Deep Learning for EEG Neurofeedback Training

An Adversarial Deep Learning Model Network for Mu-Rhythm Specific Neurofeedback Training

Master's thesis

Emmy van der Ree

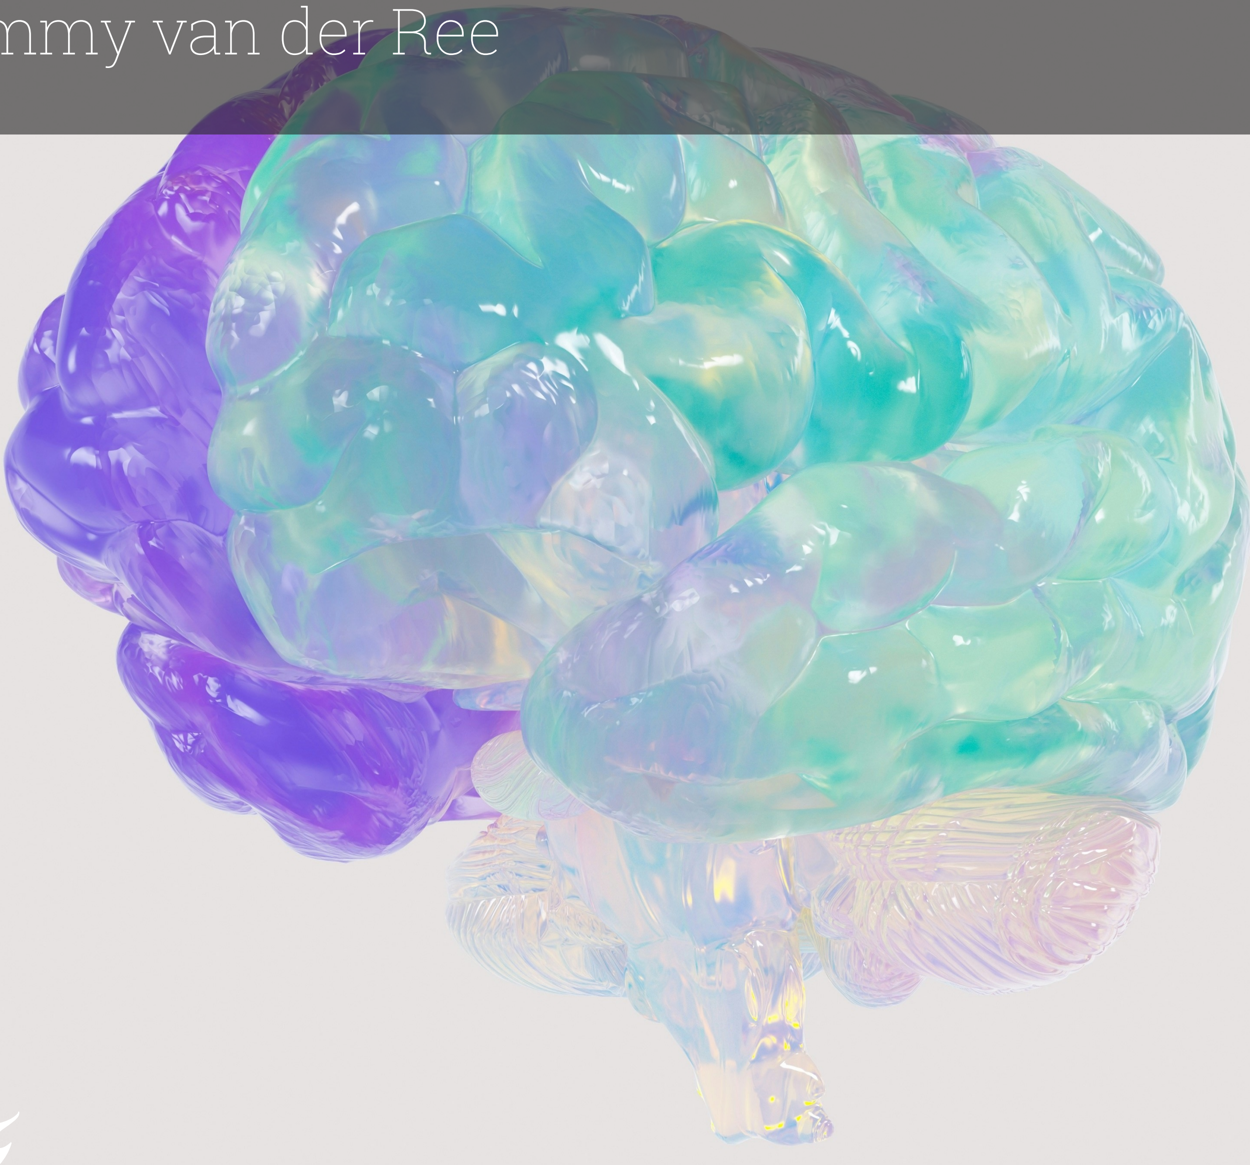

Delft University of Technology

dr.ir. Mark de Ruit & ir. Ioannis Kyriazis

An Adversarial Deep Learning Model Network for Mu-Rhythm Specific Neurofeedback Training

Emmy van der Ree
Supervisors: dr.ir. Mark de Ruit; ir. Ioannis Kyriazis

***Abstract*—This thesis presents a novel neurofeedback system for mu-rhythm modulation using an adversarial deep learning approach. The goal was to train subjects to modulate the mu-rhythm in their brain activity and to investigate the usability of this system for reflex modulation experiments. Two EEG classifiers were implemented: a Rest vs. Motor Imagery (RestMI) model and a Motor Imagery vs. Motor Movement (MIMM) discriminator. Five healthy subjects participated in five sessions of BCI training followed by a reflex assessment. During the reflex assessment the subjects had to hold a constant flexion in their wrist in order to provoke mechanical reflexes, which introduced an extra challenge for the classifiers. The RestMI model achieved a mean classification accuracy of 0.73 in the first two sessions, however performance decreased when trials with wrist flexion were introduced. The MIMM model showed a low online performance during early sessions, indicating subjects could deceive the discriminator. The reflex assessment showed mixed results, with indication of modulation of the long latency response. These findings suggest adversarial DL models can support specific mu-rhythm training in some subjects, although further work is needed with a larger sample size and more task-specific training sessions.**

***Index Terms*—EEG; Neurofeedback; Reflex modulation; Motor Imagery; Deep learning**

I. Introduction

A brain-computer interface allows a person to control a computer with their brain signals. This can be achieved using the electrical signals from the brain measured through an electroencephalogram (EEG). The subject is trained to modulate their brain activity based on feedback of the computer system. This is called neurofeedback training. Likewise, the computer system needs to adapt to the subject's unique brain physiology in order to give the appropriate feedback to the subject. Forenzo et al. (2024)

Through BCI training it is possible to train subjects to modulate a variety of brain rhythms, the rhythm this paper will be focusing on is called the mu rhythm, which ranges from 7-13Hz. This rhythm is most prominent over the sensorimotor cortex and is attenuated during motor activity or motor imagery (MI), which is the imagination of motor activity without any actual activation. The modulation of this rhythm is interesting because there is evidence that its modulation has an impact on the strength of reflexesThompson et al. (2018), as well as having possible applications in improving visuomotor performance in sportsWang et al. (2023).

Thompson et al. (2018) studied the effects of SMR modulation on the H-reflex of the flexor carpi radialis in 8 subjects (2 with spinal cord injury). The H-reflex is an electrical analog of the spinal stretch reflex. The subjects were trained over the course of 10+ sessions using a classifier based on a surface-Laplacian spatial filter, and autoregressive spectral estimation. They found that the H-reflex is significantly larger during trials in which neurologically healthy subjects up-regulated their SMR (Rest) and smaller during trials in which subjects down-regulated their SMR (with MI). This has implications in research into motor function recovery in people with CNS disorders such as traumatic spinal cord injuries or degenerative diseases like Parkinson's disease. However, training subjects to adequately modulate this rhythm takes a lot of time and requires many sessions.

Deep learning (DL) models can extract features and relationships from large and complex datasets, which makes them very suitable for decoding EEG signals and adapting to subject-specific features. Forenzo et al. (2024) Forenzo et al. (2024) tested a DL based classifier against a traditional decoder on 28 subjects in an MI-based BCI task. They retrained their classifier on all data in-between sessions. They found that the DL classifier outperformed a traditional classifier from the 3rd session onward. The subjects in this study did have at least 2 sessions of BCI training with a traditional decoder before starting the trials. Using a DL model also simplifies the data processing pipeline as artifact removal is considered a redundant process for DL because accuracy improvement is minimal and in some subjects it may remove relevant information. As such, raw signal data is used extensively with DL based MI BCI, with or without minimal preprocessing. Altaheri et al. (2023)

BCI users have a ranged ability to produce distinctive brain states. Some users (15-30%) are completely unable to control a BCI, which is called BCI illiteracy.Tibrewal et al. (2022) Allison and Neuper (2010) Some studies focus on eliminating these subjects early onJeunet et al. (2015), but this does not address any possible problems on the computer side as effective BCI control depends on synergy between subject and machine. DL has shown improved performance in low-performing users over traditional decoders. Tibrewal et al. (2022)However, this raises the question whether this improvement is due to a higher accuracy in recognizing the desired control signal, or whether the model learns to recognize spurious correlations due to confounding factors caused by the subject's (subconscious) changing of mental or physical state between trials.

This question of classification quality should be addressed when using a DL model as an EEG classifier. As the exact inner workings and training process of both the DL model

Delft University of Technology, Delft, 2628CD, The Netherlands;
e-mail: e.m.vanderree@student.tudelft.nl

and the subject's brain are unknown, they are two 'black boxes' connected to each other. DL models require a lot of training data and are usually re-trained in between sessions with the new data generated by the subject. Forenzo et al. (2024) Because the quality of the EEG data produced by the subject that is used to train the DL model is difficult to control, is impossible to determine that the DL model is being trained to only recognize a change in the mu-rhythm. If the model is trained incorrectly the subject is then subsequently trained with false feedback, generating worse data. This can result in the subject and classifier never adequately converging or converging on a control signal other than (only) mu-rhythm modulation. This could be an issue when trying to train a subject over several sessions in order to investigate the mu-rhythm's influence on reflexes.

This thesis proposes to investigate and remedy this issue by approaching the subject-computer "system" as a semi-unsupervised machine learning paradigm, and mimicking a Generative Adversarial Network (GAN) Goodfellow et al. (2014) by inserting a discriminator model. The discriminator model will be trained to distinguish between EEG signals produced during MI and actual motor movement (MM). This model will be used to control the quality of the data being used to train the classification model and its output will be used, together with the classifier, for feedback to the subject. In order for the discriminator model to be usable as a quality control, its performance should be negatively correlated with the performance of the classifier.

