


Classifying Rock Paper Scissors Hand Gestures from Optically Pumped Magnetometer Data

Karahan Yilmazer 

Department of Electrical and Computer Engineering, Technical University of Munich

 yilmazerkarahan@gmail.com

February 15, 2024

1 Introduction

For many years, electroencephalography (EEG) has been the leading recording device for brain-computer interfacing (BCI) applications [13]. EEG electrodes pick up global changes in the electrical field generated by the synchronous firing of pyramidal cells in the cortex. It is a well established recording system, with good temporal resolution which has been shown to be useful in many BCI settings [7], [9], [17]. However, a shortcoming of EEG is its low spatial resolution, which makes other recording systems like magnetoencephalography (MEG) that record the magnetic fields induced by electrical changes an attractive alternative.

So far, the most widely used MEGs were the superconducting quantum interference device MEGs (SQUID-MEGs) [12]. These are very bulky setups that restrict the user. However, the restriction in portability is not the main drawback, but rather the relatively large distance between the scalp and the recording electrodes. As the electrodes heat up quickly, they have to be actively cooled. Considering that the magnetic field decays with $\frac{1}{r^2}$ where r is the distance between the source and the recording electrode, bringing the electrodes closer to the sources would already provide a strong increase in signal quality.

This is where the relatively new optically pumped magnetometers MEGs (OPM-MEGs) come into play. These electrodes can again record magnetic fields, but the real benefit is that they can be placed very close to the scalp which should theoretically increase the signal quality and allow brain signal analyses that were not possible with EEG before. Since these electrodes are relatively new, it is important to test whether this assumption holds true.

As OPMs are assumed to have higher quality than EEG and they have been shown to be effective on simple motor imagery paradigms [6], testing them on a more complex paradigm seemed fitting. In comparison in EEG-BCIs, there have been approaches to classify even finer movements like individual fingers [8], however, the remaining majority focused on motor imagery of either the left or right hand [1]. This motivated the recording of two datasets: one where the participant performed single finger movements and one where he performed rock, paper, scissors movements, both with his dominant right hand. It is to be noted that both paradigms were tested on SQUID-MEG before [2], [16] serving as a baseline for the findings of the current study. However in the paper classifying rock, paper, scissors gestures a deep learning approach was used, which makes the model hard to interpret and thus use in the clinic.

The focus of the current study is on analyzing the second dataset. It aims to classify right hand gestures of rock, paper, scissors using explainable feature extraction methods and machine learning algorithms. Compared to typical EEG-BCI applications, which classify left vs. right motor imagery, this project presents two main challenges. First, it involves three classes (rock, paper, scissors) instead of two (left or right hand). Second, all signals originate from the same hemisphere, making spatial separation more difficult.

To overcome these issues and better understand the neural mechanisms underlying higher-level motor movements, various signal analysis methods were employed. These include analyses of motor imagery/execution involving band power features and common spatial patterns (CSP). They also include time-frequency analyses, which cover the gamma frequency range, thanks to MEG's resilience to high-frequency muscle signals. And lastly source space analysis which was made possible with the high spatial resolution of the OPMs. Subsequently, this paper will introduce the exploratory data analysis conducted on the recorded dataset.

2 Methods & Results

2.1 Dataset

The dataset comprised of four distinct experimental runs, capturing data from a total of 65 electrodes. Each electrode measured three-dimensional magnetic fields, recording along the x, y, and z axes, resulting in 195 unique data channels. With a sampling rate of 1200 Hz, the setup was capable of capturing signal components up to 600 Hz, in line with the Nyquist theorem. Each individual run spanned approximately 10 minutes, cumulating in a comprehensive dataset of 45 minutes when combined.

Complementing the OPM data, magnetic resonance imaging (MRI) scans of the participant were obtained to assist with source localization. These scans were not only crucial for identifying signal origins but also served to customize a 3D-printed helmet tailored to the participant's head seen in Figure 9. This helmet, equipped with numerous recording sites, facilitated a versatile arrangement of the OPM sensors. For this experiment, half of the electrodes were strategically placed around the left motor cortex to capture activity corresponding to the right-hand gestures. The remaining sensors were distributed across the scalp, ensuring extensive brain coverage. This electrode configuration is depicted in Figure 10.

2.2 Preprocessing

With an abundance of channels and trials, the preprocessing for further analysis was kept to a bare minimum. The goal was to have the cleanest epoched data possible after the preprocessing.

2.2.1 Bad Channel Removal

The first step of the preprocessing was the removal of bad channels. Bad channels were visually detected in two ways: once through the power spectral density (PSD) plots and then through filtered time series data inspection.

For this purpose, the PSD plots were plotted for each axis separately. These plots revealed that some channels' power values deviated strongly from others. All of these channels were marked as bad.

Afterwards, a band-pass filter between 1 and 400 Hz followed by homogeneous field compensation (HFC) [19] was applied to the data. These filtering steps took care of the strongest artifacts. The filtered data were visually inspected through time series plots to detect any additional bad channels.

It is worth noting that once a bad channel was identified, it was dropped from all four runs. This simplified further analysis by making the subset of available channels uniform along all runs. With that said, 55 out of 195 channels were marked as bad and dropped from all runs.

2.2.2 Epoching

After dropping the bad channels a band-pass filter between 7-30 Hz was applied before cutting the continuous data into epochs. However this filtering was not applied for the time-frequency analysis, as higher frequencies were also of interest. This will be motivated and explained in more detail in the Time-Frequency Analysis section.

Afterwards, continuous data was cut into 60 epochs for each hand gesture in each run, summing up to 180 epochs for each run and 720 in total. The epochs were cut starting from 500 ms before the cue presentation until 2100 ms after. Then, a linear detrending was applied to the epochs. Later, bad epochs were rejected by visual inspection of time courses and variances. Just like the bad channels, the rejection threshold for bad epochs was quite low to ensure good data quality. After this step 122 out of 720 trials were rejected. With 198 rock, 202 paper and 198 scissors trials, the dataset was still balanced and was not heavily biased towards one or multiple classes.

2.3 Analysis

The analysis and classification of the preprocessed data was done once in sensor space and once in source space. For the sensor space analysis, event-related potentials and time-frequency plots for each condition and axis

pair were extracted. Then, maximally different time and frequency ranges were distinguished through visual inspection. These ranges were then used to compute common spatial patterns features that were classified with a support vector machine (SVM) classifier.

On the source space side, having the MRI scans of the participant allowed a source reconstruction using the sLORETA algorithm [14]. The estimated source time course activities were then plotted on a 3D model of the participant's brain which served as a feasibility test for the source reconstruction. Lastly, the inverse operator was applied to the preprocessed epochs to get another set of source estimates for each condition. These were then decoded using a sliding logistic regression estimator.

2.3.1 Time-Frequency Analysis

A common way of analyzing time series M/EEG data is to use the event-related potentials (ERP). These are time courses averaged over trials. They reflect the phase-locked parts of the signal and can be very useful for discriminating different classes, like for example the existence or absence of an observed error [5]. ERPs provide great insights on the average signal change over time. However, tiny phase differences arising from time lags in recording or cue presentation can totally disrupt ERPs [3].

An example of such ERPs can be seen in Figure 1. Here, the ERPs were plotted for each condition and axis pair separately. Unfortunately, ERP patterns of all three conditions were overlapping in all three axes, which made ERPs a bad candidate for being used as features. However, all ERPs show some consistent activity within the range of 0-1 s. Since the data was extensively cleaned, the chances of this ERP feature corresponding to artifacts, such as muscle activity, was relatively low, however still possible. In any case, this time window showed promise of being correlated with the actual task at hand.

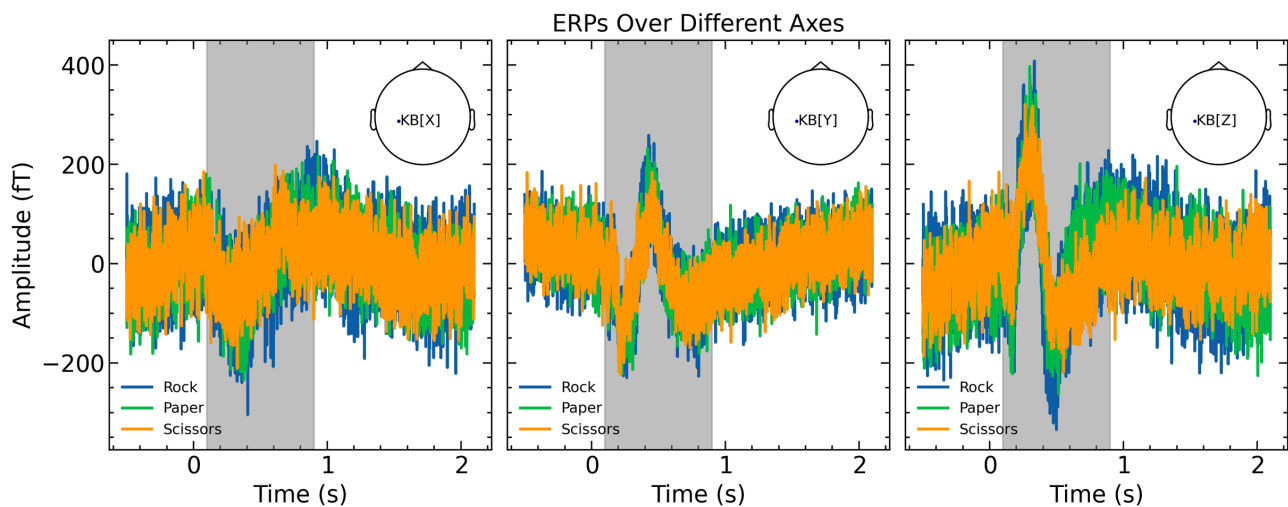


Figure 1 Event-related potentials for all conditions along all axes over a channel in the motor cortex. There seems to be some consistent activity in the time window of 0-1 s. However, the gestures cannot be separated from ERP features alone.

