

## Oscillating Mindfully: Using Machine Learning to Characterize Systems-Level Electrophysiological Activity During Focused Attention Meditation

Noga Aviad, Oz Moskvich, Ophir Orenstein, Etam Bengier, Arnaud Delorme, and Amit Bernstein

### ABSTRACT

**BACKGROUND:** There has been rapid growth of neuroelectrophysiological studies that aspire to uncover the “black box” of mindfulness and meditation. Reliance on traditional data analysis methods hinders understanding of the complex, nonlinear, multidimensional, and systemic nature of the functional neuroelectrophysiology of meditation states.

**METHODS:** Thus, to reveal the complex systemic neuroelectrophysiology of meditation, we applied a machine learning extreme gradient boosting classification algorithm and 4 complementary feature importance methods to extract systemic electroencephalography features characterizing mindful states from electroencephalography recorded during a focused attention meditation and a control mind-wandering state among 26 experienced meditators.

**RESULTS:** The algorithm classified meditation versus mind-wandering states with 83% accuracy, with an area under the receiver operating characteristic curve of 79% and F1 score of 74%. Feature importance techniques identified 10 electroencephalography features associated with increased power and coherence of high-frequency oscillations during focused attention meditation relative to an instructed mind-wandering state.

**CONCLUSIONS:** The findings help delineate the complex systemic oscillatory activity that characterizes meditation.

<https://doi.org/10.1016/j.bpsgos.2024.100423>

Given exponential progress over the past 2 decades in mindfulness research and its applications, there is fast-growing scientific interest in illuminating the “black box” of mindfulness and meditation practices and states more broadly (1).<sup>1</sup> Central to these efforts is research that aspires to uncover the functional neurophysiological substrate or correlates of altered consciousness during mindfulness and related meditation practices (2–6) and, most notably, the study of functional brain activity in real-time during mindfulness and related meditation practices and states (4,5,7–9). The high temporal resolution of electroencephalography (EEG) and its capacity to capture the real-time dynamics of brain activity and large-scale

synchronization of neural networks noninvasively has thus been instrumental (10). However, despite vast research into oscillatory activity during meditation, understanding of the complex functional electrophysiology of meditation practices and states is still limited (2,11).

A central limitation of existing research is that most EEG studies of mindfulness and meditation have largely focused on unidimensional, individual, typically spectral correlates within the EEG signal (12). This traditional approach to EEG data is not designed to empirically elucidate nonlinear, multidimensional, and complex systems-level endogenous electrophysiological activity that likely characterizes complex brain states during meditation (2,12). Furthermore, scholars have argued that variability in data analytic and experimental methods have limited comparison, integration, and generalization of findings across studies, which further hinders an integrative and complex understanding (5,12). This, in turn, translates into a further restricted basis for making informed choices regarding design elements and parameters (e.g., specific EEG features of interest) in future empirical studies.

Machine learning (ML) algorithms represent a fast-growing method to advance systems-level understanding and pattern recognition from complex data (13). ML modeling enables discovery of systems-level patterns in large amounts of high-

<sup>1</sup>The definition of mindfulness has been the subject of extensive and ongoing debate among Buddhist scholars as well as scientists. We use the term mindfulness to refer broadly to a family of mental state(s) and related meditation practices characterized by present-moment attention and awareness, nondistractedness, and a particular set of salutary attitudes toward experience (e.g., patience, acceptance, curiosity, nonjudging) (54–59). Likewise, meditation refers to a broader class of contemplative practices grounded in Buddhist as well as a range of other wisdom traditions that includes a number of families or types of practices and targets (e.g., attentional, constructive, and deconstructive families).

dimensional data with limited assumptions regarding the components of interest in the data (14). Accordingly, ML has been applied recently to EEG data (15), including preliminary efforts to classify meditation states (12,15–19). These initial studies have sought to test the potential of ML classification techniques to detect and classify meditation versus other mental states (e.g., rest with open or closed eyes, reading, listening) based on electrophysiological activity. Overall, these initial studies have demonstrated moderate to excellent classification accuracy (78%–90%), suggesting that meditative states can be distinguished from other mental states based on electrophysiological activity.

However, ML studies of meditation have yet to apply emerging post hoc explainability computational methods to empirically extract the systems-level electrophysiological features of meditation states. These methods are critical not only to test whether the EEG signal distinguishes brain activity during meditation from other mental states (e.g., mind wandering) but also to empirically characterize the functional neural architecture of that systems-level oscillatory activity within the signal that guides such differentiation (20). Accordingly, computational explainability and feature importance methods are critical for deriving informative insights from the model and improving understanding of its predictions so as to empirically characterize systems-level electrophysiology of meditation states (20,21), i.e., for seeing into the black box of meditation.

Therefore, in the current study, we aimed to advance an integrative systems-level understanding of electrophysiology during focused attention (FA) meditation (22). Relative to open monitoring and (de)constructive meditation practices, FA meditation practices are common to mindfulness as well as a number of related meditation practices and entail a narrower set of practices that may be subserved by a more unified pattern of electrophysiological activity (23). To do so, we sought to first apply ML algorithms and modeling for classification of 2 experimental states—FA meditation and mind wandering—among advanced meditators. Second, we sought to apply multiple techniques of model explainability and feature importance to extract ML classification features that most strongly distinguished between meditation and mind-wandering states. We conducted these analyses on open-access EEG data of experienced meditators wherein EEG was recorded in a mixed between-within design over 2 blocks, one of FA meditation and the other of a control comparison state of mind wandering (11).

