Resting EEG and Anxiety

Classification of state and trait anxiety using EEG recordings and evaluation of logistic regression, random forest, and support vector machine models

By Joshua Threlkeld

Abstract

Resting-state electroencephalography (EEG) offers a non-invasive window into brain activity, but its utility as a psychological biomarker (measurable indicator of mental state) can vary significantly depending on what it is used to predict. This paper investigates the degree to which frequency-domain spectral features including band power, frontal alpha asymmetry, and alpha/beta ratios can classify resting-state EEG conditions (eyes open vs. eyes closed) and predict state and trait anxiety. Using an OpenNeuro dataset with 256-channel resting-state recordings and Spielberger State-Trait Anxiety Inventory (STAI) scores, I apply logistic regression, random forest, and support vector machine classifiers across three tasks of increasing time (temporal distance) between the EEG recording and target variable. Eyes open (EO) measured against eyes closed (EC) classification achieves 65–67% accuracy, with occipital alpha power emerging as the dominant feature: consistent with established visual cortex dynamics. State anxiety classification reaches 53–55% at the subject level, or only moderately above chance. Trait anxiety classification drops to 39–43%: below chance across all models and aggregation levels. Correlation analysis reveals that state and trait anxiety engage spectral features in opposite directions at frontal electrodes, suggesting distinct neural signatures. Critically, at no point in this study did model complexity resolve classification failure. These results point to the bottleneck being feature expressiveness relative to the target variable's temporal structure. Thus, while resting-state spectral features are reliable state markers, they are insufficient at trait prediction.

Contents

1. Introduction

Machine learning in neuroscience represents a stem cell in the body of AI in healthcare. This degree of immense promise is tied to the fact that resting-state EEG is among the most accessible and cost-effective modalities for neurological research. Unlike task-based paradigms, resting-state recordings do not require stimulus delivery or behavioral response, making them practical for clinical deployment. However, the relationship between cortical oscillations and psychological constructs is often indirect, and the conditions under which spectral features carry diagnostic signals remain only partially understood.

Anxiety is an ideal medium to examine this relationship. Among the most prevalent mental health conditions, its patterns of brain electrical activity have been studied extensively with EEG. Several spectral features, such as frontal alpha asymmetry and alpha/beta ratio, have been proposed as anxiety biomarkers. Frontal alpha asymmetry is defined as the log of the difference between right and left frontal alpha electrode power. It is associated with emotional valence and approach-withdrawal motivation (Davidson, 1992). In particular, rightward asymmetry is linked with anxiety and negative affect. (Adolph & Margraf, 2017) Yet, frontal alpha asymmetry is not without limitations as a biomarker. Effect sizes across anxiety studies are modest and sensitive to outliers, and the measure's specificity to anxiety versus depression remains an open empirical question (Thibodeau et al., 2006). Additionally, the alpha/beta power ratio at parietal electrodes has been proposed as a general index of cortical relaxation, despite its components playing different roles (mechanistic and descriptive, respectively) in active processing. Alpha suppresses active processing, whereas beta describes active processing. It stands to reason that a high alpha/beta ratio is indicative of the brain actively suppressing engagement. Nonetheless, classification studies using these features report inconsistent results, and effect sizes are generally modest (Thibodeau et al., 2006). Namely, these results are parallel to established limitations in biomedical signal classification. Cavanagh (2021) — an analysis of EEG spectral features in a depression cohort dataset — similarly found weak classification performance. This too suggests the limitation is specific to resting-state spectral features rather than to purely research on anxiety.

One underexplored explanation for this inconsistency is the temporal relationship between the EEG recording and psychological construct. Resting-state EEG captures a specific snapshot of cortical activity. State anxiety, which reflects how an individual feels just before the moment of assessment, is temporally aligned with that snapshot. Trait anxiety, which reflects a dispositional tendency, is not. If spectral features encode momentary emotion more reliably than stable emotion, the time between measurement and target should directly correlate to task classification performance: this paper tests that hypothesis directly.

I use a dataset with resting-state EEG and concurrent STAI state and trait scores to build a three-task classification framework (Shalamberidze et al., 2026). The tasks are ordered by increasing temporal separation from the recording: eyes open vs. eyes closed, state anxiety, and trait anxiety. Eyes open (EO) vs. eyes closed (EC) serves as a pipeline validation task with a known neural signature, serving as a de facto performance ceiling. State and trait anxiety classification then probe whether the same feature set generalizes to psychological targets of increasing temporal removal. Critically, I apply identical preprocessing, feature extraction, and modeling procedures across all three tasks so that differences in classification performance can be solely attributed to the target variable. Rather than optimizing for predictive accuracy, I aim to characterize where resting-state spectral features succeed and fail to encode, and to identify the conditions under which they carry meaningful signals.

2. Methods

2.1 Datasets

The dataset used in this study is a publicly available resting-state EEG collection hosted on OpenNeuro (Shalamberidze et al., 2026). Participants were 51 right-handed undergraduate students (26 female, 25 male) recruited from the University of Alberta, aged 17–51 years (M = 20.4, SD = 4.9). No additional exclusions were applied beyond those made by the original authors during preprocessing. (Shalamberidze et al., 2025)

EEG was recorded using a 256-channel EGI HydroCel Geodesic Sensor Net with Net Amps amplifier at a sampling rate of 500 Hz, with online reference at Cz. Participants completed a resting-state protocol consisting of four alternating one-minute blocks in the order eyes open, eyes closed, eyes open, eyes closed (EO-EC-EO-EC), with transitions between blocks signaled by an auditory beep.

Data were preprocessed by the original authors in EEGLAB prior to public release. The pipeline included bandpass filtering (0.1–50 Hz), line noise removal via CleanLine at 60 and 120 Hz, kurtosis-based channel rejection applied twice, average re-referencing, ICA decomposition using the extended Infomax algorithm, and artifact component removal via ICLabel at a probability threshold of 0.8, followed by spherical interpolation of removed channels. No additional preprocessing was applied.

Anxiety was assessed using Spielberger State-Trait Anxiety Inventory (STAI) surveys, which yield separate continuous scores for how a participant feels at the moment of assessment (state anxiety) and for dispositional anxiety (trait anxiety). Critically, it was established that state anxiety is temporally proximate to the EEG signal, whereas trait anxiety is by definition temporally distal. Both subscales were available for all participants and used as target variables in the classification analyses described below. Additionally, both subscales were administered prior to the EEG recording, meaning state scores reflect participants' momentary anxiety immediately preceding the resting-state session.

2.2 Feature Extraction

Spectral features were extracted from six electrodes, capturing frontal and occipital dynamics relevant to anxiety and perceptual brain state: F3, F4, and Fz in the frontal region, and O1, O2, and Pz in the posterior region. Electrode labels were mapped to their corresponding EGI HydroCel 256-channel equivalents (F3 → E37, F4 → E18, Fz → E26, O1 → E123, O2 → E158, Pz → E101). Each continuous block was segmented into non-overlapping two-second epochs, yielding approximately 30 epochs per block per subject. Band power was estimated for each epoch and channel using Welch's method[1] with a window length of 1,000 samples (2 seconds at 500 Hz), equal to the epoch duration. Power was extracted across three frequency bands: theta (4–8 Hz), alpha (8–13 Hz), and beta (13–30 Hz), and log-transformed prior to modeling to normalize the distribution of power values. Two composite features were additionally computed. Frontal alpha asymmetry was defined as the log difference in alpha power between right and left frontal electrodes ($log_{F4\_alpha} - log_{F3\_alpha}$), operationalizing the approach-withdrawal framework described in the introduction (Davidson, 1992; Adolph & Margraf, 2017). The alpha/beta power ratio at Pz was computed as a continuous index of posterior cortical relaxation,

[1] A two-second epoch is used to calculate the Fourier transform, thus reducing noise.

with higher values reflecting greater inhibition of active processing. (Chang & Choi, 2023) This yielded a feature vector of 20 dimensions per epoch: 18 band power values (6 channels × 3 bands) plus frontal alpha asymmetry and the Pz alpha/beta ratio. Across all 51 subjects and four blocks, the full feature matrix contained 6,120 epochs.

