

# CNSocialDepress: A Chinese Social Media Dataset for Depression Risk Detection and Structured Analysis

Jinyuan XU<sup>1\*</sup>, Tian LAN<sup>2\*</sup>, Xintao YU<sup>3</sup>, Xue HE<sup>3,4</sup>, Hezhi ZHANG<sup>5</sup>, Ying WANG<sup>6</sup>

Pierre Magistry<sup>1</sup>, Mathieu Valette<sup>1</sup>, Lei LI<sup>7†</sup>

<sup>1</sup>Ertim Inalco, <sup>2</sup>Milkuya Studio, <sup>3</sup>Sorbonne Université, <sup>4</sup>IRD Lab,

<sup>5</sup>Faculty of Psychology, Peking University,

<sup>6</sup>Faculty of Psychology and Cognitive Science, Beijing Normal University,

<sup>7</sup>Beijing Institute of Technology

## Abstract

Depression is a pressing global public health issue, yet publicly available Chinese-language resources for depression risk detection remain scarce and largely focus on binary classification. To address this limitation, we release **CNSocialDepress**, a benchmark dataset for depression risk detection on Chinese social media. The dataset contains 44,178 posts from 233 users; psychological experts annotated 10,306 depression-related segments. CNSocialDepress provides binary risk labels along with structured, multidimensional psychological attributes, enabling interpretable and fine-grained analyses of depressive signals. Experimental results demonstrate the dataset's utility across a range of NLP tasks, including structured psychological profiling and fine-tuning large language models for depression detection. Comprehensive evaluations highlight the dataset's effectiveness and practical value for depression risk identification and psychological analysis, thereby providing insights for mental health applications tailored to Chinese-speaking populations.

**Keywords:** Depression Detection, Chinese Social Media, Benchmark Dataset, Mental Health

## 1. Introduction

Depressive disorders are among the most common mental health conditions worldwide and are characterized by persistent low mood or loss of interest in daily activities. According to a 2023 report<sup>1</sup> by the World Health Organization (WHO), approximately 280 million people worldwide live with depression. In China, a 2024 survey by the Chinese Center for Disease Control and Prevention (Wang et al., 2024a) estimates that around 95 million individuals are affected by depression. Of the approximately 280,000 suicides reported annually in China, 40% are linked to depressive disorders. Moreover, research (Wang et al., 2024a) shows that depression is strongly associated with both suicidal behavior and non-suicidal self-injury.

Motivated by the urgent need to detect depression, researchers increasingly apply machine learning (ML) and natural language processing (NLP) methods to automatically assess depression risk (Squires et al., 2023; Hasib et al., 2023; Aleem et al., 2022; Liu et al., 2024; Shi et al., 2024; Jia and Li, 2024). While early efforts have shown promising results, they remain constrained by the limitations of existing datasets. Traditional depression detection studies often rely on clinical data (Bittar et al., 2019; Fernandes et al., 2018) or on transcripts from med-

ical or psychological interviews (Shen et al., 2022; Li et al., 2022b), which are expensive to collect, limited in size and diversity, and may not reflect the informal, emotionally nuanced expressions typical of real-world online environments.

To overcome these limitations, recent research has shifted toward leveraging user-generated content on social media platforms (Bucur et al., 2025; Harrigan et al., 2021; Li et al., 2025d; Jiang et al., 2024; Cai et al., 2025a), which provides rich linguistic signals and is more readily available. Such data are also more easily anonymized, thereby alleviating some privacy concerns. Beyond these methodological advantages, social media has become an important channel for emotional expression, especially in East Asian contexts (Zhou et al., 2023; Yang and Li, 2009; Zhou et al., 2024; Yang et al., 2025), where personal emotions may be expressed more freely online than in face-to-face settings. As a result, social media can provide richer and more authentic linguistic signals for depression analysis.

However, most existing datasets focus solely on classification with binary or multi-class labels and lack structured psychological insights or professional validation. Moreover, generative models are increasingly used in mental health applications (Hu et al., 2024; Xu et al., 2024; Yang et al., 2024b; Lai et al., 2023; He and Qu, 2025a; Gu et al., 2025; Cai et al., 2025b; Xu et al., 2025). Researchers in psychology and computational linguistics emphasize the need for datasets that combine risk labels with structured analyses, such as user profiles or

\*These authors contributed equally to this work.

†Corresponding Author.

<sup>1</sup><https://www.who.int/news-room/fact-sheets/detail/depression>

survey-based explanations. Such datasets better match real-world settings. Multifaceted insights can support clinical decision-making and downstream intervention.

Some recent efforts have introduced summarization-style datasets for depression. They leverage information extraction and automated summarization techniques (Sotudeh et al., 2022, 2021; He and Qu, 2025b; Shi et al.). However, these datasets often rely on model-generated content. The content is typically validated using automatic metrics such as ROUGE (Lin, 2004). They also lack annotation by mental health professionals, which raises concerns about domain-specific reliability and accuracy.

To address these challenges, we introduce **CNSocialDepress (CNSD)**, a publicly available Chinese-language dataset for depression risk detection that can be applied to the development of early-stage depression detection tools. It pairs binary risk labels with structured psychological analyses. All annotations are drafted and validated by certified mental health professionals, ensuring domain relevance and annotation quality. **CNSD** supports multiple task paradigms. These include binary classification, structured analysis generation, summarization, and fine-tuning large language models (LLMs) for psychological reasoning.

Our contributions are fourfold:

- We release **CNSD**, a high-quality Chinese dataset for depression risk detection that combines binary labels with structured, interpretable psychological analyses.
- We propose an expert-in-the-loop annotation protocol with structured templates and quality control.
- We benchmark **CNSD** across diverse task settings, including classification, structured analysis generation, summarization, and LLM fine-tuning for psychological reasoning.
- We present and validate a pipeline for generating structured psychological analyses for depression risk (Section 4).

## 2. Related Work

### 2.1. Existing Datasets for Depression Detection

Current datasets for depression detection mainly come from English-language social media platforms such as Twitter, Reddit, Facebook, and Instagram (Shen et al., 2017; Parapar et al., 2022; Zhang et al., 2021; Raihan et al., 2024; Li et al., 2025c; Jiang et al., 2025). Common annotation strategies include identifying self-disclosure (Yates

et al., 2017; Bathina et al., 2021; Islam et al., 2018), manual coding by clinical raters (Almouzini et al., 2019; Yazdavar et al., 2020; Alhamed et al., 2024), and symptom mapping (Seabrook et al., 2018; Aldarwish and Ahmad, 2017; Zhang et al., 2022; Li et al., 2025a). Symptom definitions are often based on DSM-5 criteria (Association, 2013) and questionnaire instruments such as the PHQ-9 (Kroenke et al., 2001).

Most datasets use binary labels. Others incorporate severity scales or symptom-specific tags (Zhang et al., 2022; Mowery et al., 2017; Guan et al., 2025). Multilingual resources also exist for depression detection. They cover languages such as Spanish (Romero et al., 2024), Arabic (Maghraby and Ali, 2022), Russian (Stankevich et al., 2020), Portuguese (Santos et al., 2024), Japanese (Yuka Niimi, 2021), and Thai (Hämäläinen et al., 2021). For Chinese, Sina Weibo is the most common data source (Li et al., 2020; Shen et al., 2018; Li et al., 2023; Guo et al., 2023; Yang et al., 2021; Li et al., 2025b). Datasets such as WU3D (Wang et al., 2020a) and SWDD (Cai et al., 2023) are widely used.

### 2.2. Methodological Advancements

Early work relied on statistical representations such as TF-IDF and hand-crafted features (Yang et al., 2020; Li et al., 2022a). These features were typically paired with traditional machine learning classifiers (Cortes, 1995; Breiman, 2001; McCallum et al., 1998; Dreiseitl and Ohno-Machado, 2002; He et al., 2026). With the rise of neural approaches (Schmidhuber, 2015), representation learning became central to depression detection. Embedding methods such as Word2Vec (Mikolov et al., 2013) and neural architectures including LSTMs (Sak et al., 2014), CNNs (Kim, 2014), Transformers (Vaswani, 2017), and BERT (Devlin et al., 2019) improved performance.

Recently, large language models (LLMs) have been explored for depression detection (Lan et al., 2024; Hu et al., 2024; Xin Yan, 2023; Lai et al., 2023; Li et al., 2026; He and Qu, 2025a; Shi et al., 2025). They can produce predictions and generate natural-language rationales or structured analyses, which supports interpretability. LLM-based approaches have also been used to improve explainability in mental health analysis (Wang et al., 2024c; Yang et al., 2024b; Xu et al., 2024; Yang et al., 2023; Hu et al., 2024; Gu et al., 2025).

### 2.3. Summary and Gaps

Despite these efforts, few datasets include structured psychological analyses, especially for non-English social media. Furthermore, available resources often lack annotation by mental health

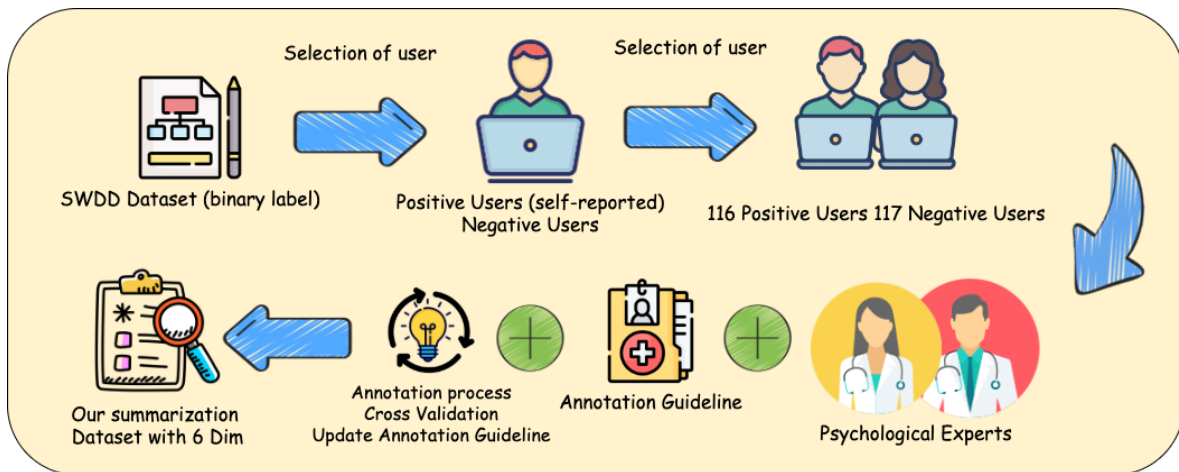


Figure 1: Dataset construction process. We sample 116 positive (depressed) users and 117 negative users from SWDD for expert annotation. Psychologists draft an initial guideline based on DSM-5 and corpus statistics, and iteratively refine it through random cross-validation. The final gold-standard data include a six-dimensional structured analysis summary for each user.

professionals, which limits practical utility. This study addresses these limitations by presenting an expert-validated Chinese-language dataset with detailed structured annotations. The dataset provides richer resources for depression risk detection and analysis.

