LEMIP-Net: A Longitudinal Exposure-Aware Deep Learning Framework for Predicting Mental-Health Impact from Pandemic-Related Social Media Content

Mental Health Prediction

1st Sujata Patil
Department of Electronics and Communication Engineering
KLE Technological University
Hubballi, India
sujata.patil@kletech.ac.in

2nd Les Sztandera
Department of Computer Information Systems
Thomas Jefferson University
Philadelphia, PA 19107, USA
les.sztandera@jefferson.edu

3rd Hemalatha K. L
Department of Computer Science and Engineering,
Sri Krishna Institute of Technology
Bengaluru, India
hemalathaklise@skit.org.in

***Abstract*—In this era, the rapid increase of health and pandemic-related information on social-media platforms has led to diverse issues that includes public anxiety, stress and behavioral volatility. Even though, prior researchers have explored various techniques still there are challenges due to the dependency on post-level analysis using static embeddings that limits exposure intensity, temporal progression and longitudinal psychological patterns. Therefore, to resolve these challenges a Longitudinal Exposure-Aware Mental-Impact Prediction Network namely LEMIP-Net is proposed, which models evolving emotional and linguistic behavior of users interacting with pandemic discourse. Initially, raw user-generated social-media posts are restructured into temporal sequences that enables observation of gradual psychological drift. Subsequently, individual post is semantically encoded by incorporating Cross-lingual Language Model with Robustly Optimized Bidirectional Encoder Representations from Transformers approach (XLM-RoBERTa). Consequently, exposure-classification module estimates probability of every post related to health or pandemic discourse, and these probabilities are aggregated to quantify user-specific exposure intensity. Further, a hybrid weak-supervision strategy refines mental-health labels through sparse self-reports and lexicon-based cues. Finally, the fused sequences are processed using a Transformer-Bidirectional Long Short Term Memory based architecture to capture global behavioral trends and short-term emotional shifts, which performs multi-task prediction of mental-health class, severity and deterioration risk. Hence, the experimental results illustrate that the proposed LEMIP-Net significantly outperforms state-of-the-art models, achieving robust and generalizable mental-health prediction by jointly modeling longitudinal behavior and exposure intensity.**

Keywords-cross-lingual language model; exposure modelling; mental-health prediction; pandemic-related social media; temporal modelling.

I. INTRODUCTION

In recent years, expanding of social media plays the massive impact on human's mental and physical health especially students. Although social media provides information of politics, education and Information Technology (IT), still impacts the human's self-esteem, mental health and sleep patterns [1]. In particular, social platforms such as twitter, Facebook and Instagram allow the users to share their posts, thoughts and ideas. In addition, studies revealed that link between negative consequences in social media does increase the stress, depression and anxiety [2]. Moreover, due to lack of offline communication, the depressed people have negative thoughts, low confidence and ambiguous issues [3]. However, using the social media negatively, impact the user's health issues such as disturbance of sleep, guilt feelings, difficulty in concentrating and suicidal thoughts [4]. Moreover, number of patients has been increasing every year with mental problems due to problems of social media and people who are already suffering with mental issues or physiological orders, will face more difficulties [5]. Also, there was a survey, where social media created development of fear and panic among the people and also females were affected mentally more than males in content of social media. [6]. To overcome these problems, early detection of stress, depression could prevent the mental health issues. In particular, the computers have the ability to express and recognize the emotions assists give better feedback to the users [7]. Further, sentiment analysis examines the people emotions, feelings, mood and attitude and one of the active types of research area in Natural Language Processing (NLP) [8]. Moreover, detecting the depression in posts has achieved important advancement in identifying the depression from social media posts. Further researchers analyzed the social media data to extract the valuable patterns and insights that related to the mental health problems. Further, by analyzing the huge information on social media, researches understand about mental health problems of users [9]. State of the art methods include Convolutional Neural Network (CNN), Transformer and Bidirectional Long Short-Term Memory (Bi-LSTM) performed the strategies to detect the depression. However, these models have huge training time and transformers models was not able to captured the important content that effect the accuracy [10]. Further, Long Short-Term Memory (LSTM) models have been utilized to examine sequential text data from social media platforms such as Twitter, where they learn contextual feature representations from documents, paragraphs and sentences. However, LSTM struggles with long term dependencies [11] and moreover, NLP techniques used for mental status of a person based on writing or speech and predicting the depression. However, most NLP techniques does not appreciate the variability [12] of depression.

The key contributions of the research are as follows:

- A Longitudinal and Exposure-Aware Mental-Health Prediction Framework (LEMIP-Net) is proposed, which jointly captures user's temporal linguistic behavior and their cumulative exposure to health and pandemic-related content. Therefore, by transforming raw social media posts into structured behavioral timelines, the proposed LEMIP-Net model allows clinically aligned assessment of mental-health risk.
- A domain-specific exposure classification and intensity modelling mechanism is employed that quantifies the likelihood and intensity of health-related content encountered by individual user, which provides an essential dimension for understanding the impact in psychological outcomes.
- The proposed LEMIP-Net improves inherently noisy user-reported mental-health labels by combining linguistic symptom markers, contextual cues and rule-based heuristics. Thus, this hybrid weak-supervision strategy systematically enhances label fidelity and significantly increases model learning robustness when compared to dependency on self-reports.