The research objective therefor is to create a pair of Deep Learning models that is usable for training subjects to modulate the mu-rhythm in their brain activity as measured in EEG. This approach will introduce a method of quality control for Deep Learning based neurofeedback training and this paper will investigate the usability of this system for experiments where the specific mu-rhythm modulation is required, in this case reflex modulation experiments. Lastly the results of the reflex experiments will be evaluated for effects of mu-rhythm modulation on the reflex strength.

II. METHODOLOGY

A. Participants

5 healthy participants were included in the experiment (aged 23-25, 4 right handed, 1 ambidextrous, 2 male), all were students at the TU Delft. At the start of the experiment subjects were read an introduction describing the purpose of the study and a description of the experiment structure and proceduresA, after which they were given time to read the full information letter and sign the consent form. This study was approved by the HREC committee of TUDelft, with approval number 5115.

B. Experimental Setup

Preliminary testing was performed on the PhysioNet EEG Motor Movement/Imagery Dataset Physionet (2009) to determine the best model and channel configuration to be used in the online experiments. Fewer EEG channels are preferred in the interest of reducing computational time, memory required and system complexity as well as preparation time during

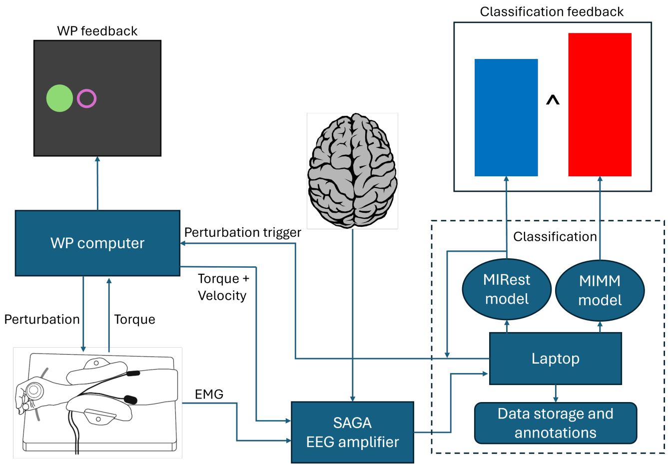

Fig. 1. The complete equipment setup during the reflex assessment. All signals (EEG, EMG and Wrist Perturbator (WP) data) are collected through the SAGA EEG amplifier for synchronization. The operator laptop receives and stores the data and timing annotations and runs the classifiers to supply the subject with feedback on their EEG signals as well as determining and sending the perturbation trigger to the RM computer. The RM computer also processes the torque acting on the handle to provide feedback to the subject on a second screen. The torque and velocity data are sent to the SAGA EEG amplifier and then combined with the EEG, EMG to be sent to the operator laptop.

electrode placement for participant comfort. It can also reduce the chance of overfitting which may be caused by including irrelevant channels. Altaheri et al. (2023) The detailed results of the preliminary testing can be found in appendix B-A. Following the preliminary testing, EEGNet was chosen as the model architecture best suited for the online experiments.

During the experiment sessions subjects were seated and had their right arm in a wrist perturbator (WP) Schouten et al. (2006) on a table. Standard EEG equipment was used (CE approved devices; TMSi SAGA Amplifier, type CF-510 K Clearance) with TMSi Infinity gel EEG caps. The ground electrode was placed on the right mastoid process, and the reference electrode was placed on the left mastoid process. EEG data was recorded from the left half + midline of a 32 electrode configuration (channel set 3), which consists of 18 electrodes total, as shown in figure 2.

Subjects received visual feedback from a classifier based on their EEG signals. The classifier of the first session is based on signal processing which does not require training data. The classifiers of subsequent sessions will be two deep learning models trained on the data of the subject's previous sessions.

During some trials, including the reflex assessment, participants were tasked to apply a small torque (maximum 5% of maximum voluntary contraction (MVC)) against the handle of the wrist perturbator. During the reflex assessment in the last session the manipulator would give perturbations to the wrist in order to provoke reflexes in the wrist flexor muscle. The perturbation consisted of a 40 ms ramp at an angular velocity of 2 rad/s resulting in a maximal amplitude of 0.08 rad. The perturbation was given in the direction of wrist extension. Giri (2022) In order to measure the reflex strength, Electromyography(EMG) was recorded with differential electrodes (Delsys Bagnoli system, 10x10mm in size), placed on the belly of the

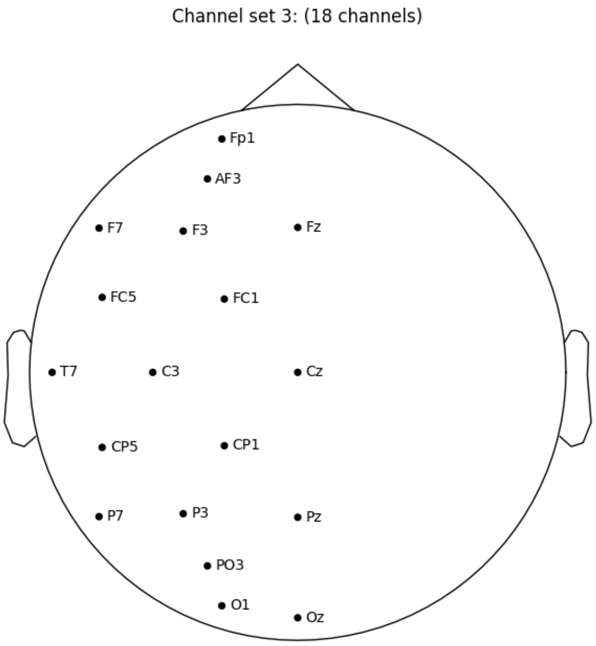

Fig. 2. Visualization of the channel montage used in the online trials.

respective muscles, the ground electrode was placed on the mastoid process.

The EEG data, torque on the handle, angular velocity of the handle and the EMG of the wrist extensor and flexor muscles were recorded at a sample rate of 500Hz. The full equipment setup used during the reflex assessment is visualized in figure 1. The online application which controlled experiment timing, classification, feedback and perturbation triggering was written specifically for this thesis project and is described in appendix C. The online application was run on a laptop with an AMD Ryzen 5 5600U (2.30 GHz) CPU with 16.0 GB of RAM. The online classification was run entirely on CPU due to the sequential nature of the operations (instead of batch-wise during model training). A single classification of a 1 second segment took 0.9 ms on this system. The classifiers were trained on the DelftBlue supercomputer Delft High Performance Computing Centre (DHPC) at TU Delft on a NVIDIA A100 GPU with 10GB of VRAM. Classifier training took between 1-4.5 hours, depending on the experiment stage.

C. Experimental Protocol

Subjects were trained over 5 1-1.5 hour sessions, taking place over the course of 3.5-6 weeks. There was a mean of 5 days (with no sessions) in-between sessions with a 1-day minimum. Due to scheduling issues and because the experiment period crossed with the university exam period, subjects 4 and 5 had a very long period in-between sessions 2-3 and 1-2 respectively, of 4 weeks for subject 4 and 3.5 weeks for subject 5. Other subjects had a maximum of two weeks in between sessions. Subject 2 did session 2 twice due to a technical error. At the beginning of each segment of session 1 subjects were read a standard description of what would be expected of them (A) and they were asked to minimize blinking and other movement, especially of the face. In the last session, or when subjects struggled with the red MIMM bar, they were told to preferably pay attention to the blue RestMI bar.

The DL classifier was used from the second session onward. To minimize the any negative influence from the subject learning to use the BCI system for the first time, the first session was structured in a way to introduce new elements one-by-one in order to get data that is as free as possible from new learning experiences. The session structures are shown in figure 3. Furthermore, each session started with a 3 trial(1 Rest, 2 Active) unrecorded practice run to allow the subject to get used to the updated classifier. This practice run was also done before the new situation of MI and wrist flexion combined. The last session was started with 2 short segments of 2 runs each in order to conduct the reflex assessment when the subject has become familiar with the new classifier and the added wrist flexion task but before the subject is possibly influenced by fatigue.

Every trial had the same timing. Trials would start with 2 seconds of preparation time where the symbol of the task would be shown. Then the feedback started, lasting for a duration of 20 seconds. Finally, there was a 3 second break. (fig. 5) 6 consecutive trials made up a Run, which adds up to a total of 150 seconds. Each session consisted of 3 segments (fig. 4) which in turn consisted of 5 runs and lasted 13.3 minutes. There were 10 second breaks between each Run and each segment was started manually by the operator so the subject could have a longer break to ask questions or drink some water. All runs consisted of an equal number of randomly distributed Active and Rest trials, the rest of the text will only mention the type of Active trial. Session 1-3 consisted of MI and MM to train the subject to control the BCI system. Session 4 consisted of two MI segments after which the MVC measurement was performed, the subject was then given a short practice run before the full segment with both MI and wrist flexion. This was done to both train the subject in this combined task and to allow the model to train on this type of data before the reflex assessment.

The feedback of the models was presented in the form of two equally-sized bars, one blue and one red. The height of the blue bar was directly determined by the RestMI model, which distinguishes between Rest(negative) and MI(positive), so this bar would become taller when the RestMI model more confidently classified the EEG signal as MI. During MI trials the height of the red bar was determined by the MIMM model, which distinguished between MI(negative) and MM(positive), so this bar would become taller when the EEG signal more closely resembled MM. Thus the subject would have the dual task of maximizing the difference between their Rest and MI state while minimizing the difference between their MI and MM state. The output of the MIMM model was first transformed by the function $((x-1)^3+1)$ to make it seem easier for the subjects to 'deceive' this model during the trials. During Rest trials both bars would be the same height, determined by the RestMI model alone.