## METHODS AND MATERIALS

### Participants

The original parent study (11) sampled meditators from 3 distinct meditative traditions as well as control participants, who were included in the study based on age, gender, and years of meditation practice. To develop an ML model capable of classifying FA meditation specifically, we used a subset of the original dataset, including samples from Vipassana and Himalayan yoga tradition practitioners only ( $n = 26$ ). These 2 traditions involve sustained FA on a specific object of attention. Vipassana meditators practiced the Vipassana meditation techniques per S.N. Goenka (24). Although the Himalayan yoga tradition involves a variety of practices, here participants were

instructed to engage in an FA practice that entailed focus on a mantra with or without awareness of the breath (see the [Supplement](#) for more details). The parent project was approved by the local Meditation Research Institute Indian ethical committee and the ethical committee of the University of California San Diego (Institutional Review Board project # 090731).

### Procedure

The parent study consisted of two 20-minute sessions, a meditation condition and an instructed mind-wandering control condition, in a within-subject design. The 2 blocks were counterbalanced to prevent order effects. In the meditation block, the first 10 minutes constituted a preparatory breath focus period during which all participants were asked to focus on their breath regardless of their particular FA meditation practice. During the final 10 minutes of the meditation block, participants were instructed to practice their regular formal FA meditation practices. In the instructed mind-wandering block, participants were instructed to remember autobiographical events from childhood to the most recent past and were explicitly told to avoid remembering emotionally charged events. See the [Supplement](#) for details regarding the 64 + 8 channel Biosemi Active-Two amplifier system and a 10–20 Headcap standard 64-channel cap used for data collection.

### Dataset and Models

**Data Selection.** We chose, a priori, to analyze data recorded from 9 midline-frontal electrodes: Fp1, Fpz, Fp2, AF3, AFz, AF4, F1, Fz, F2 (see the [Supplement](#) for a detailed rationale for electrode selection). The final dataset included 26 experienced meditators ( $n = 12$  Himalayan yoga tradition,  $n = 14$  Vipassana practitioners). The entire 10-minute meditation period and the parallel 10-minute instructed mind-wandering period analyzed in the parent study were used for classification. The 10-minute breath focus period was used for calibration (see [Calibration](#) below).

**Data Preparation for Classification.** Preprocessing of EEG data included referencing to the right mastoid, down-sampling from 1024 Hz to 256 Hz, and noise filtering and artifact rejection as detailed in the parent study (11). In addition, each recording was split into 10-second non-overlapping epochs, 2560 samples each (sample rate = 256 Hz), and discontinuous epochs were removed (see the [Supplement](#) for additional details). Finally, data from all 9 electrodes were combined into an array of shapes (2560, 9) for each epoch. Each epoch was labeled 0 or 1, representing meditation and nonmeditation states, respectively, based on the condition during which the epoch was recorded (meditation or mind wandering). Epochs were divided into training and testing sets with an 80%:20% ratio in a leave- $n$ -subjects-out manner, meaning that epochs belonging to the same participant were only used in a single set. Weights were used to account for imbalances in data between labels/conditions in calculation of the loss function. Per the within-subject design, data were centered around the mean and standardized across participants and conditions.

**Performance Evaluation.** The performance of the classification model was evaluated using the following measures: 1) classification accuracy, the most straightforward evaluation metric of ML models, is defined as the ratio of correct predictions out of the total predictions (accuracy = number of correct predictions / total number of predictions  $\times$  100%); 2) area under the receiver operating characteristic curve (AUROC) plotted using the true positive rate (i.e., model sensitivity, true positive rate = true positives / [true positives + false negatives]) as a function of false positive rate (i.e., probability of false alarm, false positive rate = false positives / [false positives + true negatives]) (see Table 1); and 3) F1 score, the harmonic mean of precision (i.e., the ratio of true positives to all positive predictions, true positives / [true positives + false positives]) and recall, another term for sensitivity or true positive rate (F1 score =  $2 \times [\text{precision} \times \text{recall}] / [\text{precision} + \text{recall}]$ ) (see Table 1). AUROC and F1 scores were calculated in addition to the more intuitive measure of accuracy because both are considered more reliable, in particular for cases of imbalanced datasets (i.e., when the number of examples in one class greatly outnumbers the examples in another class), such as the current dataset (25). We performed 50 iterations of random subsampling cross-validation in a leave-*n*-subjects-out manner such that participants were randomly split into training and testing sets, and each participant's epochs appeared on a single set only. The partition into train (80%) and test (20%) sets was repeated randomly at each of the 50 iterations; because there were 26 participants in total, 21 random participants were used for training, and 5 were used for testing/validation at each iteration. Scores were calculated on the testing set and averaged across the 50 runs of the model.

