INDIAN JOURNAL OF SCIENCE AND TECHNOLOGY

ORIGINAL ARTICLE

OPEN ACCESS

Received: 21/05/2025

Accepted: 18/08/2025

Published: 08/09/2025

Citation: Nair DS, Lalithambigai M (2025) ACSAT: An Adaptive Contextual Sentiment-Aware Transformer for Early Detection of Depression in Social Media Text. Indian Journal of Science and Technology 18(34): 2752-2760. https://doi.org/10.17485/IJST/v18i34.950

* **Corresponding author**.

divyasukumaran1989@gmail.com

Funding: None

Competing Interests: None

Published By Indian Society for Education and Environment (iSee)

ISSN
Print: 0974-6846
Electronic: 0974-5645

ACSAT: An Adaptive Contextual Sentiment-Aware Transformer for Early Detection of Depression in Social Media Text

Divya S Nair[1]*, M Lalithambigai[1]

1 Department of Computer Science, Sri Krishna Adithya College of Arts & Science (Co-Education), Coimbatore, Tamil Nadu, India

Abstract

Objectives: Prompt and correct recognition of depression from social media posts presents an economic, non-invasive method of managing mental health. However, the current models typically perform poorly on very subtle contextual and emotional hints, resulting in not-so-optimal detection, especially about severity levels. **Methods:** This study introduces ACSAT—an Adaptive Contextual Sentiment-Aware Transformer—that endorses a transformer encoder, a sentiment-aware attention model, and severity-sensitive contrastive learning, which uses margin-weighted separation and balanced sampling to capture both categorical differences and the ordinal progression of depression severity. The model is trained on user posts to label depression into three categories: None, Moderate, and Severe. The model weighs emotionally meaningful tokens adaptively and learns fine-grained distinctions in the linguistic presentation of depressive symptoms. **Findings:** The ACSAT model was tested on the Mental Health Disorder Detection Dataset from Kaggle; the data was processed using stratified fold cross validation to make for a robust experiment. The model achieved an accuracy of 96%, precision of 94%, recall of 95%, and F1-score of 92%, consistently beats all baseline models such as Support Vector Machines (SVM), Random Forests (RF), and classical NLP classifiers. The enhanced performance is mainly driven by ACSAT's dual sensitivity to both contextual semantics and emotional nuances. All baselines were tuned using the same preprocessing, features, and stratified fold splits with nested cross-validation for hyper-parameter tuning and testing, in order to compare the models fairly without test leakage. **Novelty:** In contrast to previous models, ACSAT uses contextual embeddings combined with affect-sensitive attention and applies depression severity-specific contrastive learning better to identify subtle variations in effect between social media posts. ACSAT is thus a strong, generalizable, early and scalable depression detection method.

Keywords: Depression detection; Mental health prediction; Social media mining; Contrastive learning; Text classification; Psychological disorder detection; Emotion-aware AI

1 Introduction

Depression is a disabling and prevalent mental disorder that impacts over 264 million individuals worldwide. Depression is typically underdiagnosed or diagnosed late because conventional diagnosis by clinical interview and self-report questionnaires is slow and subjective. While social media has become an increasingly popular means of expression, there lies potential in utilizing Artificial Intelligence (AI) to identify depression symptoms in a cost-effective, scalable, and non-invasive process.

Recent advancements in Deep Learning (DL) and Machine Learning (ML) have allowed for design of systems that can identify complex patterns found in textual, physiological, and clinical data. For instance, a systematic review[(1)] indicated the proliferation of ML/DL for depression detection in social media, point to the growing emphasis of research in this field. For example,[(2)] used DL models for the prediction of treatment response in patients suffering from depression, whereas[(3)] used large volumes of social media text to identify symptoms of depression. There have been other studies that investigated multimodal data fusion with audio, video, and physiological signals[(4)] or constructed prediction models on Electronic Health Records (EHRs) for predicting post-stroke depression[(5)].

The use of user-posted social media data has also been extremely promising.[(6)],[(7)] constructed DL models that successfully detect depressive behavior on tweets and Chinese microblogs with better performance than conventional classifiers. Also, anxiety and stress detection as comorbid with depression has been proposed through ensemble ML methods[(8)],[(9)]. Context-aware and transformer models, including the model developed by[(10)], also complement user context in task prediction of mental health.

