



ARTICLE

TRANSHEALTH: A Transformer-BDI Hybrid Framework for Real-Time Psychological Distress Detection in Ambient Healthcare

Parul Dubey^{1,*}, Pushkar Dubey², Mohammed Zakariah^{3,4,*}, Abdulaziz S. Almazyad⁴ and Deema Mohammed Alsekait⁵

¹Symbiosis Institute of Technology, Nagpur Campus, Symbiosis International (Deemed University), Pune, 440008, India

²Department of Management, Pandit Sundarlal Sharma (Open) University Chhattisgarh, Bilaspur, 495009, India

³Department of Computer Science and Engineering, College of Applied Studies and Community Service, King Saud University, Riyadh, 11495, Saudi Arabia

⁴Department of Computer Engineering, College of Computer and Information Sciences, King Saud University, Riyadh, 11543, Saudi Arabia

⁵Department of Information Technology, College of Computer and Information Sciences, Princess Nourah bint Abdulrahman University, P.O. Box 84428, Riyadh, 11671, Saudi Arabia

*Corresponding Authors: Parul Dubey. Email: dubeyparul29@gmail.com; Mohammed Zakariah. Email: mzakariah@ksu.edu.sa

Received: 19 April 2025; Accepted: 05 August 2025; Published: 23 September 2025

ABSTRACT: Psychological distress detection plays a critical role in modern healthcare, especially in ambient environments where continuous monitoring is essential for timely intervention. Advances in sensor technology and artificial intelligence (AI) have enabled the development of systems capable of mental health monitoring using multimodal data. However, existing models often struggle with contextual adaptation and real-time decision-making in dynamic settings. This paper addresses these challenges by proposing TRANS-HEALTH, a hybrid framework that integrates transformer-based inference with Belief-Desire-Intention (BDI) reasoning for real-time psychological distress detection. The framework utilizes a multimodal dataset containing EEG, GSR, heart rate, and activity data to predict distress while adapting to individual contexts. The methodology combines deep learning for robust pattern recognition and symbolic BDI reasoning to enable adaptive decision-making. The novelty of the approach lies in its seamless integration of transformer models with BDI reasoning, providing both high accuracy and contextual relevance in real time. Performance metrics such as accuracy, precision, recall, and F1-score are employed to evaluate the system's performance. The results show that TRANS-HEALTH outperforms existing models, achieving 96.1% accuracy with 4.78 ms latency and significantly reducing false alerts, with an enhanced ability to engage users, making it suitable for deployment in wearable and remote healthcare environments.

KEYWORDS: Psychological distress detection; transformer architecture; BDI reasoning (Belief-Desire-Intention); real-time ambient healthcare; multimodal sensor data

1 Introduction

Worldwide, mental health disorders (including anxiety, depression, and other affective disorders) are among the most important contributors to the burden of disability [1]. WHO (2023) estimates that globally, around 970 million people live with a mental health disorder, including over 280 million people suffering from depression and well over 301 million people suffering from anxiety disorders [2]. Alarming, through suicide and suicide-related comorbidities, mental health problems are responsible for more than 14.3% of



global deaths per year [3,4]. While this global crisis grows, early diagnosis and contextualized tracking are still significant challenges—particularly in resource-limited and stigmatized environments [5].

Psychological disorders are characterized by the nonlinear dynamics of a similar process of time series, as the symptoms may change dynamically (e.g., from depression to mania), and different elements can impact the process, such as environment, stress, genetics, and social context [6,7]. Using traditional methods, such as clinical interviews, self-report questionnaires, and structured diagnostic tools to diagnose disorders, are limited by poor responsiveness and sensitivity. These methods often miss early and short-duration psychological symptoms that usually occur in healthy individuals who neither display visible behavioral anomalies nor seek clinical help due to stigma.

One of the domains, ambient healthcare, which is driven by the Internet of Things (IoT), Body Area Networks (BANs), and wearables, enables a wide range of potential opportunities for discreetly observing physical and behavioral indicators [8–11]. Electroencephalographic (EEG) signals, galvanic (or skin) response (GSR), heart rate variability, and activity level are measurable sensors that prove to be rich, multimodal data streams having representative features for health, emotional, and cognitive information [12]. Isolation among elderly individuals increases risks of depression, anxiety, and Alzheimer's. A proposed system offers voice/text support, image-based pattern recognition, chatbot assistance, and personalized recommendations to enhance remote mental health care with privacy and early intervention [13].

Previous research has either applied multiagent systems (MAS), rule-based classifiers, or shallow machine learning models to detect mental illnesses; however, such techniques are typically ill-suited for adaptation in more fluid, real-world situations. Furthermore, a number of methods require a large number of handcrafted features or subjective data input, which limits their scalability and generalizability. This necessitates the need for AI systems that can move from being reactive to proactive, adaptive, explainable, and used in more autonomous ways in nonlinear, interactive, uncertain environments.

While hybrid cognitive-AI systems have been studied in the past, the novelty of TRANSHEALTH lies in its operational integration of transformer-based continuous inference with a real-time symbolic reasoning engine that adapts decisions based on evolving belief states and contextual desires. Unlike static multi-agent models or batch-based cognitive engines, TRANSHEALTH supports real-time psychological state monitoring, low-latency interventions, and adaptive alert behavior in resource-constrained, wearable environments. The model is designed to balance accuracy, interpretability, and adaptiveness, which is critical for real-world mental health applications where user trust, context-awareness, and fast response are essential. This research contributes to the field of mental healthcare and intelligent systems in the following significant ways:

- Suggested a transformer-based system that combines EEG, GSR, heart rate, and activity data to help spot psychological disorders early, using a method that improves understanding and accuracy.
- Developed a Temporal Context Encoder to model long-term dependencies and differentiate between transient stress and clinical symptoms, achieving over 96% accuracy with low-latency, energy-efficient performance suitable for real-time edge deployment.
- Enabled real-time intervention through a BDI-driven alerting mechanism that adapts to user context and notifies stakeholders when distress thresholds are breached, supporting proactive and preventive mental healthcare.

In contrast to previous works that supply behaviour, the present study introduces a novel hybrid framework, TRANSHEALTH, that augments transformer-based attention mechanisms with Belief–Desire–Intention (BDI) reasoning to enable real-time detection and intervention in psychological distress scenarios.

In contrast with traditional models, TRANSHEALTH allows for the concurrent and context-aware processing of disparate physiological modalities (e.g., EEG, GSR, heart rate and physical activity) and modulates its reasoning strategy based on patterns of user-specific historical data of activity and behavior. Such integration can provide not only more accurate prediction performance but also better interpretability and adaptability—crucial properties to ensure a successful deployment in real-world ambient healthcare systems. The core novelty lies in:

- Combining deep learning and symbolic reasoning to enable cognitive adaptability;
- Introducing a temporal context encoder tailored to differentiate transient from clinical distress;
- Employing real-time BDI-based alerting, reducing false positives and improving user engagement.

In doing so, this framework moves beyond reactive classification, enabling proactive, personalized mental healthcare support within non-intrusive ambient environments. Fig. 1 shows the figure of the abstract.

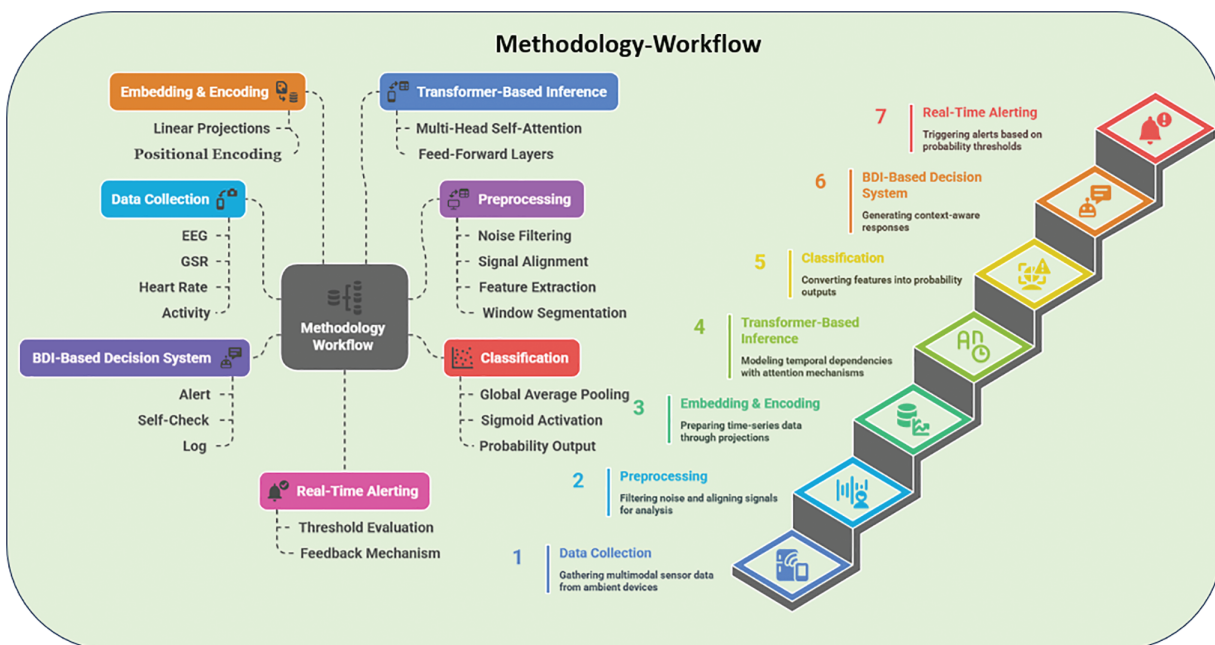


Figure 1: Visual workflow of the proposed methodology for realtime psychological disorder detection

