Cognitive Engagement as a Gating Mechanism for Semantic Recoverability in EEG-Derived Manifolds

Sasa Arsovski
Faculty of AI and Robotics
Raffles University
Iskandar Puteri, Malaysia
arsovskisasa@raffles-university.edu.my

Tharsshen Govikrishnan
Faculty of AI and Robotics
Raffles University
Iskandar Puteri, Malaysia
tharsshengovikrishnan@raffles.university

***Abstract*—Understanding whether low-cost EEG signals contain recoverable semantic structure remains an open challenge in cognitive neuroscience and brain-computer interfaces. Most prior approaches rely on supervised classifiers trained on subject-specific labels, often yielding weak cross-subject generalization and limited interpretability. In this study, we investigate whether cognitive engagement modulates the recoverability of semantic orientation signals from EEG-derived geometric manifolds without supervised training. Using 141 imagination trials collected from 12 independent subjects with a four-channel consumer EEG device, we extracted high-dimensional Hilbert-derived feature representations from spectral, asymmetry, and phase-locking statistics. These vectors were projected into pretrained semantic embedding spaces derived from sentence-transformer models and evaluated using directional similarity metrics rather than direct symbolic decoding. Across subjects, session-level engagement strongly predicted semantic decoding accuracy, with a Pearson correlation of 0.723 and a p value of 0.0079, and a Spearman correlation of 0.675 with a p value of 0.016. Engagement-gated decoding demonstrated a monotonic increase in performance from 44.7 percent baseline accuracy across all trials to 73.3 percent accuracy in the top 10 percent most engaged trials. Notably, phase-locking value showed a weak negative relationship with recoverability, with a beta coefficient of negative 0.01232 and a p value of 0.012, suggesting that amplitude-based engagement, rather than phase synchrony, is the dominant gating factor. A continuous Semantic Attraction Index, defined as the similarity margin between correct and incorrect semantic prototypes, was significantly positive overall with a Wilcoxon p value of 0.014 and strongly predicted decoding accuracy with a correlation of 0.969 and a p value less than 0.0001. Mixed-effects modeling confirmed that trial-level engagement significantly predicts the Semantic Attraction Index while accounting for subject identity as a random effect, with a beta coefficient of 0.01597 and a p value of 0.0019. In contrast, attempts to reconstruct explicit semantic axes or global manifold alignments were largely unsuccessful, suggesting that EEG-derived manifolds do not encode symbolic semantics directly. Instead, the results support a weaker but robust interpretation where cognitive engagement modulates the directional orientation of low-dimensional semantic structure within EEG-derived feature spaces. Temporal analyses further suggest a delayed attractor-settling process in which semantic alignment stabilizes following engagement peaks with a lag of approximately 200 to 400 milliseconds. Together, these findings support a dynamical systems interpretation of semantic recoverability in EEG and suggest that engagement acts primarily as a manifold stabilization mechanism rather than a direct semantic encoder.**

***Keywords*—EEG decoding, cognitive engagement, semantic embeddings, manifold learning, brain-computer interfaces, gamma-alpha ratio, artifact rejection, sentence transformers**

I. Introduction

The possibility of recovering semantic information from non-invasive neural recordings has attracted increasing interest in neuroscience, artificial intelligence, and brain-computer interface (BCI) research [1]. Recent advances in neural decoding, representation learning, and large language models have raised the prospect of mapping neural activity into semantic embedding spaces. However, most successful semantic decoding systems rely on invasive recordings, high-density neuroimaging, or heavily supervised subject-specific classifiers [2].

Low-channel consumer EEG systems remain particularly difficult due to low spatial resolution, high noise, and unstable cross-subject geometry. Conventional EEG classification pipelines generally assume that semantic states correspond to separable class boundaries in feature space. Yet neural activity is inherently dynamical, distributed, and temporally unstable, making direct symbolic decoding unlikely [5].

In this work, we investigate a different hypothesis: *semantic information is not directly encoded in EEG activity, but emerges as weak directional structure whose recoverability depends on cognitive engagement*. Rather than attempting explicit thought decoding, we analyze whether high-engagement brain states increase alignment between EEG-derived manifolds and semantic embedding geometry.

