CASE: Cortex-Aware Spatial Embedding for EEG Motor Imagery Classification

Zongyao Feng, Konstantin Markov
Department of Computer and Information Systems
University of Aizu
Aizu, Japan
{d8252101, markov}@u-aizu.ac.jp

2025 13th International Conference on Awareness Science and Technology (iCAST) | 979-8-3315-7581-6/25/$31.00 ©2025 IEEE | DOI: 10.1109/iCAST68191.2025.11300197

***Abstract*—Motor imagery (MI) classification is a core task in brain–computer interfaces (BCIs), but learning robust spatial representations from electroencephalography (EEG) is hindered by low signal-to-noise ratios and limited training data. Existing approaches, rely on purely data-driven spatial filters that may overlook well-established neuroanatomical relationships between electrodes and cortical functions. We propose Cortex-Aware Spatial Embedding (CASE), a lightweight and generalizable method that injects neuroscience priors into EEG classification models. CASE uses a large language model with a neuroscience-oriented prompt to generate concise, task-aware descriptions for each electrode. These descriptions are converted into dense semantic vectors via a frozen biomedical text encoder, aligned to the model's feature space through a small MLP, and integrated into the backbone via learnable weighted addition. Evaluated on the BCI Competition IV–2a dataset under an inter-session protocol, CASE consistently improves classification accuracy and Cohen's kappa across common EEG baselines in MI. Our results demonstrate that incorporating LLM-derived spatial priors can effectively bridge domain knowledge and data-driven learning, yielding more robust MI decoding.**

***Index Terms*—Brain–Computer Interface (BCI), Motor Imagery (MI) Classification, Electroencephalography (EEG), Spatial Representation Learning, Spatial Identity Embedding**

I. INTRODUCTION

Brain-computer interfaces (BCIs) enable direct communication between the human brain and external devices, bypassing traditional neuromuscular pathways. This technology holds profound significance for restoring motor functions in individuals with disabilities, enhancing human-computer interaction, and advancing neurorehabilitation [1], [2]. Motor imagery classification, a key paradigm in BCIs, involves decoding imagined movements from neural signals to control devices like prosthetic limbs or cursors [3]. Real-world applications span from assistive technologies for paralyzed patients to immersive gaming and augmented reality systems. Advancing motor imagery classification could transform rehabilitation outcomes and expand accessible computing [4]. Decoding electroencephalography (EEG) signals presents formidable challenges, including low signal-to-noise ratios that obscure subtle patterns [5]. High inter-session and inter-subject variations further complicate reliable performance across different recordings or users [6]. Although EEG offers high temporal resolution for capturing rapid brain dynamics, its low spatial resolution limits precise localization of activity sources. Rich information still underlies the spatial domain, where electrode placements reflect distinct functional roles in motor processing [7].

Recent advances in motor imagery classification often build on the spatial modeling strategies pioneered by EEGNet [8]. These methods typically employ spatial 1D convolutions with kernel sizes matching the number of EEG electrodes, effectively learning various weighted combinations of channel signals. Such an approach proves robust because the fixed kernels remain independent of specific inputs, mitigating some inter-session variability through consistent feature extraction. Models can generalize better across varying conditions without overfitting to transient noise. Yet this simplicity comes at a cost. With limited training data, a single convolutional layer struggles to capture the intricate spatial contexts inherent in brain topography. Complex relationships between electrodes, grounded in neuroanatomy, often elude discovery, leading to suboptimal filters that fail to leverage domain knowledge.

Motor imagery relies on coordinated activity across distinct brain regions, each contributing uniquely based on the task at hand [9], [10]. For instance, imagining hand movements typically activates contralateral sensorimotor areas, with electrodes like C3 and C4 detecting these patterns over the left and right hemispheres. Foot imagery, by contrast, engages midline regions around Cz. These reflects the somatotopic organization of the motor cortex [11]. The functional role of any given electrode thus emerges from both the mental task and its spatial position on the scalp, following the 10-20 system [12]. Models that incorporate this prior knowledge can better discern task-relevant signals amid noise, leading to more stable spatial filters and enhanced classification. Manually defining these roles demands deep neuroscience expertise, often involving time-consuming literature reviews or empirical mapping. This expertise barrier limits scalability. Large language models, however, offer a solution. Through targeted prompting, they can distill established neuroanatomical insights into electrode-specific descriptions, injecting spatial context without manual intervention.