### 3. Construction of Dataset

Our raw data come from the SWDD dataset (Cai et al., 2023), a user-level corpus collected from Sina Weibo, one of the largest Chinese-language social media platforms. The dataset includes two user groups: depressed and non-depressed. Each user has dozens to hundreds of Weibo posts. In addition to binary user labels, SWDD provides expert-annotated user-level depression features based on DSM-5 criteria<sup>2</sup> (Association, 2013).

From SWDD, we select 116 self-reported depressed (positive) users who disclosed a clinical diagnosis of depression in their posts, and 117 non-depressed (negative) users as candidates. Each selected user has at least 60 posts. The total length of a user’s posts is at least 3,000 tokens, computed using the Qwen2.5 tokenizer.<sup>3</sup>

To guide annotation, a team of psychology experts drafted an initial guideline based on DSM-5 criteria, the PHQ-9 (Kroenke et al., 2001), and prior work on linguistic markers of depression in online text (Mothe et al., 2022).

We define two levels of criteria. Primary criteria cover objective statements with higher diagnostic specificity, such as clinical symptoms, medical records, and explicit self-reports. Secondary cri-

teria capture subjective linguistic or emotional expressions with lower specificity, such as negative phrasing and emotional intensity. This hierarchy follows clinical principles that prioritize objective symptoms over subjective perceptions. The guidelines cover six dimensions:

- **Dimension 1: Depressive psychological state (primary).** This dimension reflects loss of self-worth, overwhelming guilt, and suicidal ideation. Representative spans include *inferiority*, *apology*, *death*, *self-harm*, and *suicide*.
- **Dimension 2: Medical expressions related to depression (primary).** This dimension covers mentions of medication use and clinical diagnoses, such as *taking medication*, *venlafaxine*, *side effects*, *depression*, *anxiety disorder*, *hospital*, and *doctor*.
- **Dimension 3: Clinical symptoms related to depression (primary).** This dimension captures physiological or somatic changes, including appetite or weight changes, sleep disturbance, fatigue, and physical pain. Representative spans include *appetite*, *insomnia*, *sleeping pills*, *nightmares*, *tiredness*, *headache*, and *stomachache*.
- **Dimension 4: Negative emotions (secondary).** This dimension covers adverse emotional states such as *sadness*, *grief*, *anxiety*, *loneliness*, and *despair*.
- **Dimension 5: Potential external causes of depression (secondary).** This dimension captures possible triggers such as intimate relationship issues, family conflict, major social

<sup>2</sup>DSM-5 Fact Sheets.

<sup>3</sup>Qwen2.5 tokenizer.

	Positive (Depressive)	Negative (Non- depressive)
NO. of Users	116	117
NO. of Texts	20,360	23,818
NO. of Tokens	1,024,978	1,115,951

Table 1: Dataset statistics.

events, and everyday stressors. Representative spans include *divorce*, *parents*, *school*, *teacher*, and *dropping out*.

- **Dimension 6: Language Use Patterns Related to Depression (secondary).** This dimension captures linguistic patterns such as negation, questions, and derogatory expressions. Representative spans include *I don't know*, *I don't like myself*, *I don't want to*, *I can't do it*, *why*, and *what should I do ?*.



Each user has multiple posts. We treat each post as the primary unit for review. Annotation is performed on semantic spans within each post. Each span is assigned to one of the six dimensions. A post may contain multiple spans, and different spans may belong to different dimensions.

Annotation was conducted by four senior researchers with expertise in psychology and depression scales. Annotators were blinded to the original binary labels in SWDD to ensure independent judgment. They periodically cross-checked overlapping samples to maintain consistency. Because the task relied more on qualitative judgment than on fixed quantitative criteria, the annotation guideline was not set in stone from the outset. It was revised in rounds as experts reviewed newly labeled data and resolved disagreements through discussion, and earlier annotations were corrected when the updated standards called for it. Given this evolving process, conventional inter-annotator agreement scores are not directly applicable. Since our annotation follows a different scheme from SWDD and all labeling was performed independently of the original labels, some user-level labels in our dataset differ from those in the original dataset.

This process yields the *CNSD Gold* set (100 positive and 100 negative users) and the *CNSD Test* set (16 positive and 17 negative users). Figure 1 summarizes the construction process.

Each entry (Figure 2) contains a binary label and a dimension-specific analysis. For each dimension, the analysis provides an overall assessment. It also lists post indices with brief justifications.

The overall statistics are presented in Table 1, and the 10,306 dimension-level annotations are summarized in Table 2.



Data Sample


**User id:** user 61

**User posts:**

**post1:** 'You are really planning to leave, I just don't believe it. Covering my eyes and closing my ears, pretending you will come back. Deceiving myself.'

**post2, post3. post4, post5, ....**

**Label:** Depressed 

**Dim 1: Depressive psychological state (primary).**

**Overall Assessment:** The user presents a persistent low mood accompanied by a diminished sense of self-worth.

**Post Selections:** post\_36 [self-negation], post\_37 [persistent low mood], post\_62 [self-negation], post\_152 [nocturnal depressive mood], post\_230 [self-harm ideation], post\_278 [anxiety attack], post\_275 [suicidal request], ...

**Dim 2: Medical expressions related to depression (primary).**

**Overall Assessment:** The posts mention medical help-seeking behaviors related to depression, such as clinical visits and antidepressant medication use.

**Post Selections:** post\_318 [hypnotic medication use], post\_321 [medication adherence], post\_334 [distress during clinical follow-up], ...

**Dim 3: Clinical symptoms related to depression (primary).**

**Overall Assessment:** The posts mention depression-related clinical symptoms such as insomnia and fatigue.

**Post Selections:** post\_173 [insomnia], post\_209 [hypersomnia], post\_273 [sleep disturbance], ...

**Dim 4: Negative emotions (secondary).**

**Overall Assessment:** The posts frequently express emotions such as sadness, hopelessness, and anger.

**Post Selections:** post\_1 [extreme sadness], post\_17 [hopelessness and helplessness], post\_67 [crying], ...

**Dim 5: Potential external causes of depression (secondary).**

**Overall Assessment:** The posts mention external stressors, such as stressful life events.

**Post Selections:** post\_94 [workplace criticism], post\_101 [academic/school pressure], post\_102 [social pressure], post\_112 [social aversion], ...

**Dim 6: Language Use Patterns Related to Depression (secondary).**

**Overall Assessment:** The posts include crisis-oriented statements, such as "I don't know what to do" and "I want to die."

**Post Selections:** post\_18 [suicidal ideation], post\_203 [death wish], ...










Figure 2: A user-level annotated entry from CNSD-Gold (English translations).

Dimension	Negative User	Positive User
Dim1	117	2127
Dim2	2	555
Dim3	126	768
Dim4	514	2933
Dim5	193	863
Dim6	231	1877
Total	1183	9123

Table 2: Dimension-level annotation statistics for negative and positive users. This table reports the number of expert-annotated depression-related spans in each dimension. A single text may contain multiple spans.

## 4. Automated Dataset Generation Pipeline

Expert annotation requires substantial human effort and professional psychological expertise. As a result, user-level depression-risk datasets for Chinese social media remain scarce. To narrow this gap, we develop an automated dataset generation pipeline. We use it to construct *CNSD Silver* from the SWDD binary classification dataset. The pipeline is informed by the manually annotated *CNSD Gold* dataset (Section 3).

To build *CNSD Silver*, we randomly sample 100 positive (depression-risk) users and 100 negative (non-risk) users from SWDD. Each user contributes dozens to hundreds of posts. We keep users with at least 3,000 tokens in total.

The pipeline consists of two modules.

### 4.1. Module I: Dimension-Wise Automatic Labeling

Module I uses a mid-sized model (Qwen2.5-14B (Bai et al., 2023)) to automatically assign posts to the six depression-related dimensions ( $\mathcal{D} = \{D_1, D_2, \dots, D_6\}$ ). We fine-tune the model on text-level training data curated from positive users in *CNSD Gold*. After fine-tuning, the model can assign each post to one or more dimensions.

#### Data preparation for fine-tuning Module I.

1. For each dimension  $D_k$ , we extract posts from positive users in *CNSD Gold* that contain at least one annotated span of  $D_k$ . These posts form a set  $S_k$ .
2. We identify the dimension with the fewest training posts. The minimum count is  $n_{\min} = 442$ , which corresponds to *depression-related medical expressions*:

$$n_{\min} = \min_{k \in \{1, 2, \dots, 6\}} n_k = 442.$$

3. To balance dimensions, we downsample each  $S_k$  to  $n_{\min}$  when  $n_k > n_{\min}$ . This yields balanced sets  $S'_k$ :

$$S'_k = \begin{cases} S_k, & \text{if } n_k = n_{\min}, \\ \text{RandomSubset}(S_k, n_{\min}), & \text{if } n_k > n_{\min}. \end{cases}$$

We obtain  $6 \times 442$  labeled positive posts. Each post is formatted using an instruction template such as: "This post belongs to [category] because it mentions [evidence]." We further augment the labels with 20 paraphrased expressions with the same meaning (Appendix A.2).

4. For negative training data, we split posts from negative users in *CNSD Gold* into individual texts. We remove any text that contains depression-related content in any dimension. From the remaining texts, we uniformly sample  $6 \times 442$  examples. We assign them negative label expressions indicating that there is no evidence for the target dimension (As with the positive examples, we create 20 negative label expressions to reduce overfitting; details are provided in A.3).
5. We fine-tune Qwen2.5-14B with LoRA (Hu et al., 2022) on the combined set of  $2 \times (6 \times 442)$  texts for 1 epoch. We use a learning rate of  $5 \times 10^{-5}$ , with the prompt provided in Appendix A.4.

### 4.2. Module II: Automatic Verification and Summarization

In Module II, we apply the fine-tuned Module I model to label all posts for each user. We then collect posts that are assigned to at least one dimension. Next, we use DeepSeek-R1-671B (DeepSeek-AI et al., 2025) to verify the assigned labels and the supporting evidence on a case-by-case basis. After verification, we summarize the collected posts using a few-shot prompting setup. The prompt instructs the model to act as a psychology expert and to produce structured summaries aligned with *CNSD Gold*. We also provide detailed requirements and gold-standard examples. The full prompt is shown in Appendix A.7.