These models are also bound to be limited in their ability to grasp the subtlety of emotion and linguistic fine-grained patterns characteristic of depressed language. Most of these models are geared towards binary classification (not depressed/depressed) and not towards the significance of recognizing severity levels, such as Moderate or Severe depression. Sentiment analysis that is popularly used is often also constrained to being an auxiliary task and not deeply embedded in the model's decision-making.

Previous approaches may fall short in capturing emotional nuance for different reasons. Some models use generic embeddings (e.g. word2vec or GloVe) that do not expose subtle affective cues, while other models do not leverage affective lexicons or sentiment features and are thus less capable of finding emotionally informative expressions altogether. Additionally, shallow classifiers or any basic CNN/RNN layers restrict their ability to model contextual effects around emotional change. These examples illustrate some of the reasonings behind the framework ACSAT, which uses sentiment-aware attention and severity-aware contrastive learning to capture emotional nuances in language related to depression more effectively. ACSAT applies severity-aware contrastive learning; a layer of complexity over standard contrastive learning, in that it weights separation margins according to severity levels of depression classification. Instead of treating all classes equally relative to one another, such as in standard contrastive learning (ERC time-frame), the severity-aware method means that we require greater separation between distant classes (None and Severe), and smaller separation between closer classes (e.g. None and Moderate). We are reflective and were cautioned by the ordinal aspect of depression severity level in not assuming the same equal spaces between the examples but using a more subjective approach to similar patient or facet levels together, therefore reflects the nature of how we deal with questions of depression severity.

Research Gap:

Despite impressive advancements, the following shortcomings are noted in the current literature:

- Lack of emotion integration: Current models do not have emotion-aware mechanisms that assign priority to emotive tokens essential for depression detection.
- Insufficient state granularity in classification: The Majority of methods engage in binary classification without investigating intermediate or extreme depressive states.
- Lack of contextual consideration: Models tend to lack the capability to monitor fine-grained changes in language use or implied emotion, particularly in noisy social media streams.
- Insufficiency of severity-aware representation learning: Contrastive learning approaches tailored for depression severity remain less explored in existing literature.

They collectively diminish the capability of current systems for achieving accurate, sensitive, and early-stage depression detection, particularly in real-world scenarios using multiform language expressions in addition to shifting emotional tone.

Contributions:

To overcome the issues above, this paper proposes ACSAT—an Adaptive Contextual Sentiment-Aware Transformer model that:

- Uses a sentiment-aware attention mechanism to pay attention to affectively important tokens in user-submitted social media content.

- Utilizes a severity-aware contrastive learning approach to predict between None, Moderate, and Severe depression categories.
- Outperforms baseline models like SVM, Random Forest, and standard NLP classifiers (96% accuracy, 94% precision, 95% recall, 92% F1-score).
- Provides a context-sensitive, non-invasive, and scalable solution for real-time mental health monitoring, ready for implementation in digital mental health platforms and public health programs.

These limitations together are obstacles to the provision of present systems in accurate, sensitive, and early-onset depression detection, particularly under real-world conditions involving multilingual modes of language and variable emotional tone.

Organization: This paper is structured such that Section 2 explains the dataset, preprocessing steps, and model configuration. Section 3 presents the experimental setup, outcomes, and performance analysis. Section 4 provides a discussion on results, limitations, and future work directions. In section 5, the references are cited.

Table 1. Contribution Comparison of various authors' work

Data Type	Techniques Used	Limitations	Key Contribution	How ACSAT Advances Over This Work
Neuroimaging (fMRI) (11)	Random Forest, feature extraction	Expensive, non-scalable modality; not text-based	Diagnosed with anxious depression using imaging biomarkers	Uses low-cost, scalable text input; adaptable to social media context
Clinical & Demographic (12)	ML models (unspecified)	No emotional or linguistic analysis; binary prediction	Identified postpartum depression risk	Introduces context- and sentiment-aware modeling with severity grading
Social media comments (13)	ML classifiers (e.g., NB, SVM)	Ignores context and sentiment; lacks severity-level detection	Applied ML on user comments for depression screening	Uses an advanced transformer with attention and contrastive learning
Brain Imaging (MRI) (14)	DL (segmentation networks)	Invasive method, limited to clinical settings	Detected structural markers of depression	ACSAT is non-invasive, text-based, and deployable in real-time systems
Clinical + Psychosocial (15)	ML with feature engineering	No language-based emotion recognition; no real-time use	Built a predictive model for antenatal depression	Enables early detection via daily digital footprint (social media text)
Social Media Text (16)	Transformer + Contrastive Learning	Binary classification only; lacks personalized expression modeling	Sentiment-aware transformer for depression detection	ACSAT improves with multi-class prediction and deeper sentiment-context fusion
Social Media Text (17)	Transformer + Auxiliary Features	No severity-level detection; limited emotional attention	Used auxiliary features to improve BERT performance	ACSAT integrates direct affective signal modeling with contrastive learning
Chat logs, chatbot input (18)	BERT + Psychological First Aid	Focused on chatbot design, not fine-grained classification	Detected emotional distress in conversations	ACSAT specializes in clinical-level detection from passive user text
Social Media Text (19)	Transformer-based classifiers	No contrastive learning or sentiment attention	Detected depression severity from posts	ACSAT adds an adaptive fusion of emotional and contextual semantics

