

Mapping Motor Cortex Stimulation to Muscle Responses: A Deep Neural Network Modeling Approach

Md Navid Akbar
makbar@ece.neu.edu
Northeastern University
Boston, MA, USA

Marc A. Sommer
marc.sommer@duke.edu
Duke University
Durham, NC, USA

Dana Brooks
brooks@ece.neu.edu
Northeastern University
Boston, MA, USA

Mathew Yarossi
m.yarossi@northeastern.edu
Northeastern University
Boston, MA, USA

Moritz Dannhauer
moritz.dannhauer@duke.edu
Duke University
Durham, NC, USA

Eugene Tunik
e.tunik@neu.edu
Northeastern University
Boston, MA, USA

Marc Martinez-Gost
marc.martinez.gost@estudiant.upc.edu
Technical University of Catalonia
Barcelona, USA

Sumientra Rampersad
s.rampersad@northeastern.edu
Northeastern University
Boston, MA, USA

Deniz Erdoğan
erdogmus@ece.neu.edu
Northeastern University
Boston, MA, USA

ABSTRACT

A deep neural network (DNN) that can reliably model muscle responses from corresponding brain stimulation has the potential to increase knowledge of coordinated motor control for numerous basic science and applied use cases. Such cases include the understanding of abnormal movement patterns due to neurological injury from stroke, and stimulation based interventions for neurological recovery such as paired associative stimulation. In this work, potential DNN models are explored and the one with the minimum squared errors is recommended for the optimal performance of the M2M-Net, a network that maps transcranial magnetic stimulation of the motor cortex to corresponding muscle responses, using: a finite element simulation, an empirical neural response profile, a convolutional autoencoder, a separate deep network mapper, and recordings of multi-muscle activation. We discuss the rationale behind the different modeling approaches and architectures, and contrast their results. Additionally, to obtain a comparative insight of the trade-off between complexity and performance analysis, we explore different techniques, including the extension of two classical information criteria for M2M-Net. Finally, we find that the model analogous to mapping the motor cortex stimulation to a combination of direct and synergistic connection to the muscles performs the best, when the neural response profile is used at the input.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; Supervised learning by regression; Cross-validation; *Model verification and validation*; • **Hardware** → Neural systems.

KEYWORDS

Transcranial magnetic stimulation (TMS), muscle synergy, convolutional neural network (CNN), performance-complexity analysis, autoencoder.

ACM Reference Format:

Md Navid Akbar, Mathew Yarossi, Marc Martinez-Gost, Marc A. Sommer, Moritz Dannhauer, Sumientra Rampersad, Dana Brooks, Eugene Tunik, and Deniz Erdoğan. 2020. Mapping Motor Cortex Stimulation to Muscle Responses: A Deep Neural Network Modeling Approach. In *The 13th Pervasive Technologies Related to Assistive Environments Conference (PETRA '20)*, June 30-July 3, 2020, Corfu, Greece. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3389189.3389203>

1 INTRODUCTION

Coordinated activation of groups of muscles for voluntary movement is a prominent feature of human motor control [2]. Abnormal coordination of muscles is a marker of physical impairment in numerous congenital and acquired disease states such as cerebral palsy, stroke, and Parkinson's disease. To optimize the use of neurotechnology and rehabilitation to improve impaired function, we must first gain a better understanding of cortical control of healthy multi-muscle coordination [15]. Transcranial magnetic stimulation (TMS) describes a non-invasive procedure that uses magnetic fields to stimulate nerve cells in the brain [7]. When applied to the motor cortex, TMS produces motor evoked potentials (MEPs) in a spatially selective set of muscles that can be recorded using standard surface electromyography (EMG). TMS, therefore, provides the best non-invasive approach to stimulation-recording for the study of motor control in humans. A deeper mechanistic understanding of the extent to which TMS can induce coordinated muscle activation is essential for innovation in TMS use as biomarker of disease or a means of intervention. For example, use of TMS as a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PETRA '20, June 30-July 3, 2020, Corfu, Greece

