Original article

Evaluating the efficacy of motor imagery classifiers: Linear Discriminant Analysis and a Multi-Layer Perceptron neural network

Krishnan T[1], Tivnan M[2]

Submitted: February 17, 2024, Revised: version 1, May 17, 2024
Accepted: May 20, 2024

Abstract
This study provides an analysis of two methods of electroencephalography (EEG) signal classification – Linear Discriminant Analysis (LDA) and a Multi-Layer Perceptron (MLP) neural network – in the context of Brain-Computer Interface (BCI) systems. Since EEG data is inherently high-dimensional and complex, the ability to accurately distinguish motor imagery classes is crucial for improving BCI applications. The results demonstrate that the MLP classifier achieves a significantly higher classification accuracy (99.63%) in comparison to the LDA classifier (64.15%). However, temporal differences in an individual's EEG patterns coupled with non-generalizability to other subjects may pose non-trivial impediments to the development and commercialization of the MLP model. Regardless, this substantial difference in classification accuracy highlights the effectiveness of neural networks in handling complex data and may have the potential to enhance future BCI performances, holding great promise for individuals, specifically those with motor-inhibiting neurodegenerative diseases, who depend on these systems for an improved quality of life.

Keywords
Machine Learning, Electroencephalography, Linear Discriminant Analysis, Multi-Layer Perceptron, Brain-Computer Interface, Neural network, Motor imagery

---

[1]Corresponding author: Teyj Krishnan, Tesla STEM High School, 4301 228th Ave NE, Redmond, WA 98053, USA. teyj.krishnan@gmail.com
[2]Matthew Tivnan, Harvard Medical School, 25 Shattuck St, Boston, MA 02115, USA mtivnan@mgh.harvard.edu

Introduction

In the field of neuroscience, Brain-Computer Interfaces (BCI) have emerged as a transformative technology. They enable users to establish a link between their brain signals and external devices, allowing these individuals to control such devices solely with their minds. This can prove extremely useful to those who suffer from motor disabilities, like paralysis, which can partially or completely impair limb movement (1). Among other types of BCIs, Electroencephalography (EEG) based BCIs are often utilized due to their low cost and non-invasive nature (2). These devices utilize the placement of small electrodes on the scalp of the user, which monitor and record electrical impulses, then quantify these recordings for data analytics.

One of the fundamental tasks of EEG-based BCIs is accurately classifying motor imagery, in which the device records the user's electrical activity when they intend to make a movement, then the BCI accurately decodes what the motor intent of the user was (3). If successful, the classification can allow for a means of neurorehabilitation for patients who suffer from neurodegenerative diseases that impair motor movement (4).

This paper investigates and compares two methods of EEG classification: Linear Discriminant Analysis (LDA) and a Multi-Layer Perceptron (MLP) neural network. Linear Discriminant Analysis has widely been used for EEG classification. It attempts to find a linear combination of features that maximizes the discriminability between two or more classes. In the past, LDA has shown promise due to its simplicity and efficiency. In comparison, an MLP is capable of extracting intricate features from a raw EEG dataset, possibly leading to improvements in classification accuracy (5).

Numerous studies have investigated the classification of motor imagery-based EEG signals, contributing to the evolving landscape on neuro-rehabilitative devices. In the study presented in (6), the performance of Support Vector Machine (SVM), LDA, and MLP classifiers was evaluated using a combination of eight distinct features as a vector. The study demonstrated that the SVM classifier achieved a classification accuracy of 98.8%, and this classifier was later used for the offline controlling of a prototypical arm. The study conducted in (7) proposes a dual-Convolutional neural network model that effectively integrates left and right-side electrode information. Through using the time series of both hemisphere's signals as the input to the dual-CNN, the model accurately classified motor imagery with an accuracy of 98.88%.