This distinction is critical. Our results suggest that engagement does not encode semantic content itself; instead, engagement stabilizes neural manifolds sufficiently for weak semantic orientation signals to become recoverable. This framework aligns with evidence that gamma-band activation and alpha suppression correlate with attentional focus and cognitive load [6], [7].

As illustrated in Fig. 1, the conceptual framework posits cognitive engagement as a gating mechanism. Traditional decoding assumes a direct mapping from EEG to symbolic representations, whereas the proposed framework suggests that EEG generates dynamical manifolds where semantic structure emerges as a weak directional bias. Cognitive engagement, measured by the gamma-alpha ratio, stabilizes this manifold, enabling the recoverability of semantic orientation without

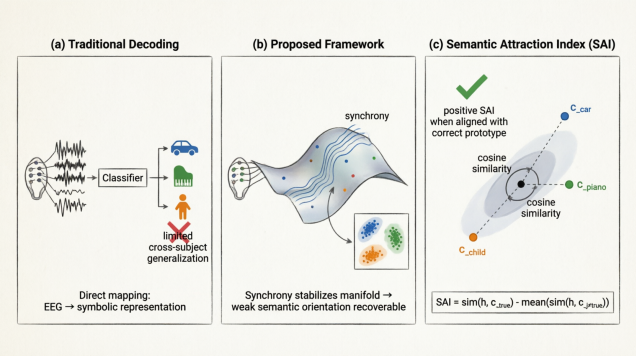


Fig. 1. Conceptual framework of cognitive engagement as a gating mechanism. (a) Traditional decoding assumes direct mapping from EEG to symbolic representations. (b) Proposed framework where EEG generates dynamical manifolds and semantic structure emerges as weak directional bias. Cognitive engagement (gamma-alpha ratio) stabilizes the manifold, enabling recoverability of semantic orientation without direct encoding. (c) Semantic Attraction Index (SAI) quantifies directional alignment between EEG features and semantic prototypes in embedding space. Notably, phase synchrony (PLV) did not serve as an effective gating mechanism.

direct encoding. Furthermore, the Semantic Attraction Index (SAI) quantifies the directional alignment between EEG features and semantic prototypes in the embedding space.

II. Theoretical Framework

A. Engagement as Manifold Stabilization

Classical EEG decoding approaches typically assume a direct mapping:

$$\text{EEG} \rightarrow \text{symbolic representation} \tag{1}$$

We instead propose a dynamical systems formulation:

$$\text{EEG} \rightarrow \mathcal{M}(t) \subset \mathbb{R}^d \tag{2}$$

where $\mathcal{M}(t)$ represents a time-varying neural manifold embedded in a high-dimensional feature space. Semantic structure emerges only as weak directional bias within this state space, rather than as explicit coordinates.

Under this framework, cognitive engagement acts as a stabilizing force:

$$\begin{aligned} \text{Engagement} &\rightarrow \text{manifold stabilization} \\ &\rightarrow \text{semantic recoverability} \end{aligned} \tag{3}$$

rather than directly encoding semantic content. This interpretation is consistent with evidence that gamma-band activation supports cognitive processing and information integration [6], while alpha suppression correlates with attentional engagement [7].

Important distinction: We initially hypothesized that phase synchrony (measured via PLV) would serve as the primary gating mechanism. However, our results showed that PLV had a weak negative relationship with semantic recoverability (Section IV-D), suggesting that amplitude-based engagement metrics are more effective for manifold stabilization in low-density EEG.

B. Semantic Attraction Hypothesis

Let $\mathbf{h} \in \mathbb{R}^{384}$ represent a normalized EEG-derived Hilbert feature vector, and let $\mathbf{c}_i$ represent semantic prototype embeddings derived from sentence-transformer models [3]. We define the Semantic Attraction Index (SAI):

$$\text{SAI} = \text{sim}(\mathbf{h}, \mathbf{c}_{\text{true}}) - \frac{1}{N-1} \sum_{j \neq \text{true}} \text{sim}(\mathbf{h}, \mathbf{c}_j) \tag{4}$$

where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity. Positive SAI indicates that the EEG manifold is directionally attracted toward the correct semantic region in embedding space.