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7773-7/20/06...\$15.00

<https://doi.org/10.1145/3389189.3389203>

diagnostic tool to identify cortical motor topography associated with abnormal muscle patterns may help identify who may benefit most from specific interventions following stroke, and track cortical changes associated with recovery [24]. Used as a tool for neurological intervention, TMS synchronized with electrical peripheral nerve stimulation, has been demonstrated to strengthen neural connection and induce partial restoration of muscle activity [19]. Thus, TMS has the ability to produce positive plastic changes in the corticospinal tract of stroke-induced paralysis patients [8, 11]. Such interventions could possibly be enhanced if precise TMS coil placement to stimulate a specific muscle group could be paired with multi-muscle electrical stimulation to enhance the functional grouping of muscles. A major challenge to both of the aforementioned use cases is to reliably locate the region of interest in the brain's motor cortex that maps to the desired muscle grouping. To address this challenge, we introduce deep learning in the modeling scenario.

Convolutional neural networks (CNNs) have already been demonstrated to outperform classical methods, such as support vector machines, for classification of EMG signals [1]. Moreover, a CNN autoencoder has the capacity to successfully learn biologically plausible features [13]. Our group has also previously described a preliminary framework for investigating cortical control of multi-muscle activation using TMS and deep learning, using a single stimulus intensity.

In this work, we first present the mathematical model of our system. Subsequently, we methodically investigate different CNN architectures and a rigorous hyper-parameter fitting scheme to create a model that is more broadly generalizable, incorporating trials from four different stimulus intensities. We call this model M2M-Net (motor cortex to muscle network). We then recommend the best model among those tested, for M2M-Net, and discuss the likely reasons governing the differences in performances. Finally, for assessing the trade-off between complexity and performance, we explore different statistical techniques, including extension of the Bayesian information criteria (BIC) and the Akaike information criteria corrected (AICc) for M2M-Net.

2 EXPERIMENTAL METHODS

2.1 Data Acquisition

A 30 year old, right-handed, healthy male, eligible for TMS, participated following informed consent. Prior to TMS, a T1-weighted image (TI=1100ms, TE=2.63ms, TR=2000ms, 256x192x160 acquisition matrix, FOV=256x192mm, 1mm³ voxels) was used for neuronavigation (Brainsight, Rogue Research). The procedure used for TMS mapping has been described by our group elsewhere [25]. Briefly, the subject was seated, with forearms supported. Surface EMG (Trigno, Delsys, 2kHz) was recorded from 15 hand-arm muscles during TMS. The TMS coil (Magstim 200, 70mm figure-8) was held tangential to the scalp with the handle posterior and 45° to midline. The right 1st dorsal interosseus (FDI) muscle hotspot was found via a coarse map of the hand knob area. Peak-to-peak EMG amplitude 20-50msec after the TMS pulse was the outcome variable (Matlab, The Mathworks). TMS intensity during mapping in proportion to the resting motor threshold (RMT), which was the minimum intensity required to elicit MEPs >50 μ V on 3/6 consecutive trials. TMS




Figure 1: Normalized E-field distributions in the brain. (A) The E-field generated on the brain as a result of stimulation by the TMS coil. (B) The high contrast intensity observed after application of the E-field to neural firing transfer function. (C) Mask of the motor cortex.

(100-300 stimuli, 4sec ISI) was delivered over a 7cm² area centered on the hotspot using stimulus intensities of 110%, 120%, 130% and 140% of RMT. For each intensity, one stimulus was delivered to each of 49 equidistant grid points, and the remaining stimuli were delivered using real time feedback of MEPs to maximize information about responsive areas [16].

2.2 Model Components

We have previously described a forward model comprised of four components to predict TMS-evoked multi-muscle activation based on stimulus parameters alone [25].

- (1) A finite element (FE) model that generates, from TMS parameters (coil geometry/position/orientation, magnetic pulse characterization, stimulus intensity), a volumetric E-field on a mesh that accurately represents tissue segmentation of the subject specific MRI [4, 18].
- (2) A nonlinear mapping of local E-fields to the expected neural ensemble firing rates, based on nonhuman primate experimental data. This transformation uses a sigmoidal fit, to model the dose-response curve for putative excitatory neurons that may project to spinal circuits [6, 14].
- (3) A two-stage neural network model we call M2M-Net that maps either Step 1 or Step 2 component outputs to multi-muscle responses. Model selection for the M2M-Net is the main subject of this paper.
- (4) A low dimensional linear structure, which we refer to as Synergy (to contrast with MEP), representing synergistic muscle activation in coherent groups, extracted using non-negative matrix factorization [12, 23] applied to recorded MEPs. Essentially, they function as a low-rank non-negative approximation of the MEP data.

