Title: Decoding personalized motor cortical excitability states from human electroencephalography

Running title: Decoding personalized excitability states

Authors: Sara J Hussain[1]* and Romain Quentin[2]

Affiliations: [1]Movement and Cognitive Rehabilitation Science Program, Department of Kinesiology and Health Education, University of Texas at Austin. Austin, TX, USA.
[2]COPHY Team, Lyon Neuroscience Research Center (CRNL), INSERM, CRNS, Université Claude Bernard Lyon 1, Lyon, France.

***Corresponding Author:**
Sara J Hussain, PhD
540 Bellmont Hall
2109 San Jacinto Blvd
Austin, TX 78712
Email: sara.hussain@austin.utexas.edu
Phone: 847-309-5189

Conflicts of Interest: The authors declare no competing financial interests.

Acknowledgements: Data included in this publication were acquired at the Human Cortical Physiology and Neurorehabilitation Section (directed by Leonardo Cohen) at NINDS.

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Abstract

OBJECTIVE: Brain state-dependent transcranial magnetic stimulation (TMS) requires real-time identification of cortical excitability states. Current state-of-the-art approaches deliver TMS during brain states shown to correlate with motor cortex (M1) excitability at the group level. Here, we hypothesized that machine learning classifiers could successfully discriminate between high and low M1 excitability states in individual subjects using power spectral features obtained from low-density electroencephalography (EEG) signals.

METHODS: We analyzed a pre-existing publicly available dataset that delivered 600 single TMS pulses to the right M1 during EEG and electromyography (EMG) recordings. Personalized multivariate pattern classification was used to discriminate between brain states during which TMS evoked small and large motor-evoked potentials (MEPs). These brain states were labeled as low and high M1 excitability states, respectively.

RESULTS: Personalized classifiers successfully discriminated between low and high M1 excitability states in every subject tested. MEPs elicited during classifier-predicted high excitability states were significantly larger than those elicited during classifier-predicted low excitability states in all subjects. Personalized classifiers trained using EEG features obtained immediately before each TMS pulse performed better than personalized classifiers trained using EEG features obtained from earlier time points. Personalized classifiers generalized weakly but significantly across subjects.

CONCLUSION: Individual subjects exhibit unique brain activity patterns that correspond to low and high M1 excitability states. These patterns can be efficiently captured using power spectral features obtained from low-density EEG signals. Deploying individualized classifiers during brain state-dependent TMS may enable effective, fully personalized neuromodulation in the future.

Keywords

machine learning, brain stimulation, transcranial magnetic stimulation, motor cortex

1. Introduction

Over the past few decades, transcranial magnetic stimulation (TMS) has emerged as a potential treatment for a variety of brain disorders, including depression (George et al. 1995; O'Reardon et al. 2007), substance use disorders (Gorelick et al. 2014; Hanlon et al. 2018; Ekhtiari et al. 2019), memory problems (Wang et al. 2014; Nilakantan et al. 2019), and motor impairment caused by neurological damage (Tazoe and Perez 2015; Hummel and Cohen 2006; Webster et al. 2006). Although early work reported that different patterns of repetitive TMS significantly alter primary motor cortex (M1) excitability (Pascual-Leone et al. 1994; Chen et al. 1997; Stefan et al. 2000; Huang et al. 2005), such interventions have since been shown to exhibit weak effect sizes and wide inter- and intra-individual variability (Ziemann and Siebner 2015; Goldsworthy et al. 2021), with in some cases less than half of tested individuals showing the expected plastic response (López-Alonso et al. 2014).

Brain state-dependent TMS may improve the magnitude and consistency of TMS effects by delivering TMS during periods of increased cortical excitability and thus increased sensitivity to LTP-like plasticity (Kraus et al. 2016; Kraus et al. 2018; Zrenner et al. 2018; Baur et al. 2020). In the motor system, brain state-dependent TMS typically involves applying stimulation during brain states that correlate with M1 excitability at the group level, like sensorimotor mu rhythm desynchronization (Haegens et al. 2011; Kraus et al. 2016; Kraus et al. 2018) or sensorimotor mu rhythm trough phases (Haegens et al. 2011; Zrenner et al. 2018; Baur et al. 2020; Hussain et al. 2021). Results from studies using this approach have so far been promising: M1 TMS interventions applied during high excitability brain states potentiate corticospinal excitability (Kraus et al. 2016; Kraus et al. 2018; Zrenner et al. 2018; Baur et al. 2020) and enhance motor memory (Hussain et al. 2021) more than identical interventions applied during low excitability states.

Nonetheless, current brain state-dependent TMS approaches pose some challenges. Most notably, a reliable high excitability brain state must first be detected across a range of individuals. This can be difficult, as shown by several recent studies reporting inconsistent relationships between the sensorimotor mu rhythm and corticospinal excitability (Sauseng et al. 2009; Berger et al. 2014; Zrenner et al. 2018; Hussain et al. 2019a; Madsen et al. 2019; Bergmann et al. 2019, see Madsen et al. 2019 for summary). Further, maximizing the beneficial effects of brain state-dependent TMS requires the target brain state to positively correlate with excitability in each individual. Excluding individuals who do not show such a correlation is also

problematic, as it limits the generalizability of results and may artificially inflate effect sizes by screening out potential “non-responders” ahead of time (Madsen et al. 2019). Other obstacles include difficulty extracting the target brain state in all individuals (Zrenner et al. 2020), a problem that will only be exacerbated in clinical studies involving patients with brain lesions (Hussain et al. 2020) or other forms of volumetric brain loss.

Here, we developed an analytical approach that overcomes these issues by exploiting the rich spatiotemporal dynamics embedded within whole-scalp EEG signals. Specifically, we aimed to identify personalized patterns of brain activity that predict the presence of high and low M1 excitability states. We hypothesized that multivariate pattern classification, a form of supervised machine learning, could successfully discriminate between high and low M1 excitability states in individual subjects using only simple power spectral features calculated from low-density EEG signals. To test our hypothesis, we trained personalized machine learning classifiers to identify brain states during which single-pulse TMS evoked either large or small motor-evoked potentials (MEPs) in 20 healthy adults. Our results show that these personalized classifiers successfully discriminated between high and low M1 excitability states in 100% of tested subjects. Further, personalized classifiers generalized weakly across subjects, indicating that individual subjects exhibit unique brain activity patterns that reflect high and low M1 excitability states. Overall, our findings provide the conceptual and methodological framework for efficient, fully personalized decoding-based brain state-dependent TMS interventions in the future.

2. Materials and Methods

2.1. Human subjects

Data analyzed in this study have been published previously (Hussain et al. 2019a; Hussain et al. 2019b) and are publicly available (https://openneuro.org/datasets/ds002094). The original study involved simultaneous single-pulse transcranial magnetic stimulation (TMS), electroencephalography (EEG), and electromyography (EMG) in 20 healthy subjects (6 F and 14 M, age 30 ± 1.59 years). Original study procedures were approved by the National Institutes of Health Combined Neuroscience Section Institutional Review Board and conformed to local regulations for human subjects research. All subjects provided their written informed consent prior to participation. The original study involved a 5-minute eyes-open resting EEG recording followed by application of 600 single TMS pulses during EEG and EMG recordings (see Hussain et al. 2019a for additional methodological details). Only EEG and EMG data acquired during TMS application are presented here.