2.3 Classification

2.3.1 Models

Three classifiers with fundamentally different boundary assumptions — logistic regression, SVM, and random forest — were used to partition the 20-dimensional spectral feature space. Logistic regression ignores outliers and draws a linear hyperplane; this is a favorable characteristic because, in data with few participants, outliers may cause overfitting. SVM with an RBF kernel is capable of learning non-linear boundaries without the variance of an ensemble method. Random forest too learns non-linear feature interactions and estimates. However, unlike SVM, it partitions the space recursively through ensemble decision trees.All three models were configured with balanced class weights to account for label imbalance across tasks. Default configurations were used throughout to avoid overfitting and ensure differences in performance can be attributed to the nature of what's being predicted, dispelling the notion of any model being limited/optimized for one task.

2.3.2. Cross-Validation

For epoch-level classification, a five-fold StratifiedGroupKFold procedure was used, preventing leakage and ensuring that all epochs were assigned to either training or test data. Stratification preserves class balance across folds, while group assignment by subject ID guarantees that a subject's epochs do not repeat in training and test splits. Epochs from the same subject are highly correlated and subject overlap between splits would inflate accuracy estimates that fail to generalize to new individuals. Feature scaling via StandardScaler[2] was applied within each fold and only fit on training data.

For subject-level classification, epochs were averaged per subject prior to modeling, reducing the dataset to one feature vector per participant. With only 51 subjects, Leave-One-Out

[2] Normalizes each feature to zero mean, thus preventing differences in spectral scale from biasing the models

cross-validation was used. This method of training on 50 subjects and testing one at each fold maximized training data and produced stable performance estimates.

2.3.3. Task Definitions

Three binary classification tasks were constructed, ordered by increasing the time (temporal distance) between the EEG recording and the target variable.

Task 1 (Eyes Open vs. Eyes Closed): Each epoch was labeled according to the block condition in which it was recorded (EO = 0, EC = 1). Effectively, this task serves as a pipeline validation benchmark task with a known neural signature.

Task 2 (Trait Anxiety): Subjects were divided into high and low trait anxiety groups using a median split of STAI-trait scores (high ≥ median = 1, low < median = 0). Median split labeling was applied at the subject level prior to epoch assignment, ensuring all epochs from a given subject carried the same label. This task was evaluated at both epoch and subject aggregation levels.

Task 3 (State Anxiety): An identical median split procedure was applied to STAI-state scores (high ≥ median = 1, low < median = 0). Because state scores were collected immediately before the EEG recording. As such, it serves as the intermediate point in the temporal distance framework.

2.4 Evaluation

Classification performance was assessed using two metrics reported for each model and task: accuracy and macro-averaged F1 score. Macro F1 averages precision and recall across both classes, giving equal weight to minority and majority class performance. Therefore, macro F1 is a more robust metric in this study because accuracy can be dominated by the majority class.

For epoch-level models, mean and standard deviation of both metrics are reported across the five cross-validation folds. For subject-level models using Leave-One-Out cross-validation, a single accuracy and F1 value is reported across all held-out predictions.

Chance performance for all binary tasks is 50%. Results consistently at or below this threshold are interpreted as an absence of recoverable signal in the feature set relative to the target variable. That is to say, the "temporal distance" hypothesis specifies that below-chance trait anxiety classification isn't a failure of the pipeline. Rather, the null result *is* the result.

In addition to classification metrics, Pearson correlations between subject-level mean features and continuous STAI scores are reported as a complementary analysis. (Al-Qazzaz et al., 2025) This threshold-free approach assesses linear relationships between individual features and anxiety severity without reducing the target to a binary label, and is used to identify which features carry directional information about anxiety even when classification performance is near chance.

3. Results

3.1 Eyes Open vs. Eyes Closed

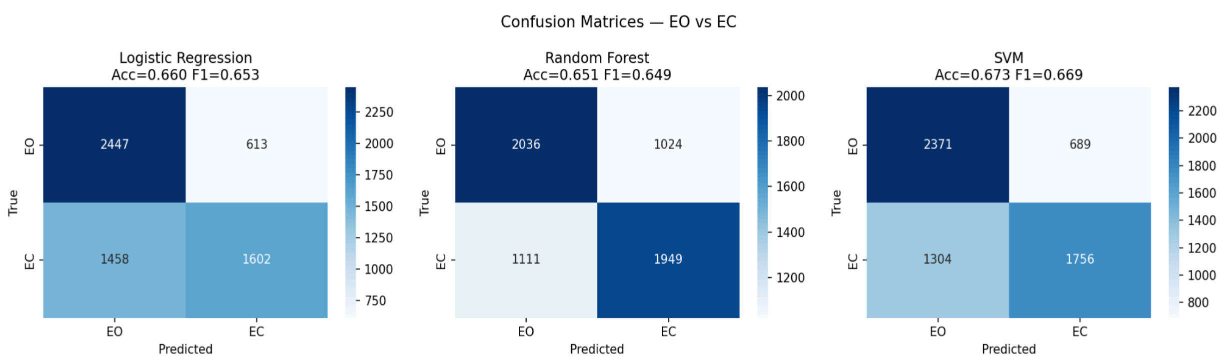

Fig. 1: Confusion matrices for the classification of EO vs EC brain states using three machine learning classifiers trained on EEG spectral features. Each matrix displays the number of correctly and incorrectly classified segments for Logistic Regression (Accuracy=0.660, F1=0.653), Random Forest (Accuracy=0.651, F1=0.649), and Support Vector Machine (Accuracy=0.673, F1=0.669). In all three models, EO segments were classified more accurately than EC segments.