## 5. Experiments

We design our experiments to answer two research questions. **First**, how effectively does our data generation pipeline produce high-quality structured annotations? **Second**, how well do various large language models (LLMs) perform on user-level depression risk summarization and structured psychological analysis?

Based on these questions, we conduct three sets of experiments:

1. To assess the quality of model-generated six-dimensional structured depression analyses. We compare models fine-tuned on *CNSD Gold* and *CNSD Silver* with baseline models and few-shot prompting.
2. To evaluate the performance of these fine-tuned models on user-level depression risk classification.
3. To benchmark leading LLMs on a structured depression analysis summarization task using *CNSD Gold*.

In Section 5.5, we also use *CNSD Gold* as a test set for binary depression classification. This further illustrates the dataset’s versatility.

### 5.1. Baseline Models

We use the following generative baselines: DeepSeek-R1-Distill-14B (DeepSeek-AI et al., 2025), Qwen2.5-14B (QwenTeam, 2024), GPT-4o<sup>4</sup>, GPT-4o-mini<sup>5</sup>, DeepSeek-R1-671B<sup>6</sup> (DeepSeek-AI et al., 2025), and Llama3-8B-Chinese-Chat (Wang et al., 2024b). We summarize key characteristics in Table A.1.

### 5.2. Experimental Setup

We use NVIDIA A800 (80GB) and A100 (80GB) GPUs. Unless otherwise specified, we set the generation temperature to 0.7. All fine-tuning experiments use the LLaMA-Factory framework (Zheng et al., 2024). We apply LoRA (Hu et al., 2022) for 1 epoch with a learning rate of  $5 \times 10^{-5}$ .

Our experiments have two parts. The first evaluates the pipeline for data generation. The second benchmarks models on the *CNSD Gold* dataset for structured summarization and analysis generation.

For classification, we report accuracy, precision, recall, and F1-score. For generation, we report BLEU (Papineni et al., 2002), ROUGE-1 (Lin, 2004), and BERTScore (Zhang et al., 2020).

### 5.3. Task I: Data Generation

We conduct experiments on the *CNSD Test* dataset. We focus on two Qwen2.5-14B models fine-tuned on *CNSD Gold* and *CNSD Silver*:


- **Qwen2.5-14B Gold:** Fine-tuned on *CNSD Gold*.
- **Qwen2.5-14B Silver:** Fine-tuned on *CNSD Silver*. *CNSD Silver* is generated using the

<sup>4</sup><https://openai.com/index/hello-gpt-4o/>

<sup>5</sup><https://openai.com/index/>

GPT-4o-mini-advancing-cost-efficient-intelligence/

<sup>6</sup><https://www.deepseek.com/>


 You are a professional in psychology and text-based sentiment analysis.
 

-----

Please determine whether the user is at risk of depression based on the following six dimensions:

- *Negative emotions (secondary criterion)*
- *Depressive psychological state (primary criterion)*
- *Depression-related clinical symptoms (primary criterion)*
- *Potential external factors causing depression (secondary criterion)*
- *Depression-related medical expressions (primary criterion)*
- *Depression-related language expression patterns (secondary criterion)*

For each dimension, please cite the corresponding text reference and provide a concise explanation.

 [Output Data Example]

- **Answer: “Yes” or “No.”**
- **Briefly explain the manifestations in each of the six dimensions, and cite the corresponding text reference.**
- **In conclusion, “The user is at risk of depression” or “The user is not at risk of depression.”**

Figure 3: Prompt used for Table 3.

pipeline in Section 4 on 100 positive and 100 negative users sampled from SWDD.

We compare these models with other baselines in two aspects:

1. Depression risk classification performance.
2. Quality of the generated six-dimensional analyses, including hallucinated or unsupported content.

Unless noted otherwise, we use user-level generation. We concatenate all posts from a user as input. We instruct the model to output a binary label and a corresponding justification.

#### 5.3.1. Six-Dimensional Structured Depression Analysis

Prior work suggests that incorporating structured knowledge can improve LLM outputs (Moiseev et al., 2022). *CNSD Gold* and *CNSD Silver* provide fine-grained structured annotations across six dimensions. They support user-level summarization and analysis of social media content. We hypothesize that fine-tuning on these datasets can reduce hallucinations. It may also improve generation quality compared to few-shot prompting and larger models.

Table 3 compares different generation strategies using automatic metrics. Our Pipeline achieves

Strategy	BERTScore	ROUGE-1	BLEU
<b>Pipeline</b>	<b>0.791</b>	<b>0.478</b>	<b>0.288</b>
FS:Qwen2.5 14B	0.649	0.076	0.075
FS:GPT-4o	0.678	0.269	0.094
FS:GPT-4o Mini	0.674	0.170	0.070
FS:DeepSeek R1 671B	0.678	0.301	0.054
FS:DeepSeek R1 Distill -14B	0.6756	0.191	0.073

Table 3: Comparison of generation strategies on text quality. In the table, FS stands for Few-Shot. Pipeline refers to the automated dataset generation pipeline proposed in Section 4, applied to our proposed dataset in Section 3. The purpose is to show that, with the same prompt, our pipeline achieves the highest text generation quality while meeting our task requirements. The prompt used in this experiment is shown in Figure 3.

the best scores on all three metrics: BERTScore (0.791), ROUGE-1 (0.478), and BLEU (0.288). These results indicate stronger semantic alignment, higher lexical overlap, and better  $n$ -gram precision. Few-shot (FS) prompting yields lower scores across models, including Qwen2.5-14B, GPT-4o, and DeepSeek-R1-671B. Notably, DeepSeek-R1-671B does not outperform the Pipeline. This suggests that the proposed Pipeline produces text that matches the reference more closely than few-shot prompting, across both smaller and larger models.

We also compare the original Qwen2.5-14B model with Qwen2.5-14B Gold and Qwen2.5-14B Silver (Table 4). We use two complementary evaluation methods.

- **Automatic metrics.** The original model achieves a BERTScore of 0.649. It increases to 0.764 for Gold (+16.9%) and 0.787 for Silver (+21.5%). ROUGE-1 rises from 0.076 to 0.217 for Gold (+172%) and 0.218 for Silver (+172%). BLEU improves from 0.075 to 0.186 for Gold (+149%) and 0.237 for Silver (+218%).
- **Human evaluation.** Two linguistic experts evaluate accuracy, coverage, and hallucination. Compared to the original Qwen2.5-14B, both Gold and Silver improve substantially. Silver increases by 33.4% in accuracy and 37.1% in coverage. Its hallucination score increases by 15.5% (higher is better). Gold achieves larger gains. Accuracy increases by 42.9% and coverage increases by 58.0%. The hallucination score increases by 26.8%. Overall, Gold performs best in human evaluation. Silver is slightly behind but remains close to Gold and clearly outperforms the original model.

In summary, Gold performs best in human evaluation. Silver performs better on automatic metrics. Both models substantially outperform the original model. These results show that fine-tuning with our dataset and pipeline can improve generation quality at semantic, lexical, and structural levels.

### 5.3.2. Classification Task

In the depression classification task, *Qwen2.5-14B Silver* achieves the best performance among all models. It reaches an accuracy of 0.944 and an F1 score of 0.941. It also surpasses *Qwen2.5-14B Gold*, which is fine-tuned on human-annotated data. Moreover, it outperforms large-scale models such as GPT-4o (accuracy = 0.917, F1 = 0.923) and DeepSeek-R1-671B (accuracy = 0.861, F1 = 0.872). These results support the effectiveness of our generation pipeline and suggest that the resulting silver data are of high quality. Detailed results are reported in Table 5.

### 5.4. Task II: Structured Summarization and Analysis Generation

Our dataset supports research on generative models for depression risk assessment from social media. It is particularly useful for user-level summarization of depressive signals. To the best of our knowledge, there is no directly comparable open-source dataset with the same annotation schema. We therefore conduct an initial benchmark of mainstream generative models on our annotated dataset of 233 users. Results are reported in Table 6.

We use the following prompt:

- **Instruction:** *Read the following texts written by one user. Then answer the question.*
- **Question:** Based on these texts, does this user exhibit a depressive mood ?
  - If the user exhibits a depressive mood, answer Yes.
  - If the user does not exhibit a depressive mood, answer No.
- **Output requirements:** Provide a brief justification based on evidence from the texts. Then produce a structured summary aligned with the annotation schema.

Model	BERTScore	ROUGE-1	BLEU	Human.acc	Human.cov	Human.hallu
Qwen2.5-14B	0.649	0.076	0.075	0.583	0.574	0.657
Qwen2.5-14B Silver	<b>0.787</b>	<b>0.218</b>	<b>0.237</b>	0.778	0.787	0.759
Qwen2.5-14B Gold	0.764	0.217	0.186	<b>0.833</b>	<b>0.907</b>	<b>0.833</b>

Table 4: Comparison of generation quality across models. Qwen2.5-14B Gold is fine-tuned on *CNSD Gold*. Qwen2.5-14B Silver is fine-tuned on *CNSD Silver*, which is generated by our pipeline (Section 4). We report automatic metrics (BERTScore, ROUGE-1, BLEU) and human evaluation metrics: *acc* (accuracy), *cov* (content coverage), and *hallu* (hallucination; higher is better). Higher values indicate better performance for all metrics.

Model	Accuracy	Recall	Precision	F1
Qwen2.5-14B	0.889	0.944	0.850	0.894
Qwen2.5-14B (FT, Silver)	<b>0.944</b>	0.889	<b>1.000</b>	<b>0.941</b>
Qwen2.5-14B (FT, Gold)	0.861	0.889	0.842	0.864
GPT-4o	0.917	<b>1.000</b>	0.857	0.923
GPT-4o-mini	0.667	<b>1.000</b>	0.600	0.750
DeepSeek-R1-671B	0.861	0.944	0.801	0.872
DeepSeek-R1-Distill-14B	0.889	<b>1.000</b>	0.818	0.900

Table 5: Comparison of Models for Generated Dataset Quality: Application to Classification Tasks.

The results vary across models. BERTScore falls in a narrow range (0.654–0.710), with GPT-4o achieving the highest score. ROUGE-1 is highest for DeepSeek-R1-671B (0.454), which exceeds the second-best model (GPT-4o at 0.346) by an absolute margin of 0.108. BLEU ranges from 0.093 (Qwen2.5-7B) to 0.152 (GPT-4o).

