IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AI-POWERED COGNITIVE FATIGUE DETECTION USING KEYSTROKE DYNAMICS

Nitin Aadithya S
Student
Department of Data Science
Kumaraguru College of Liberal Arts and Science, Coimbatore, India

Abstract: Cognitive fatigue is a significant psychological and physiological condition that impairs an individual's ability to concentrate, process information, and perform tasks effectively. In modern digital environments, prolonged computer usage across professional and academic settings has substantially increased mental workload, leading to rising concerns about undetected fatigue. Traditional detection methods rely on intrusive physiological sensors such as electroencephalography (EEG) devices, eye-tracking systems, and heart rate monitors. These approaches are expensive, uncomfortable, and unsuitable for continuous real-world deployment. This paper proposes an AI-Powered Cognitive Fatigue Detection System using Keystroke Dynamics as a non-intrusive behavioral biometric. Keystroke behavioral features including typing speed, key hold time, inter-key latency, error frequency, and pause duration are analyzed using machine learning classifiers: Logistic Regression, Random Forest, Support Vector Machine, and an Artificial Neural Network. The weighted ensemble achieves a binary classification accuracy of 92.0%, with the ANN achieving 91.3% individually. The system operates with total pipeline latency under 40 milliseconds and CPU utilization below 3%, making it suitable for deployment as a lightweight background monitoring application. A Streamlit-based interactive dashboard provides real-time fatigue scores, session analytics, and evidence-based cognitive reactivation recommendations.

Index Terms — Cognitive Fatigue, Keystroke Dynamics, Behavioral Biometrics, Machine Learning, Artificial Neural Network, Human-Computer Interaction, Real-Time Monitoring, Streamlit.

I. INTRODUCTION

The rapid advancement of digital technologies has profoundly transformed how individuals engage with computers across professional, academic, and personal contexts. Modern workplaces and educational institutions require sustained cognitive engagement over extended hours, often without adequate breaks or recovery periods. This sustained mental demand leads to a condition known as cognitive fatigue, a progressive neuropsychological state characterized by declining attention, slower information processing, degraded working memory, and impaired decision-making.

Unlike physical fatigue, cognitive fatigue is often invisible and subjectively underreported. Individuals frequently continue working in a fatigued state without being fully aware of how significantly their performance has deteriorated. The economic and safety consequences are substantial. Studies confirm that cognitive fatigue is associated with increased medical errors, traffic accidents, software defects, and occupational hazards across multiple industries.

Keystroke dynamics offers a compelling non-intrusive solution. Every individual types with a distinctive rhythm shaped by their motor habits, cognitive processing speed, and attentional state. This rhythm changes measurably as cognitive fatigue accumulates. The universal presence of keyboards in digital work environments, combined with the entirely passive nature of keystroke monitoring, positions keystroke dynamics as an ideal behavioral sensor for cognitive state assessment.

This paper proposes an AI-Powered Cognitive Fatigue Detection System that leverages keystroke dynamics as a behavioral biometric input, applies machine learning classification, and delivers real-time fatigue alerts and interactive analytics through a Streamlit dashboard. The proposed framework aligns with global sustainability objectives, particularly SDG 3 (Good Health and Well-being), SDG 4 (Quality Education), and SDG 8 (Decent Work and Economic Growth).

II. LITERATURE REVIEW

A. Cognitive Fatigue and Neuropsychological Basis

Van Dongen et al. [1] demonstrated that even moderate sleep restriction produces performance deficits equivalent to total sleep deprivation, while subjects subjectively reported only mild sleepiness, underscoring the need for objective behavioral monitoring. Boksem and Tops [2] proposed that cognitive fatigue results from progressive depletion of cognitive control resources mediated by frontal lobe mechanisms, manifesting as reduced consistency in motor outputs directly observable in keystroke behavioral metrics.

B. Keystroke Dynamics as a Behavioral Biometric

The systematic study of keystroke dynamics as a behavioral biometric was advanced by Monrose and Rubin [3] through statistical models for keystroke-based user authentication. The CMU Keystroke Dynamics Benchmark Dataset by Killourhy and Maxion [4] established key hold time, inter-key latency, and n-graph timings as primary discriminative features. Bergadano et al. [5] demonstrated that behavioral biometric patterns persist across arbitrary text content, which is critical for cognitive monitoring applications.

C. Keystroke Dynamics and Cognitive Load

Vizer et al. [6] showed that under cognitive stress, typing speed decreased, error rates increased, and inter-key interval variability increased significantly. Epp et al. [7] at CHI 2011 demonstrated that fatigue and 11 other emotional states produced statistically significant classifiable differences in typing behavior using SVMs, achieving classification accuracies of 72% to 84%. Ghosh et al. [8] confirmed feasibility of near-real-time cognitive assessment within 30-second typing windows.