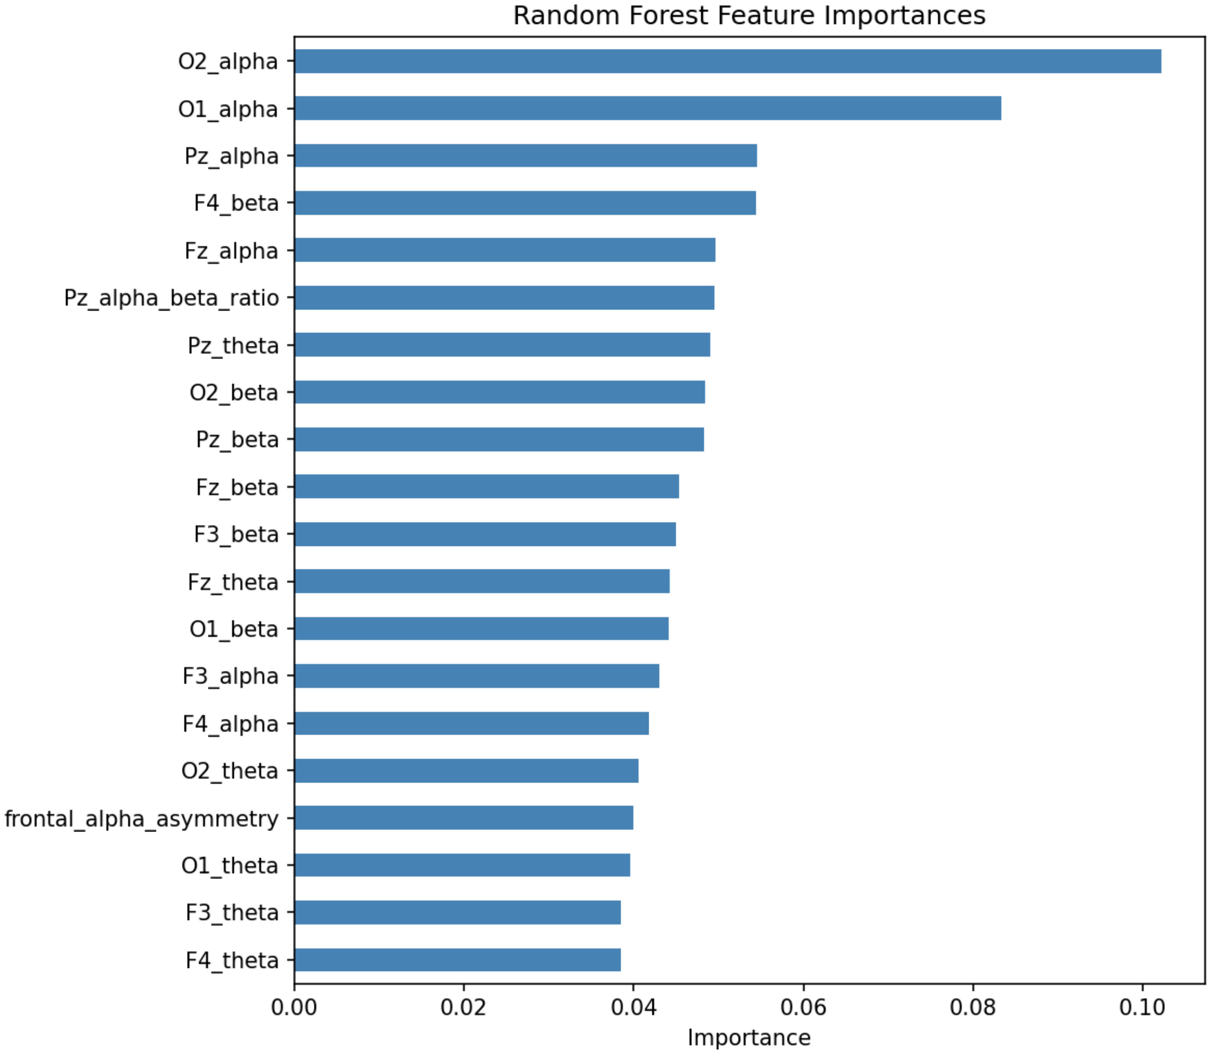

Fig. 2: Random Forest feature importances for the EO vs. EC brain state classification. Feature importance values represent the mean decrease in impurity across all decision trees, with higher values indicating greater contribution to classification performance. Occipital alpha power at electrodes O2 (importance ≈ 0.103) and O1 (importance ≈ 0.083) emerged as the two most influential features. Frontal theta (F3_theta, F4_theta) and occipital theta features (O1_theta) were ranked lowest in importance.

Epoch-level classification of EO vs. EC yielded the strongest performance across all tasks. Logistic regression achieved an accuracy of 0.660 (F1 = 0.653), random forest achieved 0.651 (F1 = 0.649), and SVM achieved 0.673 (F1 = 0.669). Performance was consistent across all three models, spanning a range of less than 0.03 in accuracy, which indicates that the result reflects a ceiling in feature expressiveness rather than the limitation of any single modeling approach.

Inspection of the confusion matrix (Figure 1) reveals that EO epochs were classified more accurately than EC epochs across all three models. This asymmetry is consistent with the spectral signature of eyes open recording: posterior alpha suppression and elevated broadband activity produce a more distinctive feature profile than eyes closed resting state.

Feature importance analysis (Figure 2) identified O2_alpha and O1_alpha as the dominant predictors, followed by Pz_alpha and F4_beta. The prominence of occipital alpha power is consistent with the well-established alpha-blocking phenomenon, in which alpha power at posterior sites is suppressed during visual engagement and elevated during eyes closed rest (Davidson, 1992). This provides direct validation that the spectral feature extraction pipeline captures physiologically meaningful signals.

3.2 Trait Anxiety

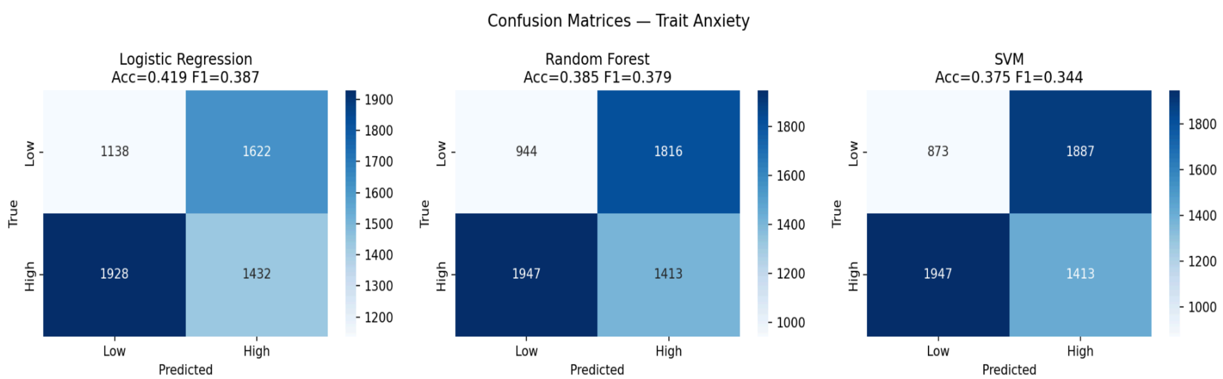

Fig. 3: Confusion matrices for the binary classification of trait anxiety (low vs. high) using logistic regression (Accuracy=0.419, F1=0.387), random forest (Accuracy=0.385, F1=0.379), and support vector machine (Accuracy=0.375, F1=0.344) models.

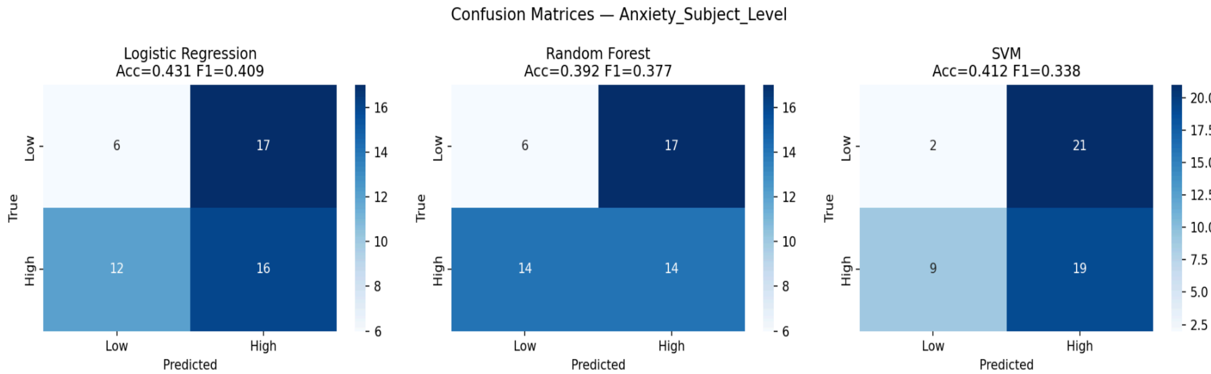

Fig. 4: Confusion matrices for the binary classification of anxiety level (low vs. high) using logistic regression (Acc=0.431, F1=0.409), random forest (Acc=0.392, F1=0.377), and support vector machine (Acc=0.412, F1=0.338) models.

