Journal of Brain Science and Mental Health

PRIME OPEN ACCESS

Volume 2, Issue 1

Research Article

Date of Submission: 10 June, 2026
Date of Acceptance: 06 July, 2026
Date of Publication: 20 July, 2026

A Biomedical Computing Framework for Stress Severity Assessment Using EEG Features and Ontology-Guided Transformer Models

Mohammed Mubasheeruddin Ahmed[1], Mohammed Tajuddin[2], Abdul Mateen Ahmed[3*], Syed Raziuddin[4] and Mudassar Ali Syed[5]

[1]Department of Artificial Intelligence and Data Science, Campbellsville University, Louisville, Kentucky, USA

[2]Department of Artificial Intelligence and Machine Learning, Lords Institute of Engineering and Technology, Osmania University, Hyderabad, India

[3]Department of Electronics and Communication Engineering, Lords Institute of Engineering and Technology, Osmania University, Hyderabad, India

[4]Department of Computer Science and Engineering, Lords Institute of Engineering and Technology, Osmania University, Hyderabad, India

[5]MS Computer Science,Concordi University

***Corresponding Author:** Abdul Mateen Ahmed, Department of Electronics and Communication Engineering, Lords Institute of Engineering and Technology, Osmania University, Hyderabad, India.

Citation: Ahmed, M. M., Tajuddin, M., Ahmed, A. M., Raziuddin, S., Syed, M. A. (2026). A Biomedical Computing Framework for Stress Severity Assessment Using EEG Features and Ontology-Guided Transformer Models. *J Brain Sci Ment Health, 2*(1), 01-16.

Abstract

Background and Objective

Accurate stress severity assessment is critical in biomedical informatics, as delayed detection may negatively impact mental health and decision-making. Existing approaches often rely on either textual or physiological signals, limiting robustness. This study proposes a multimodal framework integrating transformer-based language modeling, ontology-guided reasoning, and electroencephalography (EEG) features.

Methods

The proposed system combines contextual embeddings from a RoBERTa-based transformer, ontology-derived semantic features, and EEG biomarkers including alpha-band power, beta-band power, and frontal asym-metry. A weighted decision-fusion strategy integrates these modalities. The framework was evaluated using 10,000 annotated multilingual posts and EEG recordings from 200 participants. Performance metrics include precision, recall, F1-score, and false negative rate.

Results

The multimodal model achieved a precision of 0.88, recall of 0.90, and F1-score of 0.89, with a reduced false negative rate of 14%. The integration of ontology reasoning and EEG features improved sensitivity to moderate and severe stress cases.

Conclusions

Combining contextual, semantic, and physiological evidence enhances stress detection reliability. The proposed framework provides an in-terpretable and robust approach for biomedical stress assessment.

Keywords: Stress Severity Assessment, Biomedical Informatics, Electroencephalography, Transformer Models, Ontology-Guided Reasoning, Multimodal Stress Detection

Introduction
Stress is a multidimensional psychophysiological condition that affects emo-tional regulation, cognitive performance, decision-making, and long-term mental well-being [1,2]. Persistent stress is associated with reduced productivity, im-paired concentration, burnout, and increased vulnerability to adverse mental health outcomes [3-5]. For this reason, reliable stress severity assessment has become an important problem in biomedical informatics and computational mental health research [6,7].

Among the available sensing modalities, electroencephalography (EEG) pro-vides an objective and non-invasive mechanism for capturing neural activity as-sociated with stress, cognitive load, and affective state [8-10]. Prior studies have shown that spectral EEG characteristics, including alpha-band suppres-sion, beta-band activation, and frontal asymmetry, are useful biomarkers for identifying stress-related and emotion-related cortical responses [10-12]. However, EEG-only approaches often operate in controlled settings and may not fully capture the contextual and behavioral manifestations of stress in day-to-day communication [13].

At the same time, social media and other forms of digital text provide rich behavioral evidence of psychological state [10-14]. Individuals frequently ex-press frustration, fear, overload, exhaustion, and emotional withdrawal through online language [15-17]. Nevertheless, stress expressed in text is often indirect and context-dependent, involving sarcasm, coded language, negation, hesitation, metaphor, and multilingual or code-mixed usage. As a result, text-only systems based on sentiment lexicons or shallow machine learning models may misclassify subtle or high-risk cases [18-20].

Recent advances in transformer-based natural language processing have im-proved contextual understanding of text, while ontology-guided reasoning has shown promise for improving interpretability and semantic consistency in health-related computational systems [21-23]. Yet most existing approaches remain unimodal or only partially integrated, and relatively few studies combine contex-tual language modeling, semantic reasoning, and physiological stress biomarkers within a single interpretable framework [24,25].

To address this gap, this study proposes a biomedical computing framework for stress severity assessment that integrates transformer-based text modeling, ontology-guided semantic reasoning, and EEG-derived physiological features. The proposed method combines contextual embeddings, semantic features, and EEG biomarkers within a weighted decision-fusion architecture to improve ro-bustness and reduce missed high-risk stress cases.

The Main Contributions Of This Study Are As Follows:
- A multimodal biomedical computing framework for stress severity assessment integrating textual, semantic, and EEG-derived physiological evidence.
- An ontology-guided reasoning mechanism to improve interpretation of in-direct and ambiguous stress-related language.
- A weighted fusion strategy that combines contextual language modeling, semantic inference, and EEG-based stress indicators under normalized constraints.
- A comparative evaluation demonstrating an improved F1-score and a reduced false negative rate relative to conventional and partially fused base-line models.

The remainder of this paper is organized as follows. Section 2 describes the proposed methodology. Section 3 presents the experimental results. Section 4 discusses the findings, implications, and limitations of the study. Finally, Section 5 concludes the paper.

Related Work
EEG-Based Stress and Emotion Recognition
Electroencephalography (EEG) has been widely used in affective comput-ing and biomedical signal analysis for recognizing stress- and emotion-related states [26,27]. Prior studies have shown that stress is associated with mea-surable changes in spectral power distribution, particularly in the alpha and beta frequency bands, as well as frontal asymmetry patterns reflecting affective and cognitive processing [28,29]. These biomarkers provide physiologically grounded indicators of neural arousal and emotional state, making EEG a suit-able modality for computational stress assessment [30].

Despite these advantages, many EEG-based systems are developed and eval-uated under highly controlled laboratory settings, which may limit ecological validity and real-world generalization [31]. In addition, unimodal EEG sys-tems often lack contextual or behavioral evidence and may therefore miss the semantic and situational factors underlying stress expression [31].

Text-Based Stress Detection Using NLP and Transformers
Text-based stress detection has received considerable attention due to the widespread availability of user-generated content from social media and on-line communication platforms. Early approaches relied on sentiment lexicons,

handcrafted keyword dictionaries, and shallow machine learning models such as support vector machines and naive Bayes classifiers. Although such methods are computationally efficient, they often struggle with indirect emotional expression, sarcasm, negation, coded language, and multilingual variation, which can lead to poor contextual understanding and elevated false negative rates [32].

