Multimodal Digital Phenotyping for Early Prediction of Manic Episodes Through Keystroke Dynamics and Circadian Pattern Analysis

Krish Bhatnagar
Birla Institute of Technology and Science, Pilani
f20230970@pilani.bits-pilani.ac.in

Abstract

Manic episodes in bipolar disorder are characterized by acute behavioral escalation requiring early intervention. This research proposes a multimodal digital phenotyping framework integrating keystroke dynamics with circadian rhythm features to forecast manic episodes 3-7 days prior to clinical onset. The system leverages a hybrid architecture of temporal convolutional and recurrent neural networks with personalized adaptation. It generates risk predictions and clinically actionable alerts while ensuring user privacy through strict on-device processing and data encapsulation. This framework addresses a critical gap in mental healthcare: providing passive, unobtrusive monitoring to detect pre-onset behavioral signatures within a clinically actionable window.

Introduction

Bipolar disorder is a chronic and recurring condition marked by significant shifts in mood and energy levels, with global epidemiological studies indicating that approximately 2.4% of the world's population experiences some form of bipolar spectrum disorder (Merikangas et al. 2011). Current clinical practice, which relies on sporadic assessments and self-reporting, often misses the critical window for preventative care. Passive smartphone monitoring has emerged as a powerful tool for bridging this gap. A comprehensive synthesis of 42 peer-reviewed studies indicates that passive sensing combined with machine learning achieves mood prediction accuracy rates ranging from 67% to 97%, depending on the specific modality and clinical context (Grünerbl et al. 2015; Shen et al. 2025).

Within this domain, keystroke dynamics and circadian rhythms provide complementary signals for predicting manic onset. The BiAffect platform demonstrated that keystroke metadata correlates significantly with mania severity (R2=0.34) (Zulueta et al. 2018), while wearable monitoring of circadian phase advances has achieved predictive accuracy as high as 98% (Lim et al. 2024). This research proposes a multimodal framework integrating these distinct behavioral and physiological data streams. By fusing keystroke mechanics with circadian biomarkers, the system aims to exceed single-modality performance and enable timely clinical intervention in the critical 3–7-day pre-onset window.

Background

Keystroke dynamics in mood disorders capture typing patterns reflecting cognitive and motor changes. The BiAffect study established keystroke metadata predicts manic and depressive symptoms with clinical significance (Zulueta et al. 2018). DeepMood achieved 90.31% accuracy in mood detection using deep learning on keystroke sequences (Cao et al. 2017). Keystroke entropy correlates with cognitive dysfunction during bipolar episodes (Ajilore et al. 2025).

Circadian rhythm instability represents a hallmark biological marker of bipolar disorder. Sleep-wake disruption consistently precedes mood episodes. Lim et al. (2024) demonstrated that daily circadian phase shifts predict manic episodes with 98% accuracy in 168 patients across 44,787 observation days, with phase advances occurring 2-5 days before clinical mania. Song et al. (2024) showed circadian phase disturbances directly precede mood symptoms. Advanced neural network architectures enable accurate modeling of time-series health data, capturing non-linear relationships, and temporal dependencies characterizing psychiatric prodromal periods (Meng et al. 2022). Robust clinical prediction requires multimodal data fusion combining heterogeneous sensors, as single-modality approaches lack sufficient specificity and sensitivity for clinical deployment.

Approach

Data Collection and Privacy Protection use custom on-device keyboard applications to collect anonymized keystroke

metadata including millisecond-precision timing, key type, and accelerometer vectors capturing hand motion. No text content is recorded or stored.
Circadian features derive from ambient light sensors measuring activity exposure, motion accelerometers detecting activity patterns, and device usage logs capturing behavioral rhythms. Privacy is maintained through on-device preprocessing where all feature extraction occurs locally without transmitting raw data, end-to-end AES-256 encryption protecting data in transit. Only encrypted feature vectors transmit to secure research servers.

Feature Extraction identifies keystroke features including key press duration reflecting motor speed, inter-keystroke intervals capturing cognitive delays, typing velocity and acceleration quantifying motion dynamics, and variability metrics reflecting motor consistency. Keystroke entropy measures typing rhythm regularity while irregularity indices capture deviations from baseline (Zulueta et al. 2018). Circadian features include cosinor-derived sleep-wake timing and daily rhythm amplitude, nonparametric metrics quantifying day-to-day consistency and within-day activity fragmentation, and anomaly detection algorithms identifying unusual phase shifts relative to individual baseline (Lim et al. 2024).

Architecture and Personalization employ hierarchical machine learning where specialized modules process keystroke and circadian data independently. One module processes keystroke sequence using temporal convolutional networks with dilated causal convolutions identifying high-frequency behavioral signatures. Other processes daily circadian vectors using bidirectional LSTM networks with multi-head temporal attention mechanisms identifying temporal trends. A learned cross-modal fusion component combines both streams adaptively, dynamically adjusting keystroke versus circadian contributions based on data quality and individual differences. Two-stage transfer learning enables personalization: population-level models trained on aggregate data to provide initialization, followed by on-device fine-tuning.
Clinical Decision Support generates tiered alerts based on predicted risk. Tier 1 represents passive background monitoring with no alert generation, enabling continuous surveillance without alert fatigue. Tier 2 provides gentle user-directed nudges prompting self-monitoring behaviors when moderate risk signals emerge. Tier 3 triggers formal clinician notification including risk magnitude, temporal visualization, and predicted days-to-onset. Each alert includes SHAP-based feature attribution showing which behavioral changes such as increased typing speed combined with irregular sleep patterns triggered the alert.