The overall research is structured as follows: Section II describes the literature review, Section III demonstrates proposed LEMIP-Net framework, Section IV illustrates experimental results, discussion and Section V includes Conclusion.

II. Literature Review

The literature review is performed through an organized selection strategy that assisted to recognize relevant research on mental-health prediction from social media during large-scale health crises. Specifically, peer-reviewed journal and research articles published were retrieved from standard journals using keywords including mental health prediction, social media analytics, pandemic-related sentiment analysis and longitudinal modeling. Then, the studies were filtered based on methodological consistency and relevance with significance on Machine Learning (ML) and Deep Learning (DL) approaches that are applied to mental-health inference from user-generated textual data.

Bashar, Nayak and Balasubramaniam [13] determined a hybrid deep learning model, which integrated Semi-Supervised Neural Topic Model (SNTM) and Informed Neural Network (INN) that evaluated COVID-19 discussions happened in Australian Twitter. Specifically, 2.9 million tweets were the data acquired, which were pre-processed and applied for SNTM, INN for topic discovery and sentiment classification, respectively through lexicon-based prior knowledge. Further, evolving public quires at outbreak time were interpreted by tweet volume, dynamic topic modelling, and semantic brand scoring. Thus, this method captured topic diversity and sentiments connected with real-world actions, but this model had challenges such as dependency on keyword-filtered data, English-only tweets, and lack of multimodal context, which affects the real-world deployment.

Inamdar, Chapekar, Gite & Pradhan [14] recommended a Machine Learning (ML) based framework, which detected mental stress in Reddit posts through NLP techniques. Further, this framework utilized reddit dataset that contains approximately 2800 labelled texts, which were pre-processed by various embedding strategies. Where the pre-processing techniques included Bidirectional Encoder Representations from Transformers (BERT) tokenization, Embeddings from Language Models (ELMo), and Bag-of-Words (BoW) representations. Specifically, these features were utilized to train classifiers such as logistic regression, SVM, XGBoost, and random forest models. Subsequently, this framework determined that with the limited data, the effective stress detection was possible. However, lack of demographic context, and exclusion of multimodal indications limits generalization among various sectors.

Abbas, Munir, Raza, Samee, Jamjoom & Ullah [15] introduced a depression detection model, which was a combination of BERT contextual embeddings and probabilistic features produced by random forest approach. Specifically, a dataset that contained 20,000 labelled tweets was considered and applied pre-processing. Subsequently, extracted contextual BERT embeddings and given to random forest that generated depression-related probability features, then these enhanced features were utilized to train many classifiers. Among which logistic regression attained greater accuracy and evaluated through statistical T-tests and k-fold cross-validation. Therefore, this model improved feature quality for mental health prediction, despite advantages this approach relied more on textual content and lacks user behavioral context that limits the real-world use cases.

Villa-Pérez, Trejo, Moin & Stroulia [16] demonstrated a ML approach, which used English and Spanish Twitter communications to detect nine mental health disorders. Further, two bilingual datasets were created from collected timelines of analyzed users by strict self-report patterns and cross verified with control users. Moreover, pre-processing of tweets were performed, through which linguistic features were extracted, which included, Part-of-speech (POS) tags, Linguistic Inquiry n-grams, q-grams, Word Count (LIWC), and word embeddings. Thus, this method attained greater accuracy through n-gram features, but this method contains limitations such as dependency on unverifiable self-reports, imbalance in dataset, minimized performance for low-frequency disorders and demographic mismatching.

Radwan, Amarneh, Alawneh, Ashqar, AlSobeh & Magableh [17] suggested an advanced approach that utilized Large Language Models (LLMs), ML algorithms and Generative Pre-trained Transformer 3 (GPT-3) embeddings. Specifically, these were to detect and classify social media posts that caused stress disorders and lower the mental health of individuals. Further, through all these considered techniques, a screen tool was generated that used online textual data, whereas posts were converted into vectors by GPT-3 embeddings, which also captured linguistic nuances and semantic meaning. However, there were challenges such as model bias, limited generalizability, dataset imbalance, low performance among populations, and require to improve the pre-processing techniques, which are significant for the further process in the approach.

Recent state-of-the-art approaches signifies that transformer-based embeddings and hybrid learning frameworks assists to effectively capture psychological signals from social media posts [13], [15], [17]. However, most existing methods depend on post-level classification, static representations and sparse self-reported labels, while neglecting cumulative exposure effects and longitudinal behavioral evolution [14], [16]. Thereby, these approaches remain limited in modeling sustained mental-health trajectories and exposure-induced distress. Hence, these limitations motivate the proposed LEMIP-Net framework, which integrates exposure-aware modeling with longitudinal sequence learning that helps for the development of mental-health prediction beyond static post-level baselines.