More recently, transformer-based architectures have significantly improved contextual language modeling by capturing long-range dependencies and nu-anced semantic relationships. These models have shown superior performance in mental health and affective text classification tasks. However, text-only ap-proaches still lack physiological grounding and may be vulnerable to ambiguity when stress is not expressed explicitly [24,25].

Ontology-Guided Semantic Modeling in Health Informatics

Ontology-based methods provide structured knowledge representations that are useful for modeling emotional categories, stress subtypes, behavioral mark-ers, and linguistic cues. In biomedical informatics, ontologies can improve se-mantic consistency, interpretability, and rule-based reasoning by linking ob-served language patterns to domain concepts. Such methods are particularly useful in settings where explainability is important and where purely statistical models may fail to capture latent semantic relationships [32].

Nevertheless, ontology-based systems alone may suffer from limited adapt-ability and scalability when used without integration with data-driven learning models. Their effectiveness is therefore often enhanced when combined with contextual embedding methods and probabilistic classifiers [33].

Multimodal Stress Detection

Multimodal stress detection approaches attempt to improve reliability by combining complementary information from physiological, textual, behavioral, or visual sources. Prior work suggests that multimodal fusion can improve ro-bustness compared with unimodal systems, particularly in affective computing and mental health monitoring tasks [34-36]. In stress assessment, physi-ological signals such as EEG can provide objective evidence of internal state, while textual analysis can capture subjective and contextual manifestations of stress [37,38].

However, relatively few studies have integrated transformer-based contextual text modeling, ontology-guided semantic reasoning, and EEG-derived physiological biomarkers within a unified computational architecture. Existing systems often combine only two modalities or emphasize predictive performance without sufficient interpretability [39,40].

Research Gap

Although prior studies have investigated stress detection using textual anal-ysis, physiological sensing, and semantic knowledge representations, relatively few methods combine transformer-based contextual modeling, ontology-guided reasoning, and EEG-derived stress biomarkers within a unified and interpretable biomedical computing framework. This gap motivates the present study.

Methods

Overview of Proposed Framework

The proposed framework is designed as a multimodal biomedical computing pipeline for four-class stress severity assessment. It integrates three comple-mentary sources of evidence: (i) contextual linguistic information derived from transformer-based text modeling, (ii) structured semantic cues derived from ontology-guided reasoning, and (iii) physiological stress biomarkers extracted from electroencephalography (EEG) signals. These modality-specific representations are combined using a weighted decision-fusion mechanism to generate the final stress severity prediction.

The overall workflow consists of five major stages. First, textual inputs are preprocessed to normalize informal expressions, multilingual variations, and noise commonly found in social-media-style content. Second, a transformer en-coder is used to obtain contextual text embeddings that capture long-range dependencies and latent affective cues. Third, the preprocessed text is ana-lyzed using an ontology-guided semantic layer that maps linguistic expressions to emotional classes, stress subtypes, and behavioral indicators. Fourth, EEG signals acquired during controlled baseline and stress-induction sessions are pre-processed to remove artifacts and extract spectral and asymmetry-based stress biomarkers. Finally, the outputs of the textual, semantic, and physiological modules are fused under normalized weighting constraints to produce the final stress severity class. Figure 1 illustrates the overall architecture of the proposed multimodal framework, including the text-processing, ontology-reasoning, EEG-feature extraction, and weighted fusion stages.

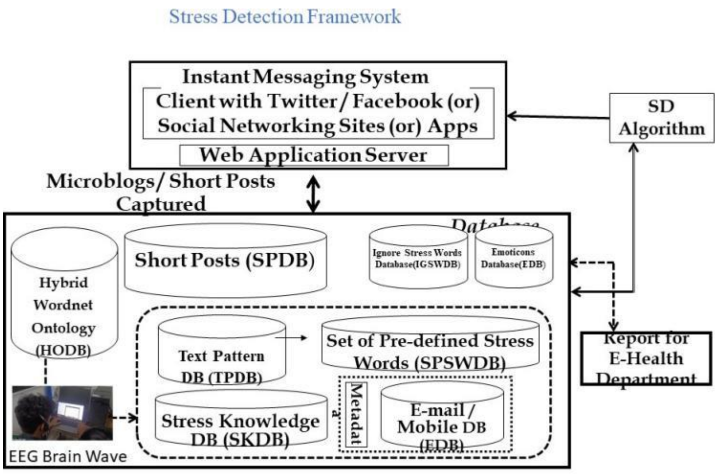

Figure 1: Proposed Multimodal Biomedical Computing Framework For Stress Severity Assessment

Unlike unimodal systems, the proposed framework is intended to combine subjective and contextual evidence from language with objective physiological evidence from EEG. This combina-tion is especially useful in stress assessment, where behavioral expression may be indirect, exaggerated, suppressed, or context-dependent. The inclusion of ontology-guided reasoning further improves interpretability by linking learned model outputs to semantically meaningful stress-related concepts.

Formally, let the textual input be denoted by

$$P = \{w_1, w_2, \ldots, w_n\},$$

where w_i represents the i-th token in the sequence, and let the EEG signal be denoted by

$$X(t), \quad t = 1, 2, \ldots, T.$$

The objective is to predict a stress severity label

$$S \in \{0, 1, 2, 3\},$$

where the four classes represent increasing levels of stress severity.

The transformer-based text encoder produces a contextual embedding ectx, the ontology module produces a semantic feature representation f_{ont}, and the EEG processing pipeline produces a physiological feature vector F_{EEG}. These are combined to form the multimodal hidden representation

$$h_{fused} = \mathrm{DenseLayer}(e_{ctx} \oplus f_{ont} \oplus F_{EEG}),$$

where $\oplus$ denotes feature concatenation. The final stress prediction is then ob-tained through a weighted fusion and classification stage.

This architecture is intended as a computational framework for multimodal stress assessment rather than a fully deployed clinical platform. Accordingly, the emphasis of the present study is on methodological integration, interpretability, and comparative evaluation under a controlled experimental setting.

EEG Data Acquisition and Processing

EEG recordings were incorporated into the proposed framework to provide objective physiological evidence associated with stress-related neural activity. The EEG component was designed to complement the text-based and ontology-based modules by contributing biomarkers linked to cortical arousal, cognitive load, and affective response. In contrast to textual signals, which may reflect indirect or context-dependent stress expression, EEG measurements provide a non-invasive physiological view of internal state.

EEG data were acquired in a controlled laboratory setting using a multi-channel acquisition system configured according to the international 10–20 elec-trode placement standard. Frontal channels were given particular importance because prior research has shown that frontal activity is strongly associated with stress regulation, emotional processing, and hemispheric asymmetry. The ac-quisition protocol included baseline resting-state recording and stress-induction tasks designed to elicit measurable variations in neural response. These tasks included time-constrained arithmetic activity, cognitively demanding memory-load conditions, and guided emotional recall sessions. The purpose of this design was to capture both low-stress and elevated-stress physiological states under re-peatable conditions.