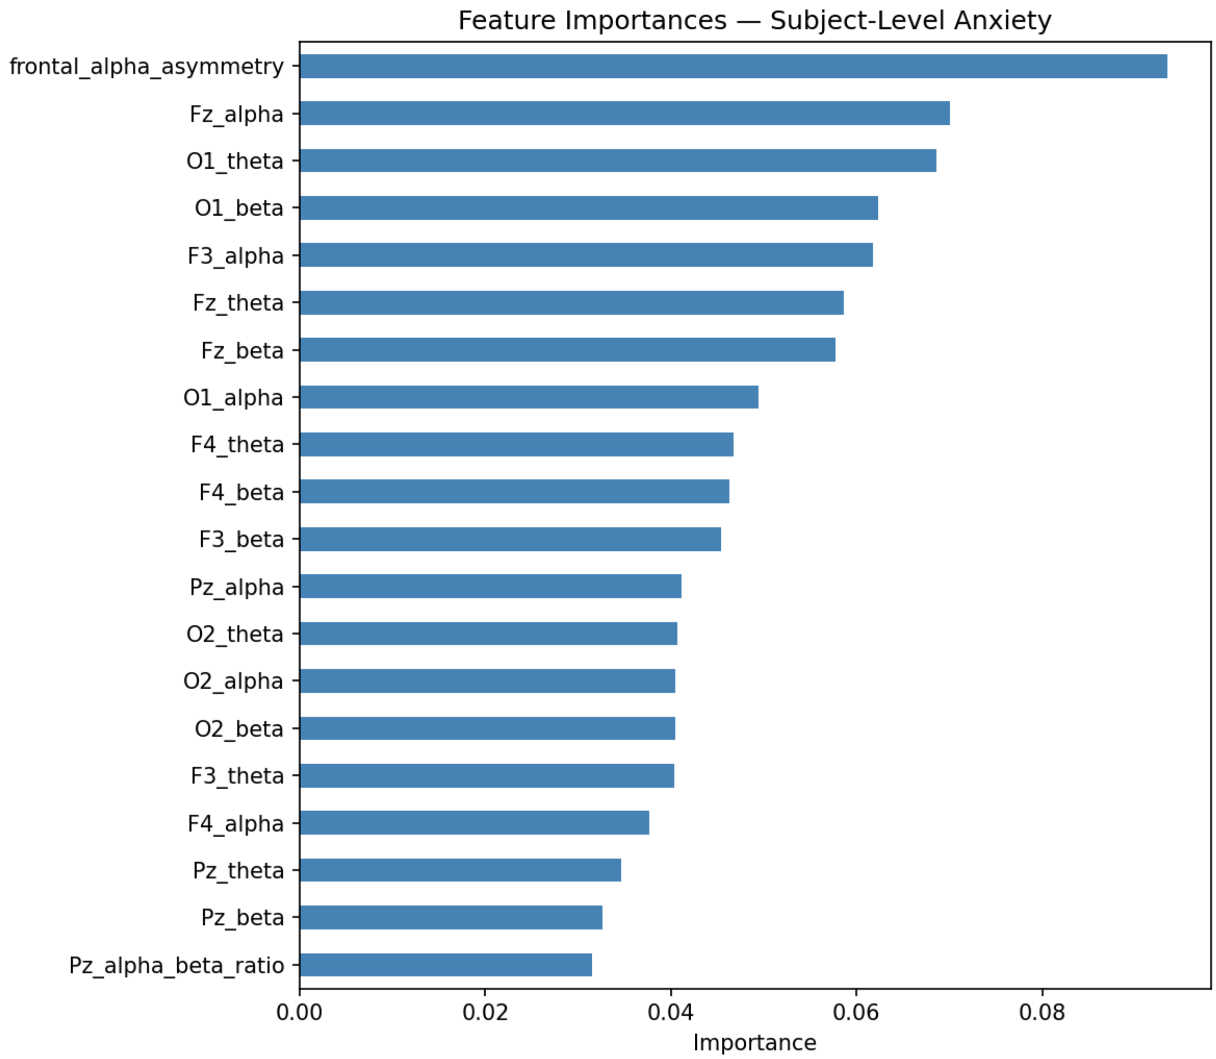

Fig. 5: Random forest feature importances for the classification of anxiety level at the subject level. Feature importance values represent the mean decrease in impurity across all decision trees in the ensemble, with higher values indicating greater contribution to classification performance. Among the 20 EEG spectral features examined, frontal alpha asymmetry (importance ≈ 0.091) and Fz_alpha (importance ≈0.069) ranked the highest. Parietal alpha-to-beta ratio and Pz_beta ranked lowest among all features.