Fig. 1 illustrates Steps 1 and 2.

2.3 System Model

The mathematical model of the system for mapping directly to the MEPs (Direct) is given by

$$\mathbf{m}_i = f_D(\mathbf{X}_i) + \epsilon_{D,i}, \quad (1)$$

where $\mathbf{m}_i$ is the $k \times 1$ ($k = 15$ muscles) observed muscle activity vector for the i -th trial, $\mathbf{X}_i \in \mathbb{R}^{l \times l \times l}$ ($l = 64$) is the input 3D matrix

corresponding to the E-field distribution or the neural firing rate, $f_D()$ represents M2M-Net for the Direct connection, $\epsilon_{D,i}$ is the $k \times 1$ residual mapping error, and $i \in 1, 2, \dots, N$ ($N = 629$ is total number of trials for training, from the four different stimulus intensities).

In the case of mapping to the synergies (Synergy), the muscle activation vector can also be expressed as follows

$$\mathbf{m}_i = \tilde{\mathbf{B}}\mathbf{a}_i + \mathbf{w}_i, \quad (2)$$

where $\tilde{\mathbf{B}} \in \mathbb{R}^{k \times r}$ ($r = 9$ synergies) is the unknown basis matrix, in which each row indicates the relative activation level of each muscle in all the synergies, $\mathbf{a}_i$ is the $r \times 1$ unknown activation vector that linearly combines synergy bases, and $\mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ is a spatiotemporal white additive Gaussian noise. Using the maximum likelihood (ML) criteria and non-negative matrix factorization (NMF) technique as described in [25], we estimate $\mathbf{B}$ to give the following relation

$$\mathbf{M} \approx \mathbf{B}\mathbf{A}, \quad (3)$$

where $\mathbf{M} \in \mathbb{R}^{k \times N}$ is the matrix of the all the muscle activations given by $\mathbf{M} = [\mathbf{m}_1 \ \mathbf{m}_2 \ \dots \ \mathbf{m}_N]$, and $\mathbf{A} \in \mathbb{R}^{r \times N}$ is the matrix of all the activation vectors given by $\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_N]$. Finally, our Synergy model takes the form

$$\mathbf{m}_i = \mathbf{B}f_S(\mathbf{X}_i) + \epsilon_{S,i}, \quad (4)$$

where $f_S()$ represents M2M-Net for the Synergy connection and $\epsilon_{S,i}$ is the $k \times 1$ residual mapping error.

For the third case, we first train M2M-Net on the activation matrix representing the synergies, given by $\mathbf{A}$. Afterwards, we use the matrix $\xi_S \in \mathbb{R}^{k \times N}$ of all the residual errors from the Synergy model, given by $\xi_S = [\epsilon_{S,1} \ \epsilon_{S,2} \ \dots \ \epsilon_{S,N}]$, to train a new Direct connection, $f_D()$. This is given by

$$\epsilon_{S,i} = f_D^*(\mathbf{X}_i) + \epsilon_{B,i}, \quad (5)$$

where $\epsilon_{B,i}$ is the $k \times 1$ residual mapping error. Lastly, the predictions obtained from the Synergy and the new Direct connections are then added to give the combinatorial model (Both)

$$\mathbf{m}_i = \mathbf{B}f_S(\mathbf{X}_i) + f_D^*(\mathbf{X}_i) + \epsilon_{B,i}. \quad (6)$$

2.4 Network Training and Testing Procedures

M2M-Net is a hybrid deep network, composed of a convolutional autoencoder used for dimensionality reduction and a separate deep CNN mapper, as seen in Fig. 2. For each stimulation, the 64-cube three-dimensional E-fields, or ensemble firing rates, serve as the input to the autoencoder. The autoencoder provides low-dimensional, 16-cube representations of the input, which are then passed on to the mapper. The mapper contains a sequence of convolutional layers. The output from these layers are then flattened and connected to fully connected dense layers. The structure of the last dense layer determines if outputs will be mapped via Direct, via Synergy, or via Both.