During the reflex assessment the subject was given MI and

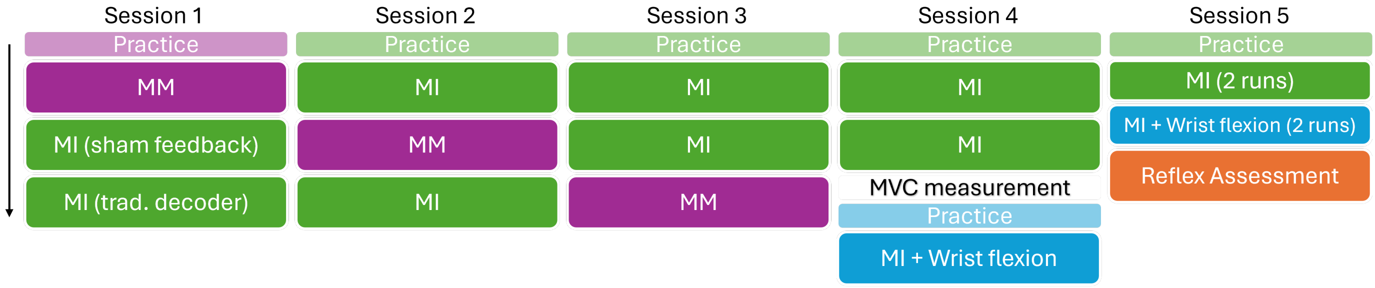

Fig. 3. Structure of the experiment sessions, indicating the Active trial type per segment (see fig 4). All MI segments from the second session onward use feedback from the DL model, and all MM segments use sham feedback only. Practice was unrecorded and consisted of 1 Rest trial and 2 Active trials, this was done before each session to allow the subject to get used to the updated classifier and before the first segment to include wrist flexion to get used to the new situation. The reflex assessment was the same as the MI + wrist flexion segments but with the wrist perturbator activated.

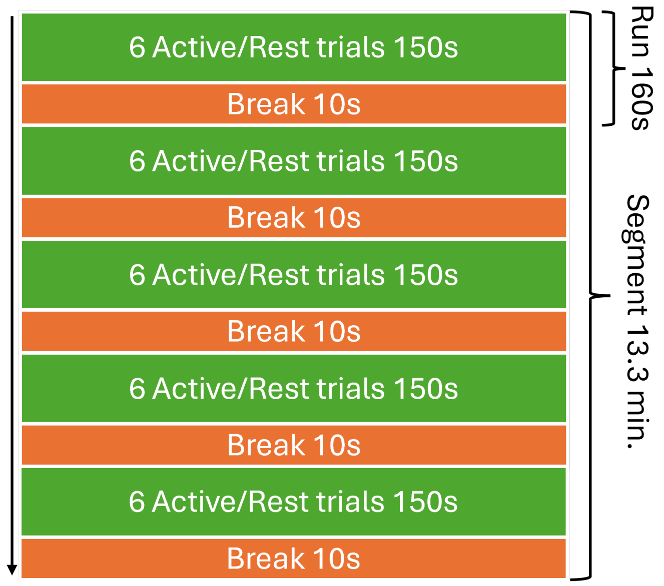

Fig. 4. Structure of each segment. The active trials could be MI or MM. There were an equal number of randomly distributed Active and Rest trials in each run. The long breaks are signified to the subject by the word 'BREAK' on screen. In between segments subjects could have as long a break as they wanted to.

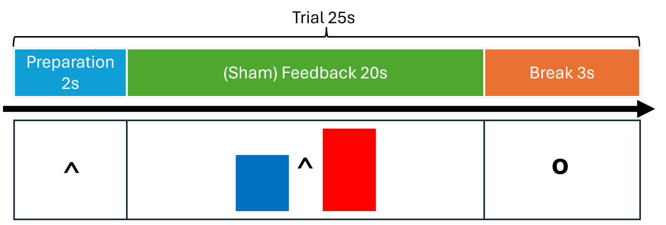

Fig. 5. Structure of an individual trial. On the top is the timing of the different parts, on the bottom is what the subject sees on the screen. The total length of a trial is 25 seconds. During the preparation the screen will show the sign of the coming task through either an upwards arrow (Active) or a downwards arrow (Rest) for 2 seconds, after which the (sham)feedback will show for 20 seconds, this period is recorded and used for model training. Finally, the subject has a 3 second break signified by a circle on the screen for them to blink or adjust position if they need to.

Rest tasks, the same as all the other MI segments. At the same time they had to apply a torque, which was 5% of their MVC as measured in session 4. While they were holding the torque, and the Rest/MI classifier determines they are in the right brain state (classifier output < 0.25 or > 0.75, respectively), the operator laptop triggers the wrist perturbator to apply a perturbation on the handle against the direction of torque that the subject is applying.

D. *Data processing*

The data was evaluated in 1s segments. For the traditional classifier a small LaPlacian spatial filter was implemented by subtracting the mean of the surrounding 4 electrodes from the signal from the C3, then the power spectral density was estimated using the welch function from the python package scipy, using a Hann window with 1 segment. After this the area under the PSD was calculated from 7-13Hz. This signal was normalized by subtracting the mean and dividing by the variance of a 30 second calibration signal collected in rest at the start of the segment.

1) DL classification: EEGNet is a compact convolutional neural network for EEG-based BCI systems. It has demonstrated exceptional classification performance across multiple datasets and is considered a benchmark model across many studies.Ke et al. (2024)

The model for each session apart from the last, was trained from scratch on data from all preceding sessions, the last session's model was trained on all of the MM data from session 1-3 and the MI data from session 2-4, the first session MI data was excluded in order to optimize the classifier for the final session, because this data could be affected by early learning effects of the subject.

A detailed description of the training conditions and training data for every session's classifier can be found in the appendix in B-B.

Before training the dataset was augmented using Gaussian noise with a 0 mean and a standard deviation of 0.01μV. The augmented multiple is 2 so the dataset will be doubled. The data is band-pass filtered between 0.1 and 40 Hz to prevent aliasing and remove line noise and low-frequency drift. The first second of every trial was cut off to account for the participant reaction speed. The data or incoming signal was classified in 1 second segments. Single trials that were disrupted by outside disturbances or small technical issues were noted during the session and discarded.

The models were trained using the AdamW optimizer from Pytorch using Cross-Entropy Loss, with a learning rate of 2e-4, a batch size of 24 and a maximum of 300 epochs. Early stopping was implemented with a patience of 30 epochs. The data was split into a train/test/validation split of size 64/20/16% respectively. During training the model was validated after each epoch. After training on the training set the optimum epoch was selected, the validation and training set added together, and the optimal epochs model was trained further until the training loss was equal or lower than the training loss of the first optimal epoch. After training the model was tested on the (training)test set.

2) Reflex assessment: The reflex data analysis was done according to the methods used in Schuurmans et al. (2009). The EMG data was band-pass filtered between 20-80 Hz using a 3rd order Butterworth filter. Segments were identified from the original data by calculating the jerk (first derivative of acceleration) from the recorded velocity data from the RM and identifying the peaks which indicate the start of the perturbation using the function find_peaks from the python package scipy.

The data was segmented from 200ms before the perturbation to 150ms after the perturbation. If the mean torque prior to perturbation deviated more than 5% from the goal torque the segment would be rejected. The segments were rectified and averaged over all segments from one subject per condition (MI/Rest).

The dimensionless magnitude of the short latency A_{M1} and the long latency A_{M2} responses were calculated from the data by calculating the mean amplitude in the time windows 20-50 ms and 55-100 ms after perturbation onset respectively. The time delay until onset of the M1 response T_{M1} was determined as the first point in time earlier than the time of maximum EMG, where the normalized EMG exceeded the value 1.0 by more than 3 times standard deviation before perturbation onset.

E. Performance Analysis

The performance of the models was evaluated during training, on an unseen test set randomly selected from the data of the previous sessions (see above), this is called the training performance from now on. The performance was also evaluated on the full dataset from the following session, during which the model was used for the online feedback, henceforth called the session performance. In the ideal situation the MIMM model's performance is lower during sessions than during training, as the goal is to have the subject 'deceive' this model.

The training performance consisted of the accuracy (using a threshold of 0.5) and the Cross-Entropy loss. Additionally, for the session performance the Area Under the Receiver-Operator Curve (AUROC) was calculated as well as the specificity to investigate bias of the model classification.

As there was no MM data in session 4 and 5, the accuracy and the AUROC of the MIMM model was only calculated for session 2-3. The session performance of the RestMI model during sessions 4 and 5 was evaluated separated by trials with and without flexion to determine the influence of wrist flexion on the quality of the EEG data.

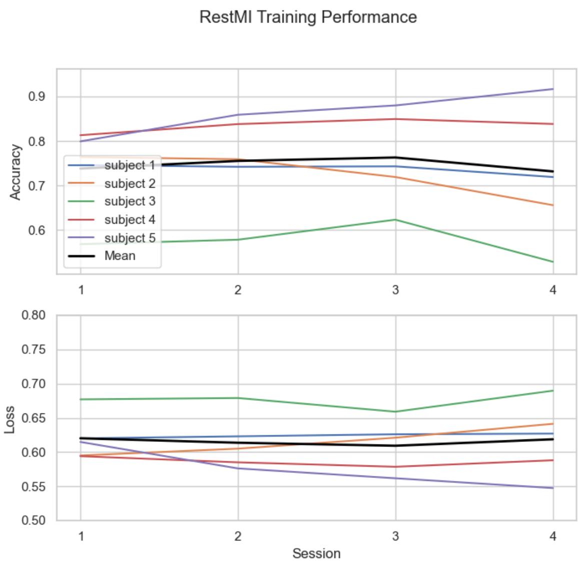

Fig. 6. Accuracy and loss after training of the RestMI model per subject by session. The mean accuracy is 0.74 after session 1, 0.76 after session 3 and 0.73 after session 4. The mean loss is 0.62 after session 1 and is the same after session 4.

F. Statistical Analysis

In order to determine if there is a correlation between the performances of the RestMI and MIMM models, the Pearson correlation between the losses of the classifiers of all subjects and sessions was calculated. Furthermore, the Pearson correlation was calculated between the training and session performance loss of the RestMI model to see how well it generalizes to new data.

Due to the low sample size, no further statistical analysis was performed.

The code written for this thesis can be found on GitHub:
https://github.com/Em-R2019/online-classification.git
https://github.com/Em-R2019/EEG-model-training.git

III. Results

All five participants completed the five training sessions. Subject 2 did session 2 twice due to technical issues. Overall subjects reporting feeling that they had control over the feedback from the DL model, but that the level to how much control varied within a session from segment to segment as well as between sessions.

A. Training performance

The training results of both the RestMI and MIMM models can be seen in figures 6 and 7. The mean performance for the RestMI model increases from session 1-3 and then drops

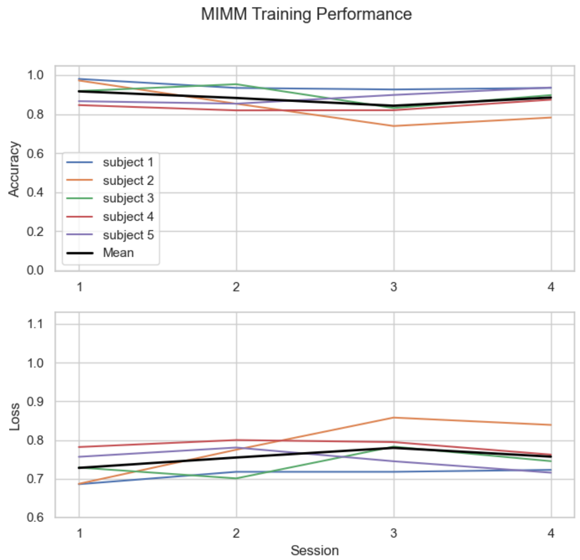

Fig. 7. Accuracy and loss after training of the MIMM model per subject by session. The mean accuracy is 0.92 after session 1, and 0.88 after session 4. The mean loss is 0.73 after session 1 and 0.75 after session 4.