2 Methodology

The research methodology used in this work is centered on designing a strong and smart depression prediction model by the adaptation of the Adaptive Contextual Sentiment-Aware Transformer (ACSAT) model. This part outlines the selection of the dataset, preprocessing, model design, and evaluation plans. Natural Language Processing (NLP) and DL transformers are combined to extract semantic context and emotional signals from user-generated text data. Baseline comparison models like SVM, RF, and plain vanilla NLP classifiers are utilized in performance benchmarking. The measure of establishing key performance indicators, like accuracy, precision, recall, and F-measure, is utilized in determining the efficiency of the proposed model to diagnose depression correctly.

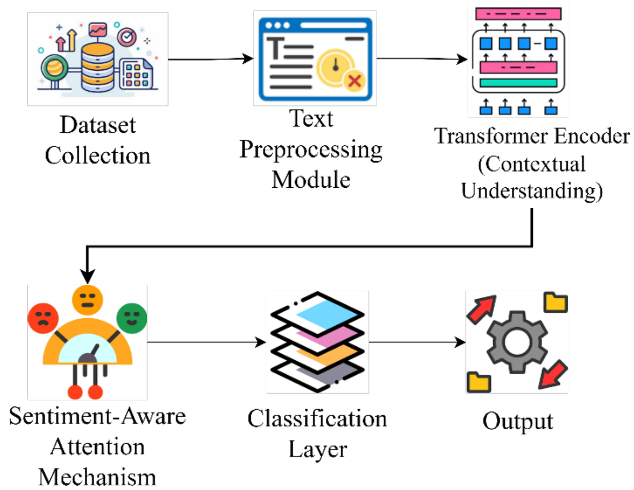

Fig 1. Overall Architecture

Figure 1 is the general structure of the Adaptive Contextual Sentiment-Aware Transformer (ACSAT) depression prediction model. The process starts with the acquisition of text data, which is fed into a text preprocessing module to preprocess and tokenize the data. The preprocessed inputs are then fed into a transformer encoder that acquires contextual semantics. Simultaneously, a sentiment-aware attention module computes the emotional importance of every token, giving more importance to highly emotionally important words. Both modules' outputs are concatenated in the depression severity level prediction classification layer (None, Moderate, or Severe) to yield the final output for every user input.

2.1 Dataset Description

Dataset Link: https://www.kaggle.com/datasets/nazmussakibrupol/mental-health-disorder-detection-dataset

The Mental Health Disorder Detection dataset (Kaggle) represents expert-annotated text samples indicating symptoms of various mental health conditions. The dataset includes seven clinical disorder categories and provides high-quality labeled data for supervised learning. This dataset has been constructed for use with Natural Language Processing (NLP) based classification tasks and enables robust model evaluation for depression and associated mental health detection. The dataset does not describe their exact annotations for the depression severity labels. Although useful for a benchmark for methodological evaluation purposes, these restrictions should be considered when making judgments on results generalizability.

There is a class imbalance in the None, Moderate, and Severe categories of the dataset. To compensate, 5-fold stratified cross-validation was used to preserve class distributions, class-weighted loss functions were implemented, and mini-batches were balanced between severity levels at training time.

2.2 Adaptive Contextual Sentiment-Aware Transformer (ACSAT)

Adaptive Contextual Sentiment-Aware Transformer (ACSAT) is a new state-of-the-art advancement in the identification of depression using social media analysis. It combines sentiment analysis and transformer-based models for better detection of depressive states based on text information. The model harnesses the strength of pre-trained language models and identifies weak contextual and affective signals implied in user posts. The model uses an adaptive process that continuously adjusts itself based on varying intensities of sentiment so that it can differentiate more effectively between non-depressed, moderately depressed, and severely depressed subjects. Adaptability is critical because it captures the depressive symptom continuum and the varied linguistic expression of symptoms. Through the addition of sentiment-aware features, ACSAT can more effectively distinguish between linguistically similar patterns that can manifest different emotional states and, therefore, reduce misclassification rates. The application of severity-aware contrastive learning enhances the capacity of the model to learn fine distinctions in depression severity. Experimental studies have validated that ACSAT is better than baseline transformer models due to the fact that it delivers higher macro F1 scores in depression severity detection tasks. The breakthrough highlights the importance of incorporating sentiment and context awareness into ML models for mental health tasks.

