

Emotion-To-Motion: A Wearable EEG–GSR Biofeedback System For Emotion-Aware Meditation

Prapawit Intun

Participatory Citizen Lab. Bangkok, Thailand participatorycitizenlab@gmail.com
<https://orcid.org/0009-0007-2482-8493>

Abstract

This paper presents Emotion-to-Motion, a wearable biofeedback system that translates physiological signals into real-time visual and kinetic feedback to support meditation and self-reflection. The experimental setup combines a consumer-grade EEG headband (Muse 2) and a wrist-worn galvanic skin response (GSR) sensor. During five-minute meditation sessions, the system streams brain activity and skin conductance, performs basic signal cleaning, extracts window-level features related to arousal and attentional state, and estimates user condition in real time. A machine learning model classifies each 5-second window into one of five states: calm, focused, excited, stressed, or neutral.

The inferred state and raw signals are displayed through a web-based dashboard that supports live monitoring and session review. In parallel, each detected state is mapped to a short text prompt that drives a depth-map-inspired generative visual pattern. The visual layer is coupled to an Arduino-controlled linear actuator, which moves upward or downward according to the recognized state, closing the loop between sensing, on-screen representation, and physical motion. A pilot study suggests that the end-to-end pipeline can operate reliably over full sessions and that embodied motion cues help participants notice and reflect on fluctuations in their internal state. The paper details the system architecture, the feature-based classification pipeline, and the command-filtering logic used to stabilize actuator behavior for a meditation context, and discusses applications in biofeedback training, interactive art, and emotion-aware installation or spatial prototypes.

Keywords— Biofeedback, electroencephalography (EEG), galvanic skin response (GSR), meditation.

1. INTRODUCTION

Meditation and self-reflection practices are widely used to manage stress, improve emotional awareness, and support well-being. However, traditional meditation relies heavily on subjective introspection, which can be difficult for beginners or for individuals whose attention and emotional state fluctuate during practice. Wearable sensing and physiological computing provide an opportunity to augment meditation with real-time biofeedback, offering an additional channel for noticing internal changes that may otherwise feel ambiguous or hard to describe.

This paper introduces Emotion-to-Motion, a wearable EEG–GSR biofeedback system that links physiological signals to interpretable, real-time feedback in two forms: (1) a web dashboard that visualizes streaming signals and inferred state, and (2) a physical actuator that provides a tangible motion cue. The system is designed as a closed loop: physiological data are collected continuously, processed into short windows, classified into a small set of meditation-relevant states (calm, focused, excited, stressed, neutral), and then translated into synchronized visual and kinetic outputs. The aim is not to claim clinical diagnosis, but to explore how low-cost sensing and embodied feedback can support awareness and reflection during short meditation sessions.

The technical motivation for combining EEG and GSR is that the two modalities capture complementary aspects of experience. EEG reflects changes associated with attention and relaxation dynamics, while GSR provides a sensitive marker of autonomic

arousal. In real-world wearable conditions, both streams are noisy and affected by fit and motion; therefore the system emphasizes practical design choices—windowing, artifact-aware preprocessing, feature-based modeling, and command filtering—to maintain stable feedback without suppressing meaningful transitions.

The contribution of this work is an end-to-end prototype that integrates sensing, classification, visualization, and actuation in a reproducible pipeline, along with a pilot evaluation that examines both system reliability (continuous operation during full sessions) and user perception (whether motion cues help noticing internal fluctuation). Beyond meditation support, the same architecture can be extended to interactive art and emotion-aware installations where physiological signals drive spatial or kinetic responses.

II. BACKGROUND

Electroencephalography (Eeg) Measures Electrical Activity In The Brain Via Electrodes Placed On The Scalp, Capturing Brainwaves Such As Alpha, Beta, Theta, And Delta Bands, Which Correlate With States Of Relaxation, Focus, Or Arousal [1]. Consumer-Grade Devices Like The Muse 2 Headband Make Eeg Accessible Outside Clinical Settings, Offering Dry Electrodes And Bluetooth Connectivity For Real-Time Data Streaming [2].

Galvanic Skin Response (Gsr), Also Known As Electrodermal Activity, Detects Changes In Skin Conductance Due To Sweat Gland Activity, Which Is Modulated By The Autonomic Nervous System And Reflects Emotional Arousal [3]. Wrist-Worn Gsr Sensors Are Noninvasive And Suitable For Wearable Applications.

Emotion models such as the circumplex model describe affect along valence and arousal dimensions, which can be operationalized into application-specific state sets. In this work, we use a five-state representation—calm, focused, excited, stressed, and neutral—chosen to reflect meditation-relevant changes in arousal and attentional engagement.

Biofeedback Involves Monitoring Physiological Signals And Providing Feedback To Help Users Self-Regulate [5]. In Meditation Contexts, It Can Guide Practitioners Toward Calmer States By Visualizing Or Physicalizing Emotional Fluctuations.

