qwen3-14B zero-shot HAMD-13 scoring:

You are a psychiatrist using HAMD-13. Score each subscore based on the summary.

Guidelines:

- Guilt (0-4): 0=absent, 4=delusions
- Suicide (0-4): 0=absent, 4=attempts
- Insomnia (0-2): 0=none, 2=nightly
- Work/Interests (0-4): 0=normal, 4=stopped
- Psychic Anxiety (0-4): 0=absent, 4=severe
- Somatic (0-2): 0=absent, 2=severe
- Hypochondriasis (0-4): 0=absent, 4=delusions
- Weight Loss (0-2): 0=none, 2=obvious
- Insight (0-2): 0=aware, 2=denies

[Summary]: {summary}

Output JSON: {"Guilt": X, "Suicide": X, ...}

### G.3 PHQ-8 Scoring Prompt

The prompt for Qwen3-14B zero-shot PHQ-8 scoring:

# Task:

Based on the given interview conversation excerpts with the patient, assign a PHQ-8 score (0-3) for each of the 8 factors below.

If a factor is not explicitly mentioned in the interview, you must still estimate a reasonable score by inferring from the overall patient situation and context.

# PHQ-8 items and scoring rules:

Over the last 2 weeks, how often has the patient been bothered by the following?

# Items:

1. NoInterest - Little interest or pleasure in doing things.
2. Depressed - Feeling down, depressed, or hopeless.
3. Sleep - Trouble falling or staying asleep, or sleeping too much.
4. Tired - Feeling tired or having little energy.
5. Appetite - Poor appetite or overeating.
6. Failure - Feeling bad about yourself, or that you are a failure.
7. Concentrating - Trouble concentrating on things.
8. Moving - Moving or speaking so slowly that others noticed, or being fidgety.

# Each PHQ-8 item is rated on a 4-point scale:

0 = Not at all

1 = Several days

1006 2 = More than half the days  
1007 3 = Nearly every day  
1008  
1009 [Transcript]: {transcript}  
1010  
1011 # Response Format:  
1012 Output all 8 factors in this format:  
1013 - NoInterest: 0  
1014 - Depressed: 0  
1015 - Sleep: 0  
1016 - Tired: 0  
1017 - Appetite: 0  
1018 - Failure: 0  
1019 - Concentrating: 0  
1020 - Moving: 0