Circuits and Mechanisms for TMS-Induced Corticospinal Waves: Connecting Sensitivity Analysis to the Network Graph

AUTHORS

Gene J. Yu[1,2], Federico Ranieri[3], Vincenzo Di Lazzaro[4,5], Marc A. Sommer[1], Angel V. Peterchev[1,2,6,7], and Warren M. Grill[1,6,7,8]

AFFILIATIONS

[1]Department of Biomedical Engineering, Duke University, Durham, NC, USA

[2]Department of Psychiatry and Behavioral Sciences, Duke University, Durham, NC, USA

[3]Neurology Unit, Department of Neuroscience, Biomedicine and Movement Sciences, University of Verona, P.Le L.A. Scuro 10, 37134 Verona, Italy

[4]Department of Medicine and Surgery, Unit of Neurology, Neurophysiology, Neurobiology and Psychiatry, Università Campus Bio-Medico di Roma, Via Alvaro del Portillo, 21 - 00128 Roma, Italy

[5]Fondazione Policlinico Universitario Campus Bio-Medico, Via Alvaro del Portillo, 200 - 00128 Roma, Italy

[6]Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA

[7]Department of Neurosurgery, Duke University, Durham, NC, USA

[8]Department of Neurobiology, Duke University, Durham, NC, USA

ABSTRACT

Transcranial magnetic stimulation (TMS) is a non-invasive, FDA-cleared treatment for neuropsychiatric disorders with broad potential for new applications, but the neural circuits that are engaged during TMS are still poorly understood. Recordings of neural activity from the corticospinal tract provide a direct readout of the response of motor cortex to TMS, and therefore a new opportunity to model neural circuit dynamics. The study goal was to use epidural recordings from the cervical spine of human subjects to develop a computational model of a motor cortical macrocolumn through which the mechanisms underlying the response to TMS, including direct and indirect waves, could be investigated. An in-depth sensitivity analysis was conducted to identify important pathways, and machine learning was used to identify common circuit features among these pathways.

Sensitivity analysis identified neuron types that preferentially contributed to single corticospinal waves. Single wave preference could be predicted using the average connection probability of all possible paths between the activated neuron type and L5 pyramidal tract neurons (PTNs). For these activations, the total conduction delay of the shortest path to L5 PTNs determined the latency of the corticospinal wave. Finally, there were multiple neuron type activations that could preferentially modulate a particular corticospinal wave.

The results support the hypothesis that different pathways of circuit activation contribute to different corticospinal waves with participation of both excitatory and inhibitory neurons. Moreover, activation of both afferents to the motor cortex as well as specific neuron types within the motor cortex initiated different I-waves, and the results were interpreted to propose the cortical origins of afferents that may give rise to certain I-waves. The methodology provides a workflow for performing computationally tractable sensitivity analyses on complex models and relating the results to the network structure to both identify and understand mechanisms underlying the response to acute stimulation.

AUTHOR SUMMARY

Understanding circuit mechanisms underlying the response to transcranial magnetic stimulation remains a significant challenge for translational and clinical research. Computational models can reconstruct network activity in response to stimulation, but basic sensitivity analyses are insufficient to identify the fundamental circuit properties that underly an evoked response. We developed a data-driven neuronal network model of motor cortex, constrained with human recordings, that reproduced the corticospinal response to magnetic stimulation. The model supported several hypotheses, e.g., the importance of stimulating incoming fibers as well as neurons within the cortical column and the relevance of both excitatory and inhibitory neurons. Following a sensitivity analysis, we conducted a secondary structural analysis that linked the results of the sensitivity analysis to the network using machine learning. The structural analysis pointed to anatomical mechanisms that contributed to specific peaks in the response. Generally, given the anatomy and circuit of a neural region, identifying strongly connected paths in the network and the conduction delays of these paths can screen for important contributors to response peaks. This work supports and expands on hypotheses explaining the response to transcranial magnetic stimulation and adds a novel method for identifying generalizable neural circuit mechanisms.

INTRODUCTION

Transcranial magnetic stimulation (TMS) can non-invasively activate superficial cortical regions to study brain functions, treat psychiatric and neurological disorders, and collect diagnostic biomarkers [1]. However, improving methodologies and developing new applications remain slow and challenging due to the uncertainties about what is activated by TMS and how this activation courses through the circuits within and beyond the stimulated region [2]. One approach to understanding these network effects in the motor cortex is via descending volleys of activity that propagate to the spinal cord in response to TMS and can be recorded epidurally as transient corticospinal waves (Fig 1). The corticospinal waves represent the activity of layer 5b pyramidal tract neurons (PTNs) that send axons into the spinal cord [3]. The shortest latency direct wave (D-wave) is widely agreed to represent the direct activation of PTNs [4]. Subsequent waves are called indirect waves (I-waves) and likely represent transsynaptic activations of PTNs resulting from the initial direct activation of PTNS, axons of afferents, and other neuron types. Understanding the neurons and circuits that produce the I-waves would provide insight into patterns of neuron activation and the circuit connections that mediate the cortical response to TMS [5].

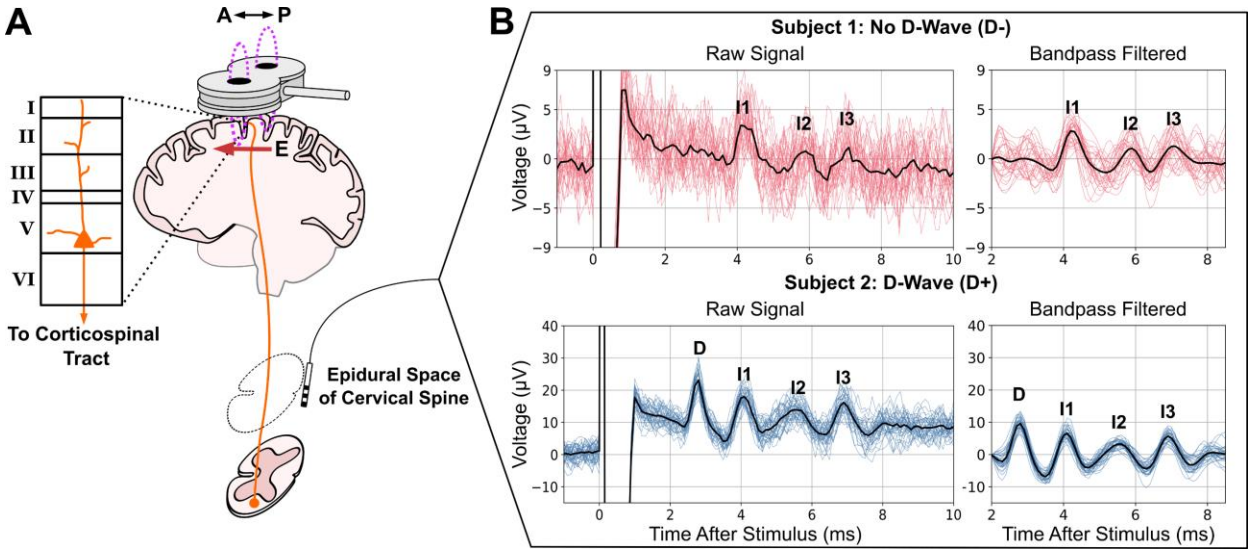

Fig 1. Descending volleys of spinal waves provide a window into motor cortical responses to TMS.

A) TMS coil with electric field (E) induced in the posterior–anterior (P–A) orientation over the motor cortex. L5 PTNs send axons into the spinal cord (corticospinal tract), and their activity is recorded epidurally at levels C1–C5. B) Epidural recordings of corticospinal waves in two human subjects. Individual trials are plotted with colored lines. The solid black lines are trial averages.

Current understanding of I-waves arises from epidural recordings combined with pharmacological interventions that identified the synaptic receptors involved in I-wave generation and broadly suggested excitatory and inhibitory mechanisms that contribute to I-waves [5,6]. These and other experimental findings were organized into conceptual frameworks to propose mechanisms that give rise to the corticospinal waves [5,6]. Two broad categories of these frameworks are I-wave generation through circuit activations and I-wave generation via intrinsic neuronal mechanisms (neural oscillator hypothesis). With circuit activation, corticocortical afferents are thought to initiate activations in different neuronal populations that propagate through the cortical circuit to L5 PTNs. Intrinsic neuronal mechanisms have also been hypothesized to allow L5 PTNs to behave as neural oscillators such that the

I-waves result from repeated spiking from the same neuron due to the dynamics following initial excitation by TMS.

Computational neuronal network models have been developed that integrate anatomical and electrophysiological details to investigate TMS-induced corticospinal waves. A model by Esser et al. represented the major layers of motor cortex using spiking point neurons and homogeneous activation of a proportion of fiber terminals across all layers to represent activation by single TMS pulses [7]. Rusu et al. developed a network model of layer 2/3 and layer 5 pyramidal neurons with realistic dendritic morphologies to investigate the effect of somatodendritic conduction and integration on I-wave generation [8]. These models generated I-wave activity that qualitatively resembled experimental findings. However, the models were not directly constrained by experimental recordings and lacked an exhaustive sensitivity analysis to investigate, among other variables, the effects of inhomogeneous activation across different neuron types.

To determine the TMS activations and neuron-to-neuron projections that contribute to I-waves, we used experimental recordings of the corticospinal response to TMS to constrain a computational model of a motor cortical macrocolumn. Starting from a reduced version of the Esser model, that could produce I-waves and is mathematically compact, we established a spiking neuronal network model of motor cortex that reproduced the features of D-waves and I-waves recorded epidurally in the cervical spine of human subjects. Next, a unified model was developed that generated responses with and without a D-wave with a change in a single parameter. A sensitivity analysis of the unified model was conducted using the two-variable-at-a-time (TVAT) method. Finally, machine learning and graph theoretical measures were used to relate the connectivity of the model to the results of TVAT analysis and identify general mechanisms producing I-waves at the circuit level. A high-level representation of the methodology is summarized in Fig 2.

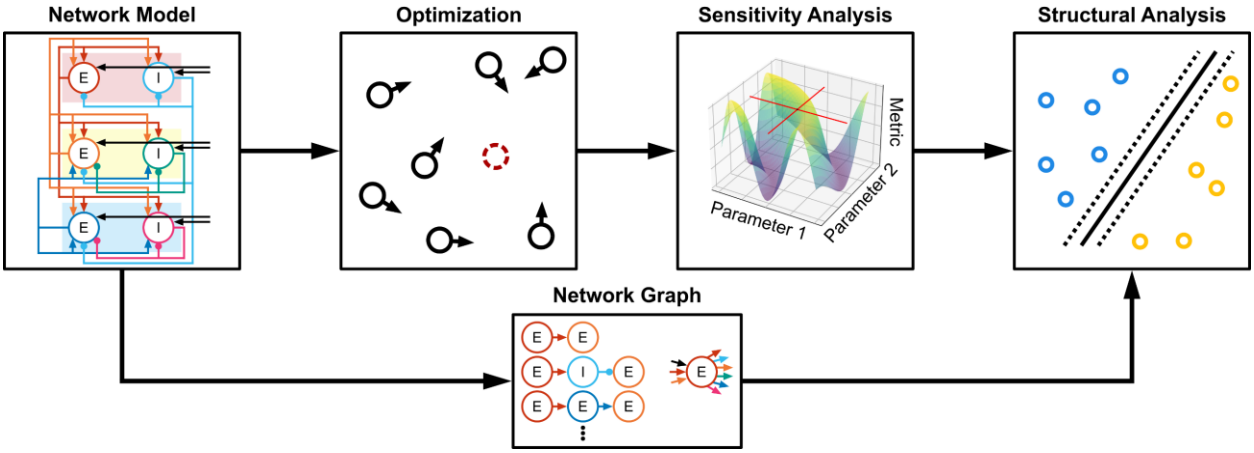

Fig 2. High level diagram of methodology.

A network model was defined, and particle swarm optimization was used to constrain parameters using experimental data. A TVAT sensitivity analysis was conducted on the optimized model, and finally the network graph was used to identify structural patterns that predict the sensitivity analysis. E: Excitatory neuron. I: Inhibitory neuron.

RESULTS

The neuronal network model used to simulate the effects of TMS represents a human cortical macrocolumn within the motor cortex and included layer (L) 2/3, L5 and L6 and is based on a model developed in Esser et al., 2005 [7] (Fig 3A). Each layer contained excitatory neurons representing pyramidal neurons and inhibitory neurons representing fast-spiking parvalbumin-positive basket cells (BC). More specifically, the layer 2/3 and layer 6 pyramidal neurons were intratelencephalic (IT) neurons with corticocortical projections, while the layer 5 pyramidal neurons were PTNs. Inhibition was mediated only by parvalbumin-positive BCs because they provide the strongest inhibition compared to somatostatin and vasoactive intestinal protein expressing interneurons [9]. Excitatory afferents (AFF) were included that targeted each of the neuron types in the motor cortical column model. The afferents non-specifically represented activity that may arise from other cortical/sub-cortical areas. Direct

activation due to TMS was represented using an input–output approach. Given a stimulus intensity as input, the output was the proportion of the population that fired an action potential in response to the TMS pulse. Both neurons and afferents could be activated, and the effect of direct activation was defined separately for each neuron and afferent type. Simulations were performed using NEURON 8.2.0+ and scripted in Python 3.8.13 [10].

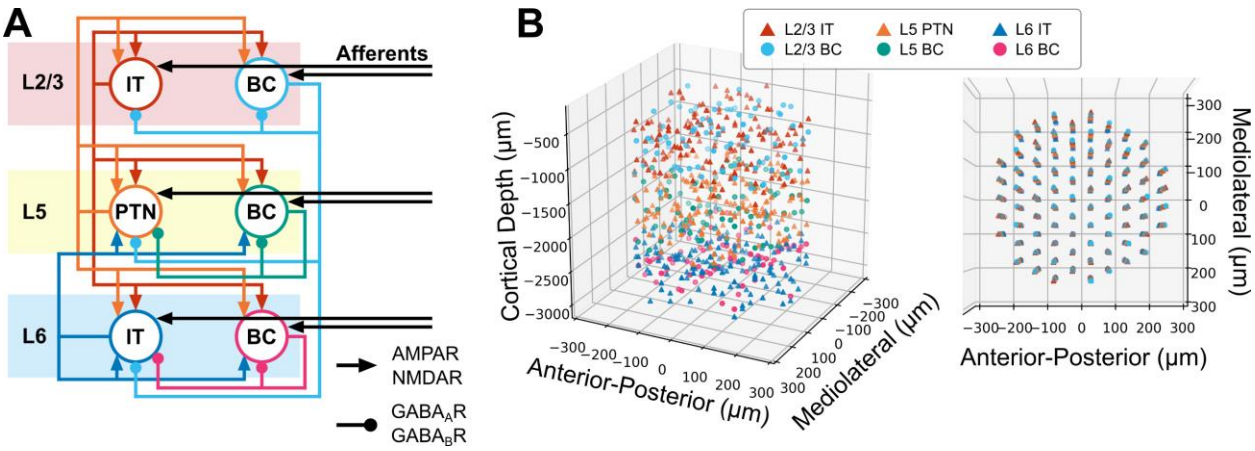

Fig 3. Overview of motor cortical macrocolumn model.

A) Block diagram of cortical connectivity. Arrowheads denote excitatory connections mediated by AMPA and NMDA receptors. Round heads denote inhibitory connections mediated by $GABA_A$ and $GABA_B$ receptors. B) Three-dimensional representation of neuron locations. (Left) Side view showing laminar distribution. (Right) Top view depicting microcolumn organization within macrocolumn. IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell.

Optimized Models Reproduce Experimental Data

Particle swarm optimization was used to identify parameters for models that responded with (D+) or without a D-wave (D-). The objective function included the firing rate of the network prior to stimulation (i.e., no stimulation) and several properties of the corticospinal response after stimulation (see Methods for a detailed description of the experimental data) including the timings and amplitudes of the peaks

and troughs. The parameters being optimized included the synaptic weights of each projection, the proportion of neurons activated by TMS, the conduction velocities for each neuron type, and the propagation delay due to stimulation of afferents. The total number of optimized parameters was 98, and the total list of parameters and their optimization ranges are described in Methods.

The final selected models had average corticospinal wave errors of 10.3% and 15.4% for the D+ and D− models, respectively (Fig 4A). The corticospinal tract activity generated by the individually optimized models captured many of the features of the experimental data (Fig 4C). The spiking responses of the models are represented using raster plots in Fig 4D, and it can be observed how the spiking activity of L5 PTNs produce the corticospinal responses in Fig 4C. The final parameter values for each of the optimized models are presented in S1 Appendix Fig A-C. To increase coverage of the parameter space and avoid local minima, multiple optimizations were executed. The convergences of total error, the distances among their solutions, and simplified Pareto front are shown in S1 Appendix Fig D-E.