III. EXPLORATORY LITERATURE REVIEW

The Integration Of Eeg And Gsr For Emotion Recognition Has Been Extensively Explored In Cognitive Neuroscience And Human-Computer Interaction, With Applications In Mental Health And Biofeedback. A Systematic Review Of Eeg-Based Emotion Recognition Highlights The Use Of Deep Learning Techniques, Such As Convolutional Neural Networks (Cnns) And Recurrent Neural Networks (Rnns), To Classify Emotions From Brain Signals, Achieving Accuracies Over 90% When Combined With Multimodal Data Like Physiological Sensors [6]. Challenges Include Real-Time Processing And Ethical Concerns Around Data Privacy.

Real-Time Eeg Systems Have Been Developed For Neurohumanities, Employing Principal Component Analysis (Pca) And Tree-Based Algorithms To Estimate Valence, Arousal, And Dominance (Vad) Every Few Seconds, With Extra-Trees Classifiers Outperforming Benchmarks At 94% Accuracy [7]. Such Models Can Classify Complex Emotions, extending to over 15 states using VAD mappings. In contrast, our prototype adopts a five-state representation (calm, focused, excited, stressed, neutral) tailored to meditation-relevant changes in arousal and attentional engagement. This choice emphasizes interpretability and stable real-time feedback rather than exhaustive affect coverage.

Empirical Studies Using Low-Cost Noninvasive Sensors, Including Eeg, Electromyography (Emg), And Gsr, Demonstrate Reliable Detection Of High Vs. Low

Valence And Arousal In Cross-Subject Settings, Eliminating The Need For Individual Model Tuning [8]. Gsr Often Complements Eeg By Providing Arousal-Specific Insights, Improving Overall Classification Performance.

Ai-Driven Systems Leveraging Eeg For Mental Well-Being Incorporate Linear Discriminant Analysis (Lda) For Real-Time Emotion Classification, Offering Personalized Interventions To Address Stress And Anxiety Gaps In Traditional Therapy [9]. These Systems Emphasize Scalability For Global Mental Health Applications.

Overviews Of Eeg-Based Biometric Identification Integrate Emotion Recognition To Enhance System Robustness, Using Machine Learning Methods Like Support Vector Machines (Svms) And Long Short-Term Memory Networks (Lstms) On Multimodal Datasets [10]. Emotional Variability In Eeg Signals Is A Key Factor, Suggesting Adaptive Biofeedback For Meditation.

The Fusion Of Artificial Intelligence And Internet Of Things (Aiot) Enables Continuous Emotion Monitoring Via Wearables, Analyzing Signals Like Eeg, Electrocardiography (Ecg), And Gsr With Both Traditional (E.G., Svm, Random Forests) And Deep Learning Approaches [11]. Healthcare Applications Focus On Personalized, Context-Sensitive Systems, Though Challenges Like Data Heterogeneity Persist.

Studies On Social Robots As Mental Health Coaches Use Gsr And Eeg To Detect Valence And Arousal During Stress Management Tasks, Revealing Patterns Of Relaxation Or Excitement That Inform User Experience In Human-Robot Interactions [12]. Physiological Signals Prove Effective For Objective Evaluation In Biofeedback Settings.

These Works Collectively Underscore The Potential Of Eeg-Gsr Fusion For Emotion-Aware Systems But Highlight Gaps In Kinetic Feedback Integration For Meditation, Which Our System Addresses.

IV.METHODOLOGY




FIGURE 2: END-TO-END EEG–GSR CLOSED-LOOP BIOFEEDBACK PIPELINE. SYSTEM OVERVIEW OF EMOTION-TO-MOTION. THE PIPELINE INCLUDES (I) DATA ACQUISITION FROM EEG (MUSE 2) AND GSR, (II) SIGNAL CONDITIONING AND PREPROCESSING (FILTERING, ARTIFACT-AWARE HANDLING, WINDOW SEGMENTATION), (III) FEATURE EXTRACTION AND INFERENCE FOR ESTIMATING USER STATE, (IV) MAPPING AND DECISION LOGIC THAT CONVERTS INFERRED STATES INTO (A) DASHBOARD UPDATES AND (B) ACTUATOR CONTROL COMMANDS, AND (V) REAL-TIME OUTPUTS INCLUDING A MONITORING DASHBOARD AND A LINEAR ACTUATOR. THE CLOSED-LOOP STRUCTURE HIGHLIGHTS HOW PHYSIOLOGICAL SENSING DRIVES SYNCHRONIZED VISUAL AND KINETIC FEEDBACK DURING A MEDITATION SESSION.

Summarizes the full closed-loop mechanism of the proposed “Emotion-to-Motion” system, from physiological sensing to real-time feedback and physical actuation. The key contribution is not only multimodal emotion inference (EEG + GSR), but the translation of inferred affect into two synchronized feedback channels: (1) an interpretable on-screen dashboard that allows users to observe state fluctuations over time, and (2) a tangible kinetic output via a linear actuator. This dual feedback design is intended to support meditation and self-reflection by externalizing otherwise internal and ambiguous emotional changes into legible patterns and movement.