As it can be seen, ERPs do not capture frequency information. The common way of looking at frequency characteristics of the data is to extract the Fourier spectrum or PSD. This is a good approach if there are strong frequency components present over longer periods in the signal. However, if these frequency components are transient and of small amplitude, they might get lost in broadband noise.

Here is where time-frequency analyses shine brightest. By plotting spectrograms of data, insights into the frequency changes over time can be gained. This way, even transient activities can be visualized and averaging over trials becomes less prone to small phase shifts.

Hoping to find any differences between different gestures, time-frequency plots of epoched data were plotted for each axis. For these plots, the unfiltered but cleaned (no bad channels or epochs) data was used. The initial idea was to check whether gamma band shows any discriminable activity.

The top row in Figure 2 shows that the frequency power in decibels phase-locked, non-phase locked and total activity. The phase locked activity corresponds to the ERP, and a very faint activity can be seen at the lowest

frequency, which is much more evident in Figure 3. The non-phase locked activity was not captured by the ERP and thus any information gained from this time-frequency matrix would be new to the analysis.

On the other hand, the bottom row represents the inter-trial phase locking/coherence (ITPC), which is a measure of uniformity of phase angles across trials [20]. However these values were not analyzed any further as there seemed to be no information of interest hidden here.

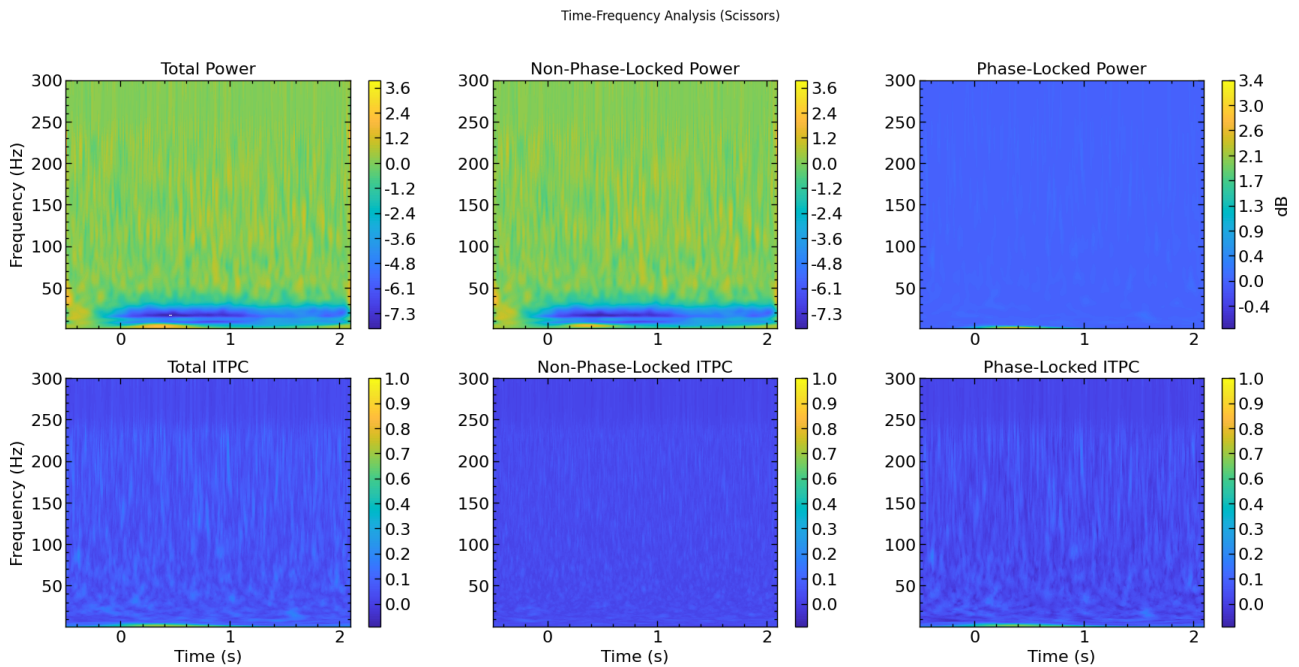


Figure 2 Time-frequency power and inter-trial phase locking (ITPC) of the scissors gesture along the z-axis. The interesting activity seems to be happening in the power of lower frequencies.

Both the total power or the non-phase-locked power showed that there was no significant gamma activity cluster and that the values in the gamma range were roughly the same for the entirety of the epochs. The cue onset did not significantly change these values and there was not phase locking as well. However, the lower frequencies showed a drop in power after cue onset, visualized with the dark blue shade. This reflects the well known phenomenon of event-related desynchronization (ERD) [10], [15]. Since the lower frequency bands offered meaningful insights, zooming into that region was the next logical step.




Figure 3 Time-frequency power plots of the z-axis for all conditions. All the plots look very similar but going from the left plot to the right reveals a larger decrease in power for 20 Hz in the range of 0.75-2 s.

Figure 3 shows that all the time-frequency plots look very similar for all conditions. However there are slight differences between the gestures within the frequency window of 7-30 Hz and time window of 0.75-2 s, especially for the scissors condition. The best way to test this hypothesis was to use band-pass filter and crop the data within this range and use it as input to the common spatial patterns.

That is exactly what was done as the next step. A least-squares linear-phase FIR band-pass filter was designed with cutoff frequencies of 7 and 30 Hz, which can be seen in Figure 13. After filtering the continuous data with this filter, it was cut into epochs again and using the indices of bad epochs from the manual rejection, this time the epochs were dropped automatically. Furthermore, the epochs were further cropped to the range of 0.75-2 s. Afterwards, MNE's CSP implementation was used to get the CSP features. However, this step occasionally threw rank deficiency errors. To solve this issue, the epoched data was transformed using principal component analysis (PCA), keeping the components explaining 99% of the variance. After PCA and CSP, lastly, SVM was used as a classifier to complete the pipeline.

Different time windows and subsets of channels were experimented with this setup. However, the maximum classification accuracy achieved was 44%. Which is not so surprising, seeing that the time-frequency plots are quite similar to each other. It is worth noting that Figure 3 show the plots for the z-axis, however the other axes looked and performed quite similarly. Not being able to discriminate the gestures in the sensor space, further analyses were conducted in the source space.

2.3.2 Source Localization

Moving from sensor space to source space requires the MRI scans of the brain of interest and fortunately, these scans were readily available for the current participant. To reconstruct the sources first different tissues were segmented using the open-source software FreeSurfer using the command `recon-all`. These segmentations were then used to create boundary element model (BEM) surfaces using the MNE function `mne.bem.make_watershed_bem()`. These surfaces are triangulations of the interfaces between different tissues needed for forward computation. Normally for EEG, three different BEM layers are necessary: inner skull, outer skull and outer skin. For MEG however, one layer of inner skull is enough. Since MNE creates all three layers automatically, all of them were used for the forward model computation.

The problem faced here was that the Watershed BEM creation using MNE did not produce good results as it can be seen in Figure 4. As an alternative, another open-source software Brainstorm [18] was used to create the same BEM surfaces. As it can be seen in Figure 5, these BEMs looked much better, capturing the correct anatomical structure. However, using these surfaces for the forward computation in MNE proved to be a major challenge due to data incompatibility.

MNE requires the BEM surfaces to be of file type `.surf`. Brainstorm can export the BEM surfaces into various data formats but `.surf` is not one of them. FreeSurfer can be used to convert some of these formats to `.surf` files, but the vertex ordering information gets lost during this conversion. This results in MNE not being able to use these surfaces for the forward model computation. More experimentation can provide a workaround to this problem but until then the Watershed BEMs of MNE were used for the following steps.

Then the surface-based source space was computed using the function `mne.setup_source_space()` with the recommended spacing of "oct6". This computation restricts the sources to be on the surface. An alternative way of doing this is to use the volume-based source space using `mne.setup_volume_source_space()`. Comparing these two methods might reveal further insights, however, for the current analysis only the surface-based source space was computed. The results of these sources on the 3D brain model can be seen in Figure 17.