To overcome the shortcomings of data-driven spatial modeling, we introduce a simple yet effective method: Cortex Aware Spatial Embedding (CASE), which extracts neuroscience priors of recorded regions by EEG electrodes from language models, then injects them into EEG models. Specifically, a neuroscience-oriented persona in a large language model generates task-aware text descriptions for each EEG electrode. These descriptions pass through a frozen text encoder to yield embeddings. A simple multi-layer perceptron (MLP) is used to align CASE with the model's low-level feature space. The embeddings are then integrated via weighted addition into standard EEG backbones like EEGNet [8] or EEG-TCNet [13], with end-to-end training preserving the efficiency. This approach enhances spatial learning under low SNR and limited data, bridging neuroscience with machine learning. Empirical evaluations of CASE demonstrate consistent gains in both accuracy and Cohen's Kappa coefficient across different baselines.

Our main contributions include the following.

- We propose CASE, a novel framework for enhancing the performance of EEG models using Cortex Aware Spatial Embedding , requiring no additional data or complex engineering.
- We show that integration of CASE enhances baseline architectures performance on a standard motor imagery benchmark, with average improvements in both accuracy and Cohen's Kappa coefficient.
- Our work highlights the potential of LLMs as knowledge bridges for neuroscience-informed AI, opening avenues for building more effective BCI applications.

II. Related Works

Motor imagery classification has been advanced by compact, domain-specific architectures that balance accuracy with efficiency on limited EEG datasets. EEGNet [8] pioneered this trend by combining temporal convolutions for frequency-specific feature extraction with depthwise spatial convolutions spanning all electrodes, enabling low-parameter learning of spectral and spatial patterns. EEG-TCNet [13] extended this design by adding dilated temporal convolutional networks, capturing long-range temporal dependencies more effectively. TCNetFusion [14] further improved performance by integrating multiple processing paths and fusing their outputs, thereby strengthening temporal modeling while maintaining focus to spatial features. While differing in temporal modeling strategies, these approaches all rely on purely data-driven spatial convolution filters learned from noisy EEG recordings, which can limit robustness under low signal-to-noise ratios and small training sets.

In other domains, incorporating explicit node identity information has proven effective for spatial learning. For example, in traffic forecasting, node identity encoding assigns each sensor a fixed embedding representing its physical or functional role in the network [15], improving forecasting performance across varying traffic conditions. This concept remains largely unexplored in EEG, despite the fact that electrode positions follow the standardized 10–20 or 10-10 system and are tied to well-characterized functional specializations in the motor cortex. Embedding this spatial prior directly into EEG models could help recover spatial patterns that better align with cortical functional interactions.

Large language models (LLMs) and domain-specific encoders provide a scalable way to generate such priors without manual annotation. Sentence transformers [16] map natural language descriptions into dense semantic embeddings, while biomedical models like BioBERT [17] capture specialized terminology and relationships from large-scale biomedical text corpora. By prompting an LLM with neuroscience knowledge, one can automatically produce electrode-specific descriptions grounded in biomedical literature. Passing these descriptions through a frozen biomedical encoder yields rich embeddings that can be injected into EEG models. However, this cross-modal strategy has rarely been seen in BCI applications, which could offer a promising bridge between established neuroscience and data-driven learning.

III. Methodology

We introduce CASE (Cortex-Aware Spatial Embedding), a novel framework that infuses neuroscience priors into EEG models via task-aware semantic vectors for each electrode. Derived from LLMs and biomedical encoders, these embeddings enrich the model's spatial feature space. As shown in Figure 1, the method involves three key stages: (1) LLM-based generation of electrode descriptions, (2) encode electrode description into dense embeddings, and (3) injection into the backbone's features. This approach is plug-and-play, requiring minimal parameters and preserving backbone efficiency.