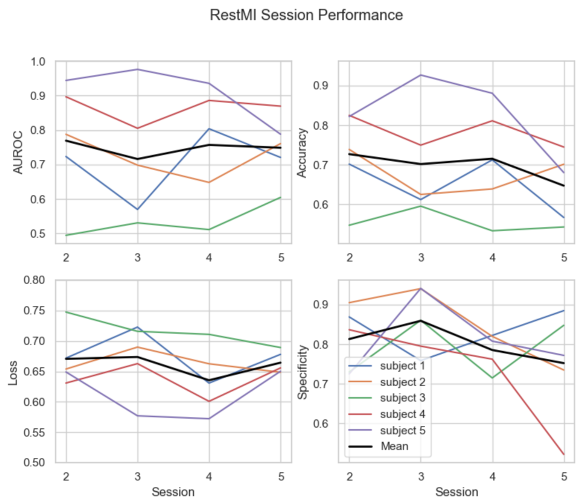

Fig. 8. Performance of the RestMI model during the sessions per subject. The mean accuracy is 0.73 for session 2 and 0.65 for session 5. The mean AUROC is 0.77 for session 2 and 0.75 for session 5. The mean loss is 0.67 for session 2 and is 0.66 for session 5. The mean specificity is 0.81 for session 2 and 0.75 for session 5.

for session 4, notably the performance per subject seems to diverge for session 4. The RestMI models for subject 4 are the lowest performers with some distance to the others, while the subject 7 and 8 RestMI models are the highest performers. The MIMM models for every subject perform better than the the RestMI model in training, with a decrease in performance over time until session 4, when the mean performance increases again.

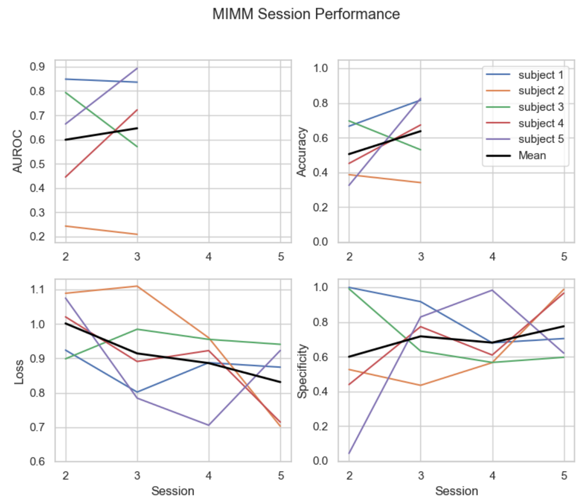

Fig. 9. Performance of the MIMM model in-session per subject. Sessions 4 and 5 did not contain any MM data, therefore the AUROC and accuracy is only calculated for sessions 2 and 3. The mean accuracy is 0.50 for session 2 and 0.64 for session 3. The mean AUROC is 0.60 for session 2 and 0.65 for session 3. The mean specificity is 0.60 for session 2 and 0.78 for session 5. The mean loss is 1.00 for session 2 and 0.83 for session 5.

B. Session performance

The in-session results of the RestMI model are summarized in figure 8. The session performance of the model is more variable over time than the training performance. The accuracy, specificity and AUROC decrease between the second and last session and the mean loss stays around the same value between the second and last session. The mean specificity is higher than the accuracy for every session meaning that the model is better at identifying (negative) Rest trials than (positive) MI trials.

The in-session results of the MIMM model are visualized in figure 9. The AUROC and accuracy are only visualized for sessions 2 and 3 as sessions 4 and 5 do not contain any MM data. Notably, the session performance is low in session 2 and 3, both in general and in comparison to the training performance. Some of the values of subjects 2 and 5 are below random classification. The mean specificity is higher than the accuracy for sessions 2 and 3, meaning the model is better at recognizing (negative) MI than (positive) MM. The in-session performance increases over time to approach the training performance in session 5.

The mean session performance of the RestMI model separated by trials with or without flexion is visualized in figure 10. In session 4 there are 6 runs worth of flexion trials to 12 no-flexion runs. In session 5 there are 8 flexion runs to 2 no-flexion runs. The performance during flexion trials is significantly lower than during no-flexion trials (0.64-0.75 and 0.63-0.71 accuracy for session 4 and 5 respectively). Performance for flexion trials increases for session 5, after the model has been retrained with flexion data, while the no-flexion performance decreases. However, no-flexion performance is still higher than flexion performance. The relations between the specificity and accuracy suggest that the model has more difficulty identifying Rest(negative) than MI(positive) trials

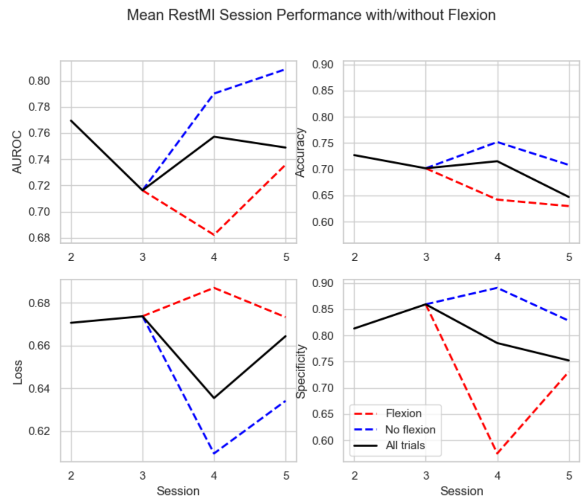

Fig. 10. Mean performance of the RestMI model during the sessions, separated by trials with or without flexion. The ratio of flexion:no-flexion trials is different between session 4 and 5: it is 1:2 for session 4 and 4:1 for session 5.

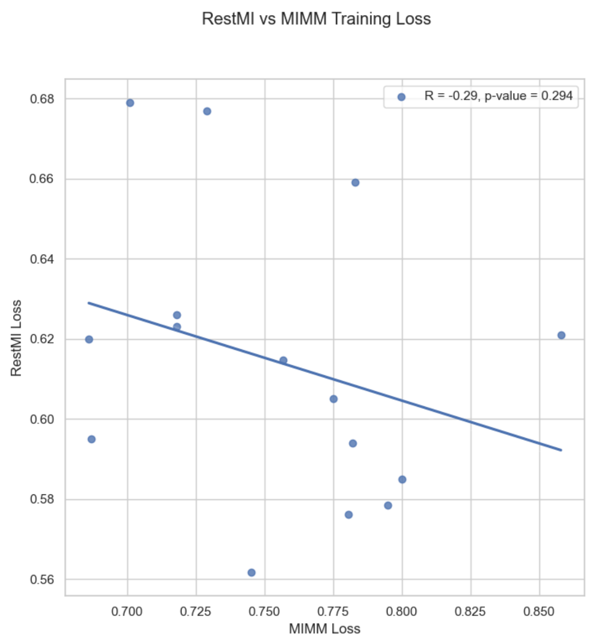

Fig. 11. Training loss of the RestMI model vs Loss of the MIMM model after session 1, 2 and 3 with a linear regression model fit. The Pearson correlation between the losses is -0.29 with a p-value of 0.29.

during session 4's flexion trials while this is the opposite for the no-flexion trials of all sessions. When the model is retrained with flexion data for session 5 that relation is restored for flexion trials.

The Pearson correlation between training and session performance of the RestMI model is 0.74 with a p-value of 0.0002.

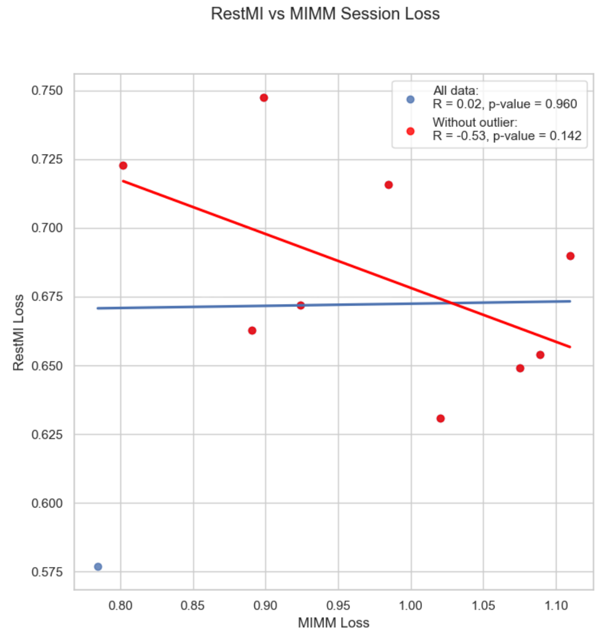

Fig. 12. In-session loss of the RestMI model vs Loss of the MIMM model of sessions 2 and 3 with a linear regression model fit with and without the potential outlier. The Pearson correlation of the losses is -0.01 with a p-value of 0.99. Without the outlier the Pearson correlation is -0.56 with a p-value of 0.12.

C. Relation RestMI and MIMM models

The relations between the RestMI and MIMM model losses are visualized in figures 11 and 12. The Pearson correlation between the training losses is -0.29 with a p-value of 0.29. This indicates a negative correlation between the RestMI and MIMM loss, but this is not significant. The Pearson correlation of the in-session losses is -0.01 with a p-value of 0.99. Indicating no correlation. When the lower left outlier (subject 5, session 3) is removed, the Pearson correlation is -0.56 with a p-value of 0.12, which again indicates a negative correlation but not statistically significant.

D. EEG Bandpower

The EEG bandpower of the Mu band in subject 4 during session 5 is visualized topographically in figure 13, separated into MI/Rest trials and with and without wrist flexion. There are high levels of power around P3, PO3 and Pz for all conditions. Power around Fz increases in the Rest conditions. During the normal Rest trials without wrist flexion there is a peak of power around C3 and CP1, but this effect is not visible in the Rest with wrist flexion condition. The topomaps of the rest of the subjects can be found in the appendix D.

