

NieNie: Adaptive Rhythmic System for Stress Relief with LLM-Based Guidance

Yichen Yu*

AYXR Research Group
North Carolina State University
Raleigh, North Carolina, USA
Georgia Institute of Technology
Atlanta, Georgia, USA
lunarsboy@gmail.com

Qiaoran Wang*

AYXR Research Group
Raleigh, North Carolina, USA
National University of Singapore
Singapore, Singapore
violetforever999@gmail.com

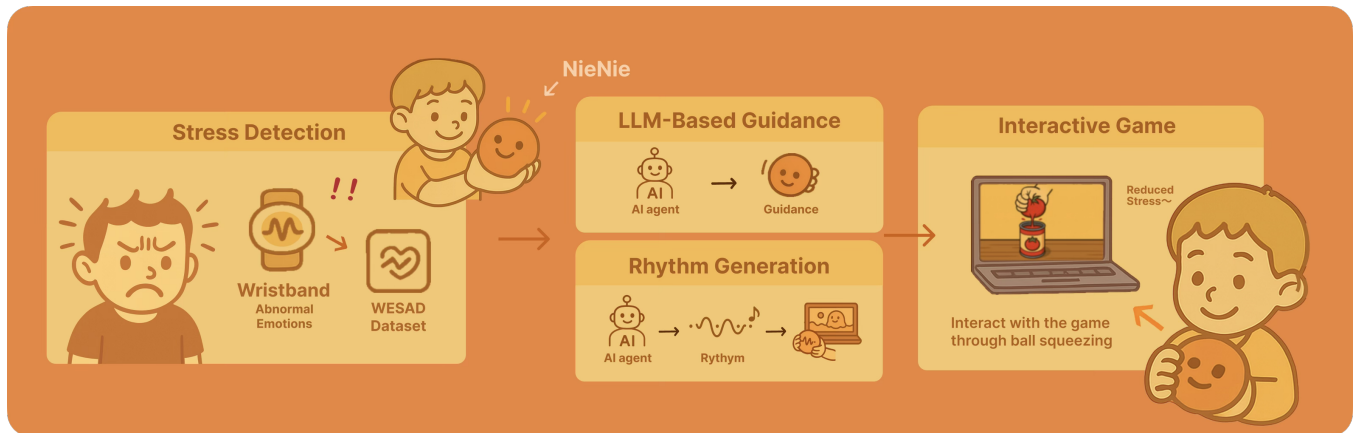


Figure 1: System overview of NieNie, combining stress detection, LLM-driven rhythmic guidance, and interactive squeezing gameplay for emotional regulation.

Abstract

We introduce *NieNie*, an immersive stress relief system that uses rhythmic squeezing and adaptive, emotionally resonant instruction to facilitate emotional self-regulation. Designed specifically to alleviate the daily stress experienced by college students, *NieNie* continuously monitors physiological signals such as heart rate variability and skin conductance, and responds through dynamic rhythmic guidance and language support generated by a large language model (LLM). In contrast to conventional stress-relief apps that rely on passive content or static scripts, *NieNie* embeds personalized, emotionally attuned feedback.

CCS Concepts

• **Human-centered computing** → Ubiquitous and mobile computing systems and tools; *User centered design*; • **Computing**

methodologies → *Cognitive science*; • **Applied computing** → *Consumer health*; *Health informatics*; *Psychology*.

Keywords

Stress regulation, Affective computing, Biofeedback, Large language models, embodied interaction, Ubiquitous computing, Mental health, Emotional support, Rhythm-based interaction, Psychological wellbeing

ACM Reference Format:

Yichen Yu and Qiaoran Wang. 2025. *NieNie*: Adaptive Rhythmic System for Stress Relief with LLM-Based Guidance. In *Proceedings of the 18th ACM Conference on Security and Privacy in Wireless and Mobile Networks (ISWC '25)*, October 12–16, 2025, Espoo, Finland. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3714394.3750586>

1 INTRODUCTION

In recent years, mental health crises among college students have reached an alarming level. Students are facing increasing academic pressure, uncertain career prospects, social isolation, and growing digital distractions, all of which contribute to heightened psychological stress. Despite growing awareness and investment in mental health infrastructure, traditional intervention strategies such as mindfulness meditation apps, counseling sessions, and journaling prompts often fail to effectively engage students during moments of cognitive or emotional overload. Their core limitation lies in their

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ISWC '25, Espoo, Finland.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1481-8/25/10
<https://doi.org/10.1145/3714394.3750586>

reliance on user-driven engagement: individuals in a state of stress may lack the appropriate channels or methods to release stress. Additionally, most existing solutions remain static, lack personalization, and are disconnected from users' physical and emotional states during moments of stress [4, 6].