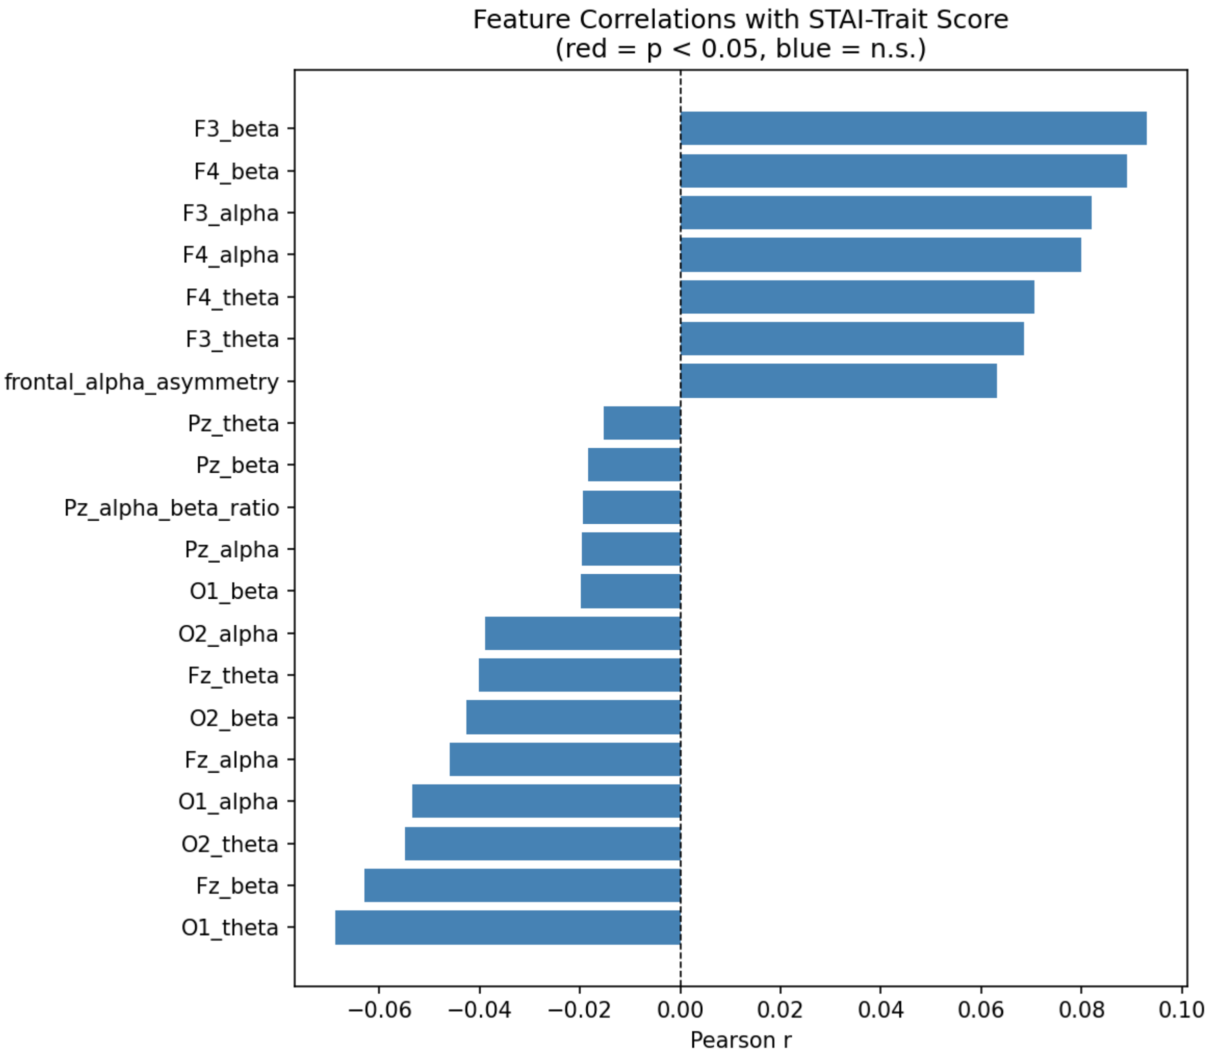

Fig. 6: Pearson correlations between EEG spectral features and STAI-Trait scores. Features are ranked by correlation magnitude, with none reaching statistical significance ($p < 0.05$). Frontal beta and alpha features (F3_beta, F4_beta, F3_alpha, F4_alpha) showed the strongest positive correlations, while occipital and midline features (O1_theta, Fz_beta, O2_theta) exhibited the largest negative correlations.

Epoch-level classification of trait anxiety (high vs. low) performed below chance across all three models. Logistic regression achieved an accuracy of 0.419 (F1 = 0.387), random forest achieved 0.385 (F1 = 0.379), and SVM achieved 0.375 (F1 = 0.344). Performance was similarly below chance at the subject level under leave-one-out cross-validation: logistic regression 0.431

(F1 = 0.409), random forest 0.392 (F1 = 0.377), and SVM 0.412 (F1 = 0.338). Confusion matrices for epoch-level and subject-level analyses are shown in Figures 3 and 4, respectively.

The pattern of below-chance performance is consistent across both levels of analysis and across all three classifiers, suggesting it reflects the absence of recoverable signal in the feature set relative to the trait anxiety label rather than a model-specific failure.

Pearson correlations between subject-level mean features and continuous STAI-trait scores confirmed this interpretation (Figure 5). No feature reached statistical significance (all $p > 0.05$). The largest positive correlations were observed at bilateral frontal electrodes (F3_beta: $r = 0.093$, F4_beta: $r = 0.089$, F3_alpha: $r = 0.082$), while the largest negative correlations appeared at occipital and frontal midline sites (O1_theta: $r = -0.069$, Fz_beta: $r = -0.063$). Effect sizes were uniformly small, with no feature exceeding $|r| = 0.10$.

3.3 State Anxiety

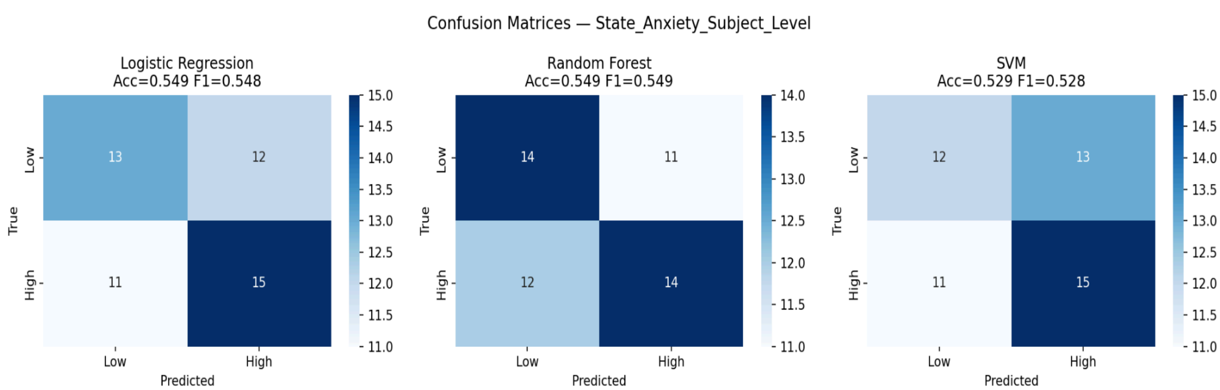

Fig. 7: Confusion matrices for the binary classification of state anxiety (low vs. high) using logistic regression (Accuracy=0.502, F1=0.491), random forest (Accuracy=0.471, F1=0.461), and support vector machine (Accuracy=0.444, F1=0.429) models.

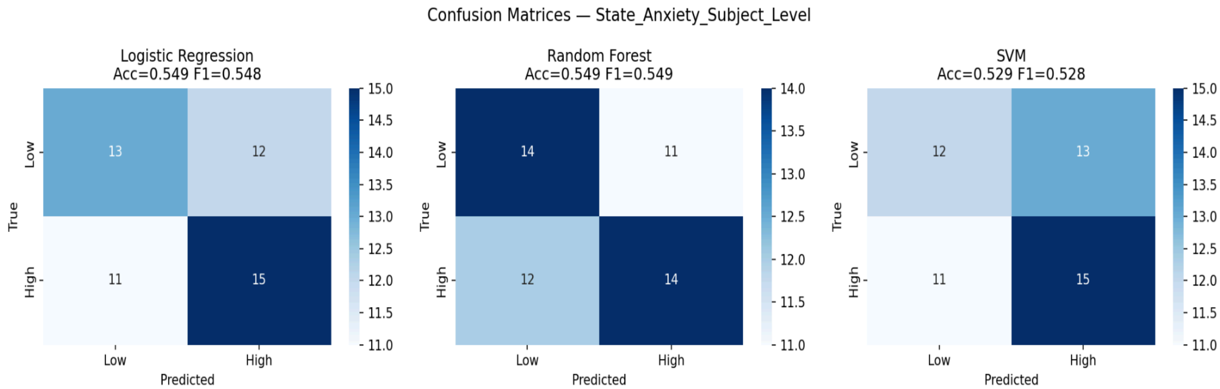

Fig. 8: Confusion matrices for the binary classification of state anxiety (low vs. high) using logistic regression (Accuracy=0.549, F1=0.548), random forest (Accuracy=0.549, F1=0.549), and support vector machine (Accuracy=0.529, F1=0.528) models.