2.2. Data acquisition

2.2.1. EEG and EMG signal acquisition. EEG signals were recorded using TMS-compatible amplifiers (BrainAmp MR+, BrainVision; sampling rate = 5 kHz, hardware filtering = DC – 1 kHz, resolution 0.5 µV). EEG impedances were maintained below 12 kΩ throughout all recordings. Bipolar EMG signals were recorded using adhesive electrodes arranged in a belly-tendon montage (Signal, Cambridge Electronic Design; sampling rate = 5 kHz, hardware filtering = 5 Hz – 2 kHz).

2.2.2. TMS. TMS was applied as previously reported (Hussain et al. 2019a; Hussain et al. 2019b). Briefly, monophasic posterior-to-anterior current TMS (MagStim 200^2, MagStim Co. Ltd) was applied over the scalp hotspot for the left first dorsal interosseous muscle (L. FDI) using a figure-of-eight coil. After determining resting motor threshold (RMT) using an automatic threshold-tracking algorithm (Adaptive PEST procedure; Awiszus and Borckhardt 2011), 6 blocks of 100 single TMS pulses were applied to the L. FDI hotspot at 120% RMT (inter-stimulus interval = 5 s with 15% jitter). Breaks were taken during each block when requested by the subject or in response to TMS coil overheating. TMS coil position was monitored for accuracy using frameless neuronavigation throughout the experiment (BrainSight, Rogue Research).

2.3. Data and statistical analysis

2.3.1. EMG processing. EMG signals acquired during single-pulse TMS delivery were used to calculate peak-to-peak MEP amplitudes, which served as an objective marker of M1 excitability. EMG data were analyzed using custom-written scripts in MATLAB combined with FieldTrip (Oostenveld et al. 2011). EMG data were segmented from 25 ms to 100 ms after TMS delivery, demeaned, and linearly detrended. Residual EMG offset was removed by subtracting the average EMG amplitude between 1 and 20 ms before TMS delivery from the entire timeseries signal. MEPs contaminated by voluntary motor activity were identified as those trials that contained (1) pre-stimulus EMG root-mean-square values exceeding an individually-defined upper limit (75th percentile + 3 * the inter-quartile range) or (2) MEP signals that were poorly correlated with the average MEP signal observed in that subject (Pearson's correlation, right-tailed, $r < 0.4$). These MEPs were excluded from further analysis (7.52 ± 1.40% of all MEPs). In each remaining trial, the peak-to-peak MEP amplitude was calculated as the difference between the maximum and minimum voltage deflection between 20 and 40 ms after each TMS pulse. Low and high M1 excitability states were defined at the individual subject level as those brain

states during which single-pulse TMS elicited an MEP with an amplitude in the top 50% or bottom 50% of all remaining MEP amplitudes, respectively (i.e., median split). MEP amplitudes were dichotomized in this way because brain state-dependent TMS requires a binary decision (i.e., whether to deliver TMS) to be made at each time step during real-time signal processing.

2.3.2. EEG preprocessing. EEG signals preceding single-pulse TMS delivery were used to calculate spectral power across a range of frequencies. Resulting spectral power values were used as features when building subject-specific classifiers. EEG data were analyzed using custom-written scripts in MATLAB combined with FieldTrip (Oostenveld et al. 2011). EEG signals acquired during single-pulse TMS delivery were segmented from 3.005 to 0.005 s before TMS delivery. Data from TP9 and TP10 were removed due to poor scalp electrode contact; thus, 28 channels remained for analysis (FP1, FP2, F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T7, T8, P7, P8, Fz, Cz, Pz, Iz, FC1, FC2, CP1, CP2, FC5, FC6, CP5, and CP6). Data for all remaining channels were re-referenced to the common average reference, demeaned, linearly detrended, notch filtered to remove line noise (zero-phase shift Butterworth filter, order = 4, stop band = 58 to 62 Hz), and downsampled to 500 Hz to reduce computational load. Power spectra were calculated for six 500 ms windows leading up to each TMS pulse (windows included 3.005 – 2.505 s, 2.505 – 2.005 s, 2.005 – 1.505 s, 1.505 – 1.005 s, 1.005 – 0.505 s and 0.505 – 0.005 s before TMS) using Welch's method (4-100 Hz with 0.25 Hz resolution, excluding line noise frequencies [58-62 Hz]). Power spectral estimates were averaged within each frequency band separately for each window. This created 168 power spectral features per window (i.e., 28 EEG channels for each of 6 different frequency bands). Features included broadband (4-100 Hz, excluding line noise frequencies [58-62 Hz]), theta (4-7 Hz), alpha (8-12 Hz), beta (13-35 Hz), low gamma (36-58 Hz), and high gamma (62-100 Hz) power.

2.3.3. Single-trial classification at the individual subject level. Multivariate pattern analysis (MVPA) was performed using the MVPA-Light toolbox (Treder et al. 2020) combined with custom-written scripts in MATLAB. We first aimed to predict whether an MEP elicited by a single TMS pulse was either small or large (see *EMG processing)* using information contained within power spectral features obtained immediately preceding each pulse (0.505 – 0.005 s before TMS; see *EEG preprocessing*). To achieve this, we used Linear Discriminant Analysis (LDA) with stratified k-fold cross-validation (k = 10).

We optimized two key criteria when fitting personalized classifiers: 1) the power spectral features included in the classifier, and 2) the lambda value (i.e., learning rate). For each subject, we optimized the features used by first ranking all 168 power spectral features calculated from EEG data immediately preceding TMS delivery by importance using chi-square tests. That is, we tested for statistical dependencies between each feature and the two M1 excitability states (i.e., high or low, see *EMG preprocessing*). The negative log of the chi-square test's p-value was calculated as the feature score; here, higher feature scores reflect features that more strongly covary with M1 excitability states. All features were ranked by importance using the feature score. We then iteratively searched for the optimal number of features by training and testing 168 different LDA classifiers, with each classifier using a feature set that had 1 more feature than the previous one. During consecutive iterations, features were added in order of importance (i.e., classifier 1 used the most important feature, classifier 2 used the two most important features, etc.). For each iteration, subject-specific LDA classifiers were fit to a training subset of EEG features (90% of all available trials per subject). This allowed the classifier to find the hyperplane that best separated the low and high M1 excitability states. The trained classifier was then used to predict whether a given trial's MEP belonged to either the low or high excitability state using a separate testing subset of the EEG features (remaining 10% of trials). These steps were performed for each fold (k=10), generating a predicted excitability state for each trial's MEP. We calculated the area under the receiver-operator curve (AUC, 1 = perfect classification, 0.5 = chance-level classification) by comparing this prediction to each MEP's ground truth class. Within each iteration, AUC values and classifier weights were averaged across folds. While optimizing the features used, we also optimized each subject's lambda value by training and testing 100 LDA classifiers for each set of features, with possible lambda values linearly spaced between 1e-50 and 1. Training and testing sets were always independent and stratification was used.