**Calibration.** We added several epochs to the training set as calibration epochs for each test participant. The epochs used for calibration included the first 30% of epochs recorded during mind wandering (approximately 1.85 minutes), as well as the first 70% of epochs recorded during breath focus (approximately 4.2 minutes). The specific number of epochs used for calibration was selected based on experimental analysis to determine the minimal number of epochs needed to achieve maximum performance (see the Supplement for experimental analysis of number of epochs for calibration). The need for such calibration stems from the well-documented challenge of high interparticipant variability in EEG data, which causes particular concern when training the model using leave-*n*-subjects-out cross-validation (26). This method of cross-validation, which is required to properly evaluate model performance on unseen participants, often leads to lower performance than the scenario where intraparticipant data is used both in training set and test set, making it more applicable to real-life and clinical scenarios, when a model would be

used in medical applications suitable for unseen participants (13). This is because using epochs of the same participant for more than a single set (e.g., some of the participant's epochs are in the training set and other epochs of the same participant are in the test set) likely leads to data leakage (27), when the model learns the characteristics of each participant in addition to the features used for classification, thus overestimating the model's performance and underestimating its classification error (28). Therefore, we used epochs recorded during breath focus as well as the first epochs recorded during mind wandering to quantify a baseline (premeditation) electrophysiological activity benchmark for each test participant (26).

**Classifier Configuration and Hyperparameters.** We sought a classification model that performs well on structured data (i.e., formatted, ordered, tabular data) and provides estimates of the importance of each electrophysiological feature during meditation from the trained predictive model. For these reasons, we chose to use XGBoost (29), a framework for decision tree ensemble models wherein predictions of multiple trees are aggregated. This framework allows for relatively straightforward retrieval of feature importance scores, calculated by the average gain in accuracy across all splits in the trees where each feature was used, and thereby interpret the model's predictions.

XGBoost models allow tuning of many hyperparameters to maximize model performance and prevent overfitting. The optimal hyperparameter values for the model were selected based on grid search results and include max\_depth, which controls the maximum depth of the tree: 10; learning\_rate, which controls the rate at which the model learns patterns in data: 0.2; subsample, which controls the number of samples supplied to a tree: 0.8; colsample\_bytree, which controls the number of features supplied to a tree: 1; min\_child\_weight, which blocks potential feature interactions to prevent overfitting: 5; and gamma, a regularization parameter to prevent overfitting: 2.

**Feature Importance.** The classifier was trained on 225 univariate features. These features were selected based on previous findings in the literature describing electrophysiology of meditation, feasibility of the computation, and the capability to directly compare and contrast results to findings reported in the majority of existing EEG studies of meditation (12). They include the power of delta (1–4 Hz), theta (4–8 Hz), alpha (8–12 Hz), beta (12–30 Hz), and gamma (30–150 Hz) frequency bands in each of the 9 electrodes, as well as the computed coherence between all combinations of 2 electrodes of the 9 and each of the 5 frequency bands. Features were calculated, as reported in the parent study to facilitate comparability, using the EEGLAB functions of Newcrossf for coherence and Spectopo for spectral power.

The features that contributed most to the model's predictions were extracted using 3 methods. First, important features were extracted using XGBoost's Gain method (average accuracy gain across all splits in which the features were used). Second, SHapley Additive exPlanations (SHAP) library (30)—a game-theoretical approach to explain the decisions made by any model—was used to validate the results, as well as interpret each feature's pattern of contribution

**Table 1. Classification Metrics**

	Actual (Positive)	Actual (Negative)	Positive Predictive Value
Predicted (Positive)	True positives	False positives	Precision
Predicted (Negative)	False negatives	True negatives	
Sensitivity	Recall		

independently. Finally, permutation importance functionality provided by Scikit-learn library (31) was used to further evaluate each feature's importance based on both the AUROC score and F1 score. In permutation importance, a single feature is permuted so that any potential correlation between the feature and label is lost. Then, model baseline performance score is compared with the performance score obtained by the model on the same dataset with the permuted feature to infer how important the feature is for classification. Using this approach, the 10 strongest features were extracted using 4 complementary methods: XGBoost feature importance, SHAP, permutation importance based on the AUROC score, and permutation importance based on the F1 score. Among the 10 strongest features, each feature was assigned a score representing its relative importance in the model, such that the most important feature received a score of 10, the next most important received a 9, and so on. This procedure was then repeated over 10 iterations of random subsampling cross-validation to attain a generalized score and avoid biases caused by specific test sets. Finally, scores of all 10 runs from all 4 methods were summed. The average of the 4 methods was computed as well, but using an average score rather than a summed score did not lead to a change in the features identified as most important for classification. Like other decisions in this study, the selection of 4, and these 4 methods in particular, is arbitrary. However, the methods were selected a priori to help ensure that feature selection was grounded both within and across permutations of participants as well as independent orthogonal feature importance methods based on fundamentally different mathematical assumptions. By doing this, we attempted to strike a balance between robustness and predictive power in feature selection (see Table 2).

## RESULTS

### Model Performance

After training, over 50 iterations of random subsampling cross-validation, the XGBoost classifier achieved mean accuracy = 83%, AUROC = 79%, and F1 score = 74% in classifying FA meditation versus mind-wandering control states based on the 225 univariate features (see the Supplement for more details).

These results provide strong evidence that these 2 experimental states were characterized by distinct electrophysiological activity

### Feature Importance

The most important features for the model's predictions were extracted by means of XGBoost, SHAP, permutation importance based on AUROC score, and permutation importance based on F1 score separately during each iteration of a 10-fold cross-validation (see Figure 1). Three features received particularly highly weighted scores, which are indicative of their overall relative importance for the model's predictions (see Table 2): 1) beta coherence between electrodes Fz and F2, such that high coherence was indicative of the meditation state, and low coherence was indicative of the nonmeditation state (see Figure 2); 2) delta power in electrode AF4, such that low power was indicative of the meditation state, and high power was indicative of the nonmeditation state (see Figure 3); and, 3) gamma coherence between electrodes Fpz and AFz, such that high coherence was indicative of the meditation state, and low coherence was indicative of the nonmeditation state (see Figure 4). Together, these key electrophysiological features point to an overall pattern of highly synchronized, high-frequency oscillations as well as reduced low delta oscillations during FA meditation in frontal and midline areas (e.g., the dorsolateral prefrontal cortex, anterior cingulate cortex).