The creation of ACSAT aligns with increasing awareness regarding social networking websites as an efficient means for information monitoring of mental health, especially an extensible and noninvasive method of early detection and intervention techniques. With specific accuracy in people at risk, ACSAT helps their timely assistance and interventions for depressive patients. The model's architecture also acknowledges the value of individualized assessment, taking into account variations in language use and affect expression between individuals. The individualized treatment makes the model more robust across populations and linguistic styles. The success of ACSAT further identifies the value of integrating state-of-the-art NLP technology with psychological knowledge to treat complex mental disorders. Future studies can investigate the application of ACSAT in other psychological illnesses to broaden its use in the mental health diagnosis and support network.

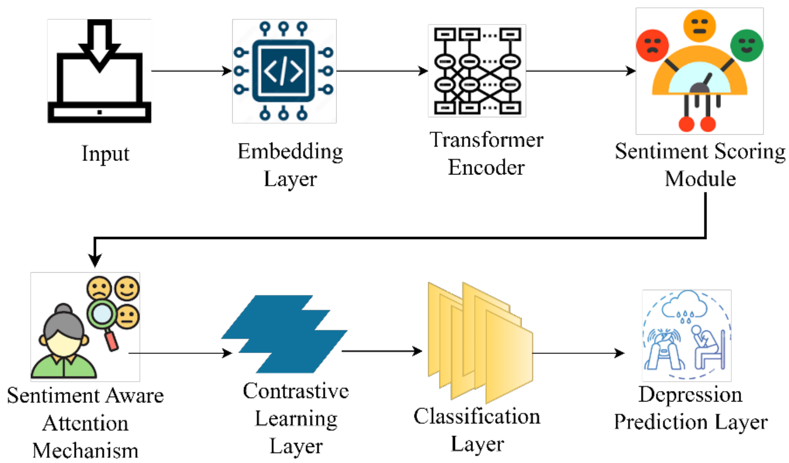

Fig 2. Architecture of Adaptive Contextual Sentiment-Aware Transformer (ACSAT)

Figure 2 depicts the in-depth structure of the Adaptive Contextual Sentiment-Aware Transformer (ACSAT) model used for predicting depression. The process begins with text input, which is converted into semantic vector space representations in the Embedding Layer. Such embeddings are given as input to a Transformer Encoder, which goes deep into contextualized relationships in the text. At the same time, the Sentiment Scoring Module provides tokens with emotional weights so that the Sentiment-Aware Attention Mechanism can attend to emotionally relevant information. The mechanism also allows the model to pay more attention to depression words more effectively. A Contrastive Learning Layer enhances representations by separating various levels of depression severity. Lastly, the processed context is input into a Classification Layer that provides the depression prediction result, labeled as none, moderate, or severe. Such a structure well integrates linguistic context and affective sentiment to enhance depression detection accuracy.

$$E_i = Embed(w_i) \tag{1}$$

Equation (1) is the conversion of a word w_i into a high-dimensional vector representation from a pre-trained embedding model. This conversion captures the semantic meaning of the word, allowing the model to understand language contextually.

$$H = Transformer(E) \tag{2}$$

Equation (2) is also called passing the word sequence embedded into a Transformer encoder. This embeds contextual relationships between words to allow the model to learn how each meaning of a word changes depending on its context.

$$\alpha_i = softmax(S \bullet H_i) \tag{3}$$

Equation (3) is a formula for sentiment-sensitive attention weight for token i. In this case, the contextual embedding H_i is combined with the sentiment score S and normalized via softmax to highlight emotive tokens for depression detection.

$$C = \sum_i \alpha_i \bullet H_i \tag{4}$$

Equation (4) computes the final context vector by summing all token representations. H_i together, each of which is weighted by its sentiment-sensitive attention score α_i. The vector C preserves the emotionally and contextually significant information of the whole text.

$$L_{contrastive} = \sum_{i,j} yij \bullet \|C_i - C_j\|^2 \tag{5}$$

Equation (5) defines contrastive loss, which pushes similar samples of the same depression severity (whereyij) to be similar shares context vectors C_i andC_j. It pushes the model to differentiate between depression levels more by maximizing inter-class distance and minimizing intra-class distance.