A. Problem Formulation

Let an EEG trial be $X \in \mathbb{R}^{C\times T}$, with C electrodes, T time samples, and label $y \in \{1, \ldots, K\}$ for K motor imagery classes. The EEG classification model can be expressed as:

$$\hat{\mathbf{p}} = g_\theta \left(f_\theta(X) \right), \tag{1}$$

where f_θ extracts low-level features, g_θ handles higher modules and classification, and $\hat{\mathbf{p}} \in [0, 1]^K$ is the softmax output. Purely data-driven models learn spatial patterns from X alone, often yielding suboptimal filters under low SNR. CASE enhances $f_\theta(X)$ with fixed spatial priors for effectiveness.

B. Electrode Description Generation

We set up Claude 4 with a neuroscience-specialized persona using the prompt in Table II, this is to ensure task-relevance and anatomical accuracy. Then for each electrode $i \in \{1, \ldots, C\}$, we ask the LLM to generate a text description in JSON format, in which the **embedding_text** captures the functional role in motor imagery, we denote the **embedding_text** as d_i. An example of the text description for motor imagery classification is presented in Table I. The outputs form $\mathcal{D} = \{d_1, \ldots, d_C\}$. This automates the prior extraction from the neuroscience domain, avoiding manual effort and the requirement of domain expertise.

C. Embedding Extraction

Descriptions $\mathcal{D}$ are then encoded using a frozen BioBERT-based sentence transformer $\phi(\cdot)$. Specifically, we use the fine-tuned BioBERT checkpoint provided in [18], which builds