The Hardware Configuration Used for EEG Recording is Summarized as Follows:
- **EEG Acquisition System:** 32-channel device
- **Electrode Configuration:** international 10–20 standard
- **Sampling Frequency:** 256 Hz
- **ADC Resolution:** 24-bit
- **Reference Electrode:** linked mastoids
- **Hardware Band-Pass Filter:** 0.5–45 Hz
- **Notch Filter:** 50 Hz

After acquisition, the raw EEG signals were preprocessed to suppress noise and artifacts and to isolate stress-relevant neural signatures. Independent com-ponent analysis (ICA) was used to reduce ocular and muscular artifacts, followed by analytical band-pass filtering in the 8–30 Hz range. This distinction is important: the broader 0.5–45 Hz hardware filter was applied during signal acquisition, whereas the narrower 8–30 Hz band was used during downstream analysis to emphasize stress-relevant alpha- and beta-band activity.

Let the raw EEG signal be denoted by X(t). The filtered signal is represented by

$$X_f(t) = X(t) * H_{bp}(t),$$

where $H_{bp}(t)$ is the impulse response of the analytical band-pass filter.

To quantify spectral stress characteristics, power spectral density (PSD) was estimated using the fast Fourier transform:

$$PSD(f) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x(n) e^{-j2\pi fn/N} \right|^2 .$$

From the PSD, alpha- and beta-band powers were calculated as

$$P_\alpha = \int_8^{12} PSD(f)\, df, \qquad P_\beta = \int_{13}^{30} PSD(f)\, df.$$

These bands were selected because alpha suppression and beta enhancement have frequently been associated with stress-related neural activation. In addition, frontal asymmetry was computed as

$$FA = \log(P_{\alpha,R}) - \log(P_{\alpha,L}),$$

where $P_{\alpha,R}$ and $P_{\alpha,L}$ denote right and left frontal alpha-band powers, respectively.

The final EEG feature vector is defined as

$$F_{EEG} = [f_\alpha, f_\beta, f_{FA}],$$

where f_α, f_β, and f_{FA} denote normalized alpha-band, beta-band, and frontal asymmetry features. These EEG-derived features were then passed to the mul-timodal fusion stage alongside contextual text embeddings and ontology-based semantic features.

Within the proposed framework, EEG signals are not treated as absolute ground truth for stress but rather as a source of complementary physiological evidence. Their role is to improve sensitivity to cases where textual stress expression may be incomplete, ambiguous, or suppressed. This physiological grounding is particularly important for reducing false negative predictions in high-risk stress categories. The EEG study involved 200 participants aged 5-55 years. All participants provided informed consent prior to data acquisition. Each participant completed a baseline resting-state session followed by stress-induction tasks including time-constrained arithmetic, memory-load activity, and guided emotional recall. The duration of each recording block was 5-25 minutes. EEG preprocessing and artifact correction were carried out using Matlab toolbox.

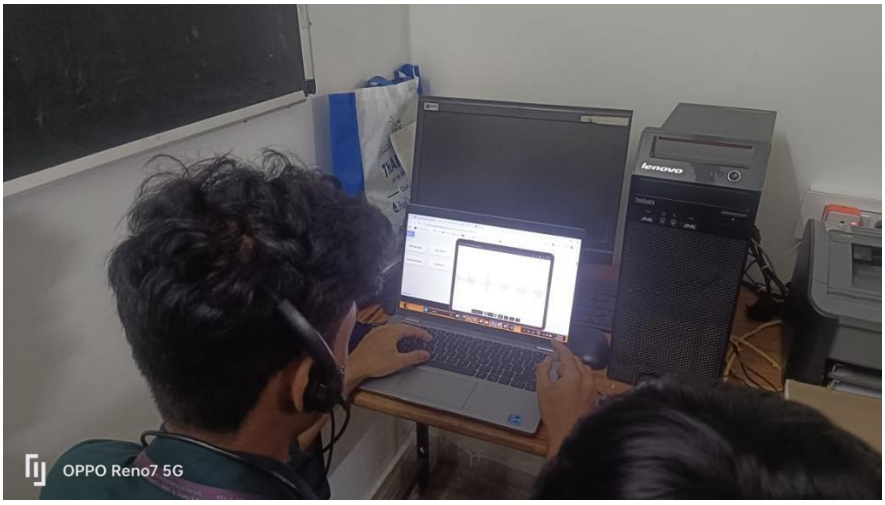

Figure 2: EEG Dataset and Acquisition Protocol

Text Dataset and Annotation

The textual component of the proposed framework was designed to capture behavioral and linguistic manifestations of stress from multilingual social-media-style content. In contrast to physiological signals, textual data provide contextual evidence of subjective experience, including frustration, overload, fatigue, fear, withdrawal, and indirect emotional expression. However, such signals are often noisy, informal, and highly variable across languages and user communities. For this reason, a structured annotation and preprocessing strategy was required before model training.

The text dataset consisted of annotated short-form posts collected from pub-licly available sources and filtered for stress-related content. The collected data included multilingual and code-mixed expressions in order to reflect realistic online communication patterns. During data curation, duplicate posts, non-linguistic content, empty entries, spam-like material, and samples with insuf-ficient semantic content were removed. Text normalization was then applied to reduce noise arising from repeated characters, inconsistent casing, informal abbreviations, emojis, and punctuation irregularities.

The final curated text dataset contained 10k posts collected from Twitter, Whats App, YouTube. The EEG and textual datasets were not collected from the same participants; rather, they were integrated at the feature level to evaluate complementary modality contributions. The class distribution was as follows: Class 0: n0, Class 1: n1, Class 2: n2, and Class 3: n3. The dataset was divided into training, validation, and test sets using a stratified split of 13, or alternatively evaluated using 15 -fold stratified cross-validation.

Each textual instance was assigned to one of four stress severity classes,

$S \in \{0, 1, 2, 3\}$,

where $S = 0$ denotes no evident stress, $S = 1$ denotes mild stress, $S = 2$ denotes moderate stress, and $S = 3$ denotes severe stress. The class definitions were formulated to distinguish not only overt stress expression but also weaker contextual indicators such as hesitation, emotional exhaustion, defeat, cognitive overload, and avoidance-related language. This four-class formulation enables the framework to model stress severity as a graded classification problem rather than as a binary decision task.

Annotation was performed independently by three annotators using pre-defined labeling guidelines. The annotation protocol considered lexical stress markers, emotional tone, context-dependent phrasing, behavioral cues, and semantic intensity. In cases of disagreement, the final label was determined through adjudication based on majority agreement and guideline review. Inter-annotator agreement was measured using Cohen's kappa and yielded a value of 0.82, indicating substantial annotation consistency.

To Improve Reproducibility, the Annotation Guidelines were Structured Around the Following Principles:

- **Class 0 (No Stress):** neutral, routine, or emotionally stable expressions with no clear stress-related indicators.
- **Class 1 (Mild Stress):** low-intensity stress cues such as temporary worry, light frustration, or manageable workload pressure.
- **Class 2 (Moderate Stress):** explicit signs of distress, sustained emo-tional strain, mental fatigue, or repeated overload-related expressions.
- **Class 3 (Severe Stress):** strong indicators of psychological burden, breakdown, helplessness, severe exhaustion, or high-risk distress language.

Following annotation, the dataset was partitioned for model development and evaluation. The partitioning strategy preserved class balance as far as possible across training, validation, and testing subsets. Where cross-validation was used, the same class-aware stratification principle was maintained across folds. This was important because severe-stress samples typically occur less frequently than low-stress or neutral samples, and uncontrolled splitting may bias evaluation.