We tested two common techniques of preprocessing methods applied to both the E-field and neural firing rate inputs: mean standardization and min-max scaling. Both of these methods have been shown to speed up the training process, as well as improve prediction accuracy [17]. For the output, activations for the individual muscles were min-max scaled to the unit interval [0,1].

Each set of inputs (E-field distributions or neural firing rates) and associated outputs (muscle activity or synergy activations) are

split into calibration and test sets, in a 9:1 ratio. The calibration set is further split into training and validation sets. A 10-fold cross validation scheme is used to tune hyperparameters, within the calibration set. The lowest normalized root mean squared error (NRMSE) performance on the validation set determines the best choice for the hyperparameters. Once the choice of hyperparameters are decided, a total of 12 different models are trained using the entire calibration set, and their respective performances are assessed on the testing set. This process is repeated three times, each time randomly selecting different subsets of the entire data for calibration and test. The NRMSE of a test set of length T ($T = 70$) is calculated as

$$\text{NRMSE} = \sqrt{\frac{\sum_{j=1}^T \|\epsilon_j\|_2^2}{\sum_{j=1}^T \|\mathbf{m}_j\|_2^2}}, \quad (7)$$

where $\|\cdot\|_2$ denotes the Euclidean norm.

During training, the inputs are processed in mini-batches of size 32. Stochastic gradient descent with momentum is selected as the optimizer for the autoencoder, while Adadelta is chosen for the mapper. A dynamic reduction in learning rate is scheduled with the increase in epoch, for each optimizer. Further details of the network layers and parameters are outlined in Section 2.5.

We select NRMSE as the performance evaluation metric for our 12 models. Later, we will also illustrate the variation of NRMSE with the number of parameters in the mapper, for the models of interest. However, to attain a more meaningful insight, we will also attempt to extend two classical statistical techniques: the Bayesian information criteria (BIC) and the Akaike information criteria corrected (AICc), for M2M-Net.

Following the model outlined in [22], our muscle activation vector may be represented in the form

$$\mathbf{m}_i = \mu(\gamma_i) + \epsilon_i, \quad (8)$$

where γ_i is an unknown parameter vector, and ϵ_i is now assumed to be Gaussian with mean zero and covariance matrix given by $E\{\epsilon\epsilon^T\} = \sigma^2 \mathbf{I}$. The ML estimate of σ^2 for N trials then becomes

$$\hat{\sigma}^2 = \frac{1}{Nk} \sum_{i=1}^N \|\mathbf{m}_i - \mu(\gamma_i)\|_2^2. \quad (9)$$

The corresponding value of log-likelihood of the probability density function of $\mathbf{M}$ then becomes $(Nk \log \hat{\sigma}^2 + \text{constant})$. For BIC and AICc, each expression includes this value of the log-likelihood (the constant is trivial) plus a characteristic penalty term involving the number of free parameters, p , to be estimated (in the mapper) [9, 22]. We can thus express

$$\text{BIC} = (Nk) \log \hat{\sigma}^2 + p \log(Nk), \quad (10)$$

and

$$\text{AICc} = (Nk) \log \hat{\sigma}^2 + \frac{Nk(p + Nk)}{Nk - p - 2}. \quad (11)$$

2.5 Layers and Parameters

In this section, we describe the methodology of parameter selection for our network shown in Fig. 2.

Convolution Layers: M2M-Net contains two sets of convolution-activation-maxpool layers in the encoder portion of the autoencoder, and three such sets in the 3-layer mapper. For the 6-layer




Figure 2: Block diagram of the system model.

mapper, two identical convolution layers are used before each activation layer. Going back to the autoencoder, there are two sets of symmetric convolution-activation-upsample layers in its decoder portion. The number of filters in subsequent convolution layers are doubled: from 32 to 64 for the encoder of the autoencoder, and from 16 to 64 for the mapper. For the decoder, it is reversed. This choice is aligned with the established understanding that while low-level filters learn basic shapes, high-level filters are capable of identifying more sophisticated patterns [20]. The convolutional layers in M2M-Net have $3 \times 3 \times 3$ filters, except the ones used to prepare the intermediate 16-cube representation and the final 64-cube reconstruction. Both of these layers have a $1 \times 1 \times 1$ filter window, and a single channel. The padding used is 1 element on each side, and the stride is 1.

