

Passive Digital Phenotyping for Longitudinal Burnout and Occupational Mental Health Surveillance: A Transformer-Based Explainable Deep Learning Approach Using Smartphone Behavioral Streams

Hardik Srivastava^{1*}, Shiraj Ahmad²

Abstract

Occupational burnout constitutes a pervasive yet chronically under-surveilled public health threat, its insidious temporal evolution rendering episodic self-report instruments structurally inadequate for early detection. This paper introduces BurnoutSense, a passive digital phenotyping framework that continuously harvests eight heterogeneous smartphone behavioral data streams encompassing application usage ecology, communication metadata, geospatial mobility, screen interaction dynamics, inferred sleep rhythmicity, keystroke kinematics, ambient noise exposure, and battery/charging cadence to construct individualized multivariate behavioral signatures sensitive to occupational mental health deterioration. At the computational core, we design the Hierarchical Temporal Transformer (HTT), a novel architecture incorporating irregular-sampling-aware positional encodings, circadian attention masking, and a dual-resolution temporal attention mechanism that simultaneously captures intra-day behavioral variability and inter-week longitudinal drift. To bridge the clinical explainability gap, we introduce the SHAP-Temporal Attribution Module (SHAP-TAM), which decomposes burnout risk predictions along both behavioral-feature and temporal-position axes, generating clinician-interpretable surveillance narratives. A 24-week longitudinal cohort study involving 387 knowledge-sector workers demonstrates that HTT achieves an Area Under the ROC Curve (AUC) of 0.931 for burnout onset prediction, with sensitivity 0.887 and specificity 0.906, significantly surpassing LSTM (AUC 0.847), Bi-GRU (AUC 0.863), Informer (AUC 0.879), and standard Transformer (AUC 0.891) baselines. The framework integrates differential privacy ($\epsilon = 0.8$, $\delta = 10^{-5}$) and federated inference to ensure sensitive behavioral data never leaves the user's device in identifiable form. BurnoutSense pioneers a clinically actionable, ethically grounded, always-on paradigm for occupational mental health surveillance.

Keywords: Digital phenotyping, occupational burnout, hierarchical temporal transformer, explainable AI, federated learning, smartphone behavioral sensing, SHAP attribution, mental health surveillance.

*Author for Correspondence

Hardik Srivastava
E-mail: hardiksatwik2004@gmail.com

Student¹, Assistant Professor², Department of CSE, Greater Noida Institute of Technology, Greater Noida, Uttar Pradesh, India

Received Date: July 04, 2026

Accepted Date: July 17, 2026

Published Date: July 18, 2026

Citation: Hardik Srivastava, Shiraj Ahmad. Passive Digital Phenotyping for Longitudinal Burnout and Occupational Mental Health Surveillance: A Transformer-Based Explainable Deep Learning Approach Using Smartphone Behavioral Streams. International Journal of Behavioral Sciences. 2026; 3(2): 44–52.

INTRODUCTION

The contemporary knowledge economy imposes sustained cognitive, emotional, and social demands upon workers at an unprecedented scale. The World Health Organization classifies burnout a syndrome arising from chronic workplace stress inadequately managed, characterised by emotional exhaustion, depersonalisation, and reduced professional efficacy as an occupational phenomenon with global economic consequences exceeding USD 322 billion annually through absenteeism, presenteeism, and voluntary attrition [1]. Despite widespread recognition, the surveillance infrastructure for

burnout detection remains anchored to psychometric self-report instruments such as the Maslach Burnout Inventory (MBI) and the Copenhagen Burnout Inventory (CBI), administered at discrete intervals [2]. These instruments, while psychometrically validated, are episodic in nature, retrospective in orientation, susceptible to social desirability bias, and architecturally incapable of capturing the granular behavioural trajectories that preceded clinically significant burnout transitions by days to weeks [3].

The pervasive adoption of personal smartphones with over 6.8 billion active devices globally as of 2024 engenders an extraordinary naturalistic sensing infrastructure. The human-smartphone interface continuously generates high-dimensional behavioural data encompassing mobility patterns, communication dynamics, application engagement ecology, nocturnal charging rhythms, typing kinematics, and acoustic environmental exposure. Collectively these streams constitute a digital behavioural phenotype uniquely sensitive to changes in affective, cognitive, and social functioning the precise dimensions that erode during burnout progression [4]. Prior research in mobile health has demonstrated that passively collected smartphone data yields predictive signals for depressive episodes, manic phase onset in bipolar disorder, and social withdrawal characteristic of schizophrenic prodrome [5]. Nevertheless, occupational burnout, which is distinguished from clinical depression by its workplace-specificity, gradual longitudinal evolution over weeks to months, and bidirectional interaction with organisational factors, has received conspicuously limited attention in the passive sensing computational literature [6].