III. METHODOLOGY

The proposed research incorporates LEMIP-Net, which is a deep sequential neural architecture designed to model the temporal progression of user's emotional and linguistic behavior while simultaneously quantifying their exposure levels to pandemic-related content. Initially, the user-generated social-media posts are pre-processed into longitudinal timelines that allows the model to capture gradual psychological patterns instead of isolated expressions. Subsequently, individual post is semantically encoded using XLM-RoBERTa, then an exposure-classification module is utilized to evaluate every post to estimate the probability that belongs to health or pandemic discourse. Consequently, these probabilities are aggregated across temporal windows to compute a quantitative health-content intensity score, which provides a key variable reflecting how frequently and intensely user interacts with pandemic information. Further, the fused sequence of semantic embeddings, exposure intensities and refined mental-health indicators is then processed through the LEMIP-Net architecture, where temporal patterns are learned using Transformer and Bi-LSTM layers. Finally, a multi-task prediction module outputs the user's mental-health status, severity score and risk of future deterioration. Hence, this integrated design as demonstrated in Fig. 1, ensures that both content exposure and temporal behavioral evolution facilitates to provide accurate context-aware prediction of mental-health impact.

A. System Model and Data Description

The research incorporates the pandemic-period mental-health dataset [17] that comprises 32,487 social-media posts, which were collected from 4,216 unique users between March 2020 and July 2022. Specifically, each record includes the post text, timestamp, engagement metadata and sparse self-reported mental-health indicators. Additionally, the dataset includes content in English and Arabic, which reflects multilingual pandemic discourse. The dataset is split into 70% training, 15% validation and 15% testing that facilitates in maintaining user-level separation to avoid leakage. In particular, user-level separation facilitates that all posts from a specified user are allocated effectively to a single subset (training, validation, or testing), which helps in preventing overlap of user-specific linguistic or behavioral patterns across splits. This avoids information leakage and enables a reliable evaluation of the model's generalization to unseen users. Hence, this dataset provides a sufficiently large and a temporally rich resource for modeling longitudinal behavior, as the dataset comprises time-stamped post sequences spanning multiple months per user. Thereby, this operation allows the analysis of psychological changes over time instead of isolated observations. Further, each user u contributes a chronological sequence of posts $P_u = \{p_1, p_2, \ldots, p_T\}$, each associated with a timestamp t_i and enhanced metadata including hashtags, mentions, engagement and contextual cues, which are related to COVID-19, health fear, vaccination, restrictions, anxiety and uncertainty. Thus, the dataset for each user is expressed using (1):

$$D_u = (P_u, T_u, M_u) \tag{1}$$

Where P_u signifies the post sequence, T_u refers to the corresponding timestamps and M_u defines the available mental-health labels or self-reports. Additionally, the textual content comprises naturally occurring pandemic-related expressions such as cases rising, quarantine, fever symptoms and vaccine fear while the mental-health indicators include user self-assessments, sentiment scores or psychological lexicon matches. Subsequently, data pre-processing is performed, which removes noise, bot-generated content, irrelevant posts and normalizes the text for stable embedding generation. Specifically, to define temporal granularity, a data-driven stability rule is applied, where the posting density is evaluated for individual user and choose the minimum window (Δt) where $\geq 80\%$ of users have at least one post. Hence, the weekly windows assist to maximize temporal continuity while preventing sparse user sequences. Henceforth, to formally select the optimal temporal window Δt, the posting-density stability is assessed through (2):

$$S(\Delta t) = \left(\frac{1}{U}\right) \Sigma_u\, I(n_u(\Delta t) \;\geq\; 1) \tag{2}$$

Where candidate window size (1,3,7,14 days) is defined as Δt, the proportion of user with ≥ 1 post in each window is demonstrated as $S(\Delta t)$, U signifies the total users and $n_u(\Delta t)$ determines the number of posts of user u in window Δt and $I(.)$ determines the indicator function. Thus, the temporal windowing is selected as $\Delta t = 7$ days since it maximizes the $S(\Delta t)$ while preventing fragmentation in low-activity users.

In particular, the raw dataset is transformed into a structured longitudinal behavioral record using a temporal aggregation technique, where user posts are chronologically organized into definite time window size that assists to reconstruct complete behavioral timelines for proposed LEMIP-Net.

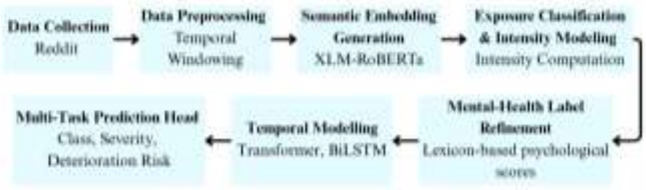

Figure 1. Flow diagram of proposed LEMIP-Net framework

Hence, this step ensures that the raw dataset, which is initially heterogeneous and sparse converted into a structured, time-sensitive representation. Hence, this operation is essential for capturing gradual psychological changes and exposure accumulation that a single-post models struggles to detect, thereby resolves a core limitation of design using only post.