2 Literature Review

Monitoring mental health for people in high-stress jobs (like IT professionals) has gained a parade of research interest over the past decade. Various methods have also been suggested, such as machine learning models, Internet of Things (IoT) networks of sensors, and even multiagent systems, to identify or forecast states such as anxiety and depression. Although these efforts represent important progress, each has certain limitations that drive the need for a more comprehensive solution. Researchers [14] used machine learning on other workplace-related factors to predict IT staff awareness of mental health. This approach resulted in helpful insights but was essentially dependent on self-reported data, which can lead to biases and decrease the reliability.

Wearable and IoT devices for mental health monitoring are another thread of research. Widiанти et al. [15] detect occupation-related stress with a framework that is based on mobile data analytics and wearable sensor systems, providing evidence of continuous stress tracking. In addition to physical sensors, some frameworks infer mental health from indirect indicators or external data sources. Authors proposed

detecting mental illness risk based on job stress indicators (e.g., workload and deadlines), but such proxies may not capture individual psychological states accurately, potentially leading to misclassification. Similarly, social media analytics has been employed in this domain: researchers mined social networking data to gauge anxiety levels among IT workers. While innovative, that approach raises ethical issues and questions about the validity of online behaviour as a mental health signal. These indirect and external data methods can complement traditional assessments, but they also illustrate the challenge of balancing privacy and accuracy in mental health monitoring.

Several studies have applied standard machine learning models to predict mental health outcomes in professional settings. Some researchers [16] developed a predictive model for employee burnout (a condition closely linked to chronic anxiety and depression), but its performance is constrained by the specific features and training data used. Recent advancements in AI have enabled precision psychiatry through models like BERT and GRU-based CNNs, which integrate behavioural, physiological, and contextual data for accurate mental health diagnosis. A proposed model achieved 97% accuracy, highlighting its potential for early detection and personalised interventions [17].

A recent study utilised the eB2 app, combining Hidden Markov Models and Transformer networks, to passively monitor psychiatric patients and forecast emotional states with 93% accuracy and 0.98 AUC. The model demonstrated strong potential for real-time risk detection, especially for predicting suicidal ideation and enhancing treatment planning [18]. Another study presents a context-aware framework leveraging multiagent systems (MAS) for psychological state recognition in complex healthcare environments [19]. By integrating cognitive modelling with belief-desire-intention (BDI) architecture, the proposed system simulates humanlike reasoning to adaptively assess nonlinear manifestations of mental distress. The agents operate collaboratively, drawing from real-time inputs such as behavioural patterns and historical context, to generate actionable mental health insights. The model emphasises dynamic knowledge representation, situational reasoning, and modular agent cooperation, making it a valuable baseline for context-driven psychological monitoring. This MASBDI framework serves as a foundational benchmark in our study, particularly for evaluating the reactivity, interpretability, and decision-making latency of real-time mental health interventions against our proposed transformer-enhanced system.

Existing approaches tend to be either device-centric, data-specific, or limited to retrospective analysis—lacking a unified, context-aware mechanism for proactive intervention. To overcome these gaps, recent attention has turned to cognitive modelling techniques and multiagent system (MAS) frameworks. The Belief-Desire-Intention (BDI) agent architecture [20] provides a basis for embedding humanlike reasoning in software agents, enabling them to interpret complex contexts and make autonomous decisions. MAS has been successfully applied in healthcare for tasks ranging from remote patient monitoring to disease management [21], demonstrating how distributed agents can collaborate and handle multimodal data [22]. TranSenseFuser, a deep learning architecture integrating temporal convolutions with multi-head attention, has been developed for stress detection using PPG signals. By enhancing sensor fusion and offering explainability through attention maps, the model effectively addresses motion artefacts and demonstrates robust performance and generalisability across subjects on benchmark datasets. DynaMentA, a dual-layer transformer framework combining BioGPT and DeBERTa with dynamic prompt engineering, addresses the limitations of general-purpose LLMs in mental health classification. By integrating biomedical cues with context-sensitive attention and a feedback-guided ensemble mechanism, it offers a scalable and interpretable solution for high-stakes mental health applications [23].

SLiTRANet, a novel EEG-based deep learning framework, combines spectral analysis, graph convolution, and Transformer architecture for real-time MDD detection within an IoMT setup. By leveraging S-transform and a customised linear graph convolution network, the model demonstrates robust and

generalised performance across diverse datasets, offering a significant advancement in automated depression diagnosis [24]. Deep learning-driven NLP techniques, such as BERT and GPT, are being used to develop digital twins in psychiatry and neurological rehabilitation by analysing clinical texts and patient language patterns. These models enable personalised care and predictive insights by integrating multimodal data while highlighting the need for ethical implementation in clinical environments [25]. A systematic review of 184 studies highlights the increasing use of multimodal, passively sensed data—such as audio, video, and smartphone inputs—for mental health detection, emphasising neural networks for effective feature fusion and modelling. The study offers a taxonomy of methodologies, aiding researchers in aligning data sources with targeted mental health conditions [26].

EEG Mind-Transformer introduces a novel architecture integrating dynamic temporal attention, hierarchical brain modelling, and spatial-temporal fusion for mental health monitoring using EEG signals. By effectively capturing complex spatiotemporal patterns, it demonstrates superior generalisability and interpretability, offering promising implications for clinical and research applications in mental health assessment [27]. Another study presents a novel approach to mental stress detection by leveraging in-ear PPG signals and Vision Transformer (ViT) models, demonstrating strong potential for accurate classification through time-frequency representations. Despite a small sample size, the method highlights promising applications for wearable stress monitoring in real-world mental health care [28].

This review highlights the complementary potential of large language models and smart physiological monitoring devices in stress management. It emphasises the integration of AI-driven language understanding with bio-signal sensing technologies, laying the groundwork for future multimodal, personalised mental health support systems [29]. FL-BERT+DO is a privacy-preserving framework that combines federated learning with data obfuscation and BERT-based sentiment analysis to forecast mental health sentiment. It ensures data remains decentralised while maintaining strong performance in emotion classification, demonstrating resilience to privacy attacks and offering a secure approach to mental health monitoring [30]. Another review examines AI's role in healthcare, focusing on cognitive and emotional analytics to enhance patient-centred care. It highlights advancements in sentiment analysis and clinical decision support while also addressing challenges like ethics, privacy, and bias [31]. Another study showed a self-supervised learning framework using multimodal physiological signals and transformer-based fusion achieves state-of-the-art emotion recognition, offering improved accuracy and robustness over supervised methods with limited data [32].

In contrast to prior reactive or single-point solutions, the proposed framework integrates ambient sensing with a BDI-driven MAS architecture to enable real-time, context-aware assessment and early warning of anxiety and depression symptoms, thus addressing the limitations identified in earlier works.

Despite growing interest in AI-driven mental health systems, several gaps remain underexplored:

- Most deep learning models lack interpretable output layers, making them unsuitable for high-stakes applications like mental health monitoring.
- Symbolic cognitive models like BDI have rarely been integrated into real-time, wearable-friendly frameworks.
- Few studies rigorously address privacy-preserving computation, consent-driven data governance, or cultural adaptability, especially in marginalized or resource-limited settings.

TRANSHEALTH directly addresses these gaps by:

- Integrating symbolic reasoning to ensure action-level traceability and user trust.
- Supporting on-device inference to minimize data transmission.

- Offering a modular design that can be deployed on wearable edge hardware while maintaining adaptability to context and user behavior.

3 Problem Statement

Challenges in the early detection of psychological disorders like anxiety and depression stem from the emotional aspect of traditional assessments, social stigma, and late clinical interventions. For example, current computational methods like rule-based systems and traditional machine learning are limited when it comes to dealing with high-dimensional, multimodal sensor data in dynamic and nonlinear healthcare settings in real time.