Operationally, the loop begins with synchronized acquisition of EEG (256 Hz) and GSR (10 Hz). These modalities capture complementary aspects of affect: EEG provides information related to cognitive state and attentional/relaxation dynamics, while GSR provides a strong proxy for autonomic arousal. After basic cleaning and segmentation, the system performs window-level inference to estimate one of five states: calm, focused, excited, stressed, or neutral. The resulting label is then mapped into (i) a short text prompt that drives the generative visual layer and (ii) actuator commands that encode a simple “up/down” motion policy. Importantly, (Figure 2 clarifies that the system is designed as a research testbed: each stage (preprocessing, feature extraction, model inference, mapping strategy) can be modified independently without changing the overall architecture, enabling future studies on alternative emotion representations, personalized calibration, and expanded actuation behaviors.




FIGURE 3: HARDWARE AND SOFTWARE INTEGRATION AND REAL-TIME COMMUNICATION: IMPLEMENTATION-LEVEL ARCHITECTURE OF SENSING, COMPUTATION, AND ACTUATION. THE FIGURE DETAILS THE CONNECTIVITY BETWEEN (I) WEARABLE/EEG DEVICES AND A HOST COMPUTER (WIRELESS LINK AND SAMPLING STREAMS), (II) A PROCESSING LAYER THAT PERFORMS BUFFERING, SYNCHRONIZATION, AND INFERENCE, AND (III) A MICROCONTROLLER/ACTUATOR LAYER (E.G., ARDUINO) THAT RECEIVES CONTROL COMMANDS IN REAL TIME. THE DIAGRAM HIGHLIGHTS CRITICAL PRACTICAL ELEMENTS FOR REPRODUCIBILITY: TIMESTAMP ALIGNMENT, MESSAGE FORMAT (CONTROL PAYLOAD), STREAMING CADENCE, AND LATENCY-SENSITIVE PATHS. THE DESIGN SUPPORTS MODULAR SUBSTITUTION OF SENSORS AND ACTUATORS WHILE PRESERVING THE SAME COMMUNICATION PROTOCOL AND CONTROL STRUCTURE.

The implementation architecture required to sustain real-time streaming and actuation during a five-minute meditation session. The integration is organized around three layers: sensing, edge computation, and actuation/feedback. The EEG headband (Muse 2) streams multi-channel signals over Bluetooth, while the GSR wristband provides electrodermal activity samples at a lower rate. Both streams must be buffered and time-aligned before feature computation; even small timing drift can degrade window-level classification when multimodal features are fused. For this reason, the processing node (Raspberry Pi 4 in the current prototype) is treated as the synchronization hub, standardizing timestamps and managing message cadence.

A second critical consideration emphasized by (see Figure 3) is latency budgeting. The system targets continuous operation with predictable update intervals so that the dashboard, visuals, and actuator reflect the same inferred state at approximately the same time. In practice, total delay includes sensor transmission latency, buffering time for the window (5 s), computation time for preprocessing/feature extraction, model inference time, and command transmission to the Arduino. Although the 5-second window introduces an inherent temporal granularity, stable feedback depends on minimizing additional jitter and ensuring that each window produces exactly one state estimate and one actuator decision. This architectural breakdown also supports modularity: additional sensors and actuators can be integrated in future iterations by following the same message protocol, timestamp alignment, and latency constraints.

V. DESIGN COMMAND FILTERING



FIGURE 4: SIGNAL PREPROCESSING AND ARTIFACT HANDLING FOR EEG AND GSR: PREPROCESSING WORKFLOW FOR PHYSIOLOGICAL SIGNALS PRIOR TO MODELING. FOR EEG, THE PIPELINE INCLUDES BANDPASS FILTERING, NOTCH FILTERING FOR POWER-LINE INTERFERENCE, SEGMENTATION INTO FIXED-LENGTH WINDOWS, AND ARTIFACT MITIGATION (E.G., MOVEMENT-RELATED SPIKES, ELECTRODE NOISE) USING REJECTION CRITERIA AND/OR ROBUST STATISTICS. FOR GSR, THE WORKFLOW SEPARATES TONIC AND PHASIC COMPONENTS, REMOVES TRANSIENT OUTLIERS, AND NORMALIZES SIGNALS TO REDUCE INTER-PARTICIPANT VARIABILITY. THE FIGURE EMPHASIZES WHICH STEPS ARE APPLIED PER MODALITY AND WHERE NORMALIZATION OCCURS TO ENSURE STABLE FEATURE DISTRIBUTIONS ACROSS SESSIONS.

The “design command filtering” stage sits between model inference and physical actuation and is introduced to ensure that the system’s outputs remain interpretable, safe, and temporally coherent during a five-minute meditation session. In the prototype, each 5-second window produces one predicted affective label (five states: calm, focused, excited, stressed, and neutral). This label is then mapped to (i) a short text prompt that drives the depth-map inspired visual layer and (ii) an Arduino command that moves the linear actuator up or down. However, raw window-level predictions can fluctuate due to physiological noise, small sensor artifacts, or borderline transitions between emotional states. If actuator commands were executed directly from every predicted label without stabilization, the actuator could “chatter” (rapidly switch directions), which reduces user trust and distracts from the reflective goal of the biofeedback experience.

