



OPEN A low-latency neural inference framework for real-time handwriting recognition from EEG signals on an edge device

Ovishake Sen¹✉, Raghav Soni², Darpan Virmani¹, Akshar Parekh³, Patrick Lehman¹, Sarthak Jena², Adithi Katikhaneni¹, Adam Khalifa¹ & Baibhab Chatterjee¹

Brain–computer interfaces (BCIs) hold significant promise for restoring communication in individuals with severe motor or speech impairments. Imagined handwriting, as a form of motor imagery, offers an intuitive paradigm for character-level neural decoding. While invasive techniques such as electrocorticography (ECoG) offer high decoding accuracy, their surgical requirements pose clinical risks and hinder scalability. Non-invasive alternatives like electroencephalography (EEG) are safer and more accessible but suffer from low signal-to-noise ratio (SNR) and spatial resolution, limiting their effectiveness in high-resolution decoding. Here, we investigate how advanced machine learning, combined with informative feature extraction, can overcome these limitations, enabling EEG-based decoding performance that approaches invasive methods, while supporting real-time inference on edge devices. We present the first real-time, low-latency, high-accuracy system for decoding imagined handwriting from non-invasive EEG signals on a portable edge device. EEG data were collected from 15 participants using a 32-channel headcap and preprocessed with bandpass filtering and artifact subspace reconstruction (ASR). We extracted 85 time domain, frequency domain, and graphical features, then applied Pearson correlation coefficient-based feature selection to reduce latency while preserving accuracy. We developed a hybrid architecture, EEdGeNet, which integrates a Temporal Convolutional Network (TCN) with a multilayer perceptron (MLP), trained on the extracted features and deployed on the NVIDIA Jetson TX2 for real-time inference. The system achieved $89.83\% \pm 0.19\%$ accuracy with 914.18 ms per-character inference latency. By selecting only ten key features, the model incurred a minimal accuracy loss of $< 1\%$, while achieving a $4.51 \times$ reduction in inference latency (202.62 ms) compared to the full 85-feature set. These findings show that non-invasive EEG, combined with efficient feature and model design, can enable accurate, real-time neural decoding on low-power edge devices, paving the way for practical, portable BCIs.

Keywords EEG decoding, Imagined handwriting, Real-time inference, Edge device, BCI, Lightweight machine learning models

Imagined handwriting offers an intuitive and expressive mechanism for translating neural intent into character-level text. This paradigm leverages well-established motor representations in the brain and enables high-throughput decoding, making it particularly promising for individuals with severe motor or speech impairments¹.

While invasive modalities such as ECoG and intracortical microelectrode arrays provide high signal fidelity and spatial resolution, their reliance on surgical implantation imposes clinical risks and limits long-term adoption. Non-invasive techniques such as EEG are safer, more accessible, and better suited for scalable applications. EEG-based motor imagery systems have shown promise in assistive domains, including speller interfaces, virtual cursor control, and thought-to-text conversion². However, EEG signals are inherently noisy, non-stationary, and spatially coarse—posing challenges for reliable decoding in fine-grained tasks like handwriting³.

Recent advances in machine learning and signal processing have begun to address these limitations. Techniques such as artifact subspace reconstruction, temporal-spectral filtering, and targeted feature extraction

¹Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA. ²Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611, USA. ³Behavioral and Cognitive Neuroscience, University of Florida, Gainesville, FL 32611, USA. ✉email: ovishake.sen@ufl.edu

can enhance signal quality and reduce dimensionality while preserving meaningful neural dynamics. When combined with compact neural architectures, these strategies enable decoding performance that increasingly approaches invasive systems, even from non-invasive data^{3,4}.

For practical deployment, BCIs must operate in real time on portable, low-power platforms. On-device inference reduces latency, improves responsiveness, and eliminates dependence on external processing, which is especially critical for users with mobility constraints^{3,6}. However, edge deployment of EEG-based decoders remains challenging due to the computational cost of deep models and the high dimensionality of raw signals^{7,8}. Efficient preprocessing and feature extraction are thus essential—not only to suppress artifacts and noise, but also to retain physiologically meaningful components that reduce computational overhead and enable robust, low-latency inference on resource-constrained hardware.

Recent works on character and speech recognition using invasive signals relying on intracortical or implanted electrode systems have utilized Utah arrays, stereoelectroencephalography (SEEG), and ECoG to achieve high-resolution recording of neural signals^{9–10}. These systems have been used for the classification of spoken or attempted speech and handwriting. Willett et al. used Utah arrays implanted in the motor cortex to decode attempted handwriting, achieving 90 characters per minute with 94.1% accuracy in real time¹¹. Other works have utilized deep learning models, including recurrent neural networks and encoder-decoder architectures, to achieve high classification accuracies for signal decoding. Although these methods demonstrate high real-time performance, they require surgical implantation and clinical oversight^{8–11}. In contrast, our approach leverages non-invasive EEG to enable real-time, low-latency decoding of imagined handwriting on edge devices, offering a safer and more accessible alternative.

For non-invasive EEG-based decoding, studies have employed diverse preprocessing and artifact removal pipelines—including usage of bandpass and notch filtering, line noise removal, baseline correction, Independent Component Analysis (ICA), Artifact Subspace Reconstruction (ASR), Multiple Artifact Rejection Algorithm (MARA), and wavelet or entropy-based transformations—to enhance signal quality for downstream classification tasks. Several works also incorporate referencing schemes, normalization (such as Z-score), and dimensionality reduction techniques and follow preprocessing with feature extraction^{3,7,12,13}. These preprocessing pipelines have been applied for diverse EEG classification tasks such as Alzheimer's stage classification¹⁴, speech and vocabulary generation¹⁵ and ischemic stroke classification¹⁶. This shows that usage of such preprocessing pipelines is critical for preparing clean and informative data for EEG classification.

In imagined character and speech recognition, most EEG-based studies have predominantly relied on offline deep learning models, achieving high classification accuracy without addressing real-time inference. Ullah et al. (2021) focused on imagined character recognition using a deep Convolutional Neural Network (CNN) trained on Morlet wavelet-transformed EEG signals¹⁷. Their model was validated on the EEGMMIDB and BCI Competition IV-2a datasets, as well as a custom visual-imagery dataset comprising 10 subjects and 26 English characters. They reported classification accuracies of 99.45% on public datasets and 95.2% on their proprietary dataset. However, the model was not evaluated for real-time inference capability. Several studies have explored EEG decoding using various other deep learning models and traditional machine learning techniques. Models and architectures such as SVMs, LSTMs, CNN-based EEGNet, pretrained models on ResNet50, and BART-based LLMs to capture spatiotemporal patterns and generalize across complex EEG data have been used for speech and character classification. Various feature extraction pipelines have also been incorporated with these models that extract specific features ranging from spectral power and coherence to wavelet-scattering and 3D topographic representations, achieving classification accuracies up to 95.5% in specific tasks^{15,18–20}. Despite having high offline performance, most of these frameworks have been designed and evaluated in non-real-time settings, limiting their current applicability for real-time deployment.

For motor imagery EEG classification, studies have primarily utilized offline deep learning frameworks such as CNNs, LSTMs, and transformers with physics-informed feature extraction, spatial-temporal attention-based learning or hybrid architectures, achieving high accuracy but limited real-time applicability. These models have been employed to classify binary, three-class, and four-class motor imagery tasks and have been validated using BCI Competition datasets, with several achieving offline accuracies up to 92%^{21–23}. While most frameworks have been developed and evaluated in offline mode, NeuroGrasp has demonstrated partial real-time feasibility using CNN-BLSTM and Siamese networks²⁴. Overall, these architectures are designed to maximize accuracy by learning intricate signal representations, though real-time applicability remains limited in most cases.

Traditional BCI paradigms like Steady-State Visual Evoked Potential (SSVEP) based BCIs achieve high information transfer rates²⁵, but they depend on sustained visual attention to flickering stimuli, which can induce fatigue and are unsuitable for patients with visual or oculomotor impairments. Event-Related Potential (ERP)/P300 spellers offer reliable performance²⁶, yet they also rely on repeated stimulus presentation, leading to slower communication speeds and reduced practicality for continuous use. Speech-attempt paradigms, including overt or attempted articulation, provide natural mapping to language but typically require invasive recordings (ECoG or intracortical arrays) to reach state-of-the-art performance^{10,11}. In contrast, imagined handwriting directly encodes rich, character-level linguistic information without reliance on external stimuli or residual motor ability.

Figure 1a illustrates the proposed framework for real-time neural inference from non-invasive EEG signals to decode imagined handwriting. A motor and communication-impaired individual is equipped with a multi-channel EEG headcap, and the system is designed to translate character-level motor intentions into text in real-time through low-latency inference pipelines deployed on edge hardware. This approach demonstrates the feasibility of portable neural decoding, offering high-speed processing and accurate predictions directly on resource-constrained devices. As shown in Fig. 1b, unlike conventional methods that depend on cloud-based infrastructure, high-performance computing platforms, or invasive neural recordings, our system enables on-device computation with non-invasive EEG data, significantly reducing latency and enhancing deployment practicality. Comparative validation against existing EEG-based models, as depicted in Fig. 1c, demonstrates

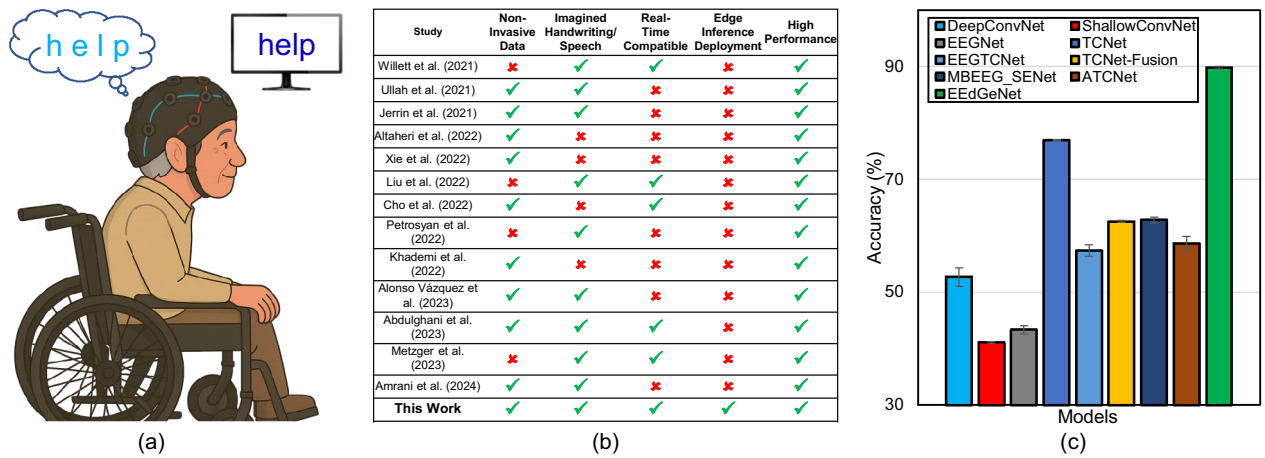


Fig. 1. Proposed framework for real-time imagined handwriting decoding from non-invasive EEG signals. **(a)** Illustration of a motor-impaired person wearing a multi-channel EEG headcap for real-time neural decoding. **(b)** Comparative analysis of prior studies based on: non-invasive signal acquisition, imagined handwriting/speech decoding, real-time inference compatibility, edge-level deployment feasibility, and classification performance. The proposed approach is the only method that meets all criteria concurrently. **(c)** Accuracy comparison with state-of-the-art EEG models on 85 features extracted from imagined handwriting data, showing the superior performance of our proposed model (EEdGeNet) for edge-based inference.

that the proposed EEdGeNet architecture achieves superior performance in decoding imagined handwritten characters from feature-extracted EEG signals.