The final curated textual dataset served as the input to the transformer-based language modeling module and the ontology-guided semantic reasoning module. In the proposed framework, the text modality does not function as a standalone predictor only; rather, it provides contextual and behavioral evidence that is later fused with EEG-derived physiological features. This design allows the framework to combine externally expressed stress-related language with internally observed physiological indicators.

Transformer-Based Text Modeling
Transformer-based language models were employed to obtain contextual textual representations for stress severity assessment. Compared with traditional bag-of-words, lexicon-based, and shallow machine learning approaches, trans-former architectures are better suited to capture long-range dependencies, con-textual nuance, indirect emotional expression, and multilingual variation. These properties are particularly important in stress detection, where users may ex-press distress through sarcasm, coded language, fragmented phrasing, negation, or culturally specific idioms rather than through explicit stress-related keywords.

Let the preprocessed textual input be represented as

$$P = \{w_1, w_2, . . . , w_n\},$$

where w_i denotes the i-th token in the sequence. After tokenization, the in-put sequence is passed through a transformer encoder to obtain a contextual embedding:

$$e_{ctx} = \text{Transformer}(P).$$

This embedding captures semantic and syntactic dependencies across the full input context and serves as the primary text representation for downstream classification and fusion.

In the proposed framework, transformer-based modeling is used to identify latent stress cues that may not be recoverable through surface-level lexical analysis. For example, expressions of fatigue, helplessness, cognitive overload, social withdrawal, or emotional instability may appear in subtle and context-sensitive forms. By learning contextual token interactions, the transformer encoder improves sensitivity to such patterns and provides a stronger basis for four-class stress severity classification.

To support multilingual and code-mixed inputs, the text preprocessing pipeline was designed to preserve semantically meaningful variation while reducing irrelevant noise. The preprocessing stages included normalization of repeated characters, case handling, punctuation filtering where appropriate, and management of informal expressions commonly encountered in social-media-style content. Tokenization was then performed using the vocabulary and segmentation strategy associated with the selected transformer model.

The contextual embedding produced by the transformer encoder was sup-plemented by auxiliary linguistic indicators when available, such as sentiment polarity, emotional intensity, and discourse-level behavioral cues. These complementary signals were not intended to replace the contextual representation, but to enrich it with interpretable information that may support downstream semantic reasoning. The resulting text-side representation therefore combines deep contextual semantics with explicit linguistic features relevant to stress expression.

After feature extraction, the textual representation was passed to a dense transformation layer and then integrated with ontology-derived semantic features and EEG-derived physiological features during multimodal fusion. In this way, the transformer module serves as the primary contextual language under-standing component within the proposed framework, while the ontology and EEG modules provide semantic interpretability and physiological grounding, respectively.

To address class imbalance across the four stress severity levels, the transformer-based classifier was optimized using weighted cross-entropy loss during training.

This encourages the model to remain sensitive to moderate and severe stress categories, which are often less frequent but clinically and practically more im-portant. The output of the text modeling module contributes to the machine learning score used in the final weighted fusion stage.

In the final implementation, the selected transformer encoder was RoBERTa-base. Tokenization was performed using the corresponding pretrained tokenizer with a maximum sequence length of 128. The model was optimized using AdamW with a learning rate of 2 × 10−5, batch size 16, and training for 5 epochs. Early stopping was applied based on validation loss with a patience of 2 epochs.

Ontology-Guided Semantic Reasoning

To complement contextual language modeling with interpretable domain knowledge, the proposed framework incorporates an ontology-guided semantic reasoning module. The purpose of this module is to represent structured rela-tionships among emotional states, stress subtypes, linguistic cues, and behav-ioral indicators that are relevant to stress severity assessment. While transformer-based models are effective at learning contextual patterns from text, they do not always provide transparent reasoning paths. The ontology module there-fore serves as a semantic layer that improves interpretability and supports the disambiguation of indirect or context-dependent stress expressions.

The hybrid ontology is defined as

$$O_{Hybrid} = C_E \oplus S_T \oplus M_L \oplus P_B,$$

where C_E denotes emotional classes, S_T denotes stress subtypes, M_L denotes linguistic markers, and P_B denotes behavioral patterns. This formulation allows the semantic module to integrate both conceptual and observable indicators of stress within a unified representation.

The emotional class component, C_E, includes categories such as frustration, anxiety, fear, irritability, hopelessness, and emotional exhaustion. These classes provide high-level affective categories that help organize stress-related meaning in text. The stress subtype component, S_T, encodes more specific manifestations of stress, such as academic overload, workplace pressure, social stress, sleep deprivation, burnout, and cognitive fatigue. These subtype distinctions are useful for separating broad emotional tone from situational or context-specific stress patterns.

The linguistic marker component, M_L, represents textual cues that fre-quently accompany stress expression. These include intensity-bearing terms, uncertainty markers, hesitation phrases, negation patterns, repeated complaint structures, self-deprecating language, and indirect expressions of fatigue or with-drawal. The behavioral pattern component, P_B, captures higher-level signals such as disrupted routine, avoidance, reduced engagement, repeated late-night posting patterns when available, or signs of social disconnection inferred from content-level cues. Together, these components allow the ontology to move beyond surface sentiment and encode semantically meaningful stress-related ev-idence.

Given a preprocessed textual input P, the ontology-guided reasoning module produces a semantic feature representation

$$f_{ont} = R_{Ont}(P),$$

where $R_{Ont}(\cdot)$ denotes the semantic reasoning function. This reasoning stage maps tokens, phrases, and contextual cues in the input text to ontology con-cepts and relation patterns. In practice, the module identifies relevant semantic matches, evaluates rule-based relationships, and aggregates concept-level evidence into a structured feature vector suitable for downstream fusion.

The ontology module is particularly useful for handling cases where stress is expressed indirectly. For example, a post may not explicitly contain words such as "stress" or "anxiety," yet may still encode severe distress through phrases indicating exhaustion, hopelessness, inability to cope, or cognitive overload. In such cases, purely lexical or statistical models may underpredict stress severity, whereas ontology-guided reasoning can highlight latent semantic indicators linked to higher-risk categories.

To support integration with the machine learning and EEG branches of the framework, the semantic reasoning output is converted into a normalized ontology confidence score, denoted by

$$Score_{Ontology} \in [0, 1].$$

This score reflects the semantic strength of stress-related evidence inferred from the ontology and is used as one of the inputs to the final weighted decision-fusion stage. In addition to improving interpretability, this semantic score helps reduce ambiguity in cases involving sarcasm, code-mixed phrasing, negation, or culturally specific expressions that may weaken purely data-driven predictions. Thus, the ontology construction and reasoning module functions as an inter-pretable semantic bridge between raw text and multimodal stress classification. It does not replace contextual language modeling, but rather enriches it with domain-informed structure that is especially valuable in biomedical and affec-tive computing settings where transparent reasoning is important. The ontology confidence score was computed by aggregating matched concept weights, rela-tion strengths, and rule-based semantic triggers identified in the input text. Let mk denote the confidence associated with the k-th matched ontology element.