E. Reflex Assessment

The results of the reflex assessment are summarized in figure 14. The mean number of segments per subject for MI was 30 (range 21-39) and 32 for Rest (range 29-39). The average A_{M1} response was the same during MI and Rest

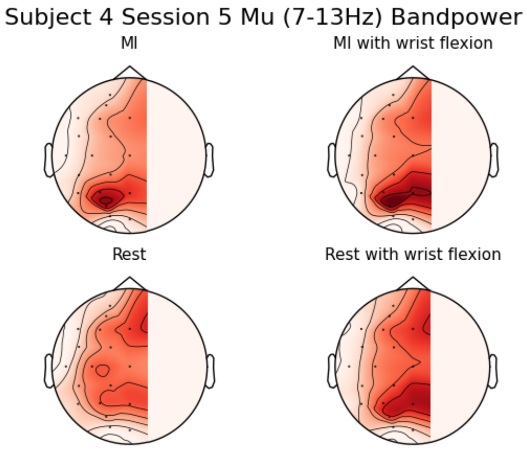

Fig. 13. Topographical map plot of the averaged Mu bandpower of subject 4 during session 5 in MI and Rest trials with and without wrist flexion. During the normal Rest trials there is a higher bandpower visible over the motor cortex, however this is not visible during the trials with wrist flexion.

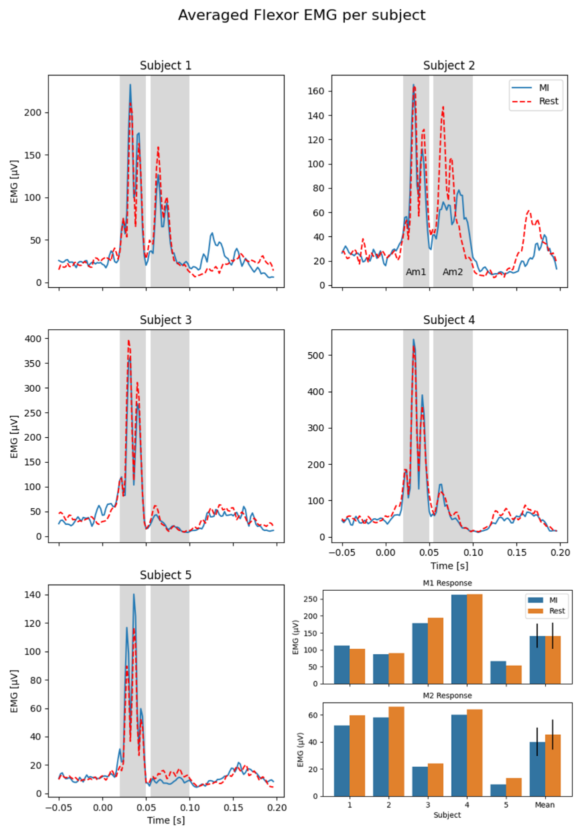

Fig. 14. Graphs: Averaged normalized EMG activation of the wrist flexor during perturbation per subject. Bar plot: A_{M1} and A_{M2} values per subject in MI and Rest. Mean A_{M1} during MI: 141 ± 36 μV, Rest: 141 ± 38 μV, mean A_{M2} during MI: 40 ± 10 μV, Rest: 45 ± 11 μV.

(MI: 141 ± 36 μV, Rest: 141 ± 38 μV) and the average A_{M2} response was higher during Rest trials (MI: 40 ± 10 μV, Rest: 45 ± 11 μV).

IV. Discussion

The design of this study incorporated a novel adversarial approach using a pair of DL models to improve the specificity of mu-rhythm-based neurofeedback training. The results do not demonstrate a statistically significant correlation between the model performances, but this should be tested further with a higher sample size.

The consistently higher specificity compared to accuracy across sessions suggests that the RestMI classifiers are more effective at recognizing Rest examples than MI examples. This bias may be the result of a relatively lower variability in the Rest state compared to the variability involved in generating MI, particularly in novice BCI users.

The observation that the MIMM model performance is significantly lower during sessions 2-3 compared to training and in some subjects performed worse than random classification, suggests that participants were indeed able to 'deceive' the discriminator to some degree, which indicates that the intended GAN-like nature of the system could have merit. Especially noting that the RestMI model did not suffer such a large drop in performance between training and session, so it is less likely that it is caused by large inter-trial discrepancies due to EEG non-stationarity or changes in MI strategy. However, over time the session performance improves, even though the training performance slightly decreases, indicating that as the generalization ability of the model rises due to having a larger training dataset, it becomes harder for the subjects to deceive it. This is in line with observations during the online trials, as several subjects indicated that it became increasingly difficult to make both bars rise at the same time.

Presenting the output of both models as equal might have led to confusion or frustration in participants. A different way of combining the two types of feedback may work better, where the RestMI model output is clearly the most important. Especially as the ultimate goal was to create a large difference between the Rest and MI brain states. Possibly the discriminator could be mainly used as a way to assess and tag data quality offline before training the RestMI model rather than use it as feedback for the subject directly. A traditional classifier or a human expert might be used for this purpose as well, especially to identify correct up-regulation during Rest trials, which the paradigm in this thesis does not have a method for controlling. Alternatively, separate training periods could be incorporated to teach the subject each task separately: first the ability to accurately down-regulate their mu rhythm using the MIMM model, and only then incorporate the Rest task into the training paradigm.

Some subjects mentioned having difficulty not thinking about their hand during the rest trials. Furthermore, some subjects found it difficult to focus on the majority white screen for long periods of time, suffering from dry or teary eyes.

Ideally, subjects would learn to control a BCI system with the traditional decoder as much as possible before starting data collection for classifier training to minimize the influence of the initial subject learning process negatively influencing the performance of the model like in Forenzo et al. (2024). Due to the short experiment period and because DL models require a lot of data to train it was decided to use the first session's data for the classifier training and start utilizing the DL classifier from the second session onward. Even though the first session was structured in such a way as to minimize the impact of early learning effects, this could have negatively impacted the model's ability to focus on mu-rhythm specific patterns rather than more general cognitive engagement signals.

The introduction of wrist flexion trials during session 4 had a clear deteriorating effect on the session and training performance of the RestMI model. Retraining the model after session 4 with wrist flexion data clearly improved the flexion performance while decreasing the no-flexion performance. This underscores the difficulty in generalizing EEG-based classifiers across even minor changes in task execution. It emphasizes the necessity of task-specific training data for creating a DL classifier for BCI.

This study aimed to train subjects to accurately down-regulate their mu rhythm by integrating the MIMM discriminator model into the BCI system. During standard Rest and MI tasks, performed without physical movement or muscle contraction, the MI condition may require more cognitive effort, as it is a deviation from the resting baseline. However, during the reflex assessment, where subjects maintained continuous wrist flexion, this paradigm was effectively inverted as MI became more akin to the new baseline, while the Rest condition, requiring the suppression of motor intent during ongoing muscle activation, might be more difficult.

The EEG bandpower data of the final session indicates that subject 4 was able to up-regulate their Mu rhythm during the rest trials, but was not able to do this while simultaneously holding wrist flexion. Possibly due to distraction from having to concentrate on two tasks and/or not having enough training in the new task. The data of the rest of the subjects are less clear and most show strong eye artifacts. This may be due to less well-fitting EEG-caps, and eye irritation from the mostly white screen. This is most likely for subject 3, who was in-between cap sizes and used the larger cap size for comfort, they also struggled the most with dry eyes from looking at the screen. 2 subjects (2 and 3) have stronger eye artifacts during the flexion trials and 2 (1 and 5) show less eye artifacts during flexion trials, meaning there is no clear effect of having to look between the two feedback screens on eye artifacts.

Giri (2022) tried to replicate the results of Thompson et al. (2018) in the spinal stretch reflex using mechanical perturbation with 5 healthy subjects. The subjects were trained over 5 to 8 sessions. This study found significant effects in the mechanical reflex in only 1 subject. The results of the reflex assessment of this thesis show that there is no difference in the magnitude of the M1 response. The M2 response, however, shows a smaller magnitude during MI in all subjects. This indicates that mu modulation affects the M1 and M2 responses differently. The M2 response has in fact previously been shown to be modulated by the cerebral cortex and can be affected by intention and habituation. It is also possible for the voluntary response to overlap with the second half of M2 (75 ms after perturbation) Lemmers (2022) Goodin et al. (1990), however this overlap is not apparent in any of the subjects' EMG data. Both the M1 response and the H-reflex are elicited through Ia-afferents Cruccu and Deuschl (2000), so following the results of Thompson et al. (2018), one would expect the M1 response to be affected by mu modulation as well. It is possible that this response would show an effect after a longer training period because this would yield both cleaner data as the classifier becomes better at recognizing mu modulation as well as subjects which can modulate their mu rhythm more strongly.

A key limitation of this study is the small sample size (n=5) which limits the generalizability of the findings. The large inter-subject variability in model performance, especially of the RestMI model could be caused by differences in mental state, cognitive strategies, fit of the EEG cap and responsiveness to neurofeedback. The reported performances are a combination of the model performance and subject performance, making it difficult to distinguish whether performance improvement is the result of model adaptation to a larger dataset or of subject learning. A channel ablation study could be performed to investigate which channels affect classifier performance the most, this would say something about which spatial features the models are learning to recognize. This information could help discern between subject learning or model adaptation effects.

Despite these challenges, this thesis demonstrates that incorporating adversarial dynamics into a neurofeedback system can be a viable strategy. However, the resulting internal model dynamics should be explored more thoroughly, for example through a channel-ablation study. A larger sample size is also needed to generate more significant results and more training sessions, specifically more task-specific training sessions, are required to improve performance. Furthermore the feedback design should be refined in order to more clearly communicate the intended goal. For the future of DL BCI, more work should be done to extricate and separate subject and model performance in order to effectively evaluate and improve the next generation of BCI systems.

V. Conclusion

This study introduced an adversarial deep learning approach to enhance the specificity of mu-rhythm neurofeedback in an MI-based BCI system. By pairing a Rest vs MI classifier with a MI vs MM discriminator, the system was designed to ensure higher data quality during training. Results showed that the MIMM model could be successfully deceived during the earlier sessions, supporting the feasibility of adversarial dynamics in neurofeedback.

The drop in classifier performance during wrist flexion trials highlights the importance of task-specific training data. Although the small sample size limits generalizability, this work demonstrates the potential of adversarial DL models to improve neurofeedback systems. Future work should refine feedback design, increase session numbers, and could potentially further separate task learning and implement data quality assessment using the discriminator model to improve training effectiveness.

The results of the reflex assessment indicate an effect of mu-rhythm modulation on the long-latency M2 reflex response. However, a study with larger sample sizes with an improved BCI system should be performed to increase the generalizability of the findings.