To address this significant gap, we have made *NieNie*, a real-time, rhythm-guided system that senses physiological indicators of stress and guides users toward calmness through interactive gameplay. Our approach is grounded in research on affective computing, embodied interaction, and narrative psychology. Unlike treating stress as a cognitive challenge requiring rational solutions, *NieNie* employs a more intuitive, body-based feedback loop. Users rhythmically squeeze a soft device under the system's guidance, while the system responds with visual, auditory, and narrative cues that adaptively adjust based on the user's physiological state. All guidance text and rhythms are not pre-written but generated in real-time by a large language model that simultaneously evaluates real-time sensor inputs and recent interaction history. This makes the experience smoother, more empathetic, and personalized, enabling the system to act as a dynamic co-regulator rather than a static guide [7, 14].

NieNie's core vision is to transform stress management from a cognitively demanding, text-heavy process into an embodied, emotionally resonant, and even playful experience. Unlike conventional breathing apps that provide countdown timers or voiceovers, *NieNie* incorporates the user's own sensorimotor feedback into the rhythm generation and response cycle. This not only helps reinforce positive behavioral loops but also fosters a sense of immersion and agency that is often missing from traditional wellness tools. Our system is designed with portability and scalability in mind, enabling deployment in all the places with minimal setup.

In the following sections, we will provide a detailed introduction to existing research related to this design, elaborate on the system architecture and implementation plan, and present the preliminary results of the pilot deployment. Finally, we will reflect on the current limitations of the system and propose future development directions and challenges.

2 RELATED WORK

Our work on *NieNie* builds upon a rich intersection of research in affective computing, biofeedback-based interaction, rhythmic entrainment, and personalized guidance via large language models. This convergence enables us to design a system that is not only reactive to stress but also engaging and adaptive to each user's unique emotional landscape.

Biofeedback games and stress sensing. Early explorations in biofeedback-based interaction have demonstrated the potential of physiological signals such as electrodermal activity (EDA), heart rate variability (HRV), and respiration to serve as real-time indicators of affective state [6, 8]. Notably, Schwartz et al.[12] introduced a game-based biofeedback framework that incorporated EDA and HRV as continuous inputs for modulating gameplay difficulty. Their work showed that embedding physiological awareness into the mechanics of a game increased both user retention and emotional regulation efficacy. Subsequent research expanded this approach

by exploring adaptive feedback loops that altered game dynamics based on biosignal fluctuations [13].

Rhythmic entrainment and embodied regulation. Teixeira et al. [14] demonstrated that rhythm-based breathing games on mobile devices significantly improved physiological indicators of calmness in users. Their work supports the broader theory of entrainment, wherein the human body and mind tend to synchronize with external rhythmic stimuli, such as pulsing lights, sounds, or haptic cues. In the context of *NieNie*, we leverage this principle not through passive exposure but via active, rhythmic interaction—specifically, the act of squeezing and releasing a soft device in coordination with system generated timing cues. This embodied interaction aligns with research on somatic therapy and sensorimotor regulation, where bodily motion becomes a medium for emotional processing.

Narrative and emotional engagement. Emotional engagement is a key predictor of adherence and effectiveness in mental health interventions [7]. In *NieNie*, we extend this by integrating real-time, adaptive narrative feedback generated by a large language model. Unlike scripted systems, our LLM-based module can adjust tone, style, and thematic elements based on the user's ongoing emotional and physiological state, enabling a more responsive and personalized interaction loop.

LLMs for mental wellness. Liu et al.[10] and Valmeekam et al.[15] demonstrate how LLMs can generate personalized prompts for emotional reflection and empathy. *NieNie* positions this capability within an embodied gameplay context, bridging LLM-driven feedback with physical regulation strategies [5, 9].