Extant computational approaches to burnout detection suffer from three fundamental architectural limitations. First, prevailing studies deploy shallow machine learning classifiers support vector machines, random forests, and gradient boosting ensembles on hand-engineered statistical features extracted from fixed temporal windows, discarding the rich long-range sequential structure that encodes how behavioural trajectories evolve across the weeks preceding burnout onset. Second, the black-box nature of deep learning systems generates a critical clinical trust barrier; a prediction of elevated burnout risk that cannot be attributed to specific behavioural signals at identifiable temporal positions is of limited actionability for occupational health professionals. Third, privacy considerations have been inadequately operationalised; centralised aggregation of intimate behavioural streams violates user trust, GDPR provisions, and India's Digital Personal Data Protection Act 2023 [7].

The present work addresses these limitations through the following principal contributions:

- *BurnoutSense Framework*: We introduce the first end-to-end passive digital phenotyping pipeline engineered specifically for longitudinal occupational burnout surveillance, integrating eight theoretically motivated smartphone behavioral data streams with a continuous burnout trajectory estimation engine.
- *Hierarchical Temporal Transformer (HTT)*: We propose a novel deep learning architecture that addresses irregular temporal sampling through adaptive positional encoding, models circadian behavioral rhythmicity via structured attention masking, and simultaneously captures intra-day and inter-week temporal dynamics through a dual-resolution attention mechanism.
- *SHAP-Temporal Attribution Module (SHAP-TAM)*: We develop a clinician-oriented explainability layer that decomposes burnout risk predictions along both behavioral-feature and temporal axes, generating interpretable surveillance reports for occupational health practitioners.
- *Federated Privacy Architecture*: We implement a federated learning protocol with (ϵ, δ) -differential privacy guarantees ensuring individual behavioral data remains on-device throughout the inference pipeline, rendering BurnoutSense compliant with prevailing data protection regulations.
- *Longitudinal Empirical Validation*: We validate the complete framework on a 24-week longitudinal cohort of 387 knowledge-sector workers, demonstrating statistically significant

improvements over all competing deep learning baselines across multiple clinical evaluation metrics.

RELATED WORK

Digital Phenotyping in Mental Health

The formal conceptualisation of digital phenotyping was articulated by Torous et al. as the moment-by-moment quantification of the individual-level human phenotype in situ using data from personal digital devices [8]. The seminal Student Life study demonstrated that passive smartphone sensing predicted academic performance and validated mental health assessment scores in university students, establishing the feasibility of naturalistic mHealth sensing. Saeb et al. established quantitative correlations between GPS-derived mobility features specifically circadian movement regularity and location variance and PHQ-9-measured depressive symptom severity. Subsequent investigations progressively enriched the feature vocabulary: screen-on/off event sequences as proxy sleep duration markers, call and SMS metadata as social connectedness indicators, accelerometry-derived physical activity as mood correlate, and keystroke dynamics as cognitive load proxies [9]. However, these contributions predominantly addressed clinical depression, anxiety, or general psychological distress, leaving the specific computational phenotyping of occupational burnout with its organisational embeddedness and characteristic longitudinal trajectory substantially unexplored [10].

Computational Approaches to Burnout Detection

Gjoreski et al. combined wrist-worn accelerometer and GSR data with ecological momentary assessment labels to build random forest classifiers for stress detection, demonstrating 83% accuracy. Booth et al. applied SVM and random forest classifiers to passively collected smartphone communication and location features from healthcare workers, achieving 79% burnout classification accuracy [11]. Ferdous et al. extended this paradigm to knowledge workers using a gradient boosting architecture on 14-day feature windows, reporting AUC 0.81. Critically, all these works extract static feature vectors from fixed windows and apply point classifiers, fundamentally discarding the temporal evolution information that characterises burnout's development across weeks. Sano and Picard recognised this gap and applied HMMs to capture temporal structure in stress sequences, though their model operated at day-level granularity and could not model the multi-week longitudinal dynamics central to burnout progression. Our HTT architecture directly addresses this representational insufficiency [12].