In this paper, we present a novel methodology for real-time inference from non-invasive EEG signals to decode imagined handwriting. The key contributions of this study are as follows:

- We demonstrate, for the first time, real-time inference of imagined handwriting from non-invasive EEG signals, achieving a character classification accuracy of $89.83\% \pm 0.19\%$ using 85 extracted features from raw EEG data.
- We implement, for the first time, a real-time EEG decoding pipeline fully deployed on a portable hardware platform (NVIDIA Jetson TX2), with an average inference latency of 914.18 ms per character.
- We select ten key features using Pearson correlation, achieving $88.84\% \pm 0.09\%$ accuracy while reducing inference latency by $4.51\times$ compared to the full feature set (202.62 ms), enabling real-time BCI with minimal computational overhead.
- We develop EEdGeNet, a hybrid TCN-MLP architecture that can be efficiently deployed on the NVIDIA Jetson TX2 and achieves superior performance in imagined handwriting recognition compared to state-of-the-art EEG classifiers.
- Our optimized feature extraction pipeline, combined with a lightweight and tuned neural model that avoids the high computational overhead typical of existing motor imagery classifiers, enables real-time inference and deployment on resource-constrained edge devices.

In essence, this work presents the first real-time, high-accuracy decoding of imagined handwriting from non-invasive EEG signals on a portable edge device, demonstrating robust generalizability across multiple participants and low-latency performance suitable for assistive communication. The proposed method achieves $89.83\% \pm 0.19\%$ accuracy across 15 participants using 85 extracted features, with an inference latency of 914.18 ms per-character on the NVIDIA Jetson TX2. Selecting only ten key features results in $< 1\%$ accuracy loss while reducing latency by $4.51\times$, demonstrating the feasibility of low-power, real-time BCI systems for assistive communication.

Materials and methods

The proposed real-time EEG-based handwriting recognition pipeline

The methodology for decoding imagined handwritten characters from EEG signals began with data acquisition from 15 participants. Raw EEG recordings were segmented into per-character windows and preprocessed to suppress noise and artifacts. From each segment, 85 features were then extracted, covering the time and frequency domains, along with graphical descriptors. To minimize computational overhead during inference, a feature selection step is performed, retaining only the most discriminative features and spatially relevant EEG channels associated with handwriting intent. The feature set is then used to train a suite of ML classifiers, with model performance assessed on a held-out test set to evaluate generalizability. The complete framework for low-latency, real-time EEG inference is illustrated in Fig. 2. Figure 2a presents the end-to-end pipeline, including bandpass filtering, ASR-based artifact removal, feature extraction, and model execution for both training (on a local machine) and inference (on local machine as well as on edge hardware). Figure 2b shows the EEG data acquisition setup with the Emotiv headset. Figure 2c illustrates the electrode placement following the

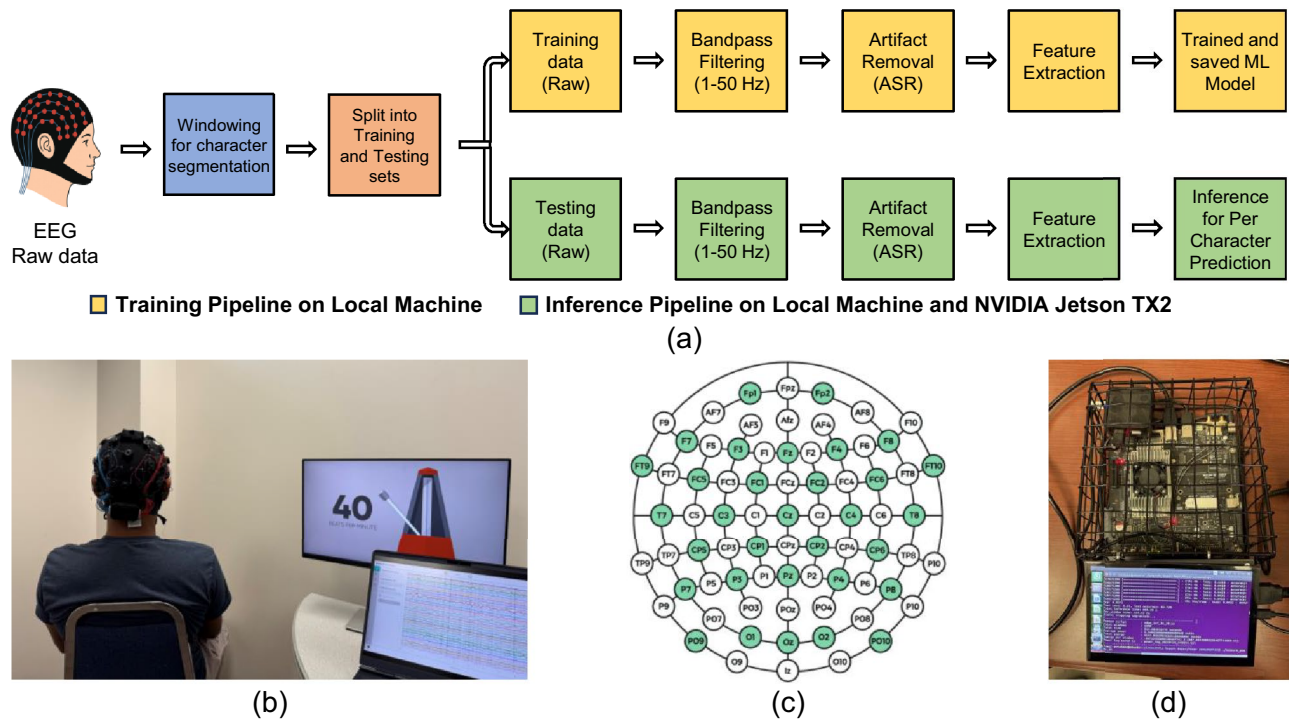


Fig. 2. Proposed methodology for low-latency neural inference on an edge device for real-time handwriting recognition from EEG signals. **(a)** The entire pipeline for training (yellow) and inference (green), including preprocessing, feature extraction, and per-character prediction. **(b)** Data collection setup showing a participant wearing the Emotiv Flex 2.0 EEG headcap during data acquisition. **(c)** Electrode layout based on the international 10–20 system, with utilized channels highlighted in green. **(d)** Real-time inference on the NVIDIA Jetson TX2 with a batch size of 1.

international 10–20 system, with the channels used in our study highlighted in green. Figure 2d demonstrates real-time deployment on the NVIDIA Jetson TX2 for per-character prediction. The following sections present a detailed overview of the key components comprising the proposed pipeline.

Participants and EEG recording

EEG signals were collected from 15 participants using the Emotiv Epoch FLEX 2.0 saline-based head cap, comprising 32 scalp electrodes with a sampling frequency of 256 Hz. All experimental protocols were approved by the University of Florida Institutional Review Board (IRB202301691) and were conducted in accordance with the relevant guidelines and regulations. Informed consent was obtained from all participants. Each participant was instructed to imagine writing all 26 lowercase English alphabet characters (a–z), in addition to one “do nothing phase” intended to capture resting-state neural activity without any intentional motor or cognitive imagery. All recording sessions were conducted in a quiet room to minimize environmental artifacts and auditory or visual distractions. Prior to data acquisition, participants were instructed to relax and avoid interaction with any electronic devices to reduce extraneous neural activity. The recording session commenced with a structured baseline protocol consisting of a 5-second preparation window, followed by 15 seconds of eyes-open relaxation, a 5-second transition period, and then 15 seconds of eyes-closed relaxation. A subsequent 20-second resting period was given before the onset of the data recording phase.

Imagined handwriting EEG data collection was synchronized to a 40 BPM metronome. Starting from the fourth beat (1 minute and 6 seconds into the session), participants were instructed to begin imagining the act of writing the target character in synchrony with each subsequent metronome beat. Data was recorded continuously until the 144th beat (4 minutes and 36 seconds). Participants were then given a short break before proceeding with the next character trial. The complete set of 26 alphabet characters and the do nothing phase was recorded over multiple sessions (typically 4–5), adjusted based on each participant’s cognitive load and relaxation level. Each session included data collection for 6 to 8 characters. Following each session, the raw EEG signals were exported as CSV files using EmotivPRO software for subsequent preprocessing and analysis.

Data preprocessing

Following data acquisition, the raw EEG signals collected from each participant were subjected to a standardized preprocessing pipeline designed to attenuate noise and artifacts while preserving the natural neural activity. To ensure compatibility with real-time inference, preprocessing techniques were applied separately to each character window after the train/test split. This design reflects the real-time constraint of obtaining per-character predictions, necessitating that preprocessing be performed independently on each window.

Accordingly, while independent component analysis (ICA), principal component analysis (PCA), discrete wavelet transform (DWT), empirical wavelet transform (EWT), and variational mode decomposition (VMD) were explored, these methods proved insufficient for fully denoising the signal likely due to the shorter window lengths inherent to character-level EEG segments^{27,28}. We therefore employed ASR, a method demonstrated to be effective for cleaning short-duration EEG recordings^{29,30}. ASR balances robustness with computational efficiency, making it particularly well-suited for real-time deployment on edge devices. The preprocessing steps are outlined as follows.

