NeuroScape: EEG-Adaptive Spatial Soundscapes for Meditation through LLM-Guided Scene Planning

Jianing Yu*
Harvard University
Cambridge, Massachusetts, USA
nomy_yu@gsd.harvard.edu

Jingyi Zoe Liu*
Harvard University
Cambridge, Massachusetts, USA
jingyi_liu@gsd.harvard.edu

Xinyan Teresa Li*
Harvard University
Cambridge, Massachusetts, USA
xinyanli@gsd.harvard.edu

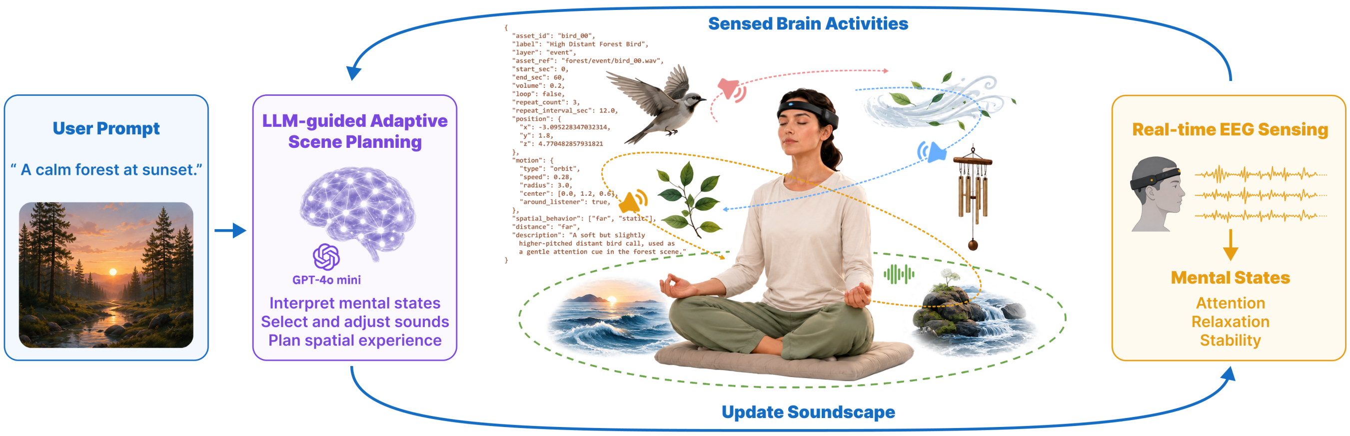

Figure 1: *NeuroScape* transforms spatial audio from a static meditation background into a responsive environment that adapts to users' internal states. The system combines user-guided scene initialization, EEG sensing, adaptive scene planning, and real-time spatial soundscape rendering in a continuous closed loop.

Abstract

Meditation soundscapes are typically pre-recorded and cannot respond to users' changing internal states. We present *NeuroScape*, an EEG-adaptive spatial soundscape system for meditation through LLM-guided scene planning. Informed by a formative study, NeuroScape organizes sounds into ambient, event, and action layers and continuously adapts their selection, density, position, and movement using real-time EEG signals. A preliminary pilot with four novice meditators demonstrates the technical feasibility of this closed-loop pipeline and shows individualized cognitive and auditory trajectories across sessions. Our work positions LLM-guided scene planning as a semantic bridge between physiological signals and adaptive auditory environments.

CCS Concepts

• **Human-centered computing** → **Human computer interaction**; *Auditory feedback*; Ubiquitous computing; User studies.

*These authors contributed equally to this research.

UbiComp/ISWC '26, Shanghai, China

ACM ISBN 978-1-4503-XXXX-X/2026/XX
https://doi.org/XXXXXXX.XXXXXXX

Keywords

Electroencephalography (EEG), brain–computer interface, neuroadaptive systems, meditation, adaptive soundscapes, spatial audio, large language models, biofeedback, closed-loop interaction.