Lastly, the forward operator was computed with `mne.make_forward_solution()`, which is commonly referred to as the gain or leadfield matrix. The forward operator was then used to compute inverse operator using `mne.make_inverse_operator()` which was then used to estimate the sources, commonly referred to as source time courses (STC). For this last step `mne.minimum_norm.apply_inverse()` was used with method set to sLORETA. This step involves projecting the MEG sensor data from the sensor space onto the source space, which is positioned on the brain. These source activations within the brain are assumed to generate the signals recorded by the sensors.

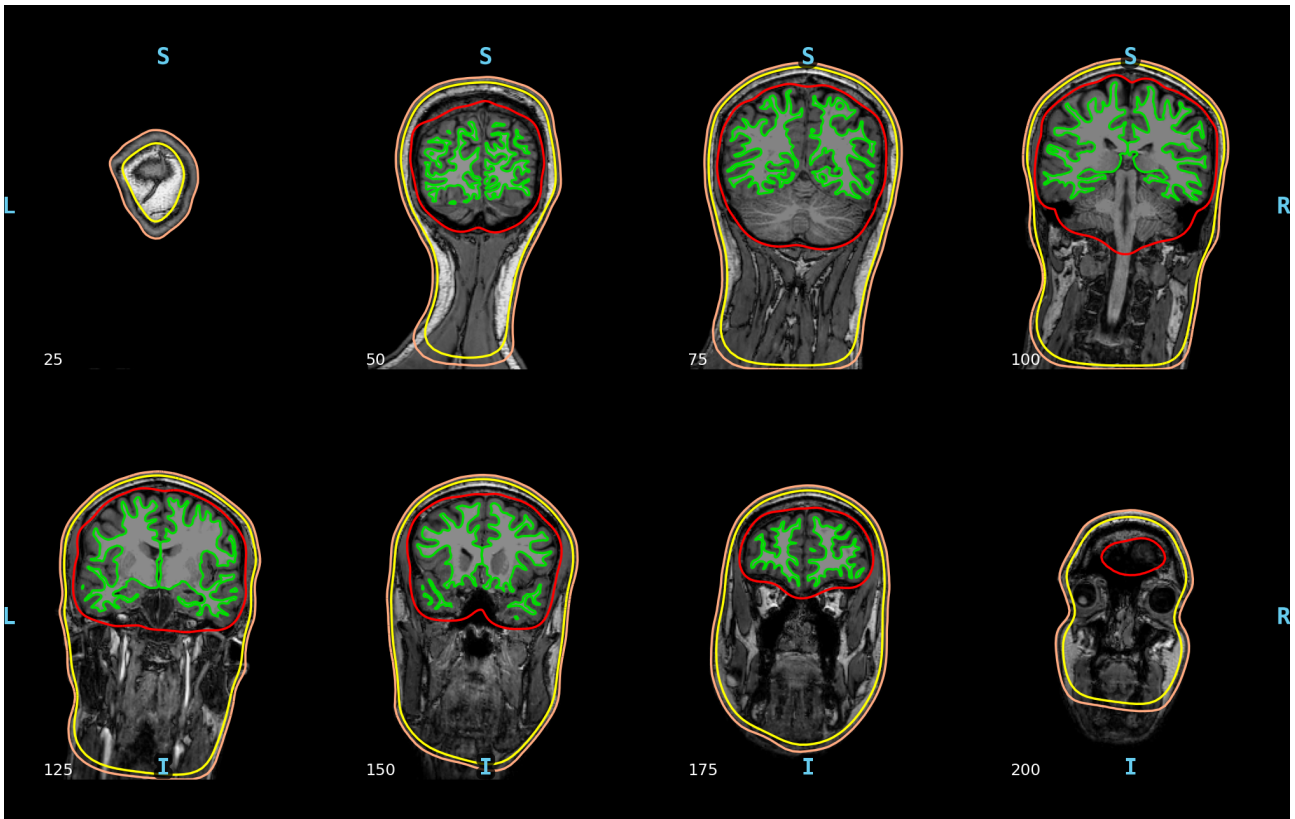


Figure 4 The BEM parcellation using the Watershed algorithm of MNE. This method fails to detect the layers correctly.

It is expected that the STC activations reflect the brain activations during the task. In other words, we would expect to see a strong activation over the motor cortex. Looking at Figure 6, it can be seen that the peak of the STC activations indicated by the blue marker can indeed be seen very close to the motor cortex. It falls more over the premotor cortex but this might be due to how the analysis was done. Furthermore, the STC in the same figure reflects the expected time course of motor activation. Following the cue at time 0s there is an activation at the motor cortex, which then decays over time. This time range also corresponds to the ERPs seen in Figure 1 and the ERD pattern in time-frequency plots of Figure 3. All together, the results indicate that the motor execution task could correctly be picked up by the OPM sensors.

After all source space objects had been calculated and saved to disk, they were then utilized to decode the gestures. To achieve this, the function `mne.minimum_norm_apply_inverse_epochs()` was employed. This function accepts any epochs object as input. However, to isolate the motor component, the epochs object prefiltered within the 7-30 Hz range was specifically used. The function calculates the estimated source activity that would produce the observed activity on the sensors, assuming the source activity is projected onto these sensors.

Unfortunately, the source estimates returned by the function take up too much memory. Which is why the epochs object had to be downsampled before passing to the function. However, once the estimated source time courses were computed, these could be used as features for the classification algorithm. Since the features were simply time series data of the estimated sources, a sliding window estimator was used to make a prediction for small overlapping time windows. The results of this estimation can be seen in Figure 7.

For the sliding window, the pipeline comprised of a standard scaler, a feature selection algorithm choosing the best 500 features based on the ANOVA f-test, followed by a logistic regression model. A peak test set accuracy of 48.31% could be achieved, which is marginally above the chance level of 33%. Again, the time window of 0-1 s showed an increase in decoding accuracy, further suggesting the underlying motor activity picked up by the sensors.




Figure 5 The BEM parcellation using the algorithm from Brainstorm. This method works very well, however the surfaces could not be read into MNE for further analysis.

2.3.3 Further Classification Approaches

On top of the time-frequency and source space analyses more conventional analyses methods were also used and to see if brute-force analysis with extensive hyperparameter search could provide better classification results. In this approach, a lot of features were extracted from the data and put into a feature matrix with very little neurophysiological guidance. Then different machine learning pipelines were used to see if any of the feature-pipeline combinations could achieve high accuracies. Unfortunately, none of the tested methods could achieve a classification significantly above chance level (33% for three classes). Still, in the following different approaches that were tested will be presented.

As mentioned, many features were extracted from the data using different approaches. First of them was the extraction of the whole PSD spectra from different channels. Then more selective features were extracted such as the relative band power and spectral peak positions and amplitudes from different frequency bands, including the gamma range. Lastly, the average signal variance from each trial was extracted.

All of these features were extracted for each trial and each electrode. This resulted in high dimensional feature matrices of course. If we take the band power features as an example, we can look at how these highdimensional feature matrices come into shape. The band power features were extracted for seven frequency bands: delta, theta, alpha, beta and gamma. But gamma was split into three subbands: low gamma from 30-100 Hz, middle gamma from 100-200 Hz and high gamma from 200-300 Hz. These seven band power values for extracted for each of the 168 channels, which resulted in 980 features for only 598 trials in the training set, which is a setting prone to overfitting.

One way to overcome this issue was to manually pick only the electrodes of interest, such as the ones over the motor cortex. However, so far only automatic approaches were tested. These automatic approaches were Principal Component Analysis (PCA) and supervised and unsupervised Uniform Manifold Approximation and Projection (UMAP) and linear discriminant analysis (LDA) as dimensionality reduction methods and Fisher Score for feature selection. All four methods provided feature matrices with maximum 10 features (number of features chosen based on 5-fold cross validation). Then the projected/selected features were used for