1. **Bandpass filtering** To remove low-frequency drifts and high-frequency noise, the raw EEG signals were passed through a bandpass filter with a passband of 1–50 Hz^{31–33}. This range captures the relevant EEG frequency bands while excluding low-frequency noise arising from sources such as head movements, electrode wire displacement, scalp perspiration, or slow drifts in the EEG signal over extended periods⁷. A fourth-order Butterworth filter was designed using the `scipy.signal.butter` function and applied uniformly across all channels for every EEG sample.
2. **ASR** After bandpass filtering, we employed a simplified variant of Artifact Subspace Reconstruction (ASR) to suppress transient, high-amplitude disturbances. This method was applied independently to each EEG channel. For each channel, the mean and standard deviation were computed, and samples exceeding a z-score threshold of 3.0 were identified as artifacts and replaced with NaN values. Missing values were subsequently reconstructed using linear interpolation, followed by backward and forward filling to ensure signal continuity at temporal boundaries. This procedure preserved the intrinsic temporal structure of the EEG signal while mitigating large, non-stationary artifacts³⁴.

Figure 3a illustrates a participant's 1.5-second EEG segment recorded during the writing of the character 'v', while Fig. 3b shows the corresponding signal following preprocessing, demonstrating the effective removal of noise and artifacts.

Feature extraction

Following preprocessing, a comprehensive set of features was extracted from each EEG window to capture both linear and non-linear signal characteristics. Specifically, we computed 12 time-domain features, 8 frequency-




Fig. 3. (a) Raw EEG signals. (b) Preprocessed EEG signals. The 32-channel EEG data from a participant's sample for the character 'v' is shown before and after preprocessing. Bandpass filtering (1–50 Hz) and Artifact Subspace Reconstruction (ASR) were applied to effectively remove noise and artifacts, improving overall signal quality.

domain features, and 65 graphical features derived from Poincaré plots of discrete wavelet transform sub-bands, as described in³⁵. The time-domain features include: mean, variance, skewness, kurtosis, Root Mean Square (RMS), Slope Sign Changes (SSC), Hjorth mobility, Hjorth complexity, Hurst exponent, Shannon entropy, sample entropy, and permutation entropy. The frequency-domain features consist of: spectral band powers in the delta, theta, alpha, beta, and gamma bands, peak frequency, mean dominant frequency, and spectral entropy. The graphical features comprised thirteen descriptors computed from each Poincaré plot: ellipse area, eccentricity, first and second eigenvalues, mean deviation from the identity line, mean distance from the origin, sequential standard deviation of vector lengths (SSVL), covariance, mean x- and y-coordinates, standard deviations along both axes, and the central tendency measure (CTM). These features were extracted independently from the approximation (A) and detail (D1–D4) sub-bands of the DWT, resulting in a total of 65 graphical features per EEG window.

Time domain features

Mean, RMS, SSC, and Variance were calculated using the NumPy library. Skewness, kurtosis, and Shannon entropy were computed using the SciPy library. Hjorth mobility and Hjorth complexity were calculated as described by³⁶. Permutation entropy was calculated as described by³⁷.

Hurst Exponent³⁸ was computed using the Eq. (1):

$$Hurst_Exponent = \frac{\log(R/S) - \log c}{\log N} \quad (1)$$

where R is the range of the cumulative sum of deviations of the signal values from the mean, S is the standard deviation, and N is the number of samples of the signal. In EEG analysis, the constant c only shifts the intercept and does not affect the Hurst exponent, so it is set to 1 to simplify computation and avoid unstable estimation in noisy signals.

Sample Entropy was computed using the Eq. 2):

$$SampEn(m, r, N) = -\ln \left(\frac{\Phi^{m+1}(r)}{\Phi^m(r)} \right) \quad (2)$$

where m is the length of the subsequence, r is the tolerance, N is the number of data points in the signal, and $\varphi_m(r)$ is the probability that two subsequences of length m are similar, where similarity is determined by the tolerance r .

The probability of a match between two subsequences within the signal is defined as:

$$\Phi^m(r) = \frac{1}{N - m + 1} \sum_{i=1}^{N-m+1} \frac{B_i^m(r)}{N - m}$$

where $B_i^m(r)$ is the count of pairs of subsequences for which the Chebyshev distance between that pair is less than the tolerance parameter r , defined as $r = 0.2 \times \sigma$, where σ is the standard deviation of the signal. Chebyshev distance between two subsequences within the signal is defined as:

$$d[u(i), u(j)] = \max_{k=0, \dots, m-1} |x_{i+k} - x_{j+k}|$$

where $u(i) = \{x_i, x_{i+1}, \dots, x_{i+m-1}\}$ is a subsequence of length m from a signal of length N

We use $N - m + 1$ instead of $N - m$ because it accurately reflects the total number of embedding vectors of length m in a time series of length N . This avoids excluding the final valid segment and prevents systematic underestimation in entropy computation.

Graphical features from Poincaré-DWT analysis

Graphical features were extracted from Poincaré plots constructed using the DWT coefficients of each EEG channel within a 1.5-second window. The EEG signal was decomposed into five sub-bands, comprising the approximation (A) and detail levels (D1–D4), using the `wavedec` function from the `PyWavelets` library. For each sub-band, time-lagged DWT coefficients were embedded in phase space to generate Poincaré plots, capturing the dynamical evolution of the signal over time. Thirteen graphical descriptors were computed from each plot using NumPy, characterizing the shape, spread, and structural complexity of the trajectory: ellipse area, eccentricity, first and second eigenvalues (representing variances along the principal axes), mean deviation from the identity line, mean distance from the origin, sequential standard deviation of vector lengths (SSVL), covariance, mean of x and y coordinates, standard deviations along both axes, and central tendency measure (CTM). Applied across all five DWT sub-bands, this resulted in 65 graphical features per EEG window. These non-linear descriptors capture morphological and recurrence-based patterns in neural activity and have been shown to enhance classification performance in a range of cognitive and neurological decoding tasks^{35,39,40}.

Feature selection

We used the Pearson correlation coefficient to select the most important features, calculated using the Eq. 3).

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \cdot \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}} \quad (3)$$

Here, r is the Pearson correlation coefficient. x_i, y_i are values of all features (85) and n is the number of channels (32). r measures the linear relationship between all the feature pairs. The coefficient r captures the strength of the linear relationship between a pair of features, where $r = 1$: perfect positive correlation, $r = -1$: perfect negative correlation, $r = 0$: no linear correlation. To identify the most relevant features, we first included negatively correlated feature pairs with a correlation less than -0.4 . Then, if any selected feature pairs exhibited a correlation greater than 0.7 , we discarded one feature from each such pair to avoid redundancy. Finally, we selected the 10 most important features that provided approximately the same accuracy as using all 85 features, while improving inference time/character by $4.51 \times$ and inference energy/character by $5.41 \times$.

Evaluation methodology

Proposed EEdGeNet model architecture

The proposed EEdGeNet model architecture comprises two principal blocks: a Temporal Convolutional Block and a Dense Transformation Block. The first block employs a Temporal Convolutional Network (TCN) to effectively capture both spatial and temporal dependencies inherent in the features extracted from raw EEG signals. The second block flattens the temporal representations and transforms them into a compact feature vector, leveraging a deep Multilayer Perceptron (MLP) to learn complex non-linear mappings. The outputs from all sliding windows are concatenated and fed into a softmax classifier to perform multi-class classification. Further architectural details of the proposed custom model are provided in the subsequent sections.

Figure 4 shows the architecture of the proposed model, comprising a Temporal Convolutional Block (TCB) and a Dense Transformation Block (DTB). The TCB captures temporal dependencies using dilated convolutions, while the DTB performs non-linear transformations via stacked dense layers.

Pre-processing and input representation

In this work, we utilize windows of EEG signals processed through a feature extraction pipeline that computes a predefined set of time-domain and frequency-domain features. The input to the model is a windowed segment $X_i \in \mathbb{R}^{C \times F}$, where $C = 32$ denotes the number of EEG channels and F represents the number of computed features per channel. The model is designed to map each windowed segment X_i to its corresponding class label y_i , given a labeled dataset $\mathcal{S} = \{(X_i, y_i)\}_{i=1}^m$, where $y_i \in \{1, \dots, n\}$ and $n = 27$ denotes the total number of classes in our task.

Temporal convolutional block

To extract spatial and temporal relationships among EEG channels, a TCN is integrated into the model architecture⁴¹. This block comprises an initial two-layer causal convolutional sub-block with Batch Normalization (BN), Exponential Linear Unit (ELU) activation, and dropout, followed by two dilated residual sub-blocks designed to capture local dependencies and feature interactions within each channel. The input to this block is a windowed segment $X_i \in \mathbb{R}^{C \times F}$, where C denotes the number of EEG channels and F represents the number of computed features per channel.

The input to the TCN block is denoted as $Z^{(0)}$, and its transformation through the block is described in Eq. 4.

$$\begin{aligned} Z^{(1)} &= \text{Dropout} \left(\phi \left(\text{BN} \left(\text{Conv1D}_{k=10, d=1} \left(\text{Dropout} \left(\phi \left(\text{BN} \left(\text{Conv1D}_{k=10, d=1} \left(Z^{(0)} \right) \right) \right) \right) \right) \right) \right), \\ Z^{(2)} &= \text{Dropout} \left(\phi \left(\text{BN} \left(\text{Conv1D}_{k=3, d=d_j} \left(\phi \left(\text{BN} \left(\text{Conv1D}_{k=3, d=d_i} \left(Z^{(j)} \right) \right) \right) \right) \right), \quad j = 1, 2; d_i = 2, d_j = 4, \end{aligned} \quad (4)$$

where $Z^{(1)}$ is the output of the initial causal convolutional sub-block, and $Z^{(2)}$ is the output from the two dilated residual sub-blocks. Here, $\phi(\cdot)$ denotes the ELU activation function, k is the kernel size, d is the dilation rate, and $\text{BN}(\cdot)$ represents Batch Normalization.

The final output $Z^{(2)}$ is subsequently flattened and passed to the fully connected layers in the Dense Transformation Block to learn complex, non-linear representations from the extracted features.




Fig. 4. Proposed EEdGeNet model architecture integrating a Temporal Convolutional Block (TCB) and a Dense Transformation Block (DTB). The TCB (highlighted in yellow) captures temporal dependencies via dilated convolutions, while the DTB (highlighted in black) performs hierarchical non-linear feature transformations. Abbreviations: BN—Batch Normalization, DR—Dropout layer, EL—ELU activation, and D —dilation rate.