As illustrated by the pipeline in this section (Emotion Classification → Generative AI to DepthMap → Arduino Command Mapping), filtering is implemented as a decision layer that converts a stream of window-level predictions into actuator commands with bounded variation. The objective is not to hide emotional changes, but to express changes in a way that is legible: movement should indicate meaningful transitions (e.g., a sustained shift from neutral to focused) rather than reacting to minor oscillations around a decision boundary. This is particularly important for wearable EEG and GSR signals recorded in realistic conditions, where drift, contact quality, and movement can introduce brief distortions even when the upstream preprocessing is carefully designed.

Command filtering should be understood as a complement to signal preprocessing. Preprocessing improves the reliability of the inputs to the model, while command filtering ensures that the output behavior remains consistent with the intended interaction. In other words, the system can be scientifically “correct” at the level of classification accuracy, yet still feel unstable or confusing if the physical actuator responds to transient prediction flips. A robust real-time experience therefore requires stabilizing logic on the control side as well.

In practice, filtering can be implemented using lightweight temporal rules that match the 5-second update cadence and respect the actuator’s physical constraints. Common strategies include: (1) majority voting across the last k windows (e.g., 2–3 windows) to reduce single-window misclassifications; (2) hysteresis / dwell time, where a new emotion label must persist for at least n consecutive windows before the actuator changes direction; and (3) rate limiting, which restricts how frequently the actuator is allowed to reverse. These strategies produce motion that is smoother and more semantically meaningful, while still allowing the system to reflect genuine shifts in user state.

The filtering stage also strengthens reproducibility and evaluation because it makes the mapping from physiology to motion explicitly testable. Beyond reporting offline metrics (e.g., confusion matrix, macro-F1, balanced accuracy), the system should document control parameters such as the voting window size, minimum dwell time, and reversal interval. These parameters directly influence prediction stability and perceived

responsiveness, and they allow future researchers to replicate the same behavior under comparable windowing and latency constraints.

Finally, because actuator motion is intended to align with the dashboard visualization and the depth-map rendering, filtering acts as a synchronization contract between software outputs and physical behavior. Each window should produce one consistent state estimate and one motion decision so that the timeline shown on screen corresponds to the same time segment expressed physically. By stabilizing command transitions and enforcing consistent timing, design command filtering helps close the loop from sensing to inference to actuation in a way that supports both human interpretability and engineering robustness. (see Figure 4).

VI. TRAINING PROCESS




FIGURE 5: FEATURE ENGINEERING AND MODEL TRAINING/VALIDATION PROTOCOL: MODEL DEVELOPMENT PIPELINE FROM WINDOWED SIGNALS TO EMOTION/STATE CLASSIFICATION. EACH WINDOW IS TRANSFORMED INTO A FEATURE VECTOR USING TIME-DOMAIN, FREQUENCY-DOMAIN, AND NON-LINEAR DESCRIPTORS (E.G., BAND POWER, STATISTICAL MOMENTS, SPECTRAL ENTROPY, PEAK-RELATED FEATURES FOR GSR). THE FIGURE PRESENTS THE TRAINING PROTOCOL, INCLUDING DATASET PARTITIONING STRATEGY, CROSS-VALIDATION SCHEME, AND HYPERPARAMETER SELECTION. PERFORMANCE REPORTING INCLUDES CLASS-BALANCED METRICS (MACRO-F1 AND

BALANCED ACCURACY) ALONGSIDE CONFUSION MATRICES TO ACCOUNT FOR CLASS IMBALANCE AND TO DIAGNOSE SYSTEMATIC MISCLASSIFICATION PATTERNS.

1. Data Collection, Windowing, and Label Construction (Ground Truth)

Each participant completed a five-minute meditation session while wearing a Muse 2 EEG headband (EEG sampled at 256 Hz) and a wrist-worn GSR sensor (sampled at 10 Hz). Both streams were recorded continuously and stored with timestamps to enable later alignment. To support real-time classification and actuator feedback, the continuous recordings were segmented into fixed-length windows of 5 seconds, yielding a sequence of consecutive samples per session. This window size is consistent with the system's live update cadence, allowing one model decision per window and direct mapping to visual and kinetic output.