Acknowledgment

I would like to thank my supervisors Mark and especially my daily supervisor Ioannis for their support and guidance throughout the sometimes rocky process of this thesis and my literature review. I would also like to thank Cyrano for pointing me towards GANs for the problem I was trying to solve. Furthermore I couldn't have done this without the love and support of my closest friends, my aunt and a special shoutout to Crina, for being my double marathon buddy;).

References

Allison, B. Z. and Neuper, C. (2010). Could anyone use a bci? In *Brain-computer interfaces: Applying our minds to human-computer interaction*, pages 35–54. Springer.

Altaheri, H., Muhammad, G., Alsulaiman, M., Amin, S. U., Altuwaijri, G. A., Abdul, W., Bencherif, M. A., and Faisal, M. (2023). Deep learning techniques for classification of electroencephalogram (eeg) motor imagery (mi) signals: A review. *Neural Computing and Applications*, 35(20):14681–14722.

Cruccu, G. and Deuschl, G. (2000). The clinical use of brainstem reflexes and hand-muscle reflexes. *Clinical Neurophysiology*, 111(3):371–387.

Delft High Performance Computing Centre (DHPC) (2024). DelftBlue Supercomputer (Phase 2). https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2.

Forenzo, D., Zhu, H., Shanahan, J., Lim, J., and He, B. (2024). Continuous tracking using deep learning-based decoding for noninvasive brain–computer interface. *PNAS nexus*, 3(4):pgae145.

Giri, N. (2022). Design and validation of a neurofeedback protocol to assess the effect of the sensorimotor rhythm on reflex modulation. Master's thesis, Delft University of Technology.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

Goodin, D. S., Aminoff, M. J., and Shih, P.-Y. (1990). Evidence that the long-latency stretch responses of the human wrist extensor muscle involve a transcerebral pathway. *Brain*, 113(4):1075–1091.

Jeunet, C., N'Kaoua, B., Subramanian, S., Hachet, M., and Lotte, F. (2015). Predicting mental imagery-based bci performance from personality, cognitive profile and neurophysiological patterns. *PloS one*, 10(12):e0143962.

Ke, S., Yang, B., Qin, Y., Rong, F., Zhang, J., and Zheng, Y. (2024). Fact-net: a frequency adapter cnn with temporal-periodicity inception for fast and accurate mi-eeg decoding. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*.

Lawhern, V. J., Solon, A. J., Waytowich, N. R., Gordon, S. M., Hung, C. P., and Lance, B. J. (2018). Eegnet: a compact convolutional neural network for eeg-based brain–computer interfaces. *Journal of neural engineering*, 15(5):056013.

Lemmers, D. J. (2022). Stretch evoked potentials and m2 reflexes: The effect of time between perturbations and intention. Master's thesis, Delft University of Technology.

Mane, R., Robinson, N., Vinod, A. P., Lee, S.-W., and Guan, C. (2020). A multi-view cnn with novel variance layer for motor imagery brain computer interface. In *2020 42nd annual international conference of the IEEE engineering in medicine & biology society (EMBC)*, pages 2950–2953. Ieee.

Miao, Z., Zhao, M., Zhang, X., and Ming, D. (2023). Lmda-net: A lightweight multi-dimensional attention network for general eeg-based brain-computer interfaces and interpretability. *NeuroImage*, 276:120209.

Physionet, E. (2009). Eeg motor movement/imagery dataset.

Schouten, A. C., de Vlugt, E., van Hilten, J. B., and van der Helm, F. C. (2006). Design of a torque-controlled manipulator to analyse the admittance of the wrist joint. *Journal of neuroscience methods*, 154(1-2):134–141.

Schuurmans, J., De Vlugt, E., Schouten, A. C., Meskers, C. G., De Groot, J. H., and Van Der Helm, F. C. (2009). The monosynaptic ia afferent pathway can largely explain the stretch duration effect of the long latency m2 response. *Experimental brain research*, 193:491–500.

Song, Y., Zheng, Q., Liu, B., and Gao, X. (2022). Eeg conformer: Convolutional transformer for eeg decoding and visualization. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:710–719.

Thompson, A. K., Carruth, H., Haywood, R., Hill, N. J., Sarnacki, W. A., McCane, L. M., Wolpaw, J. R., and McFarland, D. J. (2018). Effects of sensorimotor rhythm modulation on the human flexor carpi radialis h-reflex. *Frontiers in neuroscience*, 12:505.

Tibrewal, N., Leeuwis, N., and Alimardani, M. (2022). Classification of motor imagery eeg using deep learning increases performance in inefficient bci users. *Plos one*, 17(7):e0268880.

Wang, K.-P., Frank, C., Hung, T.-M., and Schack, T. (2023). Neurofeedback training: Decreases in mu rhythm lead to improved motor performance in complex visuomotor skills. *Current Psychology*, 42(24):20860–20871.

APPENDIX A
PARTICIPANT INSTRUCTIONS

A. Introduction

The purpose of this research study is to investigate whether deep learning can be used to promote training humans to voluntarily modulate their brain activity. We also want to understand how practicing this might influence reflexes, which are automatic responses of your muscles that are often affected following brain injury. We are interested in how reflexes can temporarily become stronger or weaker during voluntary brain activity modulation, this will be tested at the end of the research study. Participating in this research will take you 5 sessions of 1-1.5 hours to complete within 5 weeks. The brain activity that we are interested in is called the Sensori-Motor Rhythm, or SMR. This rhythm is most prominent in the central areas of the brain and is related to active movement or processing of information from your senses. The rhythm is stronger when you are at rest and becomes weaker when you are active. It also becomes weaker when you only imagine movement or sensory experiences. During this study you will train your ability to modulate the SMR without actual movement. During the experiments you will be seated. Your arm will be resting on a table while you are holding the handle of a robotic manipulator. In order to measure your brain activity you will wear an electroencephalogram (EEG) cap, this looks like a swimming cap with sensors attached to it. In the last two sessions of the study you will be asked to apply a small torque (max 10% of your maximum) on the handle of the robotic manipulator. While holding the small torque, the manipulator will provoke reflexes by short and small displacements of the position of your hand when it is active. During this, the reflex activity of your wrist muscles will be measured using adhesive electrodes on your arm. The manipulator will not be active until the very last session, and you will be made aware when it will be used. Today's session will consist of three segments, each lasting about 12 minutes. during the first segment you will be asked to actually move your wrist. During the second segment you will be asked to imagine moving your wrist and for the last segment you will be asked to imagine moving your wrist while you receive feedback from the computer on your SMR. Later sessions will be structured similarly. Your participation in this study is entirely voluntary and you can withdraw at any time. If you decide not to participate, no further action is required. You can ask the researcher questions about any aspect of the research study.

B. Start experiment (MM)

This segment will last about 12 minutes. During this segment you will receive instructions on the computer screen in front of you to move your hand $\hat{()}$ or to relax (v). We start with actual hand movement to get a baseline of your brain activity during actual movement. The tasks will last 20 seconds, in-between these tasks you will receive short breaks (o) of a couple seconds and longer breaks, signified by the word 'break' on screen, of about 15 seconds. During the 'relax' task you should really focus on relaxing, look at the screen and try not to think about anything. During a break you may do anything you want, look around, move, think about anything or scratch your face. Please try to minimize blinking or facial movements during the tasks and use the breaks for blinking or stretching your face. Try to loosely grip the handle and move it at a calm speed. During the tasks you will see two coloured bars on the screen, these are to simulate visual feedback and have no meaning during this segment. You should look at the screen during the tasks. The segment consists of 5 runs with 6 tasks each, the tasks will be ordered randomly. The segment will start with two practice trials after which you will have the ability to ask questions.

C. Start MI

This segment will again last about 12 minutes. It is structured the same as the last, except when the upwards arrow is shown you should keep you hand still and only imagine moving your wrist and the sensations associated with this movement. If you struggle to imagine just moving your wrist, you can imagine moving the handle or performing some other movement that mainly involves using your wrist, such as drawing. There will be no practice runs, the trials will start right away.

D. Start feedback (traditional decoder)

This segment will again last about 12 minutes. It will start with a 30 second calibration for the computer to adapt to your EEG signals. During this time you should simply look at the screen without blinking and not performing any task. After this it is structured the same as the others, and when the upwards arrow is shown your task will again be only imagining the movement of your wrist or hand without actually moving it. However, this time you will receive feedback on your SMR from the computer through the two coloured bars on the screen. When you are performing correctly, the bars should move in the direction of the arrow. So, when the upwards arrow is shown you should imagine movement and the bars should be as tall as possible. When the downwards arrow is shown you should relax and the bars should be as short as possible. Again, if you struggle to imagine just moving your wrist, you can imagine moving the handle or performing some other movement that mainly involves using your wrist, such as drawing. During this first session the computer is using a simple algorithm to generate the feedback which may limit its accuracy. Try to not be discouraged and focus on the imagination of moving your wrist.

APPENDIX B
METHODS

A. Preliminary testing

Several models were selected from literature to test whether they could differentiate between MI and MM. The models selected were FACT-NetKe et al. (2024), EEGNetLawhern et al. (2018), ConformerSong et al. (2022), LMDA-netMiao et al. (2023) and FBCNetMane et al. (2020). Only the first two were found to be effectively able to recognize the difference after training. These two were tested under different circumstances to see which one worked best. Only the first two were found to be able to recognize the difference after training.

The two selected models were tested with different electrode configurations on an offline dataset Physionet (2009) in order to reduce the number of electrodes needed in the online trials in the interest of participant comfort and time required. The electrode configurations tested are shown in figure 15. The results of the tests are shown in figure 16. EEGNet was selected for the online trials, using the left half + midline of a 32 electrode configuration (channel set 3), which consists of 18 electrodes total.