Finally, to test and verify the relative importance of these features for classification, the 3 primary features were removed from the dataset entirely. Model performance without the 3 primary features dropped to AUROC = 72% and F1 score = 66%. These results support the suggested importance of the features for classification but demonstrate that 3 individual features were not sufficient to characterize the electrophysiological activity during FA meditation relative to instructed mind wandering.

In addition to the 3 most central features, 7 secondary features were found to impact the model's predictions, although less strongly. Their weighted importance scores were lower than those of the 3 most dominant features (see Table 2): 1) gamma power in electrode AF3, such that high power was

**Table 2. Weighted Scores of Primary and Secondary Features**

Feature	XGBoost	SHAP	Permutation: AUROC	Permutation: F1	Average
Fz-F2 Beta Coherence	33	33	39	30	33.75 <sup>a</sup>
AF4 Delta Power	18	47	29	31	31.25 <sup>a</sup>
Fpz-AFz Gamma Coherence	22	40	19	25	26.5 <sup>a</sup>
AF3 Gamma Power	10	22	25	29	21.5
Fz-F2 Gamma Coherence	47	30	4	1	20.5
Fp1-AF3 Gamma Coherence	41	14	10	13	19.5
AF4-F2 Theta Coherence	38	14	7	9	17
AF4 Gamma Power	5	4	29	24	15.5
Fpz Gamma Power	5	18	16	16	13.75
F1-F2 Gamma Coherence	8	13	9	13	10.75


The weighted score of each feature was assigned based on 10-fold cross-validation, as well as the summed and averaged weighted scores, used to identify the most important features for classification.

AUROC, area under the receiver operating characteristic curve; SHAP, SHapley Additive exPlanations.

<sup>a</sup>These 3 features received particularly high weighted scores, indicative of their overall relative importance for the model's predictions.




## Oscillating Mindfully




**Figure 1.** Top 10 important features generated by XGBoost and SHapley Additive exPlanations (SHAP). Both figures were generated during the same run on the same training set. A direct comparison of the two is not straightforward because they were generated using 2 different approaches. Despite that, they include mostly the same features. The image on the left was generated by XGBoost and provides an ordered list of top 10 features and their calculated importance (F) scores, such that gamma coherence between Fz-F2 electrodes contributed the most to the model's predictions, etc. The image on the right was generated by SHAP and provides an ordered list of top 10 features as well as other useful information, such as the distribution of different values for each feature (represented by color, such that pink represents high values, and blue represents low values) and the distribution of the feature's impact on predictions (represented by the x-axis, such that negative values represent a mindful state prediction, and positive values represent a nonmindful state prediction). For example, high (pink) beta coherence between Fz-F2 electrodes indicates the meditation state (negative x value), whereas low (blue) coherence indicates the nonmeditation state (positive x value). Note that the low values are more indicative of the nonmeditation state than the high values that are indicative of the meditation state, as demonstrated by the length of the tail to the right or the absolute value of x.

indicative of the meditation state, and low power was indicative of the nonmeditation state; 2) gamma coherence between electrodes Fz and F2, such that high coherence was indicative of the meditation state, and low coherence was indicative of the nonmeditation state; 3) gamma coherence between electrodes Fp1 and AF3, such that high coherence was indicative of the meditation state, and low coherence was indicative of the nonmeditation state; 4) theta coherence between electrodes AF4 and F2, such that high coherence was indicative of the meditation state, and low coherence was indicative of the nonmeditation state; 5) gamma power in electrode AF4, such


that high power was indicative of the meditation state, and low power was indicative of the nonmeditation state; 6) gamma power in electrode Fpz, such that high power was indicative of the meditation state, and low power was indicative of the nonmeditation state; and 7) gamma coherence between electrodes F1 and F2, such that high coherence was indicative of



**Figure 2.** Pattern of effect for beta coherence in Fz-F2 electrodes. x Value represents level of beta coherence in electrodes Fz-F2, and y value represents impact on model predictions, such that negative values represent a mindful state prediction, and positive values represent a nonmindful state prediction. The greater the absolute value of y, the higher the estimated probability of the prediction is. Note that coherence needs to be very high to indicate mindful state prediction, whereas lower values indicate nonmindful state prediction.



**Figure 3.** Pattern of effect for delta power in electrode AF4. The x-axis represents delta power in electrode AF4, and the y-axis represents impact on model predictions, such that negative values represent a meditation state prediction, and positive values represent a nonmeditation state prediction. The greater the absolute value of y, the higher the estimated probability of the prediction is. Note that when delta power in AF4 electrode is very low, all data points are assigned negative y values, indicating a meditation state prediction. Conversely, when power is relatively high, all data points have a positive y value, indicating a nonmeditation state prediction.



**Figure 4.** Pattern of effect for gamma coherence in Fpz-AFz electrodes. The x value represents the level of gamma coherence in electrodes Fpz-AFz, and the y value represents impact on model predictions. Thus, negative values represent a meditation state prediction, and positive values represent a nonmeditation state prediction. The greater the absolute value of y, the higher the estimated probability of the prediction. Note that low coherence is indicative of a nonmeditation state.

the meditation state, and low coherence was indicative of the nonmeditation state.