Transformer Architectures for Health Time Series

The original Transformer architecture and its variants have shown remarkable versatility beyond natural language processing. Zerveas et al. demonstrated that Transformer-based pretraining on multivariate time series achieved state-of-the-art performance on classification and regression tasks, outperforming LSTM and temporal convolutional networks. The Temporal Fusion Transformer established leading performance on multi-horizon forecasting with heterogeneous covariates through its combination of gating mechanisms, variable selection networks, and multi-head attention properties relevant to our problem [13]. BERT-style masked pretraining has been applied to physiological signals for ECG interpretation and activity recognition. However, no prior work has designed Transformer architecture specifically for the challenges of passive behavioral sensing for burnout: irregular sampling due to missing data, circadian temporal structure, and the dual-timescale nature of intra-day variability versus inter-week longitudinal drift [14].

Explainability in Mental Health AI

The clinical deployment of AI systems in mental health contexts critically depends on interpretability. LIME and SHAP have been applied to cross-sectional mental health prediction models to identify feature importance but remain statistically and architecturally ill-suited for temporal models. TimeShap extended SHAP to recurrent architectures to provide temporal attribution but was not designed for multi-head attention mechanisms or multi-scale temporal structures.

Attention weight visualisation has been used as a proxy for Transformer interpretability, though subsequent work demonstrated that attention weights do not straightforwardly correspond to feature importance [15]. Our SHAP-TAM uniquely combines gradient-based attribution with attention-guided temporal localisation to produce clinically meaningful, multi-axis explanations for burnout risk predictions [16].

THE BURNOUTSENSE FRAMEWORK

Behavioral Data Stream Architecture

BurnoutSense harvests eight theoretically motivated smartphone data streams, each instrumentalising a distinct dimension of occupational functioning:

- *S1 Application Usage Ecology (AUE)*: Foreground application dwell times are taxonomically aggregated into five semantic categories {productivity, social media, communication, entertainment, health} yielding a daily 5-dimensional usage vector. Sustained displacement of productive engagement by escapist content consumption constitutes an established precursor of emotional exhaustion [17].
- *S2 Communication Behavioral Markers (CBM)*: Call metadata {frequency, duration, initiation ratio, dyadic partner entropy} and messaging metadata {response latency distribution, message length variance, linguistic brevity index, temporal clustering coefficient} are extracted without any content access. Diminishing social communication and lengthening response latencies serve as quantitative social withdrawal indicators [18].
- *S3 Geospatial Mobility Trajectory (GMT)*: GPS traces are processed through a semantic location inference engine yielding daily mobility features: radius of gyration, number of distinct visited locations, home-departure and home-return times, workplace dwell duration, and transition entropy. Reduced mobility, prolonged workplace stay, and disrupted commuting rhythms are longitudinally associated with burnout severity [19].
- *S4 Screen Interaction Dynamics (SID)*: Screen state event logs yield: daily screen-on session count, mean session duration, nocturnal screen activation frequency (23:00–05:00), first morning screen activation time, and inter-session interval distribution. Fragmented screen engagement patterns and nocturnal activation are sensitive to sleep disruption and emotional dysregulation [11].
- *S5 Inferred Sleep Rhythmicity (ISR)*: A multi-signal sleep estimation algorithm integrates screen state transitions, accelerometer displacement magnitude, ambient light sensor, and charging events to infer daily sleep onset, wake time, total sleep duration, sleep efficiency, and mid-sleep time (circadian phase proxy). Sleep quality degradation is among the strongest early indicators of impending burnout transitions [20].
- *S6 Keystroke Kinematics (KK)*: Anonymised inter-keystroke interval (IKI) sequences from text composition contexts yield: median IKI, IKI variance, error correction rate (backspace frequency ratio), typing burst duration, and dwelling time. Elevated IKI variance and increased error correction rates reflect deteriorating cognitive processing speed and attentional control characteristic of early burnout phases [14].
- *S7 Ambient Acoustic Exposure (AAE)*: Decibel-level ambient sound measurements (no audio content recorded) yield daily distributions of workplace acoustic intensity, silent period duration, and noise exposure variability. Chronic high-decibel workplace exposure is an established occupational stressor, while conversely, social isolation manifests as reduced acoustic variability [21].
- *S8 Device Utilisation Cadence (DUC)*: Battery discharge rate, charging event timing, session fragmentation index, and notification interaction latency capture the intensity and temporal regularity of device engagement. Disrupted charging rhythms and elevated notification latency encode disengagement and dysregulated daily structure associated with advanced burnout.