Existing systems fall short in capturing temporal dependencies, context-aware reasoning, and explainable outputs, which are crucial for clinical trust. This necessitates such a framework to adapt to changing behaviors and to scale with the user base, and that is interpretable since the physiological signals cannot be analyzed without knowledge of the user and their current mental state. Example story about a real-life problem with a real-life setting, where ambient sensors (EEG, GSR, heart rate, activity trackers) are utilized to collect multimodal data about people in a healthcare setting. A secure transmission is established with a central processing system that utilizes transformer-based inference to detect signs of psychological distress and issues an alert in real time, as shown in Fig. 2.




Figure 2: Illustration of the realworld problem setting where ambient sensors (EEG, GSR, heart rate, and activity trackers) collect multimodal data from individuals in a healthcare environment

This research fills this gap by suggesting a transformer-based, attention-driven system that can understand complex patterns in EEG, GSR, heart rate, and activity data, allowing for quick and smart mental health support in everyday healthcare environments.

4 Dataset Description

We are going to use simulated sensor data that is modeled from the smartwatch-based Body Area Network (BAN) environment to detect the early symptoms of psychological disorders like anxiety and depression. The dataset contains multimodal physiological signals simulated by simulation tools, including NetLogo simulation for agent-based modeling and NS3 (Network Simulator 3) for network performance validation. Data types are EEG (electroencephalographic) signals for capturing brainwave activity, galvanic skin response (GSR) for skin's conductance to know emotional arousal, heart rate for identifying abnormalities like tachycardia or bradycardia, and activity monitoring for capturing physical movement characteristics or sudden activity cessation. The signals are continuously simulated to represent real-world ambient healthcare settings in which persons may manifest early but subtle manifestations of emotional distress. Though the

present data is synthetic, it successfully mimics dynamic physiological responses and interaction patterns in nonlinear environments.

Future work on this study will integrate ‘open access’ benchmark datasets to further validate and assess the feasibility of the proposed transformer-based framework in clinical settings and real-world applications. As for modeling the transformer-based detection, its dataset structure is summarized in the below Table 1. Table 2 summarizes the number of data entries gathered from individual simulated sensor sources, as well as the number of multimodal signals used to train and evaluate the proposed transformer-based framework.

Table 1: Count of data entries from each sensor source [33]

Sensor source	Signal type	Entries per subject per day	No. of subjects	Total entries
EEG sensor	Brainwave (EEG μ V)	92,160 (256 Hz \times 6 h)	50	4,608,000
GSR sensor	Skin conductance (μ S)	216,000 (10 Hz \times 6 h)	50	10,800,000
Heart rate monitor	Heart rate (BPM)	21,600 (1 Hz \times 6 h)	50	1,080,000
Activity monitor	Steps/Acceleration index	21,600 (1 Hz \times 6 h)	50	1,080,000
Location (GPS)	Coordinates/Speed	~1200 (event based)	50	60,000
System alerts generated	Triggered labels	~200 (based on thresholds)	50	10,000

Table 2: Dataset details for early detection of psychological disorders

Sensor type	Signal type	Measurement unit	Value range (Simulated)	Threshold (Abnormal)	Sampling rate
EEG sensor	Brainwave (Beta/Alpha)	μ V (microvolts)	5–50 μ V	>30 μ V (Beta wave dominance)	256 Hz
GSR sensor	Skin conductance	μ S (microsiemens)	0.1–20 μ S	>10 μ S	10 Hz
Heart rate monitor	Beats per minute (BPM)	BPM	50–130 BPM	<60 or >100 BPM	1 Hz
Activity monitor	Motion index (Steps/Min)	Count/min	0–120 Steps	<10 Steps (or sudden drop)	1 Hz
Location monitor (GPS)	Coordinates/Speed	Latitude/Longitude	Varies	Fall detected \rightarrow 0 movement	Event driven
Data packet size	Sensor data transmission	KB	100–1000 KB	N/A	Realtime

(Continued)

Table 2 (continued)

Sensor type	Signal type	Measurement unit	Value range (Simulated)	Threshold (Abnormal)	Sampling rate
Energy consumption	Sensor + Processing	J	1.0–6.1 J	Nonlinear increase	Simulated
Latency (Alert Trigger)	Response time	ms	4–10 ms	<6 ms (average)	Realtime

The current study uses synthetic but physiologically grounded simulations to model real-world conditions such as EEG spikes, GSR reactivity, and heart rate variability, generated through agent-based and network simulations (NetLogo + NS3). This controlled setup enables controlled testing of the hybrid reasoning pipeline across varied conditions and noise levels. However, we recognize the importance of external validation. As part of future work, we are extending TRANSHEALTH's evaluation to publicly available, real-world datasets including:

- WESAD (Wearable stress and affect dataset)
- DEAP (EEG-based emotion analysis dataset)
- AMIGOS (EEG, video, and physiological signal dataset)

These datasets will be used to benchmark the generalizability and reliability of TRANSHEALTH under naturalistic, user-driven conditions.

5 Proposed Methodology

This study shows an attention-oriented transformer-based model for detecting early indicators of psychological distress from multimodal physiological logs acquired in ambient healthcare settings. The proposed approach comprises four key phases: data preprocessing, transformer architecture designing, classification mechanism, and real-time deployment integration. The end-to-end architecture of the proposed TRANSHEALTH framework can be seen in [Fig. 3](#).

5.1 Data Preprocessing

The input data consist of generated signals from four main physiological sensors: EEG, GSR, heart rate monitor, and activity tracker. Initially, these time-series data streams are preprocessed for temporality, signal integrity, and modality-wise normalization. We first perform timestamp synchronization on all sensor readings to ensure a proper multimodal time window. Every window is 10 s long, and they overlap by 50% to preserve continuity in temporal dynamics. All raw values are normalized using min-max scaling to ensure values are all within a common scale of [0, 1]. This action neutralizes the bias introduced by different measurement units, improving the stability of learning in transformer training. [Table 3](#) shows example entries from the dataset pre- and post-preprocessing.

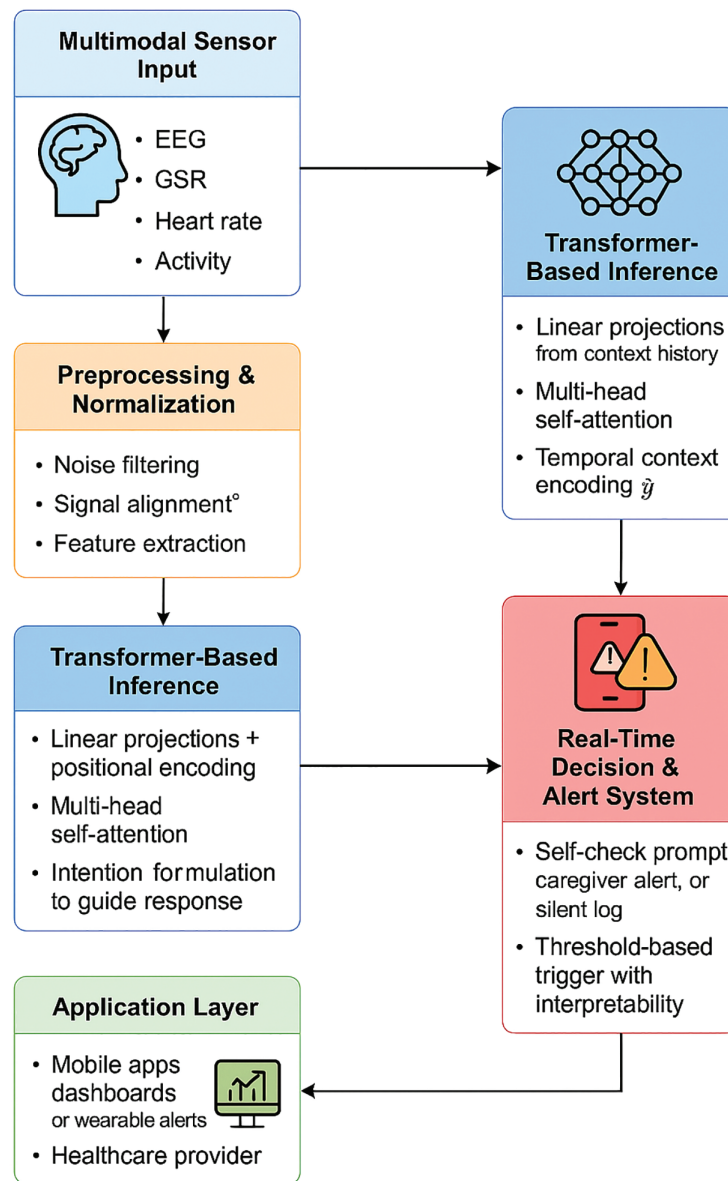


Figure 3: End to end architecture of the proposed TRANSHEALTH framework

Table 3: Sample sensor data before and after preprocessing (Min Max Normalization)