### Dimensionality Reduction Using Feature Selection

Finally, to further quantify and validate the relative importance of the primary and secondary features observed, a new dataset was created. This new dataset contained only the 10 features found to be most important, excluding all other 214 features. This was done to investigate whether the identified minimal set of features was sufficiently quantitatively comprehensive to account for the electrophysiology of the meditation state (32). Model performance on the new dataset achieved averaged scores of AUROC = 78% and F1 score = 72%, which are relatively close to the model classification values achieved with the entire dataset of 225 features. These secondary results thereby further validate the results achieved by summing feature importance scores obtained via orthogonal methods of XGBoost (Gain), SHAP, and permutation importance model explainability and derived characterization of the model's feature importance (see Figure 1). More specifically, these findings demonstrate that the identified set of features sufficiently characterize the key electrophysiological patterns that distinguish between FA meditation and mind-wandering states. Consequently, these results also indicate that the rest of the features—despite constituting the very large majority of the data—are either redundant or not important for classification purposes or for the characterization of functional electrophysiology during FA meditation relative to mind-wandering.

### DISCUSSION

We applied advances in ML modeling for classification and pattern recognition to identify systems-levels oscillatory

patterns that characterize electrophysiological brain activity during FA meditation relative to instructed mind wandering (11). We used the XGBoost decision trees ensemble ML classification algorithm to learn, with no prior assumptions, the electrophysiological patterns that differentiate between these states and then calculated XGBoost's Gain feature importance, SHAP, permutation importance based on the AUROC score and permutation importance based on F1 score methods to evaluate the relative contribution of features to the model's classification.

First, using 225 calculated EEG features, ML was able to classify experimental meditation versus mind-wandering states with 83% accuracy, with an AUROC score of 79% and an F1 score of 74%. These findings provide strong evidence that using ML, FA meditation can be robustly differentiated from an instructed mind-wandering state based on electrophysiological activity alone. Second, feature importance analyses revealed that 3 primary features and 7 secondary features largely accounted for the ML model's predictions (see Table 2). These features include, first, an increase in power and coherence of high-frequency oscillations during meditation states relative to mind wandering; second, decreased delta activity during meditation in AF4 electrode; and lastly, increased theta coherence during FA meditation relative to mind wandering. Finally, we found that the identified set of features achieved almost identical ML model prediction to the original ML model of all 225 features. Therefore, the findings indicate that the identified features distinguished between an FA meditation state and a control mind-wandering state and that over 200 other features were redundant or not similarly predictive of electrophysiological activity unique to meditation. In addition, these results, and in particular the emphasis on the role of coherence, highlight the complexity of the model, reflecting the assumed systemic complexity of the phenomenon of electrophysiological activity during meditation.

Furthermore, it is noteworthy that the current findings are consistent with the main findings reported in the parent study (11). In the original analysis of these data, Braboszcz *et al.* (11) found that elevated gamma waves were linked to trait mindfulness when looking at EEG activity across meditation and mind-wandering conditions. They reported elevation in median gamma power over parieto-occipital electrodes in experienced meditators across all tested meditation traditions relative to meditation-naïve participants. However, in their analysis, the same pattern over frontal electrodes was not found. Additionally, although Braboszcz *et al.* (11) also sought to characterize differences in gamma activity that distinguished between meditation and the control mind-wandering states, no such electrophysiological features were detected. However, in the current study of these EEG data, the ML model was able to predict and characterize the EEG signal during meditation versus nonmeditation states within participants. This is arguably strong and direct evidence of the potential utility of ML for mapping electrophysiology of complex mental states such as meditation, and perhaps more generally, when seeking to characterize mental states vis-à-vis complex, multidimensional, and nonlinear data (33).

Moreover, the observed model explainability findings provide a novel and powerful cross-validation of findings across a large number of independent and heterogeneous EEG studies

that were previously difficult to integrate (4–6,12). Specifically, first, findings of increased beta and gamma power and coherence during the meditation state are consistent with previous studies that documented increased high-frequency oscillations during meditation, mostly in advanced meditators (4,5,7,9,34,35). Beta oscillations are typically linked to sensorimotor processing but have also been associated with attention, emotion, and cognitive control (6,36), and gamma activity and synchronization have been associated with diverse cognitive functions and aspects of arousal, learning, and attention (37,38), as well as enhanced top-down control in vision (39) and audition (40). Second, reduced frontal delta activity during meditation has been reported previously (8). Moreover, the complementary pattern of increased delta activity has been previously linked to autobiographical memory activation (41) and was reported in one study when participants were asked to focus on a present distraction (7), arguably similar to the control state of mind wandering used here, suggesting a potential role of delta oscillations in attentional engagement (4). Third, the observed elevation in theta coherence during the meditation state is consistent with numerous studies that have associated elevated theta power and coherence in midline-frontal brain regions with meditation states (4,6,34,42,43). This activity has been linked to concentrative attentional engagement and internalized attention (4). Finally, although we only tested frontal midline electrodes, the feature importance results described above reflect a general pattern of high coherence and synchronization between electrodes and hemispheres during meditation, as demonstrated by increased gamma coherence in several electrode couples, including bilateral ones (e.g., F1-F2). The increased overall coherence during meditation supports the common hypothesis that meditation states likely involve highly synchronized oscillatory activity in large brain networks rather than a specific, isolated, or localized electrophysiological event (4,44).