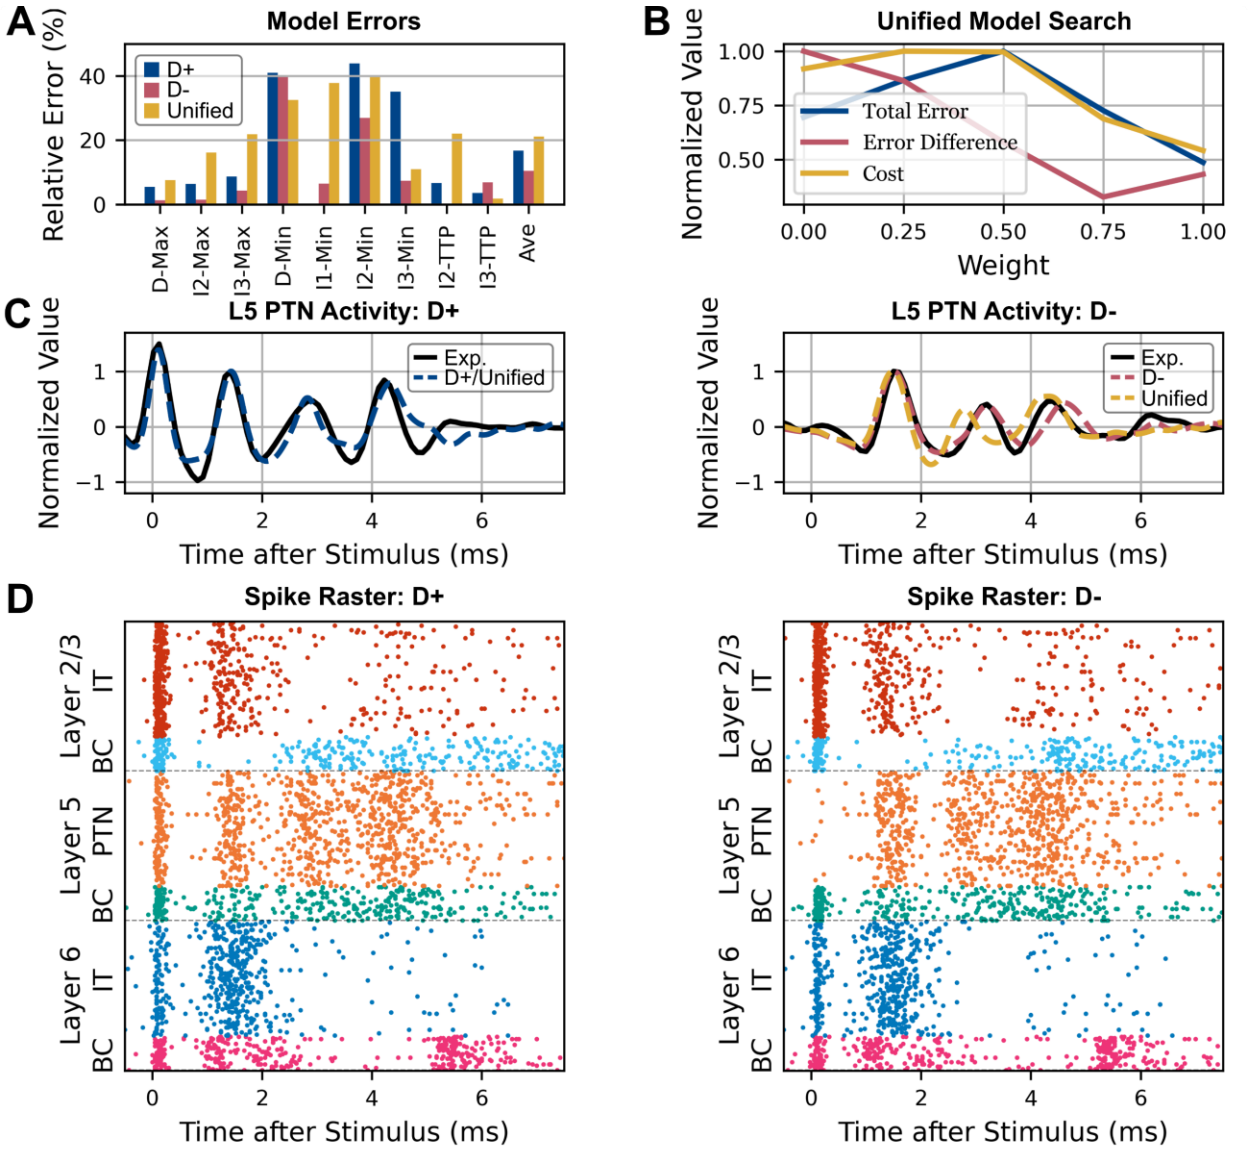

Fig 4. Optimization results and unified model.

A) Distribution of relative errors across corticospinal wave objectives for the individual best D+ and D− models and the unified D− model. Average error is plotted on the right side. B) Identification of the unified model. Weighted combinations of the parameters for the D+ and D− solutions were tested. Cost represents the sum of total error and error difference. A unified model that used the D+ parameters resulted in the lowest error across models and between models. C) Simulated epidural corticospinal activity for optimized models (dashed colored lines) compared to experimental data (solid black line). The

unified model exhibiting a D-wave (left, D+) used the same parameters as its individual best, so only a single simulation output is shown. The case without a D-wave (right, D-) has the individual best and unified model results. D) Spike raster plots for all motor cortical neuron types. A band-pass filter was applied to the activity of the Layer 5 PTN (orange) to represent the corticospinal responses shown in C.

Unified Model Accommodates Both Response Types

Despite being separately optimized, there were similarities among many of the parameters of D+ and D- models, but one of the largest differences was for direct activation of L5 PTNs (S1 Appendix Fig A), which generated the D-wave. Given the similarity between the remaining parameter values, we pursued a parsimonious model that had identical values for all parameters except the direct activation of L5 PTNs.

The unified model was generated by creating weighted combinations of the parameters of the D-wave and non-D-wave models (Fig 4B). All parameters were identical between models except direct activation of L5 PTNs, which used the respective individual optimal values. The best unified model was selected based on the total error across both models as well as the absolute difference of total error between both models to identify a model that reproduced both response types without favoring one response type over the other. The model representing the subject exhibiting a D-wave had the best generalizability to the subject without a D-wave compared to any of the weighted combinations of the model parameters, and the resulting unified D- model had average relative errors of 19.9% (compared to 10.3% for the individually optimized D- model) while the error for the unified D+ model remained unchanged (Fig 4A). The parameter values of the unified model were then used as the fixed point in a sensitivity analysis.

Sensitivity Analysis Reveals Parameters that Preferentially Contribute to Corticospinal Waves

Due to the high dimensionality of the parameter space (98 parameters), total grid search, random, or quasi-random sampling would require a prohibitively large number of simulations to characterize fully the relationships between the parameters and the corticospinal response. To reduce the computational cost, a two-variable-at-a-time (TVAT) sensitivity analysis was conducted. TVAT is a form of fixed-point analysis that varies two parameters simultaneously in a grid-search with the remaining parameters fixed at their original values. TVAT analysis is more computationally intensive than the widely used one-variable-at-a-time method, but allows characterization of pairwise interactions between variables [11,12].

TVAT analysis was performed using direct activation parameters and synaptic weights. All unique parameter pairs were varied in a grid search spanning the entire parameter range used in the optimization. The amplitudes of the simulated corticospinal waves were measured to construct amplitude maps as a function of the parameter pair involved, and polynomial regressions were used to characterize the amplitude maps. The effect sizes of a parameter for each corticospinal wave were computed using their polynomial regression coefficients if the regressions had an $r^2 \geq 0.5$ (see Methods). Fig 5A-B show examples of good and poor fits of the polynomial regressions that comprise the sensitivity analysis. The total effect sizes, computed as the sum of effect sizes across all corticospinal waves, for the 20 most influential parameters are shown in Fig 5C. Activation of L5 PTNs (TMS-L5 PTN) had the largest effect size followed by activation of afferents to L5 PTN (TMS-L5 PTN AFF). Activation of L2/3 ITs and L6 ITs had large effect sizes. Important projections included the L2/3 IT projection to L5 PTN and L2/3 IT, L5 BC projection to L5 PTN and L5 PTN to L5 PTN. All effect sizes are shown in S1 Appendix Fig F.

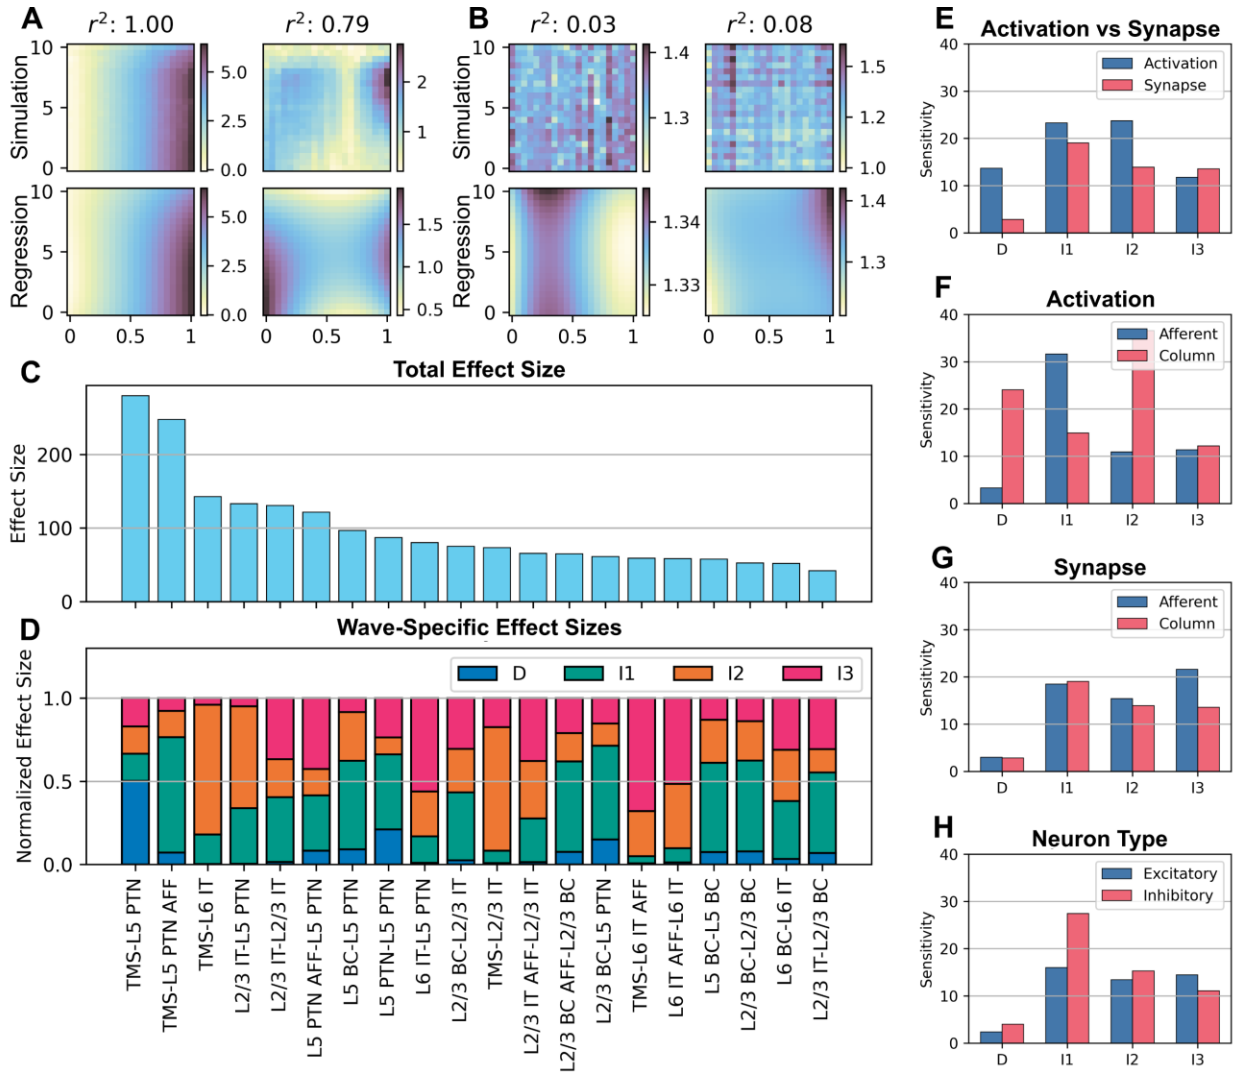

Fig 5. TVAT sensitivities, effect sizes, and their relative contributions across corticospinal waves.

A) Examples of TVAT surfaces with polynomial regressions that fit the data well. Simulation measurements are displayed on the top row; regressions are below. B) Same as A but with regressions resulting in poor fits. C) Rank sorted total effect sizes across all waves. Only the 20 largest effect sizes are shown for legibility; the full results are shown in S1 Appendix Fig F. C and D share the same x-axis. D) Relative effect sizes normalized across all waves by parameter. Parameter names were shortened and hyphenated such that the label before the hyphen corresponds to the presynaptic source and the label after the hyphen corresponds to the postsynaptic target, e.g. TMS-L6 BC indicates the activation of L6

basket cells via TMS and L2/3 BC-L5 PTN indicates the projection of L2/3 basket cells to L5 pyramidal tract neurons. E-H) Effect sizes were divided based on their contribution to a specific corticospinal wave and then grouped based on various categories. The averages within the groups are plotted. E) Compares sensitivity to activation vs. synaptic parameters. F) Sensitivity to activation of afferents vs. activation of neurons within the cortical column. G) Sensitivity to synaptic parameters related to afferents vs. neurons within the cortical column. H) Sensitivity to excitatory vs inhibitory neurons. IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.

The influence of the parameters on each individual corticospinal wave relative to the total are summarized in Fig 5D. This plot reveals that while activation of L5 PTNs substantially affected D-waves, this parameter made minimal contributions to I-waves. The activation of afferents to L5 PTNs most substantially affected the I1-wave. This analysis led to a subsequent grouping of parameters that preferentially influenced a single corticospinal wave versus parameters that affected multiple waves.

Different groupings of the total effect sizes were made to compare the average effect sizes of broader categories. The effect sizes were further subdivided based on corticospinal wave to quantify the sensitivity of the waves to the different groupings. First, the sensitivity to direct activation was compared to the sensitivity to the synaptic strengths of the network (Fig 5E). The D-wave and I1-I2 waves were highly sensitive to direct activation, and the I3-wave had an overall lower but similar sensitivity to both direct activation and synaptic strengths. Next the activation and circuit parameters were each divided between extracolumnar afferents and intracolumnar neurons (Fig 5F). The I1-wave was sensitive to activation of afferents while the D-wave and I2-wave were sensitive to activation of neurons within the column. Sensitivity levels were similar across I-waves for the synaptic effects of afferents and cortical neurons (Fig 5G). The I3-wave was more sensitive to afferents while the I1-wave was more sensitive to intracortical synaptic effects. The D-wave was not sensitive to synaptic parameters. Finally, corticospinal waves were similarly sensitive to excitatory and inhibitory neurons (Fig

5H). Sensitivities were relatively similar for I1-I3 waves to excitatory neurons while the I1-wave had the greatest sensitivity to inhibitory neurons.

Separating the effect sizes for each corticospinal wave revealed that the individual parameters could preferentially affect one wave over others (Fig 5D). A parameter was defined as having a preferential effect if the parameter's largest effect size on a corticospinal wave was at least 50% larger than its second largest effect size. The activation parameters that preferentially affected each corticospinal wave were verified by visualizing the simulations performed for the TVAT analysis (Fig 6). These visualizations demonstrate that the sensitivity analysis was consistent with the actual simulations. The analysis identified that: the D-wave was most sensitive to the activation of L5 PTNs, the I1-wave was most sensitive to direct activation of afferents to L5 PTNs, the I2-wave was most sensitive to activation of L6 Its, and the I3-wave was most sensitive to direct activation of afferents to L6 ITs.