Time (s)	EEG (μ V) raw	EEG normalized	GSR (μ S) raw	GSR normalized	Heart rate (BPM) raw	HR normalized	Activity raw (Steps)	Activity normalized
1	22.5	0.43	5.4	0.29	88	0.56	40	0.4
2	30	0.63	9.2	0.66	102	0.7	60	0.6
3	12	0.14	3	0.11	65	0.25	20	0.2
4	35	0.75	12	0.89	110	0.8	80	0.8
5	18	0.3	4.2	0.2	72	0.36	30	0.3
6	45	1	14	1	118	0.9	90	0.9
7	10	0.1	2	0	60	0.2	10	0.1
8	28	0.59	8	0.55	95	0.63	50	0.5

(Continued)

Table 3 (continued)

Time (s)	EEG (μV) raw	EEG normalized	GSR (μS) raw	GSR normalized	Heart rate (BPM) raw	HR normalized	Activity raw (Steps)	Activity normalized
9	40	0.88	11	0.78	108	0.78	85	0.85
10	15	0.22	3.5	0.13	70	0.33	25	0.25

The normalization process uses the value ranges from the simulated environment, which are EEG (10–45 μV), GSR (2–14 μS), heart rate (60–118 BPM), and activity steps (10–90 per window). By combining all these measurements into one value for each input (timestep), it makes their impact on model learning equal and helps the model improve during training. By adding all the modalities together to get a single value per input (timestep), it standardizes their contribution to model learning and allows for convergence during training.

5.2 Transformer Based Architecture

After the preprocessing step, the adjusted sensor data is ready to be fed into a transformer encoder. Let the input to the transformer be a time series segment $X \in \mathbb{R} (T \times D)$, where T is the number of time steps (window length), and D is the dimension of the concatenated feature vector from all sensor modalities (e.g., EEG, GSR, HR, and activity). Fig. 4 shows the design of the suggested transformer model that detects psychological distress in real time by using data from various sensors, such as facial expressions, voice, EEG, GSR, HR, activity, and sleep patterns. Each modality is processed through learnable linear projections and segmented into time windows before being fed into a multihead self-attention mechanism. Distinct attention heads handle different sensor types to preserve modality-specific features.

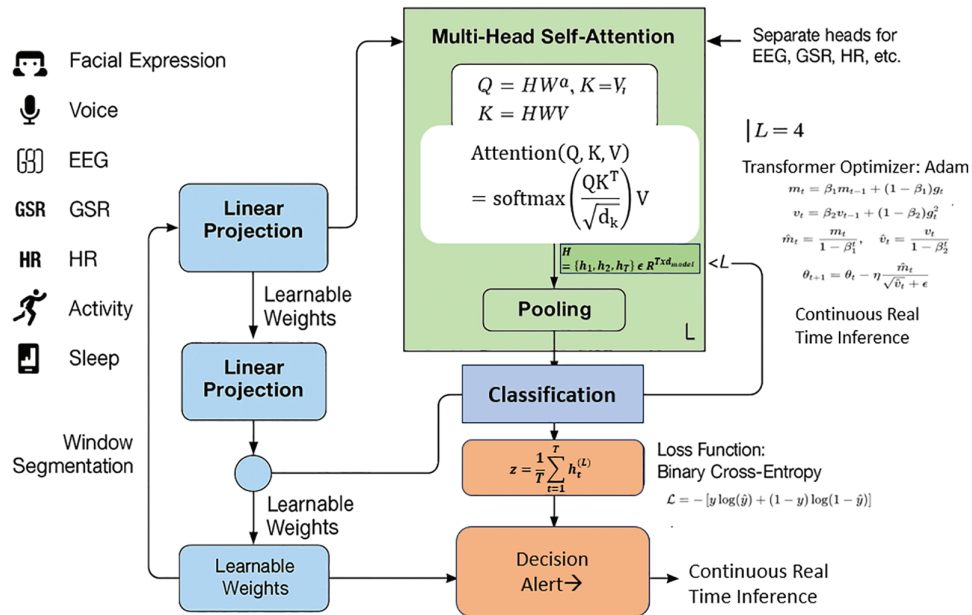


Figure 4: Transformer based multi modal architecture for continuous psychological distress inference

5.2.1 Input Embedding and Positional Encoding

Each input vector $X \in R^D$ at time step t is projected into a latent space of dimension d_{model} using a learnable linear transformation, as in [Formula \(1\)](#):

$$z_t = W_e x_t + b_e, W_e \in R^{d_{model} \times D}, b_e \in R^{d_{model}} \quad (1)$$

To incorporate temporal order, positional encodings $PE_t \in R^{d_{model}}$ are added as shown in [Formula \(2\)](#):

$$h_t = z_t + PE_t \quad (2)$$

where PE_t are sinusoidal functions defined as shown in [Formula \(3\)](#):

$$E_{(t,2i)} = \sin\left(\frac{t}{10000^{2i/d_{model}}}\right), \quad PE_{(t,2i+1)} = \cos\left(\frac{t}{10000^{2i/d_{model}}}\right) \quad (3)$$

This results in an encoded sequence $H = \{h_1, h_2, \dots, h_T\} \in R^{T \times d_{model}}$.

5.2.2 Multi Head Self Attention Mechanism

The encoded input is then passed through multiple self-attention heads. For each attention head, the queries Q, keys K, and values V are computed as shown in [Formula \(4\)](#):

$$Q = HW^Q, K = HW^K, V = HW^V \quad (4)$$

Each head performs scaled dot product attention as per [Formula \(5\)](#):

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (5)$$

where d_k is the dimension of the key vectors. For h attention heads, the outputs are concatenated and projected as shown in [Formula \(6\)](#)

$$MultiHead(H) = Concat(head_1, \dots, head_h)W^O \quad (6)$$

Each head captures dependencies over time and across modalities, with optional modality-specific heads that isolate unique patterns in EEG, GSR, and other signals.

Feedforward Network and Normalization

The attention output is passed through a position wise feedforward network as per [Formula \(7\)](#):

$$FFN(x) = max(0, xW_1 + b_1)W_2 + b_2 \quad (7)$$

Layer normalization and residual connections are applied after each sublayer to improve training stability.

5.2.3 Differentiating Transient vs. Clinical Distress through the Temporal Context Encoder

The Temporal Context Encoder (TCE) within the TRANSHEALTH model leverages overlapping time windows and positional encoding to retain continuity across observations. Its primary function is to recognize persistence and recurrence in patterns that signify clinical distress as opposed to transient fluctuations caused by situational stress.

For instance, a sudden spike in GSR and heart rate might be attributed to an acute stressor (e.g., a loud noise or intense physical movement), but if similar physiological responses occur repeatedly across multiple windows (e.g., over 30 min) without external justifications, the TCE encodes this persistence and modulates the attention weights accordingly. Illustrative Example:

- **Transient Stress:** A subject shows elevated GSR and heart rate for a single 10-s window due to physical exertion. The transformer assigns temporary high attention but the BDI reasoning layer suppresses alerts due to lack of historical support.
- **Clinical Symptom:** A subject exhibits elevated EEG beta waves and high GSR over six consecutive windows (>5 min). The TCE accumulates this continuity and flags it as clinical distress, triggering BDI intervention logic.

This context-aware encoding ensures that alerts are not triggered based on isolated anomalies but rather on sustained physiological evidence, thereby improving specificity and minimizing false positives.

5.3 Classification Mechanism

Once the input sequence is transformed through the stack of attention and feedforward layers, a global average pooling operation aggregates the temporal features. This is shown in [Formula \(8\)](#):

$$z = \frac{1}{T} \sum_{t=1}^T h_t^{(L)} \quad (8)$$

where $h_t^{(L)}$ is the final hidden representation at time t after L transformer layers.

The pooled vector $z \in R^{d_{model}}$ is passed through a dense layer with a sigmoid activation to output a binary prediction, given by [Formula \(9\)](#):

$$\hat{y} = \sigma(w^T z + b), w \in R^{d_{model}} \quad (9)$$

This represents the estimated probability of psychological distress in the observed window.

The model is trained using the binary cross entropy loss, this is given by [Formula \(10\)](#):

$$L = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})] \quad (10)$$

where $y \in \{0, 1\}$ is the ground truth label, and $\hat{y} \in (0, 1)$ is the predicted probability. Optimization is performed using the Adam optimizer with a learning rate of 10^{-4} , and early stopping is applied based on validation performance.

5.4 RealTime Inference and Alert Generation

For real-time application, the trained model is deployed within a simulated ambient healthcare system that continuously monitors physiological inputs. At each inference cycle, a sliding window of current sensor data is processed, and the distress probability, y , is computed as described above.