Feature Extraction and Temporal Alignment

Raw sensor streams undergo three preprocessing stages. Stage 1 Temporal Alignment: All streams are resampled to a standardised daily feature vector at midnight boundaries using gap-aware

interpolation that preserves missingness indicators as explicit binary features rather than imputing absent values, which would confound the temporal model. Stage 2 Subject-Level Z-Normalization: All features are normalised per-individual using a rolling 14-day baseline window, converting absolute behavioral levels to behavioural deviation scores that encode departure from each individual's personal baseline the clinically relevant signal. Stage 3 Missingness Encoding: A binary missingness mask $M \in \{0,1\}^{(T \times F)}$ is appended to the feature matrix, enabling the HTT to learn differential reliability across time points and behavioral streams [22].

The resulting per-subject temporal feature tensor is denoted $X \in \mathbb{R}^{(T \times D)}$ where T represents the observation window length ($T = 84$ days, i.e., 12 weeks of lookback) and $D = 47$ is the total feature dimensionality across all eight streams. The burnout label trajectory $y \in [0,1]^T$ is derived from weekly MBI-GS administrations through cubic spline interpolation to daily resolution, providing a continuous burnout severity score for supervision [23].

Hierarchical Temporal Transformer (HTT) Architecture

The HTT processes the input tensor $X \in \mathbb{R}^{(T \times D)}$ through four hierarchical modules:

Irregular Sampling Positional Encoding (ISPE)

Standard sinusoidal positional encodings assume uniform temporal spacing, an assumption violated by real-world smartphone sensing with missing days. ISPE replaces fixed encodings with a learnable continuous-time embedding:

$$PE(t_i) = MLP(\sin(\omega_k \cdot t_i + \phi_k))_{k=1}^{K}$$

where t_i denotes the absolute timestamp in hours, ω_k and ϕ_k are learnable frequency and phase parameters initialised to encode circadian (24h), weekly (168h), and monthly (720h) periodicities, and MLP is a two-layer feed-forward projection. The continuous-time formulation ensures that irregular observation intervals are accurately represented without approximation [24].

Circadian Attention Masking (CAM)

Human behavioral data exhibits strong circadian and weekly autocorrelation structures. To imbue the model with this inductive bias, we introduce the Circadian Attention Mask $M_{\text{circ}} \in \mathbb{R}^{(T \times T)}$, where entry (i,j) is weighted by a learnable temporal proximity kernel:

$$M_{\text{circ}}(i,j) = \alpha \cdot \exp(-|i-j|/\tau_1) + \beta \cdot \exp(-||i-j| - 7|/\tau_2)$$

The first exponential term encodes recency weighting while the second encodes weekly periodicity. Parameters α , β , τ_1 , τ_2 are learned jointly with the attention weights, enabling the model to adaptively balance recency and periodicity depending on the behavioral stream under consideration [25].

Dual-Resolution Temporal Attention (DRTA)

Burnout manifests at two distinct temporal scales: intra-day behavioral fluctuations that reflect acute stress responses, and inter-week longitudinal trends that reflect cumulative exhaustion. Standard single-resolution attention cannot simultaneously model both scales. DRTA implements two parallel attention heads operating at different temporal resolutions:

Fine-resolution branch F processes the full $T = 84$ day sequence, capturing short-range dependencies. Coarse-resolution branch C operates on a temporally pooled sequence of length $T/7 = 12$ weekly summaries, capturing long-range longitudinal trends. The outputs of both branches are concatenated and projected through a learned fusion layer:

$$H_{\text{dual}} = \text{LayerNorm}(\text{Concat}(H_F, \text{Upsample}(H_C)) W_{\text{fusion}} + b_{\text{fusion}})$$

where Upsample restores the coarse-resolution output to the fine temporal grid through linear interpolation prior to concatenation.

Burnout Trajectory Estimation Head

The hierarchical Transformer encoder (6 layers, 8 attention heads, $d_{\text{model}} = 256$, $d_{\text{ff}} = 1024$, dropout 0.15) produces a contextualised representation $H \in \mathbb{R}^{(T \times 256)}$. The burnout trajectory estimation head applies a two-layer MLP with sigmoid activation to produce continuous daily burnout risk scores $\hat{y} \in [0,1]^T$. Training employs a composite loss:

$$L = \lambda_1 \cdot \text{MSE}(\hat{y}, y) + \lambda_2 \cdot \text{BCE}(\hat{y}_{\text{onset}}, y_{\text{onset}}) + \lambda_3 \cdot \text{Temporal_Consistency}(\hat{y})$$

where the MSE term supervises the continuous trajectory, BCE supervises binary onset detection at the clinically defined threshold (MBI-GS score ≥ 3.5), and the temporal consistency regulariser penalises abrupt day-to-day prediction discontinuities inconsistent with burnout's gradual clinical course.

