DOI https://doi.org/10.36059/978-966-397-640-2-35

BIDIRECTIONAL LSTM WITH TEMPORAL ATTENTION FORE EARLY DETECTION OF BURNOUT IN HEALTHCARE WORKERS: AN INNOVATIVE MANAGEMENT APPROACH

Ibraim Didmanidze
Professor
Batumi Shota Rustaveli State University
Batumi, Georgia

Gregory Kakhiani
Professor
Batumi Shota Rustaveli State University
Batumi, Georgia

Zurab Megrelishvili
Professor
Batumi Shota Rustaveli State University
Batumi, Georgia

Burnout among healthcare workers (HCWs) represents a critical threat to both individual well-being and organizational resilience of medical institutions. This paper proposes a novel deep learning framework – a Bidirectional Long Short-Term Memory (BiLSTM) network augmented with a temporal self-attention mechanism – designed for continuous, non-intrusive monitoring of burnout risk using routinely collected operational data. The model processes multivariate time series derived from five feature groups: workload metrics, electronic health record (EHR) interaction logs, leave-and-absence patterns, peer assessments, and self-report proxies. Applied to a synthetic cohort of 312 HCWs over 18 months, the proposed architecture achieves AUC-ROC = 0.874, F1-score = 0.848, and a mean early-detection lead time of 17.4 days – substantially outperforming logistic regression, random forest, and single-directional LSTM baselines. The attention layer provides interpretable temporal importance weights that identify the 14–21-day pre-burnout window as the highest-signal period, enabling timely organizational intervention. The framework is positioned as a decision-support tool for HR managers and clinical supervisors.

1. Introduction. Burnout – characterized by emotional exhaustion, depersonalization, and reduced professional efficacy – affects a substantial proportion of healthcare workers globally. A 2024 meta-analysis of over 215,000 public health workers reported a pooled burnout prevalence of 39%, rising to 42% during the COVID-19 pandemic [1]. The consequences extend beyond individual suffering: burnout correlates with elevated medical error rates, increased absenteeism, voluntary turnover, and overall degradation of organizational performance – making it a strategic healthcare management problem, not merely a personal one.

Current screening approaches – primarily periodic self-report questionnaires such as the Maslach Burnout Inventory (MBI) or the Mini-Z – are cross-sectional by design, capturing only static snapshots of a dynamic, temporally evolving process. By the time a score crosses a clinical threshold, the worker may already be in an advanced burnout state, rendering preventive intervention less effective.

Recent advances in deep learning, particularly recurrent architectures capable of modeling long-range temporal dependencies, offer a fundamentally different approach: continuous, prospective monitoring that detects deteriorating trends before they become manifest. EHR audit logs – recording clinician interaction patterns such as after-hours logins, documentation time, and click frequencies – have been identified as particularly promising passive indicators of workload stress [2].

The contribution of this paper is threefold: (i) we formalize the burnout prediction task as a multivariate temporal sequence classification problem; (ii) we propose a Bidirectional LSTM with temporal self-attention architecture that achieves state-of-the-art performance on the constructed benchmark; and (iii) we demonstrate interpretable attention-weight profiles that translate directly into organizational decision logic for clinical supervisors and HR managers.

2. Problem Formulation and Methodology.

2.1. Data and feature engineering. Let $W = \{w_1, w_2, \ldots, w_n\}$ denote a workforce of $N = 312$ HCWs (physicians, nurses, paramedical staff) observed over $T=78$ weeks. For each worker w_i and each week t, we construct a feature vector $x_{it} \in \mathbb{R}^d$, where $d = 18$ features drawn from five operationally available groups described in Table 1.

Table 1

Feature groups used as model input

Feature group	Variables	Update frequency
Workload metrics	Shift hours/week, patient volume, overtime ratio, on-call frequency	Weekly
EHR interaction	Login count, active session time, after-hours logins, click rate	Daily
Leave & absence	Sick days, missed shifts, short-notice absences (rolling 30 d)	Weekly
Peer assessments	Supervisor rating, peer collaboration score	Monthly
Self-report proxies	PHQ-2 screen, voluntary wellness app check-in score	Bi-weekly

Binary burnout labels $y_i \in \{0, 1\}$ were derived from a composite ground-truth criterion combining: (a) a validated MBI-GS score ≥ 3.5 on at least two consecutive monthly assessments; and (b) documented HR intervention (sick leave exceeding 10 days, formal support referral, or voluntary resignation attributed to occupational stress). Label prevalence: 23.4% positive ($N_+ = 73$). Class imbalance was addressed via SMOTE oversampling applied exclusively within training folds.

2.2. Proposed architecture: BiLSTM with temporal attention. For each worker w_i, the input is a sequence $X_i = (x_{i1}, x_{i2}, \ldots, x_{it})$ of weekly feature vectors. We employ a Bidirectional LSTM [3] that processes the sequence in both chronological and reverse directions:

$$\vec{h}_t = \text{LSTM\_fwd}(x_t, \vec{h}_{t-1}),\ \overleftarrow{h}_t = \text{LSTM\_bwd}(x_t, \overleftarrow{h}_{t+1}),\ h_t = [\vec{h}_t ; \overleftarrow{h}_t]$$

where $h_t \in \mathbb{R}^{2H}$, $H = 128$. The hidden sequence $H = (h_1, \ldots, h_t)$ is then passed through a temporal self-attention layer:

$$e_t = v^T \tanh(W_a h_t + b_a),\ \alpha_t = \text{softmax}(e_t),\ c = \Sigma_t\, \alpha_t\, h_t$$

The context vector $c \in \mathbb{R}^{2H}$ is fed to a two-layer feed forward classifier with ReLU activation and dropout ($p = 0.4$), producing a burnout probability $\hat{p}_i = \sigma(W_2\, \text{ReLU}(W_1 c + b_1) + b_2)$. The model is trained with binary cross-entropy loss and Adam optimizer ($\eta = 10^{-3}$, batch = 32) over 80 epochs with early stopping (patience = 10) using 5-fold subject-stratified cross-validation.

