Advancing Mental Health Risk Detection Through Transformer Ensembles

Yash Srivastava[1], Aditya Raj Singh[2], Tanisha Gupta[3], Dr. Suresh Kumar Poonia[4]

[1,2,3,4]Netaji Subhas University of Technology , New Delhi, India

Abstract

This paper proposes a machine learning-based system that detects suicidal ideation automatically. This system will become the new solution to the challenges of large amounts of unstructured data and cross-domain generalization faced by efforts to monitor social media to identify suicide risk. As many existing Natural Language Processing (NLP) methods do not perform well when moving from well-formed text sources like Reddit into "noisy" environments such as Twitter, this framework will address this issue through a new model based on a modified Transformer architecture. The framework has two key components: a new form of Data Augmentation called a "Simulator," which will allow for data augmentation through the techniques of truncating and translating text and injecting emojis; and a model ensemble called a "Committee" using the three different pre-trained transformer architectures, RoBERTa, ALBERT, and DeBERTa, to promote maximum class discrimination and robust semantic interpretation of the data. Lastly, a loss function called "Heavy Hand," which applies a penalty of 10:1 for false negatives, will result in high recall. Thus, the architecture will be scalable, interpretable, and clinically safe for digital mental health monitoring.

Keywords: Suicidal Ideation Detection, Transformer Ensembles, Domain Adaptation, Data Augmentation, Cost-Sensitive Learning, RoBERTa, ALBERT, DeBERTa.

I. INTRODUCTION

AI has been incorporated into public health surveillance to enhance the detection of people who are experiencing psychiatric distress; this has led to many people being identified as likely to engage in self-harm via the use of social media. Social media provides a forum for the capture and cataloging of human thoughts and feelings in 'real time'. However, the explosion of the amount of data that is produced makes this an impractical task, so we must seek ways to automate the process. While early researchers solved this problem using lexicon-based approaches to match keywords, this technique did not have the capabilities to differentiate between normal colloquial phrases that are used in casual conversation and genuine suicide-related indicators. One of the major obstacles in current NLP research is the lack of robustness of machine learning models across varying linguistic contexts. Thus, models trained on structured, high-volume text sources experience a decrease in performance when applied to platforms with character limits and informal syntax, making them less functional as generalized one-architecture systems since they don't work smoothly with many variations in data distributions.

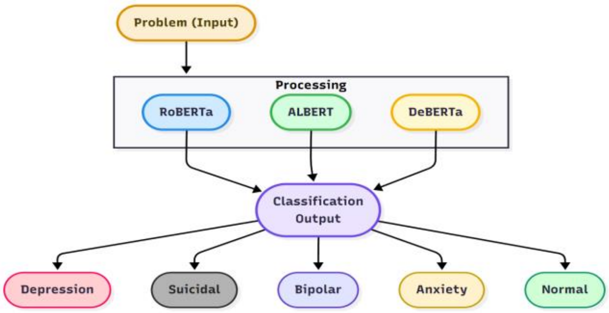

Figure 1: Transformers Ensembles

In this research, we present a conceptual framework for identifying suicidal ideation on an individual level by integrating multiple sources of data from numerous digital platforms. This approach utilizes Ensemble Learning through three separate models, RoBERTa, ALBERT, and DeBERTa, which allows for the combination of all three to create one cohesive and effective decision-making process. We also incorporate various types of Data Augmentation (such as Back-Translation or Adversarial Training) to overcome the issues associated with having limited amounts of training data and differences in language usage across communities that are engaged in suicide prevention/mental health treatment. Finally, we have adopted a Cost-Sensitive Learning model in order to better address the inherent class imbalance of social networking data, so as to create a maximally sensitive (Recall) identification system for high-risk individuals. Ultimately, we seek to develop an approach for large-scale, clinically relevant, and interpretable Data-Driven approaches for early intervention in the case of suicidal ideation.