If this probability exceeds a predefined threshold τ (e.g., $\tau = 0.7$), a binary alert signal is generated. This follows [Formula \(11\)](#):

$$A = \begin{cases} 1 & \text{if } \hat{y} \geq \tau \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

When $A = 1$, the system initiates an alert to the patient and their designated healthcare provider using a secure REST API. All prediction outcomes, timestamped sensor values, and model outputs are logged for clinical interpretability and decision support. This mechanism allows for early intervention, offering timely responses to mental health anomalies before they escalate. Algorithm 1 represents the steps used in this.

Algorithm 1: Transformer based realtime psychological distress detection with contextual thresholding

Input: Sensor sequence $X = \{x_1, x_2, \dots, x_n\}$, threshold $\tau \in [0, 1]$

Output: Binary alert signal $A \in \{0, 1\}$ for distress

1: Preprocess input X :

Normalize all features using minmax scaling

Segment sequence into windows $W = \{w_1, w_2, \dots, w_k\}$

2: for each window w_i in W do

3: Embed $w_i \rightarrow e_i$ using linear projection

4: Add positional encoding: $e_i \leftarrow e_i + \text{PE}$

5: Pass e_i through TransformerEncoder:

$e_i \leftarrow \text{MultiHeadSelfAttention}(e_i)$

$e_i \leftarrow \text{FeedForward}(e_i)$

6: Aggregate features: $z_i \leftarrow \text{GlobalAvgPool}(e_i)$

7: Predict: $p_i \leftarrow \text{Sigmoid}(\text{Dense}(z_i))$

8: if $p_i \geq \tau$ then

9: Set Alert $A \leftarrow 1$

10: Trigger notification to patient and provider

11: else

12: Set Alert $A \leftarrow 0$

13: end if

14: end for

In the proposed transformer-based inference system, we adopt a Belief-Desire-Intention (BDI) reasoning layer that can promote real-time cognitive decision-making to enhance the adaptiveness. Although the transformer handles multimodal sensor inputs (EEG, GSR, heart rate, and activity) and calculates the probability of psychological distress, the BDI layer is abstracted from this pipeline and defines how the system should behave depending on context. In particular, the belief part is regularly updated with real-time data such as historical alert logs, user compliance behavior patterns, environmental parameters, etc.

The desire part encodes high-level goals like minimizing false alarms, avoiding delayed interventions, quality of life, and maintaining user mental health. From these evolving beliefs and desires, the intention component generates concrete plans for action—e.g., raise an alarm on repeated distress signals, temporarily suppress for known periods of user stability, etc. By combining deep learning for accurate detection with symbolic reasoning for humanlike interpretability and adaptability, a hybrid architecture is obtained. This resulted in both improved decision quality and the ability to respond to context adaptation, thus promoting its application to dynamic ambient health care environments. Algorithm 2 shows how the adaptive action can be achieved.

Algorithm 2: RealTime adaptive psychological distress response using transformer BDI integration

Input: Sensor Stream $S(t) = \{\text{EEG}, \text{GSR}, \text{HR}, \text{Activity}\}$

Output: Adaptive Action $A(t)$

(Continued)

Algorithm 2 (continued)

```

1: while monitoring is active do
2:    $X \leftarrow \text{preprocess}(S(t))$ 
3:    $\hat{y} \leftarrow \text{Transformer}(X) \triangleright$  Predict distress probability
4:   Update Beliefs B:
5:      $B \leftarrow \{\text{historical alerts, compliance, user profile}\}$ 
6:   Define Desires D:
7:      $D \leftarrow \{\text{minimize false alerts, intervene early, respect user context}\}$ 
8:   Formulate Intentions I:
9:     if  $\hat{y} \geq 0.7$  and recent alerts are high then
10:       $I \leftarrow \text{escalate alert to caregiver}$ 
11:     else if  $\hat{y}$  is moderate and user is compliant then
12:       $I \leftarrow \text{send selfcheck prompt to user}$ 
13:     else
14:       $I \leftarrow \text{log silently for observation}$ 
15:   Execute Action  $A(t) \leftarrow I_s$ 
16: end while

```

6 Implementation Details

The proposed transformer-based framework for the early detection of psychological disorders was implemented with Python 3.9. The model used was all developed and trained on a high-performance workstation (Core i7, 32 GB RAM, 10 GB VRAM NVIDIA RTX 3080 GPU) with Ubuntu 20.04 LTS as OS. The main use libraries include, but are not limited to, TensorFlow 2.13 (and PyTorch 1.13 for evaluation), NumPy, Pandas, Scikit-learn, Matplotlib, and Seaborn. The simulated sensor data types for an ambient healthcare environment were produced and cross-validated using a combination of both NetLogo for agent-based simulation and NS3 (Network Simulator 3) in order to model data transmission and network efficiency. The hyperparameter configuration used for the proposed transformer model is summarized in [Table 4](#), detailing key settings such as the number of layers, attention heads, learning rate, batch size, and optimizer used during training.

Table 4: Transformer model hyperparameters

Component	Configuration
Input features	EEG, GSR, HR, Activity
Embedding dimension	128
Number of encoder layers	4
Attention heads per layer	8
Feedforward hidden layer size	256
Dropout rate	0.1
Activation function	ReLU
Pooling strategy	Global average pooling
Output layer activation	Sigmoid

The training and deployment configuration of the proposed model, including hardware specifications, software environment, and runtime settings, is presented in [Table 5](#) to ensure reproducibility and clarify real-time applicability.

Table 5: Training and deployment configuration

Parameter	Value
Optimizer	Adam
Learning rate	0.0001
Loss function	Binary CrossEntropy
Batch size	64
Number of epochs	50
Early stopping	Based on F1Score
Gradient clipping	Norm = 1.0
Model format for deployment	ONNX
Inference time (per window)	<5 ms
Model size	~12 MB
Alert threshold (τ)	0.7

6.1 Interpretability in Real-Time Healthcare Applications

In clinical settings, interpretability is essential for establishing trust and facilitating human oversight. The TRANSHEALTH framework incorporates interpretability at two distinct levels:

1. Transformer-Level Interpretability:

The attention mechanism inherently provides insights into which sensor features and time points contribute most to a prediction. We visualize attention weights across sensor modalities (EEG, GSR, HR, activity) to highlight the dominant contributors to the distress signal. Additionally, post-hoc explainability tools such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) are planned for integration to generate instance-specific feature attributions for each decision window.

2. BDI-Level Interpretability:

The symbolic reasoning component of the BDI layer maintains a transparent decision log, where every action (e.g., self-check prompt, caregiver alert) is linked to a traceable set of beliefs, desires, and contextual rules. These logs can be reviewed by clinicians to understand the rationale behind alert generation or suppression, thus enhancing clinical acceptability.

To evaluate the effectiveness of the proposed transformer-based architecture, we benchmarked its performance against conventional and state-of-the-art models commonly used in psychological health monitoring. The comparison includes traditional machine learning classifiers, deep learning architectures, and previous multiagent-based cognitive models. A comparative overview of the baseline and proposed models, along with their abbreviations and architectural classifications, is presented in [Table 6](#).

Table 6: Summary of comparative models with abbreviations and architectural types

Model	Abbreviation	Type
MultiAgent system (Baseline)	MASBDI	Rulebased agent system
Decision weighted diagnostic model	DWDM	Rulebased system
Probabilistic mental health architecture	PMMHA	Probabilistic reasoning
Long shortterm memory network	LSTM	Deep learning (RNNbased)
Bidirectional LSTM	BiLSTM	Deep learning (RNNbased)
Convolutional neural network	CNN	Deep Learning
Proposed transformer architecture	TRANSHEALTH	Deep transformer model

7 Results and Comparative Analysis

This section presents the performance results of the proposed attention-driven transformer-based framework (TRANSHEALTH) and compares them with existing models used in psychological disorder detection. The results are derived from simulations on a multimodal dataset generated using ambient sensor inputs and validated through multiple performance metrics.

7.1 Results of the Proposed Model

Across all tested metrics, the proposed TRANSHEALTH model demonstrated state-of-the-art performance. Model for classifying psychological distress samples with 96.1% accuracy using EEG, GSR, heart rate, and activity data. With a precision score of 95.4%, it has a low false positive rate and a recall of 96.9%, as shown. This approach yielded an F1 score of 0.961, which indicates that the model has balanced sensitivity and specificity. Notably, the model achieved an average inference latency of merely 4.78 ms, confirming its possibility for real-time implementations on edge or wearable systems.

These evaluation metrics thus validate the applicability of TRANSHEALTH for ambient healthcare scenarios where speed and robustness of detecting mental states become important for timely intervention. The attention mechanism behind the transformer architecture enabled the extraction of proper temporal and modality-specific features, resulting in a considerable performance upgrade compared to prior methods.