In sum, we performed grid search to identify the optimal number of features and the optimal lambda value for each subject. After grid search, each subject's best performing classifier was selected and used for all subsequent analyses. Subject-specific confusion matrices indicating the proportion of accurate and inaccurate classifications for each subject's optimized classifier were calculated and then averaged across subjects to obtain a single group-level confusion matrix. On average, optimizing each subject's personalized classifier took 7.17 ± 0.55 minutes (range = 5.23 – 11.98) on a standard laptop computer (MacBook Pro, 2.3 GHz Quad-Core Intel

i7 processor, 16 GB RAM). See Figure 1 for a schematic illustration of the classification and optimization pipeline.

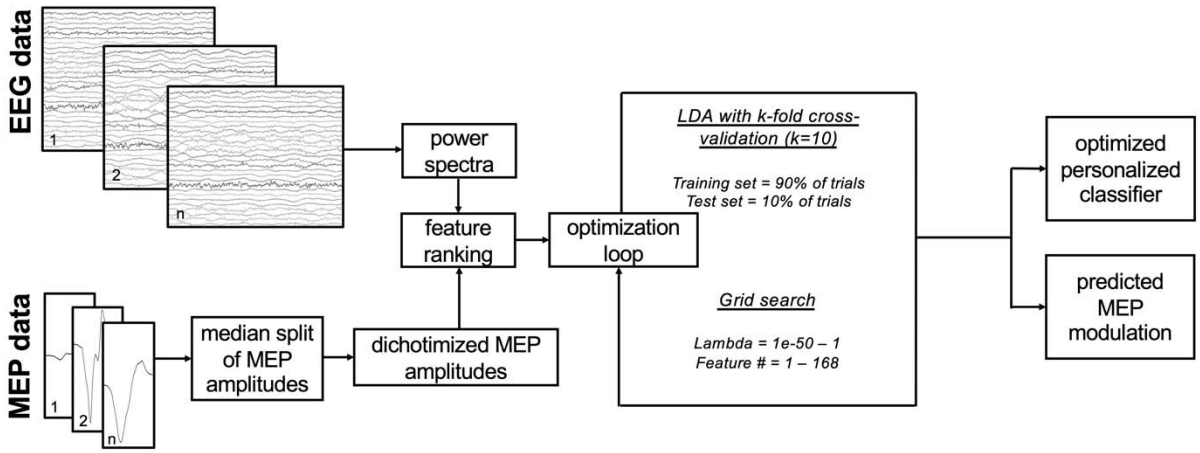

Figure 1. Analytical procedures used to build and optimize each subject's personalized classifier. Raw EEG data undergo spectral analysis. MEP amplitudes are calculated and divided into small and large MEPs via median split. Afterwards, dichotomized MEPs and power spectral features are used to rank each subject's features according to importance. Then, subject-specific grid search is performed by fitting multiple LDA classifiers for a range of lambda values and feature numbers, with features added in order of importance. Each classifier is cross-validated to prevent overfitting. Then, each subject's best performing classifier is selected and the predicted MEP modulation is calculated from predicted classes obtained during cross-validation. Overall, 16,800 LDA classifiers are trained and tested during the optimization loop, which takes ~7 minutes on a standard laptop computer.

Classification performance was evaluated statistically at the individual subject and group levels. For each subject, statistical significance was determined using permutation tests. M1 excitability state labels (high or low) were randomly shuffled and a new classifier using the optimized features and lambda value was trained to discriminate between the resulting shuffled states. Performance of this new classifier was evaluated using AUC values as described above. This process was repeated 10,000 times to obtain a subject-specific null distribution of AUC values. If a given subject's true AUC value exceeded the 95th percentile of the null distribution (analogous to a right-tailed statistical test with alpha = 0.05), above-chance classification was considered to have occurred. At the group level, statistical significance was evaluated by comparing the group-level distribution of true AUC values to the standard chance level for a two-class classification problem (0.5). Other than permutation tests, all statistical testing was performed using right-tailed Wilcoxon signed rank tests. Alpha was equal to 0.05 for all analyses.

To confirm that the optimized personalized classifiers correctly predicted the presence of low and high M1 excitability states, we also compared the amplitudes of MEPs evoked during classifier-predicted low M1 excitability states and classifier-predicted high M1 excitability states at the individual and group levels. At the individual level, MEP amplitudes were compared between predicted low and high M1 excitability states using a left-tailed Mann Whitney U-test. The U-test was used instead of the Wilcoxon signed rank test to account for potentially unequal trial numbers per predicted class. At the group level, MEP amplitudes were averaged within each classifier-predicted state per subject, pooled across subjects, and then compared using a left-tailed Wilcoxon signed rank test. Left-tailed statistical tests were used to test the *a priori* hypothesis that MEPs elicited during classifier-predicted low excitability states would be smaller than those elicited during classifier-predicted high excitability states. Alpha was equal to 0.05 for all analyses.

2.3.4. Generalization of personalized classifiers across subjects. We then tested whether personalized classifiers trained at the individual subject level using features calculated from EEG data immediately preceding TMS delivery (0.505 – 0.005 s before each TMS pulse) generalized across all other subjects. Personalized classifiers with the optimal number of features and optimal lambda value (see *Single trial classification at the individual subject level*) were used to discriminate between low and high M1 excitability states in all other subjects. Generalization performance was assessed using AUC values. Because classifiers were trained on one subject and tested on all others, training and testing sets were always independent and no stratification was used.

Generalization performance was evaluated at the individual subject and group levels. For individual subjects, AUC values reflecting the extent to which each subject's personalized classifier generalized to all other test subjects was compared to chance. At the group level, AUC values were averaged across all 19 test subjects per training subject. This created a distribution of aggregate generalization AUC values, which were then also compared to chance. Finally, aggregate generalization AUC values were compared to AUC values obtained during single-trial classification at the individual level. All statistical testing was performed using right-tailed Wilcoxon signed rank tests. Alpha was equal to 0.05 for all analyses.

2.3.5. Generalization of personalized classifiers across time. To determine if classifiers performed differently when trained and tested using power spectral features obtained from

earlier time windows preceding TMS, we applied the same classification approach described above (see *Single-trial classification at the individual subject level*). Personalized classifiers were instead trained and tested using power spectral features obtained from EEG signals obtained within five different 500 ms windows preceding TMS (3.005 –2.505 s, 2.505 – 2.005 s, 2.005 – 1.505 s, 1.505 – 1.005 s, 1.005 – 0.505 s before TMS). Each subject's optimal number of features and optimal lambda value was used for classification across time points (see *Single trial classification at the individual subject level*). Classification performance at each time point was evaluated statistically by comparing group-level AUC values to chance using right-tailed false discovery rate (FDR)-corrected Wilcoxon signed rank tests (Benjamini and Hochberg 1995). We then used pairwise FDR-corrected Wilcoxon signed rank tests to determine if each time point's AUC values exceeded the AUC values obtained from classifiers trained using EEG features obtained immediately before TMS delivery (0.505 – 0.005 s before TMS). Because results of this analysis revealed above-chance classification performance at all tested time points (see *Results*), we further asked whether the information used to achieve successful classification in earlier time windows (3.005 –2.505 s, 2.505 – 2.005 s, 2.005 – 1.505 s, 1.505 – 1.005 s, 1.005 – 0.505 s before TMS) was the same information used to achieve successful classification immediately preceding TMS (0.505 – 0.005 s before TMS). This was done by testing the extent to which personalized classifiers trained using power spectral features obtained from EEG data immediately preceding TMS generalized to earlier time points. AUC values describing the extent to which each subject's personalized classifier generalized backwards in time were thus calculated (King and Dehaene 2014). Group-level temporal generalization performance was evaluated by comparing AUC values at each time point to AUC values obtained during subject-specific classification using features calculated from EEG data within 0.505 – 0.005 s before TMS using right-tailed FDR-corrected Wilcoxon signed rank tests. Alpha was equal to 0.05 for all analyses.