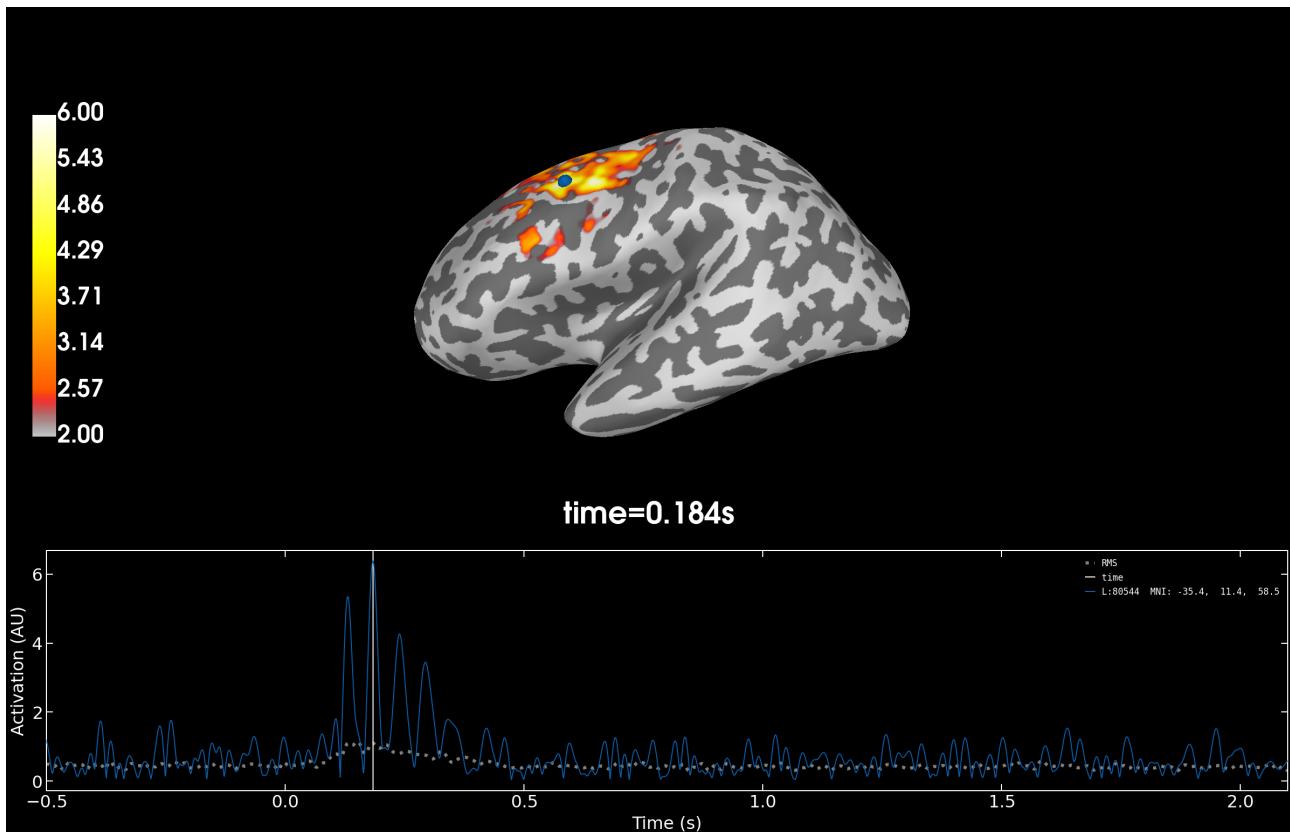


Figure 6 The STC activations computed from the inverse model. The blue marker shows the region with largest activation which falls over the premotor cortex. The time series data corresponds to average source activation of the trials.

classification using support vector machines (SVM). As mentioned before none of these combinations yielded classification accuracies significantly above chance level or accuracy levels achieved by time-frequency or source space analysis. However, one last thing that was tested was whether hyperparameter optimization could boost this performance.

For this purpose, a combination of supervised UMAP and SVM was used. A random search of hyperparameters was conducted using the Weights & Biases library. However as it can be seen from Figure 18, the hyperparameter search could also not provide a significant boost to the classifier performance.

3 Discussion

The task at hand has proven to be more challenging than an a normal motor imagery dataset analysis. The exploratory nature of the data analysis opens up a lot of possibilities for classifying different hand gestures and the results, especially those from the source localization, suggest that there is meaningful information hidden in the data. The combination of the correct methods could reveal how to best classify different hand gestures from each other. However, in the current work this could only be possible with a maximum classification accuracy of 48.31%.

There were many approaches tested on the dataset. Figure 8 shows all of these approaches and their respective performance scores to provide an overview. As it can be seen, the source space analysis performed the best out of all three approaches. However, it did not come without its limitations. First of all, a crucial step would be to find a way to use the Brainstorm BEMs in MNE. Hopefully, these BEMs will provide even more accurate forward and inverse operator computations, which might provide better discriminability in the source space.

Furthermore, decoding in source space proved to be a challenge as the computer used in this study ran out of RAM while running the script. To solve this issue, the epoched data had to be downsampled to every seventh sample. This means that six out of every seven samples were discarded which reduces the time resolution and

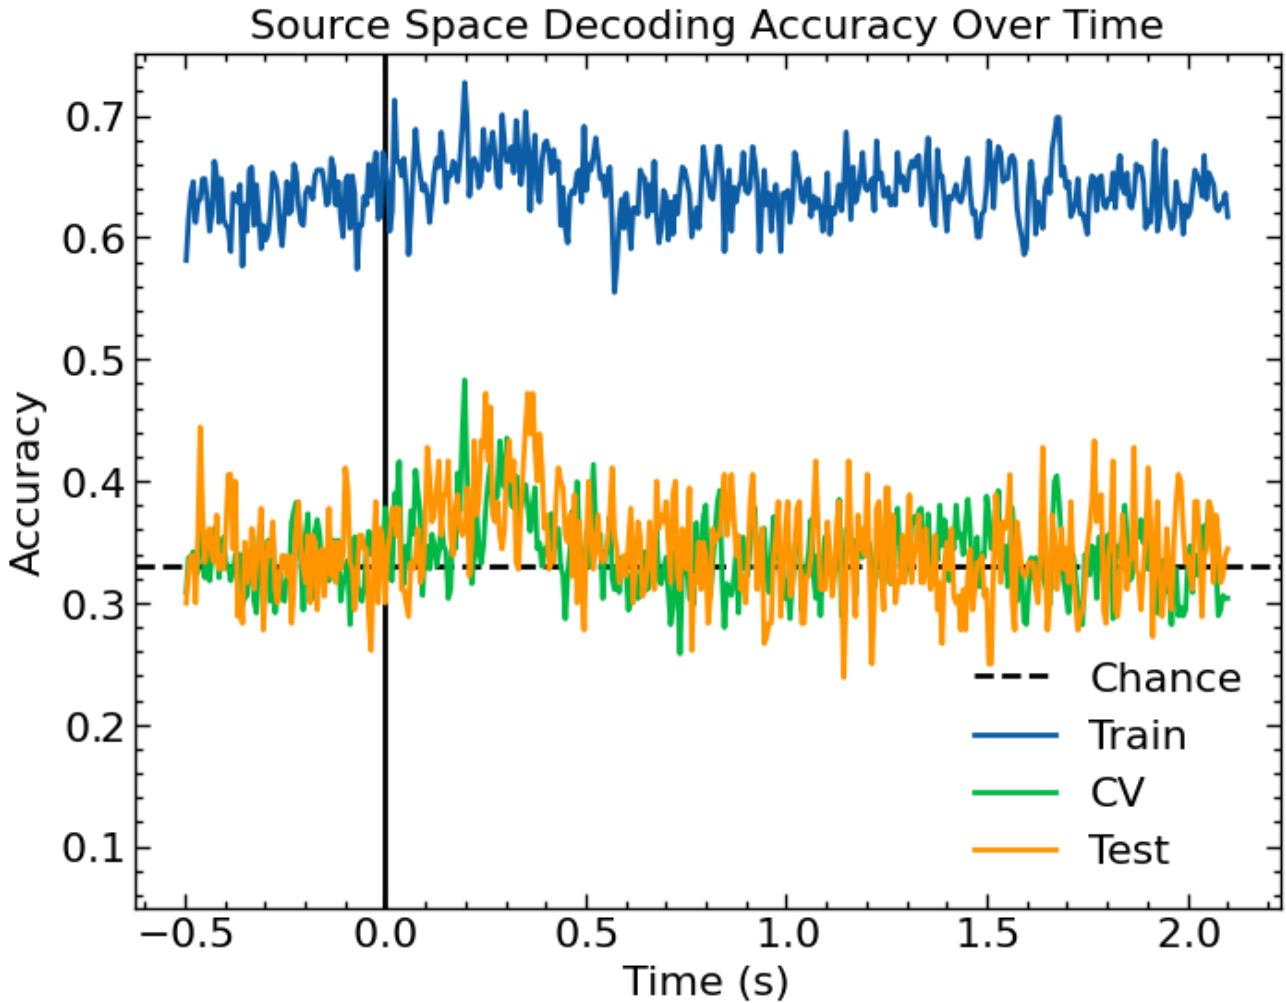


Figure 7 The STC activations computed from the inverse model. The blue marker shows the region with largest activation which falls over the premotor cortex. The time series data corresponds to average source activation of the trials.

might result in lower classification accuracies. Running the same source-space analysis with no downsampling could already provide a performance boost.

The time-frequency analysis can also be further studied. Although Figures 14, 15 and 16 show that all classes were nearly identical to the human eye, convolutional neural networks (CNNs) could reveal further insights, especially in the higher frequency ranges. CNNs are mostly used for image classification tasks and the time-frequency spectrograms are basically 2D images, so it would be very fitting to use CNNs for this purpose. This could reduce the explainability of the results but could increase the performance. One approach could be to train three different networks, one on lower frequencies, one on higher frequencies and one over the whole spectrum. Even if it would be difficult to pinpoint which channel-frequency-time points result in the best predictions, this approach could reveal which frequency bands are best for explaining higher-level motor tasks.

However, it has to be noted that in the mentioned previous work study studying finger movements with MEG, a detailed examination of the effect of frequency bands on the classification accuracy was performed [16]. They showed that the 6-11 Hz frequency band resulted in the highest decoding accuracy. But OPMs were not used in this study and the analysis included frequencies up until 100 Hz. Looking at higher frequencies with OPMs and training more sophisticated classifiers like the suggested CNNs could show different results.