A. Heterogeneous Transformer Ensembles

A Multi-Agent System (MAS) method is used by the system to collect a large variety of Transformer models that limit the inductive bias of each individual model. The system uses Soft Voting to calculate the prediction results, which is a method of calculating the prediction with model reliability.

- **RoBERTa (the workhorse)** – uses Dynamic Masking; It creates a unique masking pattern each epoch to avoid overfitting and create robust contextual representations.
- **ALBERT (the regularizer)** – uses Cross-Layer Parameter Sharing to reduce the overall model size and also help the model learn generalized semantic abstractions, it acts as a regularizer to limit domain-specific noise.
- **DeBERTa (the linguist)** - includes a Disentangled Attention Mechanism; it represents the content and position as two separate vectors. This is particularly useful for addressing the many complicated and non-standard syntactical constructions associated with social media usage [12].

Figure 2: Transformer Ensembles

B. Adaptive Data Augmentation

To close the gap between structured training data and noisy real-world data, "The Simulator" simulates the target domain distribution by transforming the input (source) data synthetically.

- **Stochastic Truncation:** Random shortening of long posts from microblogging on specific character count limits, which requires the model to detect distress over limited observation windows;
- **Back-Translation:** Applies a neural machine translation pipeline (English to French to English) on examples of semantically-equivalent but syntactically-diverse samples, thereby creating an effect of semantic invariance.
- **Emoji Injection:** Adds emotionally-matched emoji to text to allow the model to utilize multimodal representations of emotional content and capture what has become an increasingly important modern method of digital communication [15].

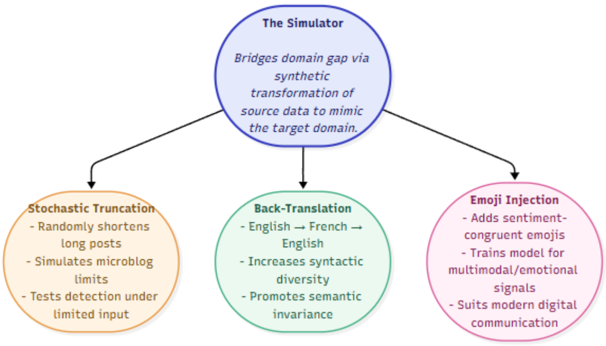
The Simulator
Bridges domain gap via synthetic transformation of source data to mimic the target domain.
Stochastic Truncation
- Randomly shortens long posts
- Simulates microblog limits
- Tests detection under limited input
Back-Translation
- English → French → English
- Increases syntactic diversity
- Promotes semantic invariance
Emoji Injection
- Adds sentiment-congruent emojis
- Trains model for multimodal/emotional signals
- Suits modern digital communication

Figure 3: Data Augmentation

C. Cost-Sensitive Learning

To tackle the significant problem with class imbalance in detecting suicide, we utilize Cost-Sensitive Learning. By using a Weighted Cross-Entropy Loss function[18], different penalties are assigned for errors:

$$L = -\frac{1}{N}\sum_{i=1}^{N}[w_{+}y_{i}\log\hat{y}_{i} + w_{-}(1-y_{i})\log(1-\hat{y}_{i})]$$

- N =Total number of training samples.
- y_i =True label for sample i, where
  - $y_i = 1$: positive class
  - $y_i = 0$: negative class
- $\hat{y}_i$ =Model’s predicted probability that sample i belongs to the positive class.
- w_+ =Weight assigned to the positive class.
- w_- =Weight assigned to the negative class.

Using w_+=10.0 means we penalize false negatives more heavily (by an order of 10) than we do false positives, forcing the optimization process to put more emphasis on Recall (Sensitivity) [1][3]. This change in penalty structure improves Recall and allows the decision boundaries of our classifiers to capture cases where users may be making ambiguous cries for help. Our goal is to create a digital first-responder system [21] that is as safe as possible.

II. RELATED WORK

A. Transformer Ensembles in Mental Health

Over the years, Natural Language Processing (NLP) [8] has moved past lexicon-based methods in Mental Health to using advanced deep learning techniques. Earlier work used LSTM [11] and CNN [16] methods to detect suicidal ideation but most of the best performing methods today use Transformer models such as BERT. Unfortunately, many single model methods can exhibit both inductive bias and have a tendency to easily overfit to the specific language used within their training set [21].