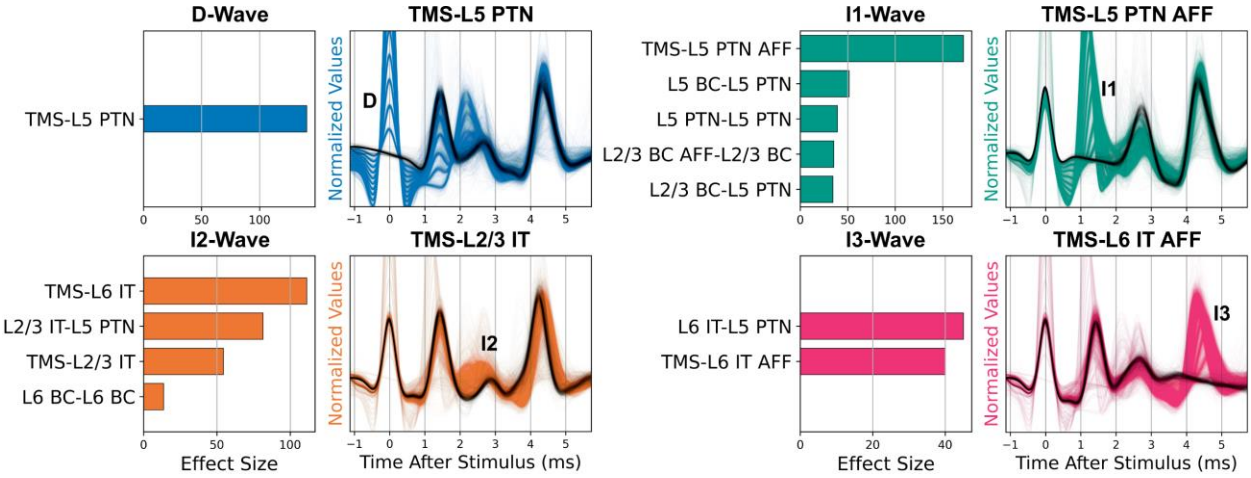

Fig 6. Effect sizes for parameters that preferentially affected a single I-wave.

For each corticospinal wave, effect sizes for parameters that preferentially affected the wave were normalized and rank sorted and visualized as bar plots. The I1-wave had 9 preferential parameters, and only 5 parameters are shown here for legibility. The remaining waves show the full numbers of preferential parameters. The full set of I1-wave preferential parameters are shown in S1 Appendix Fig G. To the right of the bar plots, the TVAT simulations involving the activation parameters with the largest effect size are shown as colored traces. The solid black line represents responses for which the parameter was set to zero. The disappearance of a wave on the solid black lines indicates that the parameter was important to the generation of that wave. The waves are labelled in the plots. IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.

Structural Parameters that Determine Preferential Influence

The sensitivity analysis predicted that multiple parameters could preferentially influence each I-wave. To identify any shared features that may predict preferential influence on the same corticospinal wave, a secondary analysis was conducted (Fig 7). The properties of the structure of the macrocolumn, such as the distances between neurons and connection probabilities, remained invariant during optimization. These invariant properties were quantified using a graph theoretical analysis, and machine learning was

used to identify patterns in the network structure that contributed to corticospinal wave generation. Because only the L5 PTNs contributed to the signal recorded in the corticospinal tract, the relationships between neuron types to the L5 PTNs were characterized by deconstructing the network graph into simple paths, i.e., paths with non-repeating nodes. All directed simple paths for all neuron types leading to L5 PTNs were characterized for analysis. See Methods for detailed descriptions of the graph characterizations.

Logistic regression with Lasso regularization was used to identify features that classified parameters with preferential versus non-preferential influence (Fig 7A). The key properties for this classification were a strong average connection probability to L5 PTNs and whether the overall effect on the L5 PTNs was excitatory or inhibitory with a validation classification accuracy of 94.6% (Fig 7B). Recursive feature elimination and support vector classification were used to identify properties of the preferential parameters that predicted which corticospinal wave they affected (Fig 7C-D). The key property was the conduction delay of the shortest path between the starting neuron and the L5 PTNs, and the validation accuracy was 87.2%.

Although the sensitivity analysis identified important circuit mechanisms (i.e., activations and projections) involved in corticospinal wave generation, the subsequent machine learning analysis identified the anatomical bases that explained how and why the circuit mechanisms had a preferential effect. This secondary structural analysis provides a method for identifying fundamental principles involved in the neural response to acute stimulation.

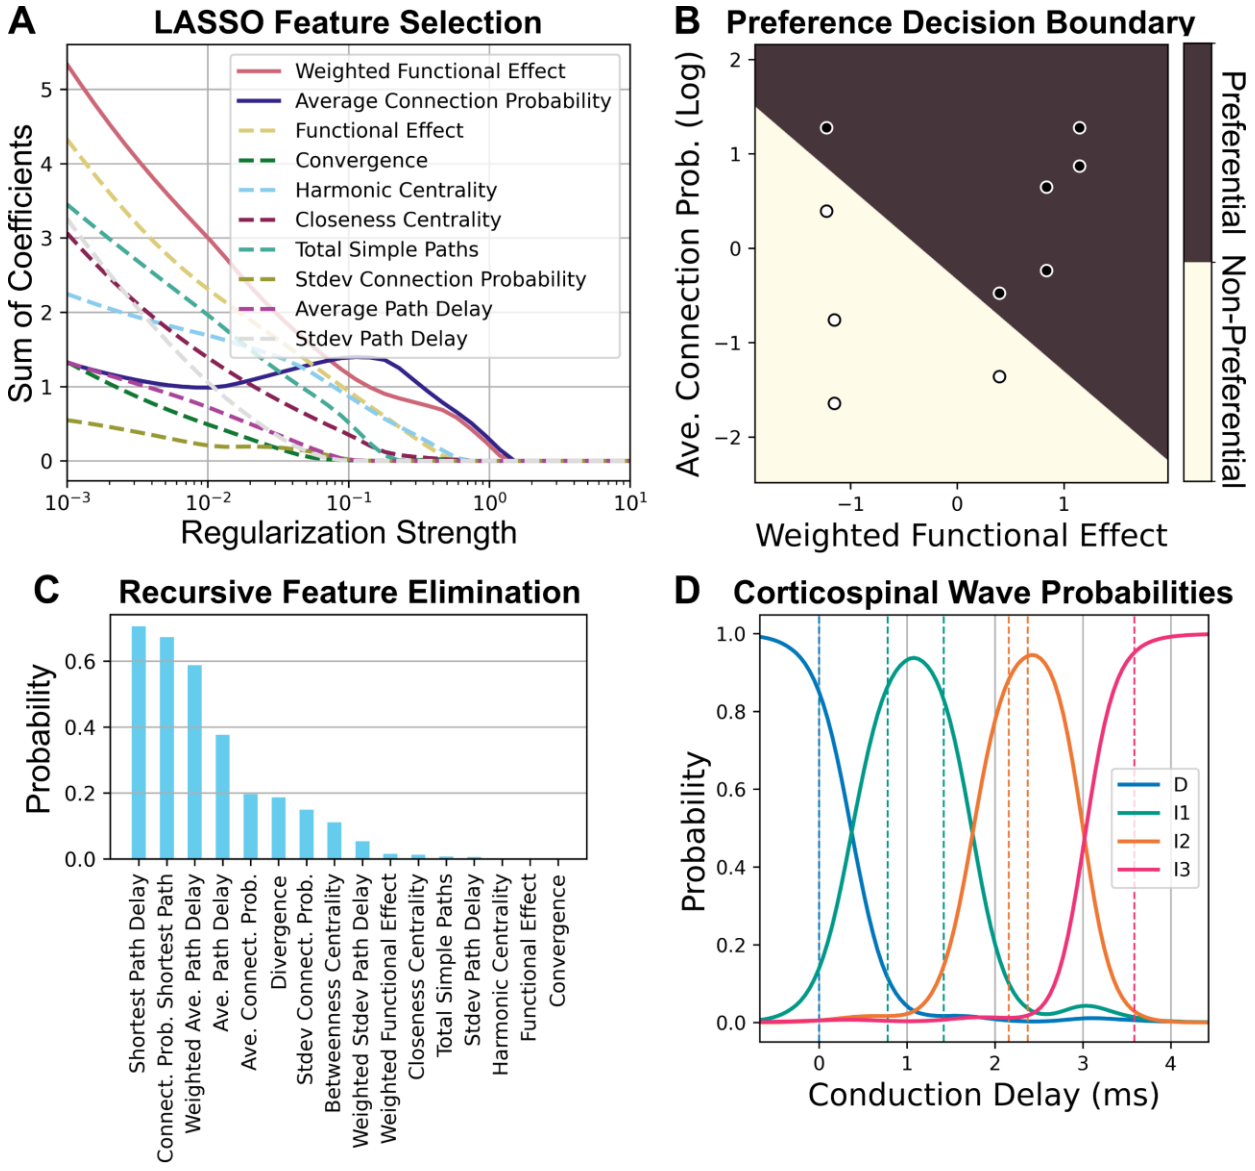

Fig 7. Classification of network features for effect size types.

A) The sum of the absolute value of the coefficients was plotted against regularization strength to identify the best parameters that classify preferential versus non-preferential parameters. Parameters that remain nonzero as regularization strength increases have better classification performance. Only the top 10 best features are shown for legibility. B) Logistic regression decision boundary for preferential parameters (dark) versus non-preferential parameters (light) using the best classification features identified in A. Dark filled dots indicate data that were preferential, and light filled dots indicate data

that were not preferential. C) Recursive feature elimination was performed to classify corticospinal wave preference for preferential parameters. The higher probabilities of remaining after elimination indicated better classification accuracy. Only 10 features are shown for legibility. D) Corticospinal wave probabilities obtained by support vector classification using the single best classification feature from C. The dashed lines represent the conduction delays of the data being classified.

DISCUSSION

We developed an experimentally-constrained model of a human motor cortical macrocolumn that generated realistic D-waves and I-waves in response to single pulse TMS. The model reproduced responses that included or excluded a D-wave by changing the direct activation of L5 PTNs, which is consistent with the mechanisms of D-wave generation [4]. TVAT sensitivity analysis, which lies between a local and global sensitivity analysis, identified the circuit pathways and TMS activations important to I-wave generation.

The results of the sensitivity analysis support the hypothesis that direct activation of the terminals of afferents to motor cortex are an important mechanism for I-wave generation but are not consistent with the hypothesis that I-waves are generated by repetitive firing of single neurons (neural oscillator hypothesis). The analysis also supports the involvement of both excitatory and inhibitory neuron types in modulating I-waves [5]. In addition, the sensitivity analysis identified afferents and neuron types endogenous to the motor cortex that can be directly activated to generate corticospinal waves. Subsequently, structural analysis identified general structural principles that allowed these activations to preferentially generate corticospinal waves. Direct activation of afferents and neuron types can preferentially contribute to single I-waves if they have a highly connected path to L5 PTNs, relative to all other paths between the activated neuron type and L5 PTNs. Finally, the latency of the I-wave that is affected by a path can be predicted by its total conduction delay to L5 PTNs.

Separate Pathways for Activation that Include Excitatory and Inhibitory Neurons

The leading hypothesis for I-wave generation proposes that 1) separate activation pathways exist for early versus late I-waves, and 2) activated pathways include both excitatory and inhibitory neurons [6,13]. The sensitivity analysis identified neural activations that preferentially modulated specific I-waves, revealed preferential activation pathways for all three I-waves, and showed that silencing their activation greatly suppressed a particular I-wave (Fig 6). The sensitivity analysis was grouped to compare the total effect sizes of excitatory and inhibitory neurons on I-wave generation and revealed that corticospinal waves exhibited comparable sensitivities to both excitatory and inhibitory neurons and that inhibitory neurons are involved in I-wave modulation (Fig 5H).

Most inhibitory neurons had non-preferential effects, i.e., affected multiple I-waves, which is consistent with experimental findings that various anesthetics, which act as allosteric modulators of $GABA_AR$, generally reduce I-wave amplitudes [6]. However, the sensitivity analysis showed that the I1-wave was most sensitive to inhibitory neurons with decreasing sensitivity for later I-waves (Fig 5H), and this is not consistent with experimental findings that show $GABA_A$ agonists affect later I-waves but not the I1-wave [14–16]. One possibility for this disagreement is the lack of inhibitory afferents in the model that could arise from adjacent cortical macrocolumns. These afferents would provide inhibition at longer latencies that would affect later I-waves. The Model Limitations and Future Directions subsection discusses this further.

Direct Activations of the Endogenous Circuit Contribute to I-Waves

The prior conceptual frameworks assumed that I-waves are initiated by activation of corticocortical fiber afferents, and the sensitivity analysis supports that the corticospinal response is most sensitive to activation of terminals of afferents. However, this analysis revealed that activation of the motor cortical circuit itself can initiate I-waves. Activation of ITs in L2/3 and L6 preferentially activated I2- and I3-waves

(Fig 6). Although designing *in vivo* TMS experiments that control for contributions of endogenous circuit elements to I-waves is difficult, the modeling results suggest activation of the endogenous circuit as another mechanism for I-wave generation, in addition to activation of afferents. Intracortical microstimulation (ICMS) studies can provide some insight into intracortical TMS effects and are further discussed below in the Comparison to Intracortical Microstimulation subsection.

Connectivity and Conduction Delay as Mechanisms for Preferential I-Wave Generation

Given that multiple mechanisms can preferentially contribute to the same I-wave, the structural analysis sought to identify the commonalities among mechanisms that yielded this response. A neuron type within the circuit could have multiple paths leading to L5 PTN with different properties for each path. Neuron types with a single path that had a high connection probability to L5 PTNs, relative to other paths starting from the same neuron type, could preferentially affect a single I-wave (Fig 7A-B). For neuron types where such a path exists, the primary mechanism for determining early versus late I-wave activation was the conduction delay of the path between the activated population and L5 PTNs (Fig 7C-D). The conduction delay defined in this study represents the combined contributions of action potential propagation along the axon, synaptic transmission, and somatodendritic propagation of the resulting postsynaptic potential. This is supported by the computational work of Rusu and colleagues who controlled conduction delay based on synaptic location within dendrites [8].

To generalize, the results of the structural analysis suggest that if the generator of a signal within a network is known, and the connection probabilities and conduction delays of the network are known, then the network elements that preferentially contribute to singular peaks of a system's impulse response can be screened by performing the following: for each neuron type 1) identify all possible paths from the neuron type to the signal generator, 2) compute the ratios of the log of the connection

probability between the most highly connected path and the remaining paths normalized by the sum of all log probabilities, and 3) obtain the latency of effect for the most highly connected path. Neuron types that have a path that is more highly connected than the remaining paths will have a preferential influence on peaks that occur during their latency of effect.

L5 PTNs as Population Oscillators but Not Neural Oscillators

Another category of hypotheses for I-wave generation is the concept of the neural oscillator. These theories were motivated by the fact that L5 PTNs can achieve firing rates that match the frequency of I-waves and led to exploration of cellular mechanisms for I-wave generation [6]. A histogram was constructed of the spike counts for each L5 PTN during the different I-waves (S1 Appendix Fig H), and L5 PTNs were most likely to contribute to a single I-wave during the corticospinal response. However, at the population level excitatory recurrent connections exist between L5 PTNs, and the sensitivity analysis demonstrated that the recurrent connections are involved in I-wave modulation as seen in Fig 5D and Fig 6. Therefore, the modeling results do not support that I-waves are generated or sustained at the neuronal level; rather, their generation appears to be a population level effect.

Comparison to Intracortical Microstimulation (ICMS) Studies

Direct cortical recordings to investigate I-waves are currently limited due to the technical challenges of suppressing the TMS artifact, which saturates recordings and prevents recovery of the activity during the period when the D-wave and I-waves occur [17,18]. ICMS in animals can generate high frequency multiunit activity with frequencies comparable to I-waves [19–21]. The results of ICMS studies can contribute to understanding the TMS response, but due to the differences in the spatial distribution and gradient of the electric field, ICMS studies cannot be used to explain fully TMS evoked I-waves [22].

ICMS applied to the primary motor cortex (M1) hand area in nonhuman primates showed that earlier peaks were elicited if the stimulation was closer to the recording site [21]. The study hypothesized that

the stimuli were activating horizontal fibers within M1, and these results support conduction delay as a mechanism determining the latencies of peaks. The horizontal fibers further represent afferents, relative to a macrocolumn, that are endogenous to M1. Single unit activity from a similar ICMS study that stimulated and recorded from M1 found minimal, sparse spiking within the time window relevant for I-waves and supports that single L5 PTNs contribute to few I-waves, if at all [19]. This corroborates the modeling predictions that I-waves represent a population response comprised of heterogeneous, sparse spiking rather than a synchronized rapid spiking response across neurons (S1 Appendix Fig H). Another ICMS study stimulated a region of the ventral premotor area F5 that sends afferents to the hand knob area of M1 [20]. Stimulation of F5 at lower intensities recruited the I1-wave first, and higher intensities eventually recruited later I-waves. Although it is known that F5 projects to M1, the laminar distribution of the terminals of F5 afferents in M1 are unknown. Nonetheless, these results are consistent with the modeling prediction that the I1-wave is most sensitive to activation of afferents. Maier and colleagues also stimulated M1 directly and found that D-waves are much less likely to be elicited than I1-waves. This finding is in line with the TMS literature [23], and the sensitivity analysis (Fig 6C) is also consistent with these experimental observations in that the I1-wave is most sensitive to stimulation of afferents compared to the D-wave, which is least sensitive.