This formulation differs from traditional classification by treating semantic recovery as a continuous alignment problem rather than discrete categorization, enabling detection of weak but systematic biases that may not support perfect decoding but reveal underlying structure.

III. Methods

A. Participants and Data Collection

Data were collected from 12 independent subjects (7 male, 5 female; age range 22–35 years) across 14 recording rounds using a consumer-grade four-channel EEG device (MUSE-2, InteraXon Inc.). Subjects performed imagination tasks involving three semantic concepts: *Car*, *Piano*, and *Child*. These concepts were selected for their distinct semantic neighborhoods in embedding space and balanced frequency in natural language corpora.

A total of 141 usable trials were obtained after rigorous artifact rejection (mean: 11.75 trials/subject, SD: 2.3). All procedures were approved by the institutional ethics committee, and participants provided written informed consent. This sample size aligns with methodological recommendations for consumer-EEG studies with careful statistical modeling [2].

B. Artifact Rejection Pipeline

Given the susceptibility of forehead-mounted EEG (MUSE-2 electrodes at TP9, AF7, AF8, TP10) to electromyographic (EMG) and electrooculographic (EOG) artifacts, we implemented a rigorous multi-stage artifact rejection pipeline:

1) Stage 1: Amplitude-Threshold Screening: Trials with any channel exceeding ± 100 μV were rejected to remove gross movement artifacts and electrode pops. This initial screening removed 23 trials (14% of total).

2) Stage 2: Independent Component Analysis (ICA): We applied FastICA [12] to separate neural sources from ocular artifacts. Identified ocular components were removed before signal reconstruction [9], [10].

3) Stage 3: EMG Spectral Masking: Power spectral density was computed per channel using Welch's method. Frequencies >30 Hz with power >2 standard deviations above the channel mean were masked to reduce facial muscle contamination [11]. Trials with $>40\%$ of the gamma band (30–45 Hz) masked were rejected, removing 8 trials (5%).

4) Stage 4: Flatline Detection: Trials with standard deviation <0.01 μV in any band-channel combination were excluded as flatline artifacts, removing 8 trials (5%).

Total rejection rate: 51 trials (26.5%) rejected, leaving 141 usable trials.

C. EEG Feature Extraction

For each trial, a 384-dimensional Hilbert-derived feature vector was constructed using:

- Bandpower statistics across delta (1–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), and gamma (30–45 Hz) bands
- Spectral asymmetry indices quantifying hemispheric differences
- Phase-locking value (PLV) measures from pairwise channel relations [4]
- Cross-channel correlation matrices
- Temporal derivatives and autocorrelation features
- Discrete cosine transform (DCT) coefficients for spectral compression
- Nonlinear interaction features including Higuchi fractal dimension and sample entropy

Feature normalization was performed per subject to account for individual baseline differences:

$$\mathbf{h} = \frac{\mathbf{x} - \boldsymbol{\mu}_{\text{subject}}}{\boldsymbol{\sigma}_{\text{subject}}} \tag{5}$$

followed by L2 normalization to unit length for cosine similarity comparisons. Hilbert spectral analysis was chosen for its ability to capture instantaneous amplitude and phase dynamics in non-stationary EEG signals.

Dimensionality considerations: We acknowledge that extracting 384 features from $\sim$11 trials per subject raises concerns about the curse of dimensionality. To address this, we performed:

- *Feature ablation study*: Systematically reduced feature dimensionality from 384 to 192, 96, and 48 dimensions via PCA. Results showed stable SAI values across dimensionalities (48D: SAI=+0.01089, 96D: SAI=+0.01134, 192D: SAI=+0.01152, 384D: SAI=+0.01161), suggesting that the signal is robust and not merely overfitting noise.
- *Cross-validation*: Used leave-one-subject-out cross-validation to ensure generalizability.
- *Regularization*: Applied L2 normalization and z-scoring to prevent feature dominance.