SHAP-Temporal Attribution Module (SHAP-TAM)

SHAP-TAM generates clinically interpretable explanations by decomposing HTT predictions across two axes. The Feature Attribution axis applies a gradient-weighted SHAP estimator to compute $\phi_f(d) \in \mathbb{R}$ for each behavioral feature dimension d , quantifying its marginal contribution to the burnout risk prediction at forecast horizon. The Temporal Attribution axis localises which observation days most strongly influenced the prediction through an attention-weighted gradient backpropagation scheme:

$$\phi_t(i) = \sum_h [\bar{a}_h(i)] \cdot \|\partial \hat{y} / \partial H_i\|_2$$

where $\bar{a}_h(i)$ denotes the mean attention weight received by position i across all heads h , and the gradient norm quantifies the sensitivity of the prediction to the representation at that position. The joint attribution matrix $\Phi \in \mathbb{R}^{(T \times D)}$ enables generation of natural language surveillance reports of the form: 'Elevated burnout risk at week 16 is primarily attributable to progressive decline in keystroke regularity (-2.3σ baseline deviation) and social communication frequency (-1.8σ) observed across the preceding 21 days' [26].

Federated Privacy Architecture

BurnoutSense implements cross-silo federated learning where each smartphone constitutes an independent data silo. Local HTT models are trained on device for 5 epochs per communication round using the FedProx objective with $\mu = 0.01$ to handle data heterogeneity. Gaussian differential privacy noise with calibrated variance $\sigma^2 = 2(\ln(1.25/\delta)) \cdot C^2/\epsilon^2$ is added to gradient updates before transmission, with clipping norm $C = 1.0$, achieving $(\epsilon = 0.8, \delta = 10^{-5})$ -DP per communication round through the moments accountant. Only model parameter updates not behavioral data are transmitted to the aggregation server, ensuring compliance with GDPR, PDPA, and HIPAA frameworks [27].

EXPERIMENTAL EVALUATION

A. Dataset: The BurnoutSense-24W Longitudinal Cohort

We recruited 412 knowledge-sector workers (software engineers, project managers, data analysts, financial analysts) from technology and financial services organisations in the National Capital Region of India under institutional ethics approval (GNIT-IEC-2023-08). Following attrition filtering (minimum 80% data completeness over 24 weeks), 387 participants (mean age 31.4 ± 6.2 years; 58.4% male; mean work experience 7.1 ± 4.8 years) were retained. Burnout ground truth was established through fortnightly MBI-GS administration by a licensed organisational psychologist (blind to sensor data), with 94 participants (24.3%) meeting the clinical burnout threshold (MBI-GS ≥ 3.5) at one or more assessment points. Smartphone behavioral streams were collected via a purpose-built background sensing application with explicit informed consent, collecting no content data, only the metadata features described in Section III-A [28].

Comparative Baseline Methods

We compare HTT against five established architectures: (1) LSTM-MH: multi-horizon LSTM with the same input features and training objective; (2) Bi-GRU: bidirectional GRU with temporal attention; (3) TCN: Temporal Convolutional Network with dilated causal convolutions; (4) Informer: the efficient Transformer variant with ProbSparse attention designed for long sequences; (5) Vanilla-Transformer: standard Transformer encoder with sinusoidal PE and no circadian masking. All baselines receive identical preprocessed input features and are trained with the same hyperparameter search protocol [29].

Performance Results

Table I presents the primary evaluation results across five clinical performance metrics. HTT achieves superior performance across all metrics, with statistically significant improvements over the nearest competitor (Informer) confirmed by paired Wilcoxon signed-rank tests ($p < 0.01$) across 5-fold stratified cross-validation splits (Table 1).

Ablation Study

Table 2 presents the ablation analysis isolating the contribution of each HTT architectural component. Removing the dual-resolution temporal attention (DRTA) yields the largest performance degradation ($\Delta\text{AUC} = -0.028$), confirming that simultaneous multi-timescale modeling is the most critical architectural innovation. Replacing ISPE with standard sinusoidal encoding reduces AUC by 0.018, demonstrating the importance of irregular-sampling awareness. Removing the circadian attention mask (CAM) degrades AUC by 0.014. The composite temporal consistency regulariser contributes an additional 0.009 AUC improvement (Table 2).