Putative Afferents for I-Wave Generation

In the present model, afferents were represented as spiking inputs that were specific for each neuron type in the model, and the effect of TMS was represented by activation of the axon terminals of these afferents within the motor cortical macrocolumn. The sensitivity analysis predicted that activation of afferents for specific neuron types could have a preferential effect on specific I-waves, so the results of the sensitivity analysis were compared to the laminar distribution of terminals of corticocortical afferents in mouse motor cortex [24] to predict the anatomical origin of afferents with preferential I-wave effects. Afferents originating from the secondary (supplementary) motor area (M2) have a high

density of terminals in the deep portion of L5 where the somata of L5 PTNs lie, and activation of M2 afferents may be a candidate for I1-wave generation. Afferents from the primary somatosensory cortex have a high density of axon terminals in L2/3 and superficial L5 and could be important for I2-wave generation. The axon terminals of the orbital cortex primarily target L6 and may contribute to I3-waves. The axon terminal distributions for lateral and anterior ventral thalamus within motor cortex were also characterized [24], but prior studies showed that lesions in those areas do not affect I-wave generation [25].

The laminar distribution of horizontal connections between columns within motor cortex have not been directly characterized. However, Narayanan and colleagues reported the laminar distribution of axon terminals endogenous to rat primary somatosensory cortex [26]. The horizontal connections of L2/3 and L5 pyramidal neurons are most dense in L2/3, which may contribute to the I2-wave. The horizontal connections of L6 pyramidal neurons are most dense in deep L5 and L6 which may contribute to I1- and I3-waves.

Model Limitations and Future Directions

An important design criterion for the modeling work was computational efficiency to enable the parameter explorations necessary for optimization and sensitivity analysis to be conducted in a reasonable time. In general, computational gains came at the expense of biological details and constraints. However, the simplified model enabled more specific and in-depth computational experiments.

Point neuron representations precluded any analyses involving dendritic processes, spatial integration of postsynaptic potentials, or ephaptic coupling. Spatially extended, i.e., morphologically realistic, neuron models [22] could accommodate these mechanisms and enable the exploration of their

contributions to modulation of I-waves but would increase execution times by a factor of approximately 1800.

Afferents were represented as spiking processes that targeted specific neuron types. More realistic representations of afferents with distributions and connectivities that matched anatomical data would more directly address the effect of specific fibers on I-waves. Nonetheless, allowing afferents to be separately variable for each neuron type provided a basis to understand their contributions.

Traditionally, L4 in motor cortex has been described as either nonexistent or very thin, which led motor cortex models to exclude L4 or represent it with inhibitory neurons only [7,27,8]. Recent evidence has identified excitatory IT neurons in L4 with projections to L2/3 [28–30] leading to more complex models of M1 [31]. The present modeling results predict that, while not included, L4 IT neurons would participate in later I-waves due to their strong projection into L2/3; therefore, future work should add L4 explicitly to the model.

A single macrocolumn comprising multiple microcolumns was modeled in this work. Communications across adjacent macrocolumns, i.e., intracortical afferents, could alter the corticospinal response to TMS as they represent “afferent” inputs to macrocolumns that arise within the motor cortex. Their interactions could further modulate I-waves through both excitatory and direct inhibitory projections, and the latencies of the feedback will likely cause adjacent macrocolumns to contribute toward late I-waves.

This work represented TMS stimulation using an input–output approach, i.e., a given stimulus intensity resulted in some proportion of neurons of a particular type to fire an action potential. The spatial distribution of activation could be constrained by modeling the induced electric field using finite element modeling [32]. However, by separating the neuron type activations from the spatial constraint, the basic properties underlying the responses to activation could be investigated with greater control.

Furthermore, the optimization included only a single stimulus intensity as a constraint. Incorporating corticospinal recordings in response to multiple stimulus intensities from the same subject would provide better constraints and allow analysis of recruitment orders for neuron types.

The predictions from the model are limited to the single pulse response and are not readily extendable to paired pulse or repetitive pulse paradigms. This is partly due to $GABA_BR$ parameters being underconstrained. $GABA_BR$ conductance was partially constrained by the baseline firing rate objective but has been shown to have no effect on I-waves [33]. However, $GABA_BR$ is important for the cortical silent period [34] and paired pulse responses [35], and these data can be incorporated as optimization constraints in future work.

Finally, experimental data from only two subjects was used with responses from a single TMS intensity. The data were representative of the two qualitative types of responses—with and without D-wave. The small dataset allowed for more rapid model development due to fewer optimization constraints, and the methods established in this work can be applied in the future to extended data from more subjects and more recordings within subject.

Conclusions

To understand the mechanisms and principles underlying a biological process, sensitivity analysis is a powerful tool. However, as the number of relevant variables increases, the analysis can become overwhelming, and conclusions become diluted. At these large numbers, degeneracy in the sensitivity analysis is possible as many mechanisms can be identified to be significant to the phenomenon of interest. However, there is also the possibility that subsets of these mechanisms share certain properties that represent a more fundamental mechanism or at least a lower-level mechanism that was previously unclear or unaccounted for. In this case, a secondary analysis can reveal more fundamental mechanisms that underly the variables that explain the phenomenon of interest. For this work, the lower-level

mechanisms were model parameters that described the anatomical structure of the network, i.e., the wiring diagram and the latencies that resulted from these anatomical constraints. The insights on how the wiring diagram and the conduction latencies affect peaks in an evoked response can be generalized and applied to areas outside the motor cortex and to stimulation modalities beyond TMS.

METHODS

Motor Cortical Column Simulations

Neuronal Network Model

The motor cortical macrocolumn model was based on the equations and parameters published by Esser et al., 2005, which specified the connectivity, somatic biophysics, and synaptic properties [7]. The model contained L2/3 ITs and BCs, L5 PTNs and BCs, L6 ITs and BCs and excitatory afferents that targeted each neuron type (i.e., six groups of afferents). The circuit describing the connectivity is shown in Fig 3A. The Esser model was chosen as a starting point due to its ability to generate I-waves and the low computational complexity of its leaky-integrate-and-fire, point neuron models. The spiking activities of the afferents were generated by a Poisson process with a mean firing rate of 0.25 Hz [36]. Noise was added to the neuron models that was independent of the synaptic drive provided by the afferents and unaffected by TMS to ensure proper baseline firing rates and reduce network synchronization. Each neuron received its own noise in the form of short, suprathreshold current injections with Poisson-distributed intervals. Although the Esser model included the thalamus and thalamocortical projections, the thalamus was omitted from the present work to further reduce computational time because it does not affect I-wave generation [25].

The macrocolumn encompassed a cylinder with a diameter of 500 µm (Fig 3B) based on anatomical studies [37]. The height of the cylinder was 2700 µm based on measurements made on human motor cortex from ex vivo brain [38]. This study also informed the total vertical thickness (i.e., depth) of the

511 layers within the macrocolumn. The cortical depth location of a neuron was uniformly and randomly
512 generated within the appropriate layer bounds. The macrocolumn was comprised of microcolumns that
513 were arranged in a triangular lattice with a spacing of 50 μm [39] resulting in 79 microcolumns and
514 matched the range of microcolumns per macrocolumn [37,40]. The microcolumns were synonymous
515 with the “topographical elements” described in the Esser model and contained 2 excitatory neurons and
516 1 inhibitory neuron per layer. With 3 neurons per layer, 3 layers per microcolumn, and 79 microcolumns
517 in the macrocolumn, there was a total of 711 neurons (Table 1).
518
519
520
521
522
523
524
525
526
527
528
529
530

Table 1. Total numbers of neurons in model.

Neuron type	Number
L2/3 IT	158
L2/3 BC	79
L5 PTN	158
L5 BC	79
L6 IT	158
L6 BC	79
L2/3 IT AFF	79
L2/3 BC AFF	79
L5 PTN AFF	79
L5 BC AFF	79
L6 IT AFF	79
L6 BC AFF	79

The conduction delay, defined as the time between the onset of an action potential and the start of the postsynaptic potential at the soma of the postsynaptic neuron, was calculated from the distance between the presynaptic and postsynaptic neuron pair and conduction velocity. The conduction velocity measured from non-human primates (570 μm/ms) was used as human measurements were not available [41].

TMS activation included only suprathreshold effects. Each stimulus activated a specified proportion of a neuron/afferent type, and neurons/afferents were randomly selected for each presentation of the stimulus. No effect was applied to neurons/afferents that were not selected. Direct activation of neurons resulted in an injection of a short suprathreshold current to elicit an action potential that was propagated orthodromically to all postsynaptically connected neurons using all relevant conduction

543 delays. Direct activation of the terminals of afferents resulted in the activation of all connected synapses
544 with the appropriate conduction delays.

545 Connectivity parameters, neuron parameters, and synaptic parameters were identical to those reported
546 in [7] with the following exceptions. Orientation selectivity-based connectivity was not included, so the
547 connectivity rules for all microcolumns were identical. Because the geometric area of the model was
548 reduced from the original, the overall synaptic drive was decreased. The subsequent optimization
549 allowed larger synaptic weights to compensate.

550 Simulation Paradigm

551 Simulations were designed to ensure that the network achieved steady state before measurements
552 were made. To reduce synchronization of the network due to simultaneous activation of afferent inputs,
553 the onsets of the Poisson spike trains of the afferents were randomly and uniformly selected between 0
554 and 200 ms. Baseline properties were measured between 500 and 2000 ms. TMS stimuli were applied at
555 2000 ms with inter-trial intervals of 200 ms with a total of five trials. This interval was selected based on
556 population averages of trials which showed no longer-term effects beyond 150 ms. Furthermore, the
557 model did not implement synaptic plasticity and thalamic connections. Analysis of the TMS response
558 was performed on the trial average. The total simulated time was 3000 ms.

559 Selecting an Appropriate Time-Step

560 The time-step was decreased from the value originally used in Esser et al., 2005, from 0.1 ms to 0.025
561 ms due to instabilities in the network during these longer simulations. The time-step was selected by
562 running single neuron simulations while log-linearly varying the time-step from 0.001 to 0.2 ms. Each
563 simulation had a length of 20 seconds, and the models received a random Poisson input with a mean
564 firing rate of 1000 Hz. The response at 0.001 ms was used as the baseline response, and the model
565 behavior were characterized using the following metrics: Number of spikes generated, mean inter-spike

566 interval (ISI), coefficient of variation of the ISI, normalized root mean square error (NRMSE) of the
567 membrane potential, and the van Rossum spike distance [42]. A time constant of 500 ms was used for
568 the spike distance because the 0.001 ms time-step case had a mean ISI of approximately 500 ms. For
569 each time-step, 50 simulations/trials were conducted. Each trial used a different random seed to change
570 the Poisson input, and the sequence of random seeds for the trials was identical across time-steps. The
571 mean of the metrics across trials for each time-step was calculated for further analysis.

572 The knee-finding Python package *Kneed* [43] was used to identify the largest time-step at which further
573 time-step increases would provide diminishing returns on the differences in metrics relative to the 0.001
574 ms time-step (Fig 8). A 5^{th} order polynomial function was fitted to the metrics as a function of the log of
575 the time-step size to provide a continuous curve to identify the knee. The smallest time-step across all
576 metrics was 0.03 ms for both mean ISI and the coefficient of variation of ISI, and a final time-step of
577 0.025 ms was conservatively selected.

578

579

580

581

582

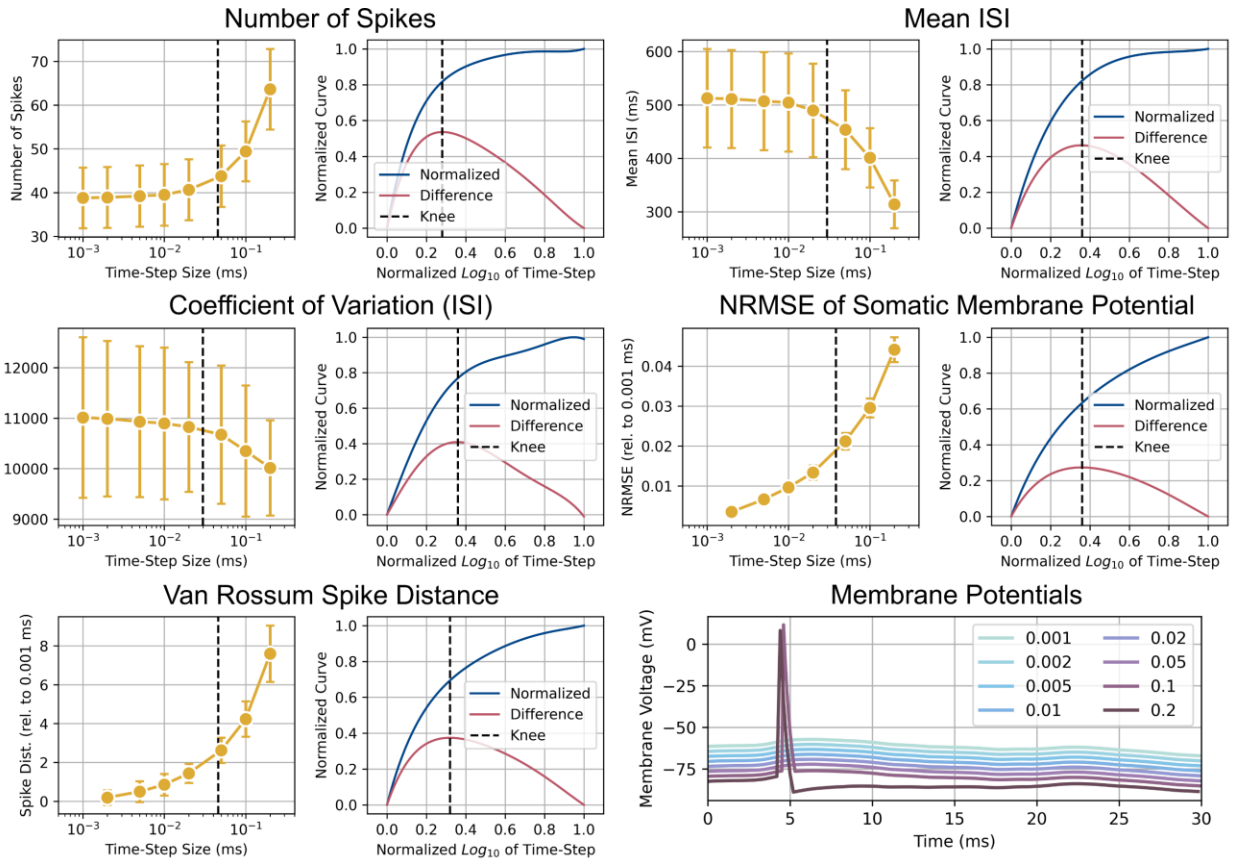

583

584 **Fig 8. Analysis to select a time-step that both minimizes computation time and is numerically stable.**

585 *Each pair of plots shows the mean of a metric as a function of the time-step size in the log_{10} scale on the*
586 *left. The right plot of the pair shows the normalized curve and difference curve used to identify the knee-*
587 *point. The vertical dashed line in the pair of plots denotes the ideal time-step. On the lower right, the*
588 *membrane potentials of the neuron model for different time-steps are shown. Offsets were added for the*
589 *y-axis to allow all lines to be distinctly seen. The plots depict a key behavior that differentiates*
590 *simulations at larger time steps. A pronounced afterhyperpolarization is seen with a 0.2 ms time-step*
591 *that is absent from other time-steps. Additionally, spikes are generated at larger time-steps (0.1 and 0.2*
592 *ms) that are absent for smaller time-steps. These dynamics contribute to the larger numbers of spikes,*
593 *lower mean ISIs, larger NRMSE, and larger spike distance observed for larger time-steps.*

Experimental Data

Experimental data were obtained from human subjects who had spinal cord stimulators implanted to treat drug-resistant dorso-lumbar pain. Data was collected in accordance with an experimental protocol that was approved by the Ethics Committee of Campus Bio-Medico University of Rome. Use of the data in this study was approved by the Institutional Review Board of the Duke University Health System.