B. *Initialization and Semantic Embedding Construction*

Further, each post p_t is converted into a dense semantic embedding by employing XLM-RoBERTa model, which is chosen based on the multilingual strength and ability to capture pandemic-related emotional and contextual features. Additionally, due to multilingual and code-switched pandemic discourse, the employed XLM-RoBERTa model applies SentencePiece tokenization with 250k vocabulary, Unicode NFKC normalization and maximum sequence length of 128 tokens where these parameters ensure multilingual consistency. Thus, the embedding process and semantic matrix is defined using (3) and (4), respectively:

$$e_i = f_{\text{XLM-R}}(p_i) \tag{3}$$

$$E = [e_1, e_2, \dots, e_T]^T \tag{4}$$

Where embedding vector of post p_i is defined as e_i, pretrained multilingual encoder is symbolized as $f_{\text{XLM-R}}$, user level embedding matrix is demonstrated as E and T signifies the number of temporal steps. In addition, by enhancing raw text with contextual semantics facilitates to improve data quality beyond the existing static LLM embeddings. Thus, this stage ensures that the model receives psychologically significant and globally representative language patterns which are related to mental-health change. Hence, the use of XLM-RoBERTa model for semantic embedding is essential as pandemic discourse is multilingual and context-sensitive. Also, the XLM-R model, which captures emotional, psychological and health-related semantics across languages which results with richer and more generalizable representations than English-only or traditional LLM embeddings. Henceforth, this semantically enhanced representation formulates the foundation upon which exposure estimation and mental-health inference depend.

C. *Exposure Classification and Health-Content Intensity Estimation*

Furthermore, to calculate the amount health-related content each user is exposed to, the proposed LEMIP-Net incorporates an exposure classifier F_{exp}, that computes each embedding e_t. Specifically, the Modelling exposure is crucial because psychological stress enhances with frequency of pandemic-related content. Thereby, without exposure modelling, the emotional variation may be misinterpreted. Thus, the classifier outputs the probability that the post concerns pandemic or health-related information as demonstrated in (5):

$$\hat{y}_i^{\varepsilon\xi\pi} = \sigma(W_{exp} e_i + b_{exp}) \tag{5}$$

Where the exposure probability is defined as $\hat{y}_i^{\text{exp}}$, W_{exp}, b_{exp} refers to classifier weight, bias respectively and σ denotes the sigmoid activation function. Hence, the cumulative exposure intensity over a window W_k is assessed through (6):

$$E_k = \frac{1}{|W_k|} \sum_{p_i \in W_k} \hat{y}_i^{\varepsilon\xi\pi} \tag{6}$$

Here, E_k refers to the normalized exposure frequency and intensity, specifically this module identifies posts discussing infection fear, symptoms, lockdown rules, rising cases, medical updates or anxiety triggers. Specifically, the training data for exposure classifier is established from manually tagged COVID-19 posts which includes 4,180 positive, 8,700 negative. Thereby, 2-layer MLP (128–64–1), learning rate = 1e−4, batch size = 32, Adam optimizer and dropout = 0.3 are utilized. In addition, the class imbalance is handled using focal loss ($\gamma = 2$) and random oversampling. Thus, the exposure classifier is trained using weighted binary cross-entropy as demonstrated in (7):

$$L_{exp} = -w_p\, y \log(p) - w_n (1 - y) \log(1 - p) \tag{7}$$

In particular, the exposure label is assigned through (8):

$$\hat{y} = \begin{cases} 1 & if\ p \geq \tau \\ 0 & otherwise \end{cases} \tag{8}$$

Where class weights for imbalance are defined as w_p, w_n respectively, y signifies true exposure label, predicted probability is defined as p and τ signifies the decision threshold where $\tau = 0.5$ during evaluation which is further optimized on validation. Subsequently, to define domain boundaries for exposure intensity, the predicted probability p is categorized into three regions based on the p range which is as follows: $p < 0.2$ signifies low-exposure, $0.2 \leq p \leq 0.8$ defines medium-exposure and $p > 0.8$ represents high-exposure. Hence, these boundaries are selected on the basis of maximizing inter-class separation on the validation set. Thus, the decision threshold $\tau = 0.5$ is considered for binary exposure assignment because it yields the highest Youden's J-statistic during classifier calibration.

Therefore, the resulting exposure time-series allows modelling the extent and persistence of health-related information a user observe. In particular, the exposure classification is employed because mental-health impact is strongly mediated by the volume, frequency and intensity of health-related content a user encounter. Hence, without explicit exposure modelling, AI systems risk conflating general emotional expression with crisis-driven psychological stress. Henceforth, by computing a quantitative exposure-intensity score for each time window, the proposed LEMIP-Net isolates the effect of health-content saturation, which allows downstream components to differentiate between natural emotional variability and exposure-induced distress.