ACM Reference Format:
Jianing Yu, Jingyi Zoe Liu, and Xinyan Teresa Li. 2026. NeuroScape: EEG-Adaptive Spatial Soundscapes for Meditation through LLM-Guided Scene Planning. In *Proceedings of Eye Tracking and Brain-Computer Interfaces in AR/VR: Progresses, Challenges and Opportunities Workshop, UbiComp/ISWC 2026 (UbiComp/ISWC '26).* ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Meditation can support attention, emotional regulation, and well-being [9]. However, novice practitioners often struggle to sustain focus, return from distraction, and remain engaged in the practice [4, 13]. While auditory support can scaffold this process without requiring visual attention or physical interaction, many meditation tools still rely on pre-recorded voice guidance, music playlists, or environmental audio that remains unchanged as users' internal states fluctuate [10, 16]. We use the term spatial soundscape to refer to an acoustic environment in which multiple sound sources are organized through direction, distance, position, and movement to create an immersive sense of place and support presence, attention, and embodied orientation [3]. Prior work has explored interactive soundscapes, EEG-modulated audio, and AI-generated auditory experiences [5, 6, 8]. However, continuously adapting a meditation soundscape presents a distinct challenge: the system must coordinate multiple sound roles, spatial behaviors, and temporal transitions while preserving environmental coherence and avoiding changes that disrupt meditation. In parallel, explicit feedback such as voice can interrupt inward attention, motivating the use of implicit physiological sensing. EEG is a non-invasive neurophysiological method that measures and records electrical activity generated by the brain using scalp-mounted electrodes [2]. EEG provides a continuous signal associated with meditation-related fluctuations in attention and relaxation.

We design *NeuroScape*, an EEG-adaptive spatial soundscape system for meditation through LLM-guided scene planning. *NeuroScape* combines user-guided scene initialization, EEG-based state estimation, LLM-based scene planning, and Unity-based spatial audio rendering in a closed loop. Rather than directly mapping a single EEG feature to an audio parameter, the system uses structured estimates of attention, relaxation, and stability as contextual input for scene-level adaptation. Sounds are organized into ambient, event, and action layers whose selection, density, position, and movement evolve while preserving semantic and spatial continuity.

This paper first presents a formative study that informed *NeuroScape*'s adaptation principles. Then it describes the system architecture, including its three-layer soundscape representation, EEG-to-scene adaptation process, knowledge-grounded LLM planner, and real-time spatial audio rendering pipeline. Finally, there is a preliminary pilot study examining the technical feasibility and session-specific behavior of the system. Through *NeuroScape*, we explore how brain-derived signals can serve as implicit context for adaptive spatial interaction, and how such signals may be translated into coherent, controllable, and non-disruptive changes in an immersive auditory environment.

2 Related Work

2.1 Auditory Support for Meditation

Meditation requires sustaining attention and repeatedly returning from distraction [2]. Auditory support can scaffold this process without requiring visual interaction, particularly for novice practitioners. Existing approaches include verbal guidance, rhythmic or musical cues, and environmental soundscapes [15]. We use soundscape to refer to an acoustic virtual environment as perceived in context [11], emphasizing the coordinated composition, spatial arrangement, movement, and continuity of multiple sounds. Natural and spatial sounds can support restoration, presence, and embodied orientation [3, 7]. Prior meditation systems have explored interactive and EEG-modulated soundscapes [4, 5], but largely rely on predefined mappings, leaving open how multiple sound roles and spatial behaviors can be coordinated as users' states evolve.

2.2 Adaptive Auditory Experiences

Recent systems demonstrate how AI can personalize auditory experiences. *Sonora* supports conversational co-creation of navigable 3D audio worlds [6]; *AImagery* combines user preferences and physiological information to generate guided imagery and spatialized sounds [1]; and *RunMe* adapts auditory support using real-time user data [17]. However, adaptive soundscape generation must coordinate multiple concurrent sources while preserving coherence across changes in sound role, density, position, movement, and timing. *NeuroScape* therefore uses an LLM as a semantic scene planner that reasons over user state, environmental context, available assets, and continuity constraints rather than directly generating audio.