The classification target consisted of five meditation-relevant states: calm, focused, excited, stressed, and neutral. Because labeling during meditation is inherently subjective, ground truth was derived from structured self-report after each session. Participants described their perceived state during the meditation, and the report was mapped into the five-state scheme. When participants indicated changes over the session, annotations were aligned to coarse temporal segments (e.g., early/middle/late) and then assigned to the corresponding 5-second windows. Windows that could not be assigned with reasonable confidence were labeled as neutral or excluded to reduce label noise. This labeling strategy prioritizes ecological validity while acknowledging that state boundaries during meditation are often gradual rather than sharply discrete. (see Figure 5)

2. Preprocessing and Feature Engineering (EEG + GSR)

2.1 EEG preprocessing

Raw EEG data were filtered to reduce noise and stabilize feature distributions across windows. A bandpass filter (0.5–50 Hz) was applied to remove slow drift and high-frequency artifacts, and a notch filter was used to suppress power-line interference. The filtered EEG signal was then segmented into 5-second windows. To reduce the impact of sensor artifacts common in consumer-grade wearables (e.g., motion spikes, electrode contact changes), a basic artifact screening step flagged windows with extreme amplitude or abnormal variance; flagged windows were either removed or robustly normalized to prevent model overfitting to noise.

2.2 GSR preprocessing

GSR signals were cleaned for transient spikes and normalized to account for strong inter-individual baseline differences. Because GSR is sensitive to arousal and exhibits slow tonic drift, the preprocessing pipeline optionally separated tonic and phasic components (or used detrending) before windowing. Outliers were clipped or removed, and per-session normalization (e.g., z-score or min-max scaling) was applied so that features computed across windows remained comparable.

2.3 Feature extraction and multimodal fusion

For each 5-second window, features were computed separately for EEG and GSR and then concatenated into a single multimodal feature vector.

- EEG features: spectral descriptors were computed to reflect brainwave dynamics relevant to relaxation and attention. Typical features include band power estimates in delta/theta/alpha/beta ranges and summary statistics of band activity within each window.

- GSR features: window-level statistics (mean, standard deviation, slope), peak-related measures, and response variability indicators were derived to capture arousal fluctuations.
- Feature fusion was performed by concatenation because it is simple, interpretable, and compatible with tree-based classifiers. All features were standardized using parameters computed from the training set only to avoid information leakage. (see Figure 5).

3. Model Training, Validation Strategy, and Performance Reporting

A Random Forest classifier was trained to predict one of the **five states** (calm, focused, excited, stressed, neutral) for each 5-second window. Random Forest was chosen for three reasons: (i) strong performance on small-to-moderate datasets, (ii) robustness to non-linear relationships and noisy features, and (iii) interpretability via feature importance analysis.

3.1 Train-validation split

The dataset was partitioned into training and validation subsets to estimate generalization. If data volume permits, a participant-aware strategy is recommended:

- Within-subject validation: train and validate on different windows from the same participant to measure session consistency.
- Cross-subject validation: leave-one-subject-out or grouped k-fold splitting by participant to test whether the model generalizes to unseen individuals.

This is important because physiological baselines differ substantially across people; cross-subject evaluation avoids overly optimistic results.

3.2 Hyperparameter selection

Key hyperparameters (e.g., number of trees, maximum depth, minimum samples per split) were selected using validation performance. Because the classes may be imbalanced (e.g., many windows labeled neutral), class-weighting or balanced sampling was applied where appropriate to prevent dominance of the majority class.

3.3 Metrics and analysis outputs

Performance was evaluated using metrics that are appropriate for imbalanced multi-class classification, where overall accuracy can be misleading if one class (e.g., neutral) dominates the dataset. Therefore, we report the confusion matrix to explicitly visualize class-wise behavior and identify systematic misclassifications, such as cases where excited and stressed (both higher-arousal states) are frequently confused, or where focused and neutral overlap due to similar low-arousal EEG-GSR patterns. In addition, we compute macro-averaged F1 and/or balanced accuracy, which treat each class more equitably by emphasizing per-class performance rather than being driven by the majority class. Beyond offline classification scores, we also assess real-time prediction stability to ensure the model behaves coherently during continuous streaming: window-to-window consistency (and, if applied, smoothing strategies such as majority voting across adjacent windows) can be quantified to demonstrate that the system avoids rapid, erratic label switching that could disrupt a meditation session. Finally, after training and validation, the model was deployed within the real-time pipeline so that each incoming 5-second window produced an emotion label that immediately triggered (i) an update in the dashboard visualization and (ii) the corresponding actuator control command, thereby closing the loop between sensing, inference, and responsive environmental feedback. (see Figure 5).




FIGURE 6: REAL-TIME DASHBOARD INTERFACE AND CONTROL SIGNALS: SER-FACING MONITORING DASHBOARD AND SYSTEM OUTPUTS IN REAL TIME. THE INTERFACE VISUALIZES MULTI-CHANNEL STREAMS (EEG/GSR), DERIVED FEATURES, INFERRED STATE PROBABILITIES, AND ACTUATOR COMMANDS OVER TIME. PANELS ARE ORGANIZED TO SUPPORT INTERPRETABILITY: RAW SIGNALS (QUALITY CONTROL), PROCESSED INDICATORS (MODEL INPUTS), AND OUTPUTS (STATE ESTIMATES AND CONTROL TRACES). THE TIMELINE VIEW ALLOWS INSPECTION OF LATENCY BETWEEN DETECTION AND ACTUATION, ENABLING DEBUGGING AND ITERATIVE TUNING OF MAPPING FUNCTIONS AND THRESHOLDS DURING PILOT SESSIONS.