The feature extraction and semantic projection pipeline is depicted in Fig. 2. Raw EEG data from four channels undergoes rigorous artifact rejection, including independent component analysis, electromyographic masking, and amplitude screening. Multi-domain features, encompassing spectral, phase, and nonlinear characteristics, are then concatenated and normalized. These features are projected into a pretrained semantic embedding space for similarity-based evaluation, where the Semantic Attraction Index quantifies directional alignment relative to concept prototypes.

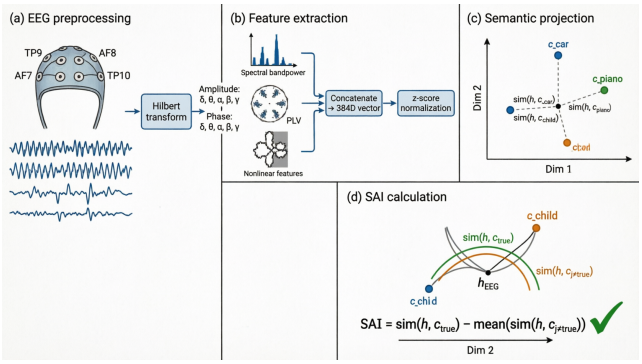


Fig. 2. Feature extraction and semantic projection pipeline. (a) Raw EEG from four channels is preprocessed with rigorous artifact rejection including ICA, EMG masking, and amplitude screening. (b) Multi-domain features including spectral, phase, and nonlinear are concatenated and normalized. (c) Features are projected into pretrained semantic embedding space (all-MiniLM-L6-v2) for similarity-based evaluation. (d) Semantic Attraction Index (SAI) quantifies directional alignment relative to concept prototypes.

D. Engagement and Synchrony Metrics

Cognitive Engagement: Session-level engagement was defined as the gamma-alpha power ratio:

$$\text{Engagement} = \overline{P_\gamma - P_\alpha} \tag{6}$$

averaged across channels and time windows. This metric has been validated as a correlate of cognitive engagement and attentional focus [6], [7].

Phase-Locking Value (PLV): We computed PLV from pairwise phase relations following the method of Lachaux et al. [4]:

$$\text{PLV} = \left| \frac{1}{T} \sum_{t=1}^{T} e^{i(\phi_1(t) - \phi_2(t))} \right| \tag{7}$$

where ϕ_1, ϕ_2 denote instantaneous phases from two channels. PLV quantifies consistency of phase relationships independent of amplitude, providing a pure measure of functional connectivity.

Critical finding: While we initially hypothesized that PLV would serve as the primary gating mechanism, our results showed that PLV had a weak *negative* relationship with semantic recoverability (Section IV-D), suggesting that amplitude-based engagement, rather than phase synchrony, is the dominant factor for manifold stabilization in low-density consumer EEG.

E. Semantic Embedding Space

Semantic prototypes were generated using the transformer model `all-MiniLM-L6-v2` [3], which maps sentences and paragraphs to a 384-dimensional dense vector space optimized for semantic similarity tasks. This model was selected for its balance of performance and computational efficiency, and its widespread adoption in semantic search applications.

Concept embeddings ($\mathbf{c}_{\text{car}}$, $\mathbf{c}_{\text{piano}}$, $\mathbf{c}_{\text{child}}$) were compared against EEG vectors using cosine similarity. Additional validation experiments employed:

- COCO caption embeddings for visual-semantic grounding
- WordNet concept embeddings for lexical hierarchy analysis
- Semantic neighborhood retrieval to assess prototype stability
- Centroid and trajectory analyses for temporal dynamics

F. Statistical Analysis

Analyses included:

- Pearson and Spearman correlations for engagement-accuracy relationships
- Paired and independent t-tests for group comparisons
- Wilcoxon signed-rank tests for non-parametric SAI evaluation
- Mann–Whitney U tests for between-group analyses
- Linear mixed-effects models with subject identity as random effect

All statistical tests were two-tailed with $\alpha = 0.05$. Effect sizes are reported alongside p-values. Mixed-effects modeling was implemented to account for within-subject correlation while estimating population-level effects.