Dense transformation block

The flattened representation $Z^{(2)}$ is passed through a hierarchical network of fully connected layers with progressively decreasing dimensionality to learn high-level discriminative representations for classification. Each dense layer is followed by BN, ELU activation, and dropout, which collectively enhance training stability and reduce the risk of overfitting.

The hidden layers are configured with the following number of neurons: $512 \rightarrow 256 \rightarrow 128 \rightarrow 64$. The transformation at each layer is defined in Eq. 5:

$$h^{(l)} = \text{Dropout} \left(\sigma \left(\text{BN} \left(W^{(l)} h^{(l-1)} + b^{(l)} \right) \right) \right), \quad l = 1, 2, 3, 4 \quad (5)$$

where $h^{(0)} = \text{Flatten}(Z^{(2)}) \in \mathbb{R}^{CF}$ denotes the flattened input vector, $\sigma(\cdot)$ is the ELU activation function, and $\text{BN}(\cdot)$ represents Batch Normalization.

The output from the final dense layer is passed through a softmax classification layer to produce the class probability distribution for subsequent label prediction.

Table 1 lists the hyperparameter configurations of the Temporal Convolutional Block and Dense Transformation Block used in the proposed EEdGeNet model. Hyperparameters were selected using scikit-learn's `GridSearchCV` with 3-fold cross-validation.

Experimental setup

Model training

The EEG dataset was used to train the EEdGeNet model locally on a personal machine. Adam optimizer was used for training the custom model along with a cosine decay learning rate schedule for learning rate reduction over a set time. The initial learning rate was set to 3×10^{-4} and was set to decay until a minimum learning rate of 1×10^{-5} . The decay steps were defined as $\text{decay_steps} = \lfloor \frac{N}{B} \rfloor \times E$, where N is the number of training samples, $B = 32$ is the batch size, and $E = 500$ is the maximum number of epochs.

To calculate the dissimilarity of the predicted probability distribution from the model and the actual class labels, sparse categorical cross-entropy loss was used. The model was trained with early stopping, monitoring the validation loss with patience of 15 epochs, and restored the weights from the epoch with the lowest validation loss. In order to address any possible class imbalance in the training set, class weights were calculated and incorporated during training to assign higher importance to minority classes. A model checkpoint callback saved the weights corresponding to the highest validation accuracy observed during training.

Local machine-based inference

Inference was performed on the test dataset using a batch size of 1, enabling character-by-character prediction. The entire pipeline, including preprocessing, feature extraction, and model evaluation, was executed in real time. The experiments were conducted locally on a personal machine equipped with an Intel Core Ultra 9-275HX processor (24 cores, 2.7 GHz), 32 GB of RAM, and an NVIDIA GeForce RTX 5060 Laptop GPU.

Inference on an edge device

To demonstrate the portability of the proposed real-time EEG-based handwritten character recognition system, inference was also conducted on an edge device. Specifically, the NVIDIA Jetson TX2 platform was used to evaluate the trained models, performing character-by-character prediction with a batch size of 1. The Jetson TX2 features a 256-core NVIDIA PascalTM GPU, a dual-core NVIDIA Denver 2 64-bit CPU, and a quad-core ARM[®]Cortex[®]-A57 MPCore. It is equipped with 8 GB of 128-bit LPDDR4 memory (1866 MHz, 59.7 GB/s bandwidth) and 32 GB of eMMC 5.1 storage. The device operates within a configurable power envelope of 7.5W to 15W.

To measure energy consumption during inference, we employed NVIDIA's `tegrastats` utility, which provides real-time board-level power readings via the `VDD_IN` rail. Power values were logged at 100 ms intervals while the inference script was executed, and the average power draw was computed over the entire inference duration. Multiplying this average by the elapsed runtime yielded the total energy, which was then normalized by the number of processed EEG windows to report the per-character inference energy cost.

| Block | Hyperparameter | Value |
|------------------------------------|------------------------|---------------------|
| Temporal convolutional block (TCB) | Dilation rates | [1, 2, 4] |
| | Kernel size | 3 |
| | Filters per Conv layer | 32 |
| | Dropout rate | 0.3 |
| Dense transformation block (DTB) | Hidden layers | [512, 256, 128, 64] |
| | Dropout rate | 0.3 |
| | L2 Regularization | 0.0008 |

Table 1. Hyperparameter configuration of the Temporal Convolutional Block (TCB) and Dense Transformation Block (DTB) used in the proposed EEdGeNet model.

Experimental results

Towards an optimal architecture for imagined handwriting from EEG

Table 2 presents a comparative evaluation of various model configurations for imagined handwritten character classification from EEG signals, considering both raw and feature-extracted inputs. A four-layer multilayer perceptron (MLP) with hidden dimensions of 512, 256, 128, and 64, when applied directly to the raw EEG signals, yields a test accuracy of $10.26\% \pm 0.13\%$, with limited discriminative capability. The use of a TCN on the raw data offers modest improvement, reaching $29.10\% \pm 1.03\%$, consistent with prior reports on TCNs outperforming recurrent models like LSTM and GRU for sequence modeling tasks⁴¹. To enhance performance, we extended the architecture by appending our proposed MLP module to the TCN backbone. This hybrid TCN-MLP model (EEdGeNet) achieves $43.63\% \pm 2.77\%$ accuracy on the raw dataset. While models using raw signals maintain low per-character inference latency on the local machine (approximately 7–11 ms), their classification performance remains suboptimal.

To overcome this limitation, we extracted a compact set of 85 features, covering the time and frequency domains, along with graphical descriptors, as described in the *Feature Extraction* section. When evaluated independently, MLP and TCN models trained on these extracted features yield accuracies of $73.82\% \pm 0.58\%$ and $72.62\% \pm 0.10\%$, respectively, demonstrating the utility of domain-informed representations. The best overall performance is obtained with the feature-based EEdGeNet model, which achieves $89.83\% \pm 0.19\%$ test accuracy, accompanied by high precision, recall, and F1-score metrics. The corresponding inference latency per character on a local machine is 41.25 ms, while inference on the NVIDIA Jetson TX2 requires 914.18 ms per-character, as the computation involves additional layers and embedded resource constraints. In subsequent analyses, we address this bottleneck by reducing the feature set to optimize on-device efficiency without significantly sacrificing classification performance.

Comparative analysis of time domain, frequency domain, and graphical features

Table 3 presents a comparative evaluation of time domain, frequency domain, and graphical feature sets for imagined handwriting classification from EEG signals using the proposed EEdGeNet model. Time-domain features achieved the highest standalone accuracy ($85.40\% \pm 0.88\%$) but incurred relatively high inference latency (36.93 ms per-character on a local machine and 792.73 ms on the NVIDIA Jetson TX2). By contrast, frequency-domain features offered the lowest latency (7.24 ms on local machine and 199.52 ms on Jetson TX2) but correspondingly yielded the lowest accuracy ($72.18\% \pm 0.16\%$). Graphical features provided an intermediate trade-off, attaining $83.26\% \pm 1.29\%$ accuracy with moderate latency (16.94 ms locally and 229.84 ms on Jetson TX2). When all 85 features were combined, overall accuracy was maximized ($89.83\% \pm 0.19\%$), albeit at the cost of increased latency (41.25 ms locally and 914.18 ms on Jetson TX2). Although the full feature set achieves the highest accuracy, the accompanying latency underscores the need for feature selection, motivating the search for a minimal subset that preserves accuracy while substantially reducing inference time. This direction is pursued in the following subsection.

Impact of feature selection on model performance

To identify the most informative features, we conducted a systematic feature selection analysis based on Pearson correlation coefficients computed across all 32 EEG channels for each character and participant. In a representative case involving one participant and one character, some of the negatively correlated feature pairs were: (Hjorth mobility vs. Hjorth complexity), (Hjorth mobility vs. delta power), (RMS vs. Hjorth mobility), (kurtosis vs. Shannon entropy), (Hjorth mobility vs. graphical_A_ellipse_area), (Hjorth mobility vs. graphical_A_std_x_axis), (Hjorth mobility vs. graphical_A_std_y_axis), (Hjorth mobility vs. graphical_A_mean_distance_origin), (Hjorth mobility vs. mean dominant frequency), (RMS vs. mean dominant frequency), (graphical_D1_ellipse_area vs. graphical_D1_covariance_x-y_coordinates), (graphical_D2_ellipse_area vs. graphical_D2_covariance_x-y_coordinates), (delta power vs. mean dominant frequency), (delta power vs. spectral entropy), and (Hjorth complexity vs. gamma power). Conversely, in the same participant-character instance, some of the most positively correlated feature pairs were: (RMS vs. variance), (Hjorth mobility vs. spectral entropy), (Hjorth mobility vs. gamma power), (graphical_A_ellipse_area vs. graphical_A_std_x-axis), (graphical_A_ellipse_area vs. graphical_A_std_y_axis), and (graphical_A_ellipse_area vs. graphical_A_mean_distance_origin).

| Model | Test Accuracy (%) | 95% CI (%) | Precision (%) | Recall (%) | F1-Score (%) | Inference Time/Character (Local Machine) | Inference Time/Character (NVIDIA Jetson TX2) |
|-------------------------------|---------------------|-----------------------|---------------------|---------------------|---------------------|--|--|
| MLP (Raw dataset) | 10.26 ± 0.13 | [10.13, 10.40] | 10.87 ± 0.13 | 10.26 ± 0.17 | 10.06 ± 0.09 | 6.70 ms | N/A |
| TCN (Raw dataset) | 29.10 ± 1.03 | [27.83, 30.38] | 30.24 ± 0.96 | 29.10 ± 1.28 | 29.15 ± 1.25 | 6.70 ms | N/A |
| EEdGeNet (Raw dataset) | 43.63 ± 2.77 | [43.19, 50.08] | 47.46 ± 3.20 | 46.63 ± 3.44 | 46.56 ± 3.41 | 10.41 ms | N/A |
| MLP (85 Features) | 73.82 ± 0.58 | [72.06, 75.59] | 74.07 ± 0.55 | 73.82 ± 0.58 | 73.88 ± 0.58 | 40.98 ms | N/A |
| TCN (85 Features) | 76.95 ± 0.10 | [76.64, 77.26] | 77.01 ± 0.08 | 76.95 ± 0.10 | 76.95 ± 0.09 | 40.98 ms | N/A |
| EEdGeNet (85 Features) | 89.83 ± 0.19 | [89.56, 90.09] | 89.87 ± 0.18 | 89.83 ± 0.19 | 89.83 ± 0.19 | 41.25 ms | 914.18 ms |