3. Results

We first determined if personalized classifiers trained using features obtained immediately before each TMS pulse (power spectral features calculated from EEG data within 0.505 to 0.005 s before TMS) successfully discriminated between low and high M1 excitability states. All 20 subjects showed above-chance classification performance ($p < 0.01$ for all). Similarly, group-level classification performance was significantly higher than chance (AUC = 0.63 ± 0.01 [range = 0.56 – 0.72], $p < 0.00005$; Figure 2a). MEP amplitudes elicited during classifier-predicted high excitability states were significantly larger than those elicited during classifier-predicted low

excitability states in all 20 subjects ($p < 0.02$ for all; Figure 2b). At the group level, MEP amplitudes were 24.8 ± 2.19% larger when evoked during classifier-predicted high versus classifier-predicted low excitability states (range = 8.25 – 47.12%, $p < 0.00005$, Figure 2c). The group-level confusion matrix confirmed that accurate classification occurred more often than inaccurate classification (Figure 2d, yellow versus blue cells). However, the features incorporated into each subject's optimized classifier varied substantially, with no features consistently included across all subjects.

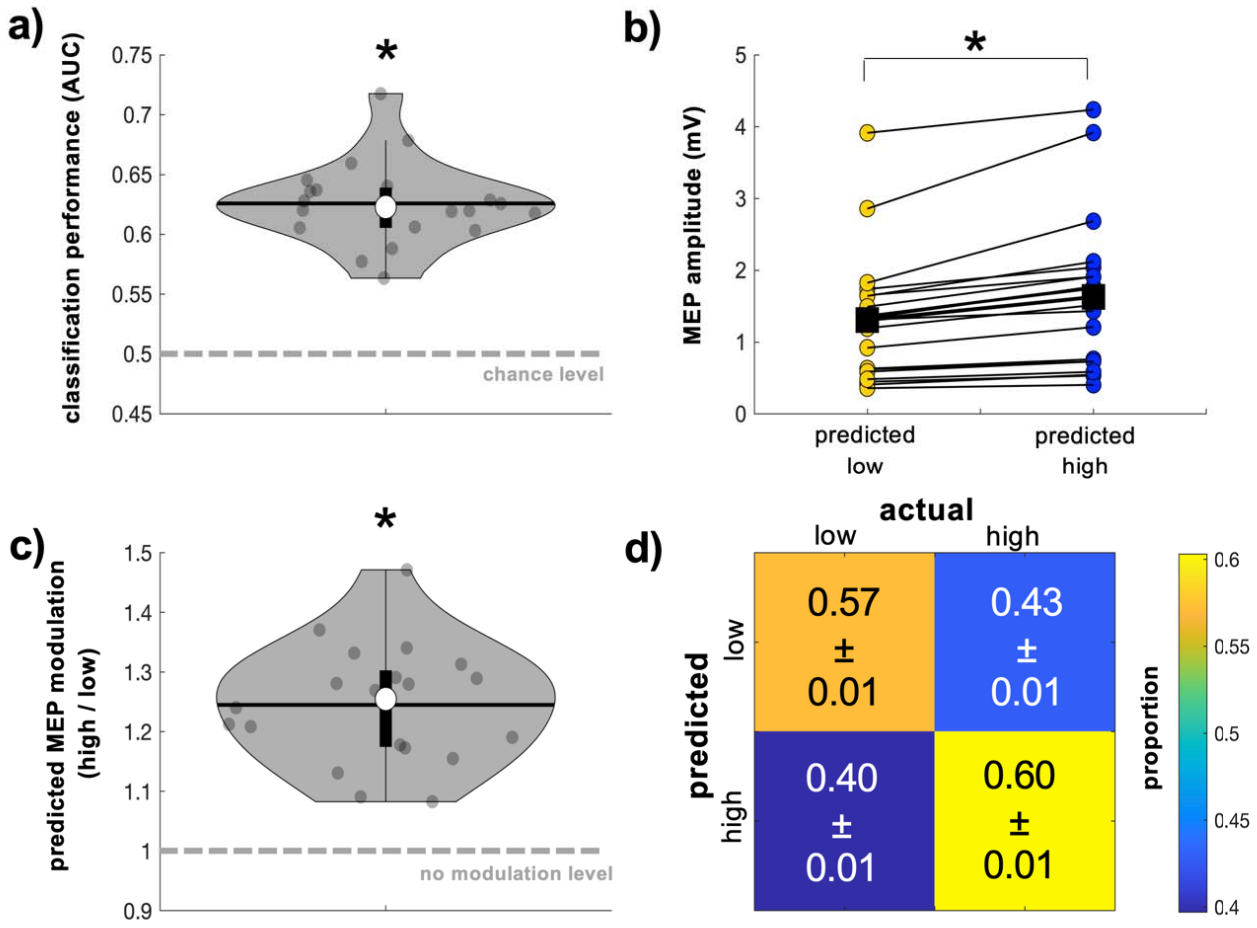

Figure 2. Personalized classification of low and high M1 excitability states. a) AUC values reflecting classification performance. Classifiers were trained at the individual subject level using personalized power spectral features obtained from EEG data immediately preceding TMS (i.e., between 0.505 and 0.005 s before TMS). These classifiers performed significantly above chance in every subject tested and at the group level. The grey dashed line indicates the standard chance level for a two-class classification problem. b) MEP amplitudes elicited during classifier-predicted low (yellow dots) and classifier-predicted high (blue dots) M1 excitability states. MEPs were significantly larger during classifier-predicted high excitability states than classifier-predicted low excitability states in every subject tested (colored dots) and at the group level (black squares). c) Modulation of MEP amplitudes between classifier-predicted high M1

excitability states and classifier-predicted low M1 excitability states. The grey dashed line indicates no MEP amplitude modulation (i.e., the amount of MEP amplitude modulation if MEPs elicited during classifier-predicted high excitability states were equal to those elicited during classifier-predicted low excitability states). d) Group-level confusion matrix describing the proportion of accurate and inaccurate classifications of M1 excitability states. Note that the proportions of true low (upper left) and true high (lower right) classifications exceed the proportions of false low (upper right) and false high (lower left) classifications. Error estimates are SEM. For a and c, grey dots = individual subject AUC values, thick whisker = interquartile range, thin whisker = full range, white dot = group median, black horizontal line = group mean. * indicates group-level statistical significance.