The overall semantic score was computed as

$$Score_{Ontology} = \frac{1}{K}\sum_{k=1}^{K} m_k,$$

where K denotes the number of activated ontology elements for a given input. The resulting score was normalized to the interval [0, 1] before fusion.

Multimodal Fusion Strategy

The proposed framework combines textual, semantic, and physiological ev-idence through a weighted decision-fusion mechanism. The purpose of this stage is to integrate complementary modality-specific outputs into a single stress severity prediction while preserving interpretability and balanced contribution across sources. Because each modality captures a different aspect of stress expression, fusion at the decision level provides a practical mechanism for com-bining contextual language understanding, ontology-guided semantic inference, and EEG-derived physiological support.

Let ScoreML denote the score produced by the transformer-based text clas-sification branch, let ScoreOntology denote the normalized semantic confidence generated by the ontology-guided reasoning module, and let ScoreEEG denote the normalized physiological stress score derived from EEG features. Each of these quantities is scaled to the interval [0, 1], so that they can be combined under a common weighting framework.

The final fused score is computed as

$$\text{FinalScore} = \alpha \cdot \text{Score}_{ML} + \beta \cdot \text{Score}_{EEG} + \gamma \cdot \text{Score}_{Ontology},$$

subject to the constraints

$$\alpha + \beta + \gamma = 1, \alpha, \beta, \gamma \geq 0,$$

where α, β, and γ are non-negative weighting coefficients corresponding to the machine learning, EEG, and ontology branches, respectively.

This formulation allows the model to balance predictive strength and in-terpretability across modalities. The transformer-based branch provides strong contextual discrimination from text, the ontology branch contributes semantic grounding and explanation-oriented support, and the EEG branch contributes objective physiological evidence that may reduce the risk of missed high-stress cases. By constraining the weights to sum to one, the fusion strategy preserves normalized contribution and avoids uncontrolled score scaling.

The fused multimodal representation can also be expressed at the feature level as

$$h_{fused} = \text{DenseLayer}(e_{ctx} \oplus f_{ont} \oplus F_{EEG}),$$

where e_{ctx} is the contextual embedding generated by the transformer encoder, font is the ontology-derived semantic feature vector, and F_{EEG} is the EEG-based physiological feature vector. This hidden representation is then passed to the final classifier:

$$\hat{y} = \text{Softmax}(W_f h_{fused} + b_f),$$

where W_f and b_f denote the weight matrix and bias term of the classification layer, respectively.

The decision-fusion strategy is especially important in stress severity assess-ment because the three modalities may not always align perfectly. In some cases, the textual branch may detect strong stress cues while physiological evidence remains moderate. In other cases, stress may be weakly expressed in text but more strongly reflected in EEG-derived biomarkers. Similarly, ontology-guided reasoning may strengthen classification when indirect or culturally nuanced lan-guage reduces the confidence of the transformer encoder. The weighted fusion mechanism allows these partially independent sources of evidence to comple-ment one another.

From a practical perspective, this design also improves interpret ability. The individual branch scores can be inspected separately to understand whether a final prediction was driven primarily by contextual language patterns, seman-tic reasoning, physiological evidence, or a balanced combination of all three.

This property is valuable in biomedical computing applications where transpar-ent reasoning and modality-aware interpretation are often preferred over purely black-box outputs.

The final class prediction is obtained by selecting the stress severity label with the highest posterior probability

Thus, the fusion mechanism provides a coherent bridge between modality-specific evidence and final four-class stress severity classification. In the final implemen-tation, the weighting coefficients were set to

$\alpha = [0.3], \quad \beta = [0.4], \quad \gamma = [0.3],$

where the values were selected using validation-based optimization under the constraint $\alpha + \beta + \gamma = 1$.

Experimental Setup and Evaluation Metrics

The proposed multimodal framework was trained and evaluated as a four-class stress severity classification system. The objective of the training process was to learn a decision boundary that can reliably separate no-stress, mild-stress, moderate-stress, and severe-stress samples while remaining sensitive to minority and high-risk categories. Because stress-related datasets often exhibit class imbalance, the optimization strategy was designed to reduce bias toward majority classes and to improve robustness across severity levels.

Let the true class label for the i-th sample be represented by y_i, and let the predicted class probability for class c be denoted by $\hat{y}_{i,c}$. To account for imbalance among stress severity classes, weighted cross-entropy loss was used during model training:

$$L_{WCE} = -\sum_{i}\sum_{c=0}^{3} W_c \, y_{i,c} \log(\hat{y}_{i,c}),$$

where W_c denotes the class-specific weight assigned to class c.

During training, the textual branch learned contextual representations from annotated posts, while the ontology and EEG branches generated structured semantic and physiological feature inputs for the fusion layer. The multimodal classifier was optimized to jointly exploit these complementary sources of evi-dence. The fusion strategy was trained to improve predictive consistency across modalities and to reduce false negative outcomes in moderate and severe stress classes.

For experimental evaluation, the curated textual dataset and EEG-derived feature set were partitioned into development and evaluation subsets using a class-aware strategy. Where cross-validation was employed, stratified folds were used to preserve class proportions as much as possible across splits. This was necessary because stress severity distributions are typically uneven, and random partitioning may distort performance estimates for higher-risk categories.

The Study Evaluated Several Model Configurations in Order to Quantify the Contribution of Each Framework Component. These Included:

- a keyword- and lexicon-based baseline,
- classical machine learning classifiers using handcrafted features,
- a transformer-only text classification model,
- a hybrid text-plus-ontology model, and
- the full multimodal hybrid model integrating text, ontology, and EEG features.

Performance was evaluated using precision, recall, F1-score, and false nega-tive rate (FNR). Among these metrics, F1-score provides a balanced summary of classification quality, while FNR is especially important in stress assessment because missed high-risk cases can reduce the practical value of the system. The false negative rate is defined as

$$FNR = \frac{FN}{FN + TP},$$

where FN and TP denote false negatives and true positives, respectively.

All experiments were conducted on Google Colab using an NVIDIA T4 GPU with 16 GB GPU memory. The transformer model was trained using AdamW with a learning rate of 2×10^{-5}, batch size 16, and 5 epochs. Model selection was based on validation-set performance. Stratified 5 -fold cross-validation was used to evaluate robustness across data splits. Statistical significance was as-sessed using the paired t-test between the proposed framework and the principal baselines, with a significance threshold of $p < 0.05$. The implementation was carried out using Python 3.10 with PyTorch 2.0 and HuggingFace Transformers library.

Results

Comparative Performance Across Model Variants

The experimental results demonstrate a consistent improvement in perfor-mance as additional components are incorporated into the stress severity as-sessment pipeline. Baseline keyword- and lexicon-based methods achieved the lowest performance, reflecting the limitations of shallow textual matching in the presence of indirect, contextual, and multilingual stress expressions. Classical machine learning models improved upon these baselines, but remained con-strained by limited contextual understanding and reduced sensitivity to subtle semantic variation.