When it comes to the overall interpretation of accuracies below the desired 60-70%, the easiest explanation would be to blame the data. As it can be seen in Figure 11, this has some truth to it, as the data from first run seems to be corrupted by an interesting artifact. On other runs, there were multiple spectral peaks, not only at the 50 Hz line noise and its harmonics, but all over the PSD spectrum, as it can be seen in Figure 12. These

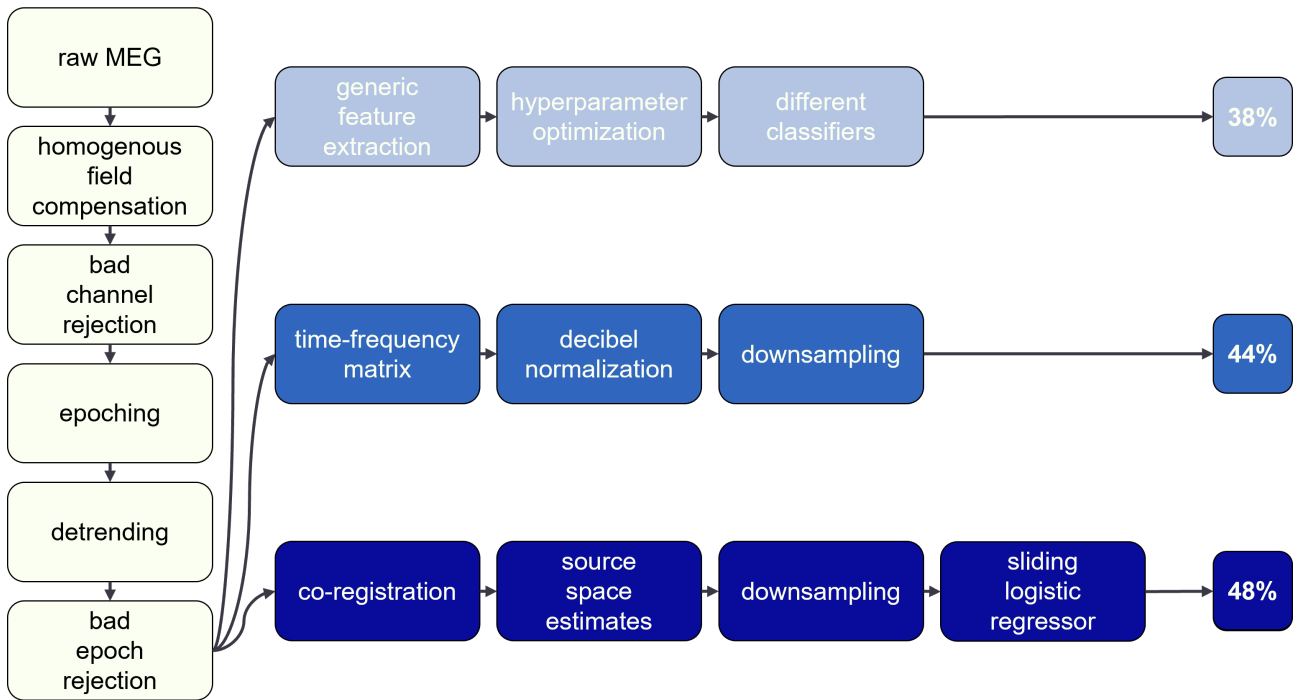


Figure 8 Preprocessing and analysis pipelines used in the current study. Three different analyses approaches were tested and the source space analysis performed the best out of all three with a classification accuracy marginally higher than chance level of 33%.

artifacts could be explained by the OPMs used in the study being highly experimental. These artifacts were partially removed with preprocessing steps like homogeneous field compensation or spectral filtering, however they still affect the data quality.

However, the machine learning aspect of the current work can definitely be reinforced with meticulous investigation of different approaches. Some possible directions for this investigation could be ensemble learning with multiple classifiers, using subsets of electrodes for feature extraction instead of relying on automatic dimensionality reduction methods, extract common event-related desynchronization or CSP features, or to use pretrained deep learning models like EEGNET [11] for performing feature extraction and classification automatically if the goal of explainable machine AI is to be abandoned.

Lastly, the feasibility of OPMs for BCI applications can be discussed. Even if it was shown that the OPMs can achieve better spatial resolutions than EEG, its main drawback is that the recordings have to be done in a magnetically shielded room as the measured magnetic fields are very small compared to earth’s magnetic field. Without solving this issue, OPMs seem to be lacking real-life applications and remaining an experimental tool.

Furthermore, one of the main benefits of using OPMs is the excellent coregistration thanks to the MRI-fitted helmets. Although this helmet is not necessary for simpler applications not involving source space analyses, designing and building these helmets for each participant is also a major challenge which might reduce participant recruitment and thus variability in datasets.

Still, OPMs prove to be a good experimental tool and proving that it can detect signals which EEG cannot might be the way to move forward. Considering that nearly 15-30% of all participants are BCI-illiterate [4], meaning that they cannot produce brain signals useful for BCI applications, training these participants, for example through reinforcement learning using neurofeedback, proves to be an important front. This process can be accelerated with the use of OPMs. When a participant proves to be BCI-illiterate, OPMs can be used to pick up brain signals more reliably than EEG to train the participant using neurofeedback.

Consider the transition from OPMs to EEG-based BCI as akin to learning to ride a bicycle equipped with training wheels. Initially, the bicycle comes equipped with supporting wheels, ensuring high-quality data recordings and thus better neurofeedback, which guide and support the participant during the initial training

phase. On top of that, these training sessions would not require the application of gel, which is another major advantage.

As proficiency increases and the user becomes more accustomed to producing more distinguishable brain signals, akin to a cyclist gaining balance and confidence, the transition to EEG could begin. By removing the training wheels the participants who were initially BCI-illiterate would hopefully be more proficient BCI users. Being able to use EEG-BCIs would be more desirable, as EEGs can be cheaper and more readily available. More importantly, they do not face the issue of having to be in a magnetically shielded room.

4 Conclusion

The results of the current study are encouraging for future research. With a maximum classification accuracy of 48.31% the results were not as high as 60-70%, yet still significantly higher than chance level of 33%. The exploratory data analysis of the current study covered both sensor and source space analysis methods. This allowed for a comparison of different approaches, yet left room for further investigation. Optimizing the current analyses could provide better classifications and thus better neurophysiological insights. Achieving this would pave the way to use OPMs as a "training ground" for BCI illiterate participants before they move to EEG-BCI applications. The potential for successfully analyzing high-level motor movements using OPMs remains high, warranting continued investigation and experimentation in this field.

References

- [1] P. Arpaia, A. Esposito, A. Natalizio, and M. Parvis, "How to successfully classify EEG in motor imagery BCI: A metrological analysis of the state of the art," *Journal of Neural Engineering*, vol. 19, no. 3, p. 031 002, Jun. 1, 2022, ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2552/ac74e0. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1741-2552/ac74e0> (visited on 07/18/2022).
- [2] Y. Bu, D. L. Harrington, R. R. Lee, *et al.*, "Magnetoencephalogram-based brain-computer interface for hand-gesture decoding using deep learning," *Cerebral Cortex*, vol. 33, no. 14, pp. 8942–8955, Jul. 5, 2023, ISSN: 1047-3211, 1460-2199. DOI: 10.1093/cercor/bhad173. [Online]. Available: <https://academic.oup.com/cercor/article/33/14/8942/7161766> (visited on 08/16/2023).
- [3] M. X. Cohen, *Analyzing Neural Time Series Data: Theory and Practice* (Issues in Clinical and Cognitive Neuropsychology). Cambridge, Massachusetts: The MIT Press, 2014, 578 pp., ISBN: 978-0-262-01987-3.
- [4] T. Dickhaus, C. Sannelli, K.-R. Müller, G. Curio, and B. Blankertz, "Predicting BCI performance to study BCI illiteracy," *BMC Neuroscience*, vol. 10, no. 1, P84, Jul. 13, 2009, ISSN: 1471-2202. DOI: 10.1186/1471-2202-10-S1-P84. [Online]. Available: <https://doi.org/10.1186/1471-2202-10-S1-P84> (visited on 03/06/2024).
- [5] G. Dornhege, J. del R. Millán, T. Hinterberger, D. J. McFarland, and K.-R. Müller, "Error-Related EEG Potentials in Brain-Computer Interfaces," in *Toward Brain-Computer Interfacing*, MIT Press, 2007, pp. 291–301, ISBN: 978-0-262-25604-9. [Online]. Available: <https://ieeexplore.ieee.org/document/6281216> (visited on 06/05/2022).
- [6] N. Fedosov, O. Shevtsov, and A. Ossadtchi, "Motor-Imagery BCI with Low-Count of Optically Pumped Magnetometers," in *2021 Third International Conference Neurotechnologies and Neurointerfaces (CNN)*, Sep. 2021, pp. 16–18. DOI: 10.1109/CNN53494.2021.9580251. [Online]. Available: <https://ieeexplore.ieee.org/document/9580251> (visited on 03/06/2024).
- [7] Á. Fernández-Rodríguez, F. Velasco-Álvarez, and R. Ron-Angevin, "Review of real brain-controlled wheelchairs," *Journal of Neural Engineering*, vol. 13, no. 6, p. 061 001, Oct. 2016, ISSN: 1741-2552. DOI: 10.1088/1741-2560/13/6/061001. [Online]. Available: <https://dx.doi.org/10.1088/1741-2560/13/6/061001> (visited on 07/18/2023).