Given that the features used to successfully discriminate between high and low M1 excitability states differed substantially across subjects, we also examined cross-subject generalization of each subject's optimized classifier. Specifically, we tested whether personalized classifiers trained using power spectral features calculated from EEG signals immediately preceding TMS (0.505 to 0.005 s before each TMS pulse) could discriminate between high and low M1 excitability states in all other subjects. Personalized classifiers trained using data from seven of 20 subjects showed above chance performance when tested on all other subjects (see Figure 3a for subject-level generalization matrix, $p < 0.03$), while those classifiers trained using data from the other 13 subjects did not perform above chance when tested on all other subjects ($p > 0.08$). This corresponded to group-level generalization performance AUC values of 0.51 ± 0.003 (range = 0.48 – 0.53, Figure 3b), which significantly exceeded chance ($p < 0.005$). Yet, generalization performance was significantly worse than personalized, subject-specific classification performance ($p < 0.000005$, see also Figure 2a).

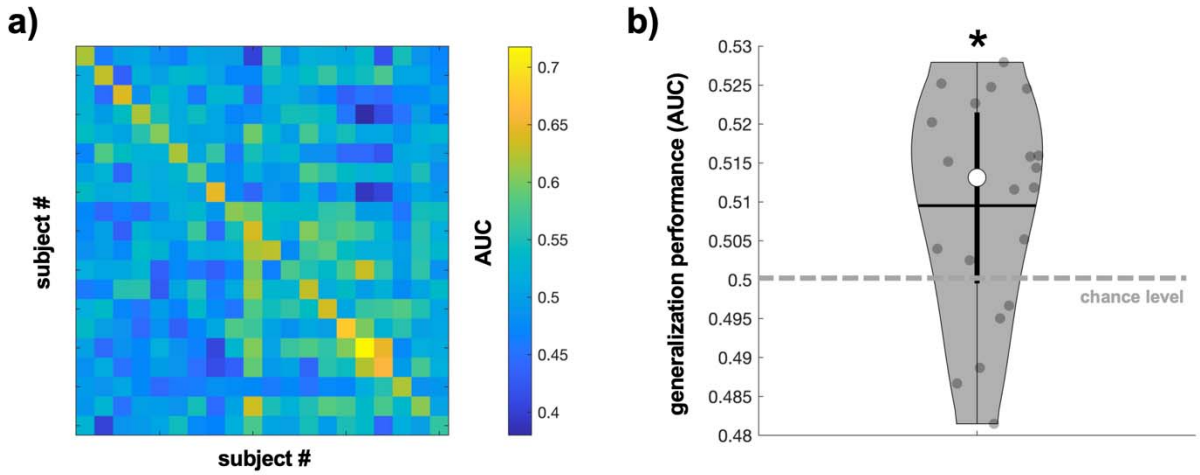

Figure 3. Cross-subject generalization of personalized classifiers. a) Matrix of AUC values reflecting cross-subject generalization performance of classifiers trained using power spectral features obtained from EEG signals immediately preceding TMS (0.505 – 0.005 s before TMS).

Warmer and cooler colors indicate higher and lower AUC values, respectively. AUC values along the diagonal indicate performance of personalized classifiers and thus are identical to those shown in Figure 2a. b) AUC values indicating generalization performance averaged across all testing subjects per training subject. AUC values exceeded chance in seven of 20 subjects and at the group level. For b, grey dots = individual subject AUC values, thick whisker = interquartile range, thin whisker = full range, white dot = group median, black horizontal line = group mean. The grey dotted line indicates the standard chance level for a two-class classification problem. * indicates statistical significance of AUC values when compared to chance.

Finally, we evaluated the temporal evolution of classification performance over 3 seconds leading up to TMS delivery. Personalized classifiers trained and tested on power spectral features extracted from EEG data within each pre-TMS window (3.005 –2.505 s, 2.505 – 2.005 s, 2.005 – 1.505 s, 1.505 – 1.005 s, 1.005 – 0.505 s before TMS) showed above-chance performance (Figure 4a, $p < 0.0002$ for all). However, classifier performance was highest immediately before TMS delivery ($p < 0.0004$ for all). Further, AUC values of personalized classifiers trained using features obtained immediately before TMS delivery (0.505 – 0.005 s before TMS) and tested on each preceding time window also exceeded chance ($p < 0.0002$ for all; Figure 4b). Classification AUC values obtained using features calculated from EEG data immediately preceding TMS delivery (i.e., 0.505 – 0.005 s before TMS) were significantly higher than generalization AUC values obtained from any other time window preceding TMS ($p < 0.0002$ for all).