Activation layers: Activation functions follow all convolutional and dense layers. The exponential linear unit (eLU) is the activation function of choice, after every convolutional layer in the autoencoder. However, for the mapper, the rectified linear unit (ReLU) is chosen as it outperformed the eLU in our cross validation rounds. The final activation function in the autoencoder is a linear function. In the mapper, it is a sigmoid for connections to the MEPs (Direct), or that to the synergies (Synergy): since the output should be constrained between 0 and 1. When mapping to the differences between the synergy activations and the ground truth via Both, the result could contain negative values as well. Thus, the second Direct connection to the residuals, indicated earlier in (5), have linear activations in the intermediate layers, and a tanh as the final activation, to constraint the output between -1 and 1.

Pooling, upsampling, and dense layers: Max-pooling and upsampling layers are used to reduce and increase the sizes of the representations, respectively. All filter windows for these layers have a common size of $2 \times 2 \times 2$, and a stride of 1. The final multi-dimensional tensor in the mapper is flattened to a one-dimensional vector, which is then fully connected to the dense layers.

Regularization: Each convolutional layer in both network is fitted with ℓ_1 weight regularization. Our input E-field distribution data are sparse. It is a well known technique to use ℓ_1 penalty in such convolutional sparse coding [3]. Batch normalization layers are used to prevent internal covariance shifts in the data [10], and follow the convolution layers in the mapper. Each dense layer is followed by a dropout layer, which helps counter against over-fitting [21]. The values of the hyper-parameters for both the dropout layers and the ℓ_1 penalty are chosen from the cross validation rounds performed within the calibration set, and are found to be 0.15 and $1e-4$, respectively.

3 RESULTS

From our experiments, min-max scaling across the entire input dataset proved to be the better of the two preprocessing techniques. It provided more distinct intermediate representations and lowered prediction errors. This scaling was performed once again on the entire set of intermediate representations, before they form input to the mapper.

Table 1 presents the arithmetic means and standard deviations of the performance evaluations for the different models. Consistent with our previous results from [25] for the TMS intensity at 110% RMT, we saw that the generalized model across four different intensities also performed better with the E-field to neural firing rate

Table 1: NRMSE values across different models.

	3 conv layer Mapper			6 conv layer Mapper		
	Direct	Synergy	Both	Direct	Synergy	Both
Raw E-field Input	0.658±0.042	0.671±0.049	0.661±0.048	0.644±0.058	0.638±0.053	0.630±0.056
Neural Firing Filter	0.519±0.048	0.505±0.041	0.500±0.043	0.518±0.045	0.516±0.049	0.508±0.049

conversion, compared to E-fields used directly as input. This could be explained in light of the neural firing rate conversion effectively thresholding and increasing the image contrast. Since the images are now more focused in the regions of interest, feature extraction becomes easier, and the model performs better. Examining different mapper combinations corresponding to neural firing, we observe that the prediction by mapping to Synergy outperformed mapping via Direct. Model accuracy was the greatest using the Both model. We also observed that the 3 layer mapper performed better compared to the 6 layer mapper. This could be a case of model complexity and overfitting, as the 6 layer models have more trainable parameters, compared to their 3 layer counterparts. Results from one experiment, out of the three conducted, is shown for the 3 layer mapper in Fig. 3.

With the models using the E-fields directly as input, we notice that the 6 layer models outperform their corresponding 3 layer counterparts. This time, since no contrast enhancing transfer function was used for these models, the network needed to identify the important zones entirely by itself. A model with more degrees of freedom, related to its number of free parameters [5], would be inclined to perform better in such a scenario, similar to what is observed. Hence, the more complex models provide the better predictions.

When comparing Fig. 3 and the row corresponding to the neural firing input in Table 1, we note that the means and the deviations appear rather high. This could be explained if we take into account the many stimulations that produce null or low muscle activations. In accordance to (7), we realize that even if the predictions by M2M-Net result in a small ϵ_j , the corresponding m_j is smaller, or even zero. As a consequence, the overall NRMSE goes up. The same factor plays a role in the deviation values. Depending on whether or not a large number of such stimulations end up in our random selections for the test sets, the NRMSE value varies, and hence a high deviation. Interesting to note though, in each of our test runs, the relative performances among the different models remained almost exactly the same. This indicates that even though the deviations might hint a cause of concern, the relative behavior of M2M-Net for the different models is preserved across the different runs.