To account for the ordering that naturally exists among depression levels, ACSAT uses a margin-weighted supervised contrastive loss. Posts with similar severity levels are encouraged to come closer in the embedding space while creating separation margins for posts when the severity classification levels vary (e.g., more separation is enforced between the None and Severe levels than None and Moderate). Mini-batches where the severity levels are balanced allow for a good representation of both positive and negative pairs. This approach ensures that the embedding space will encode both categorical contrasts and ordinal increase in depression severity; thus, providing better classification for the downstream task.

$$\hat{y} = softmax(W \bullet C + b) \tag{6}$$

Equation (6) is the last classification step in which the context vector C is fed into a linear layer with weights W and bias$b$. It is converted to probabilities in depression severity classes by applying the softmax function, and the maximum value is the prediction of the class.

Algorithm 1: Adaptive Contextual Sentiment-Aware Transformer (ACSAT)
Input: User Text (list of text posts)
Output: Depression Label (None, Moderate, Severe)
Step 1: Preprocess each post in User Text:
  Clean text (remove links, emojis, etc.)
  Tokenize into words
Step 2: For each token in each post:
  Get word embedding: E = Embed(token)
Step 3: Pass the sequence of embeddings through the Transformer:
  H = Transformer(E) // H = contextual features
Step 4: For each token:
  Get sentiment score: S = Sentiment Score (token)
Step 5: Compute sentiment-aware attention weights:
  For each i: alpha[i] = softmax (S[i] * H[i])
Step 6: Create context vector:
  C = sum(alpha[i] * H[i] for all i)
Step 7: (Optional) Use contrastive loss to better separate severity levels
Step 8: Predict depression level:
  y_pred = softmax(DenseLayer(C))
Step 9: Return DepressionLabel = argmax (y_pred)

Algorithm 1 and Figure 3 show the flow of the Adaptive Contextual Sentiment-Aware Transformer (ACSAT) depression prediction model. The process is initiated with the user-written content (UserText), which gets tokenized and embedded to map each word to a high-dimensional vector through pre-trained models. The embeddings get fed into the transformer encoder in order to acquire contextual semantics. At the same time, sentiment scores are calculated for every token, and these are then utilized to calculate sentiment-aware attention weights using a softmax function. The attention weights identify sentimentally salient tokens and produce a context vector that contains both semantic and affective information. The classification layer predicts the level of depression from the context vector by inputting the context vector and receiving the depression label as the maximum probability output.

3 Results and Discussion

This section presents a rigorous comparative evaluation of the novel ACSAT model with the conventional ML techniques SVM, RF, and NLP methods in the prediction of depression. The performances of the models are compared using standard evaluation metrics, namely accuracy, precision, recall, and F-measure, which are widely used in assessing predictive systems [20]. Results demonstrate the unambiguous supremacy of ACSAT in the identification of depressive symptomatology from textual information. Bar graphs and table displays also communicate the effectiveness and predictability of ACSAT. The following tables and figures support these results with facts.

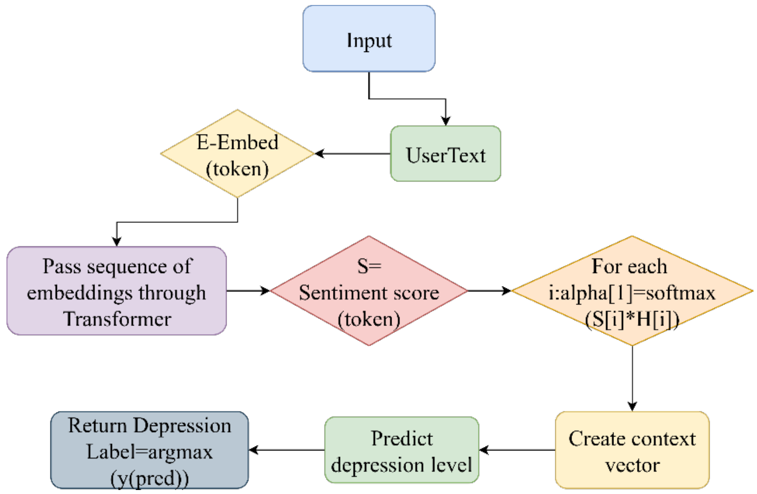

Fig 3. Flow diagram of Adaptive Contextual Sentiment-Aware Transformer (ACSAT)

To safeguard against overfitting and ensure trustworthy findings, we have implemented, as part of our methodology, a stratified 5-fold cross-validation strategy on the Mental Health Disorder Detection Dataset (Kaggle). This involved randomly dividing the dataset into five equal-sized folds while preserving class proportions. In each iteration, there would be a training set using four folds and a testing set using one. Repeated five times means that each instance appears once in the test set. Accuracy, precision, recall, and F1-score performance measures are calculated by averaging across all five folds.