Model	BERT Score	ROUGE-1	BLEU
Llama3-8B	0.654	0.194	0.122
GLM4-9B-Chat	0.694	0.208	0.140
Qwen2.5-7B	0.680	0.208	0.093
DeepSeek-R1-671B	0.698	<b>0.454</b>	0.125
GPT-4o-mini	0.696	0.166	0.126
GPT-4o	<b>0.710</b>	0.346	<b>0.152</b>

Table 6: Model performance on the depression risk summarization and analysis generation task.

### 5.5. Task III: Classification Experiments

Our dataset is primarily designed for generation tasks. We also conduct a binary classification experiment for depression risk. In addition to our dataset, we use the original SWDD and WU3D datasets. For each dataset, we randomly sample 200 positive and 200 negative users.

Besides the generative models described in Appendix A.1, we evaluate BERT-based classifiers (Devlin et al., 2019). Implementation details are provided in Appendix A.5. Results are reported in Appendix A.6.

## 6. Conclusion

We introduce **CNSocialDepress**, a new dataset for depression risk analysis from Chinese social media. To our knowledge, it is among the first to provide both user-level and text-level labels. The dataset combines binary depression indicators with fine-grained six-dimensional analyses annotated by mental health professionals. We also develop an automated data generation pipeline. It reduces manual annotation effort and supports language model fine-tuning for classification and summarization. We expect **CNSocialDepress** to be useful for a range of research and applications, supporting early risk identification and downstream mental health interventions.

## 7. Limitations

While **CNSocialDepress** provides a valuable resource for depression risk analysis on Chinese social media, several limitations should be noted. First, the dataset is sourced exclusively from Weibo. This may introduce platform-specific demographic biases, such as the underrepresentation of rural or elderly populations.

Second, linguistic diversity is not fully captured. This includes dialectal variation and metaphorical expressions that are common in Chinese depressive discourse. This limitation may reduce generalizability across regions. Third, expert annotation ensures high quality, but it limits scalability. It also results in a relatively small dataset.

In addition, the current structured analyses focus primarily on depressive symptoms. They omit co-

morbid mental health conditions (e.g., anxiety) and broader social context. Future work will expand data collection to multiple platforms. We also plan to update lexicons to capture emerging expressions. Finally, we will explore hybrid annotation frameworks (e.g., expert-guided crowdsourcing) to improve coverage and efficiency. Privacy-preserving techniques and longitudinal tracking can further strengthen ethical and practical utility.

## 8. Ethics Statement

This dataset is built upon the SWDD benchmark (Cai et al., 2023), released in early 2023. The original data consist of user-level posts from Weibo. They were annotated by psychology experts based on DSM-5 criteria.

During re-annotation, we applied a second round of de-identification to protect privacy. We remove personal information, geographic locations, and other potential identifiers from the texts. We follow recommended data protection practices Benton et al. (2017) and comply with the GDPR (General Data Protection Regulation).<sup>7</sup> We also recruit experienced researchers in depression-related psychology to provide expert annotations.

This dataset is intended for research on depression risk identification and for building assistive models. It can support multi-dimensional analysis of risk-related signals and the generation of structured summaries. It is not designed for clinical diagnosis, triage, or treatment decisions. Any outputs should be used only as references for professionals or as a self-assessment aid.

No algorithmic system can replace in-person psychiatric evaluation or provide a definitive diagnosis. If a system built on this dataset flags a high-risk case, users should be encouraged to seek professional help. The same applies to individuals with persistent mood disturbances or impaired functioning. Online self-assessment tools cannot substitute for clinical evaluation.

## 9. Bibliographical References

Maryam Mohammed Aldarwish and Hafiz Farooq Ahmad. 2017. Predicting depression levels using social media posts. In *2017 IEEE 13th International Symposium on Autonomous decentralized system (ISADS)*, pages 277–280. IEEE.

Shumaila Aleem, Noor ul Huda, Rashid Amin, Samina Khalid, Sultan S Alshamrani, and Abdullah Alshehri. 2022. Machine learning algorithms for

depression: diagnosis, insights, and research directions. *Electronics*, 11(7):1111.

Falwah Alhamed, Julia Ive, and Lucia Specia. 2024. [Classifying social media users before and after depression diagnosis via their language usage: A dataset and study](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3250–3260, Torino, Italia. ELRA and ICCL.

Salma Almouzini, Asem Alageel, et al. 2019. Detecting arabic depressed users from twitter data. *Procedia Computer Science*, 163:257–265.

Amna Amanat, Muhammad Rizwan, Abdul Rehman Javed, Maha Abdelhaq, Raed Alsaqour, Sharnil Pandya, and Mueen Uddin. 2022. Deep learning for depression detection from textual data. *Electronics*, 11(5):676.

Priyanka Arora and Parul Arora. 2019. Mining twitter data for depression detection. In *2019 international conference on signal processing and communication (ICSC)*, pages 186–189. IEEE.

American Psychiatric Association. 2013. Diagnostic and statistical manual of mental disorders. *Am Psychiatric Assoc*, 21(21):591–643.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Krishna C Bathina, Marijn Ten Thij, Lorenzo Lorenzo-Luaces, Lauren A Rutter, and Johan Bollen. 2021. Individuals with depression express more distorted thinking on social media. *Nature human behaviour*, 5(4):458–466.

Aaron T Beck, Robert A Steer, Roberta Ball, and William F Ranieri. 1996. Comparison of beck depression inventories-ia and-ii in psychiatric outpatients. *Journal of personality assessment*, 67(3):588–597.

Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. Ethical research protocols for social media health research. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 94–102, Valencia, Spain. Association for Computational Linguistics.

André Bittar, Sumithra Velupillai, Angus Roberts, and Rina Dutta. 2019. Text classification to inform suicide risk assessment in electronic health records. In *MEDINFO 2019: Health and Wellbeing e-Networks for All*, pages 40–44. IOS Press.

Leo Breiman. 2001. Random forests. *Machine learning*, 45:5–32.

<sup>7</sup><https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02016R0679-20160504>

- Peter F Brown, Vincent J Della Pietra, Peter V Desouza, Jennifer C Lai, and Robert L Mercer. 1992. Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–480.
- Ana-Maria Bucur, Andreea Moldovan, Krutika Parvatikar, Marcos Zampieri, Ashiqur Khudabukhsh, and Liviu Dinu. 2025. [Datasets for depression modeling in social media: An overview](#). In *Proceedings of the 10th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2025)*, pages 116–126, Albuquerque, New Mexico. Association for Computational Linguistics.
- Fidel Cacheda, Diego Fernandez, Francisco J Novoa, and Victor Carneiro. 2019. Early detection of depression: social network analysis and random forest techniques. *Journal of medical Internet research*, 21(6):e12554.
- Chengkun Cai, Haoliang Liu, Xu Zhao, Zhongyu Jiang, Tianfang Zhang, Zongkai Wu, John Lee, Jenq-Neng Hwang, and Lei Li. 2025a. Bayesian optimization for controlled image editing via llms. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Chengkun Cai, Xu Zhao, Haoliang Liu, Zhongyu Jiang, Tianfang Zhang, Zongkai Wu, Jenq-Neng Hwang, and Lei Li. 2025b. The role of deductive and inductive reasoning in large language models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Yicheng Cai, Haizhou Wang, Huali Ye, Yanwen Jin, and Wei Gao. 2023. Depression detection on online social network with multivariate time series feature of user depressive symptoms. *Expert Systems with Applications*, 217:119538.
- Junyeop Cha, Seoyun Kim, and Eunil Park. 2022. A lexicon-based approach to examine depression detection in social media: the case of twitter and university community. *Humanities and Social Sciences Communications*, 9(1):1–10.
- Stevie Chancellor and Munmun De Choudhury. 2020. Methods in predictive techniques for mental health status on social media: a critical review. *NPJ digital medicine*, 3(1):43.
- Corinna Cortes. 1995. Support-vector networks. *Machine Learning*.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. [Revisiting pre-trained models for Chinese natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 657–668, Online. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Stephan Dreiseitl and Lucila Ohno-Machado. 2002. Logistic regression and artificial neural network classification models: a methodology review. *Journal of biomedical informatics*, 35(5-6):352–359.
- Zhiyuan Du, Yuxian Qian, Xiao Liu, Ming Ding, Yuxian Qian, Xiaoyang Liu, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335. Association for Computational Linguistics.
- Egle Eensoo and Mathieu Valette. 2012. Sur l’application de méthodes textométriques à la construction de critères de classification en analyse des sentiments. In *TALN 2012*, volume 2, pages 367–374. GETALP-LIG.
- SAGS Evans-Lacko, Sergio Aguilar-Gaxiola, Ali Al-Hamzawi, Jordi Alonso, Corina Benjet, Ronny Bruffaerts, Wai-Tat Chiu, Silvia Florescu, Giovanni de Girolamo, Oye Gureje, et al. 2018. Socio-economic variations in the mental health treatment gap for people with anxiety, mood, and substance use disorders: results from the who world mental health (wmh) surveys. *Psychological medicine*, 48(9):1560–1571.
- Andrea C Fernandes, Rina Dutta, Sumithra Velupillai, Jyoti Sanyal, Robert Stewart, and David Chandran. 2018. Identifying suicide ideation and suicidal attempts in a psychiatric clinical research database using natural language processing. *Scientific reports*, 8(1):7426.