In Fig. 4, we can see the interplay between performance and complexity of the models with the neural firing input. As observed, the error reaches a minimum with the 3-layer mapper connected via Both, with about 2.8×10^5 parameters. Below this, the number of parameters are probably insufficient to model the complexity of the network. Above it, there are probably too many parameters, and those models likely overfit. Of the models we tested, a 3-layer mapper connected via Both probably has the optimum number of parameters.




Figure 3: Predicted muscle activation for the 3-layer mapper with the neuron firing input on a test set, for 15 muscles and 70 trials. (A) Ground truth. (B) Prediction with Direct. (C) Prediction with Synergies. (D) Prediction with Both.




Figure 4: Performance vs. complexity for the models with the neuron firing input.

To explore the trade-off between performance and complexity further, Table 2 lists the BIC and AICc values for the models indicated. We notice that both criteria select the 3-layer mapper with

Table 2: BIC and AICc values for the models with the neuron firing input.

	3-layer Mapper			6-layer Mapper		
	Dir	Syn	Both	Dir	Syn	Both
BIC (10^6)	1.221	1.219	2.501	2.552	2.550	5.164
AICc (10^4)	-7.266	-7.320	-7.269	-7.257	-7.239	-7.220

the neuron firing input as the model with the lowest score, indicating it to have the best performance-complexity value. Interestingly, they differ in the choice of the second best model, and so on. Understandably, we have to interpret the values in Table 2 with caution. CNNs are generally over-parameterized. Hence, classical statistical tools such as BIC and AICc would tend to favor the least complex model, at the cost of training performance. Again, training performance of a CNN is not necessarily a correct indicator of its test performance, or generalizability. On the other hand, it might bear useful insights to know how a model performs, taking into account (or penalizing) its complexity.

4 CONCLUSION

In this work, we have systemically explored different CNN architectures, to develop a suitable model for mapping muscle responses corresponding to varying TMS intensities. The model with the lowest NRMSE corresponded to the use of the E-field to neural firing rate conversion, a 3-layer mapper, and the use of both direct and synergy connections. This result demonstrates the value of the use of empirical knowledge at both ends of the network: in the form of firing rate conversion at the input, and the use of direct and synergistic connections at the output, reflecting corticospinal anatomy [25]. Greater generalizability, with the inclusion of the synergy layer and using a lesser complex model, is indicative of the current theoretical understanding of the purpose of muscle synergies in the motor system [2]. Given the physiological plausibility of the model proposed, future work will concentrate on the validation of this model further and its eventual incorporation as a diagnostic utility or means of intervention for neurological impairment.

ACKNOWLEDGMENTS

This project was funded in part by grants R01-NS088674 (MS), P41 GM103545-18 (DB,MD), R01-NS085122 and R01-R01HD058301 (ET), CMMI-193537 (ET, DE), CBET-1804550 (ET, DE, DB)