- [8] S. Gannouni, K. Belwafi, H. Aboalsamh, *et al.*, “EEG-Based BCI System to Detect Fingers Movements,” *Brain Sciences*, vol. 10, no. 12, p. 965, Dec. 10, 2020, ISSN: 2076-3425. DOI: 10.3390/brainsci10120965. PMID: 33321915. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7763179/> (visited on 03/06/2024).
- [9] F. P. Kalaganis, E. Chatzilari, S. Nikolopoulos, I. Kompatsiaris, and N. A. Laskaris, “An error-aware gaze-based keyboard by means of a hybrid BCI system,” *Scientific Reports*, vol. 8, no. 1, p. 13 176, 1 Sep. 4, 2018, ISSN: 2045-2322. DOI: 10.1038/s41598-018-31425-2. [Online]. Available: <https://www.nature.com/articles/s41598-018-31425-2> (visited on 08/03/2023).
- [10] J. Kalcher and G. Pfurtscheller, “Discrimination between phase-locked and non-phase-locked event-related EEG activity,” *Electroencephalography and Clinical Neurophysiology*, vol. 94, no. 5, pp. 381–384, May 1995, ISSN: 00134694. DOI: 10.1016/0013-4694(95)00040-6. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/0013469495000406> (visited on 01/04/2022).
- [11] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, “EEGNet: A Compact Convolutional Network for EEG-based Brain-Computer Interfaces,” *Journal of Neural Engineering*, vol. 15, no. 5, p. 056 013, Oct. 1, 2018, ISSN: 1741-2560, 1741-2552. DOI: 10.1088/1741-2552/aace8c. arXiv: 1611.08024 [cs, q-bio, stat]. [Online]. Available: <http://arxiv.org/abs/1611.08024> (visited on 11/21/2023).
- [12] P. K. Mandal, A. Banerjee, M. Tripathi, and A. Sharma, “A Comprehensive Review of Magnetoencephalography (MEG) Studies for Brain Functionality in Healthy Aging and Alzheimer’s Disease (AD),” *Frontiers in Computational Neuroscience*, vol. 12, 2018. DOI: 10.3389/fncom.2018.00060. PMID: 30190674. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6115612/> (visited on 11/21/2023).
- [13] T. Nazneen, I. B. Islam, M. S. R. Sajal, *et al.*, “Recent Trends in Non-invasive Neural Recording Based Brain-to-Brain Synchrony Analysis on Multidisciplinary Human Interactions for Understanding Brain Dynamics: A Systematic Review,” *Frontiers in Computational Neuroscience*, vol. 16, 2022, ISSN: 1662-5188. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fncom.2022.875282> (visited on 11/21/2023).
- [14] R. D. Pascual-Marqui, “Standardized low-resolution brain electromagnetic tomography (sLORETA): Technical details,” *Methods and Findings in Experimental and Clinical Pharmacology*, vol. 24 Suppl D, pp. 5–12, 2002, ISSN: 0379-0355. PMID: 12575463.
- [15] G. Pfurtscheller and A. Aranibar, “Evaluation of event-related desynchronization (ERD) preceding and following voluntary self-paced movement,” *Electroencephalography and Clinical Neurophysiology*, vol. 46, no. 2, pp. 138–146, Feb. 1979, ISSN: 0013-4694. DOI: 10.1016/0013-4694(79)90063-4. PMID: 86421.
- [16] F. Quandt, C. Reichert, H. Hinrichs, H. J. Heinze, R. T. Knight, and J. W. Rieger, “Single trial discrimination of individual finger movements on one hand: A combined MEG and EEG study,” *NeuroImage*, vol. 59, no. 4, pp. 3316–3324, Feb. 15, 2012, ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2011.11.053. PMID: 22155040.
- [17] S. R. Soekadar, M. Witkowski, N. Vitiello, and N. Birbaumer, “An EEG/EOG-based hybrid brain-neural computer interaction (BNCI) system to control an exoskeleton for the paralyzed hand,” *Biomedical Engineering / Biomedizinische Technik*, vol. 60, no. 3, Jan. 1, 2015, ISSN: 1862-278X, 0013-5585. DOI: 10.1515/bmt-2014-0126. [Online]. Available: <https://www.degruyter.com/document/doi/10.1515/bmt-2014-0126/html> (visited on 12/26/2021).
- [18] F. Tadel, S. Baillet, J. C. Mosher, D. Pantazis, and R. M. Leahy, “Brainstorm: A user-friendly application for MEG/EEG analysis,” *Computational Intelligence and Neuroscience*, vol. 2011, p. 879 716, 2011, ISSN: 1687-5273. DOI: 10.1155/2011/879716. PMID: 21584256.

- [19] T. M. Tierney, N. Alexander, S. Mellor, *et al.*, “Modelling optically pumped magnetometer interference in MEG as a spatially homogeneous magnetic field,” *NeuroImage*, vol. 244, p. 118 484, Dec. 1, 2021, ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2021.118484. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1053811921007576> (visited on 11/17/2023).
- [20] R. M. van Diepen and A. Mazaheri, “The Caveats of observing Inter-Trial Phase-Coherence in Cognitive Neuroscience,” *Scientific Reports*, vol. 8, no. 1, p. 2990, 1 Feb. 14, 2018, ISSN: 2045-2322. DOI: 10.1038/s41598-018-20423-z. [Online]. Available: <https://www.nature.com/articles/s41598-018-20423-z> (visited on 02/16/2024).

A Figures



Figure 9 The MRI-fitted helmet used in the experiment. The helmet has many slots, giving the experimenter the freedom to choose where to place the electrodes.

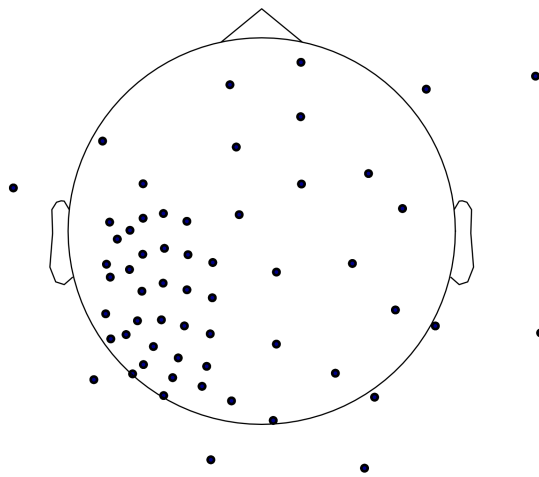


Figure 10 The electrode configuration used in the study. Half of the electrodes are focused on the right motor cortex for better coverage.




Figure 11 The PSD plot of run 1 after bad channels were dropped. There seems to be an unusual wave-like artifact in the PSD spectrum corrupting the data. Even after preprocessing, this artifact could not be removed.

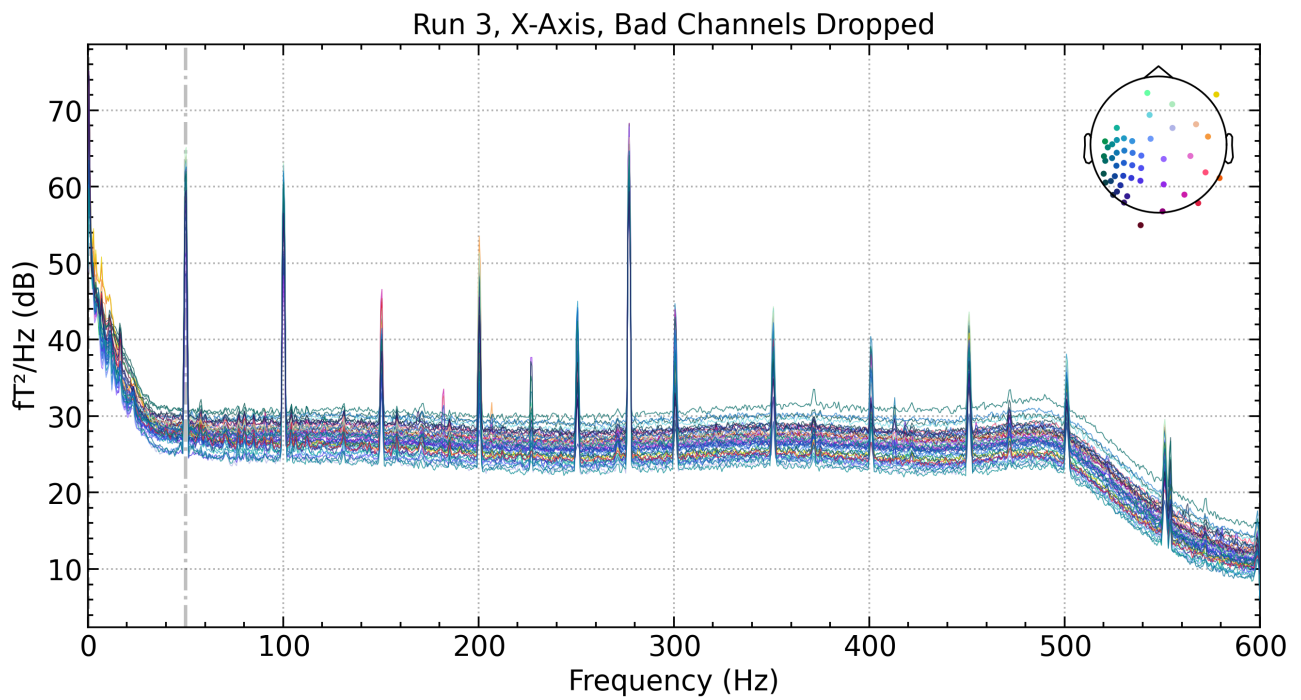


Figure 12 The PSD plot of run 3 after bad channels were dropped. There seems to be a lot of spectral peaks, not only at the 50 Hz line noise and its harmonics, but at other frequencies as well. The exact positions of these peaks shifted from one run to the other, making the standardization of the preprocessing difficult.