The real-time dashboard as an interpretability and quality-control layer, not merely a visualization. In wearable physiological computing, model predictions can only be trusted if the underlying streams are behaving plausibly; thus, the dashboard is designed to reveal signal quality, preprocessing outputs, and classification states simultaneously. A recommended organization is: (i) raw EEG and GSR plots for immediate artifact detection (e.g., sudden electrode dropouts or GSR spikes), (ii) derived indicators such as band power trajectories or arousal proxies that correspond to model inputs, and (iii) the predicted emotion probabilities/labels with timestamps. This structure allows a researcher to diagnose whether a prediction shift is driven by a real physiological change or by a sensor artifact.

In addition to interpretability, (see Figure 6). It is central to **system verification**. By plotting actuator commands alongside predicted labels, the dashboard can be used to verify temporal alignment: the movement should correspond to the same inferred window shown on screen. This matters because user perception of “responsiveness” depends on consistent timing. If feedback is delayed or misaligned, users may misattribute movements to the wrong internal state, weakening the intended reflective function of the system. The dashboard therefore serves a dual purpose: supporting user

engagement (seeing emotions externalized) and enabling rigorous engineering evaluation (latency, stability, uptime) during pilot testing.



FIGURE 7: PILOT EVALUATION RESULTS AND CLASSIFICATION BEHAVIOR: QUANTITATIVE OUTCOMES OF THE PILOT STUDY. THE FIGURE REPORTS MODEL PERFORMANCE ACROSS TARGET CLASSES (E.G., CALM/FOCUSED/EXCITED/STRESSED/NEUTRAL) USING CONFUSION MATRICES AND SUMMARY METRICS. ERROR PATTERNS ARE INTERPRETED IN TERMS OF PHYSIOLOGICAL SIMILARITY BETWEEN CLASSES AND POTENTIAL LABEL NOISE IN SELF-REPORTED GROUND TRUTH. THE RESULTS ARE ADDITIONALLY CONTEXTUALIZED AT THE SYSTEM LEVEL (ROBUSTNESS UNDER REAL-TIME STREAMING) BY REPORTING STABILITY OF PREDICTIONS ACROSS CONSECUTIVE WINDOWS AND SENSITIVITY TO ARTIFACTS.

VII. RESULTS AND DISCUSSION

This section reports findings from the pilot study at two levels: (i) system-level robustness of the end-to-end pipeline and (ii) model-level behavior of the EEG-GSR classifier when deployed in real time. The overall workflow follows the architecture presented earlier: wearable sensing and synchronization (Figures 2–3), signal cleaning and windowing (Figure 4), model inference and feature-based classification (Figure 5), and synchronized feedback through the dashboard and actuator (Figure 6). The objective of evaluation is not only to show that classification is possible, but to demonstrate that the complete loop can run continuously, produce interpretable outputs, and remain stable enough for a meditation setting.

At the system level, the prototype achieved 95% uptime across five-minute sessions, indicating that continuous streaming, buffering, inference, dashboard updating, and Arduino-based actuation can be maintained with only minimal interruption. This matters because the user experience depends on temporal continuity: a short dropout can break the alignment between the on-screen visualization and the physical actuator response, and it can also invalidate a subset of windows used for inference. The reported

uptime therefore provides evidence that the hardware–software integration and streaming schedule described in the methods are practical under pilot conditions.

At the model level, the classifier produced window-level predictions over five target states: calm, focused, excited, stressed, and neutral. The reported overall accuracy (e.g., 82% in the pilot summary) should be interpreted cautiously because the dataset can be imbalanced and because ground-truth labeling in meditation contexts is inherently challenging. For this reason, the analysis emphasizes class-wise behavior using a confusion matrix and class-balanced metrics (macro-F1 and/or balanced accuracy), rather than relying on accuracy alone. Confusion patterns are expected to occur between states with similar arousal or attentional signatures (for example, focused vs. neutral, or excited vs. stressed depending on how arousal is expressed in EEG–GSR features). These confusions are informative: they reveal which states are reliably separable with the current sensing setup and which states require either richer features, longer context, or improved labeling resolution.

A central contribution of the system is the alignment of inference outputs with two synchronized feedback channels: the dashboard and the actuator. The dashboard (Figure 6) provides interpretability by exposing signal trends and prediction trajectories over time, allowing both users and researchers to inspect whether changes in predicted state correspond to plausible changes in EEG–GSR patterns. In parallel, the actuator translates predicted state into a tangible motion cue. Pilot feedback indicates that many participants perceived the actuator as helpful for noticing internal fluctuations during meditation (e.g., 80% reporting improved awareness). In this sense, the system is evaluated not only as a classifier, but as a biofeedback interface that externalizes internal variation into a legible, embodied signal.