Materials and Methods

Dataset

We used the open-source dataset from (8), from which each classification method was tested. Its study focused on inducing motor-imagery brain activity, as well as understanding the differences in performances from subject to subject. In the study, each subject sat in front of a computer screen. For each of trials, the screen was black with a cross in the center for two seconds, then displayed text that either said "left hand" or "right hand." Upon the text being displayed, each subject visualized themselves moving the corresponding hand for three seconds, and then the screen returned to a cross on a black screen. While these motor imagery tasks were being performed, EEG and EMG data were collected simultaneously, using

64 electrodes for EEG and 4 electrodes for EMG – each electrode collected data for 358,400 time-samples, at a sampling rate of 512 samples per second. The data was subsequently organized into four broader categories: resting state, noise, motor imagery, and real hand movement. The two motor imagery arrays, imagery_left and imagery_right, are of shape (68, 358,400). Rigorous data processing techniques were applied, including indexing and rejection of bad trials, Event-Related Desynchronization/Synchronization analysis, and identification of EMG-correlated trials. Additionally, (8) calculated classification accuracy using the common spatial pattern (CSP) for feature extraction and Fisher's Linear Discriminant Analysis (FLDA) for classification, ultimately yielding a mean classification accuracy of 67.46% (±13.17%). For the study completed in this research, subject 1 was used for both training and testing of the developed models.

Data Visualization

Because of the high-dimensional nature of the dataset, the data was first visualized to facilitate the subsequent analysis. In order to visualize the dataset, a two-step preprocessing approach was applied as follows:

1. Mean Centering: First, each sensor's average signal was computed across all time stamps. Subsequently, the data was centered around the computed average by subtracting the mean from each of the signal's values. This process thereby eliminated positively and negatively skewed data points, preventing potential sources of bias.

2. Principal Component Analysis (PCA): A fundamental dimensionality reduction technique, PCA operates by identifying various linear combinations of the centered signal data, then ordering these combinations by the variance they represent, from greatest to least. This allows for a projection of the data onto a reduced set of principal components, transforming it into a lower-dimensional space while still retaining much of the original variance in the dataset.

After preprocessing was completed, the revised data was visualized as shown in Figure 1.

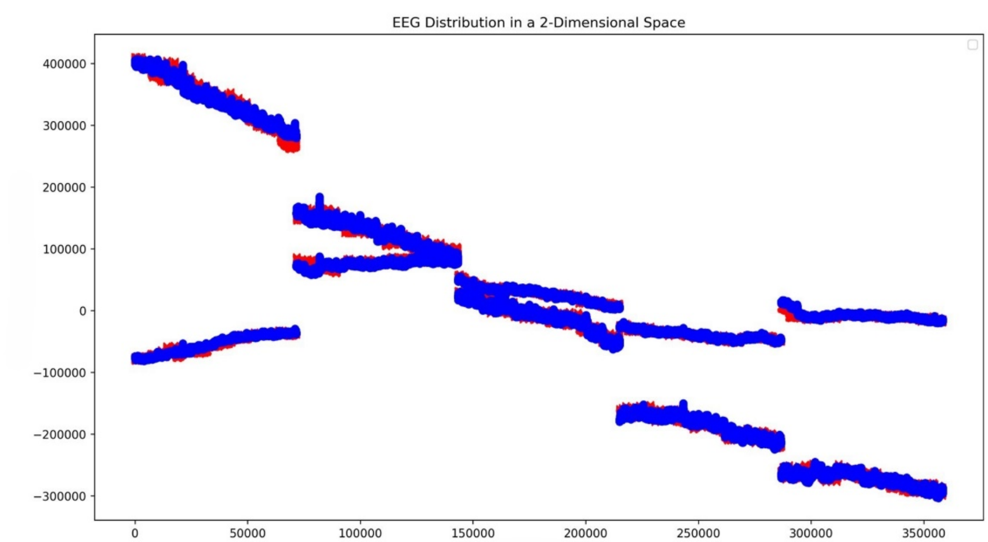

Figure 1. A scatterplot of the first two principal components of the combined left- and right- hand motor imagery EEG data, where the x-axis is the time-sample, and the y-axis is the EEG signal. The red markers represent left hand data, while the blue markers represent right-hand data.

In Figure 1, it shows that despite the two-step preprocessing steps taken, the data still retained an extremely high degree of overlap. This substantial overlap indicates that that the distinctive features within the left- and right-hand imagery datasets do not have much raw discrimination and underscores the complexity of the discriminatory task. Therefore, it would be extremely difficult to separate the left- and right-hand data with a boundary at this level of dimensionality – the data cannot be reduced to just two principal components.

Classification Methods

In this section, we will go over the two classification methods used, Linear Discriminant Analysis and Multi-Layer Perceptron. Prior to training, data was prepared and randomized using an 80-20 training-testing split. This allowed for further metric assessments of model performance, prevention of data leakage, and cross-validation capabilities.