2.3 EEG and Implicit State Sensing

Adaptive systems may use explicit feedback or implicit physiological signals. Although explicit input can communicate user intentions accurately, repeated verbal or manual interaction may disrupt meditation. Physiological sensing allows adaptation to occur while users remain engaged in the practice [12]. Signals such as heart rate, HRV, respiration, and electrodermal activity have been used to estimate arousal and autonomic regulation [1]. EEG provides a complementary signal associated with meditation-related fluctuations in attention, relaxation, and mind wandering [2]. Consumer devices such as Muse make EEG accessible to interactive systems, although the measurements remain noisy and require cautious interpretation [14, 18]. Prior work has used EEG for neurofeedback and direct sound modulation [5, 8], but less is known about using EEG-derived states as contextual input for LLM-guided scene adaptation. We therefore conducted a formative study to derive preliminary relationships between auditory structures, subjective meditation experience, and EEG-derived dynamics.

3 Formative Study

To inform the adaptation logic of *NeuroScape*, we conducted an exploratory within-subject study with six novice meditators. Each participant experienced three randomized 10-minute conditions: voice-only guidance, voice guidance with non-spatial environmental audio, and voice guidance with spatialized environmental audio. The voice-only condition served as a baseline for examining the additional effects of environmental sound and spatialization. All conditions followed the same meditation narrative. We collected EEG recordings using a Muse 2 headband, post-condition questionnaires, and semi-structured interviews after all conditions.

Subjective responses indicated a consistent preference for the spatial-audio condition. All participants preferred this condition and described it as more immersive, while five reported feeling more focused and all described a stronger sense of bodily grounding or presence within the environment. Several participants also commented on their sensitivity to the direction and movement of spatial sounds. When interpreted alongside exploratory EEG trends, localized event sounds, such as birds passing overhead, appeared to function as gentle attentional cues and were sometimes accompanied by increased beta-band activity. Continuous ambient sounds, such as wind or ocean waves, were more often associated with calmness and environmental continuity. Body action sounds, including breathing and footsteps, created a sense of movement within the scene and heightened participants' awareness of their bodies in relation to the surrounding environment.

These observations informed four design insights for adaptive spatial soundscapes. First, a persistent ambient layer should provide environmental continuity and support calmness and immersion. Second, localized spatial events may serve as occasional attentional anchors. Third, action sounds should be introduced selectively to reinforce bodily grounding. Finally, adaptations should unfold gradually to avoid abrupt changes that disrupt the meditation experience. These principles motivated our ambient-event-action sound design and conservative scene-level planning strategy.