The experimental setup is summarized in Fig 1A. For each subject, an electrode array was implanted percutaneously in the cervical epidural space, with the recording sites aligned vertically along the dorsum of the cord. Spinal potentials were recorded differentially between proximal-distal pairs of contacts (with the distal contact connected to the reference input of the amplifier), amplified and filtered (gain: 10000; bandwidth: 3 Hz to 3 kHz) by a Digitimer D360 amplifier (Digitimer Ltd., Welwyn Garden City, UK), and sampled at 10 kHz by means of a CED 1401 A/D converter (Cambridge Electronic Design Ltd., Cambridge, UK).

A figure-of-eight coil with external loop diameter of 70 mm was held over the right motor cortex at the optimal scalp position to elicit motor responses in the contralateral first dorsal interosseous (FDI) muscle with the induced current flowing in a posterior–anterior direction across the central sulcus. TMS was delivered at 120% of the resting motor threshold (RMT). Monophasic pulses were applied with a Magstim 200^2 stimulator (The Magstim Company Ltd., Whitland, UK) once every 5 seconds.

Two subjects were included in this study (Fig 1B). Subject 1 was female, 64 years old, and had a cervical epidural electrode implanted at C3–C5 level; the RMT of TMS was 34% of maximum stimulator output. Subject 2 was male, 68 years old, and had a cervical epidural electrode implanted at C1–C2 level; the RMT was 55% of maximum stimulator output. Subject 1 did not exhibit a D-wave in response to TMS (D-), while Subject 2 exhibited a D-wave (D+). Each subject received at least 30 pulses. For analysis, the responses were truncated to begin 2 ms after the TMS pulse to remove stimulation artifact. An

additional noncausal bandpass filter (second-order Butterworth, 200 Hz to 1500 Hz) was applied to remove residual stimulus artifact, potential motor artifacts, and higher frequency activity that is unrelated to the corticospinal waves. Measurements of the corticospinal response were performed on the filtered, trial-averaged signal.

Optimization of Network Model

Particle swarm optimization

Particle swarm optimization (PSO) is a metaheuristic algorithm for parameter exploration with the goal of finding parameters that satisfy one or more constraints. The particle's position represents the parameter values for the model, and a velocity term updates the position using a weighted combination of the best solution found by itself (cognitive best) and the best solution found among a particle's neighbors (social best). PSO was implemented by modifying the *inspyred* Python software package [44].

Neighborhoods were constructed using a star topology with each particle's neighborhood size being 5% of the total number of particles. There were 2048 particles and 300 iterations before the optimization was terminated. The optimization was repeated for each model four times to increase coverage of the parameter space and the likelihood of locating a global best solution. Each optimization used a different random seed that controlled the initial particle positions and their updated positions after each iteration as detailed below.

At the beginning of the optimization procedure, particle positions were initialized using Sobol sampling. Sobol sampling generates a low-dispersion quasi Monte-Carlo sequence that exhibits better coverage of the parameter space than uniform random sampling for high-dimensional spaces and has been shown to improve optimization convergence [45].

638 Particle behavior was guided by inertial velocity, cognitive velocity, social velocity, gain factor, and noise
639 [46]. Inertial weight corresponded to a particle's resistance to movement and results in a particle
640 moving towards its previous position. The cognitive weight determined a particle's preference towards
641 the position of the best solution it had found. The social weight determined a particle's preference
642 towards the position of the best solution its neighborhood had found. The cognitive and social velocities
643 were also separately modified using scalars drawn from a uniform distribution between 0 and 1. The
644 velocity was then computed as the weighted average using the inertial, cognitive, and social weights.
645 Finally, the velocity was scaled by the gain factor. For each particle coordinate, noise was sampled from
646 a zero-mean Gaussian distribution with the standard deviation controlling the strength of the noise.
647 Optimization noise is also known as mutation and was shown to be necessary for theoretical global
648 convergence of PSOs [47]. Finally, the particle position was updated using both velocity and noise.

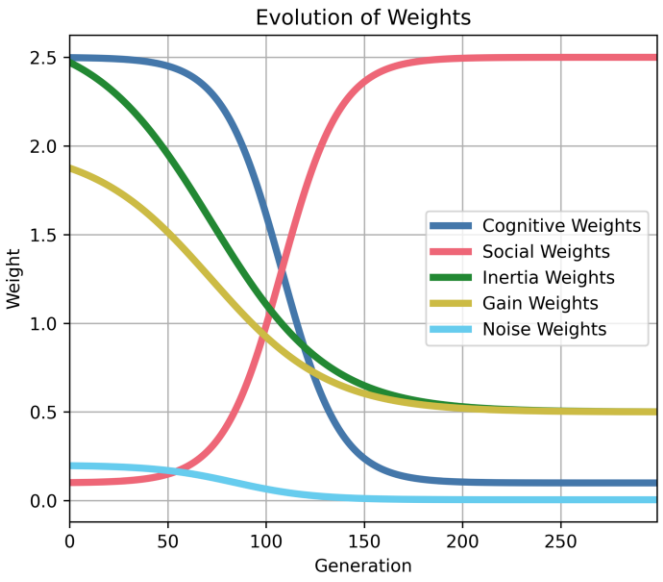

649

650 **Fig 9. Change in particle swarm optimization weights across successive iterations.**

651 *For approximately 100 iterations, optimization is exploratory with large cognitive, inertial, and gain*
652 *weights before favoring convergence with high social weights for the final 150 iterations.*

These optimization parameters were updated during optimization to switch from an initial stage of exploration to a final stage of convergence (Fig 9). During exploration, inertial weight, cognitive weight, gain factor, and noise were high, and the social weight was low. During convergence, the social weight was high, and the remaining terms were low. The progression of the parameters followed a sigmoidal function

$$y(x) = A + \frac{K}{1 + e^{(ax - bN)/N}}$$

where *x* is the current iteration of the optimization, *N* is the total number of iterations for the algorithm, *A* is the offset, *K* is the amplitude and direction of the sigmoid, *a* controls the steepness of the transition, and *b* controls the midpoint of the transition. The parameters for the sigmoidal function are reported in Table 2.

Table 2. Sigmoid function constants underlying evolution of optimization metaparameters.

Parameter	A (Minimum)	K (Amplitude/Direction)	a (Slope)	b (midpoint)
Cognitive Weight	2.5	-2.4	20	7.2
Social Weight	0.1	2.4	20	7.2
Inertial Weight	0.5	2	15	4.2
Gain Weight	0.5	1.5	10	2.4
Noise Weight	0.005	0.195	15	4.2

A damped, reflecting boundary condition was implemented on the parameter search space [48]. If a particle's position exceeded a boundary, then the particle was reflected back into the valid parameter space using the difference between the original, non-valid position and the boundary. The reflection was damped by multiplying the difference with a scalar sampled from a uniform distribution between 0 and 1.

Optimization constraints

There were four main categories of constraints: baseline activity, TMS response, synchrony, and well-behaved. The relative error was computed for each constraint except when the constraint was zero, in which case the absolute error was computed. The sum of the relative and absolute errors was used to represent the total error of a particle. Table 3 lists all constraints.

Table 3. List of Optimization Constraints

Constraints		
1. D-wave peak	18. L2/3 IT ISI	35. L5 PTN baseline CV
2. D-wave time-to-peak	19. L2/3 BC firing rate	36. L5 BC baseline CV
3. D-wave trough	20. L2/3 BC ISI	37. L6 IT baseline CV
4. D-wave time-to-trough	21. L5 PTN firing rate	38. L6 BC baseline CV
5. I1-wave peak	22. L5 PTN ISI	39. L2/3 IT population ISI std.
6. I1-wave time-to-peak	23. L5 BC firing rate	40. L2/3 BC population ISI std.
7. I1-wave trough	24. L5 BC ISI	41. L5 PTN population ISI std.
8. I1-wave time-to-trough	25. L6 IT firing rate	42. L5 BC population ISI std.
9. I2-wave peak	26. L6 IT ISI	43. L6 IT population ISI std.
10. I2-wave time-to-peak	27. L2/3 IT peak/mean ratio	44. L6 BC population ISI std.
11. I2-wave trough	28. L2/3 BC peak/mean ratio	45. L2/3 IT noise weight
12. I2-wave time-to-trough	29. L5 PTN peak/mean ratio	46. L2/3 BC noise weight
13. I3-wave peak	30. L5 BC peak/mean ratio	47. L5 PTN noise weight
14. I3-wave time-to-peak	31. L6 IT peak/mean ratio	48. L5 BC noise weight
15. I3-wave trough	32. L6 BC peak/mean ratio	49. L6 IT noise weight
16. I3-wave time-to-trough	33. L2/3 IT baseline CV	50. L6 BC noise weight
17. L2/3 IT firing rate	34. L2/3 BC baseline CV	51. Amplitude after I3-wave

The baseline state constraints included both the mean population inter-spike interval (ISI) and the mean population firing rate for the different neuron types. Both objectives were important to constrain the

679 network activity due to the nature of their calculations. Firing rate was evaluated as the number of
680 spikes elicited within a time-window. However, there was a possibility that the ISIs within the window
681 were very small due to bursting behavior. Therefore, the mean ISI was added as an additional constraint.
682 Mean ISI alone was not a good constraint for overall activity because the calculation of relative error
683 resulted in lower error for small ISIs as opposed to large ISIs, which skewed the optimization to prefer
684 smaller ISIs and therefore higher firing rates. Including both constraints balanced the difference in bias
685 between them.

686 Experimental recordings from the epidural space of the cervical spine of human subjects during single
687 pulses of TMS were used to provide constraints for the corticospinal response to TMS. The peaks,
688 troughs, and latencies (time-to-peak and time-to-minimum) for each of the corticospinal waves—D-
689 wave (if available), I1-wave, I2-wave, and I3-wave—were measured and used as constraints. An
690 additional constraint minimized the peak of the model output beyond the time-window during which
691 the I3-wave should occur to prevent additional corticospinal waves, which were not present in the
692 recordings.

693 To reduce population synchrony, the population spiking density for a neuron type was constructed and
694 smoothed with a Gaussian kernel. The ratio between the maximum and the average value and the
695 coefficient of variation of the smoothed population spiking density were used as constraints with target
696 values of one and zero, respectively.

697 A possible aberrant network behavior resulted in spiking activity of the network being dominated by
698 large firing rates in a few neurons with the remaining neurons being silent. To avoid this, the standard
699 deviation of the mean population ISI within a neuron type was minimized to prevent highly skewed
700 distributions of activity.

Optimized Parameters

There were 98 open parameters for optimization. They can be divided into the following categories: Synaptic weights scalars, conduction velocity scalars, afferent delay mean, afferent delay standard deviation, proportion activated, noise amplitude, and noise rate. These categories and their bounds for optimization are summarized in Table 4. The specific names of all parameters are listed in S1 Appendix Table A-B.

Table 4. Categories of optimized parameters.

Name	Description	Range
Synaptic Weight Scalar (N. A.) 38 parameter)	Scalar multiplied to base synaptic weights	[0.1, 10]
Conduction Velocity Scalar (N. A.) 24 parameters	Scalar multiplied to conduction velocity	[0.25, 2]
Afferent Delay Mean (ms) 6 parameters	Mean conduction delay between afferent and postsynaptic neuron	[0.2, 2]
Afferent Delay Stdev. (ms) 6 parameters	Standard deviation of conduction delay between afferent and postsynaptic neuron	[0.1, 1]
Proportion Activated (N. A.) 12 parameters	Proportion of population made suprathreshold due to application of TMS	[0, 1]
Noise Amplitude (nA) 6 parameters	Amplitude of current to generate spiking activity due to independent noise	[1, 50]
Noise Rate (N. A.) 6 parameters	Scalar multiplied with the desired firing rate to determine the mean of the Poisson process used to generate noise	[0, 1]

Characterizing Optimization Robustness

The optimization was repeated four times with different random seeds to increase coverage of the parameter space and avoid local minimum solutions. Optimizations approached similar total error (S1 Appendix Fig D). To quantify the similarity of best solutions (i.e., lowest total error) found for each optimization run, the distance among parameters for the best solutions were computed using Euclidean

714 distance, normalized by the maximal possible distance (S1 Appendix Fig D) with overall distances being
715 17.4 to 19.6% from each other for D+ and D-, respectively. The relatively low distance (i.e., large
716 similarity) indicated that solutions lie within similar regions of the parameter space.

717 When identifying a dominating front, the large number of constraints resulted in every solution being
718 considered dominating. Therefore, the constraints were grouped by category and summed together to
719 reduce the dimensionality of the dominating front to four dimensions. The categories and the
720 corresponding objectives (based on the numbering from Table 3) are the following: Corticospinal wave
721 (1–16), spiking activity (17–26), synchrony (27–38), and well-behaved (39–51). The Pareto front is
722 visualized in S1 Appendix Fig E. The category error is plotted as a function of total error and showed that
723 corticospinal wave and baseline activity objectives were opposed. Generally, a solution that better
724 matched the experimentally-recorded corticospinal waves had a worse match with the desired baseline
725 activity.

726 Sensitivity Analysis between Model Parameters and Corticospinal Waves

727 The TVAT analysis investigated the synaptic weight and activation parameters for 42 total parameters
728 with 21 equally spaced values between 0 and the maximum boundary resulting in 861 unique
729 parameter-pairs with 441 values per pair. The total number of simulations for the sensitivity analysis
730 was 344,400. For each pair, the relationships between the two parameters and the amplitudes for each
731 corticospinal wave were approximated using linear regression with elastic net regularization and a third-
732 order polynomial model that included third-order interaction terms. Prior to the linear regression, the
733 corticospinal wave amplitudes were *standardized*, i.e., the mean was subtracted, and the variance
734 normalized to one. Because they were uniformly distributed across a grid, the parameters were
735 *normalized*, i.e., the minimum was subtracted, and the values divided by the parameter boundary range.
736 Regularization is a method of embedded feature selection that determines feature importance during

737 coefficient estimation and prevents overfitting. The optimal regularization parameters were determined
738 using 10-fold cross-validation. The open-source *scikit-learn* Python package was used to conduct the
739 regression and cross-validation [49].

740 The partial effect size of a parameter for a corticospinal wave was represented as the sum of the
741 absolute values of the coefficients of the polynomial models that involved the parameter. The total
742 effect size for a corticospinal wave was calculated as the sum of the effect sizes across all polynomial
743 models, i.e., across all pair-wise interactions, that included the parameter. Poor polynomial fits,
744 indicating that there may be little or no correlation between the parameters and the corticospinal wave
745 amplitude, were excluded from the summation. Only models with a coefficient of determination greater
746 than or equal to 0.5 were included.

747 Structural Analysis between Model Circuit and Corticospinal Wave Sensitivity

748 The cortical column circuit at the neuron population level can be represented as a weighted directed
749 graph with neuron types as nodes and connection between neuron types as edges. Given the effect sizes
750 revealed by the TVAT analysis, classifiers were used to identify any similarities in graph properties that
751 may exist to explain groupings of effect sizes, i.e., preferential versus non-preferential and corticospinal
752 wave preference. The goal was to identify the minimum set of features that would separate preferential
753 vs non-preferential nodes and then identify the corticospinal wave to which a preferential node had the
754 greatest effect.

755 Graph Metrics

756 Edge weights were specified using a variety of properties such as conduction delay and the log of the
757 connection probability. Because the relevant output of the network model was generated by the L5
758 PTNs, graph analysis was conducted using these neurons as a target or reference node. Graph analysis
759 was performed using the open-source *networkx* Python package [50]. All simple paths between a

starting node and the target node (L5 PTNs) were identified. Simple paths were defined as the sequence of nodes between the start and target that do not include repeat nodes along the path. The total conduction delay from a node to the target was computed as the sum of all conduction delays between nodes along the simple path, including synaptic transmission delays (0.2 ms). The total connection probability was computed as the sum of the logs of all connection probabilities between nodes along the simple path. Averages and standard deviations were also computed for these metrics. The out-degree (divergence), in-degree (convergence), and three centrality measures were calculated as well. Finally, the overall functional effect of the simple path was computed by first determining whether the simple path would have an overall excitatory effect (+1) or inhibitory effect (−1) on the L5 PTNs by multiplying successive functional effects along the simple path. The functional effects of each simple path were then weighted by the log of the path connection probability to compute the weighted average used to represent the overall functional effect of a node to the L5 PTNs. A summary and description of these metrics are in Table 5.

Training Classifiers

Two types of classifiers were used. Logistic regression was used to identify preferential vs non-preferential nodes. Support vector classification (SVC) with a radial basis function was used to perform a multiclass prediction to identify the corticospinal wave to which a preferential node had the greatest effect [51]. Classification, cross-fold validation, and regularization were performed using the *scikit-learn* Python package [49].

Input data consisted of the graph metrics described in Table 5. The inputs were *standardized*, i.e., the means were removed, and the variance was normalized to one. This was necessary for regularization during model estimation.