Table 2. Performance comparison of model configurations applied to raw and feature-extracted EEG signals for imagined handwritten character recognition. Among the evaluated architectures, the combination of TCN and MLP (EEdGeNet) on 85 extracted features achieves the highest classification accuracy and lowest per-character latency.

| Feature Set | Test Accuracy (%) | 95% CI (%) | Precision (%) | Recall (%) | F1-Score (%) | Inference Time/Character (Local Machine) | Inference Time/Character (NVIDIA Jetson TX2) |
|-------------------------------|-------------------|----------------|---------------|--------------|--------------|--|--|
| Time Domain Features (12) | 85.40 ± 0.88 | [84.48, 86.32] | 86.18 ± 0.20 | 85.89 ± 0.19 | 85.88 ± 0.19 | 36.93 ms | 792.73 ms |
| Frequency Domain Features (8) | 72.18 ± 0.16 | [71.69, 72.67] | 72.39 ± 0.19 | 72.18 ± 0.16 | 72.22 ± 0.17 | 7.24 ms | 199.52 ms |
| Graphical Features (65) | 83.26 ± 1.29 | [81.46, 85.06] | 83.52 ± 1.29 | 83.26 ± 1.29 | 83.26 ± 1.31 | 16.94 ms | 229.84 ms |
| All Features (85) | 89.83 ± 0.19 | [89.56, 90.09] | 89.87 ± 0.18 | 89.83 ± 0.19 | 89.83 ± 0.19 | 41.25 ms | 914.18 ms |

Table 3. Comparison of results across selected time-domain, frequency-domain, and graphical features evaluated using the EEdGeNet model. Frequency-domain features yield lower test accuracy but substantially reduced per-character latency during inference. In contrast, time-domain features achieve higher accuracy across the full feature set but incur greater latency. Graphical features provide a balance between accuracy and efficiency, while combining all 85 features delivers the best overall performance for per-character prediction on the NVIDIA Jetson TX2.

These associations revealed redundancy among several features. For example, RMS and Variance were strongly correlated, suggesting that Variance may be omitted without substantial information loss. Similarly, spectral entropy and gamma power exhibited strong correlations with Hjorth mobility, indicating redundancy. Likewise, graphical_A_ellipse_area could substitute for graphical_A_std_x_axis, graphical_std_y_axis, and graphical_A_mean_distance_origin. Extending this analysis across all remaining feature pairs, we identified a consistent set of features that remained distinct: Hjorth mobility, RMS, Hjorth complexity, delta power, graphical_A_ellipse_area, mean dominant frequency, graphical_D3_ellipse_area, graphical_D1_ellipse_area, graphical_D2_ellipse_area, and graphical_D4_ellipse_area. Importantly, this ordering was preserved across all 27 character classes and 15 participants, highlighting the robustness and generalizability of our feature selection procedure.

Figure 5a–c illustrates the trade-offs among selected feature subsets in terms of classification accuracy on the proposed EEdGeNet model, along with per-character inference latency and energy consumption measured on the NVIDIA Jetson TX2. Feature subsets were derived using Pearson correlation-based ranking, as described in the *Feature Selection* section. Employing only the most discriminative feature, Hjorth mobility, achieves an accuracy of 65.18%, with an inference latency of 184.71 ms and energy consumption of 796.2 mJ per-character. Incremental inclusion of RMS, Hjorth complexity, delta power, graphical_A_ellipse_area, mean dominant frequency, graphical_D3_ellipse_area, graphical_D1_ellipse_area, graphical_D2_ellipse_area, and graphical_D4_ellipse_area progressively enhances performance, ultimately achieving $88.84\% \pm 0.09\%$ accuracy, 202.62 ms latency, and 1052.42 mJ energy usage. Notably, this 10-feature configuration delivers accuracy within $\sim 1\%$ of the full 85-feature set ($89.83\% \pm 0.19\%$), while reducing inference time and energy consumption by $4.51\times$ and $5.14\times$, respectively. These results highlight an effective trade-off between computational efficiency and predictive performance, making the system well-suited for real-time inference on resource-constrained edge devices.

Ablation study of preprocessing, signal decomposition, feature, and classifier variants

We performed an extensive ablation analysis to disentangle the contributions of preprocessing, feature extraction, and classifier choices within our proposed real-time EEG-based handwriting recognition pipeline. The comparative results, summarized in Table 4, demonstrate how each design choice influences classification accuracy, inference latency, and energy efficiency on the NVIDIA Jetson TX2.

Ablation study of preprocessing and signal denoising techniques

To assess the impact of different preprocessing and denoising strategies, we replaced the simplified ASR module with conventional approaches. The standard ASR pipeline led to a notable decline in performance, yielding an average accuracy of $80.80\% \pm 0.15\%$. ICA caused a severe performance collapse ($12.03\% \pm 0.15\%$), highlighting its ineffectiveness for imagined handwriting decoding due to excessive component distortion and signal loss. In contrast, Multiscale Principal Component Analysis (MSPCA) achieved competitive accuracy ($87.12\% \pm 1.46\%$) and the lowest inference latency (828.93 ms), yet still fell short of the simplified ASR in the overall accuracy–latency trade-off. These results confirm that the simplified ASR provides the most balanced preprocessing strategy for robust and efficient EEG decoding.

Ablation study of signal decomposition techniques

We further examined decomposition-based preprocessing approaches to evaluate their suitability for imagined handwriting decoding. VMD⁴² and EWT⁴³ achieved lower accuracies of $82.96\% \pm 1.28\%$ and $79.67\% \pm 0.82\%$, respectively than the proposed pipeline. Furthermore, both methods introduced substantial computational overhead, with inference times ranging from 2212–2295 ms, and energy consumption between 7102–7369 mJ per character. Despite their modest accuracy gains, the pronounced latency and energy costs render these decomposition techniques impractical for real-time edge deployment compared with the simplified ASR-based preprocessing pipeline.

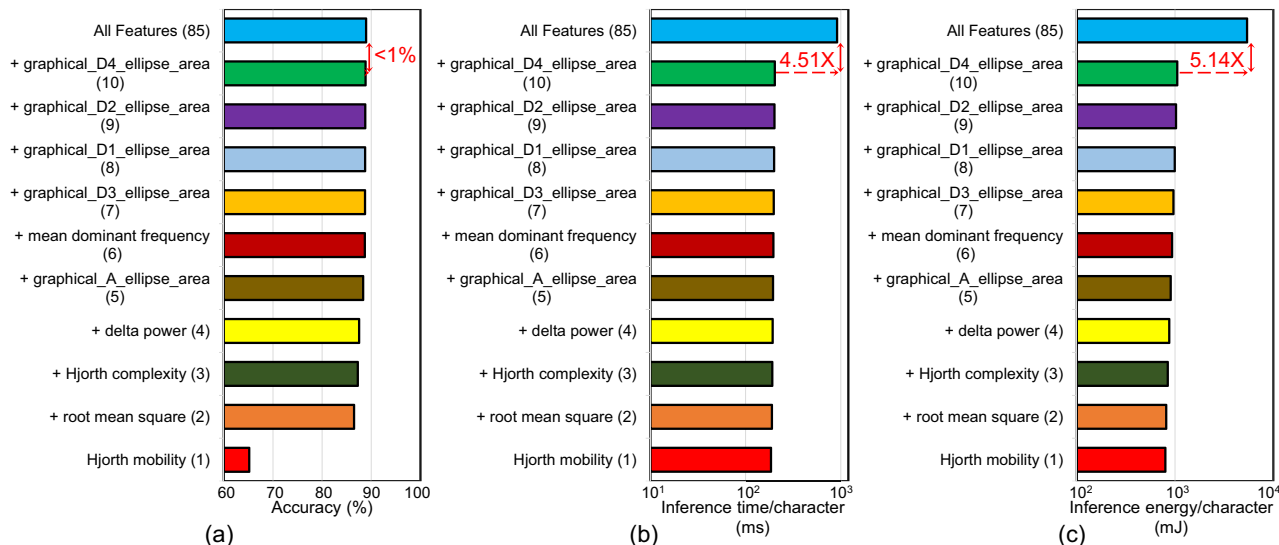


Fig. 5. Impact of feature selection using Pearson correlation coefficients on classification accuracy, inference latency, and energy consumption on the NVIDIA Jetson TX2. A compact feature set comprising Hjorth mobility, RMS, Hjorth complexity, delta power, graphical_A_ellipse_area, mean dominant frequency, graphical_D3_ellipse_area, graphical_D1_ellipse_area, graphical_D2_ellipse_area, and graphical_D4_ellipse_area results in (a) a ~1% reduction in accuracy relative to the full 85 feature set evaluated on the proposed EEdGeNet model, (b) a 4.51× improvement in per-character inference time, and (c) a 5.14× reduction in per-character inference energy. This trade-off highlights an efficient balance between computational cost and predictive performance, well-suited for edge deployment on the Jetson TX2.

| Study | Ablation Study Pipeline | Test Accuracy (%) | Inference Time/Character (NVIDIA Jetson TX2) | Inference Energy/Character (NVIDIA Jetson TX2) |
|------------------|---|---------------------|--|--|
| This work | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → Features (85) → EEdGeNet | 89.83 ± 0.19 | 914.18 ms | 5406.84 mJ |
| 1 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Standard ASR → Features (85) → EEdGeNet | 80.80 ± 0.15 | 1501.48 ms | 5051.92 mJ |
| 2 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → ICA → Features (85) → EEdGeNet | 12.03 ± 0.15 | 5492.51 ms | 18454.83 mJ |
| 3 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → MSPCA → Features (85) → EEdGeNet | 87.12 ± 1.46 | 828.93 ms | 5153.42 mJ |
| 4 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → VMD → Features (85) → EEdGeNet | 82.96 ± 1.28 | 2295.77 ms | 7369.42 mJ |
| 5 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → EWT → Features (85) → EEdGeNet | 79.67 ± 0.82 | 2212.58 ms | 7102.38 mJ |
| 6 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → Reconstructed Phase-Space Features (13) → EEdGeNet | 82.32 ± 1.07 | 202.10 ms | 1324.31 mJ |
| 7 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → Fast Fractional Fourier Descriptors (26) → EEdGeNet | 73.55 ± 0.10 | 200.99 ms | 1317.04 mJ |
| 8 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → Poincaré–DWT Graphical Features (65) → EEdGeNet | 83.26 ± 1.29 | 229.84 ms | 1506.09 mJ |
| 9 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → SODP–Binary PSO Graphical Features (26) → EEdGeNet | 85.05 ± 1.25 | 171.75 ms | 1015.80 mJ |
| 10 | Raw Data → Train/Test Split → Bandpass (1–50 Hz) → Simplified ASR → Features (85) → 1D CNN–LSTM | 77.55 ± 0.88 | 969.24 ms | 5447.83 mJ |
| 11 | Raw Data → Train/Test Split → MSPCA → CWT → ShuffleNet (pretrained on ImageNet) | N/A | N/A | N/A |

Table 4. Ablation study of the proposed real-time EEG-based handwriting recognition pipeline, spanning preprocessing, feature extraction, and classifier variants. Alternatives such as MSPCA, reconstructed phase-space features (13), fast fractional Fourier descriptors (26), SODP–Binary PSO Graphical Features (26), and Poincaré–DWT graphical features (65) reduce inference time and energy consumption on the NVIDIA Jetson TX2 but incur notable losses in accuracy. By contrast, the proposed pipeline with simplified ASR, the full 85 handcrafted features, and the EEdGeNet model achieves the highest per-character classification accuracy while maintaining a favorable trade-off between accuracy, latency, and energy efficiency.