Linear Discriminant Analysis

The goal of Linear Discriminant Analysis (LDA) is to determine a linear boundary for decision rule for binary classification. LDA additionally works at maximizing the between-class scatter and minimizing the within-class scatter. This is especially relevant to this particular dataset, as it contains 68 different dimensions, and only two classes. Therefore, with this model and other LDA models, the aim is to find the linear combinations of features in the dataset that achieve the greatest separation between classes, represented by a linear boundary.

The features used for the linear model were the imagery_left and imagery_right subsets after mean centering was applied. The model was trained on 80% of the dataset, and subsequently tested for classification on the remaining 20%.

For this task, no additional parameters were passed, and the model was trained and evaluated on the raw EEG data.

Neural Network Classifier

For this study, a Multi-Layer Perceptron (MLP) neural network classifier was utilized. A fundamental component of machine learning techniques, MLP neural networks are often utilized for their ability to model complex, nonlinear relationships of data. As outlined in Figure 2, an MLP consists of an input layer, which receives the raw EEG data in this study, a series of hidden layers, allowing for feature extraction, non-linearity, and learning hierarchies, and an output layer, which produces the result of left- or right- hand classification for each signal.

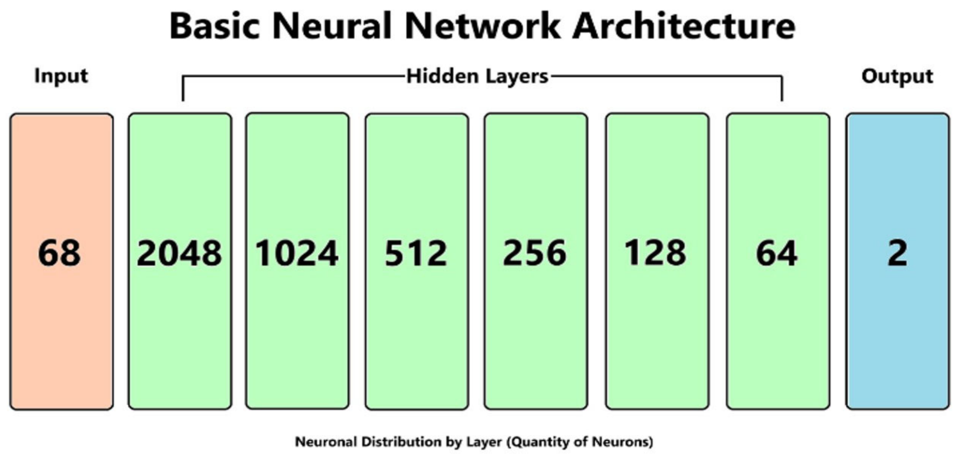

Figure 2. The architecture for the Multi-Perceptron Classifier, which contains an input layer, six hidden layers, and an output layer for classification. The values within each layer represent to the quantity of nodes, where the 68 input nodes correspond to the EEG and EMG channels for input, while the two output nodes correspond to left-hand and right-hand classification.

The architecture resembles a feedforward neural network, in which each subsequent layer contains a more precise number of neurons. The model utilizes the same features as the LDA, using imagery_left and imagery_right after the application of mean centering. The same 80-20 train-test split was implemented. The input layer takes the two motor imagery arrays as features, and the output layer is the binary classifier. Within the training architecture, a series of hyperparameters were used, as outlined below.

1. Learning Rate: The learning rate of the model, set at $1x10^{-4}$, determined the step size at which the model's parameters were updated during training. As a relatively small learning rate, a multitude of benefits were obtained, including the model's stability (optimal parameter overshoot prevention) and convergence (slower convergence for more precise solution)

2. Adam Optimizer: The Adam Optimizer, a popular optimization algorithm for neural network training, was implemented. It combines the advantages of two extensions of stochastic gradient descent (SGD): AdaGrad and RMSProp. The optimizer utilizes moving averages of the gradients to adapt learning rates for each parameter individually.

3. Negative Log Likelihood Loss Function (NLLLoss): The loss function used was the Negative-Log-Likelihood Loss (NLLLoss), which measures the dissimilarity between the predicted class labels and the true class labels. It was used as a metric of training over time, demonstrating the continuous improvement of the model during training.

Results and Discussion

After applying LDA to Subject 1, the classification accuracy yielded was ~64.15%, validating the classification accuracy calculated in (8). However, when the Multi-Layer Perceptron neural network was applied to Subject 1, the classification accuracy was significantly higher, at 99.63%. The confusion matrices for both LDA and MLP for Subject 1 are shown in Figure 3. From each matrix's values, we can calculate additional metrics to measure each model's validity, as seen in Table 1.