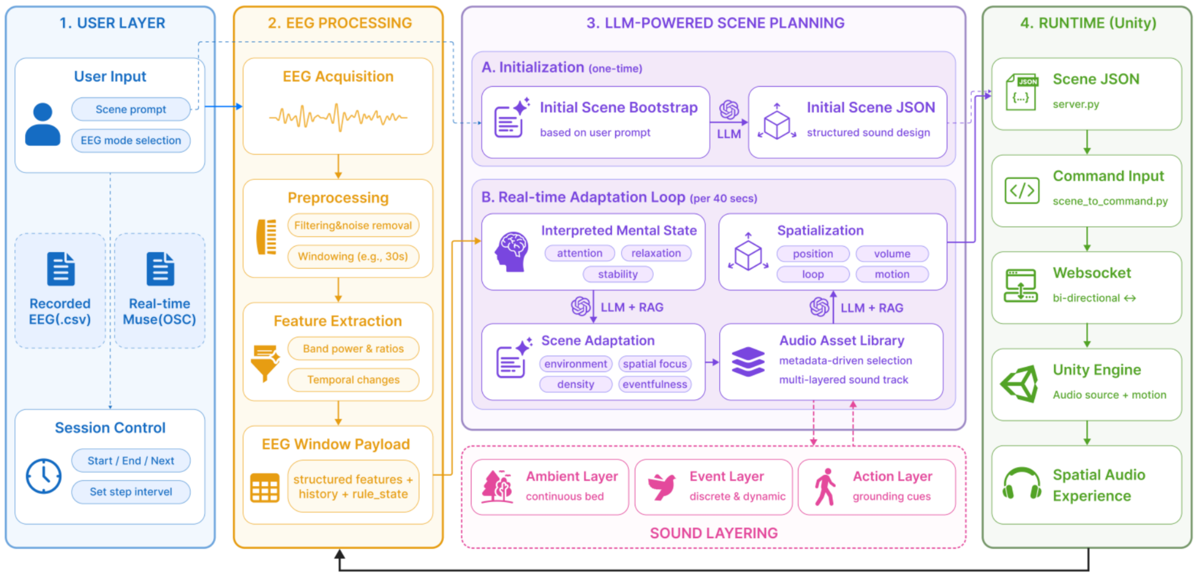

Figure 2: *NeuroScape*'s closed-loop architecture combines user-guided scene initialization, EEG-based state estimation, LLM-guided scene planning, and real-time spatial audio rendering in Unity.

4 NeuroScape

4.1 Overall System Architecture

NeuroScape continuously transforms EEG-derived brain activity into adaptive spatial soundscapes. Figure 2 illustrates the overall architecture of *NeuroScape*. Before meditation, users describe a desired natural environment using a short prompt, such as "a peaceful forest after rain" or "a quiet beach at sunset". GPT-4o mini first generates an initial scene blueprint describing the environmental context, candidate sound sources, spatial layout, and initial sound-layer configuration.

During meditation, EEG signals are streamed from a Muse 2 headband through Mind Monitor into a Python pipeline, where they are interpreted as contextual indicators of the user's evolving meditation state. Together with the current scene configuration, these indicators are passed to an LLM-based scene planner, which generates structured updates for how the soundscape should evolve.

The resulting scene plan is serialized as a JSON object and transmitted to Unity through a persistent WebSocket connection. Unity validates the updates and performs spatial audio rendering by modifying sound sources, spatial positions, movement trajectories, playback timing, and transition behaviors. The updated auditory environment then feeds back into the user's meditation state, completing the continuous EEG-adaptive loop.

4.2 EEG Interpretation

NeuroScape acquires EEG signals using a Muse 2 headband and streams them into Python via Mind Monitor. Incoming signals are grouped into overlapping 60-second windows with updates every 40 seconds.

For each window, the system extracts relative spectral power and temporal variability from the five canonical EEG frequency bands (delta, theta, alpha, beta, and gamma). These features are normalized and combined using predefined heuristic rules to estimate three operational indicators: relaxation, attention, and stability.

Rather than representing clinical measures, these indicators provide a stable and interpretable representation of the user's meditation state for scene planning, reducing the effects of transient fluctuations and noise in raw EEG signals.

4.3 Knowledge-Grounded LLM Scene Planning

Rather than directly mapping EEG-derived indicators to predefined sound effects, *NeuroScape* employs GPT-4o mini as a knowledge-grounded semantic scene planner. The planner operates over a structured scene knowledge base that defines the available sound assets, their semantic and spatial properties, and the constraints governing how they may be combined. This grounding allows the model to generate contextually appropriate scene updates within an executable design space.

The scene knowledge base annotates each sound asset with its environmental role (ambient, event, or action), affective characteristics, compatible meditation conditions, intensity range, spatial behavior, transition constraints, and compatibility with other scene elements. These annotations allow the planner to treat sounds as components of a coordinated spatial environment rather than isolated audio clips. LLM can therefore select sounds, arrange them around the listener, assign motion trajectories, and control how they enter, persist, or fade. In this way, it constructs an evolving auditory scene from a finite library of grounded assets rather than producing unconstrained descriptions from natural-language prompts alone.