- Shang Gao, Mohammed Alawad, M Todd Young, John Gounley, Noah Schaefferkoetter, Hong Jun Yoon, Xiao-Cheng Wu, Eric B Durbin, Jennifer Doherty, Antoinette Stroup, et al. 2021. Limitations of transformers on clinical text classification. *IEEE journal of biomedical and health informatics*, 25(9):3596–3607.
- GeminiTeam, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- GLMTeam, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadao Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, et al. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#).
- Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- Ruocheng Gu, Sen Jia, Yule Ma, Jinqin Zhong, Jenq-Neng Hwang, and Lei Li. 2025. Mocount: Motion-based repetitive action counting. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 9026–9034.
- Yunchuan Guan, Yu Liu, Ke Zhou, Zhiqi Shen, Serge Belongie, Jenq-Neng Hwang, and Lei Li. 2025. Learning to learn weight generation via local consistency diffusion. *arXiv preprint arXiv:2502.01117*. Accepted to CVPR 2026.
- Sharath Chandra Guntuku, David B Yaden, Margaret L Kern, Lyle H Ungar, and Johannes C Eichstaedt. 2017. Detecting depression and mental illness on social media: an integrative review. *Current Opinion in Behavioral Sciences*, 18:43–49.
- Zhihua Guo, Nengneng Ding, Minyu Zhai, Zhenwen Zhang, and Zepeng Li. 2023. Leveraging domain knowledge to improve depression detection on chinese social media. *IEEE Transactions on Computational Social Systems*, 10(4):1528–1536.
- Mika Hämäläinen, Pattama Patpong, Khalid Alnajjar, Niko Partanen, and Jack Rueter. 2021. [Detecting depression in Thai blog posts: a dataset and a baseline](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 20–25, Online. Association for Computational Linguistics.
- Keith Harrigan, Carlos Aguirre, and Mark Dredze. 2021. [On the state of social media data for mental health research](#). In *Proceedings of the Seventh Workshop on Computational Linguistics and Clinical Psychology: Improving Access*, pages 15–24, Online. Association for Computational Linguistics.
- Khan Md Hasib, Md Rafiqul Islam, Shadman Sakib, Md Ali Akbar, Imran Razzak, and Mohammad Shafiul Alam. 2023. Depression detection from social networks data based on machine learning and deep learning techniques: An interrogative survey. *IEEE Transactions on Computational Social Systems*, 10(4):1568–1586.
- Qi He and Chunyu Qu. 2025a. [Modular landfill remediation for ai grid resilience](#). *arXiv preprint*.
- Qi He and Chunyu Qu. 2025b. [Waste-to-energy-coupled ai data centers: Cooling efficiency and grid resilience](#). *arXiv preprint*.
- Qi He, Rui Shan, Chunyu Qu, Yuan Tang, and Yue Zou. 2026. [Rare-earth exposure and bottlenecks in ai data centers: Cost and schedule risk](#). *SSRN Electronic Journal*. Available at SSRN.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Jinpeng Hu, Tengting Dong, Luo Gang, Hui Ma, Peng Zou, Xiao Sun, Dan Guo, and Meng Wang. 2024. [Psychollm: Enhancing llm for psychological understanding and evaluation](#).
- Md Rafiqul Islam, Muhammad Ashad Kabir, Ashir Ahmed, Abu Raihan M Kamal, Hua Wang, and Anwaar Ulhaq. 2018. Depression detection from social network data using machine learning techniques. *Health information science and systems*, 6:1–12.
- Shaoxiong Ji, Tianyu Zhang, Laith Ansari, Jie Fu, Piyush Tiwari, and Erik Cambria. 2022. Mentalbert: Publicly available pretrained language models for mental healthcare. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 7184–7190. European Language Resources Association.
- Sen Jia and Lei Li. 2024. Adaptive masking enhances visual grounding. *arXiv preprint arXiv:2410.03161*.
- Zhongyu Jiang, Wenhao Chai, Lei Li, Zhuoran Zhou, Cheng-Yen Yang, and Jenq-Neng Hwang. 2025. Unihpr: Unified human pose representation via singular value contrastive learning. *arXiv preprint arXiv:2510.19078*.

- Zhongyu Jiang, Zhuoran Zhou, Lei Li, Wenhao Chai, Cheng-Yen Yang, and Jenq-Neng Hwang. 2024. Back to optimization: Diffusion-based zero-shot 3d human pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Herve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Heysem Kaya, Dmitrii Fedotov, Denis Dresvyan-skiy, Metehan Doyran, Danila Mamontov, Maxim Markitantov, Alkim Almila Akdag Salah, Evrim Kavcar, Alexey Karpov, and Albert Ali Salah. 2019. Predicting depression and emotions in the cross-roads of cultures, para-linguistics, and non-linguistics. In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, pages 27–35.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Kurt Kroenke, Robert L Spitzer, and Janet BW Williams. 2001. The phq-9: validity of a brief depression severity measure. *Journal of general internal medicine*, 16(9):606–613.
- Pierre Lafon. 1980. Sur la variabilité de la fréquence des formes dans un corpus. *Mots. Les langages du politique*, 1(1):127–165.
- Tin Lai, Yukun Shi, Zicong Du, Jiajie Wu, Ken Fu, Yichao Dou, and Ziqi Wang. 2023. Psy-llm: Scaling up global mental health psychological services with ai-based large language models. *arXiv preprint arXiv:2307.11991*.
- Xiaochong Lan, Yiming Cheng, Li Sheng, Chen Gao, and Yong Li. 2024. Depression detection on social media with large language models. *arXiv preprint arXiv:2403.10750*.
- Aziliz Le Glaz, Yannis Haralambous, Deok-Hee Kim-Dufoir, Philippe Lenca, Romain Billot, Taylor C Ryan, Jonathan Marsh, Jordan Devylder, Michel Walter, Sofian Berrouguet, et al. 2021. Machine learning and natural language processing in mental health: systematic review. *Journal of medical Internet research*, 23(5):e15708.
- Genghao Li, Bing Li, Langlin Huang, Sibing Hou, et al. 2020. Automatic construction of a depression-domain lexicon based on microblogs: text mining study. *JMIR medical informatics*, 8(6):e17650.
- Hao Li, Ju Dai, Xin Zhao, Feng Zhou, Junjun Pan, and Lei Li. 2025a. Wav2sem: Plug-and-play audio semantic decoupling for 3d speech-driven facial animation. *arXiv preprint arXiv:2505.23290*.
- Hao Li, Ju Dai, Feng Zhou, Kaida Ning, Lei Li, and Junjun Pan. 2025b. Au-blendshape for fine-grained stylized 3d facial expression manipulation. *arXiv preprint arXiv:2507.12001*.
- Lei Li, Sen Jia, and Jenq-Neng Hwang. 2026. Multiple human motion understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 6297–6305.
- Lei Li, Sen Jia, Jianhao Wang, Zhaochong An, Jiaang Li, Jenq-Neng Hwang, and Serge Belongie. 2025c. Chatmotion: A multimodal multi-agent for human motion analysis. *arXiv preprint arXiv:2502.18180*.
- Lei Li, Sen Jia, Jianhao Wang, Zhongyu Jiang, Feng Zhou, Ju Dai, Tianfang Zhang, Zongkai Wu, and Jenq-Neng Hwang. 2025d. Human Motion Instruction Tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lei Li, Tianfang Zhang, Zhongfeng Kang, and Xikun Jiang. 2022a. Mask-fpan: Semi-supervised face parsing in the wild with de-occlusion and uv gan. *arXiv preprint arXiv:2212.09098*.
- Mingzheng Li, Haojie Xu, Weifeng Liu, and Jiangwei Liu. 2022b. Bidirectional lstm and attention for depression detection on clinical interview transcripts. In *2022 IEEE 10th International Conference on Information, Communication and Networks (ICICN)*, pages 638–643. IEEE.
- Zepeng Li, Zhengyi An, Wenchuan Cheng, Jiawei Zhou, Fang Zheng, and Bin Hu. 2023. Mha: a multimodal hierarchical attention model for depression detection in social media. *Health information science and systems*, 11(1):6.
- Chenhao Lin, Pengwei Hu, Hui Su, Shaochun Li, Jing Mei, Jie Zhou, and Henry Leung. 2020. Sensemood: depression detection on social media. In *Proceedings of the 2020 international conference on multimedia retrieval*, pages 407–411.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Libin Liu, Shen Chen, Sen Jia, Jingzhe Shi, Zhongyu Jiang, Can Jin, Wu Zongkai, Jenq-Neng Hwang, and Lei Li. 2024. Graph canvas for controllable 3d scene generation. *arXiv preprint arXiv:2412.00091*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. [Roberta: A robustly optimized bert pre-training approach](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ashwag Maghraby and Hosnia Ali. 2022. Modern standard arabic mood changing and depression dataset. *Data in Brief*, 41:107999.
- Anshu Malhotra and Rajni Jindal. 2022. Deep learning techniques for suicide and depression detection from online social media: A scoping review. *Applied Soft Computing*, 130:109713.
- Keshu Malviya, Bholanath Roy, and SK Saritha. 2021. A transformers approach to detect depression in social media. In *2021 International Conference on Artificial Intelligence and Smart Systems (ICAIS)*, pages 718–723. IEEE.
- Andrew McCallum, Kamal Nigam, et al. 1998. A comparison of event models for naive bayes text classification. In *AAAI-98 workshop on learning for text categorization*, volume 752, pages 41–48. Madison, WI.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#).
- Fedor Moiseev, Zhe Dong, Enrique Alfonseca, and Martin Jaggi. 2022. [SKILL: Structured knowledge infusion for large language models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1581–1588, Seattle, United States. Association for Computational Linguistics.
- Josiane Mothe, Faneva Ramiandrisoa, and Md Zia Ullah. 2022. [Comparison of machine learning models for early depression detection from users' posts](#). In *Early Detection of Mental Health Disorders by Social Media Monitoring: The First Five Years of the eRisk Project*, volume 1018 of *Studies in Computational Intelligence book series (SCI)*, pages 111–139. Springer International Publishing.
- Danielle Mowery, Hilary Smith, Tyler Cheney, Greg Stoddard, Glen Coppersmith, Craig Bryan, Mike Conway, et al. 2017. Understanding depressive symptoms and psychosocial stressors on twitter: a corpus-based study. *Journal of medical Internet research*, 19(2):e6895.
- Ahmed Hussein Orabi, Prasadith Buddhitha, Mahmoud Hussein Orabi, and Diana Inkpen. 2018. Deep learning for depression detection of twitter users. In *Proceedings of the fifth workshop on computational linguistics and clinical psychology: from keyboard to clinic*, pages 88–97.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Javier Parapar, Patricia Martín-Rodilla, David E. Losada, and Fabio Crestani. 2022. [Overview of erisk at clef 2022: Early risk prediction on the internet \(extended overview\)](#). In *CLEF 2022: Conference and Labs of the Evaluation Forum, September 5–8, 2022, Bologna, Italy*, Bologna, Italy. CEUR Workshop Proceedings. CC BY 4.0; ISSN: 1613-0073.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Bénédicte Pincemin. 2022. Sémantique textométrique. *La sémantique au pluriel. Théories et méthodes*, pages 373–396.
- QwenTeam. 2024. [Qwen2.5: A party of foundation models](#).
- Nishat Raihan, Sadiya Sayara Chowdhury Puspo, Shafkat Farabi, Ana-Maria Bucur, Tharindu Ranasinghe, and Marcos Zampieri. 2024. [MentalHelp: A multi-task dataset for mental health in social media](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11196–11203, Torino, Italia. ELRA and ICCL.
- Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, volume 242, pages 29–48. Cite-seer.