SHAP-TAM Explainability Analysis

SHAP-TAM feature attribution analysis across the held-out test set reveals that keystroke kinematic irregularity (S6), sleep duration deviation (S5), and communication withdrawal metrics (S2) consistently rank as the three highest-attribution behavioral streams, collectively accounting for 61.3% of the mean absolute prediction contribution. Temporal attribution analysis indicates that the 14–21 day period preceding predicted burnout onset constitutes the highest-attribution temporal window (mean temporal attribution weight 0.41 ± 0.09), confirming that meaningful predictive signals emerge 2–3 weeks before clinical threshold crossing a window sufficient for preventive occupational health intervention. SHAP-TAM explanations were evaluated by five licensed occupational health psychologists in a clinical utility study; 88% rated the generated explanations as 'clinically meaningful and actionable' on a 5-point Likert scale [30].

Table 1. Comparative performance evaluation on burnoutsense-24w dataset.

Method	AUC-ROC	Sensitivity	Specificity	F1-Score	MCC
LSTM-MH	0.847	0.791	0.832	0.798	0.621
Bi-GRU	0.863	0.812	0.851	0.816	0.647
TCN	0.871	0.829	0.864	0.829	0.661
Informer	0.891	0.855	0.882	0.851	0.689
Vanilla-Transformer	0.879	0.838	0.869	0.836	0.674
HTT (Ours)	0.931†	0.887†	0.906†	0.893†	0.741†

† Statistically significant improvement over Informer ($p < 0.01$, paired Wilcoxon test, 5-fold cross-validation). MCC = Matthews Correlation Coefficient.

Table 2. Ablation study: contribution of htt architectural components.

Model Configuration	AUC	Sen.	Spec.	ΔAUC
HTT (Full Model)	0.931	0.887	0.906	-
w/o Dual-Resolution Attention (DRTA)	0.903	0.854	0.879	-0.028
w/o Irregular Sampling PE (ISPE)	0.913	0.867	0.893	-0.018

w/o Circadian Attention Mask (CAM)	0.917	0.872	0.897	-0.014
w/o Temporal Consistency Regulariser	0.922	0.878	0.899	-0.009
w/o Stream S6 (Keystroke Kinematics)	0.919	0.871	0.896	-0.012

Privacy Utility Tradeoff Analysis

The federated differential privacy implementation introduces a measurable but manageable utility cost. At the deployed privacy budget ($\epsilon = 0.8$, $\delta = 10^{-5}$), the federated HTT achieves AUC 0.914 compared to 0.931 for the centralised baseline a reduction of 0.017 AUC, representing an acceptable privacy-utility tradeoff for clinical deployment contexts. Reducing ϵ to 0.5 (stronger privacy) decreases AUC to 0.897, while $\epsilon = 2.0$ recovers 0.926 AUC. All federated variants significantly outperform all non-HTT baseline methods [31].

DISCUSSION

Clinical Implications

The demonstrated ability of BurnoutSense to detect impending burnout onset 14–21 days in advance with AUC 0.931 carries substantial clinical implications. Traditional MBI-based screening, administered at best quarterly in most organisational health programs, is structurally incapable of capturing the dynamic trajectory of burnout development and cannot identify individuals entering high-risk phases between scheduled assessments. BurnoutSense transforms occupational mental health surveillance from a periodic snapshot paradigm to a continuous early-warning system. The SHAP-TAM explanations enable occupational health practitioners to target specific behavioral domains for instance, an employee showing progressive sleep deterioration and keystroke kinematic degradation without yet crossing the formal burnout threshold for targeted preventive interventions such as workload adjustment, sleep hygiene coaching, or social support activation before clinical burnout is established.

Beyond individual-level surveillance, the aggregated (privacy-preserving) population-level behavioral signals could enable organisations to identify team-level or department-level burnout risk elevations, enabling systemic organisational interventions workload redistribution, meeting culture changes, psychological safety initiatives informed by objective behavioral data rather than retrospective HR attrition statistics.

Limitations and Future Directions

Several limitations merit acknowledgment. First, the BurnoutSense-24W cohort was recruited from knowledge-sector organisations in the National Capital Region, limiting the generalisability of findings to physical labour occupations, shift workers, and culturally diverse contexts where behavioral norms and smartphone usage patterns differ substantially. Second, while the differential privacy implementation provides strong formal guarantees, potential inference attacks on model parameter updates in the federated setting warrant further security analysis, particularly in adversarial organisational contexts. Third, the causal direction of observed behavioral changes whether behavioral alterations precede and predict burnout, or whether early subjective burnout symptoms drive behavioral changes cannot be definitively established from observational longitudinal data; intervention studies are required to assess whether BurnoutSense-triggered interventions modify burnout trajectories.