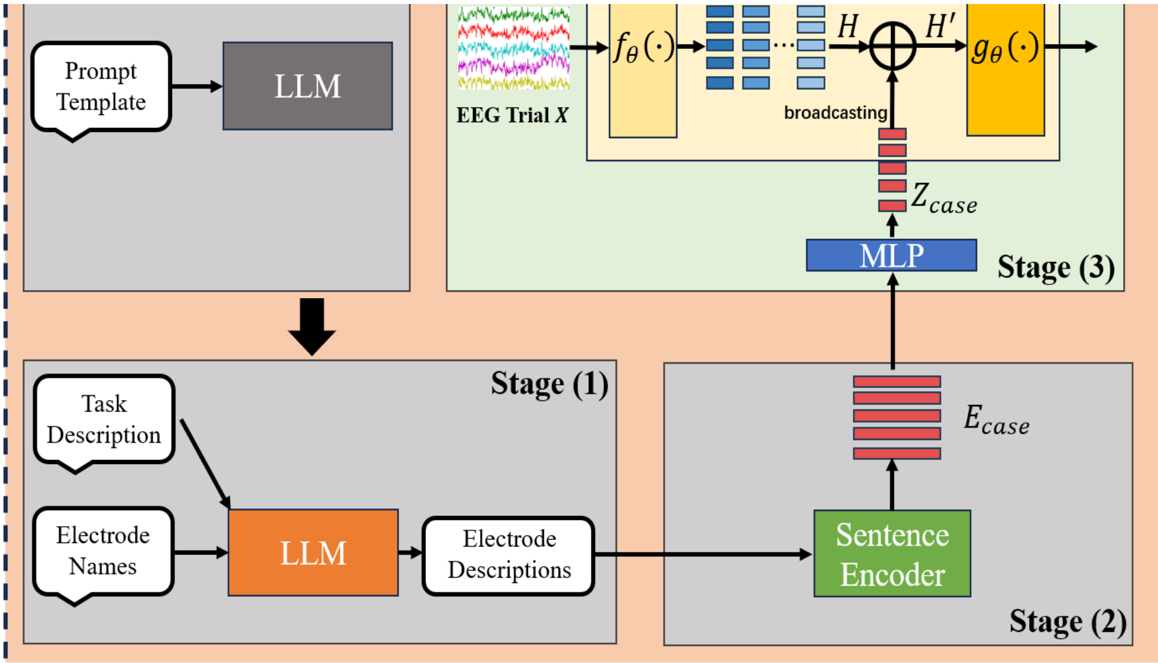

Fig. 1. Overview of the proposed Cortex-Aware Spatial Embedding (CASE) pipeline for EEG motor imagery classification. The pipeline consists of three stages: **Stage (1)** — Electrode Description Generation (grey, frozen): A large language model (LLM) is set up with a neuroscience-oriented prompt template. For each electrode in the standard 10–20/10-10 montage, the LLM produces concise, task-aware text descriptions capturing functional relevance for motor imagery. **Stage (2)** — Embedding Extraction (grey, frozen): The electrode descriptions are encoded into dense semantic vectors E_{case}, via a text encoder (BioBERT-based) to capture domain-specific semantics. **Stage (3)** — Injection into EEG Backbone (green, trainable): A MLP maps E_{case} to the backbone's low-level feature space, producing aligned spatial priors Z_{case}. The aligned priors are broadcast along the temporal dimension and integrated into the backbone's low-level feature maps $H = f_\theta(X)$ via learnable weighted addition $H' = H + \lambda Z_{\text{case}}$, where λ is a trainable scalar controlling prior influence. The enriched features are then processed by the remaining backbone layers $g_\theta(\cdot)$ to yield the class probabilities.

TABLE I
EXAMPLE OF GENERATED ELECTRODE DESCRIPTION FOR ELECTRODE C1

Description for Electrode C1
electrode: *C1*   **hemisphere:** *left*   **lobe:** *central*   **approx_region:** *left medial motor cortex*   **primary_functions:** *["motor control", "lower limb representation"]*   **task_relevance:** *Captures medial motor cortex activity. Shows ERD for feet MI due to proximity to leg motor area. Contributes to feet vs. hand discrimination.*   **embedding_text:** *Left paramedian central over medial motor cortex. Enhanced mu/beta ERD during feet motor imagery due to lower limb representation proximity. Important for feet class detection in MI-BCI.*

on BioBERT-base and is further adapted as a sentence-transformer for biomedical semantic similarity tasks. This choice leverages BioBERT's biomedical domain knowledge to capture nuanced embeddings:

$$e_i = \phi(d_i) \in \mathbb{R}^E. \tag{2}$$

Stacking them then yields $E_{\text{case}} = [e_1^\top; \dots; e_C^\top] \in \mathbb{R}^{C\times E}$.

D. Embedding Alignment and Injection into Backbone

A learnable MLP : $\mathbb{R}^E \to \mathbb{R}^F$, which consists of 2 layers with Hardswish activation, aligns E to the backbone's feature dimension F:

$$z_i = \text{MLP}(e_i), \quad Z_{\text{case}} = [z_1^\top; \dots; z_C^\top] \in \mathbb{R}^{C\times F}. \tag{3}$$

This alignment projects semantic priors into the model's low-level feature space during training. Let $H = f_\theta(X) \in \mathbb{R}^{C\times T\times F}$ be features after initial processing (e.g., post-temporal convolution in EEGNet). CASE injects via a weighted addition, broadcasting Z_{case} over time T:

$$H' = H + \lambda Z_{\text{case}} \tag{4}$$

where $\lambda \in \mathbb{R}$ is a learnable scalar that controls the prior's influence, which is initialized to 0.1. This empirical choice offered stable training behavior across subjects in practice. The enriched representation H' is then processed by g_θ to yield final prediction:

$$\hat{\mathbf{p}} = g_\theta(H'). \tag{5}$$

This injection aligns spatial priors with learned features while preserving the remaining processing of the backbone model.

IV. EXPERIMENTS

A. Dataset

We evaluate our method on the BCI Competition IV–2a dataset [19], a widely used benchmark for motor imagery classification. The dataset contains EEG recordings from nine subjects performing four motor imagery (MI) tasks: left hand,

TABLE II
PROMPT FOR SETTING UP THE LLM TO GENERATE ELECTRODE DESCRIPTIONS

Prompt

"You are a neuroscientist specializing in human EEG, the 10–20/10–10 electrode systems, and cortical functional organization (sensorimotor, premotor/SMA, somatosensory, visual, auditory, prefrontal, parietal association). Your job: given an EEG task and one or more electrode labels, produce compact, task-relevant descriptions per electrode that are useful for modeling (not clinical advice). You must be precise, concise, and conservative: if uncertain, say 'unknown' or lower your confidence—do not invent anatomy or functions."