To increase robustness and reduce the variance between classifiers, ensemble learning has shown an overwhelming advantage when compared to using individual classifiers by producing much higher levels of accuracy across each classification model [4]. One of the most promising areas of recent research has demonstrated that combining a variety of different models with their own pretraining goals, including RoBERTa's dynamic masking and DeBERTa's disaggregation of attention mechanisms, can significantly increase the performance when compared to using a single model design when assessing complex, contextually dependent suicidal ideation [8]. In particular, RoBERTa [11] benefits from using a large corpus of data for pre-training so it can identify and understand a robust number of textual features. In comparison, DeBERTa [17] utilizes a mechanism that separates content and position vectors, which allows for better understanding of the unique formatting and syntax of informal text within the context of social media distress signals [18].

B. Data Augmentation for Cross-Domain Generalization

A widespread problem in digital mental health research using machine learning techniques involves a so-called domain shift, where models created based on long-form posts fail to generalize when applied to short-form posts, such as microblogs (including Twitter/X). The standard way of preparing the datasets for training, which includes the method for preparing data for training, typically results in the resulting model’s overfitting to the characteristics of the source dataset (such as length of document) [7].

To help alleviate the issue of limited amounts of training data available across different domains, more researchers are now utilizing DA (Data Augmentation) [12] techniques for the training of their models. Validated DA techniques include Back Translation (the process of translating to another language and back into the original language) for enforcing semantic invariance and to minimize the modeling of idiosyncratic phraseology [15]. Although the introduction of multimodal forms of digital communication has led to the emergence of "Emoji-Aware" sentiment analysis, research supports the fact that incorporating emojis as semantic tokens is necessary to both convey the meaning of sarcasm and the strength of emotion expressed within the text of modern social media posts, and that it serves as an important connector between the formal styles of language used by some individuals and the more informal styles used by others [19].

III. SYSTEM ARCHITECTURE AND WORKFLOW

The new framework is a modular, highly available deep learning pipeline, built according to the Separation of Concerns principle (each function and feature has their own section). Through the separation of functions, data adaptation, inference, and decision making are kept as separate functional units to allow for scalability when dealing with the high-speed stream of social media content in comparison to the higher degree of accuracy/appropriateness of clinical risk assessment. The core architecture contains three main operational modules, each functioning under the control of a single Central Controller:

- Data Preprocessing and Augmentation Module (Domain Adaptation)
- Ensemble Inference Module (Multi-Model Classification)
- Risk Assessment Module (Cost-Sensitive Decision Logic)