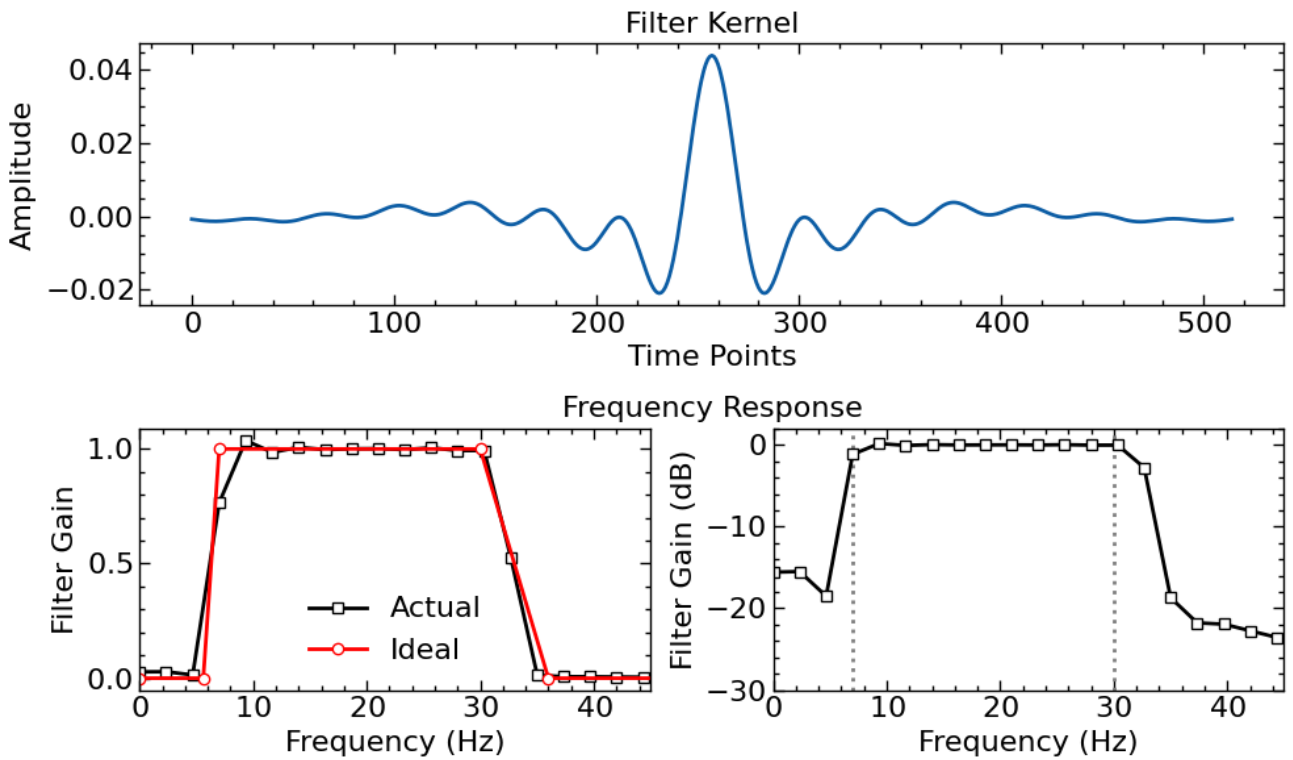


Figure 13 Least-squares linear-phase FIR band-pass filter used for filtering before extracting the CSP features.

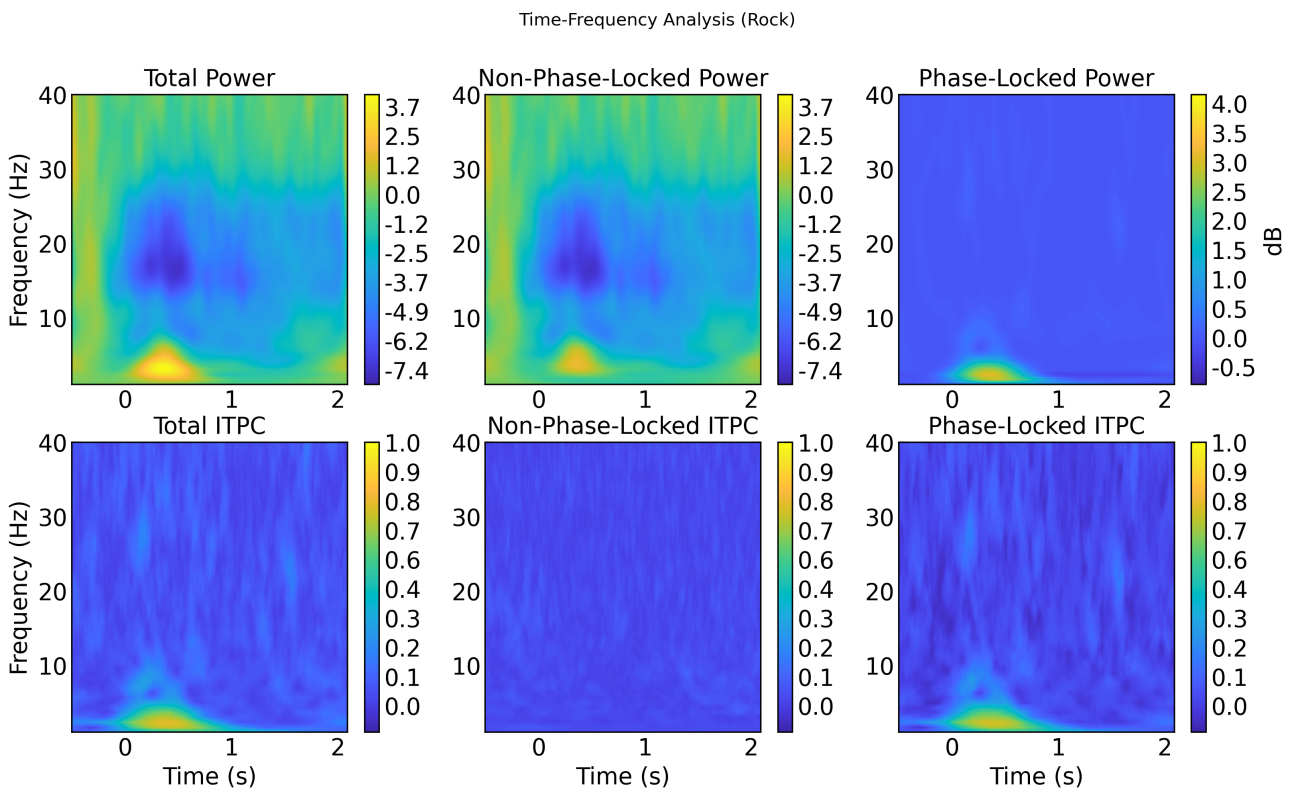


Figure 14 Time-frequency plots of rock.