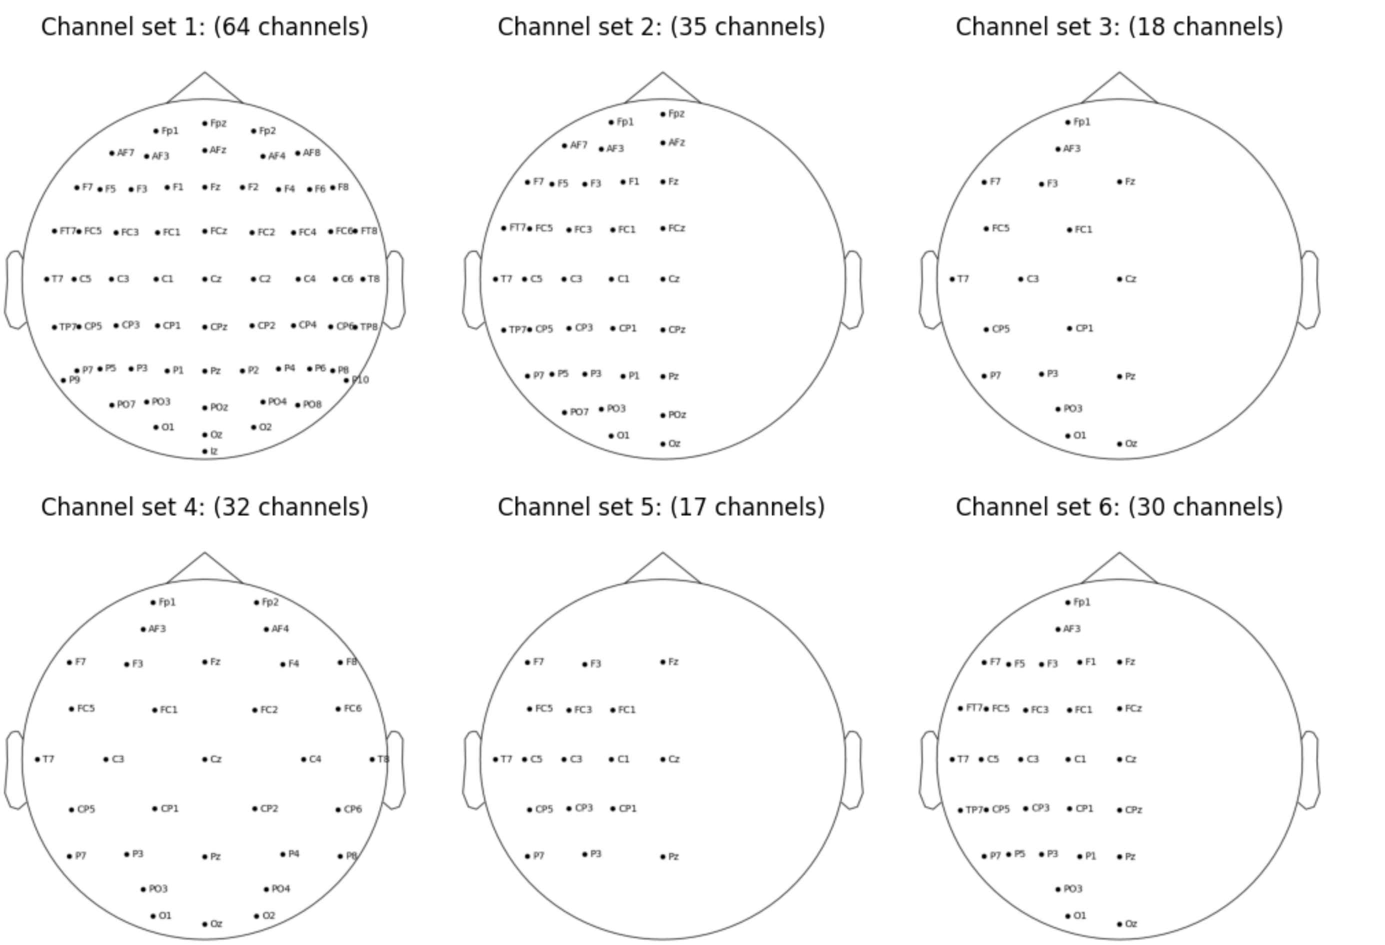

Fig. 15. Visualization of the channel sets tested offline.

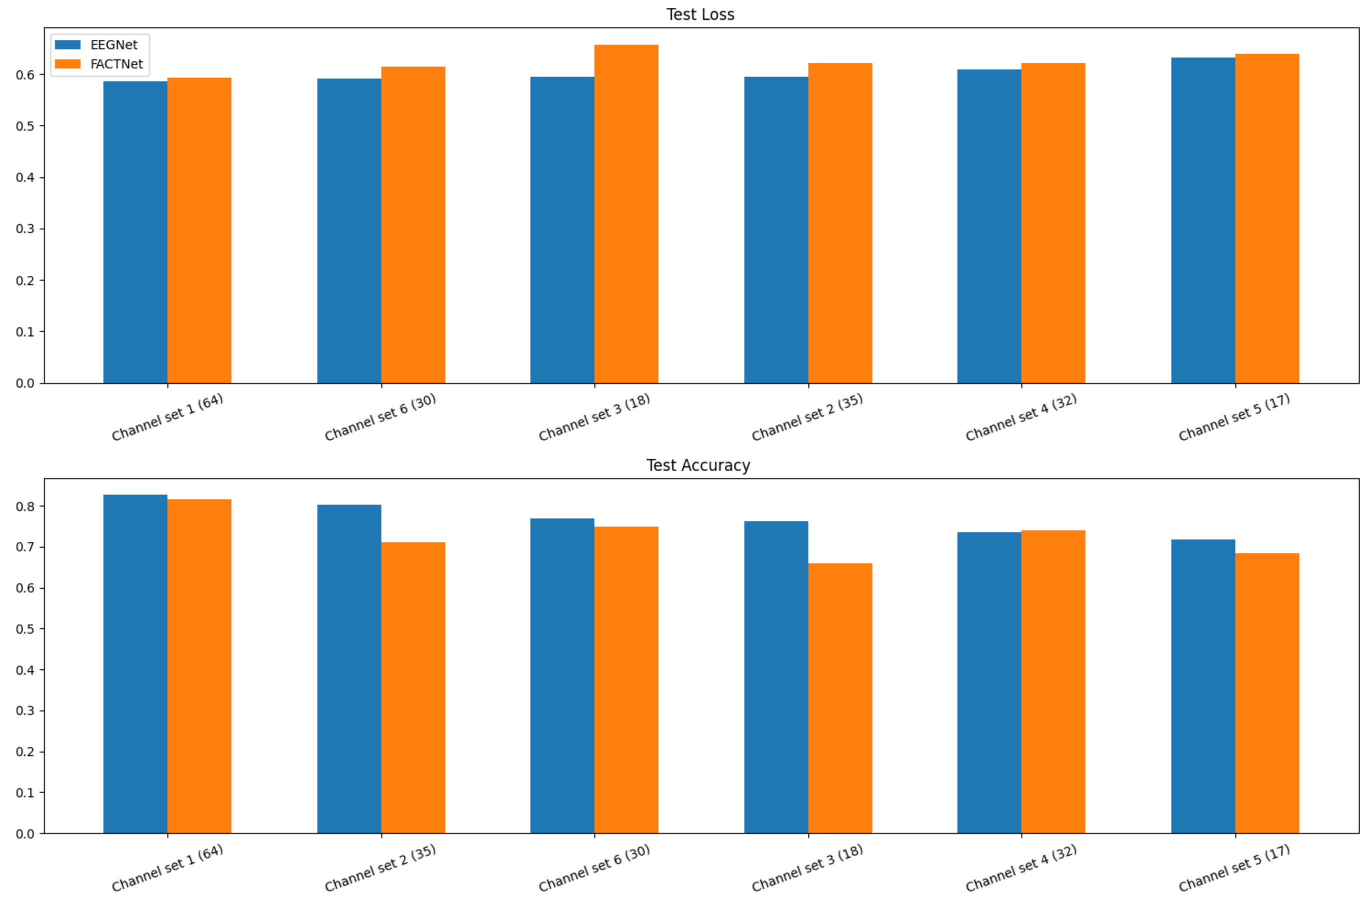

Fig. 16. Testing loss and accuracy of FACTNet and EEGNet on different channel sets for MM vs MI classification.

B. Data

The tables in this section describe the trial conditions and training data for the classifiers.

TABLE I
TRIAL CONDITIONS

Condition	**Main task**	**Feedback**	**Perturbations**
MM	Moving the right hand back and forth (flexion-extension) while holding the handle of the WP	sham	no
Rest	Mental rest while holding the handle of the WP	DL classifier (unless specified otherwise)	no
MI	Imagining moving the wrist while holding the handle of the WP	DL classifier (unless specified otherwise)	no
MI with wrist flexion	Imagining moving the wrist while holding a constant torque (5% of mvc) against the handle of the WP (= wrist flexion)	DL classifier	Only during reflex assessment
Rest with wrist flexion	Mental rest while holding a constant torque (5% of mvc) against the handle of the WP (= wrist flexion)	DL classifier	Only during reflex assessment

TABLE II
DATA USED FOR TRAINING & TESTING CLASSIFIERS. GRAY CELLS: SESSION DATA, WHITE CELLS: TRAINING DATA

Data \ **Classifiers**	**Session 1**	**Session 2**	**Session 3**	**Session 4**	**Session 5**
Session 1	15 MM; 15 sham feedback MI; 15 trad. classifier MI; 30 sham feedback Rest; 15 trad. classifier Rest	all	all	all	MM only
Session 2		15 MM; 30 MI; 15 sham feedback Rest; 30 Rest	all	all	all
Session 3			15 MM; 30 MI; 15 sham feedback Rest; 30 Rest	all	all
Session 4				30 MI; 30 Rest; 15 MI + wrist flexion; 15 Rest + wrist flexion	all
Session 5					6 MI; 21 MI + wrist flexion; 6 Rest; 21 Rest + wrist flexion