- Alba M Mármol Romero, Adrián Moreno Muñoz, Flor Miriam Plaza Del Arco, M Dolores Molina-González, María Teresa Martín Valdivia, L Alfonso Urena Lopez, and Arturo Montejo Ráez. 2024. Mentalrises: A new corpus for early detection of mental disorders in spanish. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11204–11214.
- Haşim Sak, Andrew Senior, and Françoise Beaufays. 2014. Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition. *arXiv preprint arXiv:1402.1128*.
- André Salem. 1986a. [Segments répétés et analyse statistique des données textuelles](#). *Histoire & Mesure*, 1:5–28.
- André Salem. 1986b. Segments répétés et analyse statistique des données textuelles. *Histoire & mesure*, pages 5–28.
- Wesley Ramos dos Santos, Rafael Lage de Oliveira, and Ivandré Paraboni. 2024. Setembro: a social media corpus for depression and anxiety disorder prediction. *Language Resources and Evaluation*, 58(1):273–300.
- Jürgen Schmidhuber. 2015. [Deep learning in neural networks: An overview](#). *Neural Networks*, 61:85–117.
- Elizabeth M Seabrook, Margaret L Kern, Ben D Fulcher, and Nikki S Rickard. 2018. Predicting depression from language-based emotion dynamics: longitudinal analysis of facebook and twitter status updates. *Journal of medical Internet research*, 20(5):e168.
- Faisal Muhammad Shah, Farzad Ahmed, Sajib Kumar Saha Joy, Sifat Ahmed, Samir Sadek, Rimon Shil, and Md Hasanul Kabir. 2020. Early depression detection from social network using deep learning techniques. In *2020 IEEE region 10 symposium (TENSYP)*, pages 823–826. IEEE.
- Guangyao Shen, Jia Jia, Liqiang Nie, Fuli Feng, Cunjun Zhang, Tianrui Hu, Tat-Seng Chua, Wenwu Zhu, et al. 2017. Depression detection via harvesting social media: A multimodal dictionary learning solution. In *IJCAI*, pages 3838–3844.
- Tiancheng Shen, Jia Jia, Guangyao Shen, Fuli Feng, Xiangnan He, Huanbo Luan, Jie Tang, Thanassis Tiropanis, Tat Seng Chua, and Wendy Hall. 2018. Cross-domain depression detection via harvesting social media. In *Proceedings of the International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence.
- Ying Shen, Huiyu Yang, and Lin Lin. 2022. Automatic depression detection: An emotional audio-textual corpus and a gru/bilstm-based model. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6247–6251. IEEE.
- Jingzhe Shi, Qinwei Ma, Hongyi Liu, Hang Zhao, Jeng-Neng Hwang, and Lei Li. 2025. Explaining context length scaling and bounds for language models. *arXiv preprint arXiv:2502.01481*.
- Jingzhe Shi, Qinwei Ma, Huan Ma, and Lei Li. 2024. Scaling law for time series forecasting. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Pengcheng Shi, Jiawei Chen, Jiaqi Liu, Xinglin Zhang, Tao Chen, and Lei Li. Medal s: Spatio-textual prompt model for medical segmentation. In *CVPR 2025: Foundation Models for 3D Biomedical Image Segmentation*.
- Sajad Sotudeh, Hanieh Deilamsalehy, Franck Deroncourt, and Nazli Goharian. 2021. [TLDR9+: A large scale resource for extreme summarization of social media posts](#). In *Proceedings of the Third Workshop on New Frontiers in Summarization*, pages 142–151, Online and in Dominican Republic. Association for Computational Linguistics.
- Sajad Sotudeh, Nazli Goharian, and Zachary Young. 2022. [MentSum: A resource for exploring summarization of mental health online posts](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2682–2692, Marseille, France. European Language Resources Association.
- Matthew Squires, Xiaohui Tao, Soman Elangovan, Raj Gururajan, Xujuan Zhou, U Rajendra Acharya, and Yuefeng Li. 2023. Deep learning and machine learning in psychiatry: a survey of current progress in depression detection, diagnosis and treatment. *Brain Informatics*, 10(1):10.
- Maxim Stankevich, Ivan Smirnov, Natalia Kiselnikova, and Anastasia Ushakova. 2020. Depression detection from social media profiles. In *Data Analytics and Management in Data Intensive Domains: 21st International Conference, DAM-DID/RCDL 2019, Kazan, Russia, October 15–18, 2019, Revised Selected Papers 21*, pages 181–194. Springer.
- Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: Liwc and

- computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Marcel Trotzek, Sven Koitka, and Christoph M Friedrich. 2018a. Utilizing neural networks and linguistic metadata for early detection of depression indications in text sequences. *IEEE Transactions on Knowledge and Data Engineering*, 32(3):588–601.
- Marcel Trotzek, Sven Koitka, and Christoph M Friedrich. 2018b. Word embeddings and linguistic metadata at the clef 2018 tasks for early detection of depression and anorexia. In *CLEF (working notes)*.
- Qamar Un Nisa and Rafi Muhammad. 2021. Towards transfer learning using bert for early detection of self-harm of social media users. *Proceedings of the Working Notes of CLEF*, pages 21–4.
- A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Jifei Wang, Zhenping Zhao, Jing Yang, Limin Wang, Mei Zhang, and Maigeng Zhou. 2024a. The association between depression and all-cause, cause-specific mortality in the chinese population—china, 2010–2022. *China CDC Weekly*, 6(40):1022.
- Shenzhi Wang, Yaowei Zheng, Guoyin Wang, Shiji Song, and Gao Huang. 2024b. [Llama3-8b-chinese-chat \(revision 6622a23\)](#).
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Zuyi Bao, Jiangnan Xia, Liwei Peng, and Luo Si. 2019. Structbert: Incorporating language structures into pre-training for deep language understanding. *arXiv preprint arXiv:1908.04577*.
- Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. 2016. Attention-based lstm for aspect-level sentiment classification. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 606–615.
- Yiding Wang, Zhenyi Wang, Chenghao Li, Yilin Zhang, and Haizhou Wang. 2020a. [A multimodal feature fusion-based method for individual depression detection on sina weibo](#). In *2020 IEEE 39th International Performance Computing and Communications Conference (IPCCC)*, pages 1–8.
- Yiding Wang, Zhenyi Wang, Chenghao Li, Yilin Zhang, and Haizhou Wang. 2020b. [A multimodal feature fusion-based method for individual depression detection on sina weibo](#). In *2020 IEEE 39th International Performance Computing and Communications Conference (IPCCC)*, pages 1–8.
- Yuxi Wang, Diana Inkpen, and Prasadith Kirinde Gamaarachchige. 2024c. Explainable depression detection using large language models on social media data. In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 108–126.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#).
- JT Wolohan, Misato Hiraga, Atreyee Mukherjee, Zeeshan Ali Sayyed, and Matthew Millard. 2018. Detecting linguistic traces of depression in topic-restricted text: Attending to self-stigmatized depression with nlp. In *Proceedings of the first international workshop on language cognition and computational models*, pages 11–21.
- Min Yen Wu, Chih-Ya Shen, En Tzu Wang, and Arbee LP Chen. 2020. A deep architecture for depression detection using posting, behavior, and living environment data. *Journal of Intelligent Information Systems*, 54(2):225–244.
- Dong Xue\* Xin Yan. 2023. Mindchat: Psychological large language model. <https://github.com/X-D-Lab/MindChat>.
- Jinyuan Xu, Tian Lan, Mathieu Valette, Pierre Magistry, and Lei Li. 2025. [TinyMentalLLMs enable depression detection in Chinese social media texts](#). In *Proceedings of the 15th International Conference on Recent Advances in Natural Language Processing - Natural Language Processing in the Generative AI Era*, pages 1352–1363, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Xuhai Xu, Bingsheng Yao, Yuanzhe Dong, Saadia Gabriel, Hong Yu, James Hendler, Marzyeh Ghassemi, Anind K Dey, and Dakuo Wang. 2024. Mental-Ilm: 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, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,

- Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, et al. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Juhua Yang and Lulu Li. 2009. Intergenerational dynamics and family solidarity: A comparative study of mainland china, japan, korea and taiwan. *Sociological Studies*, 3:26–53.
- Kailai Yang, Shaoxiong Ji, Tianlin Zhang, Qianqian Xie, Ziyang Kuang, and Sophia Ananiadou. 2023. [Towards interpretable mental health analysis with large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6056–6077, Singapore. Association for Computational Linguistics.
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024b. [Mentallama: interpretable mental health analysis on social media with large language models](#). In *Proceedings of the ACM on Web Conference 2024*, pages 4489–4500.
- Tingting Yang, Fei Li, Donghong Ji, Xiaohui Liang, Tian Xie, Shuwan Tian, Bobo Li, and Peitong Liang. 2021. Fine-grained depression analysis based on chinese micro-blog reviews. *Information Processing & Management*, 58(6):102681.
- Wenhao Yang, Jianguo Wei, Wenhuan Lu, and Lei Li. 2025. You only speak once to see. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Xingwei Yang, Rhonda McEwen, Liza Robee Ong, and Morteza Zihayat. 2020. A big data analytics framework for detecting user-level depression from social networks. *International Journal of Information Management*, 54:102141.
- Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. A survey on large language model (llm) security and privacy: The good, the bad, and the ugly. *High-Confidence Computing*, page 100211.
- Ziyu Yao, Xuxin Cheng, Zhiqi Huang, and Lei Li. 2025. CountLLM: Towards Generalizable Repetitive Action Counting via Large Language Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. [Depression and self-harm risk assessment in online forums](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2968–2978, Copenhagen, Denmark. Association for Computational Linguistics.
- Amir Hossein Yazdavar, Mohammad Saeid Mahdavinejad, Goonmeet Bajaj, William Romine, Amit Sheth, Amir Hassan Monadjemi, Krishnaprasad Thirunarayan, John M Meddar, Annie Myers, Jyotishman Pathak, et al. 2020. Multi-modal mental health analysis in social media. *Plos one*, 15(4):e0226248.
- Hongbin Ye, Tong Liu, Aijia Zhang, Wei Hua, and Weiqiang Jia. 2023. [Cognitive mirage: A review of hallucinations in large language models](#).
- Yutaka Miyaji Yuka Niimi. 2021. [Machine learning approach for depression detection in Japanese](#). In *Proceedings of the 35th Pacific Asia Conference on Language, Information and Computation*, pages 346–353, Shanghai, China. Association for Computational Linguistics.
- Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer.
- Wei Zhai, Hongzhi Qi, Qing Zhao, Jianqiang Li, Ziqi Wang, Han Wang, Bing Yang, and Guanghui Fu. 2024. [Chinese MentalBERT: Domain-adaptive pre-training on social media for Chinese mental health text analysis](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 10574–10585, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net.
- Ye Zhang and Byron Wallace. 2015. A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification. *arXiv preprint arXiv:1510.03820*.
- Yipeng Zhang, Hanjia Lyu, Yubao Liu, Xiyang Zhang, Yu Wang, and Jiebo Luo. 2021. Monitoring depression trends on twitter during the covid-19 pandemic: observational study. *JMIR infodemiology*, 1(1):e26769.
- Zhiling Zhang, Siyuan Chen, Mengyue Wu, and Kenny Zhu. 2022. [Symptom identification for](#)

interpretable detection of multiple mental disorders on social media. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9970–9985, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand. Association for Computational Linguistics.