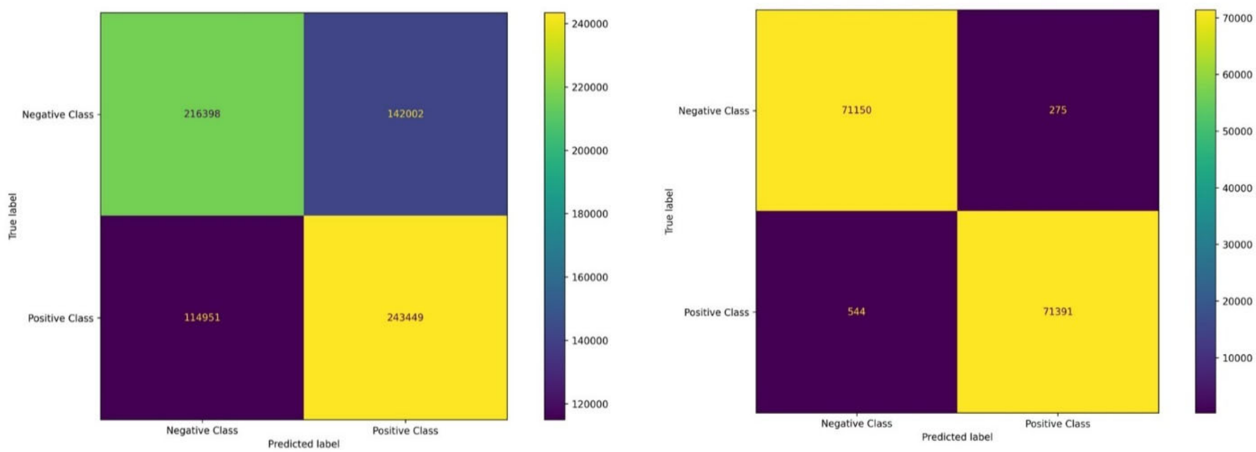

Figure 3. Confusion matrices for LDA (left) and MLP (right) for Subject 1. In each matrix, the top left value is true negatives (TN), the right value is false positives (FP), the bottom left value is false negatives (FN), and the bottom right value is true positives (TP).

Table 1. Accuracies of the three models as a function of percent labeled data.

Model	Recall	Precision	F1 Score	Accuracy (%)
LDA	0.679	0.632	0.655	64.15
MLP	0.992	0.996	0.994	99.63

Additionally, computation times were calculated on a time per second basis for both models. Across 71.680 times samples, at a sample rate of 512 samples per second, the LDA model averaged ~1.196 milliseconds to make a decision for one second of motor imagery, while the MLP averaged 41.027 milliseconds per second of motor imagery.

The results of this study demonstrate that the Multi-Layer Perceptron model for EEG motor imagery classification achieves a significantly higher classification accuracy than that of the Linear Discriminant Analysis, a commonly used classification technique used for BCI data classification. This notable difference in accuracy reflects the more powerful discriminative capabilities of neural networks in processing the complex and high-dimensional EEG data, relative to the LDA model which struggled to effectively separate the two classes with a high degree of accuracy.

However, it is important to discuss potential limitations to this study. The first stems from the fact that EEG data generally reflects the recorded individual's level of concentration, meaning that high-focus tasks elicit high frequency EEG waveforms, and low-focus tasks elicit low frequency EEG waveforms. As a result of this, performing motor imagery tasks may result in temporal differences in EEG patterns, as an individual's concentration will likely never remain stagnant. In order to account for such temporal differences, it may be necessary to either develop a larger subset of training data which is gathered across varying degrees of alertness, or to modify the model to adjust all EEG signals to a standardized set of values. Further, these temporal differences may lead to inaccuracies for patients who were able to induce left-hand control at one time, but unable to do so at another time. The nuanced thoughts that allowed for left-hand control would need to be practiced repeatedly in order to ensure near-complete control, across all time domains. Due to the nature of the MLP classifier, it is almost certain that the model experienced significant overfitting – that is, it corresponds highly to its training data – and will not be generalizable to other subjects. Therefore, in order to apply this model for classification, it would be necessary to train the model for each subject, essentially calibrating it for each subject's unique EEG data. This calibration, although it would improve personalization of a BCI device for each subject, could increase costs toward production and implementation.