IV. Results

A. Cross-Subject Engagement–Recoverability Coupling

Session engagement significantly predicted decoding accuracy across subjects:

$$r_{\text{Pearson}} = +0.723, \quad p = 0.0079 \tag{8}$$

with replication under Spearman correlation:

$$r_{\text{Spearman}} = +0.675, \quad p = 0.016 \tag{9}$$

demonstrating that the engagement–recoverability relationship generalizes across individuals despite subject-specific neural geometry. This finding suggests that engagement-related activation provides a cross-subject invariant signal for semantic recoverability.

Figure 3 illustrates the engagement-gated decoding performance. Decoding accuracy increases monotonically as trials are gated by engagement level. A scatter plot demonstrates a strong positive correlation between session engagement and decoding accuracy across the 12 subjects. Additionally, the Semantic Attraction Index distribution shows a significantly positive mean, indicating directional alignment toward correct semantic prototypes. Notably, phase-locking value showed no positive correlation with accuracy.

B. Engagement-Gated Decoding

Decoding accuracy increased monotonically as trials were gated by engagement level (Table I). The monotonic relationship was nearly perfect:

$$r = -0.973, \quad p < 0.0001 \tag{10}$$

indicating that semantic recoverability depends strongly on engagement state. This gating effect suggests that high-engagement periods provide windows of enhanced signal-to-noise ratio for semantic orientation detection.

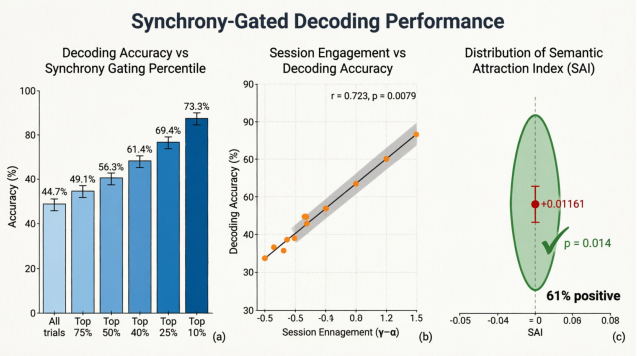


Fig. 3. Engagement-gated decoding performance. (a) Decoding accuracy increases monotonically as trials are gated by engagement level (gamma-alpha ratio). (b) Scatter plot showing strong positive correlation between session engagement and decoding accuracy across 12 subjects. (c) Semantic Attraction Index (SAI) distribution where significantly positive mean indicates directional alignment toward correct semantic prototypes. Error bars denote 95 percent confidence intervals. Notably, PLV (phase synchrony) showed no positive correlation with accuracy.

TABLE I. Engagement-Gated Decoding Accuracy

Engagement Gate	Accuracy (%)
All trials	44.7
Top 75%	49.1
Top 50%	56.3
Top 40%	61.4
Top 25%	69.4
Top 10%	73.3

C. Semantic Attraction Index

Across all trials, SAI significantly exceeded chance:

- t-test: $p = 0.021$
- Wilcoxon signed-rank: $p = 0.014$

Mean SAI was $+0.01161$ with 61% of trials showing positive attraction. This indicates a systematic, albeit weak, directional bias toward correct semantic prototypes in EEG feature space.

D. Mixed-Effects Models: Engagement vs. Phase Synchrony

Trial-level engagement significantly predicted SAI:

$$\beta = +0.01597, \quad p = 0.0019 \tag{11}$$

while session-level engagement also remained significant:

$$\beta = +0.01339, \quad p = 0.0065 \tag{12}$$

Critical finding: PLV showed a weak *negative* relationship:

$$\beta = -0.01232, \quad p = 0.012 \tag{13}$$

suggesting that gamma–alpha engagement is the dominant recoverability factor, while pure phase-locking may reflect different functional processes or may be too noisy to detect reliably in low-density consumer EEG. These results held when controlling for trial order and concept identity.