Han-yu Zhou, Wen-qi Zhu, Wen-yi Xiao, Ya-ting Huang, Kang Ju, Hong Zheng, and Chao Yan. 2023. Feeling unloved is the most robust sign of adolescent depression linking to family communication patterns. *Journal of Research on Adolescence*, 33(2):418–430.

Zhuoran Zhou, Zhongyu Jiang, Wenhao Chai, Cheng-Yen Yang, Lei Li, and Jenq-Neng Hwang. 2024. Efficient domain adaptation via generative prior for 3d infant pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.

## A. Appendix

### A.1. Baseline Models

- **DeepSeek-R1** is a large-scale model designed for complex reasoning. It performs well on tasks such as mathematics and programming, while remaining competitive on general-purpose benchmarks. We include it as a strong reasoning-oriented baseline for structured psychological analysis and depression risk assessment.
- **DeepSeek-R1-Distill-14B** is a distilled reasoning model derived from Qwen2.5-14B via knowledge distillation. It aims to retain key reasoning capabilities at lower computational cost. We include it as an efficient baseline for depression risk analysis.
- **Qwen2.5-14B** is a large language model from the Qwen series. It provides strong instruction following and long-form generation. We include it as a Chinese-capable baseline for generating structured psychological analyses from social media posts.
- **GPT-4o** is a multimodal large language model released by OpenAI. It supports multilingual generation, including Chinese. We include it as a strong general-purpose baseline for depression detection and psychological analysis.
- **GPT-4o-mini** is a smaller model in the GPT-4o family. It is designed for efficiency and lower inference cost. We include it as a lightweight baseline for depression risk identification and structured summarization.
- **Llama3-8B-Chinese-Chat**, derived from Meta-Llama-3-8B, is an instruction-tuned model for Chinese dialogue. It supports both Chinese and English inputs. We include it as a Chinese dialogue-oriented baseline for user-level depression risk analysis and generation.

### A.2. 20 Randomly Generated Expression for the Positive Users

1. Carefully read the complete user Weibo text below and extract key information across six specified dimensions according to the content. Also pay attention to the user’s depressive state, ensuring that each dimension’s explanation is consistent with this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if evidence exists in the text, provide the text index (e.g., text\_xx, 文本\_xx) and a brief note

- (e.g., text\_xx[sad]). - If no evidence is found for a dimension, state clearly that no relevant evidence was found.  
Note: Only output text index and explanation, not the actual text.
2. Read the entire user Weibo text below and, based on its content, extract core information under the following six dimensions. Also consider the user's depressive status, ensuring that each dimension's description matches this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if relevant information is detected, provide the text index (e.g., text\_xx, 文本\_xx) and a short note (e.g., text\_xx[unhappy]). - For dimensions without evidence, explicitly state no relevant evidence found.  
Note: Output should include only text index and explanation, not the actual text.
  3. Read through the complete user Weibo text below and extract key information across six dimensions based on its content. Pay close attention to the user's depressive state, ensuring that the explanations are consistent with it.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If relevant evidence is found in a dimension, provide the text index (e.g., text\_xx, 文本\_xx) and a short note (e.g., text\_xx[sorrow]). - If no evidence is found, indicate explicitly.  
Note: Provide only indices and explanations, not the original content.
  4. Carefully read the full Weibo text below and extract key evidence across six dimensions according to the text. At the same time, take into account the user's depressive state, ensuring that explanations for each dimension align with this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if there is corresponding evidence, output the text index (e.g., text\_xx, 文本\_xx) and a related note (e.g., text\_xx[frustrated]). - If no evidence is found, state that no evidence exists.  
Note: Only provide text indices and notes, not the raw text.
  5. Please read the full Weibo text below and extract key information from six dimensions. At the same time, pay attention to the user's depressive state, ensuring that your explanations match this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if evidence is detected, output the text index (e.g., text\_xx, 文本\_xx) and a short description (e.g., text\_xx[sadness]). - If no evidence is found, specify that no evidence was found.  
Note: Only output text indices and explanations, not the original text.
  6. Read the complete Weibo text below carefully and extract key information across six dimensions based on the text. Pay attention to the user's depressive state and ensure that your descriptions are consistent with it.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If relevant evidence is found, provide the text index (e.g., text\_xx, 文本\_xx) and a brief description (e.g., text\_xx[low mood]). - If no evidence is found, clearly state that no evidence exists.  
Note: Only provide text indices and explanations, not the original text.
  7. Please read through the complete user Weibo text below and extract key information across six dimensions based on its content. Pay attention to the user's depressive state, ensuring that each explanation is aligned with this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if evidence exists, provide the text index (e.g., text\_xx, 文本\_xx) and a short note (e.g., text\_xx[loss]). - If no evidence is found, clearly state no relevant evidence.  
Note: Only provide text indices and notes, not the original text.
  8. Read the Weibo text below and extract key evidence across six dimensions according to the content, while considering the user's depressive state to ensure explanations are consistent.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If a dimension contains relevant evidence, provide the text index (e.g., text\_xx, 文本\_xx) and a short note (e.g., text\_xx[low mood]). - If no evidence is found, state explicitly that no evidence exists.  
Note: Output only text indices and explanations, not the original text.
  9. Read the complete Weibo text carefully and extract key evidence across six dimensions based on the text. Pay attention to the user's depressive state, ensuring that explanations match this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.

Weibo text: {text}.

**[Output requirements]** - For each dimension, if relevant evidence is found, output the text index (e.g., text\_xx, 文本\_xx) and a brief explanation (e.g., text\_xx[depressed]). - If no evidence is found, state explicitly that there is none.

Note: Provide only text indices and notes, not the original text.

10. Read through the entire Weibo text and extract core information across six dimensions according to the content. Pay attention to the user's depressive state and ensure that the descriptions are consistent with this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if relevant evidence exists, output the text index (e.g., text\_xx, 文本\_xx) and a short explanation (e.g., text\_xx[pessimistic]). - If no evidence is found, state clearly that no evidence exists.  
Note: Only output text indices and explanations, not the original content.
11. Please read the user Weibo text below completely and extract essential information across six dimensions. Ensure that the explanations for each dimension align with the depressive state given.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If relevant evidence is detected, list the text index (e.g., text\_xx, 文本\_xx) and a short description (e.g., text\_xx[melancholy]). - If no evidence is found, specify that no evidence was detected.  
Note: Output should only contain indices and notes, not the actual text.
12. Read carefully the full Weibo text provided and extract key evidence from six dimensions. Pay attention to the user's depressive state, ensuring that outputs remain consistent with this state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If a dimension contains evidence, output the text index (e.g., text\_xx, 文本\_xx) and a concise note (e.g., text\_xx[disheartened]). - If no evidence exists, specify so.  
Note: Output indices and explanations only, not original content.
13. Please go through the complete Weibo text below and identify critical information across six dimensions. Ensure the notes correspond with the user's depressive condition.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - Where evidence is found, list the text index (e.g., text\_xx, 文本\_xx) with a short explanation (e.g., text\_xx[downhearted]). - If no evidence is detected, indicate explicitly.  
Note: Provide only text indices and notes, excluding original text.
14. Read carefully the following Weibo text and extract evidence across six preset dimensions. Make sure the explanation for each aligns with the depressive state.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If evidence exists for a dimension, output the text index (e.g., text\_xx, 文本\_xx) and a short remark (e.g., text\_xx[low mood]). - If no evidence is found, specify clearly.  
Note: Output only indices and remarks, not the raw text.
15. Read through the following Weibo text and extract major evidence across six given dimensions. Ensure the output is consistent with the depressive state specified.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If evidence is found, provide the text index (e.g., text\_xx, 文本\_xx) with a brief note (e.g., text\_xx[hopeless]). - If no evidence appears, state clearly.  
Note: Only provide indices and notes, not original text.
16. Please carefully review the Weibo text below and extract key information from six preset dimensions. All notes must align with the depressive condition of the user.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - If there is evidence, output the corresponding text index (e.g., text\_xx, 文本\_xx) and a short remark (e.g., text\_xx[emotional slump]). - If not found, specify clearly that none exists.  
Note: Only include indices and notes, no original content.
17. Please go through the complete Weibo text and extract the critical information across six given dimensions. Make sure your notes are consistent with the depressive state provided.  
**[Input]** User depressive state: {label}, Full Weibo text: {text}.  
**[Output requirements]** - For each dimension, if evidence is found, provide the text index (e.g., text\_xx, 文本\_xx) with a concise explanation

(e.g., text\_xx[feeling low]). - If no evidence exists, indicate so.

Note: Only output indices and explanations, not actual text.

18. Carefully read the Weibo text provided and extract important evidence across six specified dimensions. Keep the explanations aligned with the depressive state.

**[Input]** User depressive state: {label}, Full Weibo text: {text}.

**[Output requirements]** - If a dimension contains evidence, state the text index (e.g., text\_xx, 文本\_xx) and a short remark (e.g., text\_xx[upset]). - If not, state that no evidence is found.

Note: Output only indices and remarks, not text.

19. Read the Weibo text completely and extract relevant evidence from six dimensions. Ensure each dimension's note is consistent with the depressive state given.

**[Input]** User depressive state: {label}, Full Weibo text: {text}.

**[Output requirements]** - If relevant evidence is detected, provide the text index (e.g., text\_xx, 文本\_xx) and short description (e.g., text\_xx[distressed]). - If not found, state clearly no evidence.

Note: Only indices and notes, exclude text itself.

20. Please read through the following full Weibo text and extract essential evidence across six preset dimensions. Ensure all notes are consistent with the depressive state.

**[Input]** User depressive state: {label}, Full Weibo text: {text}.

**[Output requirements]** - If evidence is found, provide the text index (e.g., text\_xx, 文本\_xx) and a short explanation (e.g., text\_xx[worry]). - If no evidence is detected, state that explicitly.

Note: Only indices and explanations, not original text.

### A.3. 20 Distinct Label Expressions for the Negative Users

1. "This text does not contain any expressions of negative emotions, nor does it involve any dimensions of depression."
2. "There are no descriptions related to negative emotions in the text, so it does not fit any depression-related dimensions."
3. "This text does not include expressions of negative emotions, and therefore it is not classified under any depression dimensions."

4. "The content does not reflect negative emotions and does not belong to any depression indicators."