D. Research Gap

The literature reveals that most keystroke dynamics studies focus on user authentication rather than cognitive state monitoring. Existing fatigue detection systems rely primarily on physiological hardware. Few systems integrate keystroke-based behavioral monitoring with machine learning classification, personalized baseline normalization, deviation analysis, and real-time interactive dashboards within a single unified deployable platform. The proposed system addresses all these gaps.

III. SYSTEM DESIGN AND ARCHITECTURE

The proposed AI-Powered Cognitive Fatigue Detection System is designed as a hybrid, non-intrusive cognitive state monitoring framework. Unlike traditional physiological monitoring systems, it requires no specialized hardware beyond the standard keyboard. The architecture consists of five major functional components.

A. Keystroke Data Acquisition Module

This module runs as a background daemon thread and captures keyboard events including key press and key release with millisecond-precision timestamps using Python's pynput library. Only temporal patterns are recorded, not the textual content of what is typed, ensuring complete privacy. Keystroke data is buffered in an in-memory circular deque with configurable maximum size.

B. Feature Extraction Module

Raw keystroke event streams are processed in sliding 60-second time windows with 30-second overlap to extract a 15-dimensional behavioral feature vector. Features include: mean and standard deviation of key hold time, mean and standard deviation of inter-key latency, typing speed in characters per minute, error rate, pause frequency and duration statistics, burst typing metrics, coefficient of variation measures, bigram latency statistics, long pause ratio, and intra-session typing speed trend slope.

C. Machine Learning Classification Module

Pre-trained classification models including Logistic Regression, Random Forest, Support Vector Machine, and an Artificial Neural Network receive normalized feature vectors and produce fatigue state predictions. A weighted ensemble combining predictions from all four models improves overall robustness. Models are trained offline using labeled keystroke datasets with Karolinska Sleepiness Scale (KSS) ground truth, then serialized for real-time inference.

D. Hybrid Decision Engine

The Hybrid Decision Engine integrates the ensemble fatigue probability, deviation percentage from personal baseline, session duration, and time since last alert. It applies a hysteresis mechanism requiring three consecutive threshold crossings before alert generation, and trajectory analysis to detect rapidly worsening fatigue trends. Fatigue deviation is computed as: Fatigue Deviation (%) = ((Current Score − Baseline Score) / Baseline Score) × 100.

E. Interactive Streamlit Dashboard

The complete framework is deployed as a unified web-based dashboard using Streamlit, integrating login authentication, real-time fatigue probability visualization via Plotly gauges, typing metrics display, session summary statistics, historical trend charts, alert configuration, and exportable session reports. Technical outputs are presented in a user-friendly and interpretable format accessible to non-technical users.

IV. EXPERIMENTAL RESULTS

A. Dataset and Experimental Setup

The system was trained using a combination of primary data collected from 25 participants across three-hour monitoring sessions with KSS ratings collected every 15 minutes as ground truth fatigue labels. KSS scores 1-5 were labeled Alert (class 0) and scores 6-9 were labeled Fatigued (class 1) for binary classification. Statistical analysis confirmed significant differences in all 15 features across fatigue states (ANOVA $p < 0.001$ for each feature). Data was split into 70% training, 15% validation, and 15% test partitions using session-level stratified sampling.

B. Classification Model Performance

Four classifiers were trained and evaluated. All metrics are reported on the held-out test set. The table below summarizes performance results:

Table I: Machine Learning Model Performance Metrics

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	82.3%	80.1%	81.4%	80.7%	0.891
Support Vector Machine	86.1%	84.3%	85.7%	85.0%	0.921
Random Forest	89.2%	87.8%	88.4%	88.1%	0.944
Artificial Neural Network	91.3%	89.7%	90.4%	90.1%	0.961
Weighted Ensemble	92.0%	90.5%	91.2%	90.8%	0.968

The Artificial Neural Network achieves the highest individual model performance at 91.3% binary classification accuracy. The weighted ensemble achieves 92.0%, demonstrating the complementarity of different model architectures. Statistical analysis confirmed a 7.2 percentage point accuracy improvement from personalized baseline normalization compared to population-level standardization.