Moreover, the current study also builds on initial efforts to apply ML models to classification of mindfulness and related meditation states (16–18). Crucially, to the best of our knowledge, this is the first study to attempt to increase model transparency and apply emerging explainability and feature importance methods to characterizing systems-level electrophysiological activity from ML-based classification models of meditation.

We argue that the current complex and systemic understanding of mindfulness and related meditation mechanisms broadly, and electrophysiology specifically, may have various implications. First, the findings illustrate the potential utility or relative value proposition of ML to help field-wide efforts to characterize the complex, multidimensional, nonlinear, systemic neural activity that underlies meditation states more accurately (12). Second, such knowledge is more likely to guide the development of robust and effective translational innovations and applications. For example, neuromodulation protocols (e.g., transcranial alternating current stimulation) more closely grounded in comprehensive and accurate knowledge of the systemic electrophysiology of meditation may prove more effective and robust (45).

The current study has several limitations. First, findings are limited to a single dataset of 26 participants. Although most EEG studies of meditation among experienced FA meditators have included 10 to 22 participants (6,46–48), the chosen

dataset balances sample size with study design choices such as EEG spatial resolution and control condition to permit generalization by the ML model. To determine the robustness of ML and explainability findings, it is important that additional studies test the reproducibility of the observed findings in other datasets, among larger samples, and with respect to additional meditation traditions and practices. Future investigations could also examine the generalizability of the ML-based model and identified features with and without calibration, which we used in the current study to address the challenge of high inter-participant variability in EEG data, which causes particular concern when training the model using leave-*n*-subjects-out cross-validation (26). Second, future studies could likewise examine the robustness and generalizability of ML-based classification and selected EEG features in the same 9 midline-frontal electrode array selected in this investigation, alternative electrode arrays, and whole-scale electrode sampling. Third, the current study is limited by various design choices made during data preparation and model configuration (e.g., epoch duration, model hyperparameters). Future studies could examine the value of ML generally and XGBoost in particular for classification compared with the performance of various other univariate and multivariate classification approaches. Finally, we focused on 225 univariate features, including spectral and coherence features, to facilitate comparison to past work (12). However, future ML and explainability studies may better characterize the dynamic, multidimensional, and complex nature of meditation by including complexity measures of entropy (49,50) and Higuchi fractal dimension (50), graph-theory based measures of small-worldness and clustering (51), whole-brain network connectivity analysis and interaction with heartbeat evoked potential (52), as well as long-range temporal correlation of neuronal oscillations (53).

## ACKNOWLEDGMENTS AND DISCLOSURES

NA was supported by IDIT-PhD Program for Outstanding Social Sciences Researchers, The Herta & Paul Amir Faculty of Social Sciences, University of Haifa. AB acknowledges support from Israel Science Foundation (Grant No. 1120/21) and Mind and Life PEACE grants (Grant No. 2020-2062321192).

We acknowledge generous access to these open-access data shared by Professor Arnaud Delorme and his co-authors including Braboszcz, Cahn, Levy, and Fernandez. We also acknowledge that the current article is based on NA's master's thesis.

All electroencephalography data from the parent study are available on Zenodo: <http://doi.org/10.5281/zenodo.57911>. The ML and feature importance code that were used to analyze the study data are available on Github: [https://github.com/n0gaa/oscillating\\_mindfully](https://github.com/n0gaa/oscillating_mindfully).

NA, OM, OO, and AB were responsible for conceptualization of the study questions and design. AD was responsible for data collection and design of the parent study. NA, OM, EB, and AB were responsible for data analysis and interpretation. NA and AB were responsible for writing the original draft of the article. NA, OM, OO, EB, AD, and AB were responsible for reviewing and editing the article.

OO is the chief scientific officer and OM is former data scientist of a start-up (NYX) that is developing transcranial alternating current neurostimulation; however, they have no known conflict of interest to disclose with respect to the study or study findings. All other authors report no biomedical financial interests or potential conflicts of interest.

## ARTICLE INFORMATION

From the Observing Minds Laboratory, School of Psychological Science, University of Haifa, Haifa, Israel (NA, AB); NYX Technologies, Haifa, Israel

(OM, OO); Rachel and Selim Benin School of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem, Israel (EB); Swartz Center for Computational Neuroscience, University of California, San Diego, La Jolla, California (AD); Centre de Recherche Cerveau et Cognition, Toulouse III University, Toulouse, France (AD); and Center for Healthy Minds, Department of Psychology, University of Wisconsin-Madison, Madison, Wisconsin (AB).

Address correspondence to Amit Bernstein, Ph.D., at [abernstein@psy.haifa.ac.il](mailto:abernstein@psy.haifa.ac.il).

Received Jan 25, 2024; revised Oct 18, 2024; accepted Oct 28, 2024.

Supplementary material cited in this article is available online at <https://doi.org/10.1016/j.bpsgos.2024.100423>.