3 SYSTEM DESIGN

NieNie consists of four interlinked modules designed to detect stress and deliver personalized, rhythmic interventions through multi-modal sensing and real-time feedback. Each module plays a distinct role while maintaining tight integration within a continuous feedback loop.

Stress Detection. At the foundation of *NieNie* lies a real-time stress detection mechanism based on physiological sensing. Using an Apple Watch, the system captures heart rate and heart rate variability (HRV) via built-in PPG sensors. These signals are used to infer stress levels continuously during interaction, without requiring additional sensors or manual calibration. These three biomarkers commonly linked to autonomic nervous system activation. These raw signals are processed through an LSTM-based model trained on the WESAD dataset, which has been widely used for wearable stress classification [6, 11]. The model operates on short, overlapping time windows and outputs a continuous stress probability score ranging from 0 (calm) to 1 (high stress). This probability is not treated as a fixed label but rather as a modulating signal that influences the downstream rhythm pacing and narrative tone. To accommodate inter-user variability, the system includes a lightweight calibration phase to adapt baseline thresholds to each user's physiology. This ensures that feedback is not triggered unnecessarily, and that changes in physiological state rather than absolute values drive the adaptive logic of the experience.

WESAD Dataset. To enable real-time stress inference in *NieNie*, we trained a lightweight LSTM classifier using the WESAD

dataset [6], a benchmark for wearable-based stress detection. We focused on three physiological signals: electrodermal activity (EDA), skin temperature (TEMP), and heart rate (HR), sourced from the Empatica E4 sensor streams.

Each signal was preprocessed and truncated to equal length, then stacked into a multivariate time series. We segmented the data using a sliding window approach with a window size of 40 and a step size of 20, resulting in fixed-length input sequences of shape (40, 3). To simulate the three affective states—baseline, stress, and amusement—we assigned labels based on temporal segmentation of the recording sessions.

We trained a two-layer LSTM classifier implemented in PyTorch, using 128 hidden units and a final softmax output layer for three-class classification. The model was trained using cross-entropy loss and optimized with the Adam optimizer. A stratified 80/20 train-test split was used, and the model achieved strong performance on the held-out test set.

This LSTM classifier was later converted to TensorFlow Lite and quantized for efficient deployment on mobile devices. During runtime, the model consumes overlapping 40-frame windows of real-time HRV and temperature data to produce a continuous probability estimate of user stress. This signal drives the dynamic adaptation of rhythm pacing and narrative tone within the *NieNie* system.

LLM-Based Guidance. *NieNie* uses large language models (LLMs) to generate real-time narrative prompts. Unlike systems that rely on fixed scripted messages, our LLM-driven module generates dynamic, personalized guidance based on rhythm adherence, pressure trajectories, and performance trends. For example, when a user maintains a steady squeezing motion in sync with the rhythm, the system might respond: “Keep up this steady rhythm, you’re in sync with the rhythm. Imagine you’re squeezing this fruit rhythmically.” During periods of irregular input or increased pressure, it may offer supportive prompts like: “Try slowing down to find your rhythm again.” These messages are concise, context-aware, and seamlessly integrated into the game narrative. The tone of feedback (supportive, neutral, or motivational) is adjusted in real-time to ensure it is both useful and non-intrusive. By combining physiological signals with adaptive language generation, this component transforms raw sensor data into emotionally resonant and actionable feedback [9, 10, 15].

This feedback mechanism directly supports subsequent modules, allowing users to interact with soft, fruit-shaped virtual objects via squeezing gestures. Haptic metaphors reinforce rhythm regulation, creating a closed-loop system between emotional states, narrative cues, and embodied interaction.

Rhythm Generation. Based on the inferred stress level, *NieNie* generates squeeze-release rhythm patterns to guide the user’s tactile interaction. This module sets time rhythm instructions (e.g., “squeeze 5 times continuously”), which are adjusted according to the user’s current autonomic nervous system state. When stress levels rise, the rhythm may begin with shorter, faster cycles to match the user’s physiological arousal state, then gradually lengthen the cycles to promote parasympathetic activation and relaxation. Unlike systems relying on static patterns, the rhythm generator is dynamically adjusted in real-time by the system’s large language

model (LLM) component, which analyzes trends in user engagement and physiological metrics. The generated rhythm serves as a non-verbal, embodied intervention consistent with research on synchronization and breath-based emotional regulation[2, 14]. Visual and auditory cues embedded within the Unity environment reinforce the rhythm, enabling users to synchronize their bodily actions with system guidance in a fluid and intuitive manner.