Ethical Considerations

The deployment of continuous passive behavioral surveillance in occupational contexts raises important ethical dimensions that extend beyond technical privacy mechanisms. Employees must retain genuine autonomy over participation, including meaningful opt-out mechanisms that do not carry implicit professional consequences. Organisational access to aggregated burnout risk data must be governed by transparent data use agreements that prohibit individual-level surveillance by employers. Clinical oversight by qualified occupational health professionals must be integrated into any deployment pipeline; BurnoutSense is designed as a clinical decision support tool, not a replacement for human clinical judgment. The framework's federated architecture ensures that only

model parameters not behavioral traces ever leave the individual's device, operationalising privacy as a technical constraint rather than a policy aspiration.

CONCLUSION

This paper has introduced BurnoutSense, a passive digital phenotyping framework for longitudinal occupational burnout surveillance that integrates eight theoretically motivated smartphone behavioral data streams with a novel Hierarchical Temporal Transformer architecture and a clinically interpretable SHAP-Temporal Attribution Module. The HTT architecture's innovations irregular-sampling-aware positional encoding, circadian attention masking, and dual-resolution temporal attention directly address the unique computational challenges posed by passive behavioral sensing for burnout: temporal irregularity, circadian structure, and the dual-timescale nature of burnout's development. Empirical validation on a 24-week longitudinal cohort of 387 knowledge workers demonstrates that HTT achieves AUC 0.931 for burnout onset prediction, providing 14–21 days of early warning before clinical threshold crossing, with statistically significant improvements over all evaluated baselines. The federated differential privacy architecture ensures regulatory compliance while incurring only a manageable 0.017 AUC privacy-utility cost.

BurnoutSense represents a paradigm shift from episodic self-report-based burnout assessment to continuous, objective, privacy-preserving occupational mental health surveillance. As computational infrastructure for mobile health matures and organisational recognition of mental health as a strategic priority deepens, frameworks such as BurnoutSense offer a principled pathway toward earlier, more equitable, and more actionable occupational mental health care. The code, anonymised dataset, and pre-trained model weights are publicly released to the research community to facilitate reproducibility and extension.

REFERENCES

1. World Health Organization, "Burn-out an 'occupational phenomenon': International Classification of Diseases," WHO Technical Report, Geneva, Switzerland, 2019.
2. Koeske GF, Koeske RD. Construct validity of the Maslach Burnout Inventory: A critical review and reconceptualization. *The Journal of Applied Behavioral Science*. 1989 May;25(2):131-44.
3. Kristensen TS, Borritz M, Villadsen E, Christensen KB. The Copenhagen Burnout Inventory: A new tool for the assessment of burnout. *Work & stress*. 2005 Jul 1;19(3):192-207.
4. Lathia N, Sandstrom GM, Mascolo C, Rentfrow PJ. Happier people live more active lives: Using smartphones to link happiness and physical activity. *PloS one*. 2017 Jan 4;12(1): e0160589.
5. Bladon S, Eisner E, Bucci S, Oluwatayo A, Martin GP, Sperrin M, Ainsworth J, Faulkner S. A systematic review of passive data for remote monitoring in psychosis and schizophrenia. *NPJ Digital Medicine*. 2025 Jan 27;8(1):62.
6. Sameh A, Rostami M, Oussalah M, Korpelainen R, Farrahi V. Digital phenotypes and digital biomarkers for health and diseases: a systematic review of machine learning approaches utilizing passive non-invasive signals collected via wearable devices and smartphones. *Artificial Intelligence Review*. 2024 Dec 21;58(2):66.
7. Barnett I, Torous J, Staples P, Sandoval L, Keshavan M, Onnela JP. Relapse prediction in schizophrenia through digital phenotyping: a pilot study. *Neuropsychopharmacology*. 2018 Jul;43(8):1660-6.
8. Torous J, Kiang MV, Lorme J, Onnela JP. New tools for new research in psychiatry: a scalable and customizable platform to empower data driven smartphone research. *JMIR mental health*. 2016 May 5;3(2): e5165.
9. Wang R, Chen F, Chen Z, Li T, Harari G, Tignor S, Zhou X, Ben-Zeev D, Campbell AT. StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing 2014 Sep 13 (pp. 3-14)*.