Time-Frequency Analysis (Paper)

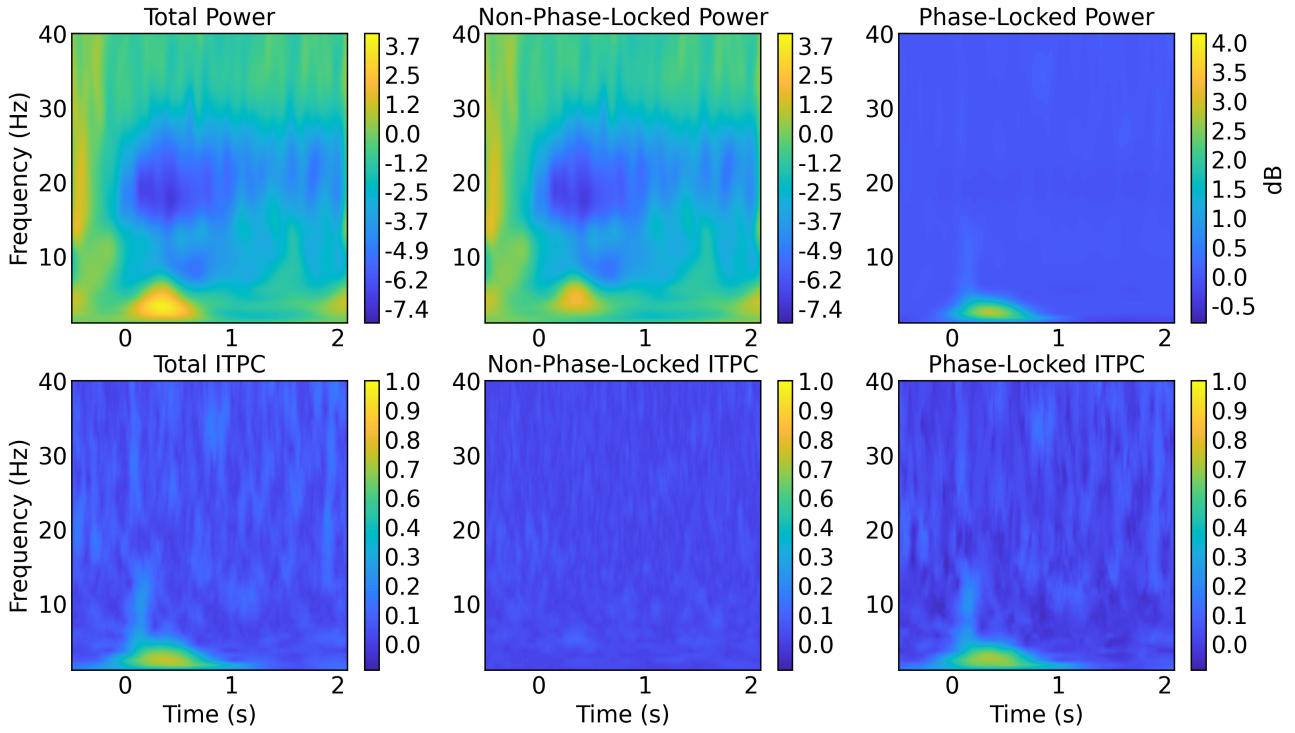


Figure 15 Time-frequency plots of paper.

Time-Frequency Analysis (Scissors)

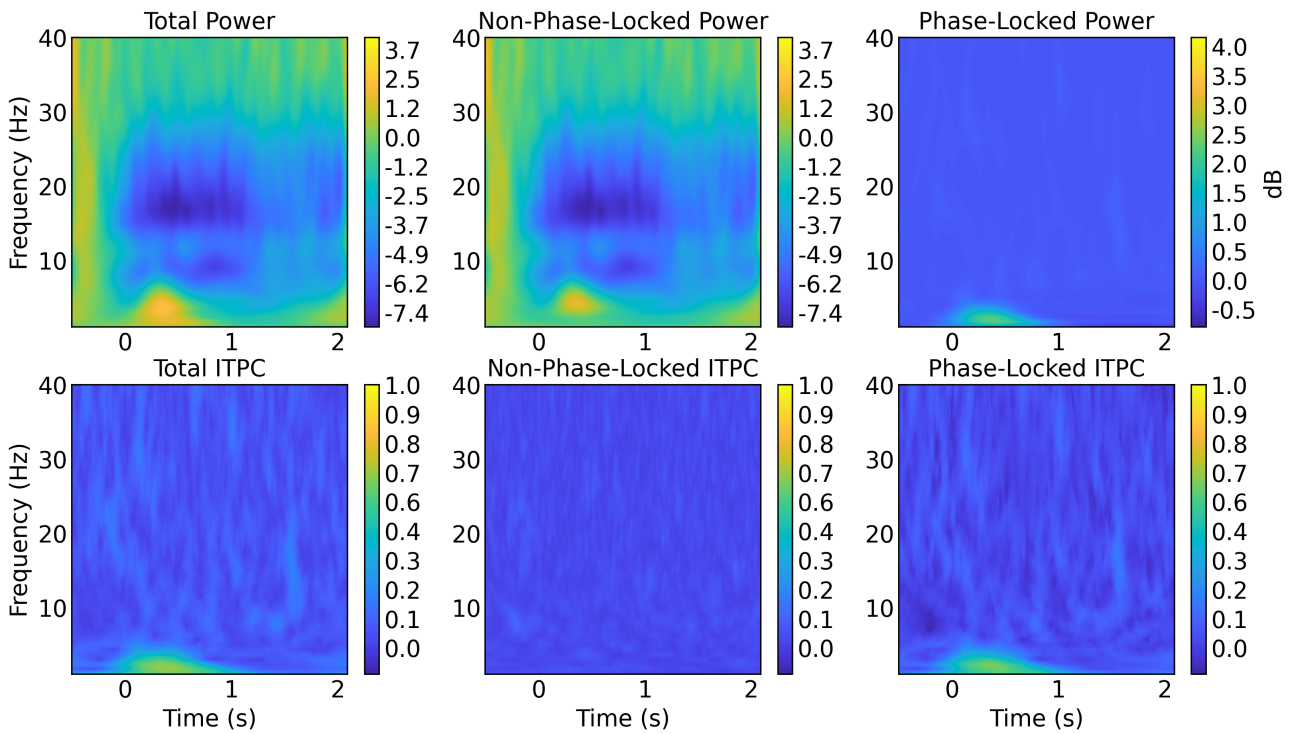


Figure 16 Time-frequency plots of scissors.




Figure 17 The surface-based sources on the 3D brain model of the participant.

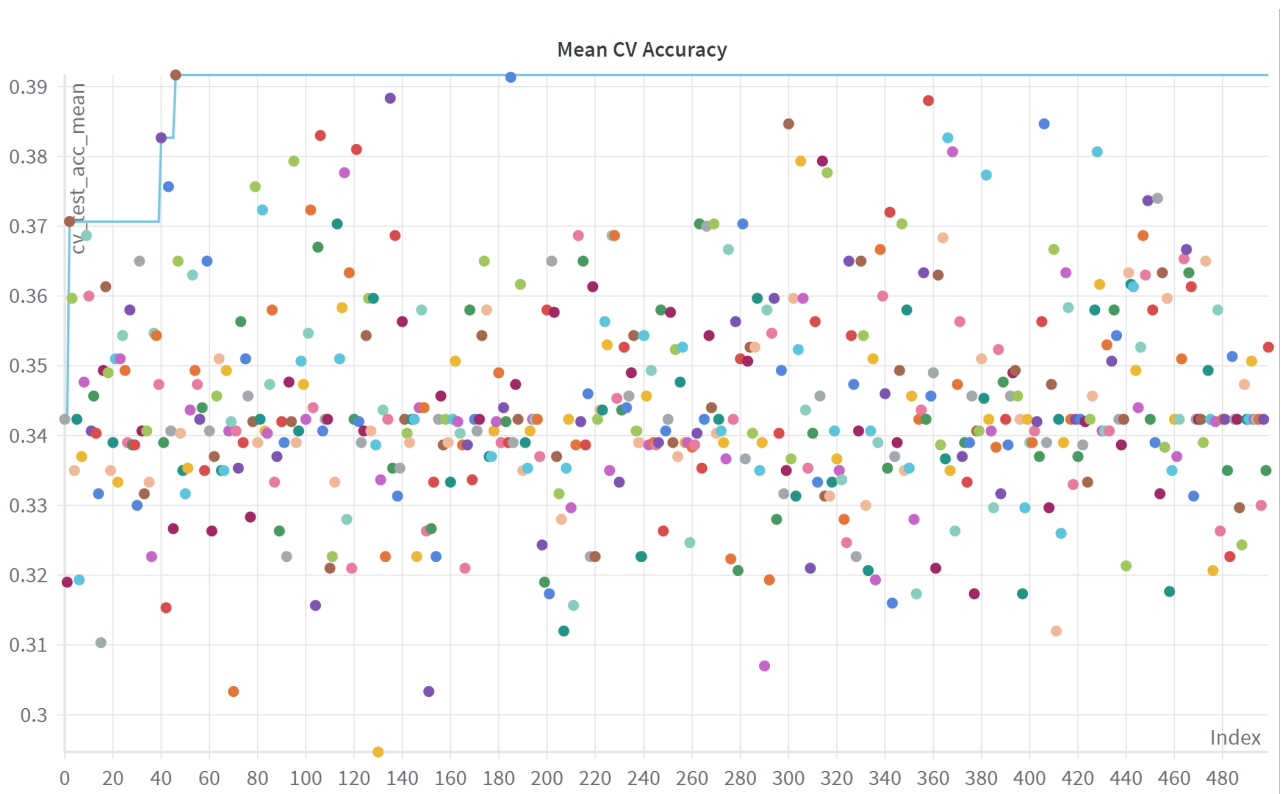


Figure 18 The results of hyperparameter sweep using Weights & Biases. Each dot represents a different configuration and the y-axis values are the CV accuracy scores. Hyperparameter optimization did not provide any significant improvement to classification accuracy.