The design command filtering stage is essential for translating window-level predictions into stable physical behavior. Because predictions are generated every 5 seconds, small fluctuations in features or borderline decisions can cause rapid switching between adjacent states if commands are sent directly. Filtering methods—such as enforcing a minimum dwell time, majority voting across consecutive windows, and rate limiting for direction changes—reduce “command chatter” and ensure that actuator motion reflects sustained trends rather than transient classification noise. This design choice directly links the modeling layer (Figure 5) to the interaction layer (Figure 6): it preserves responsiveness while maintaining the calm, non-distracting character required for meditation.

Several limitations define the scope of current claims. First, sample size is small and label assignment may be temporally coarse when based on post-session reporting, which can blur boundaries between similar states (e.g., neutral and focused). Second, consumer-grade wearable signals can vary with fit, contact quality, and motion, which affects both feature stability and real-time reliability even after preprocessing. Third, generalization across individuals remains a key challenge; future evaluations should explicitly specify participant-aware splitting and report class-balanced metrics to avoid optimistic estimates driven by within-subject similarity. These limitations motivate the outlook in the next section, particularly around calibration and more rigorous validation protocols.

VIII. CONCLUSION AND OUTLOOK



FIGURE 8: DEPLOYMENT ROADMAP AND SCALABILITY OF PARTICIPATORY NEURO-RESPONSIVE ENVIRONMENTS: PROPOSED PATHWAY FROM PILOT PROTOTYPE TO ROOM-SCALE DEPLOYMENT. THE ROADMAP OUTLINES (I) SCALING SENSING FROM SINGLE-USER TO MULTI-USER PARTICIPATION, (II) AGGREGATION STRATEGIES (E.G., WEIGHTED VOTING, ROBUST POOLING OF PHYSIOLOGICAL INDICATORS), (III) ADAPTATION MECHANISMS (PERSONAL CALIBRATION AND DOMAIN SHIFT HANDLING), AND (IV) EVALUATION CHECKPOINTS INCLUDING USER STUDY DESIGN, LONGITUDINAL MONITORING, AND ETHICS CONSIDERATIONS (PRIVACY, CONSENT, AND DATA MINIMIZATION). THE FIGURE CLARIFIES HOW THE PROTOTYPE CAN EVOLVE INTO A RELIABLE RESEARCH PLATFORM FOR PARTICIPATORY URBAN/ARCHITECTURAL INTERVENTIONS.

This paper presented Emotion-to-Motion, a wearable EEG–GSR biofeedback system that converts physiological signals into synchronized visual and kinetic feedback for meditation and self-reflection. The system implements a complete closed-loop workflow: streaming acquisition and synchronization (Figures 2–3), preprocessing and artifact-aware windowing (Figure 4), feature-based classification using a Random Forest model (Figure 5), and an interpretable dashboard paired with actuator output (Figure 6). By keeping each stage modular, the platform supports iterative improvement—preprocessing, feature design, model choice, or actuation mapping can be upgraded without redesigning the entire pipeline.

The pilot study demonstrates feasibility at three linked levels. At the engineering level, the pipeline sustained real-time operation with high reliability (e.g., 95% uptime), showing that continuous sensing, inference, visualization, and actuation can be deployed in a stable session format. At the model level, the classifier produced window-level predictions across five states—calm, focused, excited, stressed, and neutral—providing an initial basis for mapping multimodal physiological patterns to interpretable affective/attentional categories under wearable constraints. At the interaction level, participant feedback suggests that actuator motion can help make internal shifts more noticeable, supporting the intended reflective function of the system.

A key takeaway is that performance is not determined solely by classification scores. The system's usefulness depends on whether predictions can be translated into feedback that is stable, interpretable, and aligned in time. For this reason, the design command filtering layer is a core methodological contribution: it bridges modeling and embodiment by stabilizing actuator behavior while preserving responsiveness. This ensures that motion cues communicate meaningful trends rather than reacting to short-term noise or borderline decisions.

Looking ahead, three directions follow directly from the current method. First, the evaluation protocol should be strengthened by explicitly specifying participant-aware train/validation splits and reporting confusion matrices alongside class-balanced metrics, ensuring that the reported performance reflects generalization rather than within-subject similarity. Second, robustness can be improved through calibration strategies and adaptation to individual baselines, which is especially important for GSR amplitude variability and wearable EEG sensitivity to fit and motion. Third, the actuation vocabulary can expand beyond a single linear actuator toward multi-channel environmental feedback (light, sound, soft robotics), enabling the same inference-and-filtering framework to support larger-scale responsive installations.

Overall, Emotion-to-Motion provides a practical testbed for studying how physiological computing can be integrated into interactive environments. By combining multimodal sensing, interpretable visualization, and embodied motion within a reproducible closed-loop architecture, the work lays a foundation for future research in meditation support, interactive art, and emotion-aware spatial or architectural experiences.