Appendix C
Online App

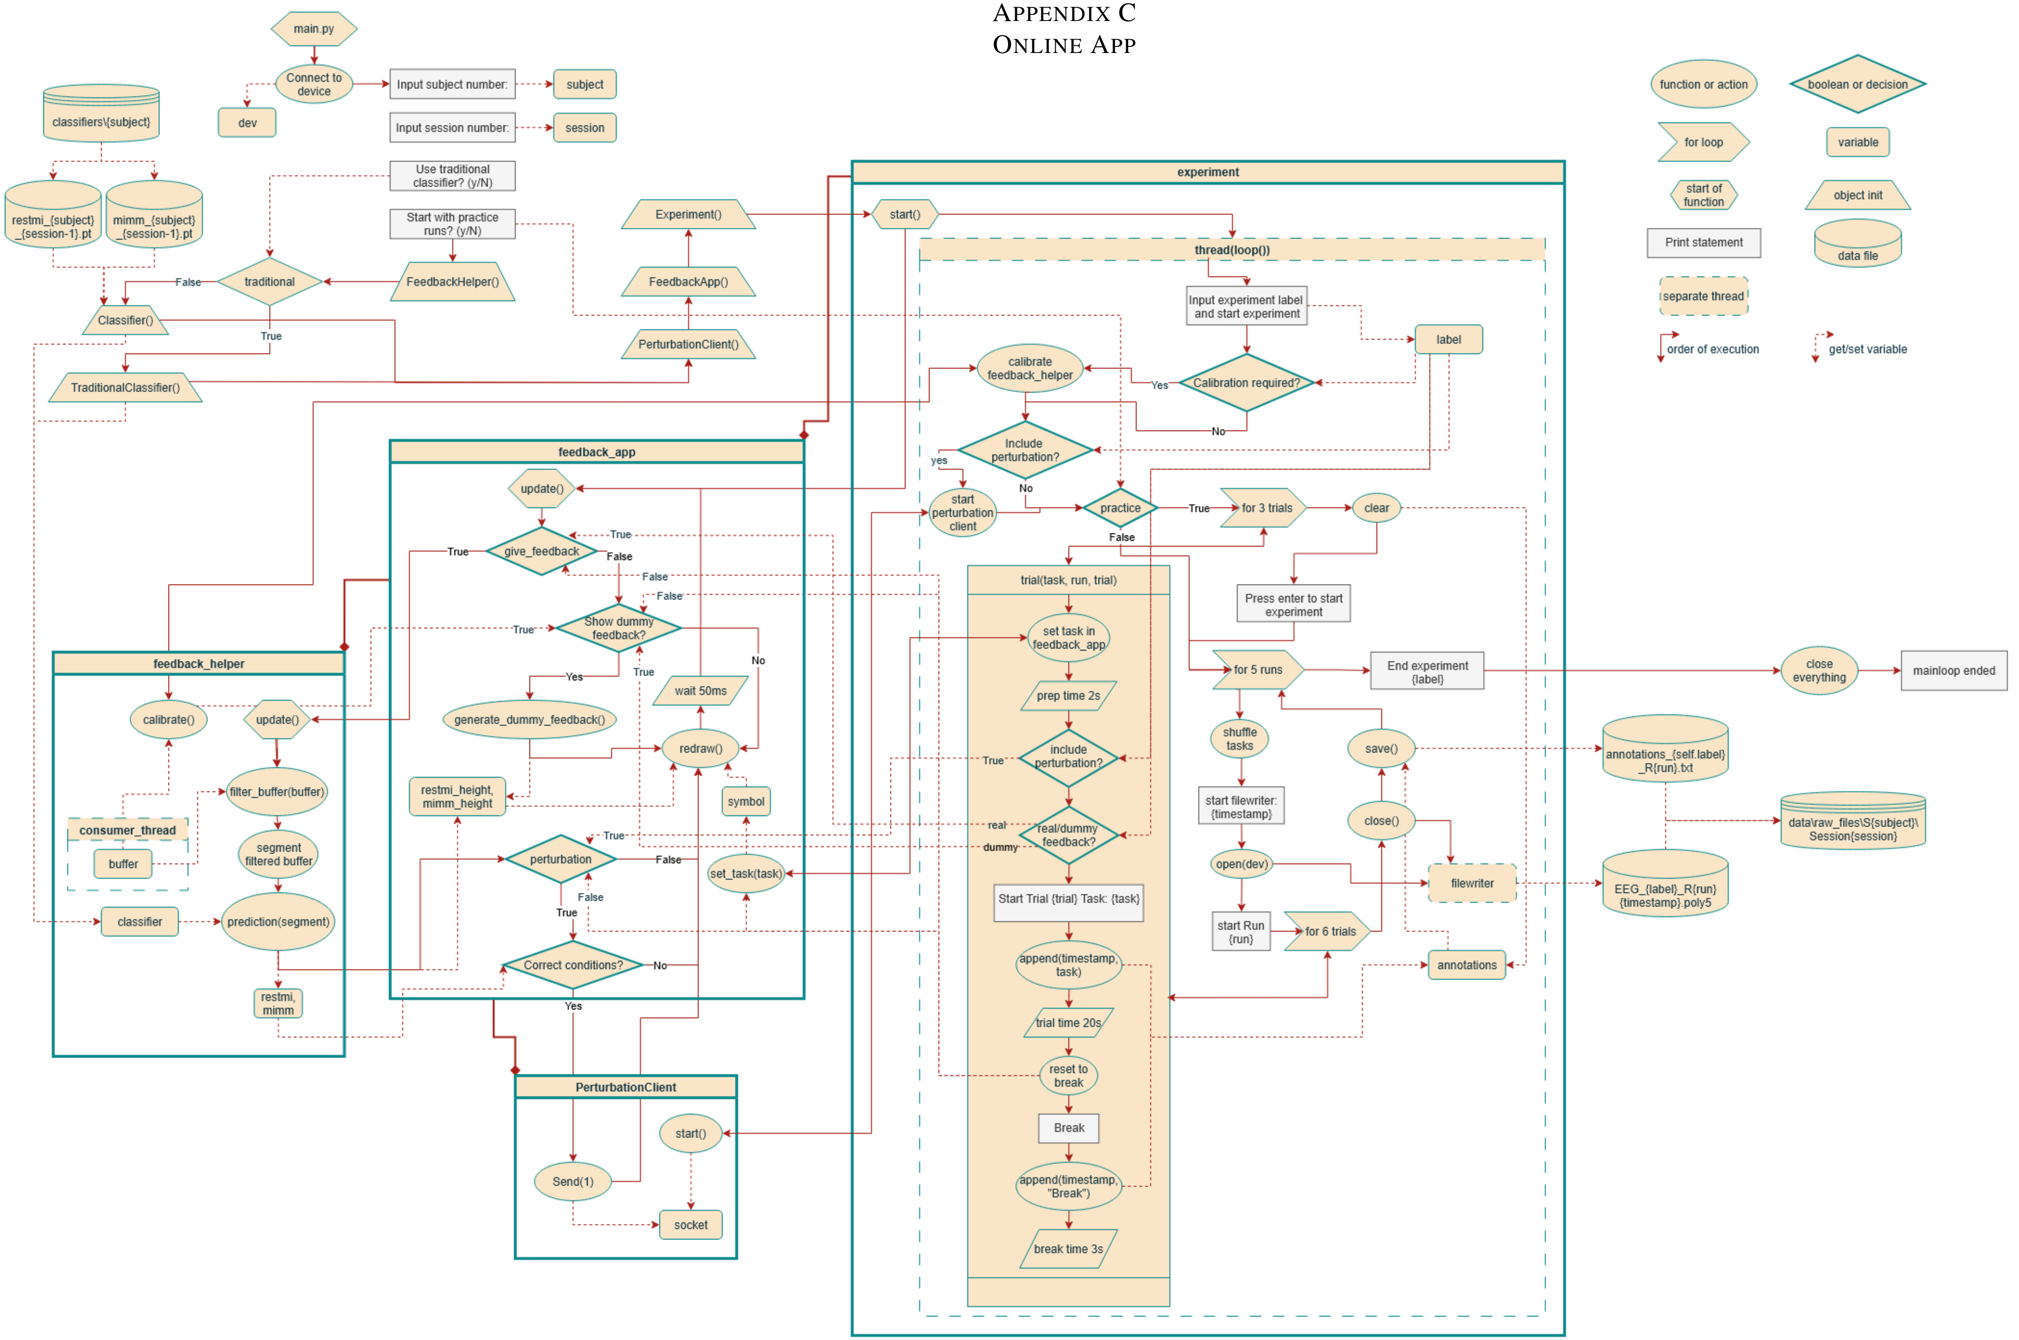

APPENDIX D
TOPOGRAPHIC PLOTS

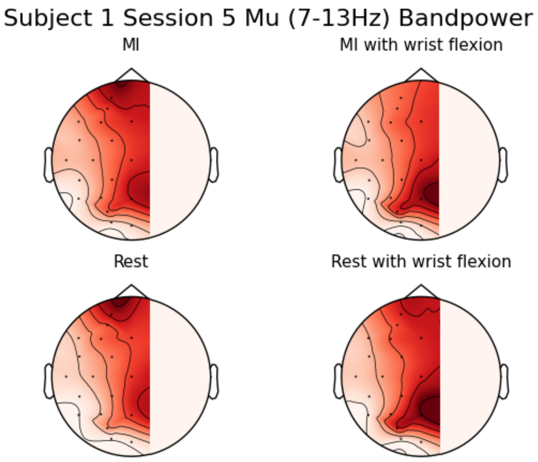

(a)

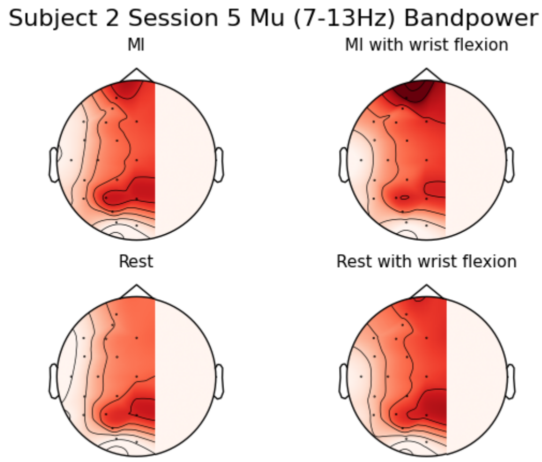

(b)

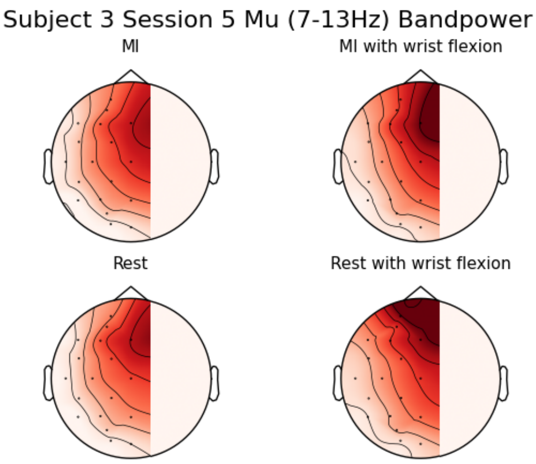

(c)

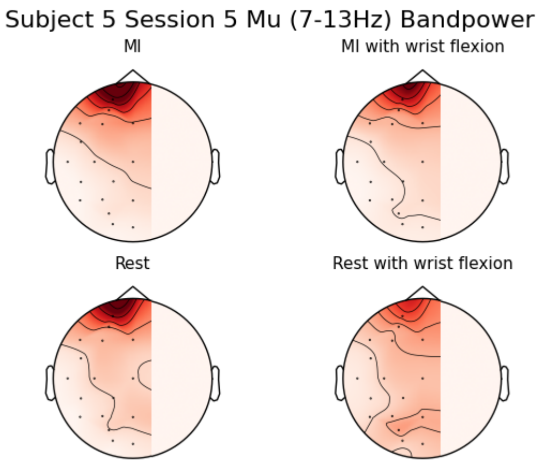

(d)

Fig. 17. Topographical map plots of the averaged Mu bandpower of subjects during session 5 in MI and Rest trials with and without wrist flexion. (d): One bad channel (FC1) was removed from the data of subject 5.