For each planning cycle, the LLM planner jointly considers the user's recent physiological trajectory and the current environmental state to determine whether the scene should remain stable or evolve. For example, reduced attention may prompt the introduction of a localized, low-salience event sound, while high stability

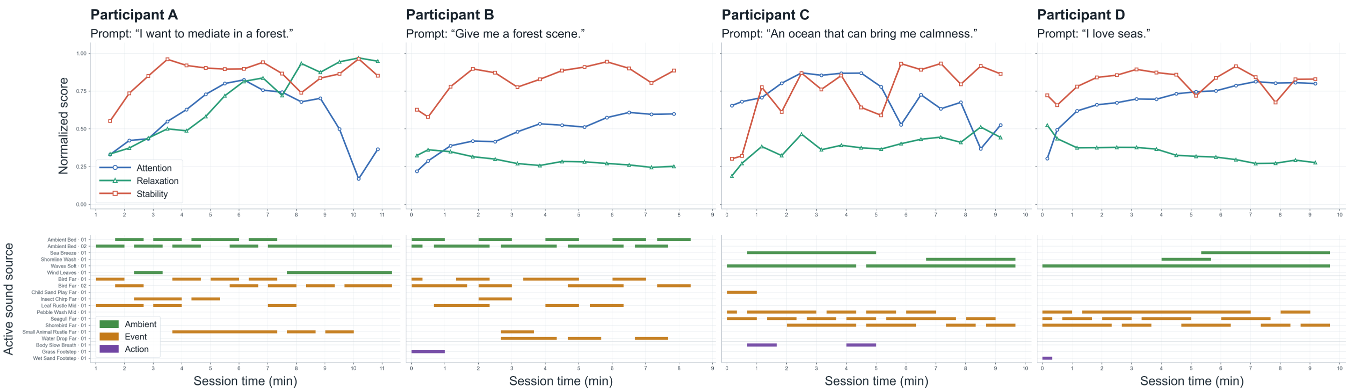

Figure 3: Cognitive-state trajectories and adaptive soundscape evolution across four 10-minute pilot meditation sessions. The upper panel shows EEG-derived attention, relaxation, and stability estimates over time. The lower panel shows the activation timeline of ambient, event, and action sound layers generated by NeuroScape.

may lead the planner to preserve the current scene. Sustained relaxation may instead support richer ambient textures and fewer salient events. These context-sensitive decisions enable scene-level adaptation beyond deterministic one-to-one mappings.

The planner outputs a structured JSON scene description specifying sound selection, spatial placement, movement trajectories, event timing, transition duration, and the relative balance among layers. Constraining the output to a predefined schema ensures reliable communication between the planning module and the rendering engine.

4.4 Spatial Audio Rendering

The generated scene plan is transmitted to Unity through a persistent WebSocket connection. The current prototype supports forest and coastal scene using a curated JSON sound library. Each sound asset is annotated with its role, affective characteristics, recommended usage, spatial behavior, and movement type. Supported behaviors include stationary, localized random, drifting, orbiting, overhead passing, and approach–recede movement. Based on the LLM-generated scene plan JSON, Unity Engine creates, updates, and removes sound sources with smooth cross-fading while controlling spatial positioning, distance attenuation, looping, movement interpolation, playback intensity, and timing.

5 Pilot Evaluation

We conducted a preliminary pilot with four novice meditators to examine the technical behavior of *NeuroScape* and the adaptation patterns produced by its closed-loop pipeline. Participants experienced either a forest or coastal soundscape generated from their initial prompts, with each session lasting approximately 10 minutes. This pilot was intended to assess feasibility and system behavior rather than effectiveness relative to a control condition.