5. "This passage contains no expressions of negative emotions and is not applicable to any depression dimensions."

6. "There are no descriptions of negative emotions in the text, so it is not classified under any depression-related dimensions."

7. "This content does not exhibit negative emotions and therefore does not involve any depression dimensions."

8. "The text does not contain any expressions related to negative emotions, nor does it meet the criteria for depression dimensions."

9. "There are no signs of negative emotions in this text, so it does not belong to any depression dimension."

10. "The text does not describe negative emotions and does not cover any depression-related dimensions."

11. "This passage does not contain any expressions related to negative emotions and is not within the scope of depression dimensions."

12. "The content does not include expressions of negative emotions, so it does not fall into any depression dimensions."

13. "There are no expressions of negative emotions in the text, nor does it meet the criteria for depression dimensions."

14. "This text does not exhibit any negative emotions and does not involve any depression dimensions."

15. "This content does not include descriptions of negative emotions and therefore is not classified under any dimensions of depression."

16. "The text does not show any negative emotions and is not applicable to depression-related dimensions."

17. "No expressions of negative emotions are seen in this passage, so it does not meet any depression dimension standards."

18. "There are no expressions of negative emotions in the text, so it does not belong to any depression dimension."

19. "This text does not display any expressions related to negative emotions and does not cover the dimensions of depression."

20. "The content lacks descriptions of negative emotions and is therefore not classified under any depression dimensions."

#### **A.4. Instructions for the Fine-Tuning Process**

- **Complete Instruction = Task Instruction + Supplementary Instruction**

##### **Task Instruction:**

This task involves the following 6 dimensions of depression:

- Potential External Causes of Depression (Secondary Judgment Criterion)
- Depression-Related Clinical Symptoms (Primary Judgment Criterion)
- Depression-Related Language Expression Patterns (Secondary Judgment Criterion)
- Depression-Related Medical Expressions (Primary Judgment Criterion)
- Depressive Psychological State (Primary Judgment Criterion)
- Negative Emotions (Secondary Judgment Criterion)

##### **Supplementary Instructions (randomly select one):**

1. "Please read the following text segment and determine whether it contains any of the above 6 expressions of depressive emotion, and provide a brief explanation; if none is present, please indicate that the text does not belong to any depression dimension."
2. "Read the following text and analyze whether it demonstrates any one of the aforementioned 6 depressive emotion expression dimensions, and provide a brief explanation; if such expression is absent, please indicate that the text does not correspond to any dimension."
3. "Please carefully read the following text segment and confirm whether any one of the 6 depressive emotion expression dimensions mentioned above exists, and include a brief explanation; if not, please state that the text does not belong to any dimension."
4. "Please read the following content and determine whether it embodies any one of the above 6 depressive emotion expression dimensions, and provide a brief explanation; if not, please indicate that the text does not cover any depression dimension."
5. "Read the following text segment and determine whether it contains any one of the 6 depressive emotion expressions mentioned above, and provide a concise explanation; if not, please state that the text does not belong to any depression dimension."
6. "Please read the following text and determine whether any one of the above 6 depressive emotion expression dimensions appears, and include a brief explanation; if not, please indicate that the text does not meet any dimension."
7. "Read the following text and check whether any one of the above 6 depressive emotion expression dimensions is presented, and provide a brief explanation; if not, please state that the text does not belong to any dimension."
8. "Please review the following text segment and determine whether any one of the above 6 depressive emotion expression dimensions exists, and provide a brief explanation; if not, then indicate that the text does not belong to any dimension."
9. "Read the following text and confirm whether any one of the above 6 depressive emotion expression dimensions is embodied, and provide a brief explanation; if not, please state that the text does not involve any depression dimension."
10. "Please read the following content and determine whether it contains any one of the 6 depressive emotion expression dimensions mentioned earlier, and provide a brief explanation; if not, please indicate that the text does not belong to any dimension."
11. "Read the following text segment and verify whether it demonstrates any one of the above 6 depressive emotion expression dimensions, and include a brief explanation; if not, please state that the text does not meet any dimension."
12. "Please read the following text and check whether it exhibits any one of the above 6 depressive emotion expression dimensions, and provide a brief explanation; if not, please indicate that the text does not cover any depression dimension."
13. "Read the following text and confirm whether any one of the 6 depressive emotion expressions mentioned above exists, and include a brief explanation; if such a feature is absent, please state that the text does not belong to any dimension."

14. "Please carefully read the following text segment and determine whether it embodies any one of the above 6 depressive emotion expression dimensions, and provide a concise explanation; if not, please indicate that the text does not correspond to any dimension."
15. "Read the following text segment and confirm whether it contains any one of the above 6 depressive emotion expression dimensions, and provide a brief explanation; if not, please state that the text does not involve any depression dimension."
16. "Please read the following text segment and determine whether any one of the 6 depressive emotion expression dimensions mentioned above appears in the text, and include a brief explanation; if not, please indicate that the text does not belong to any dimension."
17. "Read the following content and check whether it contains any one of the above 6 depressive emotion expression dimensions, and provide a brief explanation; if not, please state that the text does not belong to any depression dimension."
18. "Please read the following content and determine whether any one of the above 6 depressive emotion expression dimensions exists, and provide a brief explanation; if not, please indicate that the text does not involve any dimension."
19. "Read the following text segment and confirm whether it demonstrates any one of the above 6 depressive emotion expression dimensions mentioned above, and provide a brief explanation; if not, please state that the text does not meet any dimension."
20. "Please read the following text and determine whether it contains any one of the above 6 depressive emotion expression dimensions, and provide a concise explanation; if not, please indicate that the text does not belong to any depression dimension."

**Chinese-RoBERTa-wwm-ext** (Cui et al., 2019) is based on RoBERTa (Liu et al., 2019b) and uses whole word masking (WWM). It masks complete words rather than individual characters during MLM. This design can better capture word-level semantics in Chinese. The extended pre-training ("ext") on large Chinese corpora further improves robustness across NLP tasks.

**StructBERT-mental** (Wang et al., 2019) builds on StructBERT, which incorporates structural objectives during pre-training. It emphasizes word- and sentence-level structure. In addition, it is adapted for mental health text, which can benefit mental health-related classification.

## A.5. BERT-based Models

**BERT-base-chinese**<sup>8</sup> follows the BERT-Base architecture (Devlin et al., 2019) and is pre-trained on large-scale Chinese corpora. It uses masked language modeling (MLM) and next sentence prediction (NSP). The model learns bidirectional representations for Chinese text, typically at the character level.

<sup>8</sup><https://huggingface.co/google-bert/bert-base-chinese>

## A.6. Classification Performance Comparison

Model/Test Dataset	Swdd Origin		Wu3d		CNSD (Gold)	
	Acc	F1	Acc	F1	Acc	F1
RoBERTa Chinese (trained by swdd origin)	0.89	0.90	0.76	0.80	0.88	0.87
RoBERTa Chinese (trained by wu3d)	0.88	0.87	0.90	0.91	0.82	0.78
Bert based Chinese (trained by swdd origin)	0.89	0.90	0.76	0.80	0.85	0.82
Bert based Chinese (trained by wu3d)	0.88	0.87	0.90	0.91	0.83	0.80
StructBERT-mental (trained by swdd origin)	0.81	0.84	0.69	0.76	0.85	0.86
StructBERT-mental (trained by wu3d)	0.85	0.85	<b>0.92</b>	<b>0.92</b>	0.83	0.80
Llama3-8b	0.86	0.86	0.79	0.79	0.60	0.63
Glm4-9b-chat	0.76	0.81	0.77	0.81	0.79	0.82
Qwen2.5 7b	0.88	0.89	0.84	0.85	<b>0.94</b>	<b>0.94</b>
GPT-4o-Mini	0.65	0.74	0.65	0.74	0.53	0.68
GPT-4o	<b>0.92</b>	<b>0.92</b>	0.91	0.92	0.83	0.86
DeepSeek-R1 671b	0.86	0.87	0.89	0.90	0.81	0.84

Table 7: Model Classification Performance Comparison on Different Test Datasets.

## A.7. Module II Prompt

You are a professional psychologist and text sentiment analysis expert. Below is a set of sentence-by-sentence analysis results generated by a smaller model concerning all of the user's Weibo posts, each result presented in the format "Text Label: Analysis Content." Please use your professional knowledge to review, filter, and correct these analysis contents, and then produce a final comprehensive analysis report.

### [Task Requirements]

1. **Review and filter** the sentence-by-sentence analysis provided by the small model, selecting those text entries that are useful for generating the final report.
2. **Check** whether there is any classification error within these entries. If you discover that the small model has incorrectly assigned certain texts to an incorrect dimension (for example, classifying "unhappy" incorrectly under depression-related medical expressions), please **correct** it based on the actual semantics and psychological knowledge, reassign it to a more reasonable dimension, or exclude that erroneous information from the final report.
3. The user's depression status is [already known depression label]. Please ensure that the final comprehensive analysis report is consistent with this label:
  - If the label is "depressed," each dimension in the report should **fully reflect** evidence related to depression;
  - If the label is "non-depressed," then the report should **not contain** extensive mentions of depressive features.
4. In your descriptions, please **avoid using specific time quantifiers** (such as "≥2 weeks"), and instead use a neutral description such as "multiple occurrences."
5. The final report should be organized under the following six dimensions:
  - Negative emotions (secondary judgment criterion)
  - Depressive psychological state (primary judgment criterion)
  - Depressive clinical symptoms (primary judgment criterion)
  - Potential external causes of depression (secondary judgment criterion)
  - Depression-related medical expressions (primary judgment criterion)
  - Depression-related language expression patterns (secondary judgment criterion)

In each dimension, please cite the corresponding text label and provide a concise explanation, ensuring that all classifications in the final report are your corrected judgments. The output must not contain any references to internal corrections (such as "already removed" or "already fixed") or references to specific error rates.

### [Reference Gold Standard Answers Example]

#### [Positive Example]

{positive\_example}

#### [Negative Example]

{negative\_example}

### [Sample Input Data]

User's known depression label: **{{label}}**

Small model's sentence-by-sentence analysis results:

{input}

### [Task]

Please, based on the requirements above, **review, filter, and correct** the sentence-by-sentence analysis provided by the small model, and then produce a **final comprehensive analysis report**. The report must be **organized under six dimensions**, and in each dimension, please cite the corresponding text label along with a concise explanation. **Ensure** that the final report, after your corrections, is consistent with the user's known label **[[already known depression label]]**, and that it **directly reflects** the corrected chain of evidence and judgment conclusions, **without** including internal correction statements or specific time descriptors.

Now, please produce the final comprehensive analysis report.

Figure 4: Module II Prompt.