Constraints:

Assume standard adult scalp EEG with international 10–20/10–10 names (e.g., Fz, Cz, C3, C4, Pz, POz, CP3).

Use approximate cortical mappings (e.g., "over primary motor cortex hand area"); do not claim exact localization.

Emphasize aspects that matter for the given task.

Keep each electrode's description $\leq$ 60 words and informative enough for a text encoder.

Output fields (per electrode):

electrode: label (e.g., C3)

hemisphere: left — right — midline — unknown

lobe: frontal — central — parietal — temporal — occipital — unknown

approx_region: short phrase (e.g., primary motor cortex hand area / SMA / somatosensory cortex / parietal association / visual cortex)

primary_functions: short phrases (task-agnostic roles)

task_relevance: 1–2 sentences focused on the provided task (mechanisms, expected patterns).

embedding_text: $\leq$ 60-word compact summary optimized for text embeddings; include task keywords if helpful.

right hand, both feet, and tongue movements. For each subject, two sessions were recorded on different days. Each session includes 48 trials per class, yielding 288 trials per session.

EEG signals were recorded from $C = 22$ electrodes placed according to the international 10–20 system at a sampling rate of 250 Hz, and were bandpass filtered between 0.5 Hz and 100 Hz. Following the standard MI timing paradigm, we extract the 2–6 s interval from each trial, corresponding to the motor imagery period. This results in $T = 1000$ time points per trial. No additional preprocessing such as artifact removal or bandpassing is applied.

We adopt an inter-session evaluation protocol, for each subject, the first session is used for training and validation, and the second session is reserved for testing. The first session is split into 80% training and 20% validation sets, with class-balanced sampling in each split.

B. Settings

All experiments are implemented in TensorFlow and executed on a workstation equipped with an NVIDIA RTX 4090 GPU and an AMD EPYC 7351P 16-core CPU. Models are trained using the Adam optimizer with an initial learning rate of 0.001. A ReduceLROnPlateau scheduler monitors the validation loss and decreases the learning rate by a factor of 0.9 if no improvement is observed for 20 consecutive epochs. The batch size is set to 64, and training is capped at a maximum of 1000 epochs. To mitigate overfitting, we apply early stopping with a patience of 40 epochs based on validation loss. The model checkpoint with the lowest validation loss is retained for final testing.

C. Baselines and Evaluation Metrics

We assess the effectiveness of CASE by integrating it into three representative motor imagery classification backbones: EEGNet [8], EEG-TCNet [13], and TCNetFusion [14]. These models are chosen for their strong performance on BCI benchmarks, as detailed in Section II. We reproduced all baselines using the default hyperparameters reported in their respective publications, applied consistently to every subject in the BCI Competition IV–2a dataset [19]. For each backbone, we compare the performance of the original model with its CASE–enhanced counterpart, keeping all other training settings identical.

Performance is evaluated for each subject using two metrics: classification accuracy and Cohen's kappa coefficient. Cohen's kappa [20] accounts for agreement occurring by chance and is defined as:

$$\kappa = \frac{p_o - p_e}{1 - p_e}, \tag{6}$$

where p_o is the observed agreement between predictions and ground truth, and p_e is the expected agreement by chance. Unlike accuracy, κ provides a chance-corrected measure of reliability, which is particularly relevant for multi-class problems with balanced datasets.

V. RESULTS

Table III reports classification accuracy and Cohen's κ for all subjects and models, with and without CASE. Across all three backbones, integrating CASE consistently improves the average performance. EEGNet achieves a mean accuracy of 73.64% and κ of 0.65 in its vanilla form, which increases to 74.96% and 0.67 with CASE. EEG-TCNet improves from 74.07% and 0.65 to 76.08% and 0.68 with CASE. The largest gains are observed for TCNetFusion, where the average accuracy rises from 75.73% to 78.36% and κ from 0.68 to 0.71.