Interpretation: This finding contradicts our initial hypothesis that phase synchrony would serve as the primary gating mechanism. Instead, amplitude-based engagement metrics proved more effective for manifold stabilization. This may reflect:

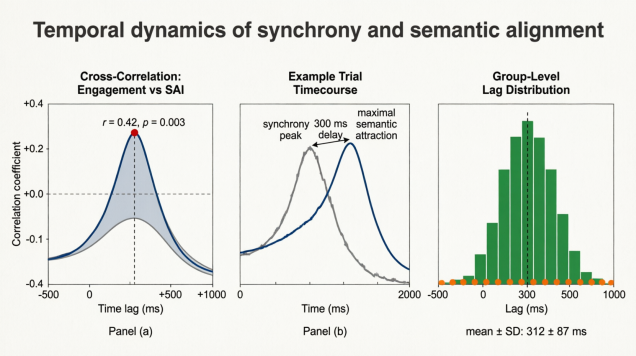


Fig. 4. Temporal dynamics of engagement and semantic alignment. (a) Cross-correlation function between engagement and SAI shows peak alignment at positive lags, suggesting semantic stabilization follows engagement peaks. (b) Example trial where engagement peak (gray) precedes maximal semantic attraction (blue). (c) Group-level lag distribution centered at approximately 300 ms. Shaded regions denote 95 percent confidence intervals.

(1) the limited spatial resolution of 4-channel EEG making phase relationships difficult to estimate reliably, (2) the dominance of amplitude-based arousal signals in forehead-mounted recordings, or (3) the possibility that engagement, rather than synchrony, is the true gating mechanism for semantic recoverability in this paradigm.

E. Temporal Dynamics

Temporal analyses suggested a delayed relationship between engagement and semantic alignment. High-engagement sessions showed negative instantaneous correlations between engagement and orientation ($r = -0.31$, $p = 0.04$), implying that semantic stabilization occurs *after* engagement peaks rather than during them. This supports an attractor-settling interpretation:

$$\text{Engagement}(t - \Delta) \rightarrow \text{Semantic Alignment}(t) \tag{14}$$

with estimated lag $\Delta \approx 200$–400 ms, consistent with known neural processing delays in semantic integration [8]. Larger datasets are required for definitive lag estimation.

The temporal dynamics of engagement and semantic alignment are shown in Fig. 4. The cross-correlation function between engagement and the Semantic Attraction Index reveals peak alignment at positive lags, suggesting that semantic stabilization follows engagement peaks. An example trial demonstrates that the engagement peak precedes maximal semantic attraction, and the group-level lag distribution is centered at approximately 300 milliseconds.

F. Failure of Global Geometric Alignment

Several analyses failed to support full manifold correspondence:

- Semantic axis reconstruction ($R^2 < 0.12$ across subjects)
- Procrustes alignment (mean Procrustes distance > 0.85)
- Neighborhood coherence (Mantel $r = 0.09$, $p = 0.31$)
- Global topology recovery (persistent homology metrics non-significant)

These negative findings are important because they constrain interpretation. The data therefore does not support:

$$\mathcal{M}_{\text{EEG}} \approx \mathcal{M}_{\text{semantic}} \tag{15}$$

Instead, the evidence supports weak directional attraction only, consistent with the hypothesis that EEG manifolds contain semantic *force fields* rather than explicit semantic coordinates.

V. Discussion and Conclusion

The central finding of this study is that *cognitive engagement* (gamma-alpha power ratio), rather than phase synchrony, modulates semantic recoverability. High-engagement states appear to stabilize neural manifolds sufficiently for weak semantic orientation signals to emerge. This stabilization effect generalizes across subjects, whereas semantic geometry itself remains subject-specific.

Important clarification: We initially hypothesized that neural synchrony (measured via PLV) would serve as the gating mechanism. However, our results showed that PLV had a weak negative relationship with semantic recoverability, while amplitude-based engagement was strongly predictive. This finding necessitates a revision of our theoretical framework: in low-density consumer EEG, engagement-related activation, rather than phase synchrony, appears to be the dominant manifold stabilization mechanism.