D. *Mental-Health Label Refinement Through Hybrid Weak Supervision*

In addition, self-reported mental-health labels in social-media data are infrequent, therefore to obtain dense and usable supervision, the proposed LEMIP-Net integrates self-reported scores s_t with lexicon-derived psychological features l_t using (9):

$$y_t^{\rho\varepsilon\phi} = \alpha y_t^{\sigma\rho} + (1-\alpha) y_t^{\lambda\varepsilon\xi}, 0 \le \alpha \le 1 \tag{9}$$

Where the refined label is represented as y_t^{ref} and α stands for confidence weight based on self-report presence. For instance, if a user reports self-stress score $s = 0.6$ but the lexicon score is the refined label which is demonstrated in Equation (10):

$$y_t^{\rho\varepsilon\phi} = 0.7(0.6) + 0.3(0.3) = 0.51 \tag{10}$$

Specifically, if a user provides stress or anxiety self-ratings, the system preserves them. In particular, when such ratings are absent, psychological lexicons detect emotional features related to worry, fear, exhaustion and distress. Hence, these contradictions across self-reports and lexicon features are determined using a noise-aware correction rule as demonstrated where if both signals disagree, confidence-weighted averaging is used, missing values use only lexicon-based features and also lexicon noise is decrease through minimum-support threshold ($\ge 3\ symptom\ terms$).

Thus, the hybrid label facilitates every temporal step in the user's sequence which carries a mental-health estimate. Hence, this step assists to mitigate label sparsity which is a major limitation in the existing research by creating a stable ground truth that enhances learning and prevents temporal gaps in mental-state representation. Henceforth, hybrid label refinement through weak supervision is justified by the inherent sparsity and inconsistency of self-reported mental-health scores in real social-media datasets. Also, the manual labels are insufficient for training deep temporal models. Thus, combining self-reports with lexicon-based psychological indicators provides dense, consistent supervision signals which allows proposed LEMIP-Net the model to learn stable mental-health patterns without oversensitivity to annotation gaps.

E. *Longitudinal Sequence Construction and Temporal Feature Encoding*

Subsequently, each time step is defined by concatenating semantic embedding, exposure intensity and refined mental-health score. Specifically, for users with missing posts in window t, a padding vector is applied using (11):

$$x_t = [\vec{0}, 0, \mu_y] \tag{11}$$

Here, the mean refined label of user is defined as μ_y and thereby the padding operation prevents temporal discontinuities. Further, the transformer assists to capture long-range global behaviour, whereas Bi-LSTM helps to capture local fluctuations using (12):

$$x_t = \left[e_t \| \mathrm{I_t} \| y_t^{\rho\varepsilon\phi}\right] \tag{12}$$

Here, the fused vector is denoted by x_t, embedding is represented as e_t, ‖ refers to concatenation, exposure intensity is defined as I_t and y_t^{ref} signifies the refined label. Further, the transformer layer models long-range behavioral dependencies and hence the fused vector first fed through the transformer, whose output is then sequentially embedded into the Bi-LSTM as illustrated in (13):

$$H^{tr} = Transformer(X) \tag{13}$$

Where the transformer output is denoted as H^{tr}, X signifies sequence of all x_t. Thus, this operation facilitates global-to-local feature flow that includes global patterns first then short-term variations. Subsequently, a Bi-LSTM layer captures short-term emotional fluctuations and bidirectional mental-health evolution as demonstrated in (14):

$$H^{lstm} = \mathrm{BiLSTM}(H^{tr}) \tag{14}$$

Here, the BiLSTM output is defined as H^{lstm}. Thus, the fused vectors determine a complete psychological snapshot at individual time step. In particular, the Transformer captures global trends such as steadily increasing anxiety, while the Bi-LSTM models instant changes influenced by daily exposure. Thereby, this longitudinal modelling resolves the inability to account for temporal mental-health progression. Hence, the use of Transformer and Bi-LSTM layers is essential where the transformers learn global behavioral patterns, such as persistent anxiety themes or sustained exposure to crisis information, while Bi-LSTM assists to capture fine-grained emotional shifts between adjacent time steps. Henceforth, this integration ensures that both long-term mental-health evolution and short-term fluctuations are effectively modelled.

F. *Multi-Task Mental-Health Impact Prediction*

Finally, a multi-task prediction head is employed which processes temporal features to estimate three outcomes as illustrated in (15) – (17):

$$\hat{c} = Softmax(W_c h_T + b_c) \tag{15}$$

$$\hat{s} = W_s h_T + b_s \tag{16}$$

$$\hat{r} = \sigma(W_r h_T + b_r) \quad (17)$$

Here, predicted class is defined as $\hat{c}$, severity score is represented as $\hat{s}$, deterioration risk is symbolized as $\hat{r}$, the final Bi-LSTM state is demonstrated as h_T, the classifier weights refer to W_c, W_s, W_r and biases signifies b_c, b_s, b_r. Therefore, by jointly predicting class, severity and risk, the proposed LEMIP-Net model, which captures both immediate mental-health state and future vulnerability which allows a comprehensive assessment. Thus, the Multi-task training loss is demonstrated using (18):