REFERENCES

- [1] Manfredo Atzori, Matteo Cognolato, and Henning Müller. 2016. Deep learning with convolutional neural networks applied to electromyography data: A resource for the classification of movements for prosthetic hands. *Frontiers in neurorobotics* 10 (2016), 9.
- [2] Emilio Bizzi and Vincent CK Cheung. 2013. The neural origin of muscle synergies. *Frontiers in computational neuroscience* 7 (2013), 51.
- [3] Hilton Bristow, Anders Eriksson, and Simon Lucey. 2013. Fast convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 391–398.
- [4] Moritz Dannhauer, Ayla Khan, Daniel White, Seyhmus Guler, Spencer Frisby, Rob S MacLeod, and Dana H Brooks. 2017. Brainstimulator: A Flexible Extensible Software Tool for Modeling and Optimizing Transcranial Brain Stimulation. *Brain Stimulation: Basic, Translational, and Clinical Research in Neuromodulation* 10, 1 (2017), e12–e13.
- [5] Tianxiang Gao and Vladimir Jovic. 2016. Degrees of freedom in deep neural networks. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 232–241.
- [6] EM Grigsby, MJ Koval, MV Smith, JK Mueller, ZD Deng, A Peterchev, WM Grill, and MA Sommer. 2015. Neural Effects of rTMS: Single Neuron Recordings From a Rhesus Macaque. *JOURNAL OF ECT* (2015).
- [7] Mark Hallett. 2007. Transcranial magnetic stimulation: a primer. *Neuron* 55, 2 (2007), 187–199.
- [8] Erik H Hoyer and Pablo A Celnik. 2011. Understanding and enhancing motor recovery after stroke using transcranial magnetic stimulation. *Restorative neurology and neuroscience* 29, 6 (2011), 395–409.
- [9] Clifford M Hurvich and Chih-Ling Tsai. 1989. Regression and time series model selection in small samples. *Biometrika* 76, 2 (1989), 297–307.
- [10] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning*. 448–456.
- [11] Nathalie Kubis. 2016. Non-invasive brain stimulation to enhance post-stroke recovery. *Frontiers in neural circuits* 10 (2016), 56.
- [12] Daniel D Lee and H Sebastian Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 6755 (1999), 788–791.
- [13] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. 2011. Stacked convolutional auto-encoders for hierarchical feature extraction. In *International Conference on Artificial Neural Networks*. Springer, 52–59.
- [14] Jerel K Mueller, Erinn M Grigsby, Vincent Prevosto, Frank W Petraglia III, Hrishikesh Rao, Zhi-De Deng, Angel V Peterchev, Marc A Sommer, Tobias Egner, Michael L Platt, et al. 2014. Simultaneous transcranial magnetic stimulation and single-neuron recording in alert non-human primates. *Nature neuroscience* 17, 8 (2014), 1130.
- [15] Zhen Ni and Robert Chen. 2015. Transcranial magnetic stimulation to understand pathophysiology and as potential treatment for neurodegenerative diseases. *Translational neurodegeneration* 4, 1 (2015), 22.
- [16] Eini Niskanen, Petro Julkunen, Laura Säisänen, Ritva Vanninen, Pasi Karjalainen, and Mervi Könönen. 2010. Group-level variations in motor representation areas of thenar and anterior tibial muscles: Navigated Transcranial Magnetic Stimulation Study. *Human brain mapping* 31, 8 (2010), 1272–1280.
- [17] K. K. Pal and K. S. Sudeep. 2016. Preprocessing for image classification by convolutional neural networks. In *2016 IEEE International Conference on Recent Trends in Electronics, Information Communication Technology (RTEICT)*. 1778–1781. <https://doi.org/10.1109/RTEICT.2016.7808140>
- [18] SCIRun. [n.d.]. SCIRun: A Scientific Computing Problem Solving Environment, Scientific Computing and Imaging Institute (SCI). <http://www.scirun.org>.
- [19] Anastasia Shulga, Pantelis Lioumis, Aleksandra Zubareva, Nina Brandstäck, Linda Kuusela, Erika Kirveskari, Sarianna Savolainen, Aarne Ylinen, and Jyrki P Mäkelä. 2016. Long-term paired associative stimulation can restore voluntary control over paralyzed muscles in incomplete chronic spinal cord injury patients. *Spinal cord series and cases* 2, 1 (2016), 1–9.
- [20] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [21] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [22] Petre Stoica and Yngve Selen. 2004. Model-order selection: a review of information criterion rules. *IEEE Signal Processing Magazine* 21, 4 (2004), 36–47.
- [23] Matthew C Tresch, Vincent CK Cheung, and Andrea d'Avella. 2006. Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *Journal of neurophysiology* 95, 4 (2006), 2199–2212.
- [24] Mathew Yarossi, Jigna Patel, Qinyin Qiu, Supriya Massood, Gerard Fluet, Alma Merians, Sergei V Adamovich, and Eugene Tunik. 2019. The association between reorganization of bilateral m1 topography and function in response to early intensive hand focused upper limb rehabilitation following stroke is dependent on ipsilesional corticospinal tract integrity. *Frontiers in Neurology* 10 (2019), 258.
- [25] Mathew Yarossi, Fernando Quivira, Moritz Dannhauer, Marc A Sommer, Dana H Brooks, Deniz Erdoğan, and Eugene Tunik. 2019. An experimental and computational framework for modeling multi-muscle responses to transcranial magnetic stimulation of the human motor cortex. In *2019 9th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 1122–1125.