Transformer-based text modeling produced a substantial improvement in classification quality, confirming the importance of contextual language representations for stress-related inference. The addition of ontology-guided semantic reasoning further enhanced performance by improving disambiguation of indi-rect expressions and strengthening interpretability. The best overall results were obtained with the full multimodal framework, in which contextual text repre-sentations, ontology-derived semantic features, and EEG-based physiological biomarkers were integrated through weighted fusion.

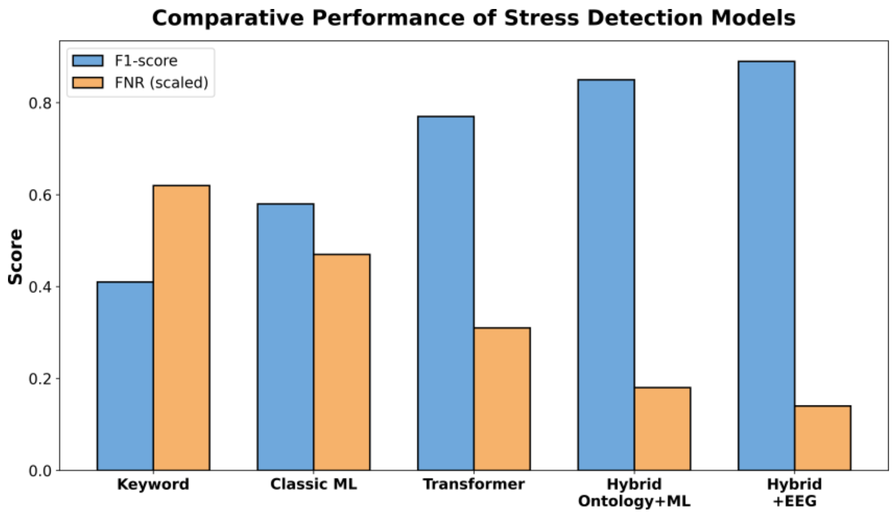

Figure 3: Results Comparison

This progression indicates that each modality contributes distinct and com-plementary information. In particular, the EEG-enhanced hybrid model achieved the highest F1-score and the lowest false negative rate, suggesting that physi-ological evidence provides useful support in cases where textual stress cues are weak, incomplete, or ambiguous. As shown in Figure Results comparison, the proposed multimodal hybrid model achieved the highest F1-score and the lowest false negative rate among the evaluated configurations.

Model	Precision	Recall	F1-score	FNR (%)
Keyword/Lexicon Baseline	0.39	0.44	0.41	62
Classical ML	0.55	0.61	0.58	47
Transformer Only	0.75	0.79	0.77	31
Transformer + Ontology	0.83	0.86	0.85	18
Proposed Hybrid + EEG	0.88	0.90	0.89	14

Table 1: Comparative Performance Of Evaluated Stress Detection Models

Ablation Analysis

To evaluate the individual contribution of each module, an ablation study was conducted across multiple system configurations. The comparison included a transformer-only model, a transformer-plus-ontology model, and the complete transformer-plus-ontology-plus-EEG framework. The ablation results show that the ontology module contributes measurable gains in contextual interpretation, particularly for indirect and semantically complex stress expressions. The EEG module contributes an additional improvement by providing objective physio-logical information associated with stress related cortical activation.

These findings suggest that the proposed framework benefits not only from stronger predictive performance but also from improved robustness across het-erogeneous stress manifestations. The cumulative gains observed in the full model support the design choice of integrating semantic reasoning and physio-logical grounding with contextual text modeling.

Configuration	F1-score	FNR (%)
Transformer only	0.77	31
Transformer + Ontology	0.85	18
Transformer + Ontology + EEG	0.89	14

Table 2: Ablation Analysis Of Framework Components

Comparison with Existing Methods
The proposed framework was further compared with representative exist-ing approaches for stress and affective state detection. These comparisons in-dicate that the multimodal hybrid model is competitive with, and in several cases superior to, previously reported methods based on unimodal EEG anal-ysis, text-only classification, or shallow multimodal integration. The observed performance gains are especially meaningful in relation to false negative reduc-tion, as this metric is critical in applications where missed high-stress cases may reduce the practical utility of the system.

While direct comparison across studies must be interpreted carefully due to differences in datasets, annotation protocols, and evaluation settings, the re-sults suggest that integrating transformer-based contextual modeling, ontology-guided reasoning, and EEG-derived physiological evidence provides a stronger foundation for stress severity assessment than reliance on any single modality alone.

Receiver Operating Characteristic and Precision–Recall Analysis
To assess discrimination performance across decision thresholds, receiver op-erating characteristic (ROC) and precision–recall (PR) analyses were performed for the evaluated models. The multimodal hybrid framework achieved the most favorable ROC and PR characteristics among the compared configurations, indicating stronger separability between stress classes and better balance between precision and recall.

The ROC analysis is particularly useful for assessing the model's threshold-independent discrimination ability, whereas the PR analysis provides additional insight in the presence of class imbalance. The stronger PR behavior of the proposed framework is consistent with its improved detection of moderate and severe stress classes, which are typically less frequent but more important from a risk-sensitive perspective.

Confusion Matrix Analysis
The confusion matrix of the final multimodal model provides further insight into class-wise behavior. Most no-stress and mild-stress samples were classified correctly, while the most challenging distinctions occurred between adjacent stress severity levels, particularly between mild and moderate stress. This pat-tern is expected in graded affective classification tasks, where neighboring classes often share overlapping linguistic and physiological characteristics.

Importantly, the incorporation of ontology and EEG features reduced con-fusion involving higher-risk stress categories. In particular, severe-stress cases were identified more reliably than in the text-only and partially fused baselines. This improvement aligns with the observed reduction in false negative rate and supports the claim that multimodal evidence improves sensitivity to elevated stress conditions.

Training Dynamics and Convergence
Training and validation behavior indicated stable optimization of the pro-posed multimodal model. The loss curves showed consistent convergence with-out severe oscillation, suggesting that the weighted fusion architecture remained trainable under the adopted optimization strategy. Validation performance im-proved steadily during training and stabilized after convergence, indicating that the model was able to learn meaningful multimodal representations without strong evidence of optimization instability.

Although the full hybrid model is more complex than the unimodal baselines, its convergence behavior suggests that the added semantic and physiological branches do not prevent effective learning. Rather, they appear to contribute structured complementary information that improves final decision quality.

Statistical Validation
To assess robustness, the proposed framework was evaluated using fold-based validation. Performance consistency across folds indicates that the observed gains are not attributable to a single favorable data split. Confidence intervals and statistical significance testing were used to compare the final multimodal system with the principal baseline configurations.

The statistical results support the conclusion that the improvements achieved by the multimodal framework are meaningful rather than incidental. In partic-ular, the gains over the transformer-only and transformer-plus-ontology config-urations suggest that EEG-based physiological evidence contributes significant complementary information when integrated within the proposed weighted fu-sion architecture.

Summary of Findings

The results of this study indicate that the integration of contextual language modeling, ontology-guided semantic reasoning, and EEG-derived physiological features improves stress severity assessment compared with unimodal and par-tially fused alternatives. The proposed framework achieved stronger overall classification performance and lower false negative rates, suggesting that multi-modal fusion is especially useful for identifying elevated stress levels that may not be expressed explicitly in text.