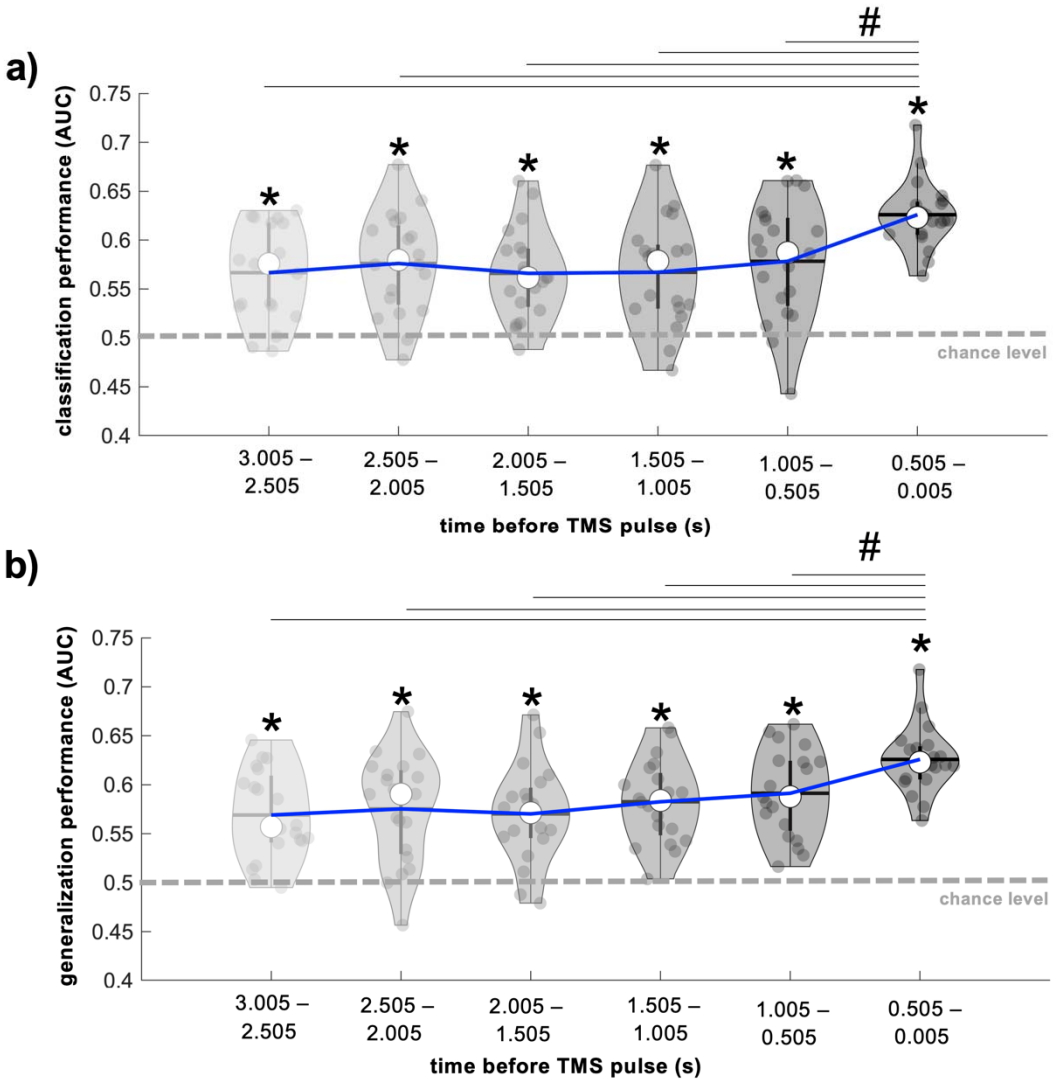

Figure 4. Evaluation of personalized classification performance over time. a) AUC values reflecting performance of classifiers trained and tested using power spectral features obtained from EEG data in different time windows preceding TMS delivery. Classification AUC values for the time point closest to TMS delivery (0.500 – 0.005 s before TMS) were significantly higher than classification AUC values obtained for all other time points. b) AUC values reflecting the extent to which personalized classifiers trained using power spectral features obtained from EEG data immediately before TMS delivery (0.500 – 0.005 s, far right) generalized to earlier time points. Classification AUC values for the time point closest to TMS delivery (0.500 – 0.005 s before TMS) were significantly higher than generalization AUC values for all other time points. Note that for both panels, the data depicted on the right-most violin plot reflects personalized classification AUC values (also depicted in Figure 2a). Time windows used to calculate power spectral features are depicted on the x axes. Blue lines indicate group-level AUC values approaching TMS delivery. Grey dots = individual subject AUC values, thick whisker = interquartile range, thin whisker = full range, white dot = group median, grey/black horizontal lines = group mean. The grey dotted lines indicate the standard chance level for a two-class classification problem. * statistical significance of AUC values when compared to chance. # statistical significance of pairwise comparisons.

4. Discussion

In this study, we used multivariate pattern classification to identify personalized M1 excitability states from human EEG recordings. We found that personalized classifiers successfully discriminated between low and high M1 excitability states in every subject tested. MEPs evoked during classifier-predicted high M1 excitability states were ~25% larger than those evoked during classifier-predicted low M1 excitability states, and personalized classifiers generalized weakly but significantly across subjects. Further, classifiers performed best when subject-specific EEG features obtained immediately before each TMS pulse were used to predict M1 excitability states. Overall, our findings demonstrate that individual subjects exhibit unique brain activity patterns that successfully predict M1 excitability states and that these patterns can be efficiently captured using simple power spectral features obtained from low-resolution EEG recordings.

Historically, TMS has been delivered to the brain irrespective of ongoing brain activity. Early studies showed that TMS protocols using different stimulation rates and burst patterns induce either LTP- or LTD-like plasticity within human M1 (Pascual-Leone et al. 1994; Chen et al. 1997; Stefan et al. 2000; Huang et al. 2005). However, more recent reports suggest that the plastic effects of TMS interventions may depend more strongly on the excitability state at the time of stimulation than the specific pulse rate or pattern (Zrenner et al. 2018; Baur et al. 2020), consistent with findings in animal models (Pavlides et al. 1988; Huerta and Lisman 1993; Huerta and Lisman 1995; Zanos et al. 2018). For example, whether hippocampal stimulation induces LTP or LTD critically depends on the theta phase (which is closely tied to hippocampal excitability, Buzsáki 2002) at the exact moment of stimulation (Pavlides et al. 1998; Huerta and Lisman 1993; Huerta and Lisman 1995). Accumulating evidence thus suggests that TMS protocols are most likely to induce LTP- and LTD-like plasticity when delivered during high and low excitability states, respectively, motivating development of effective and generalizable methods capable of predicting the presence of such states (Kraus et al. 2016; Kraus et al. 2018; Zrenner et al. 2018; Madsen et al. 2019; Shirinpour et al. 2020). Along these lines, current state-of-the-art methods use oscillatory properties that correlate with excitability at the group level (Zrenner et al. 2018; Madsen et al. 2019; Baur et al. 2020; Hussain et al. 2020; Hussain et al. 2021). However, adapting these methods to individual subjects can be challenging (Zrenner et al. 2020), and properly applying them in patients with brain lesions or other forms of structural brain damage raises additional issues (Hussain et al. 2020). Here, we introduce a

straightforward, personalized decoding-based analytical framework that successfully predicts the occurrence of high and low M1 excitability states in individual subjects.

The decoding-based approach described here holds several advantages over other recent brain state detection algorithms (Zrenner et al. 2018; Madsen et al. 2019; Baur et al. 2020; Hussain et al. 2021; Metsomaa et al. 2021). First, our approach is entirely personalized and thus does not require *a priori* knowledge of the relationship between cortical excitability and EEG activity. Instead, it discovers patterns of brain activity that correspond to high and low M1 excitability states using supervised machine learning at the individual level. Compared to "top down" hypothesis-driven studies that investigate the relationship between pre-defined brain states and M1 excitability, our data-driven "bottom up" approach increases the likelihood of successfully detecting high and low excitability states. Indeed, the strength of our method is clearly illustrated by our ability to successfully identify M1 excitability states in every subject tested. To our knowledge, no other personalized M1 excitability state detection algorithm has been successful in 100% of tested subjects. Second, because it is fully personalized, the decoding-based approach presented here in principle does not require any specific cortical morphometry. This characteristic makes it well-suited for brain state-dependent TMS in patients with brain damage that can alter the relationship between neuronal activity and scalp-level EEG signals (i.e., cortical stroke; Baumann et al. 1997; Hussain et al. 2020). Third, our classification approach requires little EEG preprocessing. Personalized classifiers built using highly preprocessed EEG signals (Metsomaa et al. 2021) are not likely to generalize well to noisy real-world EEG signals regardless of their performance during cross-validation. For this reason, we intentionally did not include time-intensive analytical steps like independent components analysis in our preprocessing approach. We therefore expect our method to perform well in a real-time setting. A third advantage is that we used all available trials to build personalized classifiers. In contrast, a recent study built personalized classifiers using only those trials resulting in the most extreme MEP amplitudes (i.e., the largest and smallest; Metsomaa et al. 2021), reporting classification accuracy of ~67%. Although retaining only those trials associated with extreme MEP amplitudes improves performance during cross-validation, classifiers trained in this way may not generalize well to a real-time setting where brain states outside of these extremes are common. Consistent with this, we found that retaining only those trials associated with extreme MEP amplitudes in our dataset artificially boosted classifier performance (AUC = ~ 0.69 ± 0.02, range = 0.60 – 0.83).