## REFERENCES

- Bernstein A, Vago DR, Barnhofer T (2019): Understanding mindfulness, one moment at a time: An introduction to the special issue. *Curr Opin Psychol* 28:vi–vx.
- Schoenberg PL, Vago DR (2019): Mapping meditative states and stages with electrophysiology: Concepts, classifications, and methods. *Curr Opin Psychol* 28:211–217.
- Brandmeyer T, Delorme A, Wahbeh H (2019): The neuroscience of meditation: Classification, phenomenology, correlates, and mechanisms. *Prog Brain Res* 244:1–29.
- Lee DJ, Kulubya E, Goldin P, Goodarzi A, Girgis F (2018): Review of the neural oscillations underlying meditation. *Front Neurosci* 12:178.
- Lomas T, Ivtzan I, Fu CHY (2015): A systematic review of the neurophysiology of mindfulness on EEG oscillations. *Neurosci Biobehav Rev* 57:401–410.
- Cahn BR, Polich J (2006): Meditation states and traits: EEG, ERP, and neuroimaging studies. *Psychol Bull* 132:180–211.
- Cahn BR, Delorme A, Polich J (2013): Event-related delta, theta, alpha and gamma correlates to auditory oddball processing during Vipassana meditation. *Soc Cogn Affect Neurosci* 8:100–111.
- Cahn BR, Delorme A, Polich J (2010): Occipital gamma activation during Vipassana meditation. *Cogn Process* 11:39–56.
- Lutz A, Greischar LL, Rawlings NB, Ricard M, Davidson RJ (2004): Long-term meditators self-induce high-amplitude gamma synchrony during mental practice. *Proc Natl Acad Sci U S A* 101:16369–16373.
- Cacioppo JT, Tassinary LG, Berntson GG (2007): *Handbook of Psychophysiology*. Cambridge, England: Cambridge University Press.
- Braboszcz C, Cahn BR, Levy J, Fernandez M, Delorme A (2017): Increased gamma brainwave amplitude compared to control in three different meditation traditions. *PLoS One* 12:e0170647.
- Deolindo CS, Ribeiro MW, Aratana MA, Afonso RF, Irmischer M, Kozasa EH (2020): A critical analysis on characterizing the meditation experience through the electroencephalogram. *Front Syst Neurosci* 14:53.
- Roy Y, Banville H, Albuquerque I, Gramfort A, Falk TH, Faubert J (2019): Deep learning-based electroencephalography analysis: A systematic review. *J Neural Eng* 16:051001.
- Bzdok D, Altman N, Krzywinski M (2018): Statistics versus machine learning. *Nat Methods* 15:233–234.
- Kora P, Meenakshi K, Swaraja K, Rajani A, Raju MS (2021): EEG based interpretation of human brain activity during yoga and meditation using machine learning: A systematic review. *Complement Ther Clin Pract* 43:101329.
- Hagad JL, Fukui K, Numao M (2020): Deep visual models for EEG of mindfulness meditation in a workplace setting. In: Shaban-Nejad A, Michalowski M, editors. *Precision Health and Medicine*. Cham, Germany: Springer, 129–137.
- Hinterberger T, Kamei T, Walach H (2011): Psychophysiological classification and staging of mental states during meditative practice. *Biomed Tech (Berl)* 56:341–350.
- Ahani A, Wahbeh H, Nezamfar H, Miller M, Erdogmus D, Oken B (2014): Quantitative change of EEG and respiration signals during mindfulness meditation. *J Neuroeng Rehabil* 11:87.
- Shang B, Duan F, Fu R, Gao J, Sik H, Meng X, Chang C (2023): EEG-based investigation of effects of mindfulness meditation training on state and trait by deep learning and traditional machine learning. *Front Hum Neurosci* 17:1033420.
- Arrieta AB, Díaz-Rodríguez N, Del Ser J, Benetot A, Tabik S, Barbado A, et al. (2020): Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion* 58:82–115.
- Lipton ZC (2017): The myths of model interpretability. *arXiv* <https://doi.org/10.48550/arXiv.1606.03490>.
- Lippelt DP, Hommel B, Colzato LS (2014): Focused attention, open monitoring and loving kindness meditation: Effects on attention, conflict monitoring, and creativity – A review. *Front Psychol* 5:1083.
- Lutz A, Slagter HA, Dunne JD, Davidson RJ (2008): Attention regulation and monitoring in meditation. *Trends Cogn Sci* 12:163–169.
- Hart W (1987): *The Art of Living: Vipassana Meditation as Taught by S. N. Goenka*. San Francisco, CA: HarperCollins.
- Jeni LA, Cohn JF, De La Torre F (2013): Facing imbalanced data recommendations for the use of performance metrics. *Int Conf Affect Comput Intell Interact Workshops* 2013:245–251.
- Krumpe T, Baumgärtner K, Rosenstiel W, Spüler M (2017): Non-stationarity and inter-subject variability of EEG characteristics in the context of BCI development. In: *Graz Brain-Computer Interface*.
- Jones DT (2019): Setting the standards for machine learning in biology. *Nat Rev Mol Cell Biol* 20:659–660.
- Tougui I, Jilbab A, El Mhamdi JE (2021): Impact of the choice of cross-validation techniques on the results of machine learning-based diagnostic applications. *Healthc Inform Res* 27:189–199.
- Chen T, Guestrin C (2016): XGBoost: A scalable tree boosting system. In: *KDD 2016: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data*, San Francisco, CA, USA:785–794.
- Lundberg S, Lee SI (2017): A unified approach to interpreting model predictions. *arXiv* <https://doi.org/10.48550/arXiv.1705.07874>.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. (2011): Scikit-learn: Machine learning in python. *J Mach Learn Res* 12:2825–2830.
- Pan F, Converse T, Ahn D, Salvetti F, Donato G (2009): Feature selection for ranking using boosted trees. In: *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, San Francisco, CA: 2025–2028.
- Bzdok D (2017): Classical statistics and statistical learning in imaging neuroscience. *Front Neurosci* 11:543.
- Fell J, Axmacher N, Haupt S (2010): From alpha to gamma: Electrophysiological correlates of meditation-related states of consciousness. *Med Hypotheses* 75:218–224.
- Thomas J, Jamieson G, Cohen M (2014): Low and then high frequency oscillations of distinct right cortical networks are progressively enhanced by medium and long term Satyananda yoga meditation practice. *Front Hum Neurosci* 8:197.
- Symons AE, El-Dereby W, Schwartze M, Kotz SA (2016): The functional role of neural oscillations in non-verbal emotional communication. *Front Hum Neurosci* 10:239.
- Varela F, Lachaux JP, Rodriguez E, Martinerie J (2001): The BrainWeb: Phase synchronization and large-scale integration. *Nat Rev Neurosci* 2:229–239.
- Fries P (2005): A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci* 9:474–480.
- Koch SP, Werner P, Steinbrink J, Fries P, Obrig H (2009): Stimulus-induced and state-dependent sustained gamma activity is tightly coupled to the hemodynamic response in humans. *J Neurosci* 29:13962–13970.
- Schadow J, Lenz D, Dettler N, Fründ I, Herrmann CS (2009): Early gamma-band responses reflect anticipatory top-down modulation in the auditory cortex. *Neuroimage* 47:651–658.
- Imperator C, Brunetti R, Farina B, Speranza AM, Losurdo A, Testani E, et al. (2014): Modification of EEG power spectra and EEG connectivity in autobiographical memory: A sLORETA study. *Cogn Process* 15:351–361.
- Hudak J, Hanley AW, Marchand WR, Nakamura Y, Yabko B, Garland EL (2021): Endogenous theta stimulation during meditation predicts reduced opioid dosing following treatment with Mindfulness-