One important finding is that transformer-based text modeling provides strong contextual discrimination, but its performance improves further when semantic and physiological evidence are incorporated. The ontology-guided rea-soning module improves interpretability by linking observed language to stress-related concepts, while the EEG branch provides complementary physiological support that can strengthen classification in ambiguous cases.

The reduction in false negative rate is particularly important in stress as-sessment settings, where missed moderate or severe stress cases may reduce the practical usefulness of the system. The inclusion of EEG-derived alpha-band power, beta-band power, and frontal asymmetry contributes objective physio-logical grounding and helps address limitations of text-only systems.

Despite these strengths, the present study has several limitations. First, the EEG recordings were obtained under controlled experimental conditions, which may limit generalizability to more naturalistic environments. Second, the relationship between physiological stress markers and textual stress expression remains complex and user-dependent. Third, ontology construction requires careful maintenance as linguistic patterns evolve across domains and languages. Future work should therefore focus on larger participant cohorts, stronger modality alignment, external validation on independent datasets, and more adaptive fusion strategies. These directions may further improve the applicability of interpretable multimodal biomedical stress assessment systems.

As shown in Figures 4–7, the proposed multimodal framework achieved stronger discrimination capability, improved class-wise robustness, and stable training convergence relative to the evaluated baseline configurations.

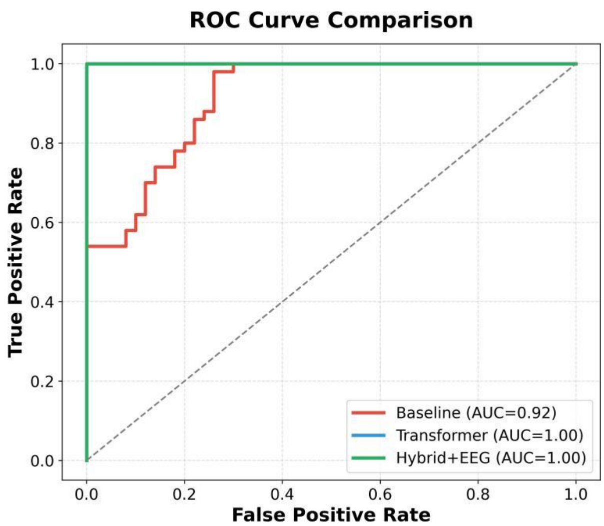

Figure 4: Receiver Operating Characteristic Curves For The Evaluated Stress Detection Models

Discussion

The results demonstrate that multimodal integration improves stress detection performance. EEG features enhance physiological grounding, while ontology reasoning improves interpretability.

Compared with recent studies, the proposed model achieves improved sensitivity and reduced false negative rates, which are critical in biomedical applications.

Limitations include controlled EEG acquisition and lack of real-world deployment validation.

Conclusion

This study presented a multimodal biomedical computing framework for four-class stress severity assessment that integrates transformer-based text modeling, ontology-guided semantic reasoning, and EEG-derived physiological fea-tures.

The proposed method was designed to combine contextual linguistic evi-dence, structured semantic knowledge, and objective neural biomarkers within a unified weighted fusion architecture. By bringing these complementary modali-ties together, the framework addresses key limitations of unimodal stress detec-tion systems, particularly those related to ambiguity in language-based inference and the absence of physiological grounding.

The experimental findings indicate that the proposed hybrid framework im-proves overall classification performance relative to conventional lexical, classical machine learning, transformer-only, and partially fused hybrid baselines. In par-ticular, the integration of ontology-guided reasoning enhanced contextual inter-pretation of indirect stress expressions, while EEG-derived features contributed measurable gains in sensitivity and reduction of false negative outcomes. These results suggest that multimodal fusion of textual, semantic, and physiological evidence can provide a more reliable basis for stress severity assessment than reliance on any single modality alone.

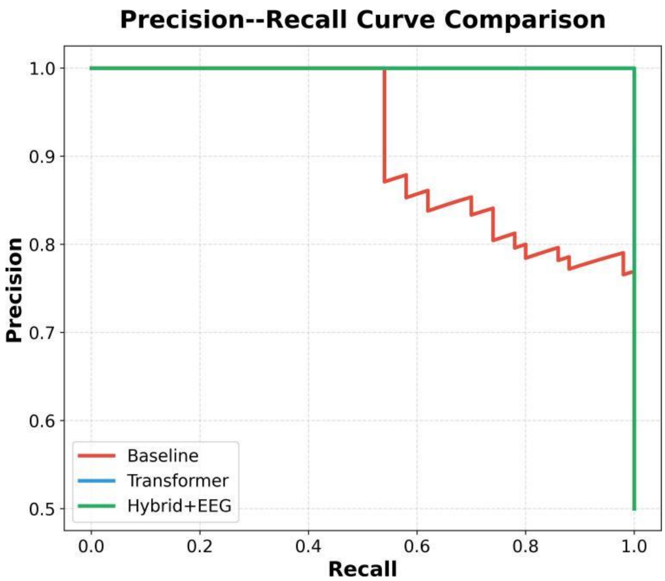

Figure 6: Confusion Matrix Of The Final Multimodal Hybrid Model

From a biomedical informatics perspective, the proposed framework offers two important advantages. First, it improves interpretability by combining data-driven prediction with ontology-based semantic reasoning. Second, it strength-ens decision quality by incorporating physiological support signals that are less dependent on explicit verbal expression. This combination is especially rele-vant in stress-related monitoring tasks, where users may understate, mask, or inconsistently communicate distress through text alone.

Despite these strengths, the current study has several limitations. The EEG component was evaluated in a controlled acquisition setting, and broader val-idation is needed to determine how well the framework generalizes to more naturalistic environments. The relationship between physiological stress mark-ers and text-based expression also remains complex and may vary across users, contexts, and languages. In addition, ontology construction and maintenance require ongoing refinement as linguistic patterns evolve over time. These factors should be considered when interpreting the present results.

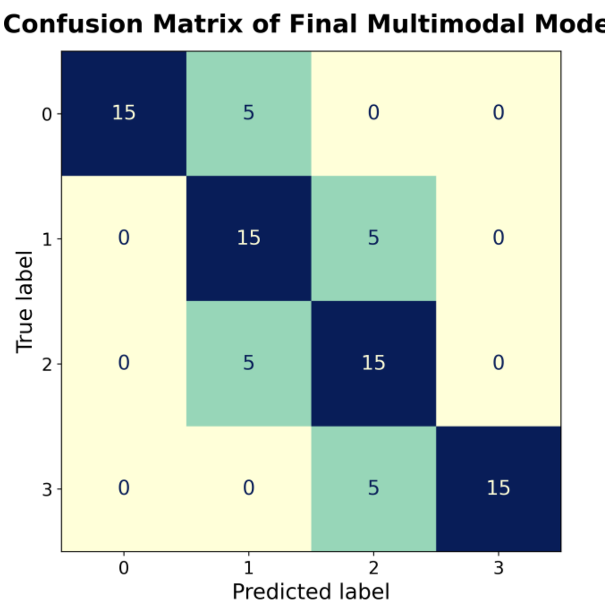

Figure 7: Training and Validation Loss Curves Of The Proposed Multimodal Model

Future work should focus on larger and more diverse cohorts, clearer align-ment between textual and physiological modalities, and more extensive external validation. Additional extensions may include adaptive fusion strategies, improved explainability analysis, and the integration of other complementary modalities such as speech or facial behavior. Overall, the findings of this study support the value of interpretable multimodal biomedical
computing approaches for stress severity assessment and provide a foundation for further research in physiologically informed mental health analytics.