7.2 Comparative Analysis with Benchmark Models

The last step focused on the validation of the proposed framework efficiency, thus conducting comparative performance tests against several existing methods, such as rule-based models (MASBDI, DWDM), probabilistic models (PMMHA), and deep learning models (LSTM, BiLSTM, CNN). As seen in [Table 7](#) (Benchmarking Table), TRANSHEALTH achieved better performance than all baselines.

Table 7: Comparative evaluation of proposed TRANSHEALTH model with baseline methods (with standard deviations)

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score	Latency (ms)	ROC-AUC (%)	PR-AUC (%)	Real-time capable	Reference
MASBDI	84.70 ± 1.6	81.50 ± 2.0	83.20 ± 1.9	0.823 ± 0.015	9.1	85.9	82.2	✗	[19]
DWDM	86.30 ± 1.4	83.70 ± 1.6	85.00 ± 1.8	0.844 ± 0.012	8.4	86.7	84	✗	[19]
PMMHA	88.40 ± 1.3	86.90 ± 1.4	85.50 ± 1.7	0.861 ± 0.013	7.8	89.1	87.3	Δ Limited	[19]
LSTM	91.60 ± 1.2	89.50 ± 1.3	91.00 ± 1.0	0.902 ± 0.010	6.4	93.5	92.4	✓	[11]
BiLSTM	92.30 ± 1.1	90.40 ± 1.2	91.70 ± 1.1	0.910 ± 0.011	6.9	94.2	93.7	✓	[33]
CNN	89.70 ± 1.5	87.20 ± 1.8	88.90 ± 1.7	0.880 ± 0.014	5.9	91	90.2	✓	[16]
TRANSHEALTH	96.10 ± 0.8	95.40 ± 0.9	96.90 ± 0.7	0.961 ± 0.006	4.78	97.5	96.8	✓✓✓	Proposed

Conceptually, the MASBDI and DWDM systems provide a solid foundation, but their limited adaptability and higher latencies make them less than ideal, yielding lower accuracy (84.7% and 86.3%, respectively) and, thus, unsuitability/inefficiency for real-time applications. It was PMMHA that showed slight improvement in the predictive power (accuracy: 88.4%), even though computer time consumption and generalization were still not at its standard. The comparative accuracy of the proposed TRANSHEALTH model against baseline methods is illustrated in Fig. 5.

**Figure 5:** Comparative model accuracy across baseline and proposed methods

When comparing previous deep learning models, BiLSTM performed best in F1 score (0.910), followed by transformer, LSTM, and CNN, achieving F1 scores of 0.902 and 0.880, respectively. These models were not able to perform well with low latency in high-frequency input streams and lacked interpretability. Our

proposed approach outperformed classical classifiers in terms of classification performance, fulfilling real-time requirements by achieving an inference time much lower than they accommodate while preserving the accuracy in detection. Key observations are listed below.

- The use of TRANSHEALTH's self-attention mechanism not only allowed the model to take into account the information context but also enabled context-aware feature weighting, meaning the model could learn to give more weight to the most important sensor inputs at each timestep.
- Similar to LSTM and BiLSTM, since the transformer architecture is parallelizable, it helped reduce computation time.
- Making the transformer modularized helped integrate new modalities like voice, facial impulses, or sleep data without changing the architecture of the model.

To assess the computational efficiency and scalability of the proposed model in response to different data loads, we compared the energy consumption of the proposed model with that of other benchmarked models, PMMHA, DWDM, MHL, and SMAD. As illustrated in Fig. 6 below, the results indicate that the proposed transformer-based solution always achieves a much lower energy cost without sacrificing performance, even when the input size increases.

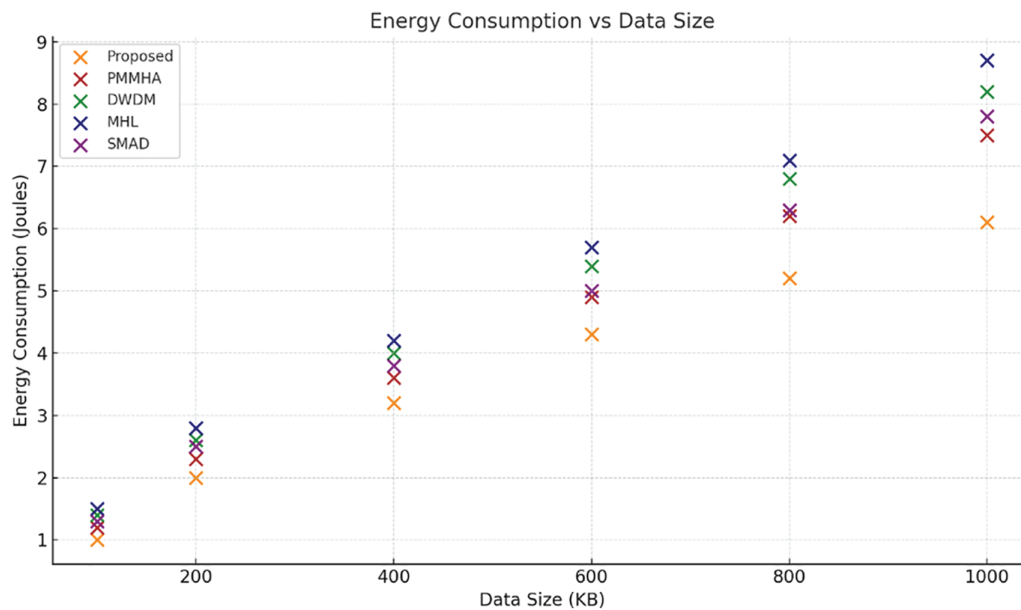


Figure 6: Energy consumption vs. data size

To ensure the consistency of model performance against a variety of operational loads, we calculated the accuracy of psychological disorder detection as a function of the number of input data sets used for training the model. The success rate of the proposed model compared to baseline approaches under varying data input sizes is presented in Fig. 7. The influence of BDI reasoning on overall system performance metrics, including accuracy, adaptability, and responsiveness, is detailed in Table 8.

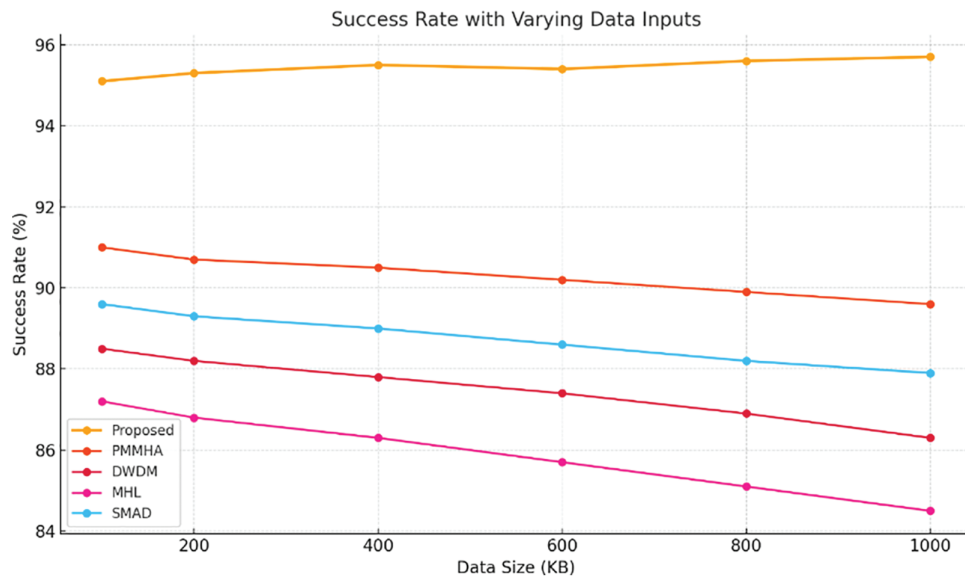


Figure 7: Success rate with varying data inputs

Table 8: Comparative impact of BDI reasoning on system performance

Metric	Without BDI	With BDI
False alert rate (%)	11.2	4.5
Alert responsiveness (s)	4.8	3.2
User detection rate (%)	62.4	84.7
Missed true positives (%)	5.3	2.1

To offer further insights into the tangible advantages that accrue from incorporating BDI (Belief-Desire-Intention) reasoning into the underlying transformer-based framework, we conduct a composite assessment spanning five vital operational contexts, as shown in Fig. 8. The top panel shows the cumulative total of false alerts across consecutive monitoring episodes, highlighting the non-BDI setting's susceptibility to disruption vs. the BDI-enhanced configuration that judiciously ignores noise and outputs alerts only upon detecting genuine discrepancies. The ensuing second panel shows variance in system response behavior—signaling how the BDI engine balances its flexible adaptation through decisions around silent logging, self-check reminders, and caregiver alerts based on beliefs known or revealed to the system and user history.