Due to the low numbers of samples for each class, the data was augmented by concatenating noisy versions of the original data. Noise was drawn from a normal distribution with zero mean and a standard deviation of 0.3, which represents 30% of the standard deviation of the standardized data.

Stratified 10-fold validation with 5 repeats was used to generate training and test sets for validation of the models. Stratified k-fold validation was chosen to allow for a balanced sampling of classes. The model performance was quantified using accuracy, computed as the number of true positives and true negatives divided by the total number of predictions. This validation strategy was performed for all the model evaluations described below.

Feature selection was performed using different methods for logistic regression versus SVC. Logistic regression used an embedded method, Lasso regularization, to eliminate non-predictive features. Lasso regularization minimizes the sum of the absolute value of all coefficients in addition to the mean squared error during model estimation which can result in the elimination of features as their coefficients drop to zero [52]. The weight of the Lasso regularization term was determined by grid-search and cross-validation. The remaining features were then used with Ridge regularization to perform the final classification. The weight of the Ridge regularization term was also determined by grid-search and cross-validation.

804 **Table 5. Description of graph metrics used to characterize the network.**

Name	Description
Convergence	In-degree of nodes / number of connected presynaptic neuron types.
Divergence	Out-degree of nodes / number of connected postsynaptic neuron types.
Total Simple Paths	Total number of unique simple paths for a node to L5 PTN.
Shortest Path Delay	Conduction delay of shortest path from node to L5 PTN.
Average Path Delay	Average path delay of all simple paths from a node to L5 PTN.
Weighted Average Path Delay	Weighted average of path delay of all simple paths from a node to L5 PTN using the log of the connection probability of the simple paths as weights.
Standard Deviation Path Delay	Standard deviation of path delays of all simple paths from a node to L5 PTN.
Weighted Standard Deviation Path Delay	Weighted standard deviation of path delays of all simple paths from a node to L5 PTN using the log of the connection probability of the simple paths as weights.
Connection Probability of Shortest Path	Connection probability of shortest path from a node to L5 PTN.
Average Connection Probability (Log)	Average of the log of the connection probabilities of all simple paths from a node to L5 PTN.
Standard Deviation Connection Probability (Log)	Standard deviation of the log of the connection probabilities of all simple paths from a node to L5 PTN.
Functional Effect	Overall excitatory/inhibitory effect of node on L5 PTN. For each simple path the excitatory/inhibitory effect of a node on the next node was represented as a +1 or -1. The effects of successive nodes were multiplied.
Weighted Functional Effect	Weighted average of the functional effect using the log of the connection probability of the simple paths as weights.
Closeness Centrality	Reciprocal of the average distance of the shortest paths between the node and all other nodes.
Betweenness Centrality	Ratio of the total number of shortest paths in the network to the number of paths that include node but do not end on the node.
Harmonic Centrality	Sum of the reciprocal of the sum of the shortest path distances between the node and all other nodes.

SVC does not support Lasso regularization, so recursive feature elimination was performed to identify the most predictive features [51]. During this procedure, an initial random sample of features was chosen, and the model was trained and evaluated. Then, models were trained while leaving one feature out. The model with the smallest decrease in performance indicated that the removed feature was not predictive and was eliminated from the feature set. This process was repeated with the remaining features until a single feature remained. Features were ranked by the number of times a feature was the sole remainder after the elimination process and divided by the total number of times the feature was included in the initial random sample. Recursive feature elimination was repeated 100 times with 5 random features chosen for each iteration. The regularization weight and the scale factor for the radial basis functions were determined using grid search and cross-validation. The final classifier was trained using Ridge regularization.

ACKNOWLEDGEMENTS

The authors thank Dr. Aman Aberra for preliminary work on the I-wave model, and the Duke Compute Cluster team for computational support.

REFERENCES

1. Lefaucheur JP, Aleman A, Baeken C, Benninger DH, Brunelin J, Di Lazzaro V, et al. Evidence-based guidelines on the therapeutic use of repetitive transcranial magnetic stimulation (rTMS): An update (2014–2018). Clinical Neurophysiology. 2020 Feb 1;131(2):474–528.

2. Siebner HR, Funke K, Aberra AS, Antal A, Bestmann S, Chen R, et al. Transcranial magnetic stimulation of the brain: What is stimulated? – A consensus and critical position paper. Clinical Neurophysiology. 2022 Aug 1;140:59–97.

827 3. Nudo RJ, Masterton RB. Descending pathways to the spinal cord, III: Sites of origin of the
828 corticospinal tract. Journal of Comparative Neurology. 1990;296(4):559–83.

829 4. Di Lazzaro V, Ziemann U, Lemon RN. State of the art: Physiology of transcranial motor cortex
830 stimulation. Brain Stimulation. 2008;1(4):345–62.

831 5. Di Lazzaro V, Profice P, Ranieri F, Capone F, Dileone M, Oliviero A. I-wave origin and modulation.
832 Brain Stimulation. 2012;5(4):512–25.

833 6. Ziemann U. I-waves in motor cortex revisited. Exp Brain Res. 2020 Aug 1;238(7):1601–10.

834 7. Esser SK, Hill SL, Tononi G. Modeling the Effects of Transcranial Magnetic Stimulation on Cortical
835 Circuits. Journal of Neurophysiology. 2005 Jul;94(1):622–39.

836 8. Rusu CV, Murakami M, Ziemann U, Triesch J. A Model of TMS-induced I-waves in Motor Cortex.
837 Brain Stimulation. 2014 May;7(3):401–14.

838 9. Pfeffer CK, Xue M, He M, Huang ZJ, Scanziani M. Inhibition of inhibition in visual cortex: the logic of
839 connections between molecularly distinct interneurons. Nat Neurosci. 2013 Aug;16(8):1068–76.

840 10. Carnevale NT, Hines ML. The NEURON Book [Internet]. Cambridge: Cambridge University Press;
841 2006 [cited 2023 Aug 15]. Available from: https://www.cambridge.org/core/books/neuron-
842 book/7C8D9BD861D288E658BEB652F593F273

843 11. Saltelli A. Sensitivity analysis: Could better methods be used? Journal of Geophysical Research:
844 Atmospheres. 1999;104(D3):3789–93.

845 12. Saltelli A, Annoni P. How to avoid a perfunctory sensitivity analysis. Environmental Modelling &
846 Software. 2010 Dec 1;25(12):1508–17.

13. Ziemann U, Rothwell JC. I-Waves in Motor Cortex. Journal of Clinical Neurophysiology. 2000 Jul;17(4):397.

14. Woodforth IJ, Hicks RG, Crawford MR, Stephen JPH, Burke D. Depression of I Waves in Corticospinal Volleys by Sevoflurane, Thiopental, and Propofol. Anesthesia & Analgesia. 1999 Nov;89(5):1182.

15. Di Lazzaro V, Oliviero A, Meglio M, Cioni B, Tamburrini G, Tonali P, et al. Direct demonstration of the effect of lorazepam on the excitability of the human motor cortex. Clinical Neurophysiology. 2000 May 1;111(5):794–9.

16. Di Lazzaro V, Rothwell J, Capogna M. Noninvasive Stimulation of the Human Brain: Activation of Multiple Cortical Circuits. Neuroscientist. 2018 Jun 1;24(3):246–60.

17. Li B, Virtanen JP, Oeltermann A, Schwarz C, Giese MA, Ziemann U, et al. Lifting the veil on the dynamics of neuronal activities evoked by transcranial magnetic stimulation. Rotenberg A, editor. eLife. 2017 Nov 22;6:e30552.

18. Mueller JK, Grigsby EM, Prevosto V, Petraglia FW, Rao H, Deng ZD, et al. Simultaneous transcranial magnetic stimulation and single-neuron recording in alert non-human primates. Nat Neurosci. 2014 Aug;17(8):1130–6.

19. Hao Y, Riehle A, Brochier TG. Mapping Horizontal Spread of Activity in Monkey Motor Cortex Using Single Pulse Microstimulation. Frontiers in Neural Circuits [Internet]. 2016 [cited 2023 Aug 15];10. Available from: https://www.frontiersin.org/articles/10.3389/fncir.2016.00104

20. Maier MA, Kirkwood PA, Brochier T, Lemon RN. Responses of single corticospinal neurons to intracortical stimulation of primary motor and premotor cortex in the anesthetized macaque monkey. Journal of Neurophysiology. 2013 Jun 15;109(12):2982–98.

21. Yun R, Mishler JH, Perlmutter SI, Rao RPN, Fetz EE. Responses of Cortical Neurons to Intracortical Microstimulation in Awake Primates. eNeuro [Internet]. 2023 Apr 1 [cited 2023 Jul 3];10(4). Available from: https://www.eneuro.org/content/10/4/ENEURO.0336-22.2023

22. Aberra AS, Peterchev AV, Grill WM. Biophysically realistic neuron models for simulation of cortical stimulation. J Neural Eng. 2018 Oct;15(6):066023.

23. Di Lazzaro V, Rothwell JC. Corticospinal activity evoked and modulated by non-invasive stimulation of the intact human motor cortex. The Journal of Physiology. 2014;592(19):4115–28.

24. Hooks BM, Mao T, Gutnisky DA, Yamawaki N, Svoboda K, Shepherd GMG. Organization of cortical and thalamic input to pyramidal neurons in mouse motor cortex. Journal of Neuroscience. 2013 Jan 9;33(2):748–60.

25. Amassian VE, Stewart M, Quirk GJ, Rosenthal JL. Physiological basis of motor effects of a transient stimulus to cerebral cortex. Neurosurgery. 1987 Jan 1;20(1):74–93.

26. Narayanan RT, Egger R, Johnson AS, Mansvelder HD, Sakmann B, de Kock CPJ, et al. Beyond Columnar Organization: Cell Type- and Target Layer-Specific Principles of Horizontal Axon Projection Patterns in Rat Vibrissal Cortex. Cerebral Cortex. 2015 Nov 1;25(11):4450–68.

27. Chadderdon GL, Mohan A, Suter BA, Neymotin SA, Kerr CC, Francis JT, et al. Motor Cortex Microcircuit Simulation Based on Brain Activity Mapping. Neural Computation. 2014 Jul 1;26(7):1239–62.

887 28. García-Cabezas MÁ, Barbas H. Area 4 has layer IV in adult primates. European Journal of
888 Neuroscience. 2014;39(11):1824–34.

889 29. Yamawaki N, Borges K, Suter BA, Harris KD, Shepherd GMG. A genuine layer 4 in motor cortex with
890 prototypical synaptic circuit connectivity. Nelson SB, editor. eLife. 2014 Dec 19;3:e05422.

891 30. Callaway EM, Dong HW, Ecker JR, Hawrylycz MJ, Huang ZJ, Lein ES, et al. A multimodal cell census
892 and atlas of the mammalian primary motor cortex. Nature. 2021 Oct 6;598(7879):86–102.

893 31. Dura-Bernal S, Neymotin SA, Suter BA, Dacre J, Moreira JVS, Urdapilleta E, et al. Multiscale model
894 of primary motor cortex circuits predicts in vivo cell-type-specific, behavioral state-dependent
895 dynamics. Cell Reports. 2023 Jun 27;42(6):112574.

896 32. Aberra AS, Wang B, Grill WM, Peterchev AV. Simulation of transcranial magnetic stimulation in
897 head model with morphologically-realistic cortical neurons. Brain Stimulation. 2020 Jan
898 1;13(1):175–89.

899 33. Ziemann U, Tergau F, Wischer S, Hildebrandt J, Paulus W. Pharmacological control of facilitatory I-
900 wave interaction in the human motor cortex. A paired transcranial magnetic stimulation study.
901 Electroencephalography and Clinical Neurophysiology/Electromyography and Motor Control. 1998
902 Aug 1;109(4):321–30.

903 34. Paulus W, Classen J, Cohen LG, Large CH, Di Lazzaro V, Nitsche M, et al. State of the art:
904 Pharmacologic effects on cortical excitability measures tested by transcranial magnetic
905 stimulation. Brain Stimulation. 2008 Jul 1;1(3):151–63.

906 35. Premoli I, Rivolta D, Espenhahn S, Castellanos N, Belardinelli P, Ziemann U, et al. Characterization
907 of GABAB-receptor mediated neurotransmission in the human cortex by paired-pulse TMS–EEG.
908 NeuroImage. 2014 Dec 1;103:152–62.

909 36. Yamashita T, Pala A, Pedrido L, Kremer Y, Welker E, Petersen CCH. Membrane Potential Dynamics
910 of Neocortical Projection Neurons Driving Target-Specific Signals. Neuron. 2013 Dec
911 18;80(6):1477–90.

912 37. Mountcastle VB. The columnar organization of the neocortex. Brain. 1997 Apr 1;120(4):701–22.

913 38. Wagstyl K, Larocque S, Cucurull G, Lepage C, Cohen JP, Bludau S, et al. BigBrain 3D atlas of cortical
914 layers: Cortical and laminar thickness gradients diverge in sensory and motor cortices. PLOS
915 Biology. 2020 Apr 3;18(4):e3000678.

916 39. Peters A. The Morphology of Minicolumns. In: Blatt GJ, editor. The Neurochemical Basis of Autism:
917 From Molecules to Minicolumns [Internet]. Boston, MA: Springer US; 2010 [cited 2023 Jun 22]. p.
918 45–68. Available from: https://doi.org/10.1007/978-1-4419-1272-5_4

919 40. Favorov OV, Diamond ME. Demonstration of discrete place-defined columns—segregates—in the
920 cat SI. Journal of Comparative Neurology. 1990;298(1):97–112.

921 41. Muller L, Reynaud A, Chavane F, Destexhe A. The stimulus-evoked population response in visual
922 cortex of awake monkey is a propagating wave. Nat Commun. 2014 Apr 28;5(1):3675.

923 42. Houghton C, Kreuz T. On the efficient calculation of van Rossum distances. Network: Computation
924 in Neural Systems. 2012 Mar 1;23(1–2):48–58.

925 43. Satopaa V, Albrecht J, Irwin D, Raghavan B. Finding a “Kneedle” in a Haystack: Detecting Knee
926 Points in System Behavior. In: 2011 31st International Conference on Distributed Computing

Systems Workshops [Internet]. 2011 [cited 2023 Nov 16]. p. 166–71. Available from: https://ieeexplore.ieee.org/abstract/document/5961514

44. Tonda A. Inspyred: Bio-inspired algorithms in Python. Genet Program Evolvable Mach. 2020 Jun 1;21(1):269–72.

45. Kucherenko S, Albrecht D, Saltelli A. Exploring multi-dimensional spaces: a Comparison of Latin Hypercube and Quasi Monte Carlo Sampling Techniques [Internet]. arXiv; 2015 [cited 2023 Aug 15]. Available from: http://arxiv.org/abs/1505.02350

46. Gad AG. Particle Swarm Optimization Algorithm and Its Applications: A Systematic Review. Arch Computat Methods Eng. 2022 Aug 1;29(5):2531–61.

47. Xu G, Luo K, Jing G, Yu X, Ruan X, Song J. On convergence analysis of multi-objective particle swarm optimization algorithm. European Journal of Operational Research. 2020 Oct 1;286(1):32–8.

48. Xu S, Rahmat-Samii Y. Boundary Conditions in Particle Swarm Optimization Revisited. IEEE Transactions on Antennas and Propagation. 2007 Mar;55(3):760–5.

49. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. J Mach Learn Res. 2011 Nov 1;12(null):2825–30.

50. Hagberg A, Swart PJ, Schult DA. Exploring network structure, dynamics, and function using NetworkX [Internet]. Los Alamos National Laboratory (LANL), Los Alamos, NM (United States); 2008 Jan [cited 2023 Aug 15]. Report No.: LA-UR-08-05495; LA-UR-08-5495. Available from: https://www.osti.gov/biblio/960616

51. Guyon I, Weston J, Barnhill S, Vapnik V. Gene Selection for Cancer Classification using Support Vector Machines. Machine Learning. 2002 Jan 1;46(1):389–422.