Interactive Game. All feedback and perceptions converge in *NieNie*’s Unity-based interactive game environment. Users interact by rhythmically squeezing a Bluetooth-connected soft device with a built-in pressure sensor. The game world responds in real time to the user’s input and physiological state. Successful synchronization with the rhythm triggers positive in-game changes. These metaphors externalize users’ internal states, helping them view their stress levels and no longer view them as hidden metrics but as malleable, interactive forces. For example, by following different in-game rhythm metaphors, users can squeeze in-game tomatoes to create tomato cans that meet rhythm requirements (see Figure 2). Conversely, disconnection or imbalance encourages users to reengage. Unity handles all rendering and logic processing, synchronized with sensor inputs and language generated by large language models (LLMs). This integration creates a deeply immersive experience where every squeeze, pause, and narrative cue forms part of an ongoing feedback loop, supporting emotional regulation through presence, rhythm, and playfulness[7].




Figure 2: Users squeeze tomatoes into the can according to the rhythm and frequency of the cans appearing. After squeezing, new tomatoes will appear. The filled cans slide to the left, while new empty cans slide in from the right.

4 IMPLEMENTATION

The *NieNie* system is implemented using off-the-shelf wearable hardware, a custom-built soft interaction device, a Unity-based game engine, and a cloud-connected large language model API. This modular architecture enables real-time physiological sensing, multimodal feedback, and responsive gameplay within a lightweight mobile application [1, 8].

Physiological data is collected using an Apple Watch, which directly provides heart rate and heart rate variability (HRV) measurements via built-in photoplethysmography (PPG) sensors (see Figure 3). These metrics are streamed in real time to a paired iPhone using Apple’s HealthKit and CoreBluetooth APIs, enabling continuous physiological monitoring without requiring manual signal

processing or feature extraction. A compact LSTM classifier trained on the WESAD dataset continuously infers stress probabilities from overlapping time windows. The model is deployed via TensorFlow Lite for efficient on-device inference, achieving sub-30ms latency on modern smartphones [3, 11].

The soft interaction system consists of a flexible pressure sensor embedded in a squeezable silicone sleeve. A microcontroller transmits normalized squeeze intensity values at 30Hz to Unity, which synchronizes the input with a rhythm clock that governs the pacing of squeeze–release cycles. This temporal alignment is used to calculate rhythm accuracy and engagement patterns in real time [13].

Unity renders an interactive game environment where users engage in rhythmic tasks such as squeezing virtual tomatoes in sync with a pacing rhythm. This visual interaction is complemented by natural language messages generated by an LLM (e.g., GPT-4o), delivered through a backend server with optimized prompt templates. These messages are brief, adaptive, and emotionally supportive, enriching the experience without requiring textual input from the user.

The entire system runs as an application, ensuring portability and ease of deployment. Interaction logs, rhythm metrics, and LLM feedback history can be optionally recorded for further evaluation under IRB-compliant protocols.



Figure 3: Apple Watch used in *NieNie* for real-time heart rate and HRV sensing, with data transmitted to the Unity game system.

5 Limitations and Future Work

Although *NieNie* demonstrates the potential of integrating physiological sensing, rhythmic interaction, and feedback based on large language models (LLMs) in mobile systems, several limitations remain, which provide guidance for future research directions.

Accurately detecting stress from physiological data remains a challenge. Individual baseline differences, inconsistent sensor placement, and the transient nature of emotions can all impact inference accuracy. While current models utilize heart rate variability (HRV), their generalizability across broader populations has yet to be validated. Future research will explore user-specific calibration and the incorporation of additional sensor modalities to enhance detection.

The emotional consistency of generated feedback also introduces design complexity. While large language models can provide rich and adaptable information, they occasionally generate outputs that are contextually inconsistent, overly abstract, or emotionally inappropriate. Enhancing user autonomy (e.g., allowing message filtering or style customization) can help ensure more resonant and respectful interactions. In terms of privacy and trust, even in systems that do not continuously store user data, continuous monitoring and personalized responses may raise concerns. Some users may perceive the system as intrusive or opaque, especially when content generated by large language models (LLMs) lacks clear attribution.