IX. FUTURE WORK

Future work will extend the current prototype from a single-user meditation device toward an architectural-scale, multi-user system. On the hardware side, we plan to integrate additional biosensing modalities (e.g., heart rate variability and respiration) and multiple actuators, including light, sound, and soft robotic or kinetic elements. This will enable richer mappings between emotional states and spatial responses, moving beyond a single linear actuator toward room-scale cues such as changing illumination, ambient soundscapes, and responsive surfaces.

On the software side, the emotion recognition pipeline will be upgraded to support adaptive and personalized models. Deep learning-based classifiers and transfer learning will be explored to better handle inter-individual variability and noisy, in-the-wild data. Multi-user sessions will be supported by a web-based dashboard that can visualise collective affective states over time and coordinate responses across a network of actuators. This will allow experiments in shared meditation or group reflection, where the space itself reacts to aggregated emotional patterns.

These developments are part of a broader research agenda on Symbiotic Reality, in which built environments and human occupants continuously co-adapt through sensing, computation, and feedback. In future doctoral work in computation in architecture, the author aims to develop a hardware-software toolkit that allows architects to design, prototype, and evaluate emotionally responsive spaces. The Emotion-to-Motion system presented here provides an initial testbed for this agenda, demonstrating how low-cost wearable sensing and simple actuation can be combined into a reproducible workflow that can later be embedded into architectural installations, public spaces, and everyday environments.

ACKNOWLEDGMENTS

The author thanks the Participatory Citizen Lab for support and participants in the pilot test.

REFERENCES

- [1] Y. Zhao, Y. Jiang, X. Liu, Q. Xu, L. Zhao, and Y. Peng, "Workload electroencephalogram characteristics in cognitive tasks based on intra- and inter-hemispheric asymmetry," *IEEE Access*, vol. 11, pp. 42377–42388, 2023.
- [2] O. E. Krigolson, C. Williams, A. Norton, C. Hassall, L. J. Colino, and K. R. Krigolson, "Evaluating the feasibility of a consumer-grade wearable EEG headband for assessing meditation and mindfulness in a real-world head-to-head comparison with established laboratory equipment," *Front. Neurosci.*, vol. 15, p. 634147, 2021.
- [3] I. Daly, D. Nicolaou, S. J. Nasuto, and K. Warwick, "A review, current challenges, and future possibilities on emotion recognition using machine learning and physiological signals," *IEEE Access*, vol. 7, pp. 140990–141020, 2019.
- [4] J. A. Russell, "A circumplex model of affect," *J. Pers. Soc. Psychol.*, vol. 39, no. 6, pp. 1161–1178, Dec. 1980.
- [5] P. Remis-Serna, M. A. Ramírez-Moreno, J. de J. Lozoya-Santos, M. Cebral-Loureda, and C. S. Vélez-Saboyá, "A systematic review of technology-aided stress management systems: Automatic measurement, detection and control," *IEEE Access*, vol. 11, pp. 106352–106374, 2023.
- [6] E. Gkintoni, A. Aroutzidis, H. Antonopoulou, and C. Halkiopoulos, "From neural networks to emotional networks: A systematic review of EEG-based emotion recognition in cognitive neuroscience and real-world applications," *Brain Sci.*, vol. 15, no. 2, p. 186, 2025.
- [7] M. A. Blanco-Ríos, M. O. Candela-Leal, C. Orozco-Romo, P. Remis-Serna, C. S. Vélez-Saboyá, J. de Jesús Lozoya-Santos, M. Cebral-Loureda, and M. A. Ramírez-Moreno, "Real-time EEG-based emotion recognition for neurohumanities: Perspectives from principal component analysis and tree-based algorithms," *Front. Hum. Neurosci.*, vol. 18, Art. no. 1319574, 2024.
- [8] D. Girardi, F. Lanubile, and N. Novielli, "Emotion detection using noninvasive low cost sensors," in *Proc. 7th Int. Conf. Affective Comput. Intell. Interact. (ACII)*, San Antonio, TX, USA, 2017, pp. 125–130.
- [9] J. M. Gavhane, G. Shirke, S. Modak, J. Prasad, and S. Raman, "AI-driven emotion recognition meets EEG for next-generation mental well-being," 2025. [Online]. Available: <https://www.researchgate.net/publication/389753730>
- [10] M. A. Becerra, C. Duque-Mejia, A. Castro-Ospina, L. Serna-Guarín, C. Mejía, and E. Duque-Grisales, "EEG-based biometric identification and emotion recognition: An overview," *Computers*, vol. 14, no. 8, p. 299, 2025.
- [11] M. J. A. Awan, O. Mazhar, O. Mazhar, A. Y. Zappa, and R. Masood, "AIoT for human emotion recognition: Potentials, challenges, and healthcare applications," *Comput. Sci. Rev.*, vol. 52, Art. no. 100683, 2024.
- [12] K. Kłęczek, A. Rice, and M. Alimardani, "Robots as mental health coaches: A study of emotional responses to technology-assisted stress management tasks using physiological signals," *Sensors*, vol. 24, no. 13, p. 4301, 2024.