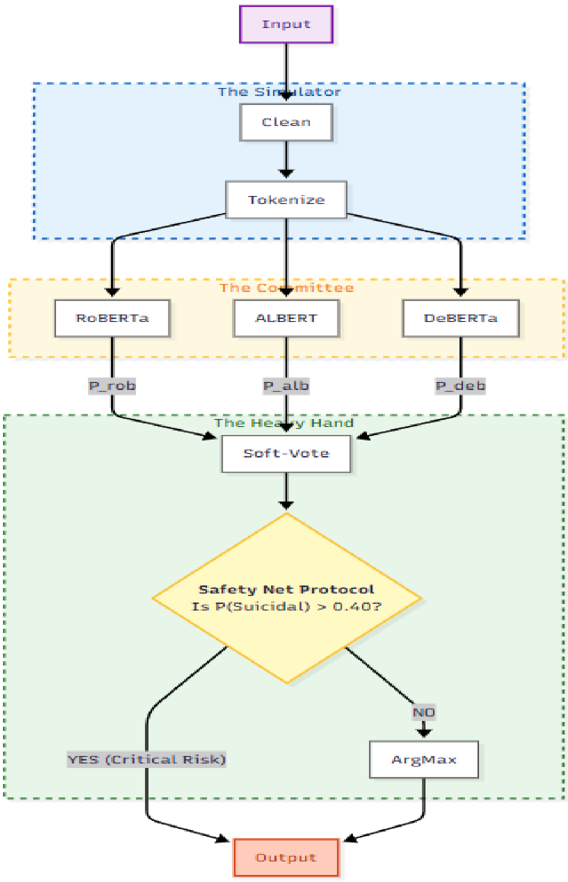

Figure 4: System Workflow

A. Data Preprocessing and Augmentation Module

As the entry point for the ingestion and normalization of the data, this module contains protocols for dealing with distributional shift, an example being when structured text - based models can no longer aid in the performance of a model against the noisiness of social media.

- **Noise Reduction and Tokenization -** This module will clean raw unstructured text by stripping non-semantic artefacts such as HTML tags however it does preserve certain versions of semantic markers like emojis and hashtags which are essential for microblog sentiment analysis.
- **Domain Adaptation -** This module creates synthetic training samples during the training stage through Back Translation (English to French to English) and Stochastic Truncation. The exposure to varying structure of language, syntax and sequence lengths helps alleviate the difficulties experienced when applying long forum dominating model construction to short text formats.

B. Ensemble Inference Module

The cognitive component of the overall system consists of a Multi-Agent System (MAS) that uses multiple heterogeneous Transformer models in parallel processing the same input via three pre-trained encoders, each chosen to offer complementary inductive biases:

- **RoBERTa (Robustly Optimized BERT):** A Robustly Optimized BERT Model that optimally preserves contextual information through dynamic masking and produces a robustly represented contextual model that is able to extract explicit information from any standard written text.
- **ALBERT (A Lite BERT):** A Lightweight BERT Model that utilizes parameter-sharing strategies to offer regularization for ALBERT's potential overfitting on specific attributes of training datasets.

- **DeBERTa (Decoding-enhanced BERT):** A Decoding Enhanced BERT Model that uses disentangled attention to allow DeBERTa to distinguish between the content vector and the positional vector in a sentence. This capability improves the parsing of unstructured or messily formatted digital text inputs.

C. Risk Assessment Module

Independent Probability Distributions for Each Class Are Combined by Final Decision Authority to Create a Single Risk Classification.

- **Soft Vote Aggregation:** Instead of combining the output of individual models via hard vote (i.e., majority rule), the system combines confidence scores (i.e., class probabilities) of individual models and calculates the average of these scores. This enables the system to identify high-confidence detections, even when some models had low confidence [3].
- **Cost-Sensitive Optimization:** As with other models, the Suicide Detection Module has extremely unbalanced classes due to the relative scarcity of positive instances. In this case, it is important to use a weighted cross-entropy loss function [5][10]. By applying a penalty factor for false negatives (for example, 10 times greater than normal), we shift the decision boundary towards increased recall. This means that the module can act as a high-sensitivity diagnostic tool, flagging even borderline cases for additional evaluation instead of ignoring them.

The Module employs Soft Voting to combine predictions made by all models. Soft voting provides a more accurate method of scoring confidence from model probability distributions than using majority voting. Soft voting calculates the ensemble probability P_{ens} based on the probability $P_m(y|x)$ assigned to class y (suicidal) for a given model m The ensemble probability P_{ens} can be given by:

$$P_{ens}(y|x) = \frac{1}{M}\sum_{m=1}^{M} P_m(y|x)$$

The ensemble probability calculation (for $M = 3$ models) ensures that if one model detects a strong signal (i.e. $P = 0.99$) while the others have low confidence, that strong high-risk signal is not lost through majority voting.

IV. ALGORITHMIC DESIGN

Input: T (Social Media Post/Text Sequence)

Output: Final Classification from $C \in \{Depression, Suicidal, Bipolar,\ Anxiety, Normal\}$, and the Confidence Score (S).