The improvements, however, are not uniform across subjects. For EEGNet, CASE yields higher accuracy in 4 out of 9 subjects, with the largest gain of +8.68% for Subject 2. For EEG-TCNet, gains are observed in 7 subjects, most notably Subject 6 (+13.19% accuracy, +0.18 κ) and Subject 5 (+3.13% accuracy, +0.04 κ). TCNetFusion benefits the most from CASE, with consistent improvements across 7 subjects, and a +2.98% gains on the averaged accuracy across 9 subjects .

Overall, these results indicate that CASE enhances a variety of MI classification backbones by enriching spatial representations with neuroscience priors. Notably, the relative benefit is greatest for TCNetFusion, followed by EEG-TCNet and EEGNet, suggesting that spatial priors are most effective when combined with models that employ stronger temporal processing while placing emphasis on the spatial domain.

TABLE III
PERFORMANCE COMPARISON ON BCI COMPETITION IV–2A

Subject	EEGNet†		EEGNet+CASE		EEG-TCNet†		EEG-TCNet+CASE		TCNetFusion†		TCNetFusion+CASE	
	ACC(%)	κ	ACC(%)	κ	ACC(%)	κ	ACC(%)	κ	ACC(%)	κ	ACC(%)	κ
1	78.12	0.71	77.78	0.70	82.29	0.76	84.38	0.79	81.60	0.75	86.46	0.82
2	48.61	0.31	57.29	0.43	57.99	0.44	52.08	0.36	60.42	0.47	61.46	0.49
3	88.89	0.85	90.62	0.88	92.36	0.90	93.40	0.91	90.62	0.88	93.40	0.91
4	74.31	0.66	73.26	0.64	65.28	0.54	68.06	0.57	67.01	0.56	68.75	0.58
5	64.58	0.53	70.14	0.60	73.61	0.65	76.74	0.69	69.44	0.59	75.35	0.67
6	61.81	0.49	59.03	0.45	47.57	0.30	60.76	0.48	62.15	0.50	61.81	0.49
7	84.38	0.79	82.64	0.77	85.76	0.81	87.15	0.83	87.50	0.83	87.15	0.83
8	82.64	0.77	81.25	0.75	77.78	0.70	81.60	0.75	81.60	0.75	85.42	0.81
9	79.51	0.73	82.64	0.77	84.03	0.79	80.56	0.74	81.25	0.75	85.42	0.81
Avg	73.64	0.65	74.96	0.67	74.07	0.65	76.08	0.68	75.73	0.68	78.36	0.71

† Reproduced.

VI. CONCLUSION

In this work, we introduced CASE, a simple yet effective approach for enhancing spatial learning in EEG-based motor imagery classification by integrating neuroscience priors into existing backbone architectures. By leveraging large language models to generate task-specific electrode descriptions and a frozen biomedical text encoder to produce semantic embeddings, CASE provides a plug-and-play means of embedding cortical functional context without manual annotation or additional recordings. Through weighted feature injection, CASE enriches low-level representations, leading to consistent improvements in both accuracy and Cohen's kappa coefficient across diverse architectures.

Our experiments on the BCI Competition IV–2a dataset confirm that CASE improves all tested backbones, with the most substantial gains observed for TCNetFusion, followed by EEG-TCNet and EEGNet. These results indicate that CASE eases the challenge of spatial learning in EEG classification. Moreover, the consistent trend suggests that CASE is particularly effective when combined with architectures capable of modeling the temporal evolution of spatial patterns.

For the future works, we look forward to see CASE being extended to other EEG task, architectures, or multi-modal learning, and being further enriched with task-general neuroscience knowledge or other fusion schemes.