Evaluation

Building on existing digital phenotyping datasets, preliminary analyses will refine the feature set, assess signal stability, and establish methodological foundations for integrating keystroke dynamics with circadian patterns. Pending institutional ethical approval, a prospective observational study will evaluate the system's ability to anticipate manic episodes through continuous passive monitoring combined with monthly clinical assessments using the Young Mania Rating Scale (YMRS). Primary metrics compare multimodal prediction against single-modality baselines using area under the receiver operating characteristic curve with DeLong statistical significance testing, while secondary metrics include sensitivity, specificity, detection latency, false alarm rates, and calibration metrics, with qualitative feedback evaluating real-world acceptability.

Discussion

The multimodal approach is expected to outperform single-modality systems by capturing complementary behavioral and physiological precursors to mania that keystroke dynamics and circadian monitoring provide independently. By fusing these distinct data streams through learned attention mechanisms, the system is anticipated to gain robustness across individual differences and sensor variability, enabling robust early detection that supports preemptive clinical interventions and reduces hospitalizations. Privacy-preserving on-device processing protects patient data throughout monitoring while interpretable alerts support clinician decision-making.

Conclusion

This proposal presents a multimodal digital phenotyping framework integrating keystroke dynamics and circadian rhythm data to predict manic episodes 3-7 days before clinical onset, addressing a critical unmet need in bipolar disorder management. The system combines temporal convolutional networks for keystroke sequences and LSTM networks for circadian patterns, weighted dynamically via learned attention mechanisms, and validated through preliminary refinement followed by prospective clinical evaluation. Successful implementation is expected to enable timely interventions to reduce hospitalizations and healthcare costs, providing a scalable blueprint for precision psychiatry tools extending to other episodic psychiatric conditions.

References

Ajilore, O.; Leow, A.; Gokkaya, M.; Zhan, L.; GadElkarim, J.; Zhang, A.; Yang, S.; Kumar, S.; Phan, K. L.; and Zulueta, J. 2025. Assessment of Cognitive Function in Bipolar Disorder with Passive Smartphone Keystroke Metadata: A BiAffect Digital Phenotyping Study. Frontiers in Psychiatry 16: 1430303. doi.org/10.3389/fpsyt.2025.1430303.

Cao, B.; Zheng, L.; Zhang, C.; Yu, P. S.; Piscitello, A.; Zulueta, J.; Ajilore, O.; Ryan, K.; and Leow, A. D. 2017. DeepMood: Modeling Mobile Phone Typing Dynamics for Mood Detection. In Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 747-755. New York: Association for Computing Machinery. https://doi.org/10.1145/3097983.3098086

Grünerbl, A.; Osmani, V.; Bahle, G.; et al. 2015. Smartphone-Based Recognition of States and State Changes in Bipolar Disorder Patients. IEEE Journal of Biomedical and Health Informatics 19(1): 140–148.

Lim, D.; Jeong, J.; Song, Y. M.; Cho, C.; Yeom, J. W.; Lee, T.; Lee, J.; Lee, H.; and Kim, J. K. 2024. Accurately Predicting Mood Episodes in Mood Disorder Patients Using Wearable Sleep and Circadian Rhythm Features. npj Digital Medicine 7: 324. doi.org/10.1038/s41746-024-01333-z.

Meng, C.; Trinh, L.; Xu, N.; Enouen, J.; and Liu, Y. 2022. Temporal Convolutional Networks and Data Rebalancing for Clinical Length of Stay and Mortality Prediction. Scientific Reports 12: 21247. doi.org/10.1038/s41598-022-25472-z.

Merikangas, K. R.; Jin, R.; He, J. P.; Kessler, R. C.; Lee, S.; Sampson, N. A.; Viana, M. C.; Andrade, L. H.; Hu, C.; Karam, E. G.; Ladea, M.; Medina-Mora, M. E.; Ono, Y.; Posada-Villa, J.; Sagar, R.; Wells, J. E.; and Zarkov, Z. 2011. Prevalence and Correlates of Bipolar Spectrum Disorder in the World Mental Health Survey Initiative. Archives of General Psychiatry 68(3): 241–251. doi.org/10.1001/archgenpsychiatry.2011.12.

Shen, S. Y.; Qi, W.; Zeng, J.; et al. 2025. Passive Sensing for Mental Health Monitoring Using Machine Learning With Wearables and Smartphones: Scoping Review. JMIR Mental Health 12: e65143. doi.org/10.2196/65143.

Song, Y. M.; Cho, C. H.; Lee, H. J.; and Kim, J. K. 2024. Causal Dynamics of Sleep, Circadian Rhythm, and Mood Symptoms in Patients with Major Depression and Bipolar Disorder: Insights from Longitudinal Wearable Device Data. eBioMedicine 103: 105094. doi.org/10.1016/j.ebiom.2024.105094.

Zulueta, J.; Piscitello, A.; Rasic, M.; Easter, R.; Babu, P.; Langenecker, S. A.; McInnis, M.; Ajilore, O.; Nelson, P. C.; Ryan, K.; and Leow, A. D. 2018. Predicting Mood Disturbance Severity with Mobile Phone Keystroke Metadata: A BiAffect Digital Phenotyping Study. Journal of Medical Internet Research 20(7): e241. doi.org/10.2196/jmir.9775.