Additionally, it is important to note that because the LDA model was trained without any parameters and solely on the mean-centered data (to optimize efficiency), a higher classification accuracy may have been compromised.

Ultimately, the results from this paper provide substantial clinical significance, particularly in the field of BCIs. The superior performance of the MLP classifier suggests its capability to enhance the accuracy and reliability of BCI systems. This enhancement can lead to large improvements in quality of life for individuals with neurologically induced motor disabilities, who may rely on BCIs for accessibility and control over their life. Additionally, the acknowledgement of overfitting suggests a method of neurorehabilitation in which clinicians can develop unique models for each patient, with the developed framework, in order to generate the highest classification accuracy from patient to patient.

Conclusion

This study investigated the difference in classification capabilities between two classifiers – Linear Discriminant Analysis and a Multi-Layer Perceptron – as they pertain to motor imagery. The research found a significant difference in classification accuracy, with LDA achieving 64.15% and MLP achieving 99.63%. These results demonstrate that for the millions of paralysis patients around the globe, the implementation of an MLP could be utilized in prosthetic devices to allow for more precise user control. However, there are limitations to this, including the MLP's lack of generalizability, as well as the potential for temporal and spatial differences in EEG patterns to lead to unpredictable model behavior. Therefore, future research in the following three areas can prove to be highly beneficial.

1. Training Data: Investigating methods to gather a more diverse set of training data – one that accounts for individual differences in EEG patterns – could allow for a generalizable MLP model, while retaining many of the original features that permit high-accuracy classification.

2. Temporal Differences: Research into the variance of EEG patterns across a time span would allow for the minimization of model unpredictability when it comes to classification during differing levels of individual concentration.

3. Reinforcement Learning with Human Feedback (RLHF): A standardized model, one that is generalizable to all users, could be coupled with the implementation of RLHF, by which each user can further improve their own model's performance through rewarding algorithms. Additionally, the RLHF techniques could be integrated into the natural usage of the proposed model, to reduce the need for pre-calibration.

In conclusion, MLP is a highly promising technique for the classification of motor imagery EEG signals and opens up exciting new possibilities for the rehabilitation of paralysis patients. Although not yet generalizable, this study demonstrated that MLP holds significant potential, and is worthwhile for future research and investment.

References

1. Lebedev M. A., Nicolelis M. A. (2006). Brain–machine interfaces: past, present and future. TRENDS in Neurosciences, 29(9), 536-546. https://doi.org/10.1016/j.tins.2006.07.004

2. Yadav H, Maini S., Electroencephalogram based brain-computer interface: Applications, challenges, and opportunities. Multimed Tools Appl. 2023 May 4:1-45 https://doi.org/10.1007/s11042-023-15653-x

3. Pfurtscheller G., Neuper C. (2001). Motor imagery activates primary sensorimotor area in humans. Neuroscience Letters, 239(2-3), 65-68.] https://doi.org/10.1016/s0304-3940(97)00889-6

4. Daly J. J., Wolpaw J. R., Wolpaw, E. W. (2008). Brain-computer interfaces in neurological rehabilitation. The Lancet Neurology, 7(11), 1032-1043.]. https://doi.org/10.1016/S1474-4422(08)70223-0

5. Schirrmeister R. T., Springenberg J. T., Fiederer L. D. J., Glasstetter M., Eggensperger K., Tangermann M., et al. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. Human brain mapping, 38(11), 5391-5420. https://doi.org/10.1002/hbm.23730

6. Narayan, Yogendra. (2021). Motor Imagery EEG Signals Classificationusing SVM, MLP, and LDA Classifiers. Turkish Journal of Computer and Mathematics Education. 12(2), 3339-3344. https://turcomat.org/index.php/turkbilmat/article/view/2393/2102

7. Lun X., Jianwei L., Zhang Y., Hao Z., Hou Y. (2022). A Motor Imagery Signals Classification Method via the Difference of EEG Signals between Left and Right Hemispheric Electrodes. Frontiers Neuroscience. https://doi.org/10.3389/fnins.2022.865594

8. Cho H., Ahn M., Ahn S., Kwon M., Jun S. C. (2017). EEG datasets for motor imagery brain-computer interface. GigaScience,6(7),1–8. https://doi.org/10.1093/gigascience/gix034