Ablation study of feature extraction techniques

We next explored alternative feature representations to evaluate the trade-off between model accuracy and computational efficiency. Reconstructed phase-space descriptors (13 features)⁴⁰, fast fractional Fourier descriptors (26 features)⁴⁴, Second-order Differential Plot (SODP) Binary Particle Swarm Optimization (PSO) graphical features (26 features)⁴⁵, and Poincaré–DWT graphical features (65 features)³⁵ significantly reduced inference latency (200–400 ms) and per-character energy consumption. However, their accuracies ranged only from $73.55\% \pm 0.10\%$ to $85.06\% \pm 1.25\%$, notably lower than the proposed pipeline. These findings suggest that aggressive feature compaction or reliance on handcrafted graphical descriptors enhances computational efficiency on the Jetson TX2 but does so at the cost of predictive performance. Overall, the full 85-feature set provides the optimal balance between classification accuracy and inference latency, offering the best trade-off for real-time edge deployment.

Ablation study of classifier techniques

Finally, we assessed alternative classifier architectures to examine their suitability for real-time edge deployment. A 1D CNN–LSTM trained on the full handcrafted feature set achieved moderate accuracy ($77.55\% \pm 0.88\%$) with an inference latency of 969.24 ms. In contrast, an image-based ShuffleNet pretrained on ImageNet⁴⁶ proved impractical for real-time operation: generating canonical wavelet transform (CWT) images from raw EEG required 438.20 ms per character on our local machine, compared with only 41.25 ms per character for the handcrafted feature pipeline. Thus, TF-image generation was approximately $10.6\times$ slower than direct feature extraction, rendering such approaches unsuitable for per-character, real-time inference on edge platforms. Overall, the proposed EEdGeNet architecture demonstrates the most favorable balance between decoding accuracy, computational efficiency, and deployability on low-power embedded systems.

Taken together, these ablations underscore that the proposed pipeline, simplified ASR followed by 85 physiologically grounded handcrafted features and the EEdGeNet model, delivers the most favorable balance between accuracy, latency, and energy efficiency. This configuration consistently outperformed all alternative pipelines, establishing it as the most effective architecture for real-time imagined handwriting recognition on edge devices.

Comparative evaluation with existing methods

Performance evaluation of the custom model against state-of-the-art architectures

We conducted a comparative evaluation of our proposed EEdGeNet model against widely used EEG decoding architectures, including DeepConvNet⁴⁷, ShallowConvNet⁴⁷, EEGNet⁴⁸, TCNet⁴¹, EEGTCNet⁴⁹, TCNet-Fusion⁵⁰, Attention Multi-Branch CNN⁵¹, and ATCNet²¹. The comparative results, summarized in Table 5, highlight the performance of these models on both raw EEG signals and feature-extracted data.

On raw EEG data, accuracies ranged from $26.17 \pm 1.47\%$ (EEGNet) to $51.42 \pm 1.14\%$ (Attention Multi-Branch CNN), reflecting the inherent difficulty of decoding raw EEG recordings. Although our EEdGeNet model achieved $43.63 \pm 2.77\%$ under this setting, none of the evaluated architectures, including our own, delivered performance sufficient for reliable classification of imagined handwritten characters directly from raw EEG.

When trained on the curated set of 85 extracted features, however, performance improved across nearly all models (with the exception of ShallowConvNet), though the degree of improvement varied. For example, EEGNet reached $43.37 \pm 0.74\%$, while TCNet achieved $76.95 \pm 0.10\%$, illustrating the benefit of incorporating meaningful features. Nevertheless, existing architectures remained limited in their ability to exploit the structured feature space effectively. By contrast, our proposed EEdGeNet model, explicitly designed for feature-based inputs, achieved the highest test accuracy ($89.83 \pm 0.19\%$) and consistently outperformed all baselines across precision, recall, and F1-score. Unlike state-of-the-art EEG architectures, which are optimized for raw time-series data where convolutional and pooling layers target oscillatory dynamics, EEdGeNet preserves feature semantics and emphasizes inter-channel dependencies. Its first stage employs a Temporal Convolutional Network (TCN) block to capture temporal relationships across features, followed by a Dense Transformation Block that leverages a deep Multilayer Perceptron (MLP) to learn non-linear mappings and progressively refine discriminative representations. This architecture enables alignment with the structure of extracted EEG features, promoting subject-generalizable learning and superior decoding accuracy.

Overall, these findings establish EEdGeNet as a feature-aware architecture that not only achieves state-of-the-art accuracy but also delivers stable and generalizable performance, setting a new benchmark for EEG-based imagined handwriting recognition.

Comparison of the proposed framework with existing approaches

Table 6 presents a comparative analysis of recent studies on imagined motor and speech EEG decoding alongside our proposed approach. Our method constitutes the first demonstration of real-time, high-accuracy decoding of imagined handwriting from non-invasive EEG signals on a portable edge device. It exhibits strong generalizability across multiple participants and achieves low-latency performance, making it suitable for practical real-time assistive communication applications. Moreover, while prior works such as Skomrock et al.⁵³ and Willett et al.¹¹ reported real-time decoding delays of about 821 ms and 1000 ms, respectively for different use cases with invasive neural recordings, our non-invasive pipeline achieves per-character inference in 202.62 ms with $88.84\% \pm 0.09\%$ accuracy. This establishes our framework as a safer and more practical alternative for real-world assistive BCI applications.

Discussion

This study establishes the feasibility of real-time decoding of imagined handwriting from non-invasive EEG recordings through a lightweight pipeline designed for efficient deployment on edge devices. While we

| Model | Raw Data | | | Features (85) | | |
|--|---------------------|-----------------------|--|---------------------|-----------------------|--|
| | Test Acc. (%) | 95% CI | Precision (P) (%) Recall (R) (%) F1-Score (F) (%) | Test Acc. (%) | 95% CI | Precision (P) (%) Recall (R) (%) F1-Score (F) (%) |
| DeepConvNet ⁴⁷ | 35.89 ± 1.42 | [34.13, 37.64] | P: 40.64 ± 1.30 R: 35.89 ± 1.42 F: 34.55 ± 1.45 | 52.71 ± 1.64 | [48.64, 56.78] | P: 54.29 ± 1.22 R: 52.71 ± 1.64 F: 52.58 ± 1.66 |
| ShallowConvNet ⁴⁷ | 43.49 ± 0.57 | [42.78, 44.20] | P: 43.87 ± 0.72 R: 43.49 ± 0.57 F: 43.24 ± 0.64 | 41.10 ± 0.18 | [40.65, 41.55] | P: 41.43 ± 0.28 R: 41.10 ± 0.18 F: 40.98 ± 0.24 |
| EEGNet ⁴⁸ | 26.17 ± 1.47 | [24.13, 28.22] | P: 29.99 ± 1.05 R: 26.17 ± 1.47 F: 24.99 ± 1.55 | 43.37 ± 0.74 | [41.12, 45.62] | P: 44.42 ± 0.59 R: 43.37 ± 0.74 F: 42.76 ± 0.77 |
| TCNet ⁴¹ | 29.10 ± 1.03 | [27.83, 30.38] | P: 30.24 ± 0.96 R: 29.10 ± 1.28 F: 29.15 ± 1.25 | 76.95 ± 0.10 | [76.64, 77.26] | P: 77.01 ± 0.08 R: 76.95 ± 0.10 F: 76.95 ± 0.09 |
| EEGTCNet ⁴⁹ | 28.92 ± 3.48 | [24.09, 33.74] | P: 32.22 ± 3.33 R: 28.92 ± 3.48 F: 28.05 ± 3.70 | 57.35 ± 1.02 | [54.26, 60.45] | P: 58.14 ± 1.34 R: 57.35 ± 1.02 F: 57.26 ± 1.07 |
| TCNet-Fusion ⁵⁰ | 27.65 ± 1.11 | [26.28, 29.03] | P: 28.97 ± 0.78 R: 27.65 ± 1.11 F: 27.52 ± 1.09 | 62.54 ± 0.17 | [62.11, 62.96] | P: 62.60 ± 0.13 R: 62.54 ± 0.17 F: 62.49 ± 0.16 |
| Attention Multi Branch CNN ⁵¹ | 51.42 ± 1.14 | [50.01, 52.83] | P: 52.38 ± 1.32 R: 51.42 ± 1.14 F: 51.52 ± 1.18 | 62.80 ± 0.52 | [61.52, 64.08] | P: 62.91 ± 0.52 R: 62.80 ± 0.52 F: 62.77 ± 0.52 |
| ATCNet ²¹ | 34.53 ± 2.66 | [30.85, 38.22] | P: 37.44 ± 3.34 R: 34.53 ± 2.66 F: 33.49 ± 2.48 | 58.64 ± 1.28 | [54.73, 62.54] | P: 59.65 ± 1.19 R: 58.64 ± 1.28 F: 58.34 ± 1.38 |
| EEdGeNet (This work) | 43.63 ± 2.77 | [43.19, 50.08] | P: 47.46 ± 3.20 R: 46.63 ± 3.44 F: 46.56 ± 3.41 | 89.83 ± 0.19 | [89.56, 90.09] | P: 89.87 ± 0.18 R: 89.83 ± 0.19 F: 89.83 ± 0.19 |