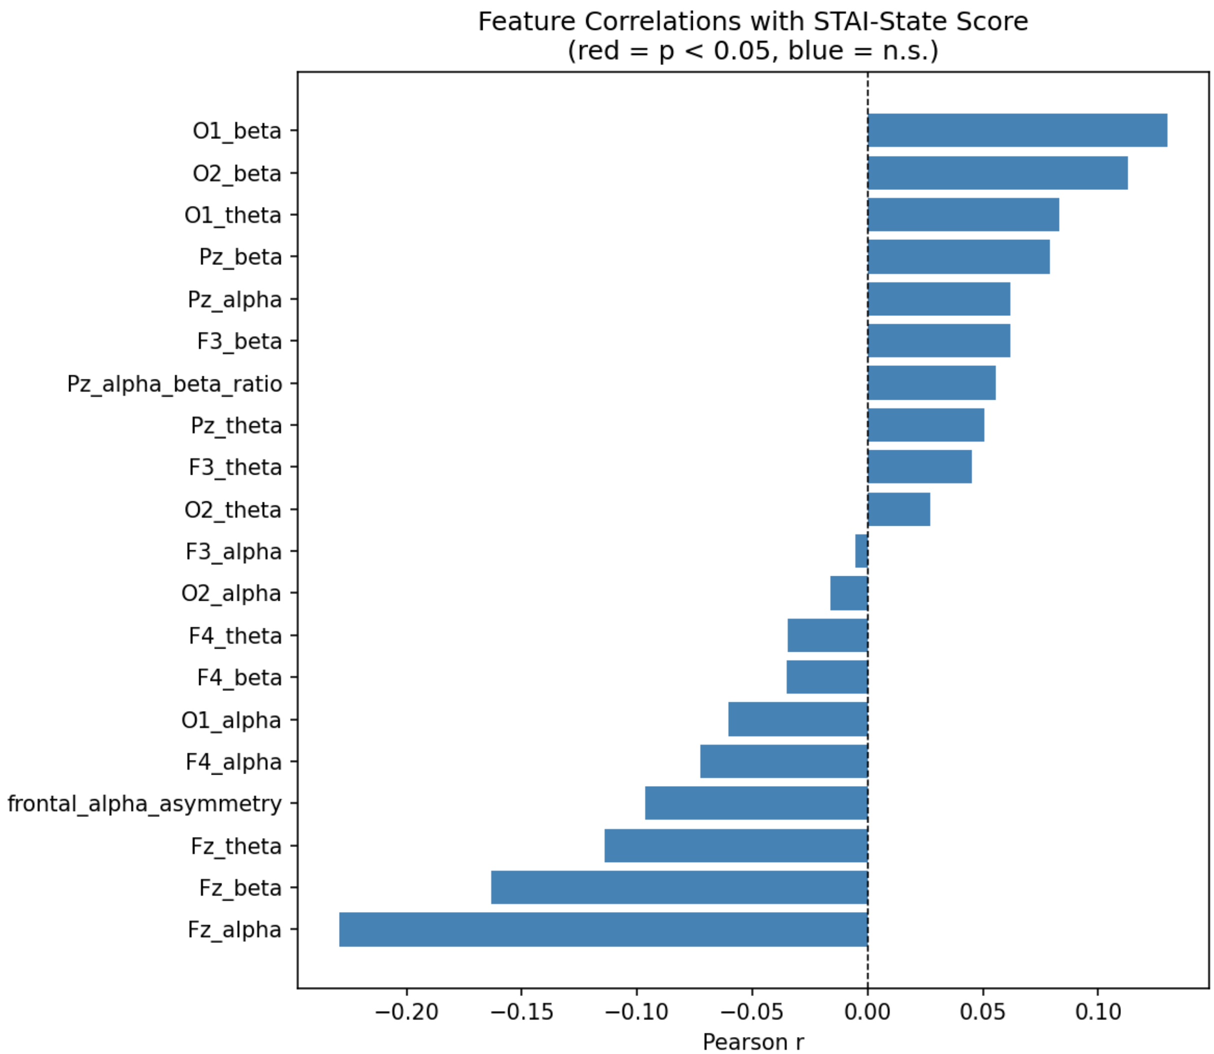

Fig. 9: Pearson correlations between EEG spectral features and STAI-State scores. Features are ranked by correlation magnitude, with all bars displayed in blue indicating that no correlations reached statistical significance ($p < 0.05$). Occipital beta and alpha features (O1_beta, O2_beta) showed the strongest positive correlations, while midline frontal features (Fz_alpha, Fz_beta) exhibited the largest negative correlations.

Epoch-level classification of state anxiety (high vs. low) produced near-chance performance. Logistic regression achieved an accuracy of 0.502 (F1 = 0.491), random forest achieved 0.471 (F1 = 0.461), and SVM achieved 0.444 (F1 = 0.429). Confusion matrices are shown in Figure 6.

At the subject level, performance improved modestly but consistently across all three classifiers: logistic regression 0.549 (F1 = 0.548), random forest 0.549 (F1 = 0.549), and SVM 0.529 (F1 = 0.528), as shown in Figure 7. The convergence of all three models above chance at the subject level and, notably, the agreement between logistic regression and random forest suggests a weak but stable signal. This signal is bolstered with aggregation over epochs.

Pearson correlations between subject-level mean features and continuous STAI-state scores are shown in Figure 8. As with trait anxiety, no feature reached statistical significance (all $p > 0.05$). The strongest negative correlations were observed at frontal midline electrodes: Fz_alpha ($r = -0.229$), Fz_beta ($r = -0.163$), and Fz_theta ($r = -0.114$). The strongest positive correlations appeared at posterior occipital sites: O1_beta ($r = 0.130$) and O2_beta ($r = 0.113$). The directional contrast between frontal and posterior features is discussed further in Section 5.

3.4 Summary of Classification Performance

Task	Model	Epoch Acc.	Epoch F1	Subject Acc.	Subject F1
EO vs. EC	Logistic regression	0.660 ± 0.042	0.653 ± 0.040	—	—
	Random forest	0.651 ± 0.034	0.649 ± 0.034	—	—
	SVM	0.673 ± 0.039	0.669 ± 0.039	—	—
Trait Anxiety	Logistic regression	0.419 ± 0.109	0.387 ± 0.124	0.431	0.409
	Random forest	0.385 ± 0.035	0.379 ± 0.028	0.392	0.377
	SVM	0.375 ± 0.071	0.344 ± 0.053	0.412	0.338
State Anxiety	Logistic regression	0.502 ± 0.112	0.491 ± 0.120	0.549	0.548
	Random forest	0.471 ± 0.051	0.461 ± 0.051	0.549	0.549
	SVM	0.444 ± 0.086	0.429 ± 0.105	0.529	0.528

Tab. 1: Classification performance across three tasks (EO vs. EC, Trait Anxiety, State Anxiety) for logistic regression, random forest, and SVM models. Epoch-level accuracy and F1 are reported as mean ± standard deviation across cross-validation folds. Subject-level accuracy and F1 reflect aggregated epoch predictions per participant and are omitted for the EO vs. EC task, which was not subject-stratified. Chance performance = 0.50 for all tasks. Subject-level EO vs. EC was not performed as the task label is epoch-level by design.