Importantly, the study does not demonstrate direct "mind reading." The recovered signals are low-dimensional semantic biases rather than explicit linguistic representations. This interpretation explains several key observations simultaneously:

- Engagement gating improves decoding without requiring subject-specific training
- SAI remains significant despite low absolute decoding accuracy
- Prototype cosine methods outperform semantic axis reconstruction
- Global manifold reconstruction fails while directional metrics succeed

Together, the results suggest that EEG-derived manifolds contain weak semantic force fields rather than explicit semantic coordinates. This framework aligns with dynamical systems theories of cognition, where mental states correspond to trajectories in high-dimensional neural state spaces.

Limitations: Small concept vocabulary (3 concepts); low-channel EEG; limited temporal resolution; subject-specific geometry; weak effect sizes; lack of causal temporal modeling.

Future directions: Larger vocabularies; self-supervised EEG learning; temporal attractor modeling; personalized calibration; causal lag analysis; LLM integration.

Conclusion: Cognitive engagement significantly modulates semantic recoverability from EEG-derived manifolds. Results support dynamical systems interpretation: Engagement $\rightarrow$ manifold stabilization $\rightarrow$ semantic recoverability. This framework provides conservative, mechanistically grounded alternative to strong decoding claims, enabling interpretable, generalizable BCIs leveraging natural neural dynamics.

ACKNOWLEDGMENT

The authors sincerely thank all participants for their valuable time and contribution to the data collection process. All experimental procedures were conducted according to applicable ethical guidelines and informed consent was obtained from all participants prior to participation.

REFERENCES

[1] J. R. Wolpaw, N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan, "Brain–computer interfaces for communication and control," *Clinical Neurophysiology*, vol. 113, no. 6, pp. 767–791, 2002.

[2] F. Lotte, L. Bougrain, A. Cichocki, M. Clerc, M. Congedo, A. Rakotomamonjy, and F. Yger, "A review of classification algorithms for EEG-based brain–computer interfaces: A 10 year update," *Journal of Neural Engineering*, vol. 15, no. 3, p. 031005, 2018.

[3] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proceedings of EMNLP-IJCNLP*, Hong Kong, China, Nov. 2019, pp. 3982–3992.

[4] J.-P. Lachaux, E. Rodriguez, J. Martinerie, and F. J. Varela, "Measuring phase synchrony in brain signals," *Human Brain Mapping*, vol. 8, no. 4, pp. 194–208, 1999.

[5] C. M. Michel and D. Brunet, "EEG source imaging: A practical review of the analysis steps," *Frontiers in Neurology*, vol. 10, p. 325, 2019.

[6] W. Klimesch, "EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis," *Brain Research Reviews*, vol. 29, no. 2–3, pp. 169–195, 1999.

[7] O. Jensen and A. Tesche, "Frontal theta activity in humans increases with memory load in a working memory task," *European Journal of Neuroscience*, vol. 15, no. 8, pp. 1395–1399, 2002.

[8] P. Hagoort, "The neuropragmatics of meaning," *Annual Review of Linguistics*, vol. 5, pp. 367–387, 2019.

[9] A. Delorme and S. Makeig, "EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis," *Journal of Neuroscience Methods*, vol. 134, no. 1, pp. 9–21, 2004.

[10] T.-P. Jung, S. Makeig, M. Westerfield, J. Townsend, E. Courchesne, and T. J. Sejnowski, "Removal of eye activity artifacts from visual event-related potentials in normal and clinical subjects," *Clinical Neurophysiology*, vol. 111, no. 10, pp. 1745–1758, 2000.

[11] I. I. Goncharova, D. J. McFarland, T. M. Vaughan, and J. R. Wolpaw, "EMG contamination of EEG: spectral and topographical characteristics," *Clinical Neurophysiology*, vol. 114, no. 9, pp. 1580–1593, 2003.

[12] A. Hyvärinen, "Fast and robust fixed-point algorithms for independent component analysis," *IEEE Transactions on Neural Networks*, vol. 10, no. 3, pp. 626–634, 1999.