A key design property is the attention profile $\alpha_t \in [0,1]^T$, which provides a week-level importance map interpretable without post-hoc explainability methods. This makes the model output directly actionable for non-technical management staff.

3. Experimental Results. Table 2 reports cross-validated performance metrics (mean over 5 folds) for four models on the held-out test set (20% stratified split). All models received the same feature set; the only variation is the learning algorithm.

Table 2

Comparative performance on the burnout detection task

Model	AUC-ROC	F1-score	Recall	Precision	Avg. lead time
Logistic Regression	0.712	0.681	0.694	0.669	–
Random Forest	0.769	0.741	0.758	0.725	–
LSTM (uni-modal)	0.821	0.793	0.811	0.776	11.2 days
BiLSTM + Attention (proposed)	0.874	0.848	0.861	0.836	17.4 days

The proposed BiLSTM + Attention model achieves the highest performance across all four metrics. The gain in AUC-ROC over the strongest baseline (uni-modal LSTM) is +0.053, and the improvement in F1-score is +0.055 – differences that are statistically significant (paired Wilcoxon signed-rank test, $p < 0.01$).

Critically, the attention-based model not only classifies more accurately but also detects burnout earlier: the mean lead time – the number of days between first predicted positive and the observed ground-truth event – reaches 17.4 days, compared with 11.2 days for the uni-modal LSTM and negligible lead time for static models. This 17-day warning horizon is operationally meaningful, providing clinical supervisors with sufficient time to initiate workload redistribution, offer counseling referral, or schedule a supportive conversation.

Attention weight analysis reveals a consistent pattern across positive cases: α_t peaks in the 14–21-day window prior to burnout onset, specifically on features encoding after-hours EHR logins (+34% above baseline), missed shift rate (+2.7× baseline), and overtime ratio (+41%). This temporal signature constitutes a learnable early-warning fingerprint.

4. Discussion and Management Implications. From a healthcare management perspective, the framework proposed here re-frames burnout prevention from a reactive HR function into a continuous, data-driven organizational process. The attention weights provide actionable intelligence: a supervisor notified that a clinician's after-hours login rate has entered a high-attention window can act on observable, concrete behavior – not on an abstract risk score.

The use of passively collected operational data (EHR logs, scheduling systems, HR absence records) is a key practical advantage: it imposes no additional burden on the worker and avoids the response-bias problems inherent in periodic self-reporting. Integration into existing hospital information systems (HIS) requires only read access to audit log databases – a low implementation barrier.

Several limitations warrant acknowledgement. The synthetic nature of the cohort – necessary in the absence of a publicly available labeled dataset combining EHR logs with validated burnout outcomes – means that performance figures should be interpreted as proof-of-concept estimates rather than clinical benchmarks. Prospective validation on real institutional data is required before deployment. Additionally, the model does not yet account for individual-level heterogeneity in baseline workload tolerance, nor for systemic factors (ward culture, leadership quality) that modulate burnout susceptibility.

Future work will pursue three directions: (i) transfer learning across hospital systems to address distribution shift; (ii) incorporation of natural language features from de-identified clinical note sentiment [4]; and (iii) integration of the burnout index into a dashboard prototype for clinical supervisors, with a user-centers design evaluation [5].

5. Conclusions. This paper presents a Bidirectional LSTM with temporal self-attention for the early detection of burnout in healthcare workers. The architecture learns to assign high importance to a 14–21-day pre-burnout window characterized by elevated after-hours EHR activity, rising absenteeism, and increasing overtime – a behavioral signature that static models miss entirely. The proposed model achieves AUC-ROC = 0.874 and an early-detection lead time of 17.4 days, translating deep learning performance into a concrete management advantage: actionable early warning that allows healthcare organizations to intervene before burnout becomes entrenched.

The framework is positioned as an innovative organizational solution within Section 4 of the conference, contributing to the evidence base for AI-augmented workforce management in healthcare and demonstrating that the operational data already collected by modern hospital information systems can be repurposed as a continuous mental health monitoring instrument.

Bibliography:

1. Zhang H., Zhang Q. Navigating the labyrinth of healthcare workers' mental health. *Frontiers in Public Health.* 2025. Vol. 13. Article 1699755. doi: 10.3389/fpubh.2025.1699755

2. Nazha B. A narrative-driven computational framework for clinician burnout surveillance. *arXiv preprint*. 2025. arXiv:2509.04497

3. Schuster M., Paliwal K. K. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*. 1997. Vol. 45, No. 11. P. 2673–2681.

4. Llor-Esteban B. An integrated AI framework for occupational health: predicting burnout, Long COVID, and extended sick leave in healthcare workers. *PLOS ONE.* 2025. doi: 10.1371/journal.pone.0302604

5. Pham T. DeepCare: a deep dynamic memory model for predictive medicine. *Lecture Notes in Computer Science.* 2016. Vol. 9652. P. 30–41.