4. Discussion

4.1 The Temporal Distance Gradient

The central finding of this study is a monotonic decrease in classification performance as the temporal distance between the EEG recording and the target variable increases. EO vs. EC classification describes the brain state at the exact moment of measurement and achieved accuracies of 65–67%. State anxiety classification, where the label reflects how the participant reported feeling immediately before recording, produced near-chance performance at the epoch level (44–50%) with a modest, above-chance signal at the subject level (53–55%). Trait anxiety classification, where the label describes a stable personality disposition that exists independently of any single recording session, fell below chance at both the epoch level (38–42%) and subject level (39–43%).

This gradient is not parsimoniously explained by task difficulty or class imbalance. Class distributions were approximately balanced across all three tasks, and all models were trained with balanced class weights. The convergence of logistic regression, random forest, and SVM at nearly identical performance levels within each task, despite fundamental differences in their decision boundaries, further indicates that the limiting factor is the structure of the feature space relative to the label, not the expressive capacity of any individual model. Together, these patterns support the interpretation that resting-state spectral features reliably encode momentary brain state, but carry a diminishingly recoverable signal as the temporal scope of the target variable scales beyond the recording window.

4.2 Pipeline Validation and the Role of EO vs. EC

The EO vs. EC result functions as an internal validation of the pipeline. A feature extraction procedure that failed to distinguish EO from EC would offer no interpretive foundation for null results in the anxiety tasks; below-chance performance could simply reflect a broken pipeline. The fact that occipital alpha power emerges as the dominant discriminating feature, consistent with the well-documented alpha-blocking phenomenon (Davidson, 1992), demonstrates that the pipeline captures physiologically grounded signals. The null results for trait anxiety and the marginal results for state anxiety can therefore be attributed to the relationship between the features and the target variable, not to a deficiency in measurement.

This logic mirrors the use of positive controls in experimental design. Just as a biological assay requires a positive control to confirm that a negative result reflects the absence of the effect rather than a failure of the assay, EO vs. EC confirms that spectral features are recoverable from this dataset before asking whether they are diagnostic of anxiety.

4.3 Trait Anxiety: A Null Result as Evidence

Below-chance performance on trait anxiety classification is interpretable as evidence about the nature of what trait anxiety is. Trait anxiety, as operationalized by the STAI-T subscale, reflects a dispositional tendency to experience anxiety across situations and time. A two-minute resting-state recording captures the participant's neural state on a single occasion and has no particular reason to encode a variable defined by its stability across occasions.

This interpretation is consistent with prior literature. Thibodeau et al. (2006), in a meta-analysis of frontal EEG asymmetry and anxiety, found that effect sizes are generally modest and highly variable across studies, suggesting that no single spectral feature reliably discriminates anxious from non-anxious individuals. The present results extend this finding to a multivariate machine learning setting: even when 20 features spanning multiple frequency bands and electrode sites are considered jointly, trait anxiety remains unrecoverable from a single resting-state session. This is consistent with Cavanagh (2021), whose analysis of resting-state spectral features similarly found weak classification performance, suggesting the limitation may be a property of resting-state spectral features generally rather than a problem specific to this dataset or methodology.

The subject-level trait accuracy of 39–43%, notably below chance, warrants brief comment. Rather than reflecting the models learning an inverted signal, this pattern most likely reflects the interaction between a very small test set (51 subjects under LOO-CV) and high variance across folds; a single misclassified subject shifts accuracy meaningfully at this sample size. The large standard deviations observed at the epoch level (e.g., ±0.109 for logistic regression accuracy) further support the interpretation that performance is unstable and centered near chance rather than systematically inverted.

4.4 State Anxiety: A Weak but Interpretable Signal

The state anxiety results occupy a theoretically meaningful middle position between EO/EC and trait anxiety. At the epoch level, performance is near chance; within-session variability of individual two-second epochs appears to obscure any signal present. At the subject level, averaging across all epochs produces a stable representation that yields above-chance classification (53–55%) across all three models. The agreement between logistic regression and random forest at exactly 0.549 is particularly notable, as it indicates the result is not due to a single model's inductive bias.

The Pearson correlation analysis illuminates the feature-level structure underlying this result. Fz_alpha shows the largest negative correlation with STAI-state score ($r = -0.229$), followed by Fz_beta ($r = -0.163$) and Fz_theta ($r = -0.114$). This pattern of frontal midline suppression is consistent with theoretical accounts linking elevated frontal midline theta to internalized attention and the cognitive load characteristic of anxious states (Cavanagh & Shackman, 2015). Simultaneously, posterior occipital beta (O1_beta: $r = 0.130$, O2_beta: $r = 0.113$) correlates positively with state anxiety scores, suggesting a pattern of heightened posterior broadband activity alongside frontal suppression. Neither pattern reaches statistical significance at $n = 51$.

4.5 Divergent Spectral Signatures of State and Trait Anxiety

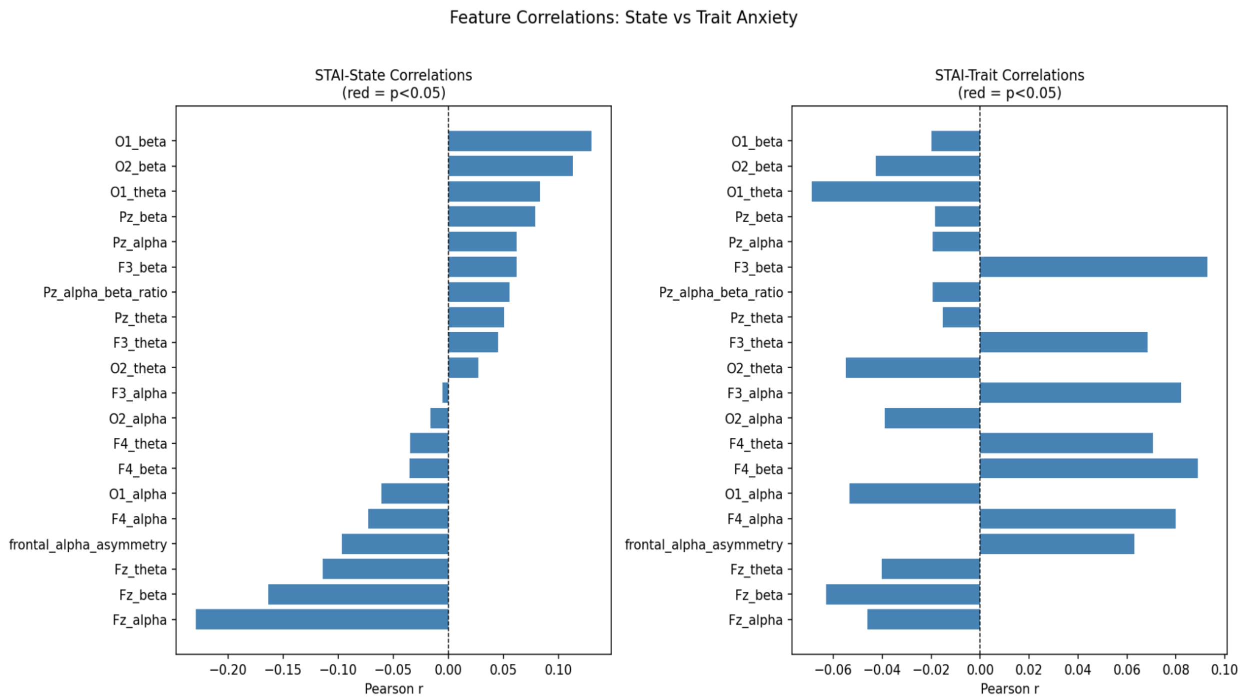

Fig. 10: Pearson correlations between EEG spectral features and STAI-State (left) and STAI-Trait (right) scores. Features are ranked by correlation magnitude, with none reaching statistical significance ($p < 0.05$). For STAI-State, occipital beta features (O1_beta, O2_beta) showed the strongest positive correlations, while midline frontal features (Fz_alpha, Fz_beta) exhibited the largest negative correlations. For STAI-Trait, frontal beta and alpha features (F3_beta, F4_beta, F3_alpha) showed the strongest positive correlations, while occipital theta (O1_theta) exhibited the largest negative correlation.