In addition to successfully discriminating between low and high M1 excitability states in 100% of tested subjects, the brain states identified here lasted for seconds at a time. Specifically, our results showed that EEG data obtained up to 3 seconds before TMS delivery successfully discriminated between low and high M1 excitability states. This result has clear practical implications for brain state-dependent TMS: because these brain states are relatively sustained over time, TMS could in the future simply be triggered upon their detection. This triggering-based approach contrasts starkly with computationally-intensive phase-dependent TMS algorithms that use autoregressive forward prediction to estimate the instantaneous oscillatory phase at some future time point (Zrenner et al. 2018; Madsen et al. 2019; Baur et al. 2020; Shirinpour et al. 2020; Hussain et al. 2020; Hussain et al. 2021) and are much easier to implement (see Bergmann et al. 2012). Further, the MEP modulation between predicted high and low M1 excitability states seen here (~25%) closely matches that observed when using phase-dependent M1 TMS (Zrenner et al. 2018; Hussain et al. 2019a; Bergmann et al. 2019; Baur et al. 2020), suggesting similar efficacy despite a reduced computational load.

Personalized classifiers generalized significantly but weakly across subjects. Despite group-level statistical significance, significant generalization occurred in only 35% of subjects and only exceeded chance at the group level by a very small amount (mean cross-subject generalization AUC = ~ 0.51, chance level = 0.50). Further, classification at the individual subject level significantly outperformed cross-subject generalization, indicating the importance of personalized EEG features when predicting M1 excitability levels. This notion is further supported by the fact that each subject's optimized features varied substantially across subjects, with no features being consistently included in all subjects' optimized classifiers. Previous work showed that personalized features can explain up to half of the variance in time-varying brain activity (Gratton et al. 2018) and can reliably identify individual subjects (Finn et al. 2015). Given these prior reports, it is perhaps not surprising that personalized brain activity patterns also predict M1 excitability states. Still, whether the brain states identified here are stable across multiple recording sessions has yet to be established. Determining whether personalized classifiers successfully generalize across multiple sessions will be important for multi-session decoding-based brain state-dependent neuromodulation.

Although the personalized classifiers used here performed above chance in every subject tested, their overall performance is lower than classifiers used for other real-time EEG analysis applications, like brain-computer interfaces (Blankertz et al. 2002). However, despite the modest

performance levels achieved here, classifier-predicted high and low M1 excitability states were associated with an ~25% difference in MEP amplitudes. Previous brain state-dependent TMS studies showed that applying TMS during high excitability brain states during which TMS evokes MEPs that are ~14-36% larger than those evoked during low excitability states substantially enhances the plastic effects of TMS interventions (Zrenner et al. 2018; Baur et al. 2020; Hussain et al. 2021). Based on these findings, we expect that translating our approach to a real-time setting will be sufficient to strengthen the effects of M1 TMS interventions. Regardless, future work should focus on further optimizing classifier performance for personalized brain state-dependent TMS applications to maximize their effects.

In conclusion, we report that individual subjects exhibit unique brain activity patterns that predict M1 excitability levels. Crucially, these patterns were successfully decoded in every subject tested using only simple power spectral features obtained from low-density EEG signals. MEP amplitudes elicited during classifier-predicted high M1 excitability states were also ~25% larger than those elicited during classifier-predicted low M1 excitability states. Finally, these brain states lasted for seconds at a time, making them ideal targets for efficient brain state-dependent TMS. Overall, the approach presented here constitutes a key step towards efficient, flexible, and fully personalized decoding-based brain state-dependent TMS in the future.

References

1. Baumann SB, Wozny DR, Kelly SK, Meno FM. (1997). The electrical conductivity of human cerebrospinal fluid at body temperature. *IEEE Trans Biomed Eng* 44(3):220-223.
2. Baur D, Galevska D, Hussain S, Cohen LG, Ziemann U, Zrenner C. (2020). Induction of LTD-like corticospinal plasticity by low-frequency rTMS depends on pre-stimulus phase of sensorimotor μ-rhythm. *Brain Stimul* 13(6):1580-1587.
3. Benjamini Y, Hochberg Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J Royal Stat Soc: Methodological* 57(1):289-300.
4. Berger B, Minarik T, Liuzzi G, Hummel FC, Sauseng P. (2014). EEG oscillatory phase-dependent markers of corticospinal excitability in the resting brain. *BioMed Res Intl* 2014:936096.
5. Bergmann TO, Lieb A, Zrenner C, Ziemann U. (2019). Pulsed facilitation of corticospinal excitability by the sensorimotor μ-alpha rhythm. *J Neurosci* 39(50):10034-10043.
6. Bergmann TO, Mölle M, Schmidt MA, Lindner, Marshall L, Born J, Siebner HR. (2012). EEG-guided transcranial magnetic stimulation reveals rapid shifts in motor cortical excitability during the human sleep slow oscillation. *J Neurosci* 32(1):243-253.
7. Blankertz B, Curo G, Müller K-R. (2002). Classifying single-trial EEG: towards brain computer interfacing. *Adv Neural Inf Proc Sys*, 157-164.
8. Buzsáki G. (2002). Theta oscillations in the hippocampus. *Neuron* 33(3):325-340.
9. Chen RM, Classen J, Gerloff C, Celnik P, Wassermann E, Hallett M, Cohen LG. (1997). Depression of M1 excitability by low-frequency transcranial magnetic stimulation. *Neurology* 48(5):1398-1403.
10. Ekhtiari H, Tavakoli H, Addolorato G, Baeken C, Bonci A, Campanella S, Castelo-Branco L, Challet-Bouju G, Clark VP, Claus E, Dannon PN, Felice AD, den Uyl T, Diana M, di Giannantonio M, Fedota JR, Fitzgerald P, Gallimberti L, Hanlon CA. (2019). Transcranial electrical and magnetic stimulation (tES and TMS) for addiction medicine: a consensus paper on the present state of the science and the road ahead. *Neurosci Biobehav Rev* 104:118-140.
11. Finn ES, Shen X, Scheinost D, Rosenberg MD, Huang J, Chun MM, Papademetris X, Constable TR. (2015). Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nat Neurosci* 18(11):1664-1671.
12. George MS, Wassermann EM, Williams WA, Callahan A, Ketter TA, Basser P, Hallett M, Post RM. (1995). Daily repetitive transcranial magnetic stimulation (rTMS) improves mood in depression. *NeuroReport* 6(14):1853-1856.