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\chi\lambda\sigma} + \lambda_2 \mathcal{L}_{\sigma\varepsilon\varpi} + \lambda_3 \mathcal{L}_{\rho\iota\sigma\kappa} \quad (18)$$

Where classification loss is defined as $\mathcal{L}_{\mathrm{cls}}$, severity regression loss is symbolized as $\mathcal{L}_{\mathrm{sev}}$ and risk prediction loss is determined as $\mathcal{L}_{\mathrm{risk}}$ and $\lambda_1, \lambda_2, \lambda_3$ signifies chosen weights $(0.5, 0.3, 0.2)$ respectively. Specifically, the multi-task head is optimized using AdamW with learning rate of $1e-5$, 0.01 weight decay and at 1.0 gradient clipping. Therefore, to stabilize multi-task optimization, an equalized gradient scaling g_i' is applied which is expressed through Equation (19):

$$g_i' = \frac{g_i}{\|g_i\|_2} \quad (19)$$

Here, the gradient contribution of each task i is denoted by g_i, specifically the gradient normalization facilitates that no single task dominates the optimization. Thereby, tach task-specific gradient g_i is scaled by its L2-norm $\|g_i\|_2$, which provides a balanced contribution during joint training. In particular, the single-task classification struggles to capture the subtle gradations of mental-health decline or quantify future vulnerability. Hence, the multi-task outputs provide clinically significant insights and enable predictive interpretations aligned with psychological theory.

IV. Experimental Setup

The proposed LEMIP-Net framework is executed using Python with PyTorch and the Hugging Face Transformers library for model development and fine-tuning. Specifically, data pre-processing and evaluation are performed using standard scientific-computing packages such as NumPy, pandas and scikit-learn. Thus, the experiments are implemented on a system with at least 32–64 GB RAM and a multi-core CPU that assists to handle sequence construction and exposure modelling efficiently. Hence, all experiments are executed in a controlled environment with fixed random seeds to ensure reproducibility and the complete training pipeline from embedding generation to multi-task prediction is performed on the same hardware and software configuration. Hence, the proposed LEMIP-Net framework is evaluated in terms of accuracy, precision, recall and F1-Score as demonstrated in the (20) – (23), respectively:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (20)$$

$$Precision = \frac{TP}{TP+FP} \quad (21)$$

$$Recall = \frac{TP}{TP+FN} \quad (22)$$

$$F1-Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (23)$$

Here, TP is true positive, TN is true negative, FP is false positive, and FN is false negative, respectively.

A. Performance Analysis

To evaluate the effectiveness of proposed LEMIP-Net, performance analysis is performed against recent state-of-the-art deep-learning and transformer-based models which are used for mental-health prediction on social-media datasets. Specifically, these models include BERT, RoBERTa and XLNet which are strong baselines for emotional and psychological signal extraction. Thus, each model is fine-tuned under identical experimental conditions and evaluated across the standard metrics as illustrated in Fig. 2.

From Fig. 2, it is depicted that the proposed LEMIP-Net outperforms all the benchmark transformer models across every evaluation metric. Although, the conventional models such as BERT, RoBERTa and XLNet achieved better results due to their robust contextual encoding capabilities, still lacks explicit mechanisms that assists to model exposure intensity or temporal emotional drift which are both essential factors in mental-health prediction. Hence, these results illustrate that integrating exposure signals and temporal dynamics resulted with more reliable and clinically significant mental-health risk estimation.

B. Ablation Study

To compute the individual contribution of each architectural component in the proposed LEMIP-Net, an ablation study is performed conducted by incrementally integrating the major modules into a shared baseline. Specifically, all variants are trained under identical conditions and evaluated as presented in the Table I.

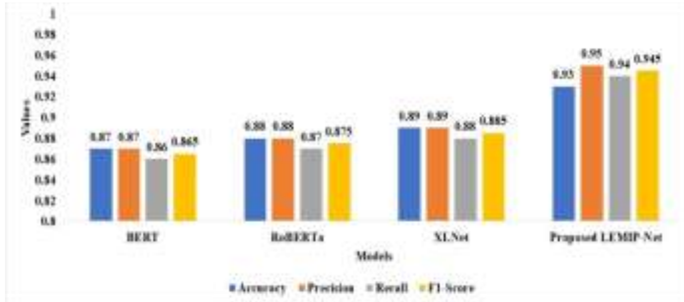

Figure 2. Performance analysis of proposed LEMIP-Net with conventional models

TABLE I. ABLATION STUDY OF PROPOSED LEMIP-NET ACROSS DIFFERENT VARIANTS

Model Variant	Accuracy	Precision	Recall	F1-score
Only XLM-RoBERTa embeddings	0.84	0.84	0.85	0.85
XLM-RoBERTa -Hybrid Weak Supervision (HWS)	0.86	0.86	0.87	0.87
XLM-RoBERTa -HWS-Transformer–Bi-LSTM	0.88	0.89	0.89	0.89
XLM-R + HWS + Temporal (No Exposure)	0.90	0.91	0.91	0.91
Proposed LEMIP-Net	0.93	0.95	0.94	0.945