948 52. Friedman J, Hastie T, Tibshirani R. Regularization Paths for Generalized Linear Models via

949 Coordinate Descent. J Stat Softw. 2010;33(1):1–22.

950

951

952

953

954

955

S1 APPENDIX

Fig A

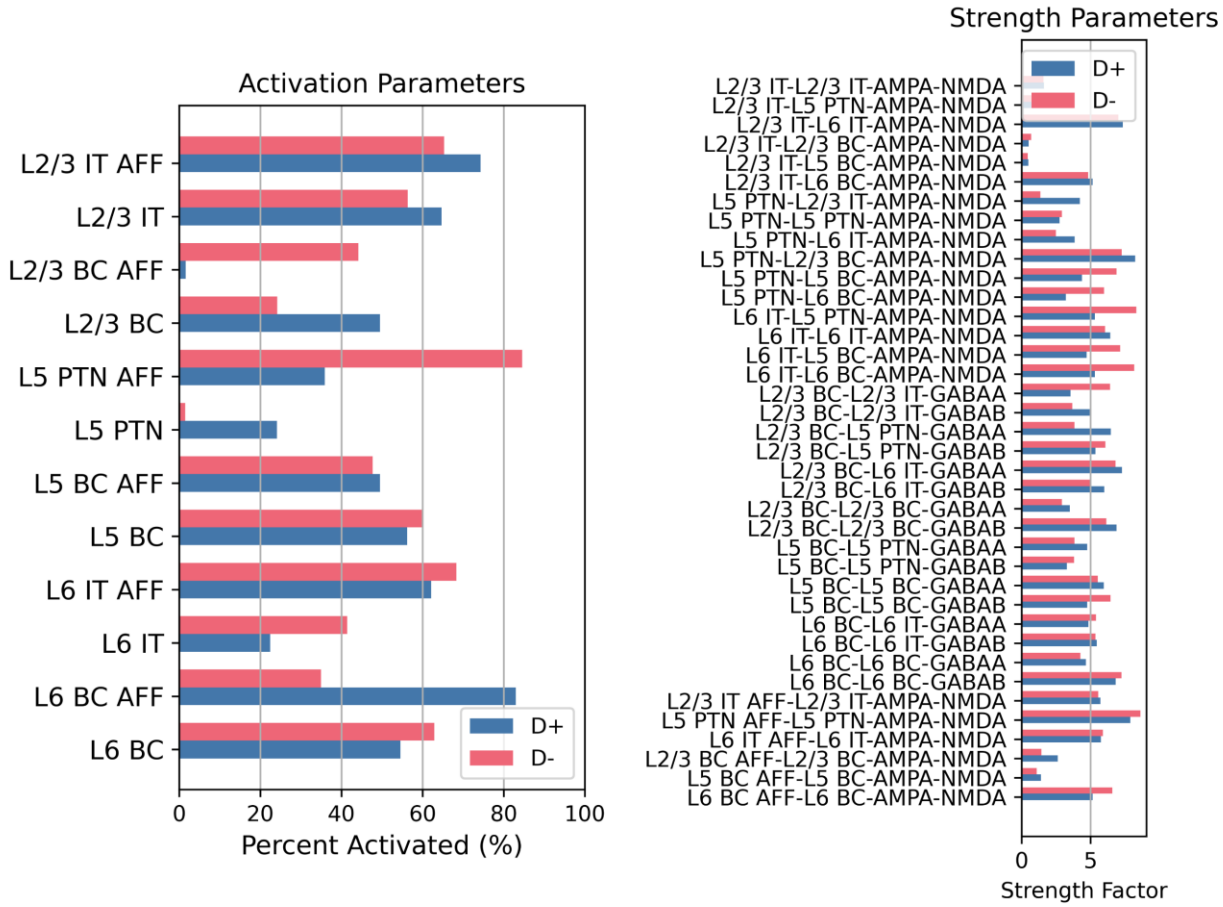

Fig A. Activation and synaptic weight scalar parameters for best D+ and D− models.

IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.

962 Fig B

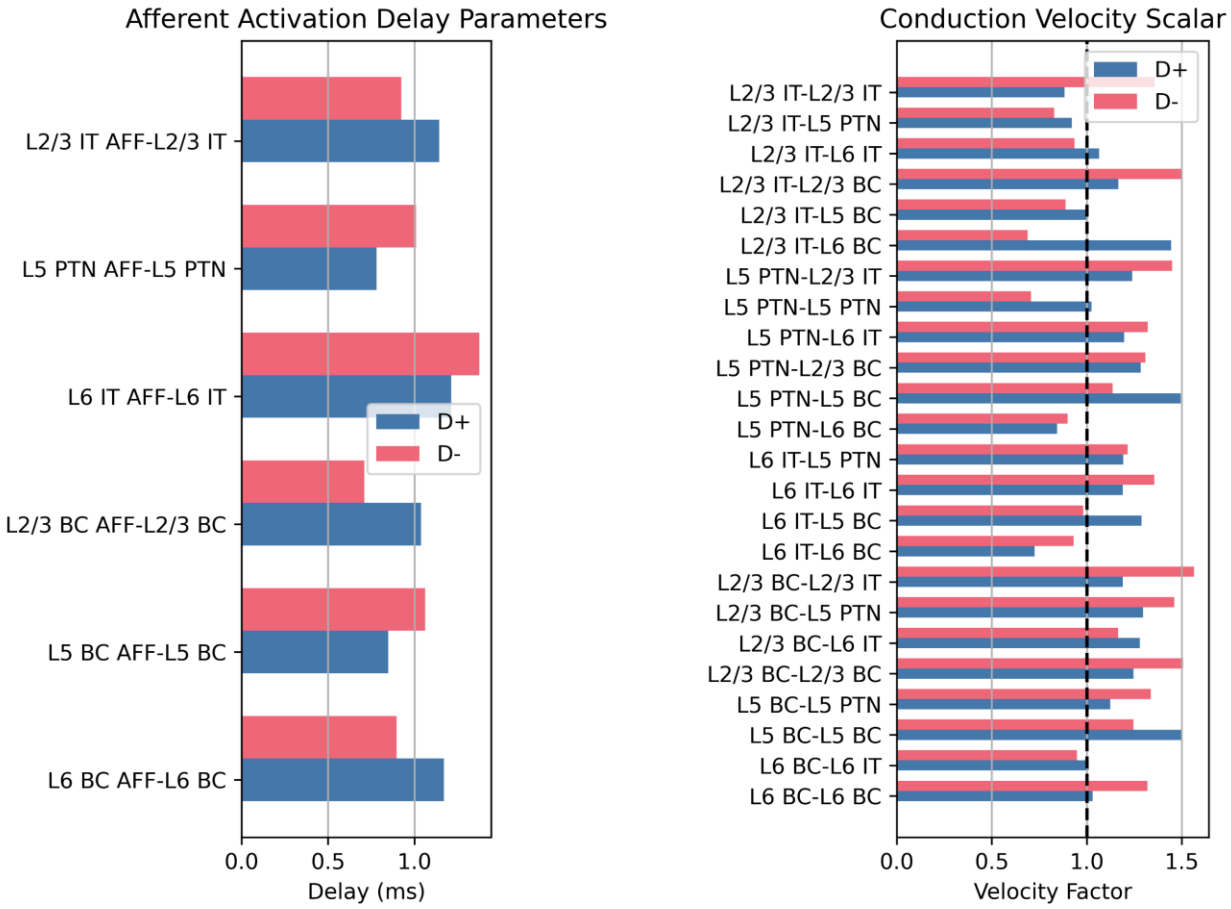

963

964 **Fig B. Afferent activation delay and conduction velocity scalars for best D+ and D− models.**

965 *IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.*

966

967

968

969

970

971 Fig C

972

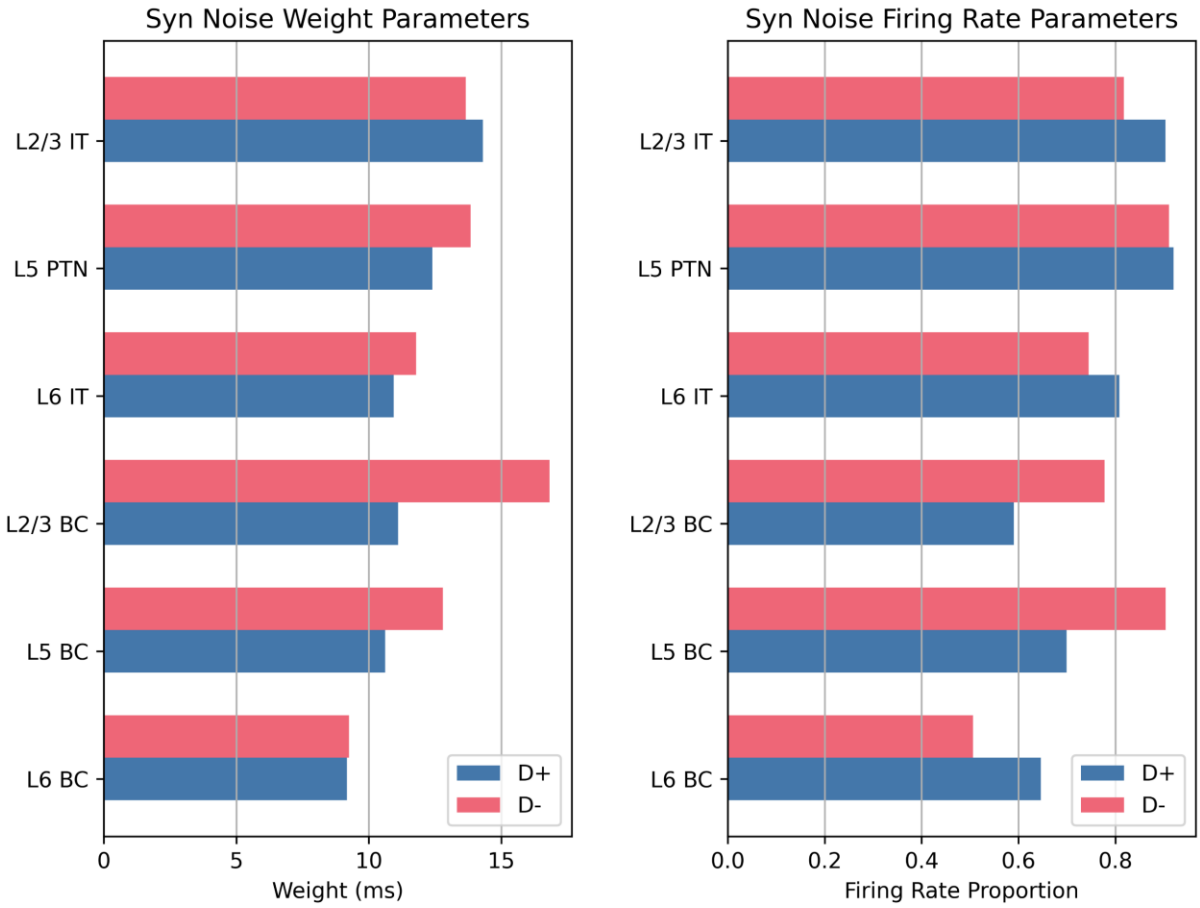

973

974 **Fig C. Synaptic noise weight and synaptic noise firing rates for best D+ and D− models.**

975 *IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.*

976

977

978

979

Fig D

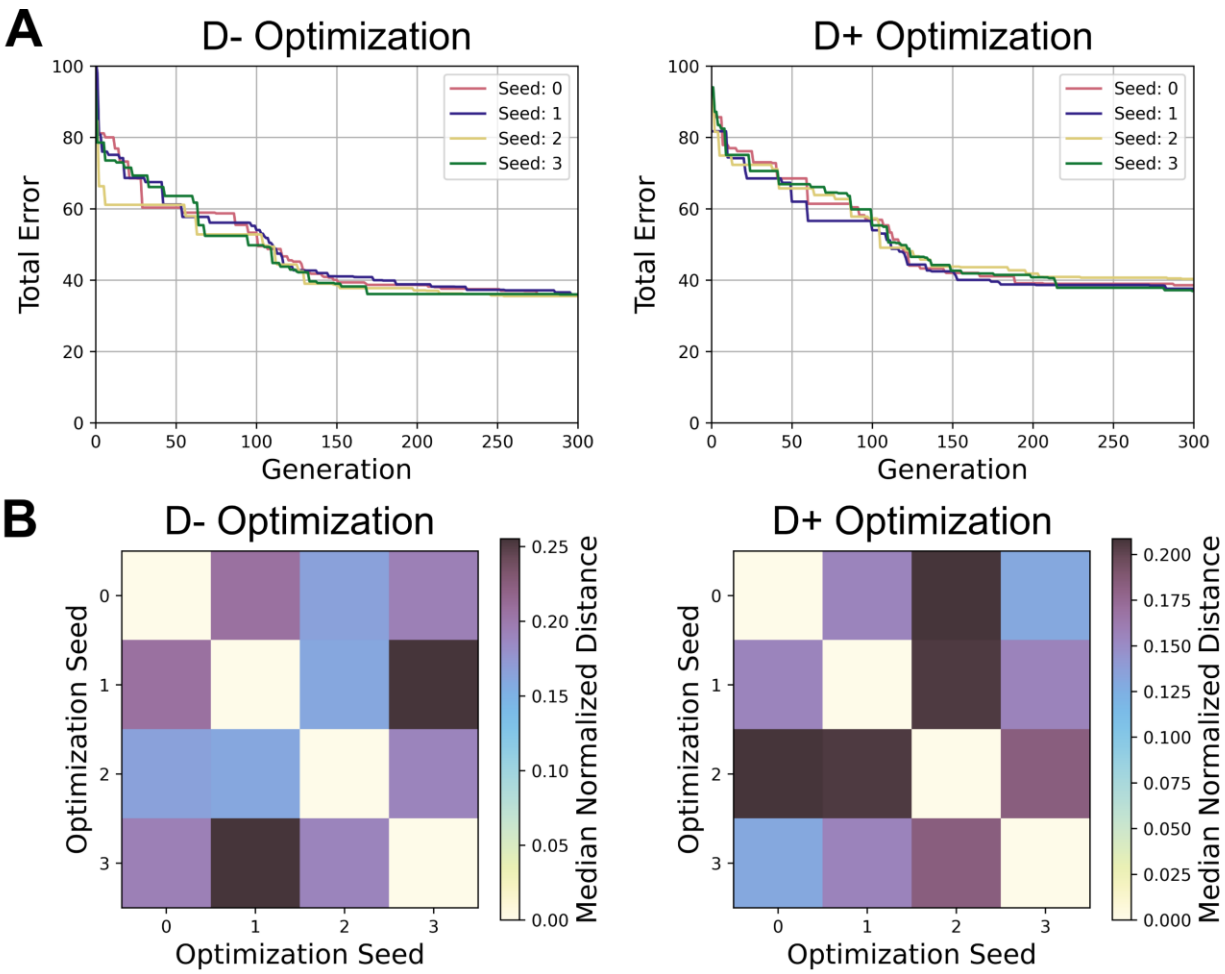

Fig D. Characterizations of convergence from optimization.

A) The cumulative lowest total error is plotted as a function of optimization iteration. Four evolution seeds were run for each responder type (D^+ and D^-). All seeds converged to errors of similar magnitude. B) The normalized distances of the parameters of the best solutions for each optimization run. The Euclidean distance of the best solutions was normalized by the maximum possible distance given the bounds of the explored parameter space. The diagonals of the matrix are zero because they represent the distance between a solution and itself. The overall normalized distances were 17.4 to 19.6% for D+ and D-, respectively.