C. Real-Time Monitoring Performance

Beyond classification accuracy, the system was evaluated for real-time performance under live monitoring conditions. The following metrics were measured during live sessions on standard consumer hardware:

Table II: Real-Time System Performance Metrics

Performance Metric	Measured Value
Keystroke capture latency	< 0.5 ms
Feature extraction time per window	< 15 ms
Model inference time (ensemble)	< 25 ms
Total pipeline latency (capture to prediction)	< 40 ms
Dashboard refresh latency	< 200 ms
Memory footprint (full system)	< 180 MB RAM
CPU utilization during monitoring	< 3% (single core)

The system meets all real-time performance requirements. The total pipeline latency of under 40 milliseconds is well within the 30-second prediction window interval. CPU utilization below 3% and memory footprint below 180 MB confirm negligible overhead, allowing concurrent use with any resource-intensive application.

D. Hybrid Decision Engine Validation

In an experimental three-hour programming session, the system successfully detected progressive fatigue beginning approximately 90 minutes into the session. A first-level alert was delivered at the 95-minute mark (mild fatigue probability 0.67), escalating to a severe fatigue alert at the 148-minute mark (probability 0.84). The system correctly identified alertness recovery following a 10-minute break, validating the hysteresis mechanism and trajectory-based assessment.

V. CONCLUSION

This paper presented an AI-Powered Cognitive Fatigue Detection System using Keystroke Dynamics that successfully achieves non-intrusive, real-time cognitive state monitoring. The system integrates keystroke behavioral monitoring, machine learning classification, personalized baseline normalization, and a hybrid decision engine within a unified Streamlit dashboard. The weighted ensemble achieves 92.0% binary classification accuracy with total pipeline latency under 40 milliseconds and CPU utilization below 3%.

The incorporation of personalized baseline normalization provides a 7.2 percentage point accuracy improvement over population-level standardization. The hysteresis-based alert mechanism prevents spurious alerts while ensuring timely detection. By processing only temporal keystroke metadata locally without external data transmission, the system preserves user privacy and complies with data protection regulations.

Future work will explore Long Short-Term Memory and Transformer-based sequence models to capture fatigue accumulation dynamics across consecutive windows, multi-modal behavioral signal fusion incorporating mouse movement and application switching patterns, federated learning for enterprise-scale deployment, and reinforcement learning for personalized recommendation optimization.

ACKNOWLEDGMENT

The authors gratefully acknowledge the participants who contributed their time and data to the experimental sessions. This research was conducted in fulfillment of the academic requirements of [Institution Name].

REFERENCES

[1] H. P. A. Van Dongen, G. Maislin, J. M. Mullington, and D. F. Dinges, "The cumulative cost of additional wakefulness: dose-response effects on neurobehavioral functions," Sleep, vol. 26, no. 2, pp. 117-126, 2003.

[2] M. A. S. Boksem and M. Tops, "Mental fatigue: costs and benefits," Brain Research Reviews, vol. 59, no. 1, pp. 125-139, 2008.

[3] F. Monrose and A. D. Rubin, "Keystroke dynamics as a biometric for authentication," Future Generation Computer Systems, vol. 16, no. 4, pp. 351-359, 2000.

[4] K. S. Killourhy and R. A. Maxion, "Comparing anomaly-detection algorithms for keystroke dynamics," in Proc. 39th IEEE/IFIP Intl. Conf. on Dependable Systems and Networks, 2009, pp. 125-134.

[5] F. Bergadano, D. Gunetti, and C. Picardi, "User authentication through keystroke dynamics," ACM Trans. Inf. Syst. Security, vol. 5, no. 4, pp. 367-397, Nov. 2002.

[6] L. Vizer, L. Zhou, and A. Sears, "Automated stress detection using keystroke and linguistic features," Intl. J. Human-Computer Studies, vol. 67, no. 10, pp. 870-886, 2009.

[7] C. Epp, M. Lippold, and R. Mandryk, "Identifying emotional states using keystroke dynamics," in Proc. CHI Conf. on Human Factors in Computing Systems, 2011, pp. 715-724.

[8] S. Ghosh, E. Papoutsaki, P. Ngao, and J. Huang, "Predicting cognitive load from an activity log during online learning," Computers in Human Behavior, vol. 69, pp. 163-174, 2017.

[9] M. Rastgoo et al., "Automatic driver stress level classification using multimodal deep learning," Expert Systems with Applications, vol. 138, p. 112793, 2019.

[10] A. Sahayadhas, K. Sundaraj, and M. Murugappan, "Detecting driver drowsiness based on sensors: A review," Sensors, vol. 12, no. 12, pp. 16937-16953, 2012.

[11] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Machine Learning Research, vol. 12, pp. 2825-2830, 2011.

[12] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in 12th USENIX Symp. on OS Design and Implementation, 2016, pp. 265-283.

[13] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5-32, 2001.

[14] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273-297, 1995.

[15] T. Akerstedt and M. Gillberg, "Subjective and objective sleepiness in the active individual," Intl. J. Neuroscience, vol. 52, no. 1-2, pp. 29-37, 1990.