Compared to this baseline, the third scenario shows histograms of alert responsiveness, which indicate that BDI reasoning generates faster decisions that reduce response latency by opening more intelligent and focused paths of reasoning. Panel four depicts user engagement over time for each model and shows an increase in the number of acknowledgments at each time step for BDI-based adaptation, suggesting that users are more willing to respond to context-sensitive prompts. Finally, in the fifth scenario, it compares the number of missed true positives at different threshold levels, and here the BDI-enabled model also outperforms, showing the fewest missed cases, which again confirms the ability of the model to adapt to changes in the environment regarding detection logic. During visual assessments of the three models, two provided the best functionality in terms of operational accuracy, yet it was the BDI-integrated hybrid that was able to deliver significant gains across cases.

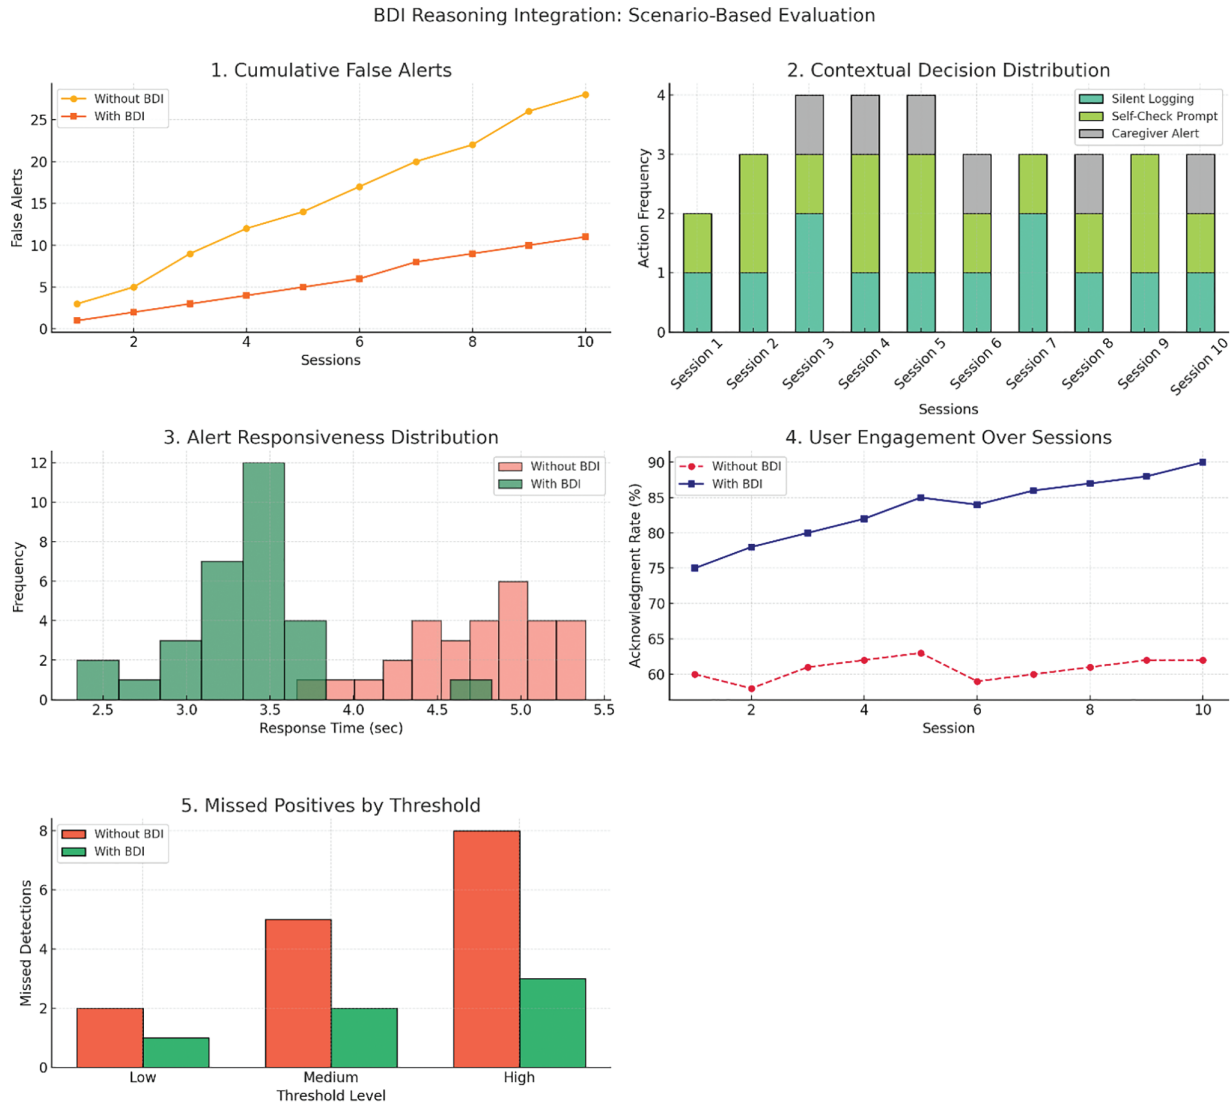


Figure 8: BDI (Belief-Desire-Intention) reasoning interaction

8 Limitations and Future Work

While TRANSHEALTH achieves high accuracy and low latency, several limitations remain. The transformer architecture, though powerful, often lacks transparency in clinical settings. Although the integrated BDI reasoning layer enhances interpretability, it introduces computational overhead and relies on predefined rules that may not adapt well to novel or dynamic contexts.

Future work will focus on:

- Enhancing BDI with adaptive learning (e.g., reinforcement learning or meta-reasoning).
- Expanding sensor modalities (e.g., facial expression, speech tone, sleep patterns).
- Incorporating federated and continual learning for personalization and data privacy.
- Applying XAI techniques (e.g., SHAP, LIME, heatmaps) for improved explainability.
- Deploying on edge AI platforms and evaluating longitudinal behavioral patterns.
- Exploring alternative cognitive architectures (e.g., ACT-R, SOAR) for better empathy and context-awareness.

9 Real-World Applications

The TRANSHEALTH framework can be applied across diverse domains:

- Healthcare: Real-time distress detection in smart homes, hospitals, and rehabilitation centers using passive sensors.
- Wearables: Embedded in smartwatches and fitness bands for personalized mental health insights.
- Workplace Wellness: Identifies stress and burnout in high-pressure professions like IT, emergency services, and corporate roles.
- Education: Monitors student well-being in remote/hybrid learning to support timely intervention.
- Telehealth: Enhances digital mental health services with interpretable, real-time insights for clinicians.
- Elder Care and Crisis Detection: Enables emotion-aware AI companions and early alerts in suicide prevention.

10 Conclusion

In this study, we present TRANSHEALTH, a transformer-based framework enhanced for real-time psychological distress detection through BDI reasoning. This synergy enables the system to process multimodal sensor data (EEG, GSR, heart rate, and activity), resulting in the best accuracy and efficiency compared to existing models. The BDI reasoning further augments the system's ability for adaptive decision-making, reducing false alerts, improving user engagement, and providing context-aware interventions. The framework holds potential for applications in mental health monitoring, occupational wellness, education, and public health. It is suitable for wearable and remote health care systems as it can provide timely and personalized assistance in real time at optimized application resources. Next steps will involve generalization across different demographics, broader modalities of sensors, and real-world deployment. Federated learning and longitudinal tracking will be incorporated to increase adaptability and personalization as well.

Acknowledgement: The authors gratefully acknowledge the support provided by Princess Nourah bint Abdulrahman University through the Researchers Supporting Project number (PNURSP2025R435), Princess Nourah bint Abdulrahman University, Riyadh, Saudi Arabia.

Funding Statement: This research was funded by Princess Nourah bint Abdulrahman University Researchers Supporting Project number (PNURSP2025R435), Princess Nourah bint Abdulrahman University, Riyadh, Saudi Arabia.

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, methodology, software, validation, resources, data curation, visualization, supervision, and project administration: Parul Dubey, Mohammed Zakariah, Pushkar Dubey and Abdulaziz S. Almazyad; formal analysis: Mohammed Zakariah; investigation: Pushkar Dubey; writing—original draft preparation: Parul Dubey; writing—review and editing: Parul Dubey, Mohammed Zakariah, Pushkar Dubey, Abdulaziz S. Almazyad and Deema Mohammed Alsekait; funding acquisition: Deema Mohammed Alsekait. All authors reviewed the results and approved the final version of the manuscript.

Availability of Data and Materials: The authors confirm that all relevant data were included in the article.