Maintaining long-term engagement after the novelty wears off remains uncertain. While the system’s gamified feedback loops can drive initial interaction, sustaining user motivation over the long term, especially in unstructured daily scenarios is a known challenge. Mechanisms such as goal setting, environmental variation, and narrative progression will be explored to promote long-term use and emotional impact.

Overall, addressing these limitations requires interdisciplinary collaboration across fields such as affective computing, human-computer interaction, behavioral psychology, and privacy research to ensure that *NieNie* provides safe, inclusive, and meaningful solutions for mental health support in the real world.

References

- [1] Mohamed Abd Al-Alim, Roaa Mubarak, Nancy M. Salem, and Ibrahim Sadek. 2024. A machine-learning approach for stress detection using wearable sensors in free-living environments. *Computers in Biology and Medicine* 179 (Sept. 2024), 108918. doi:10.1016/j.combiomed.2024.108918
- [2] Alaa Abd-alrazaq, Mohannad Alajlani, Reham Ahmad, et al. 2024. The Performance of Wearable AI in Detecting Stress Among Students: Systematic Review and Meta-Analysis. *Journal of Medical Internet Research* 26 (2024). doi:10.2196/52622
- [3] Md Santo Ali, Sapnil Sarker Bipro, et al. 2025. TEANet: A Transpose-Enhanced Autoencoder Network for Wearable Stress Monitoring. In *arXiv preprint arXiv:2503.12657*.
- [4] Jane M Burns, Tracey A Davenport, Lachlan A Durkin, Georgie M Luscombe, and Ian B Hickie. 2011. Mobile phone-based mental health support: A review of the evidence. *The Medical Journal of Australia* 194, 4 (2011), S20–S22.
- [5] Neo Christopher Chung, George Dyer, and Lennart Brocki. 2023. Challenges of Large Language Models for Mental Health Counseling. *arXiv preprint arXiv:2311.13857* (2023).
- [6] Martin Gjoreski, Hristijan Gjoreski, Mitja Luštrek, and Matjaž Gams. 2018. Wearable affect and stress recognition: A review of machine learning methods and current challenges. (2018), 214–219.
- [7] Jennifer Gonzalez, Erin Miller, et al. 2020. Emotional metaphors in VR: grounding metaphorical feedback to support mental health reflection. In *CHI EA '20*.
- [8] Javier Hernandez and Ehsan Hoque. 2014. BioGlass: Physiological parameter estimation using a head-mounted wearable device. In *UbiComp*. 25–34.
- [9] Paul Pu Liang, Zhe Wu, et al. 2023. Holistic evaluation of chatgpt in mental health: Adherence to clinical guidelines, hallucination, and bias. *arXiv preprint arXiv:2302.06656* (2023).
- [10] Meng Liu, Mengdi Yin, Yuancheng Wang, and Jason I. Hong. 2023. MindfulBot: Exploring Large Language Models for Interactive Mindfulness and Behavior Change. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–24. doi:10.1145/3610896
- [11] Naful Rashid, Luke Chen, et al. 2021. Feature Augmented Hybrid CNN for Stress Recognition Using Wrist-based Photoplethysmography Sensor. *arXiv preprint arXiv:2108.03166* (2021).
- [12] K. Schwartz, H. Klapperich, and S. Kopp. 2017. GBF: Game Biofeedback Framework. In *ACM International Conference on Multimodal Interaction (ICMI)*. 514–518.
- [13] A. Tammami, J. Mendoza, R. Muñoz, and L. E. Bruni. 2023. Design and Evaluation of Adaptive Biofeedback Games for Relaxation. In *CHI Conference on Human Factors in Computing Systems*. 1–13.
- [14] Andre Teixeira, Satvik Kumar, Verena Hafner, and Ana Paiva. 2023. Gaming away stress: Designing breathing games for emotion regulation. *International Journal of Human-Computer Studies* 174 (2023), 103136.

- [15] Kumar Valmееkam, Emily Rho, and Timothy Bickmore. 2023. Adaptive Empathetic Agents Using Large Language Models for Emotional Support. In *Proceedings*

of the 2023 CHI Conference on Human Factors in Computing Systems. ACM, 1–15. doi:10.1145/3544548.3581031