Table 5. Comparative performance analysis of state-of-the-art models and the proposed custom architecture on the raw imagined handwriting EEG dataset and 85 features extracted data. The proposed EEdGeNet model trained on 85 extracted features achieves the highest test accuracy and F1-score, outperforming existing methods across all evaluated metrics.

evaluated EWT and VMD signal decomposition techniques^{42,43} within our ablation studies, our proposed preprocessing and feature extraction pipeline achieves superior performance because it is specifically tailored to imagined handwriting EEG rather than conventional motor imagery paradigms. Replacing the simplified ASR with conventional denoising methods produced inconsistent outcomes. Standard ASR and ICA degraded performance, while MSPCA⁵⁴ offered competitive but suboptimal results. While phase-space, fast fractional fourier descriptors SODP-Binary PSO, and Poincaré-DWT graphical features have shown effectiveness in clinical EEG classification^{35,40,44,45}, our results indicate that these methods are less suited for real-time imagined handwriting decoding. While CNN-LSTM architectures have shown strong performance in binary EEG classification tasks⁵⁵, our experiments demonstrate that the proposed EEdGeNet model offers superior accuracy and efficiency for real-time, multi-class imagined handwriting decoding. While pretrained CNNs such as AlexNet and VGG have shown strong performance for offline EEG classification^{46,56}, our analysis reveals that their high preprocessing latency makes them unsuitable for real-time edge deployment, where the proposed handcrafted feature-based EEdGeNet offers a far more efficient alternative. Overall, the proposed pipeline-simplified ASR with extracted features and the EEdGeNet model-provides the most balanced trade-off between accuracy, latency, and energy efficiency for real-time imagined handwriting decoding.

As illustrated in Fig. 7, the 2D t-SNE visualizations reveal how EEdGeNet progressively learns to separate the 27 imagined-characterclasses across its hierarchical blocks. Figure 7a depicts the feature-space distribution after handcrafted feature extraction, where each point represents a single trial(420 trials per class, aggregated across all 15 participants). Figure 7b shows the latent representation after the TCB block, while Figure 7c presents the final t-SNEembedding after the DTB block, where the 27 classes form well-defined and distinct clusters. Together, these visualizations demonstrate that EEdGeNet effectivelydisentangles the neural representations of all classes through its successive processing stages.

As shown in Fig. 6, most imagined characters are correctly classified, with strong diagonal dominance in the confusion matrix. Certain letters, such as 'b', achieve consistently high classification accuracy, whereas others, such as 'y', remain challenging and exhibit higher rates of misclassification. Occasional confusions with neighboring letters suggest overlapping neural representations, motivating future work on class-specific feature refinement and adaptive modeling strategies. While within-subject performance is strong, cross-subject generalization remains limited. Leave-one-subject-out (LOSO) experiments yielded an average accuracy of only $3.55 \pm 1.15\%$, with precision and recall showing similarly low values, reflecting substantial inter-subject variability. Addressing this limitation will be a key direction in the future, with emphasis on transfer learning, adaptive calibration, and more flexible architectures to improve subject-independent decoding.

Experimental results and comparative evaluations further highlight the trade-offs between accuracy, latency, and energy efficiency. The proposed pipeline for imagined handwriting recognition from EEG, simplified

| Study | Application | Dataset Description | Model Type | Classification Accuracy | Real-Time Compatible | Edge Inference | Latency |
|--|---|---|--|---|----------------------|-------------------------|------------------|
| Sadiq et al. (2019) ⁴³ | Motor imagery classification | BCI Competition III (IVa, IVb, IVc), 118 channels, 5 subjects | MEWT + MSPCA + LS-SVM/MLP/LMT | 98% (subject-dependent), 91.4% (subject-independent) | No | No | N/A |
| Ullah et al. (2021) ¹⁷ | Imagined character recognition | EEGMMIDB, BCI IV-2a, 10 subjects, 26 characters | CNN + Morlet wavelet | 99.45% (4-class), 95.2% (custom - 26 characters) | No | No | N/A |
| Jerrin et al. (2021) ²⁰ | Imagined speech classification | ASU dataset, 64-channel EEG | Transfer learning (ResNet50) | 92.5% (2-long words), 80% (3-short words), 79.7% (3-vowels), 91.4% (2-short/long words) | No | No | N/A |
| Altaheri et al. (2022) ²¹ | Motor imagery classification | BCI Competition IV 2a | ATCNet (attention + TCN + physics-informed) | 85.38% (subject - dependent), 70.97% (subject - independent), (4-class) | No | No | N/A |
| Xie et al. (2022) ²² | Motor imagery | BCI datasets (raw EEG) | Transformer with spatial-temporal attention | 83.31% (2-class), 74.44% (3-class), 64.22% (4-class) | No | No | N/A |
| Cho et al. (2022) ²⁴ | Grasp classification from motor imagery | 12 participants, 20 channels, 4 grasp types | CNN-BLSTM + Siamese networks | 86% (offline), 79% (online) - 2-class; 68% (offline), 65% (online) (4-class) | Yes | No | Not reported |
| Khademi et al. (2022) ²³ | Motor imagery classification | BCI Competition IV 2a | Hybrid CNN-LSTM (Inception-v3, ResNet-50) | Up to 92% (4-class) | No | No | N/A |
| Yu et al. (2023) ²² | Motor and mental imagery detection | BCI Competition datasets, 118 channels, multiple subjects | Hybrid framework (signal decomposition + feature selection + ML classifiers) | Up to 98.5% (varies with task/dataset) | No | No | N/A |
| Abdulghani et al. (2023) ¹⁹ | Imagined speech recognition | 8-channel EEG | LSTM + wavelet scattering | 92.5% (4-commands) | No | No | N/A |
| Alonso Vázquez et al. (2023) ¹⁸ | Spoken word recognition | 32-channel EEG, 28 subjects, 5 Spanish words | SVM, EEGNet | 91.04% (binary), 81.2% (5-words) | No | No | N/A |
| Amrani et al. (2024) ¹⁵ | EEG-to-text decoding | ZuCo v1.0 and v2.0 | EEG encoder + BART + GPT-4 | BLEU-1: 42.75%, ROUGE-1-F: 33.28%, BERTScore-F: 53.86% | No | No | N/A |
| This work | Imagined handwriting decoding | 32-channel EEG, 15 subjects | TCN-MLP-Softmax | 88.84% ± 0.09% (27-class) | Yes | Yes (Jetson TX2) | 202.62 ms |

Table 6. Comparative analysis of recent imagined motor and speech EEG decoding methods with our proposed work. Our approach represents the first demonstration of real-time, high-accuracy decoding of imagined handwriting from non-invasive EEG signals, implemented on a portable edge device.

ASR with 85 handcrafted features, and the EEdGeNet model achieved the highest accuracy ($89.83 \pm 0.19\%$) while maintaining favorable inference time and energy balance, underscoring its suitability for real-time edge deployment. Although challenges remain in cross-subject generalization and decoding of inherently ambiguous characters, the proposed framework sets a benchmark for accuracy, efficiency, and deployability, paving the way toward practical neural communication systems.

In future work, we plan to substantially expand the participant pool to enhance model robustness and population-level generalization. Larger, multi-site datasets, similar to those employed in prior large-scale EEG studies^{52,57} will be collected to better capture inter-subject variability and improve cross-subject transferability of imagined handwriting decoding.

Conclusion

This work addresses two fundamental challenges in non-invasive neural decoding: the low signal quality of EEG and the difficulty of deploying real-time models on resource-constrained edge devices. We presented the first real-time demonstration of high-accuracy decoding of imagined handwriting from non-invasive EEG signals using a portable edge device. Our proposed framework combines efficient feature engineering with a hybrid TCN-MLP architecture to enable low-latency inference on resource-constrained hardware, achieving $89.83\% \pm 0.19\%$ classification accuracy with a per-character inference latency of only 914.18 ms on the NVIDIA Jetson TX2. Through comprehensive evaluation across multiple participants, we demonstrated the robustness and generalizability of the system, highlighting its potential for scalable deployment in practical BCI applications. Importantly, our analysis shows that a carefully selected subset of just ten features, Hjorth mobility, RMS, Hjorth complexity, delta power, graphical_A_ellipse_area, mean dominant frequency, graphical_D3_ellipse_area, graphical_D1_ellipse_area, graphical_D2_ellipse_area, and graphical_D4_ellipse_area, can deliver nearly equivalent accuracy to the full 85-feature set while reducing inference latency by $4.51\times$. This trade-off between accuracy and efficiency underscores the viability of our system for real-time, on-device neural decoding in assistive communication settings. However, our results stem from an optimized preprocessing and feature extraction pipeline combined with a lightweight neural model (EEdGeNet), reducing the computational

overhead typical of motor imagery classifiers—though outcomes may vary across different applications. Future directions will include expanding the dataset to encompass more participants and diverse recording conditions, which will be critical for enhancing robustness and generalizability. Additional efforts will focus on incorporating hybrid feature sets and adaptive deep learning strategies to improve cross-subject decoding performance.

Data availability

The datasets and code used in this study are available from the corresponding author upon reasonable request.