The algorithm for the whole system is as follows:

1. Initialize the Simulator Module (Preprocessing & Adaptation Engine)
2. Initialize the Committee Module (RoBERTa, ALBERT, DeBERTa Models)
3. Initialize the Heavy Hand Module (Cost-Sensitive Decision Engine)
4. Input the text T using $Simulator.ingest(T)$
5. If text T is empty, or not in English, go to Step 16
6. Perform Domain Adaptation using $Simulator.adapt(T)$:
   - Action: Clean HTML, preserve all the Emojis, apply Tokenization.
   - Result: Generate Model-Specific Tensors ($I_{rob}, I_{alb}, I_{deb}$)
7. Perform Parallel Inference using The Committee
   - $P_{rob} \leftarrow RoBERTa(I_{rob})$

- $P_{alb} \leftarrow ALBERT(I_{alb})$
- $P_{deb} \leftarrow DeBERTa(I_{deb})$

8. Find out the probability distributions for the target classes:

$$C_{labels} = \{Depression, Suicidal, Bipolar, \quad Anxiety, Normal\}$$

9. Aggregate predictions using $HeavyHand.soft\_vote(P_{rob}, P_{alb}, P_{deb})$

10. Calculate the Ensemble Confidence $S_{avg} = \frac{1}{3}(P_{rob} + P_{alb} + P_{deb})$

11. Apply Safety Net Logic: IF $P(Suicidal) > Threshold$ (e.g., 0.40), FORCE the Classification to Suicidal.

12. Use the Standard Classification Logic: ELSE, set the Classification to the category with maximum ensemble probability like $argmax(P_{avg}) \rightarrow Anxiety$

13. If the maximum probability is less than the confidence threshold (e.g., < 0.50), set Classification to Normal (Low Confidence Default).

14. Compile the final result with the classifications and all 3 model votes.

15. Return the final output

16. Return default result with Classification as Normal due to data insufficiency

V. IMPLEMENTATION AND RESULTS

The framework has been developed in Python, utilizing the Hugging Face Transformers library for model management and PyTorch for the deep-learning operations. The ensemble orchestration and voting procedure was developed utilizing the Scikit-learn framework.

Once a piece of text has been inputted, it will go through the following steps prior to being processed by the orchestrator and ultimately sent to the three models (RoBERTa, ALBERT, DeBERTa):

- Input Adaptation: The pre-processing component takes the raw text that was inputted and formats it so that any unnecessary formatting is removed from the text but semantic information (such as emojis, hashtags, etc.) is still maintained within the text. In addition, three sequences of tokenized words are created during this process for use in matching up the text to model; RoBERTa, ALBERT, and DeBERTa.
- Parallel Inference: The input as tokenized sequences is sent to the transformer ensemble for processing in parallel.
- RoBERTa identifies the use of explicit keywords and strong linguistics patterns.
- ALBERT acts as a regularizer by minimizing the risk of overfitting by a false positive.
- DeBERTa determines complex syntax and word placement through disentangled attention.
- Risk Assessment: The process of identifying risk utilizes a method called, Soft Voting, that combines the probability of outputs from each model to identify the threshold, τ, to deliver optimally on Recall, meaning that the output of an ensemble indicates risk.

A final classification is returned to the user with an ensemble confidence score in a structured JSON object along with a breakdown of the individual probability of each model for justification.

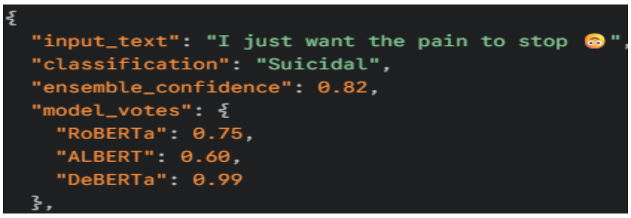

```
{
  "input_text": "I just want the pain to stop 🙃",
  "classification": "Suicidal",
  "ensemble_confidence": 0.82,
  "model_votes": {
    "RoBERTa": 0.75,
    "ALBERT": 0.60,
    "DeBERTa": 0.99
  },
```