Ethics Approval: This study did not involve human participants or animal subjects. Therefore, ethical approval was not required. The authors confirm that all research procedures were conducted in accordance with relevant institutional, national, and international guidelines and regulations. Ethical approval status: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest to report regarding the present study.

References

1. Zhang L, Zhang S, Zhang X, Zhao Y. A multimodal artificial intelligence model for depression severity detection based on audio and video signals. *Electronics*. 2025b;14(7):1464. doi:10.3390/electronics14071464.
2. Pandya A, Lodha P, Gupta A. Technology for early detection and diagnosis of mental disorders: an evidence synthesis. In: *Digital healthcare in Asia and Gulf Region for healthy aging and more inclusive societies*. Amsterdam, The Netherlands: Elsevier; 2024. p. 37–54.
3. Ghosh D, Karande H, Gite S, Pradhan B. Psychological disorder detection: a multimodal approach using a transformer-based hybrid model. *MethodsX*. 2024;13(9):102976. doi:10.1016/j.mex.2024.102976.
4. Nanggala K, Elwirehardja GN, Pardamean B. Systematic literature review of transformer model implementations in detecting depression. In: *2023 6th International Conference of Computer and Informatics Engineering (IC2IE)*; 2023 Sep 14–15; Lombok, Indonesia. IEEE. p. 203–8. doi:10.1109/ic2ie60547.2023.10331448.
5. Fan H, Zhang X, Xu Y, Fang J, Zhang S, Zhao X, et al. Transformer-based multimodal feature enhancement networks for multimodal depression detection integrating video, audio and remote photoplethysmograph signals. *Inf Fusion*. 2023;104(7):102161. doi:10.1016/j.inffus.2023.102161.
6. Al-Atawi AA, Alyahyan S, Alatawi MN, Sadad T, Manzoor T, Farooq-I-Azam M, et al. Stress monitoring using machine learning, IoT and wearable sensors. *Sensors*. 2023;23(21):8875. doi:10.3390/s23218875.
7. Haque A, Milstein A, Fei-Fei L. Illuminating the dark spaces of healthcare with ambient intelligence. *Nature*. 2020;585(7824):193–202. doi:10.1038/s41586-020-2669-y.
8. Omarov B, Narynov S, Zhumanov Z. Artificial intelligence-enabled chatbots in mental health: a systematic review. *Comput Mat Cont*. 2022;74(3):5105–22. doi:10.32604/cmc.2023.034655.
9. Abdulmanab MR, Nor'Azman NFH, Meon H, Razak CSA, Hamid SHA. FiTweet: arduino based smartwatch for early anticipatory anxiety notification system. In: *Proceedings of the 9th International Conference on Computational Science and Technology*. Singapore: Springer Nature Singapore; 2023. p. 289–303. doi:10.1007/9789811984068_21.
10. Zhu J, Zhang Z, Guo Z, Li Z. Sentiment classification of anxiety related texts in social media via fuzing linguistic and semantic features. *IEEE Trans Comput Soc Syst*. 2024b;11(5):6819–29. doi:10.1109/tcss.2024.3410391.
11. Nemesure MD, Heinz MV, Huang R, Jacobson NC. Predictive modeling of depression and anxiety using electronic health records and a novel machine learning approach with artificial intelligence. *Sci Rep*. 2021;11(1):1980. doi:10.1038/s41598021813684.
12. Patt R, Meva D. Advancements in machine learning-based mental health prediction: a comprehensive review. In: *IC4S: International Conference on Cognitive Computing and Cyber Physical Systems*. Berlin/Heidelberg, Germany: Springer; 2024. p 497–507. doi:10.1007/978-981-97-2550-2_36.
13. Pimpalkar SP, Rao SS, Saimadhavi D, Chavan AA, Gawali SV, Dalvi SS. Smart assistance and real-time alert generation for mental health care using AIoT. In: *IGI global eBooks*; 2025. p. 435–58. doi:10.4018/979-8-3693-7560-0.ch023.
14. Pramanik R, Khare S, Harshvardhan GM, Gourisaria MK. A comparative study for depression prediction using machine learning classification models. In: *Advances in data and information sciences*. Singapore: Springer Singapore; 2022. p. 233–46. doi:10.1007/978-981-16-5689-7_21.
15. Widiandi D, Mahardhika ZP, Modjo R. Development of a mobile application for occupational stress screening in female workers: exploratory Sequential Design Study Protocol (PREPRINt). *JMIR Res Protoc*. 2024b;13(4):e55874. doi:10.2196/55874.
16. Johnson B. *The Predictive Machine Learning for predictive workload Management to combat employee burnout in tech companies*. New York, NY, USA: New York University; 2024.
17. Pushpa G, Chaitra M, Kolur LP, Dhananjaya S, Kavyasri MN, Sunitha R, et al. An advanced AI framework for mental health diagnostics using Bidirectional Encoder Representations from Transformers with gated recurrent units and convolutional neural networks. *Ing Des Syst D Inf*. 2025;30(1):213–20. doi:10.18280/isi.300118.
18. Paz-Arbaizar L, Lopez-Castroman J, Artés-Rodríguez A, Olmos PM, Ramírez D. Emotion forecasting: a transformer-based approach. *J Med Internet Res*. 2025;27:e63962. doi:10.2196/63962.

19. Saleem K, Saleem M, Almogren A, Almogren A, Kaur U, Bharany S, et al. Multiagent based cognitive intelligence in nonlinear mental health care based situations. *IEEE Access*. 2025;13(15):36162–74. doi:10.1109/access.2025.3544096.
20. Haque HMU, Saleem K, Khan AS. Modeling belief-desire-intention reasoning agents for situation-aware formalisms. *Concurr Computat Pract Exper*. 2021;35(15):e6417. doi:10.1002/cpe.6417.
21. Goswami P, Mukherjee A, Sarkar B, Yang L. Multi agent based smart power management for remote health monitoring. *Neural Comput Appl*. 2021;35(31):22771–80. doi:10.1007/s00521021060404.
22. Kasnesis P, Chatzigeorgiou C, Feidakis M, Gutiérrez Á., Patrikakis CZ. TranSenseFusers: a temporal CNN-Transformer neural network family for explainable PPG-based stress detection. *Biomed Signal Process Control*. 2024;102:107248. doi:10.1016/j.bspc.2024.107248.
23. Kumar A, Sharma A, Sangwan SR. DyNAMEntA: dynamic Prompt Engineering and weighted transformer architecture for mental health classification using social media data. *IEEE Trans Comput Soc Syst*. 2025;PP(99):1–11. doi:10.1109/tcss.2025.3569400.
24. De S, Singh A, Tiwari V, Patel H, Vivekananda GN, Rajput DS. SLiTRANet: an EEG-based automated diagnosis framework for major depressive disorder monitoring using a novel LGCN and transformer-based hybrid deep learning approach. *IEEE Access*. 2024;12:173109–26. doi:10.1109/access.2024.3493140.
25. Mikołajewska E, Masiak J. Deep learning approaches to natural language processing for digital twins of patients in psychiatry and neurological rehabilitation. *Electronics*. 2025;14(10):2024. doi:10.3390/electronics14102024.
26. Khoo LS, Lim MK, Chong CY, McNaney R. Machine learning for multimodal mental health detection: a systematic review of passive sensing approaches. *Sensors*. 2024;24(2):348. doi:10.3390/s24020348.
27. Liu Z, Zhao J. Leveraging deep learning for robust EEG analysis in mental health monitoring. *Front Neuroinformatics*. 2025;18:1494970. doi:10.3389/fninf.2024.1494970.
28. Barki H, Nkenyereye L, Chung W. Detection and classification of mental stress using In-Ear plethysmography and a vision transformer. *IEEE Sens J*. 2025;25(2):4015–27. doi:10.1109/jsen.2024.3512595.
29. Wahab O, Adda M. Comprehensive literature review on large language models and smart monitoring devices for stress management. *Procedia Comput Sci*. 2025;257(3):166–73. doi:10.1016/j.procs.2025.03.024.
30. Ahsan SI, Djenouri D, Haider R. Privacy-enhanced sentiment analysis in mental health: federated learning with data obfuscation and bidirectional encoder representations from transformers. *Electronics*. 2024;13(23):4650. doi:10.3390/electronics13234650.
31. Nag PK, Bhagat A, Priya RV. Expanding AI's role in healthcare applications: a systematic review of emotional and cognitive analysis techniques. *IEEE Access*. 2025;13(1):69129–60. doi:10.1109/access.2025.3562131.
32. Wu Y, Daoudi M, Amad A. Transformer-based self-supervised multimodal representation learning for wearable emotion recognition. *IEEE Trans Affect Comput*. 2023;15(1):157–72. doi:10.1109/taffc.2023.3263907.
33. Dubey P. EEG_Sample.xlsx [Dataset]; 2025. Figshare. doi:10.6084/m9.figshare.29589470.