Received: 8 August 2025; Accepted: 16 October 2025

Published online: 20 November 2025

References

- Zhang, H. et al. Brain-computer interfaces: the innovative key to unlocking neurological conditions. *Int. J. Surg.* **110**, 5745–5762 (2024).
- Liu, X.-Y. et al. Recent applications of EEG-based brain-computer-interface in the medical field. *Mil. Med. Res.* **12**, 14 (2025).
- Sun, J. et al. A hybrid deep neural network for classification of schizophrenia using EEG data. *Sci. Rep.* **11**, 4706 (2021).
- Ahmadlou, M., Adeli, H. & Adeli, A. Fractality analysis of frontal brain in major depressive disorder. *Int. J. Psychophysiol.* **85**, 206–211 (2012).
- Janapa Reddi, V. et al. Mlperf mobile inference benchmark: An industry-standard open-source machine learning benchmark for on-device AI. *Proceedings of Machine Learning and Systems* **4**, 352–369 (2022).
- Haddidi, R., Cao, J., Ryoo, M. S. & Kim, H. Toward collaborative inferencing of deep neural networks on internet-of-things devices. *IEEE Internet Things J.* **7**, 4950–4960 (2020).
- Safari, M., Shalhaf, R., Bagherzadeh, S. & Shalhaf, A. Classification of mental workload using brain connectivity and machine learning on electroencephalogram data. *Sci. Rep.* **14**, 9153 (2024).
- Petrosyan, A. et al. Speech decoding from a small set of spatially segregated minimally invasive intracranial EEG electrodes with a compact and interpretable neural network. *J. Neural Eng.* **19**, 066016 (2022).
- Liu, D. et al. Intracranial brain-computer interface spelling using localized visual motion response. *Neuroimage* **258**, 119363 (2022).
- Metzger, S. L. et al. A high-performance neuroprosthesis for speech decoding and avatar control. *Nature* **620**, 1037–1046 (2023).
- Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M. & Shenoy, K. V. High-performance brain-to-text communication via handwriting. *Nature* **593**, 249–254 (2021).
- Mishra, A. R. et al. SignEEG v10: multimodal dataset with electroencephalography and hand-written signature for biometric systems. *Sci. Data* **11**, 718 (2024).
- Zhao, X. et al. Classification of sleep apnea based on EEG sub-band signal characteristics. *Sci. Rep.* **11**, 5824 (2021).
- Kim, S.-K., Kim, H., Kim, S. H., Kim, J. B. & Kim, L. Electroencephalography-based classification of alzheimer's disease spectrum during computer-based cognitive testing. *Sci. Rep.* **14**, 5252 (2024).
- Amrani, H., Micucci, D. & Napoletano, P. Deep representation learning for open vocabulary electroencephalography-to-text decoding. *IEEE J. Biomed. Health Inform.* (2024).
- Islam, M.S., Hussain, I., Rahman, M.M., Park, S.J. & Hossain, M.A. Explainable artificial intelligence model for stroke prediction using EEG signal (2022).
- Ullah, S. & Halim, Z. Imagined character recognition through EEG signals using deep convolutional neural network. *Med. Biol. Eng. Comput.* **59**, 1167–1183 (2021).
- Alonso-Vázquez, D., Mendoza-Montoya, O., Caraza, R., Martínez, H. R. & Antelis, J. M. EEG-based classification of spoken words using machine learning approaches. *Computation* **11**, 225 (2023).
- Abdulghani, M. M., Walters, W. L. & Abed, K. H. Imagined speech classification using EEG and deep learning. *Bioengineering* **10**, 649 (2023).
- Panachakel, J. T. & Ganesan, R. A. Decoding imagined speech from EEG using transfer learning. *IEEE Access* **9**, 135371–135383 (2021).
- Altaheri, H., Muhammad, G. & Alsulaiman, M. Physics-informed attention temporal convolutional network for EEG-based motor imagery classification. *IEEE Trans. Industr. Inf.* **19**, 2249–2258 (2022).
- Xie, J. et al. A transformer-based approach combining deep learning network and spatial-temporal information for raw EEG classification. *IEEE Trans. Neural Syst. Rehabil. Eng.* **30**, 2126–2136 (2022).
- Khademi, Z., Ebrahimi, F. & Kordy, H. M. A transfer learning-based CNN and LSTM hybrid deep learning model to classify motor imagery EEG signals. *Comput. Biol. Med.* **143**, 105288 (2022).
- Cho, J.-H., Jeong, J.-H. & Lee, S.-W. Neurograsp: Real-time EEG classification of high-level motor imagery tasks using a dual-stage deep learning framework. *IEEE Trans. Cybern.* **52**, 13279–13292 (2021).
- Nakanishi, M. et al. Enhancing detection of sseps for a high-speed brain speller using task-related component analysis. *IEEE Trans. Biomed. Eng.* **65**, 104–112 (2017).
- Farwell, L. A. & Donchin, E. Talking off the top of your head: Toward a mental prosthesis utilizing event-related brain potentials. *Electroencephalogr. Clin. Neurophysiol.* **70**, 510–523 (1988).
- Huotari, N. et al. Sampling rate effects on resting state fMRI metrics. *Front. Neurosci.* **13**, 279 (2019).
- Beckmann, C. F., DeLuca, M., Devlin, J. T. & Smith, S. M. Investigations into resting-state connectivity using independent component analysis. *Philosoph. Trans. Royal Soci. B: Biol. Sci.* **360**, 1001–1013 (2005).
- Chang, C.-Y., Hsu, S.-H., Pion-Tonachini, L. & Jung, T.-P. Evaluation of artifact subspace reconstruction for automatic artifact components removal in multi-channel EEG recordings. *IEEE Trans. Biomed. Eng.* **67**, 1114–1121 (2019).
- Blum, S., Jacobsen, N. S., Bleichner, M. G. & Debener, S. A riemannian modification of artifact subspace reconstruction for EEG artifact handling. *Front. Hum. Neurosci.* **13**, 141 (2019).
- Yeh, C. & colleagues. Corticomuscular coherence study. *Scientific Reports* (2025). Band-pass filtering EEG 1–50 Hz with zero-phase FIR filter.
- Mahadevan, R. & colleagues. Eeg–fmri emotion recognition pipeline. *Frontiers in Computational Neuroscience* (2024). EEG band-pass filtered 1–50 Hz before artifact removal.
- Choi, J. & colleagues. Ear-EEG brain-computer interface methods. *Sensors* **18** (2018). 4th-order Butterworth band-pass 1–50 Hz.
- Mullen, T. R. et al. Real-time neuroimaging and cognitive monitoring using wearable dry EEG. *IEEE Trans. Biomed. Eng.* **62**, 2553–2567 (2015).
- Akbari, H. et al. Recognizing seizure using poincaré plot of EEG signals and graphical features in dwt domain. *Bratislava Med. J./ Bratislavske Lekarske Listy* **124**, 12–24 (2023).
- Hjorth, B. EEG analysis based on time domain properties. *Electroencephalogr. Clin. Neurophysiol.* **29**, 306–310 (1970).
- Bandt, C. & Pompe, B. Permutation entropy: A natural complexity measure for time series. *Phys. Rev. Lett.* **88**, 174102 (2002).
- Pernagallo, G. Random walks, hurst exponent, and market efficiency: G. pernegallo. *Qual. Quantity* **59**, 1097–1119 (2025).

39. Sadiq, M. T., Akbari, H., Siuly, S., Li, Y. & Wen, P. Alcoholic EEG signals recognition based on phase space dynamic and geometrical features. *Chaos, Solit. Fract.* **158**, 112036 (2022).
40. Akbari, H. et al. Depression recognition based on the reconstruction of phase space of EEG signals and geometrical features. *Appl. Acoust.* **179**, 108078 (2021).
41. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. [arXiv:1803.01271](https://arxiv.org/abs/1803.01271) (2018).
42. Sadiq, M. T. et al. Motor imagery BCI classification based on multivariate variational mode decomposition. *IEEE Trans. Emerg. Top. Comput. Intell.* **6**, 1177–1189 (2022).
43. Sadiq, M. T. et al. Motor imagery EEG signals decoding by multivariate empirical wavelet transform-based framework for robust brain-computer interfaces. *IEEE Access* **7**, 171431–171451 (2019).
44. Sadiq, M. T., Yousaf, A., Siuly, S. & Almogren, A. Fast fractional Fourier transform-aided novel graphical approach for EEG alcoholism detection. *Bioengineering* **11**, 464 (2024).
45. Akbari, H. et al. Depression detection based on geometrical features extracted from SODP shape of EEG signals and binary PSO. *Traitement du Signal* **38**, 13–26 (2021).
46. Sadiq, M. T. et al. Exploiting pretrained CNN models for the development of an EEG-based robust BCI framework. *Comput. Biol. Med.* **143**, 105242 (2022).
47. Schirrmester, R. T. et al. Deep learning with convolutional neural networks for EEG decoding and visualization. *Hum. Brain Mapp.* **38**, 5391–5420 (2017).
48. Lawhern, V. J. et al. Eegnet: A compact convolutional neural network for EEG-based brain-computer interfaces. *J. Neural Eng.* **15**, 056013 (2018).
49. Ingolfsson, T. M. et al. Eeg-tcnnet: An accurate temporal convolutional network for embedded motor-imagery brain-machine interfaces. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2958–2965 (IEEE, 2020).
50. Musallam, Y. K. et al. Electroencephalography-based motor imagery classification using temporal convolutional network fusion. *Biomed. Signal Process. Control* **69**, 102826 (2021).
51. Altuwajri, G. A., Muhammad, G., Altaheri, H. & Alsulaiman, M. A multi-branch convolutional neural network with squeeze-and-excitation attention blocks for EEG-based motor imagery signals classification. *Diagnostics* **12**, 995 (2022).
52. Yu, X., Aziz, M. Z., Sadiq, M. T., Fan, Z. & Xiao, G. A new framework for automatic detection of motor and mental imagery EEG signals for robust BCI systems. *IEEE Trans. Instrum. Meas.* **70**, 1–12 (2021).
53. Skomrock, N. D. et al. A characterization of brain-computer interface performance trade-offs using support vector machines and deep neural networks to decode movement intent. *Front. Neurosci.* **12**, 763 (2018).
54. Sadiq, M. T., Yu, X., Yuan, Z. & Aziz, M. Z. Motor imagery BCI classification based on novel two-dimensional modelling in empirical wavelet transform. *Electron. Lett.* **56**, 1367–1369 (2020).
55. Hussain, W., Sadiq, M. T., Siuly, S. & Rehman, A. U. Epileptic seizure detection using 1 d-convolutional long short-term memory neural networks. *Appl. Acoust.* **177**, 107941 (2021).
56. Sadiq, M. T., Yu, X. & Yuan, Z. Exploiting dimensionality reduction and neural network techniques for the development of expert brain-computer interfaces. *Expert Syst. Appl.* **164**, 114031 (2021).
57. Sadiq, M. T. et al. Toward the development of versatile brain-computer interfaces. *IEEE Trans. Artif. Intell.* **2**, 314–328 (2021).

Acknowledgements

The authors gratefully acknowledge Mr. Brennan Borchert, Ph.D. student in the Department of Electrical and Computer Engineering at the University of Florida, for his valuable feedback on language refinement and manuscript presentation. His careful suggestions helped improve the clarity, readability, and overall quality of this work.

Author contributions

B.C. and O.S. conceptualized the idea. O.S., R.S., and D.V. were responsible for participant recruitment and data acquisition. O.S. and D.V. led the feature extraction and selection process. R.S. and O.S. performed machine learning model analysis and performance evaluation. D.V. and A.P. contributed to channel selection and spatial electrode configuration. S.J. conducted dataset-level analyses. P.L. provided input on experimental design. A.Ka. and A.Kh. contributed to the technical discussions. O.S. and R.S. wrote the main manuscript. B.C. supervised the project and provided overall research direction and guidance. All authors reviewed and approved the final manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to O.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025