REFERENCES

[1] J. Ognard, G. El Hajj *et al.*, "Advances in Endovascular Brain Computer Interface: Systematic Review and Future Implications," *Journal of Neuroscience Methods*, vol. 420, p. 110471, 2025.
[2] X.-Y. Liu, W.-L. Wang *et al.*, "Recent Applications of EEG-based bBrain-computer-interface in The Medical Field," *Military Medical Research*, vol. 12, no. 1, p. 14, 2025.
[3] I. Hameed, D. M. Khan *et al.*, "Enhancing Motor Imagery EEG Signal Decoding Through Machine Learning: A Systematic Review of Recent Progress," *Computers in Biology and Medicine*, vol. 185, p. 109534, 2025.
[4] A. Zimmermann-Schlatter, C. Schuster *et al.*, "Efficacy of Motor Imagery in Post-stroke Rehabilitation: A Systematic Review," *Journal of Neuroengineering and Rehabilitation*, vol. 5, no. 1, p. 8, 2008.
[5] M. Rashid, N. Sulaiman *et al.*, "Current Status, Challenges, and Possible Solutions of EEG-based Brain-computer Interface: A Comprehensive Review," *Frontiers in Neurorobotics*, vol. 14, p. 25, 2020.
[6] S. Saha and M. Baumert, "Intra-and Inter-subject Variability in EEG-based Sensorimotor Brain Computer Interface: A Review," *Frontiers in Computational Neuroscience*, vol. 13, p. 87, 2020.
[7] W.-K. Tam, T. Wu *et al.*, "Human Motor Decoding from Neural Signals: A Review," *BMC Biomedical Engineering*, vol. 1, no. 1, p. 22, 2019.
[8] V. J. Lawhern, A. J. Solon *et al.*, "EEGNet: A Compact Convolutional Neural Network for EEG-based Brain-computer Interfaces," *Journal of Neural Engineering*, vol. 15, no. 5, p. 056013, 2018.
[9] R. M. Hardwick, S. Caspers *et al.*, "Neural Correlates of Motor Imagery, Action Observation, and Movement Execution: A Comparison Across Quantitative Meta-analyses," *BioRxiv*, p. 198432, 2017.
[10] M. Kern, S. Bert *et al.*, "Human Motor Cortex Relies on Sparse and Action-specific Activation during Laughing, Smiling and Speech Production," *Communications Biology*, vol. 2, no. 1, p. 118, 2019.
[11] C. A. Stefano Filho, R. Attux, and G. Castellano, "Motor Imagery Neurofeedback: From System Conceptualization to Neural Correlates," *Current Behavioral Neuroscience Reports*, vol. 11, no. 2, pp. 78–98, 2024.
[12] G. H. Klem, "The Ten-twenty Electrode System of The International Federation. The International Federation of Clinical Neurophysiology," *Electroencephalogr. Clin. Neurophysiol. Suppl.*, vol. 52, pp. 3–6, 1999.
[13] T. M. Ingolfsson, M. Hersche *et al.*, "EEG-TCNet: An Accurate Temporal Convolutional Network for Embedded Motor-imagery Brain-machine Interfaces," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2020, pp. 2958–2965.
[14] Y. K. Musallam, N. I. AlFassam *et al.*, "Electroencephalography-based Motor Imagery Classification Using Temporal Convolutional Network Fusion," *Biomedical Signal Processing and Control*, vol. 69, p. 102826, 2021.
[15] H. Liu, Z. Dong *et al.*, "Staeformer: Spatio-temporal Adaptive Embedding Makes Vanilla Transformer SOTA for Traffic Forecasting," *arXiv preprint arXiv:2308.10425*, 2023.
[16] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. [Online]. Available: https://arxiv.org/abs/1908.10084
[17] J. Lee, W. Yoon *et al.*, "BioBERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
[18] P. Deka, A. Jurek-Loughrey *et al.*, "Evidence Extraction to Validate Medical Claims in Fake News Detection," in *International Conference on Health Information Science*. Springer, 2022, pp. 3–15.
[19] C. Brunner, R. Leeb *et al.*, "BCI Competition 2008–Graz data set A," *Institute for Knowledge Discovery (Laboratory of Brain-computer Interfaces), Graz University of Technology*, vol. 16, no. 1-6, p. 34, 2008.
[20] M. L. McHugh, "Interrater Reliability: The Kappa Statistic," *Biochemia medica*, vol. 22, no. 3, pp. 276–282, 2012.