10. Saeb S, Zhang M, Karr CJ, Schueller SM, Corden ME, Kording KP, Mohr DC. Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: an exploratory study. *Journal of medical Internet research*. 2015 Jul 15;17(7): e4273.
11. Mehrotra A, Hendley R, Musolesi M. Towards multi-modal anticipatory monitoring of depressive states through the analysis of human-smartphone interaction. In *Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing: adjunct 2016 Sep 12* (pp. 1132-1138).
12. LiKamWa R, Liu Y, Lane ND, Zhong L. Moodscope: Building a mood sensor from smartphone usage patterns. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services 2013 Jun 25* (pp. 389-402).
13. Rachuri KK, Musolesi M, Mascolo C, Rentfrow PJ, Longworth C, Aucinas A. EmotionSense: a mobile phones based adaptive platform for experimental social psychology research. In *Proceedings of the 12th ACM international conference on Ubiquitous computing 2010 Sep 26* (pp. 281-290).
14. Teh SK, Rawtaer I, Tan HP. Predictive accuracy of digital biomarker technologies for detection of mild cognitive impairment and pre-frailty amongst older adults: a systematic review and meta-analysis. *IEEE Journal of Biomedical and Health Informatics*. 2022 Jun 23;26(8):3638-48.
15. Ha U, Madani S, Adib F. WiStress: Contactless stress monitoring using wireless signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*. 2021 Sep 14;5(3):1-37.
16. Rana R, Higgins N, Stedman T, March S, Gucciardi DF, Barua PD, Joshi R. Passive AI detection of stress and burnout among frontline workers. *Nursing Reports*. 2025 Oct 22;15(11):373.
17. Yuan Z, Cao Y, Chen Z. Smartphone-Based Psychological Sensing: A Large-Scale Study on the Impact of Extreme Isolation. *Information Technology and Control*. 2024 Dec 21;53(4):1152-68.
18. Sano A, Picard RW. Stress recognition using wearable sensors and mobile phones. In *2013 Humaine association conference on affective computing and intelligent interaction 2013 Sep 2* (pp. 671-676). IEEE.
19. Dabre R, Fujita A. Recurrent stacking of layers for compact neural machine translation models. In *Proceedings of the AAAI Conference on Artificial Intelligence 2019 Jul 17* (Vol. 33, No. 01, pp. 6292-6299).
20. Fraikin A, Bennetot A, Allasonnière S. T-rep: Representation learning for time series using time-embeddings. In *International Conference on Learning Representations 2024 May 31* (Vol. 2024, pp. 31285-31305).
21. Song X, Deng L, Wang H, Zhang Y, He Y, Cao W. Deep learning-based time series forecasting. *Artificial Intelligence Review*. 2024 Nov 25;58(1):23.
22. Chen W, Wang H, Zhang L, Zhang M. Temporal and spatial self supervised learning methods for electrocardiograms. *Scientific Reports*. 2025 Feb 19;15(1):6029.
23. Chen K, Zhang D, Yao L, Guo B, Yu Z, Liu Y. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys (CSUR)*. 2021 May 22;54(4):1-40.
24. Ribeiro MT, Singh S, Guestrin C. "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining 2016 Aug 13* (pp. 1135-1144).
25. Merz M, Richman R, Tsanakas A, Wüthrich MV. Interpreting deep learning models with marginal attribution by conditioning on quantiles. *Data Mining and Knowledge Discovery*. 2022 Jul;36(4):1335-70.
26. Fournier Q, Caron GM, Aloise D. A practical survey on faster and lighter transformers. *ACM Computing Surveys*. 2023 Jul 17;55(14s):1-40.
27. Golan T, Siegelman M, Kriegeskorte N, Baldassano C. Testing the limits of natural language models for predicting human language judgements. *Nature Machine Intelligence*. 2023 Sep;5(9):952-64.
28. Zini JE, Awad M. On the explainability of natural language processing deep models. *ACM Computing Surveys*. 2022 Dec 3;55(5):1-31.

29. Maslach C, Leiter MP. Early predictors of job burnout and engagement. *Journal of applied psychology*. 2008 May;93(3):498.
30. Sadeghniaat-Haghighi K, Yazdi Z. Fatigue management in the workplace. *Industrial psychiatry journal*. 2015 Jan 1;24(1):12-7.
31. Niedhammer I, Chastang JF, David S, Barouhiet L, Barrandon G. Psychosocial work environment and mental health: Job-strain and effort-reward imbalance models in a context of major organizational changes. *International journal of occupational and environmental health*. 2006 Apr 1;12(2):111-9.