Figure 3 shows the temporal evolution of the EEG-derived attention, relaxation, and stability indicators together with the activation of ambient, event, and action sound layers. Across the four sessions, the system produced different indicator and soundscape trajectories rather than a uniform adaptation pattern. Stability generally increased during the early portion of the sessions, while attention and relaxation varied across participants and over time.

The soundscape updates also reflected combinations of indicators rather than a simple one-to-one mapping. In several sessions, periods of lower estimated attention or stability coincided with more frequent event-sound activation, whereas more stable and relaxed periods were accompanied by sustained ambient sound and fewer discrete events. In some cases, later declines in attention did not trigger additional event sounds when the overall state remained stable, suggesting that the planner preserved continuity instead of reacting to every local fluctuation.

These observations demonstrate the technical feasibility of coupling EEG-derived state estimates with knowledge-grounded LLM scene planning to produce session-specific spatial soundscape trajectories. However, the small sample, short sessions, and absence of a matched non-adaptive condition prevent further conclusions about meditation effectiveness or the superiority of the current planning strategy. The EEG-derived indicators should also be regarded as operational estimates rather than definitive measures of mental state, and the current prototype remains limited to two environments and a curated asset library. Our next steps are to compare NeuroScape with matched non-adaptive soundscapes, improve state estimation through individual calibration and complementary physiological signals, and evaluate alternative planning strategies in terms of coherence, responsiveness, consistency, latency, and perceived intrusiveness.

6 Conclusion

NeuroScape is a neuroadaptive auditory system that transforms EEG-derived state estimates into evolving spatial soundscapes through knowledge-grounded LLM scene planning and real-time Unity rendering. Its layered sound representation and constrained planning pipeline demonstrate how brain signals can serve as implicit context for adaptive immersive interaction. A formative study and preliminary pilot informed the design and established the feasibility of producing coherent, session-specific scene adaptations.

Acknowledgments

We thank Dr. Katarina Richter-Lunn at the Harvard University Graduate School of Design for her insightful guidance and constructive feedback throughout the development of *NeuroScape*.

References

[1] Judith Amores, Kael Rowan, Javier Hernandez, and Mary Czerwinski. 2024. AImagery: A Multisensory Approach to Anxiety Reduction with AI, Olfactory Stimuli, and Biofeedback-Enhanced Guided Imagery. In *2024 12th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 378–393.

[2] Paul Barrows, William Van Gordon, and Paul Gilbert. 2024. Current Trends and Challenges in EEG Research on Meditation and Mindfulness. *Discover Psychology* 4, Article 148 (2024). doi:10.1007/s44202-024-00269-5

[3] Hyunsung Cho, Alexander Wang, Divya Kartik, Emily Liying Xie, Yukang Yan, and David Lindlbauer. 2024. Auptimize: Optimal Placement of Spatial Audio Cues for Extended Reality. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) *(UIST '24)*. Association for Computing Machinery, New York, NY, USA, Article 126, 14 pages. doi:10.1145/3654777.3676424

[4] Karen Cochrane, Lian Loke, Andrew Campbell, and Naseem Ahmadpour. 2021. Mediscape: Preliminary Design Guidelines for Interactive Rhythmic Soundscapes for Entraining Novice Mindfulness Meditators. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction* (Sydney, NSW, Australia) *(OzCHI '20)*. Association for Computing Machinery, New York, NY, USA, 379–391. doi:10.1145/3441000.3441052

[5] Karen Anne Cochrane, Lian Loke, Caitilin de Bérigny, and Andrew Campbell. 2018. Sounds in the moment: designing an interactive EEG nature soundscape for novice mindfulness meditators. In *Proceedings of the 30th Australian Conference on Computer-Human Interaction* (Melbourne, Australia) *(OzCHI '18)*. Association for Computing Machinery, New York, NY, USA, 298–302. doi:10.1145/3292147.3292215