- Oriented Recovery Enhancement. *Neuropsychopharmacology* 46:836–843.
43. Garland EL, Hanley AW, Hudak J, Nakamura Y, Froeliger B (2022): Mindfulness-induced endogenous theta stimulation occasions self-transcendence and inhibits addictive behavior. *Sci Adv* 8:eabo4455.
44. Tang YY, Hölzel BK, Posner MI (2015): The neuroscience of mindfulness meditation. *Nat Rev Neurosci* 16:213–225.
45. Battleday RM, Muller T, Clayton MS, Cohen Kadosh R (2014): Mapping the mechanisms of transcranial alternating current stimulation: A pathway from network effects to cognition. *Front Psychiatry* 5:162.
46. Tanaka GK, Peressutti C, Teixeira S, Cagy M, Piedade R, Nardi AE, *et al.* (2014): Lower trait frontal theta activity in mindfulness meditators. *Arq Neuropsiquiatr* 72:687–693.
47. Lomas T, Cartwright T, Edginton T, Ridge D (2015): A qualitative analysis of experiential challenges associated with meditation practice. *Mindfulness* 6:848–860.
48. Saggar M, King BG, Zanesco AP, MacLean KA, Aichele SR, Jacobs TL, *et al.* (2012): Intensive training induces longitudinal changes in meditation state-related EEG oscillatory activity. *Front Hum Neurosci* 6:256.
49. Vivot MR, Pallavicini C, Zamberlan F, Vigo D, Tagliazucchi E (2020): Meditation increases the entropy of brain oscillatory activity. *Neuroscience* 431:40–51.
50. Kakumanu RJ, Nair AK, Venugopal R, Sasidharan A, Ghosh PK, John JP, *et al.* (2018): Dissociating meditation proficiency and experience dependent EEG changes during traditional Vipassana meditation practice. *Biol Psychol* 135:65–75.
51. Hauswald A, Übelacker T, Leske S, Weisz N (2015): What it means to be Zen: Marked modulations of local and interareal synchronization during open monitoring meditation. *Neuroimage* 108:265–273.
52. Jiang H, He B, Guo X, Wang X, Guo M, Wang Z, *et al.* (2020): Brain-heart interactions underlying traditional Tibetan Buddhist meditation. *Cereb Cortex* 30:439–450.
53. Irmischer M, Houtman SJ, Mansvellder HD, Tremmel M, Ott U, Linkenkaer-Hansen K (2018): Controlling the temporal structure of brain oscillations by focused attention meditation. *Hum Brain Mapp* 39:1825–1838.
54. Anālayo B (2019): Adding historical depth to definitions of mindfulness. *Curr Opin Psychol* 28:11–14.
55. Dunne JD (2015): Buddhist styles of mindfulness: A heuristic approach. In: *Handbook of Mindfulness and Self-Regulation*. New York, NY: Springer, 251–269.
56. Gethin R (2011): On some definitions of mindfulness. *Contemp Buddhism* 12:263–279.
57. Choi E, Farb N, Pogrebtsova E, Gruman J, Grossmann I (2021): What do people mean when they talk about mindfulness? *Clin Psychol Rev* 89:102085.
58. Nilsson H, Kazemi A (2016): Reconciling and thematizing definitions of mindfulness: The big five of mindfulness. *Rev Gen Psychol* 20:183–193.
59. Wielgosz J, Goldberg SB, Kral TRA, Dunne JD, Davidson RJ (2019): Mindfulness meditation and psychopathology. *Annu Rev Clin Psychol* 15:285–316.