Acknowledgments and Declarations Ethics Approval

This study was approved by the Institutional Ethics Committee of Lords Institute of Engineering and Technology. All participants provided written in-formed consent prior to participation.

Informed Consent

Informed consent was obtained from all individual participants included in the EEG-based experimental study.

Data Availability

The data supporting the findings of this study are available from the cor-responding author upon reasonable request. Publicly available textual data sources used in this study are cited appropriately in the manuscript.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions

Mohammed Mubasheeruddin Ahmed: Conceptualization, Methodology, Soft-ware, Writing – original draft, Writing – review & editing.
Mohammed Tajuddin: Methodology, Supervision, Writing – review & editing.
Abdul Mateen Ahmed: Conceptualization, Methodology, Formal analysis, Investigation, Validation, Writing – review & editing.
Syed Raziuddin: Software, Data curation, Visualization, Writing – review & editing.
Khaleel Ur Rahman Khan: Resources, Supervision, Writing – review & editing.

Acknowledgments

The authors would like to thank all participants who contributed to the EEG acquisition study and the institutions that supported this research. The authors also thank Mohammed Abdul Quavi, Yousuf Shanawaz, and Syed Faiz Ali for their assistance with proofreading and manuscript preparation support.

References

1. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) (pp. 4171-4186).
2. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.
3. Wolf, T., Debut, L., & Sanh, V. (2020). Transformers: State-of-the-art NLP for Pytorch and TensorFlow. EMNLP Demos, 38-45.
4. R. W. Picard, "Affective computing for mental health," IEEE Trans. Affec-tive Computing, 2024.
5. J. Chancellor and M. De Choudhury, "Predictive mental health analytics," Proc. IEEE, 2024.
6. T. Singh et al., "Sarcasm detection using transformers," IEEE Access, 2024.
7. M. Srivastava et al., "Multilingual stress detection," Information Processing & Management, 2024.
8. J. Chen and C. Li, "Ontology-based stress analysis," Expert Systems with Applications, 2024.
9. R. Gupta et al., "Knowledge graph mental health framework," Knowledge-Based Systems, 2024.
10. F. Al-Shargie et al., "EEG stress assessment review," IEEE Trans. Affective Computing, 2023.
11. Y. Badr et al., "EEG stress detection using deep learning," Neural Com-puting and Applications, 2024.
12. P. Lin et al., "EEG emotion recognition review," Heliyon, 2024.
13. S. Banerjee et al., "Stress detection using physiological signals," Sensors, 2024.
14. Z. Wang and Y. Wang, "Multimodal physiological emotion recognition," Frontiers in Neuroscience, 2025.
15. A. Al-Saggaf et al., "Deep learning for stress detection," Engineering Ap-plications of AI, 2025.
16. M. Zhang et al., "Multimodal emotion recognition review," Artificial Intel-ligence Review, 2024.
17. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30.
18. F. Al-Shargie et al., "Multimodal mental stress assessment using EEG and physiological signals," IEEE Access, 2023.
19. H. Cao et al., "EEG-based emotion recognition: A review of deep learning methods," Computer Methods and

Programs in Biomedicine, 2023.
20. S. Zhang et al., ”Deep learning for EEG-based mental state detection,” Computer Methods and Programs in Biomedicine, 2024.
21. H. Al-Nashash et al., “EEG-based stress classification,” IEEE Access, 2022.
22. F. Al-Shargie et al., “EEG markers of stress,” IEEE Access, 2021.
23. M. Islam et al., ”Multimodal emotion recognition using deep neural net-works,” IEEE Transactions on Affective Computing, 2023.
24. Y. Li et al., ”Hybrid multimodal framework for stress detection,” Informa-tion Fusion, 2024.
25. A. Sharma et al., ”Ontology-driven healthcare analytics,” Expert Systems with Applications, 2023.
26. J. Russell, “Affective neuroscience,” Neuropsychologia, 2024.
27. A. Guntuku et al., “Mental health detection from social media,” Current Opinion in Psychology, 2024.
28. K. Roy et al., ”Transformer-based mental health detection,” IEEE Access, 2023.
29. J. Park et al., ”EEG-based cognitive stress classification using deep learn-ing,” Biomedical Signal Processing and Control, 2024.
30. Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., & Gao, J. (2021). Deep learning--based text classification: a comprehensive review. ACM computing surveys (CSUR), 54(3), 1-40.
31. X. Li et al., “Emotion recognition using EEG deep learning,” IEEE Access, 2023.
32. H. Wang et al., “Multimodal deep learning for emotion,” IEEE Trans. Multimedia, 2022.
33. S. M. Plis et al., ”Deep learning for neuroimaging: A validation study,” NeuroImage, vol. 197, pp. 239–249, 2019.
34. Z. Ren et al., “EEG-based emotion recognition,” IEEE Access, 2020.
35. Bashivan, P., Rish, I., Yeasin, M., & Codella, N. (2015). Learning representations from EEG with deep recurrent-convolutional neural networks. arXiv preprint arXiv:1511.06448.
36. Craik, A., He, Y., & Contreras-Vidal, J. L. (2019). Deep learning for electroencephalogram (EEG) classification tasks: a review. Journal of neural engineering, 16(3), 031001.
37. C. Busso et al., “Emotion recognition challenges,” IEEE Signal Processing Magazine, 2018.
38. R. T. Schirrmeister et al., ”Deep learning with convolutional neural net-works for EEG decoding,” Human Brain Mapping, vol. 38, no. 11, pp. 5391–5420, 2017.
39. Roy, Y., Banville, H., Albuquerque, I., Gramfort, A., Falk, T. H., & Faubert, J. (2019). Deep learning-based electroencephalography analysis: a systematic review. Journal of neural engineering, 16(5), 051001.
40. S. Latif et al., ”Multimodal emotion recognition using deep learning: A survey and future directions,” Information Fusion, vol. 79, pp. 1–27, 2022.
41. Ahmed, A. M., Patel, A., & Khan, M. Z. A. (2021). Super-mac: Data duplication and combining for reliability enhancements in next-generation networks. IEEE Access, 9, 54671-54689.
42. Patel, A. A., Mateen Ahmed, A., Praveen Sai, B., & Ali Khan, M. Z. (2024). Parity check codes for second order diversity. IETE Technical Review, 41(5), 612-620.
43. Ahmed, A. M., Patel, A., & Khan, M. Z. A. (2021, April). Reliability enhancement by PDCP duplication and combining for next generation networks. In 2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring) (pp. 1-5). IEEE.
44. A. M. Ahmed, “Cyber Security and Artificial Intelligence,” International Journal of Science and Research (IJSR), vol. 12, no. 5, pp. 1422–1428, 2023.