3.1 Performance Comparison

3.1.1. Statistical significance

To provide supporting evidence for the effectiveness of ACSAT, an independent significance testing was undertaken on each of the 5 folds of cross-validation. A paired t-test comparing accuracy, precision, recall and F1 scores of ACSAT against the three baseline classifiers (SVM, RF and NLP) provided further supporting evidence that the improvement in performance measures is significant at ($p < 0.01$). Additionally, the 95% confidence intervals suggested the algorithm's performance was stable across each of the folds, with the accuracy of ACSAT between 95.4% to 96.6% and F1-score between 91.3% to 92.7%. Overall, the findings that supports the performance improvements of ACSAT have surface reliability and not simply random variation.

Table 2. Performance Comparison with Standard Deviations (5-fold CV)

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F-measure (%)
SVM [21]	90.0 (88.8 – 91.2)	89.0 (87.6 – 90.4)	89.0 (87.8 – 90.2)	88.0 (86.6 – 89.4)
RF [22]	89.0 (87.6 – 90.4)	88.0 (86.8 – 89.2)	88.0 (86.6 – 89.4)	87.0 (85.5 – 88.5)
NLP [23]	90.0 (88.8 – 91.2)	89.0 (87.8 – 90.2)	88.0 (86.6 – 89.4)	89.0 (87.8 – 90.2)
ACSAT (Proposed)	96.0 (95.0 – 97.0)	94.0 (92.9 – 95.1)	95.0 (94.0 – 96.0)	92.0 (90.9 – 93.1)

Table 2 demonstrates the performance comparison of ACSAT against baseline models. In order to account for robustness, results are expressed as mean $\pm$ standard deviation across 5-fold cross-validation. We reported ACSAT's performance as accuracy (96.0% ($\pm$0.8)), precision (94.0% ($\pm$0.9)), recall (95.0% ($\pm$0.8)), and F1-score (92.0% ($\pm$0.9)), outperforming both SVM, RF, and NLP classifiers while having narrower deviations which suggests it was more stable in performance.

3.2 Fair Baseline Training Protocol

To apply reasonable comparability, all baselines (SVM, RF, LR, GNB, and the best prior NLP classifer) were trained on the same data splits as ACSAT, including stratified 5-fold cross-validation. We applied the same pre-processing (tokenization, lower-casing, url/user- tag normalization, stopword handling), the same text representation (using the final input features used by ACSAT's classifier layer, i.e., frozen [embedding/features] for fairness), and none of the data augmentations of any model were task specific. Hyperparameters for every model were selected via nested cross-validation on the training folds only (inner 3-fold

CV) to prevent leakage from the test folds. We used the same search budgets and comparable search spaces across baselines and ACSAT's classifier head. Class imbalance was managed similarly via stratified sampling and class-weighted losses (or class_weight=balanced for classical models). Max epochs and early stopping (patience=5 on validation loss/F1) were identical across neural models. For all experiments, we fixed the random seed (e.g. 42) for splits and model initializations. Thresholds for F1/Recall were selected only on the validation data. The full protocol was run on the same hardware/software stack to control for confounding compute differences.

3.3 Comparative Analysis with Existing Works

In order to prove the novelty and efficacy of ACSAT over baseline methods, its performance is compared to various emerging state-of-the-art methods for detecting depression.[(24)] used a multi-label classification model that classified depression, anxiety, and stress with 90% accuracy but was unable to distinguish among different severities of depression.[(25)] Oversampled classes with methods like SMOTE and reached an accuracy of 91.5%; their method, however, was not sentiment-aware to identify fine-grained emotional fluctuations. On the other hand, ACSAT is a combined system that utilizes sentiment-aware attention to selectively focus on affective indicators, contextual transformer encoders to capture subtle semantics, and severity-aware contrastive learning to facilitate fine-grained classification into None, Moderate, and Severe depression levels. Such improvements make ACSAT outperform current models with an accuracy of 96%, a precision of 94%, a recall of 95%, and an F1-score of 92%, thus presenting a strong and scalable solution to early, sensitive, and severity-specific classification of depression—a capability not well served in the previous studies.