[6] Fernanda M De La Torre, Javier Hernandez, Andrew D Wilson, and Judith Amores. 2025. Sonora: Human-AI Co-Creation of 3D Audio Worlds and its Impact on Anxiety and Cognitive Load. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 308, 22 pages. doi:10.1145/3706598.3713316

[7] Lia R. V. Gilmour, Isabelle Bray, Chris Alford, and Paul R. Lintott. 2024. Natural Soundscapes Enhance Mood Recovery amid Anthropogenic Noise Pollution. *PLOS ONE* 19, 11 (2024), e0311487. doi:10.1371/journal.pone.0311487

[8] Marzieh Haghighi, Mohammad Moghadamfalahi, Murat Akcakaya, and Deniz Erdogmus. 2018. EEG-assisted modulation of sound sources in the auditory scene. *Biomedical Signal Processing and Control* 39 (2018), 263–270. doi:10.1016/j.bspc.2017.08.008

[9] Shian-Ling Keng, Moria J. Smoski, and Clive J. Robins. 2011. Effects of mindfulness on psychological health: A review of empirical studies. *Clinical Psychology Review* 31, 6 (2011), 1041–1056. doi:10.1016/j.cpr.2011.04.006

[10] Madhavan Mani, David J Kavanagh, Leanne Hides, and Stoyan R Stoyanov. 2015. Review and Evaluation of Mindfulness-Based iPhone Apps. *JMIR mHealth uHealth* 3, 3 (19 Aug 2015), e82. doi:10.2196/mhealth.4328

[11] Andrew Mitchell, Francesco Aletta, Tin Oberman, and Jian Kang. 2023. How Do We Define Soundscape?. In *Proceedings of the 10th Convention of the European Acoustics Association (Forum Acusticum 2023)*. Turin, Italy, 695–699. doi:10.61782/fa.2023.0359

[12] Clara Moge, Katherine Wang, and Youngjun Cho. 2022. Shared User Interfaces of Physiological Data: Systematic Review of Social Biofeedback Systems and Contexts in HCI. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3491102.3517495

[13] Julio Rodriguez-Larios and Kaat Alaerts. 2021. EEG alpha–theta dynamics during mind wandering in the context of breath focus meditation: An experience sampling approach with novice meditation practitioners. *European Journal of Neuroscience* 53, 6 (2021), 1855–1868. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ejn.15073 doi:10.1111/ejn.15073

[14] Joseph Sabio, Nadia S. Williams, Genevieve M. McArthur, and Nicholas A. Badcock. 2024. A scoping review on the use of consumer-grade EEG devices for research. *PLOS ONE* 19, 3 (2024), e0291186. doi:10.1371/journal.pone.0291186

[15] Marina Saskovets, Irina Saponkova, and Zilu Liang. 2025. Effects of Sound Interventions on the Mental Stress Response in Adults: Scoping Review. *JMIR Mental Health* 12 (2025), e69120. doi:10.2196/69120

[16] Katrin Schwartz, Fabienne Marie Ganster, and Ulrich S Tran. 2023. Mindfulness-Based Mobile Apps and Their Impact on Well-Being in Nonclinical Populations: Systematic Review of Randomized Controlled Trials. *J Med Internet Res* 25 (4 Aug 2023), e44638. doi:10.2196/44638

[17] Peng Tan, Xiaofei Zhu, Tao Bi, and Xiangshi Ren. 2025. RunMe: An adaptive sound system for running meditation. *ACM Transactions on Computer-Human Interaction* 32, 4 (2025), 1–15.

[18] Isaac Treves, Zia Bajwa, Keara D. Greene, Paul A. Bloom, Nayoung Kim, Emma Wool, Simon B. Goldberg, Susan Whitfield-Gabrieli, and Randy P. Auerbach. 2025. Consumer-Grade Neurofeedback With Mindfulness Meditation: Meta-Analysis. *Journal of Medical Internet Research* 27 (2025), e68204. doi:10.2196/68204