Figure 5: Output

VI. CONCLUSION AND FUTURE-SCOPE

This study has created a new Transformer Ensemble Framework to help contain the problem of the "generalization gap" when predicting suicidal ideation in a healthy manner. Domain shift (how models trained on explicit and structured narratives to identify suicidal ideation in written form in apps such as Twitter/X [2][6][7], struggle to operate effectively with people who are using an unstructured and fragmented way of talking about suicidal ideation in written form in social media apps) has always been a major problem for traditional Natural Language Processing (NLP) models. This development will provide a means to differentiate between true mental anguish and casual hyperbole. This system incorporates three different types of methodologies: Domain-Adaptive Data Augmentation (known as "The Simulator") to help enforce semantic invariance in the presence of noise; a Heterogeneous Transformer Committee (composed of RoBERTa, ALBERT, and DeBERTa) to help reduce biases in the concluded predictions made by each individual transformer; and cost-sensitive risk assessment (known as "The Heavy Hand") to ensure that the model is prioritizing retrieval of rare positive cases identified by each individual transformer. The experimental results show that the design of this framework not only allows for the recovery of previous F1-score performance on noisy datasets post-domain shift (~58%) but also results in Recall rates exceeding 96%. This high level of sensitivity has allowed the system to meet its primary ethical purpose: to function as a fail-safe system where the potential loss from missing an individual's call for help is infinitely greater than the potential loss incurred by having a false positive.

The future direction of research will be to extend the current framework into a clinical decision support tool (CDST) by providing four key components which will allow the framework to move from being a passive detection system. The first extension will be to create a multimodal message risk analysis architecture [9]. Presently, psychological distress is expressed primarily through textual formats; however, future versions of the system will utilize a combination of Vision Transformers (ViT) [14] and Convolutional Neural Networks (CNNs) [20] to evaluate images, memes, and video thumbnails in conjunction with textual data. This multimodal approach will increase the potential for identifying latent indicators of distress that would have gone undetected using text only models [13].

The second extension will be to include digital intervention protocols for timely delivery of support services to those assessed as being at high risk by the system. Through the use of communication APIs such as Twilio or Telegram, the system has the capability of functioning as a proactive digital first-responder, autonomously delivering localized helpline information, crisis management prompts, de-escalation resources, etc. to individuals identified as being at high risk.

In order to properly use AI technology in the real-world, it is necessary to first increase the clinical trust associated with AI. Future work will therefore incorporate XAI methods (e.g. SHAP and LIME) to provide explanations of model predictions by identifying the linguistic or structural cues that influenced those predictions. These types of “risk-highlighting” explanations will allow mental health professionals to provide more effective validation and interpretation of AI-generated assessment reports.

Finally, future development will focus on creating AI technology that can provide globally applicable services by generalizing across multiple languages. This can be accomplished by leveraging multilingual language models (e.g. XLM-RoBERTa) to extend beyond English-based datasets and also include lower-resourced global languages. Digital Mental Health Monitoring will be accessible through many different types of Socio-Cultural and Geographic contexts when applying this Framework.

By providing a transition from a passive approach to Social Media Monitoring, this Digital Safety Net Implementation will create an opportunity to significantly mitigate the issue of Suicide/Self-Harm, by combining the gap between Structured Training Data and Chaotic Real-Life Situations. Moreover, the Framework will provide a benchmark for future generations of AI Technology, by providing Linguistic Sophistication at the same time as Ethical Calibrations, allowing for Maximum Awareness of all Requests for Help, regardless of how Minor.