13. Gorelick DA, Zangen A, George MS. (2014). Transcranial magnetic stimulation (TMS) in the treatment of substance addiction. *Ann New York Acad Sci* 1327(1):79.
14. Gratton C, Laumann TO, Nielsen AN, Greene DJ, Gordon EM, Gilmore AW, Nelson SM, Coalson RS, Snyder AZ, Schalggar BL, Dosenbach NU. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron* 98(2):439-452.
15. Goldsworthy MR, Hordacre B, Rothwell JC. (2021). Effects of rTMS on the brain: is there value in variability? *Cortex* 139:43-59.
16. Haegens S, Nácher V, Luna R, Romo R, Jensen O. (2011). α-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proc Natl Acad Sci* 108(48):19377-19382.
17. Hanlon CA, Dowdle LT, Henderson JS. (2018). Modulating neural circuits with transcranial magnetic stimulation: implications for addiction treatment development. *Pharmacol Rev* 70(3):661-683.
18. Huang YZ, Edwards MJ, Rounis E, Bhatia KP, Rothwell JC (2005). Theta burst stimulation of the human M1. *Neuron* 45(2):201-206.
19. Hummel FC, Cohen LG. (2006). Non-invasive brain stimulation: a new strategy to improve neurorehabilitation after stroke? *The Lancet Neurology* 5(8):708-712.
20. Huerta PT, Lisman JE. (1993). Heightened synaptic plasticity of hippocampal CA1 neurons during a cholingergically induced rhythmic state. *Nature* 364(6439):723-725.
21. Huerta PT, Lisman JE. (1995). Bidirectionally synaptic plasticity induced by a single burst during cholingergic theta oscillation in CA1 in vitro. *Neuron* 15(5):1053-1063.
22. Hussain SJ, Claudino L, Bönstrup M, Norato G, Cruciani G, Thompson R, Zrenner C, Ziemann U, Buch ER, Cohen LG. (2019a). Sensorimotor oscillatory phase–power interaction gates resting human corticospinal output. *Cereb Cortex* 29(9):3766-3777.
23. Hussain, SJ, Cohen LG, Bönstrup M. (2019b). Beta rhythm events predict corticospinal motor output. *Sci Rep* 9(1):1-10.
24. Hussain SJ, Hayward W, Fourcand F, Zrenner C, Ziemann U, Buch ER, Cohen LG. (2020). Phase-dependent transcranial magnetic stimulation of the lesioned hemisphere is accurate after stroke. *Brain Stimul* 13(5):1354-1357.
25. Hussain SJ, Vollmer MK, Stimely J, Norato G, Zrenner CZ, Ziemann U, Buch ER, Cohen LG. (2021). Phase-dependent offline enhancement of human motor memory. *Brain Stimul* 14(4):873-883.

26. Kraus D, Naros G, Bauer R, Leão MT, Ziemann U, Gharabaghi A. (2016). Brain–robot interface driven plasticity: distributed modulation of corticospinal excitability. *NeuroImage* 125*:*522-532.
27. Kraus D, Naros G, Guggenberger R, Leão MT, Ziemann U, Gharabaghi A. (2018). Recruitment of additional corticospinal pathways in the human brain with state-dependent paired associative stimulation. *J Neurosci* 38(6):1396-1470.
28. King JR, Dehaene S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method: *Trends Cog Sci* 18(4):203-210.
29. López-Alonso V, Cheeran B, Río-Rodríguez D, Fernández-del-Olmo M. (2014). Inter-individual variability in response to non-invasive brain stimulation paradigms. *Brain Stimul* 7(3):372-380
30. Madsen KH, Karabanov AN, Krohne LG, Safeldt MG, Tomasevic L, Siebner HR. (2019). No trace of phase: Corticomotor excitability is not tuned by phase of pericentral mu-rhythm. *Brain Stimul* 12(5):1261-1270.
31. Metsomaa J, Belardinelli P, Ermolova M, Ziemann U, Zrenner C. (2021). Causal decoding of individual cortical excitability states. *NeuroImage,* 245:118652.
32. Nilakantan AS, Mesulam MM, Weintraub S, Karp EL, VanHaerents S, Voss JL. (2019). Network-targeted stimulation engages neurobehavioral hallmarks of age-related memory decline. *Neurology* 92(20):e2349-e2354.
33. O'Reardon JP, Solvason HB, Janicak PG, Sampson S, Isenberg KE, Nahas Z, McDonald WM, Avery D, Fitzgerald PB, Loo C, Demitrack MA, George MS, Sackeim H. A. (2007). Efficacy and safety of transcranial magnetic stimulation in the acute treatment of major depression: a multisite randomized controlled trial. *Biol Psychiatry* 62(11): 1208-1216.
34. Oostenveld R, Fries P, Maris E, Schoffelen JM. (2011). FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Compu Intell Neurosci* 2011:156869.
35. Pascual-Leone A, Valls-Solé J, Wassermann EM, Hallett M. (1994). Responses to rapid-rate transcranial magnetic stimulation of the human M1. *Brain* 117(4):847-858.
36. Pavlides C, Greenstein YJ, Grudman M, Winson J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ-rhythm. *Brain Res* 439(1-2):383-387.

37. Sauseng P, Klimesch W, Gerloff C, Hummel FC. (2009). Spontaneous locally restricted EEG alpha activity determines cortical excitability in M1. *Neuropsychologia* 47(1):284-288.
38. Shirinpour S, Alekseichuk I, Mantell K, Opitz A. (2020). Experimental evaluation of methods for real-time EEG phase-specific transcranial magnetic stimulation. *J Neural Eng* 17(4):046002.
39. Stefan K, Kunesch E, Cohen LG, Benecke R, Classen J. (2000). Induction of plasticity in the human M1 by paired associative stimulation. *Brain* 123(3):572-584.
40. Tazoe T, Perez MA. (2015). Effects of repetitive transcranial magnetic stimulation on recovery of function after spinal cord injury. *Arch Phys Med Rehabil* 96(4):S145-S155.
41. Treder MS. (2020). MVPA-Light: a classification and regression toolbox for multi-dimensional data. *Front Neurosci* https://doi.org/10.3389/fnins.2020.00289
42. Wang JX, Rogers LM, Gross EZ, Ryals AJ, Dokucu ME, Brandstatt KL, Hermiller MS, Voss JL. (2014). Targeted enhancement of cortical-hippocampal brain networks and associative memory. *Science* 345(6200):1054-1057.
43. Webster BR, Celnik PA, Cohen LG. (2006). Noninvasive brain stimulation in stroke rehabilitation. *NeuroRx* 3(4):474-481.
44. Zanos S, Rembado I, Chen D, Fetz EE. (2018). Phase-locked stimulation during cortical beta oscillations produces bidirectional synaptic plasticity in awake monkeys. *Curr Biol* 28(16):2515-2526.
45. Ziemann U, Siebner HR. (2015). Inter-subject and inter-session variability of plasticity induction by non-invasive brain stimulation: boon or bane? *Brain Stimul* 8(3):662-663.
46. Zrenner C, Desideri D, Belardinelli P, Ziemann U. (2018). Real-time EEG-defined excitability states determine efficacy of TMS-induced plasticity in human M1. *Brain Stimul* 11(2):374-389.
47. Zrenner C, Galveska D, Nieminen JO, Baur D, Stefanou MI, Ziemann U. (2020). The shaky ground truth of real-time phase estimation. *NeuroImage* 214:116761.