990 Fig E

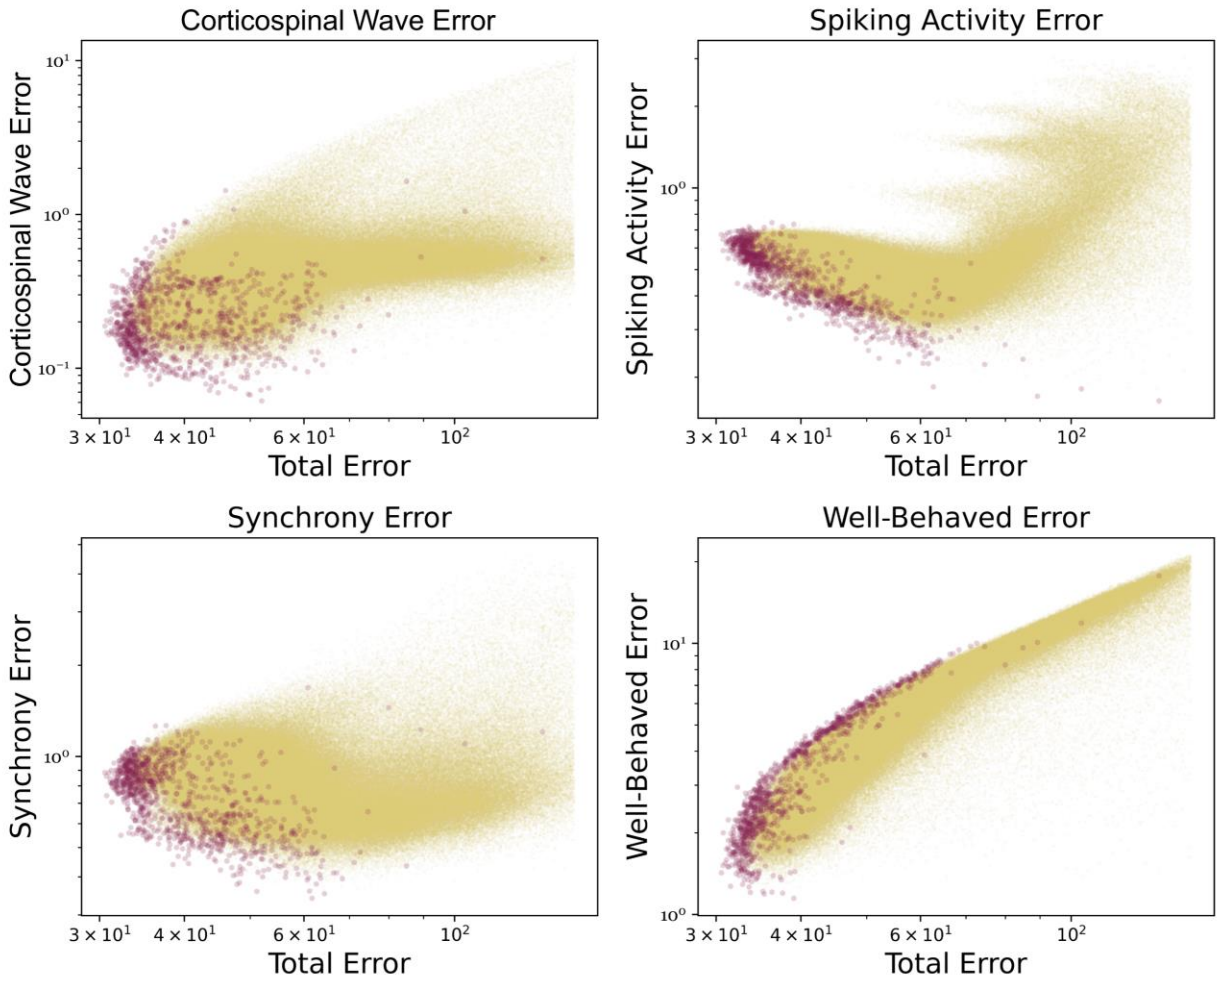

991

992 **Fig E. Visualization of reduced pareto front.**

993 *Objectives were grouped into categories, and their combined error was plotted against the total error.*

994 *Red dots indicate particles that were pareto dominant, and yellow particles indicate the remaining*

995 *particles. The correlation between the category error and the total error is a representation of the pareto*

996 *front and how the category error changes as total error is minimized.*

997 Fig F

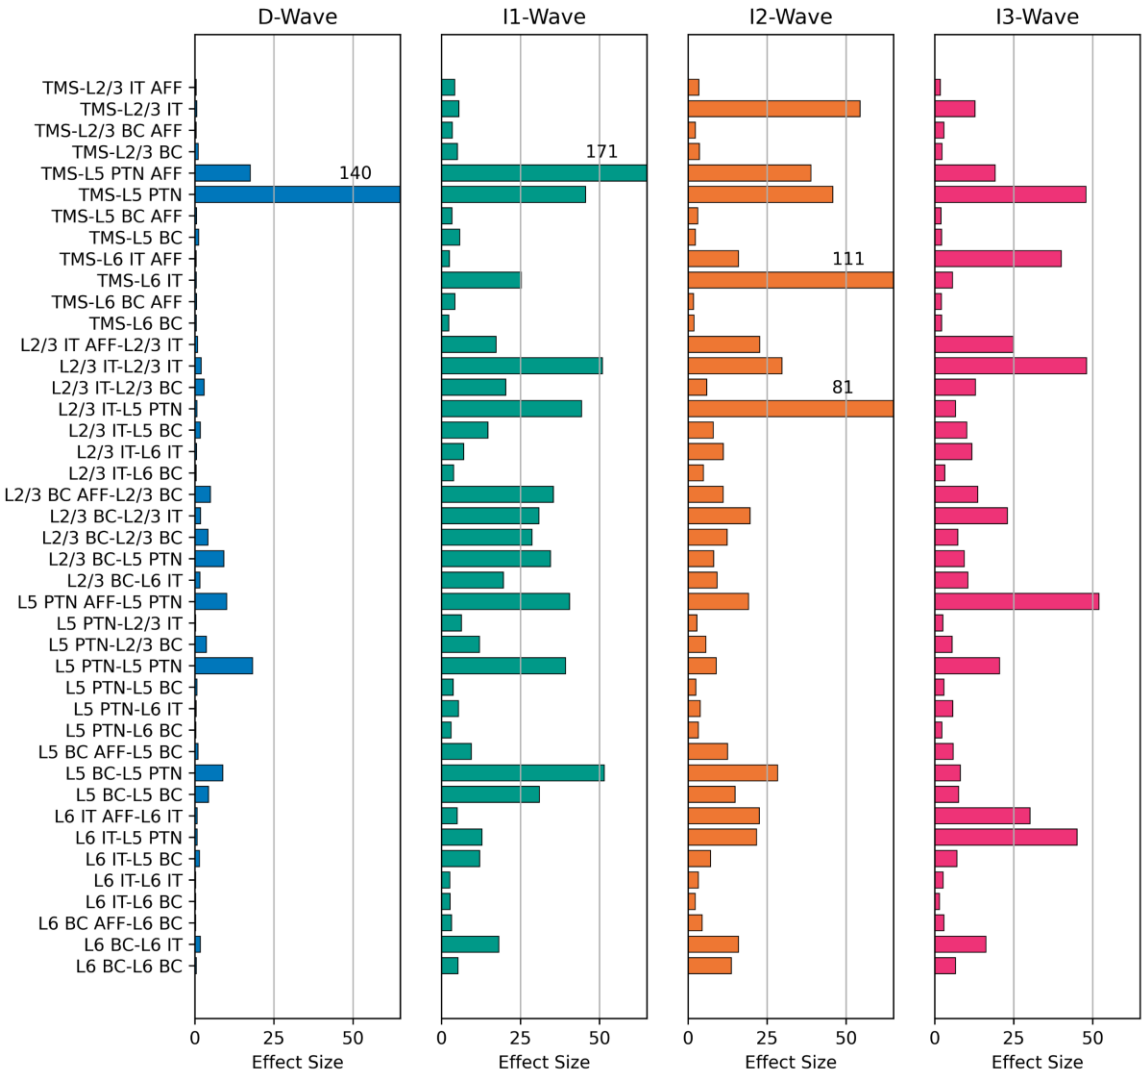

998

999 **Fig F. Effect sizes for all parameters and all waves.**

1000 *Effect size was calculated as the integrals of the absolute values of the partial derivatives of the*

1001 *polynomial fits to the TVAT surfaces. Effect sizes were not normalized, and the x-axis maximum was*

1002 *chosen to allow visualization of the smaller effect sizes. Effect sizes that are greater than the x-axis*

1003 *maximum have their values listed above their corresponding bars. Y-axis labels are shared across*

1004 *subplots. IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.*

1005 Fig G

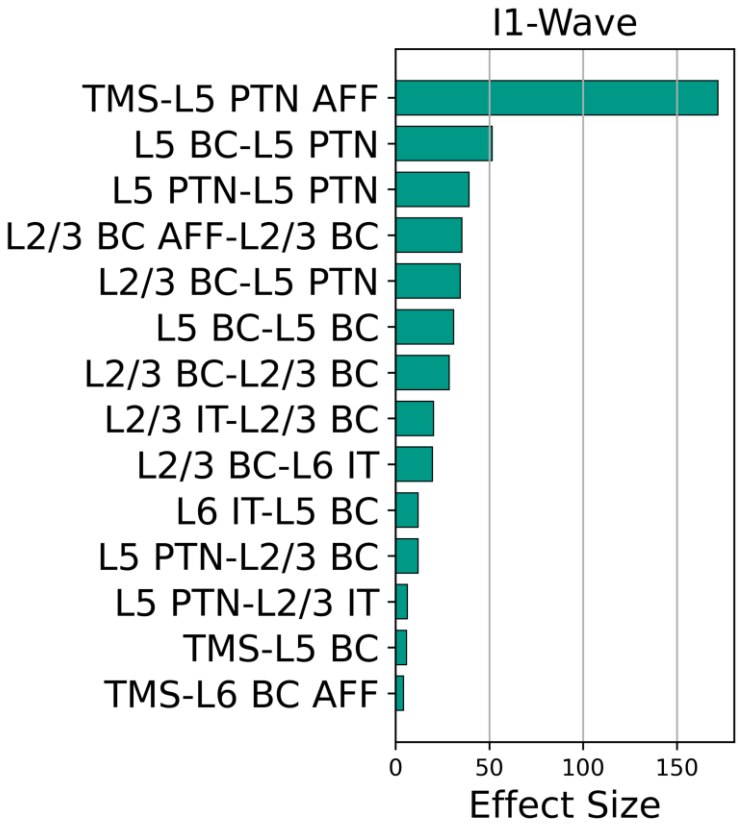

1006

1007 **Fig G. Effect size ranking for all parameters that preferentially affected I1-wave amplitude.**

1008 *IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.*

1009

1010

1011

1012

1013 Fig H

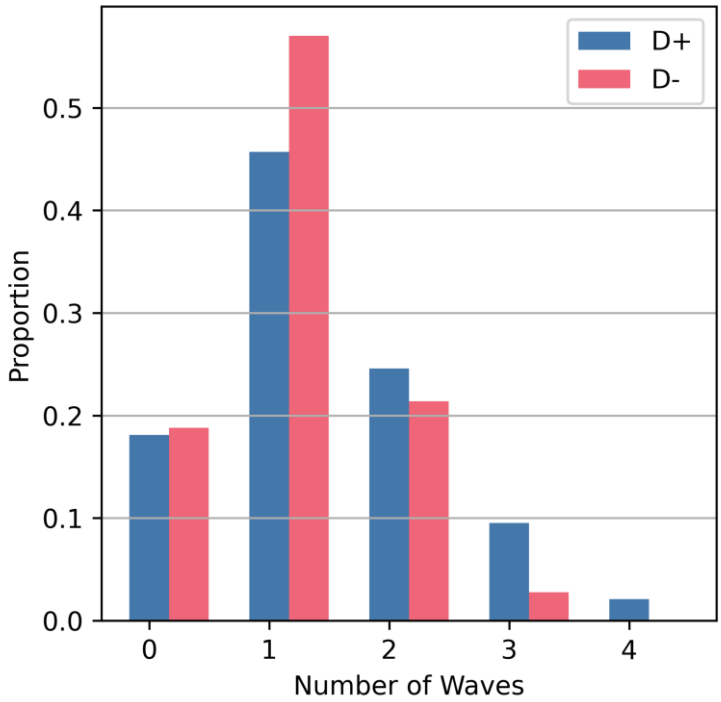

1014

1015 **Fig H. Histogram of number of waves for which L5 PTNs contributed a spike.**

1016 For each stimulus presentation, the spikes generated by each L5 PTN were divided based on the time

1017 windows for each corticospinal wave, and the total number of wave time windows during which spiking

1018 occurred was counted.

1019

1020

1021

1022

1023

1024

Table A

Table A. List of Optimized Synaptic Weight Parameters

Synaptic Weight Parameters	
L2/3 IT to L2/3 IT (AMPAR/NMDAR)	L2/3 BC to L5 PTN ($GABA_BR$)
L2/3 IT to L5 PTN (AMPAR/NMDAR)	L2/3 BC to L6 IT ($GABA_AR$)
L2/3 IT to L6 IT (AMPAR/NMDAR)	L2/3 BC to L6 IT ($GABA_BR$)
L2/3 IT to L2/3 BC (AMPAR/NMDAR)	L2/3 BC to L2/3 BC ($GABA_AR$)
L2/3 IT to L5 BC (AMPAR/NMDAR)	L2/3 BC to L2/3 BC ($GABA_BR$)
L2/3 IT to L6 BC (AMPAR/NMDAR)	L5 BC to L5 PTN ($GABA_AR$)
L5 PTN to L2/3 IT (AMPAR/NMDAR)	L5 BC to L5 PTN ($GABA_BR$)
L5 PTN to L5 PTN (AMPAR/NMDAR)	L5 BC to L5 BC ($GABA_AR$)
L5 PTN to L6 IT (AMPAR/NMDAR)	L5 BC to L5 BC ($GABA_BR$)
L5 PTN to L2/3 BC (AMPAR/NMDAR)	L6 BC to L6 IT ($GABA_AR$)
L5 PTN to L5 BC (AMPAR/NMDAR)	L6 BC to L6 IT ($GABA_BR$)
L5 PTN to L6 BC (AMPAR/NMDAR)	L6 BC to L6 BC ($GABA_AR$)
L6 IT to L5 PTN (AMPAR/NMDAR)	L6 BC to L6 BC ($GABA_BR$)
L6 IT to L6 IT (AMPAR/NMDAR)	L2/3 IT AFF to L2/3 IT (AMPAR/NMDAR)
L6 IT to L5 BC (AMPAR/NMDAR)	L2/3 BC AFF to L2/3 BC (AMPAR/NMDAR)
L6 IT to L6 BC (AMPAR/NMDAR)	L5 PTN AFF to L5 PTN (AMPAR/NMDAR)
L2/3 BC to L2/3 IT ($GABA_AR$)	L5 BC AFF to L5 BC (AMPAR/NMDAR)
L2/3 BC to L2/3 IT ($GABA_BR$)	L6 IT AFF to L6 IT (AMPAR/NMDAR)
L2/3 BC to L5 PTN ($GABA_AR$)	L6 BC AFF to L6 BC (AMPAR/NMDAR)

IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.

1032 # Table B

1033 **Table B. List of Optimized Delay, Activation, and Noise Parameters**

Parameters	
Conduction velocity scalar L2/3 IT to L2/3 IT	Activation propagation delay stdev. L2/3 IT AFF
Conduction velocity scalar L2/3 IT to L5 PTN	Activation propagation delay stdev. L2/3 BC AFF
Conduction velocity scalar L2/3 IT to L6 IT	Activation propagation delay stdev. L5 PTN AFF
Conduction velocity scalar L2/3 IT to L2/3 BC	Activation propagation delay stdev. L5 BC AFF
Conduction velocity scalar L2/3 IT to L5 BC	Activation propagation delay stdev. L6 IT AFF
Conduction velocity scalar L2/3 IT to L6 BC	Activation propagation delay stdev. L6 BC AFF
Conduction velocity scalar L5 PTN to L2/3 IT	Proportion activated L2/3 IT
Conduction velocity scalar L5 PTN to L5 PTN	Proportion activated L2/3 BC
Conduction velocity scalar L5 PTN to L6 IT	Proportion activated L5 PTN
Conduction velocity scalar L5 PTN to L2/3 BC	Proportion activated L5 BC
Conduction velocity scalar L5 PTN to L5 BC	Proportion activated L6 IT
Conduction velocity scalar L5 PTN to L6 BC	Proportion activated L6 BC
Conduction velocity scalar L6 IT to L5 PTN	Proportion activated L2/3 IT AFF
Conduction velocity scalar L6 IT to L6 IT	Proportion activated L2/3 BC AFF
Conduction velocity scalar L6 IT to L5 BC	Proportion activated L5 PTN AFF
Conduction velocity scalar L6 IT to L6 BC	Proportion activated L5 BC AFF
Conduction velocity scalar L2/3 BC to L2/3 IT	Proportion activated L6 IT AFF
Conduction velocity scalar L2/3 BC to L5 PTN	Proportion activated L6 BC AFF
Conduction velocity scalar L2/3 BC to L6 IT	Noise weight L2/3 IT
Conduction velocity scalar L2/3 BC to L2/3 BC	Noise weight L2/3 BC
Conduction velocity scalar L5 BC to L5 PTN	Noise weight L5 PTN
Conduction velocity scalar L5 BC to L5 BC	Noise weight L5 BC
Conduction velocity scalar L6 BC to L6 IT	Noise weight L6 IT
Conduction velocity scalar L6 BC to L6 BC	Noise weight L6 BC
Activation propagation delay L2/3 IT AFF	Noise rate L2/3 IT
Activation propagation delay mean L2/3 BC AFF	Noise rate L2/3 BC

Activation propagation delay mean L5 PTN AFF	Noise rate L5 PTN
Activation propagation delay mean L5 BC AFF	Noise rate L5 BC
Activation propagation delay mean L6 IT AFF	Noise rate L6 IT
Activation propagation delay mean L6 BC AFF	Noise rate L6 BC

IT: Intratelencephalic neuron. PTN: Pyramidal tract neuron. BC: Basket cell. AFF: Afferent.