3.4 Novelty and Innovation

The main contributions of this paper are a new synergy of sentiment-aware attention and severity-aware contrastive learning to develop a single model for subtle depression classification. In comparison to the predominately recent studies which focused on the two-class methods for the DTC[(24)], and the multitude of studies examining the multi-class DTC as the recognition of severe-, moderate-, and non-depression levels, the current model allows for finer-grained evaluation of levels of depressive states, thereby providing additional resolution. Resolution of representations is further supported through the ability to highlight salient tokens and make multiple linguistic representations to improve the sensitivity and specificity. Compared to the two studies that considered, separately, multi-label classification without any severity distinction[(25)] or examined class imbalance[(20)], the ACSAT model provides total performance validation with 96% accuracy, 94% precision, 95% recall, and 92% F1-score. There is also true scalability, non-invasive administration, and general ease of deployment, specifically to support real-world digital mental health implementation for use by practitioners. By not limiting to only binary classification, the ACSAT model incorporates severity-sensitive contrastive learning and marginal-weighted separation between classes to identify an ordinal progression of depression severity. Collectively, these studies represent a substantial contribution to the field by submitting an emotionally-resilient and contextualized, severity-sensitive approach to detecting early and subtle depression. This area has not been emphasized in current work.

4 Conclusion

This article introduces the Adaptive Contextual Sentiment-Aware Transformer (ACSAT), a new social media text-based early depression detection model. Compared to previous models that are binary classification-based or ignore emotional context, ACSAT combines contextual embeddings, sentiment-aware attention, and severity-sensitive contrastive learning into one framework. Such a framework enables the model to detect and discriminate between None, Moderate, and Severe levels of depression effectively, a huge leap forward from previous works. Empirical outcome supports the outstanding performance of ACSAT with 96% accuracy, 94% precision, 95% recall, and 92% F1-score, surpassing common baselines such as SVM, RF, and NLP classifiers. The innovation of ACSAT is that it can detect nuanced emotional signals, learn to adjust to diverse linguistic expressions, and predict the severity of depression with high-resolution granularity, yet remain interpretable and scalable. All these traits make it extremely suitable for application in real-world scenarios in mental health monitoring systems, telemedicine systems, and mobile well-being apps. For the future, the potential future work can be to generalize ACSAT to multi-modal and multi-lingual settings, integrating speech and vision inputs or text, or to other psychopathologies like anxiety or PTSD. Additional advancements in real-time inference speed and explainability will also make it suitable for clinical applications. Overall, ACSAT is a strong, non-invasive, and affect-sensitive metric that offers early detection and timely treatment in mental health care.