A notable finding from the correlation analysis is that state and trait anxiety are associated with distinct and, in some cases, opposing spectral patterns across the same electrode sites (Figure 9). Frontal midline alpha (Fz_alpha) correlates negatively with state anxiety ($r = -0.229$) but negligibly with trait anxiety ($r = -0.046$). Conversely, right frontal beta (F3_beta) shows its largest positive correlation with trait anxiety ($r = 0.093$) and a weaker positive correlation with state anxiety ($r = 0.062$). At F4_alpha, F4_beta, F3_alpha, state and trait correlations point in opposite directions.

Rather than reflecting different magnitudes of the same underlying neural process, state and trait anxiety appear to engage distinct and partially opposing frequency-band dynamics at overlapping electrode sites. This warrants replication in larger samples before firm conclusions can be drawn. Whether this reflects genuinely different neural generators or simply the consequence of measuring two constructs on different timescales remains an open question beyond the scope of this study. However, the pattern reinforces the argument that classifiers trained on features correlated with one anxiety construct cannot be expected to generalize to the other — an important caveat for applied EEG-based anxiety detection systems that conflate the two.

5. Limitations

The most fundamental constraint of this study is its small sample size. With only 51 participants, a single misclassified subject produces a non-trivial shift in reported accuracy. This problem is compounded under leave-one-out cross-validation, where each fold tests on exactly one subject. Performance estimates are highly sensitive to individual outliers. The subject-level results, thus, cannot be fully considered *definitive* without replication with a larger cohort.

A related limitation concerns the target variable itself. STAI scores are self-reported, and self-report measures of anxiety are subject to biases including introspective inaccuracy and response anchoring. That is to say, participants may not accurately perceive or report their momentary anxiety state, introducing irreparable noise into the labels. This measurement imprecision is distinct from the temporal distance argument developed in the discussion; it applies even when label and recording are temporally aligned, as in the state anxiety task.

The binary classification approach also introduces an artificial, superficially-imposed boundary along continuous distributions of both EEG activity and STAI scores. Median-split labeling collapses individual differences near the threshold; participants on opposite sides of the median are treated as categorically distinct. Regression-based approaches that preserve the continuous structure of STAI scores would be a more sensitive framework for future work.

Finally, this study is restricted to frequency-domain spectral features extracted from six electrodes. While this feature set was deliberately constrained to test a specific hypothesis about temporal distance and feature expressiveness, it does nothing to acknowledge the contributions of temporal dynamics, cross-frequency coupling, and functional connectivity. More expressive representations may recover signals that band power features cannot. Nonetheless, resolving the fundamental temporal mismatch between a single resting-state session and a dispositional trait construct remains a promising open question.

6. Conclusion

This study investigated whether resting-state EEG spectral features could classify perceptual brain state, momentary state anxiety, and dispositional trait anxiety, using a three-task framework explicitly ordered by increasing temporal distance between the neural recording and the target variable. The results support a coherent and interpretable pattern: classification performance degrades monotonically as that temporal distance grows.

Eyes open versus eyes closed classification achieved 65–67% accuracy across all three models, with occipital alpha power emerging as the dominant discriminating feature: a result directly consistent with the well-established alpha-blocking phenomenon. This serves as confirmation that the pipeline captures physiologically meaningful signals, grounding the interpretation of subsequent null results.

State anxiety classification produced near-chance performance at the epoch level but a modest, consistent above-chance signal at the subject level (53–55%). This suggests that averaging across epochs surfaces an equally weak and stable relationship between spectral features and momentary anxiety. Trait anxiety classification fell below chance across all models and both levels of analysis, which this paper uses as evidence about the nature of trait anxiety: a dispositional construct defined by its stability across time that a single resting-state session fails to accurately predict.

Correlation analysis revealed that state and trait anxiety engage spectral features in partially opposing directions at frontal electrode sites, which suggests they are more complicated than being different magnitudes of the same underlying neural process. This dissociation carries a practical implication for applied EEG-based anxiety detection: systems that conflate state and trait constructs risk training on a signal that is orthogonal or even contrary to what they intend to measure.

Taken together, these results reframe the question of what resting-state EEG can do. Spectral features are reliable markers of momentary brain state. Their ability to index psychological constructs diminishes as those constructs become more temporally removed from the recording window. The bottleneck in EEG-based anxiety classification is therefore not model complexity, as no classifier in this study meaningfully outperformed the others, but feature expressiveness relative to the target variable's timescale. Future work should consider whether richer feature representations, longitudinal recording designs, or alternative anxiety operationalizations can bridge this gap.

Citations

Adolph, D., & Margraf, J. (2017). The differential relationship between trait anxiety, depression, and resting frontal α-asymmetry. *Journal of Neural Transmission, 124*(3), 379–386. https://doi.org/10.1007/s00702-016-1664-9

Al-Qazzaz, N. K., Ali, S. H. B. M., & Ahmad, S. A. (2025). Evaluating the Severity of Autism Spectrum Disorder from EEG: A Multidisciplinary Approach Using Statistical and Artificial Intelligence Frameworks. *Bioengineering (Basel)*, *12*(11), Article 1225. https://doi.org/10.3390/bioengineering12111225

Cavanagh, J. F., & Shackman, A. J. (2015). Frontal midline theta reflects anxiety and cognitive control: Meta-analytic evidence. *Journal of Physiology, Paris*, *109*(1–3), 3–15. https://doi.org/10.1016/j.jphysparis.2014.04.003

Cavanagh, J. F. (2021). *EEG: Probabilistic selection and depression* [Dataset]. OpenNeuro. https://doi.org/10.18112/openneuro.ds003474.v1.1.0

Chang, J., & Choi, Y. (2023). Depression diagnosis based on electroencephalography power ratios. *Brain and Behavior*, *13*(8), Article e3173. https://doi.org/10.1002/brb3.3173

Davidson, R. J. (1992). Anterior cerebral asymmetry and the nature of emotion. *Brain and Cognition, 20*(1), 125–151. https://doi.org/10.1016/0278-2626(92)90065-T

Shalamberidze, T., Nash, K., & Caplan, J. B. (2026). *Resting-state EEG and trait anxiety* [Dataset]. OpenNeuro. https://doi.org/10.18112/openneuro.ds007609.v1.0.0

Shalamberidze, T., Nash, K., & Caplan, J. B. (2025). Rhythmic activity in resting-state EEG predicts trait anxiety. *Imaging neuroscience (Cambridge, Mass.)*, *3*, IMAG.a.44. https://doi.org/10.1162/IMAG.a.44

Thibodeau, R., Jorgensen, R. S., & Kim, S. (2006). Depression, anxiety, and resting frontal EEG asymmetry: A meta-analytic review. *Journal of Abnormal Psychology, 115*(4), 715–729. https://doi.org/10.1037/0021-843X.115.4.715