From Table I, it is observed that the ablation results demonstrates that each component contributes significantly to performance improvement. Specifically, the base model provides only moderate accuracy which signifies that only text embeddings are insufficient. Therefore, by adding HWS, enhances label quality and further introducing temporal modelling improves performance by capturing behavioral changes across time. Additionally, included a no-exposure variant that helps to validate the independent effect of exposure modelling. Thus, the full proposed LEMIP-Net model incorporates exposure classification and intensity modelling that assists to obtain the highest accuracy and F1-Score which determines exposure-aware and longitudinal signals are essential for reliable mental-health impact prediction.

C. Comparative Results

To assess the robustness of the proposed LEMIP-Net framework, the performance is compared against the existing models which are widely used in mental-health prediction from social-media content. Specifically, the models include traditional machine-learning classifiers (SVM), lexicon-augmented gradient boosting (LIWC+XGB) and LLM-enhanced models (GPT-3 + SVM) as demonstrated in Table II. Hence, all comparative models are re-processed with identical tokenization, sequence length and filtering to ensure fair comparison.

Specifically, the existing models such as SVM [14] demonstrates moderate predictive capability, LIWC+XGB [16] and GPT-3 + SVM [17] obtains stronger performance, but still struggles due to the inability to incorporate temporal dynamics and exposure intensity. Hence, the proposed LEMIP-Net outperforms these models by incorporating longitudinal behavioral patterns, refined supervision and explicit modelling of pandemic-related exposure which results with higher accuracy of 0.93, 0.95 precision, 0.94 recall and 0.945 F1-score. Henceforth, these results demonstrates that modelling both the semantic evolution and exposure context significantly improves mental-health prediction compared to static or post-level baselines.

D. Discussion

The experimental results demonstrate that the proposed LEMIP-Net effectively resolves the major limitations in existing post-level mental-health prediction models. Specifically, the existing approaches primarily depended on isolated text embeddings or classical machine-learning classifiers which limited their ability to capture the cumulative psychological effects of prolonged exposure to pandemic-related content. In contrast, the proposed LEMIP-Net integrates diverse approaches including label refinement, temporal behavioral modelling and exposure-intensity quantification which results model that is both context-sensitive and longitudinally effective. Thus, the performance observed in both benchmark comparisons and ablation results determine that mental-health risk is predicted effectively when user behavior is considered as a dynamic trajectory rather than a set of independent posts. Additionally, the results illustrate that incorporating exposure intensity provides substantial predictive advantage over transformer baselines such as BERT, RoBERTa and XLNet. Hence, this defines that psychological distress in digital environments is strongly influenced by the frequency and severity of health-related information encountered which evaluates the conceptual foundation of exposure-aware modelling. Furthermore, HWS significantly enhances label quality which illustrates those self-reports alone lack reliability and benefit from linguistic signal enhancement. In particular, on analyzing the errors, it is determined that existing transformer-based approaches misclassify posts with high exposure but neutral tone, whereas the proposed LEMIP-Net model correctly incorporates exposure signals to avoid such false negatives. Henceforth, the proposed LEMIP-Net framework not only outperforms existing models but also provides a methodology aligned with psychological and behavioral science insights.

TABLE II. COMPARATIVE RESULTS OF PROPOSED LEMIP-NET WITH EXISTING MODELS

Models	Accuracy	Precision	Recall	F1-score
SVM [14]	0.74	0.70	0.74	0.76
LIWC+XGB [16]	0.823	0.997	0.802	0.881
GPT-3 + SVM [17]	0.86	0.84	0.83	0.84
Proposed LEMIP-Net	0.93	0.95	0.94	0.945

V. CONCLUSION

In this research, the proposed LEMIP-Net, which is an exposure-aware and longitudinal deep learning framework that is designed for mental-health prediction. Specifically, pandemic-related social media content is impacting every individual mental health. Further, in the proposed LEMIP-Net model raw user posts are converted into structured behavioral timelines, through which this model captures linguistic evolution, emotional drift, and cumulative effect of exposure to pandemic. Hence, the experimental results demonstrates that proposed LEMIP-Net consistently outperforms the conventional models and transformer-based approaches. This determines the requirement of combining exposure and temporal dimensions into mental health prediction model. Additionally, the ablation analysis shows that each component such as hybrid weak supervision, temporal encoding, and exposure modelling in the proposed LEMIP-Net influences the overall performance of model. In particular, the capability of the proposed LEMIP-Net in integration of refined labels, time-dependent behavioral patterns and quantified exposure signals determined as robust approach for mental-health risk assessment. Thus, the proposed LEMIP-Net outperformed existing model with which results with higher accuracy of 0.93, 0.95 precision, 0.94 recall and 0.945 F1-score. Henceforth, the proposed approach resolves key challenges of existing approaches and contributes a scalable and robust framework for predicting mental-health risks in digital ecosystems. In the future, the proposed LEMIP-Net will explore multi-modal integration including images or engagement behavior, demographic conditioning and real-time deployment for public-health surveillance.