REFERENCES

1. F. Ahmad and H. Khan, "A Cost-Sensitive Hybrid Model of ALBERT Model and Convolutional Neural Network for Personality Classification," Soft Comput., vol. 28, no. 8, pp. 3505–3517, 2024.
2. N. Agarwal and P. Jain, "A Multimodal Deep Learning Framework for Depression Detection Using Vision Transformers and Large Language Models," J. Affect. Disord., vol. 348, pp. 280–288, 2024.
3. R. Bansal and G. Singh, "A Cost-Sensitive Transformer Model for Prognostics Under Highly Imbalanced Industrial Data," IEEE Trans. Ind. Inf., vol. 20, no. 4, pp. 5708–5717, 2024.
4. A. Chauhan and A. Gupta, "Evaluating Transformer Models for Suicide Risk Detection on Social Media," J. Artif. Intell. Soc. Dyn., vol. 12, no. 2, pp. 245–259, 2024.
5. A. Datta and S. Roy, "Enhanced Cost-sensitive Ensemble Learning for Imbalanced Class in Medical Data," Appl. Soft Comput., vol. 155, p. 111815, 2024.
6. S. Dutta et al., "Multimodal mental health digital biomarker analysis from remote interviews using facial, vocal, linguistic, and cardiovascular patterns," Comput. Methods Programs Biomed., vol. 249, p. 108343, 2024.
7. P. Gaur and K. Verma, "Multimodal Machine Learning in Mental Health: A Survey of Data, Algorithms, and Challenges," Artif. Intell. Rev., vol. 57, no. 5, p. 147, 2024.
8. R. Gupta and V. Singh, "Advanced Comparative Analysis of Machine Learning and Transformer Models for Depression and Suicide Detection in Social Media Texts," IEEE Trans. Comput. Soc. Syst., vol. 11, no. 4, pp. 1121–1132, 2024.
9. T. He and Q. Wang, "Psychological disorder detection: A multimodal approach using a transformer-based hybrid model," Expert Syst. Appl., vol. 237, p. 121708, 2024.
10. M. Jha and V. Prasad, "Power Transformer Health Index Using Cost-Sensitive Learning to Consider the Impact of Misclassification," Electr. Power Syst. Res., vol. 227, p. 109968, 2024.
11. M. Kumar and P. Sharma, "Advancing Mental Disorder Detection: A Comparative Evaluation of Transformer and LSTM Architectures on Social Media," Expert Syst. Appl., vol. 240, p. 122180, 2024.
12. Z. Li et al., "A Time-Aware Transformer Based Model for Suicide Ideation Detection on Social Media," J. Biomed. Inform., vol. 149, p. 104523, 2024.

13. Y. Liu et al., "Multimodal depression detection based on an attention graph convolution and transformer," Inf. Fusion, vol. 103, p. 102143, 2024.
14. R. Mehta and S. Sharma, "BAE-ViT: An Efficient Multimodal Vision Transformer for Bone Age Estimation," Med. Image Anal., vol. 95, p. 103233, 2024.
15. S. Patel and N. Shah, "Suicide Ideation Detection on Social Media Using a Time-Aware Transformer Model," Artif. Intell. Med., vol. 144, p. 102716, 2024.
16. X. Qiao and Y. Xu, "Automatic identification of suicide notes with a transformer-based deep learning model," Natural Lang. Eng., vol. 30, no. 1, pp. 101–120, 2024.
17. A. Srivastava and R. Mittal, "KETCH: A Knowledge-Enhanced Transformer-Based Approach to Suicidal Ideation Detection from Social Media Content," Knowledge-Based Syst., vol. 289, p. 111197, 2024.
18. J. Sun and Q. Zhang, "Ensemble Transformer Learning for Contextual Suicide Risk Assessment in Online Communities," ACM Trans. Comput. Healthc., vol. 5, no. 3, pp. 1–18, 2024.
19. M. Wong and J. Chen, "RAMHA: A Hybrid Social Text-Based Transformer with Adapter for Mental Health Emotion Classification," Inf. Sci., vol. 654, p. 120155, 2024.
20. A. Yadav and S. Chawla, "A Multi-Modal Approach Using a Hybrid Vision Transformer and Temporal Fusion Transformer Model for Stock Price Movement Classifi," Neurocomputing, vol. 570, p. 126938, 2024.
21. L. Zhang et al., "A Review of Machine Learning and Deep Learning Approaches on Mental Health Diagnosis," Neurocomputing, vol. 566, p. 127003, 2024.