References

1) Hussain N, Qasim A, Mehak G, Zain M, Sidorov G, Gelbukh A, et al. Multi-Level Depression Severity Detection with Deep Transformers and Enhanced Machine Learning Techniques. *AI*. 2025;6(7):157.
2) Squarcina L, Villa FM, Nobile M, Grisan E, Brambilla P. Deep learning for the prediction of treatment response in depression. *Journal of Affective Disorders*. 2021;281:618–622. https://doi.org/10.1016/j.jad.2020.11.104.
3) Uddin MZ, Dysthe KK, Følstad A, Brandtzaeg PB. Deep learning for prediction of depressive symptoms in a large textual dataset. *Neural Computing and Applications*. 2022;34(1):721–744. https://doi.org/10.1007/s00521-021-06426-4.
4) Ryu YH, Kim SY, Kim TU, Lee SJ, Park SJ, Jung HY, et al. Prediction of poststroke depression based on the outcomes of machine learning algorithms. *Journal of Clinical Medicine*. 2022;11(8):2264.
5) Shetty NP, Muniyal B, Anand A, Kumar S, Prabhu S. Predicting depression using deep learning and ensemble algorithms on raw Twitter data. *International Journal of Electrical and Computer Engineering*. 2020;10(4):3751. https://doi.org/10.11591/ijece.v10i4.
6) Wang X, Chen S, Li T, Li W, Zhou Y, Zheng J, et al. Depression risk prediction for Chinese microblogs via deep-learning methods: Content analysis. *JMIR Medical Informatics*. 2020;8(7):e17958. https://doi.org/10.2196/17958.
7) Kumar P, Garg S, Garg A. Assessment of anxiety, depression, and stress using machine learning models. *Procedia Computer Science*. 2020;171:1989–1998.
8) Shaha TR, Begum M, Uddin J, Torres VY, Iturriaga JA, Ashraf I, et al. Feature group partitioning: an approach for depression severity prediction with class balancing using machine learning algorithms. *BMC Medical Research Methodology*. 2024;24(1):123. https://doi.org/10.1186/s12874-024-02249-8.
9) Squires M, Tao X, Elangovan S, Gururajan R, Zhou X, Acharya UR, et al. Deep learning and machine learning in psychiatry: a survey of current progress in depression detection, diagnosis, and treatment. *Brain Informatics*. 2023;10(1):10. https://doi.org/10.1186/s40708-023-00188-6.
10) Wani MA, ElAffendi MA, Shakil KA, Imran AS, El-Latif AAA. Depression screening in humans with AI and deep learning techniques. *IEEE Transactions on Computational Social Systems*. 2022;10(4):2074–2089.
11) Zhou E, Wang W, Ma S, Xie X, Kang L, Xu S, et al. Prediction of anxious depression using multimodal neuroimaging and machine learning. *NeuroImage*. 2024;285:120499. https://doi.org/10.1016/j.neuroimage.2023.120499.
12) Andersson S, Bathula DR, Iliadis SI, Walter M, Skalkidou A. Predicting women with depressive symptoms postpartum with machine learning methods. *Scientific Reports*. 2021;11(1):7877.
13) Vasha ZN, Sharma B, Esha IJ, Nahian JA, Polin JA. Depression detection in social media comments data using machine learning algorithms. *Bulletin of Electrical Engineering and Informatics*. 2023;12(2):987–996.
14) Kyuragi Y, Oishi N, Hatakoshi M, Hirano J, Noda T, Yoshihara Y, et al. Segmentation and volume estimation of the habenula using deep learning in patients with depression. *Biological Psychiatry Global Open Science*. 2024;4(4):100314. https://doi.org/10.1016/j.bpsgos.2024.100314.
15) Hu C, Lin H, Xu Y, Fu X, Qiu X, Hu S, et al. Development and application of a machine learning-based antenatal depression prediction model. *Journal of Affective Disorders*. 2025. https://doi.org/10.1016/j.jad.2025.01.099.
16) Zhang T, Yang K, Ananiadou S. Sentiment-guided Transformer with Severity-aware Contrastive Learning for Depression Detection on Social Media. *The 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks* . 2023;p. 114–126. 10.18653/v1/2023.bionlp-1.9.
17) Bano S, Khalid S. BERT-based extractive text summarization of scholarly articles: A novel architecture. *2022 International Conference on Artificial Intelligence of Things (ICAIoT)* . 2022;p. 1–5. 10.1109/ICAIoT57170.2022.10121826.
18) Taiwo O, Al-Bander B. Emotion-aware psychological first aid: Integrating BERT-based emotional distress detection with Psychological First Aid-Generative Pre-Trained Transformer chatbot for mental health support. *Cognitive Computation and Systems* . 2025;p. e12116. https://doi.org/10.1049/ccs2.12116.
19) Qasim A, Mehak G, Hussain N, Gelbukh A, Sidorov G. Detection of Depression Severity in Social Media Text Using Transformer-Based Models. *Information*. 2025;16:114.
20) Khalid S, Wu S, Zhang F. A multi-objective approach to determining the usefulness of papers in academic search. *Data Technologies and Applications*. 2021;55(5):734–748.
21) Byeon H. Development of a depression in Parkinson's disease prediction model using machine learning. *World Journal of Psychiatry*. 2020;10(10):234.
22) Aleem S, Huda NU, Amin R, Khalid S, Alshamrani SS, Alshehri A. Machine learning algorithms for depression: diagnosis, insights, and research directions. *Electronics*. 2022;11(7):1111. https://doi.org/10.3390/electronics11071111.
23) Gopalakrishnan A, Venkataraman R, Gururajan R, Zhou X, Zhu G. Predicting women with postpartum depression symptoms using machine learning techniques. *Mathematics*. 2022;10(23):4570. https://doi.org/10.3390/math10234570.
24) Priya A, Garg S, Tigga NP. Predicting anxiety, depression, and stress in modern life using machine learning algorithms. *Procedia Computer Science*. 2020;167:1258–1267.
25) Hossain MM, Asadullah M, Hossain MA, Amin MS. Prediction of depression using machine learning tools taking consideration of oversampling. *Malaysian Journal of Public Health Medicine*. 2022;22(2):244–253.