ACKNOWLEDGMENT

The authors thank the Department of Computer Information Systems at Thomas Jefferson University for providing the academic environment and support essential for this research. L. Sztandera thanks Thomas Jefferson University, Philadelphia, PA, USA, for continued guidance and institutional resources that contributed to the successful completion of this work.

REFERENCES

[1] R. Mahevish, A. Khan, H. R. Mahmood, S. Qazi, H. M. Fakhoury, and H. Tamim, “The impact of social media on the physical and mental well-being of medical students during the COVID-19 pandemic,” J. Epidemiol. Global Health, vol. 13, pp. 902-910, November 2023.

[2] F. Benrouba and R. Boudour, “Emotional sentiment analysis of social media content for mental health safety,” Social Network Anal. Min., vol. 13, p. 17, January 2023.

[3] M. H. Al Banna, T. Ghosh, M. J. Al Nahian, M. S. Kaiser, M. Mahmud, K. A. Taher, M. S. Hossain, and K. Andersson, “A hybrid deep learning model to predict the impact of COVID-19 on mental health from social media big data,” IEEE Access, 11, pp. 77009-77022, July 2023.

[4] S. Zhou and M. Mohd, “Mental Health Safety and Depression Detection in Social Media Text Data: A Classification Approach Based on a Deep Learning Model,” IEEE Access, vol. 13, pp. 63284-63297, April 2025.

[5] İ. Aygün, B. Kaya, and M. Kaya, “Identifying patients in need of psychological treatment with language representation models,” Multimedia Tools Appl., vol. 84, pp. 397-418, 2025.

[6] M. E. Lelisho, D. Pandey, B. D. Alemu, B. K. Pandey, and S. A. Tareke, “The negative impact of social media during COVID-19 pandemic,” Trends in Psychology, vol. 31, pp. 123-142, May 2022.

[7] M. Kyrou, I. Kompatsiaris, and P.C. Petrantonakis, “Deep learning approaches for stress detection: A survey,” IEEE Trans. Affective Comput., vol. 16, pp. 499-517, September 2024.

[8] S. A. Alanazi, A. Khaliq, F. Ahmad, N. Alshammari, I. Hussain, M. A. Zia, M. Alruwaili, A. Rayan, A. Alsayat, and S. Afsar, “Public’s mental health monitoring via sentimental analysis of financial text using machine learning techniques,” Int. J. Environ. Res. Public Health, vol. 19, p. 9695, August 2022.

[9] M. Kerasiotis, L. Ilias, and D. Askounis, “Depression detection in social media posts using transformer-based models and auxiliary features,” Social Network Anal. Min., vol. 14, p. 196, September 2024.

[10] L. Ilias, S. Mouzakitis, and D. Askounis, “Calibration of transformer-based models for identifying stress and depression in social media,” IEEE Trans. Comput. Social Syst., vol. 11, pp. 1979-1990, June 2023.

[11] A. Sharaff, T. R. Chowdhury, and S. Bhandarkar, “Lstm based sentiment analysis of financial news,” SN Comput. Sci., vol. 4, p. 584, July 2023.

[12] K. Milintsevich, K. Sirts, and G. Dias, “Towards automatic text-based estimation of depression through symptom prediction,” Brain Inf., vol. 10, p. 4, February 2023.

[13] M.A. Bashar, R. Nayak, and T. Balasubramaniam, “Deep learning based topic and sentiment analysis: COVID19 information seeking on social media,” Social Network Analysis and Mining, vol. 12, no. 1, p. 90, 2022.

[14] S. Inamdar, R. Chapekar, S. Gite, and B. Pradhan, “Machine learning driven mental stress detection on reddit posts using natural language processing,” Hum.-Centric Intell. Syst., vol. 3, pp. 80-91, March 2023.

[15] M. A. Abbas, K. Munir, A. Raza, N. A. Samee, M. M. Jamjoom, and Z. Ullah, “Novel transformer based contextualized embedding and probabilistic features for depression detection from social media,” IEEE Access, vol. 12, pp. 54087-54100, April 2024.

[16] M. E. Villa-Pérez, L. A. Trejo, M. B. Moin, and E. Stroulia, “Extracting mental health indicators from english and spanish social media: A machine learning approach,” IEEE Access, vol. 11, pp. 128135-128152, November 2023.

[17] A. Radwan, M. Amarneh, H. Alawneh, H. I. Ashqar, A. AlSobeh, and A. A. A. R. Magableh, “Predictive analytics in mental health leveraging LLM embeddings and machine learning models for social media analysis,” Int. J. Web Serv. Res (IJWSR), vol. 21, pp. 1-22, January 2024.