Title

Targeting neuroplasticity to improve motor recovery after stroke: an artificial neural network model.

Journal

Brain Communications, 4(6)

Authors

Norman, Sumner
Wolpaw, Jonathan
Reinkensmeyer, David

Publication Date

2022

DOI

10.1093/braincomms/fcac264

Peer reviewed

BRAIN COMMUNICATIONS

Targeting neuroplasticity to improve motor recovery after stroke: an artificial neural network model

Sumner L. Norman,[1,2] Jonathan R. Wolpaw[3] and David J. Reinkensmeyer[4]

After a neurological injury, people develop abnormal patterns of neural activity that limit motor recovery. Traditional rehabilitation, which concentrates on practicing impaired skills, is seldom fully effective. New targeted neuroplasticity protocols interact with the central nervous system to induce beneficial plasticity in key sites and thereby enable wider beneficial plasticity. They can complement traditional therapy and enhance recovery. However, their development and validation is difficult because many different targeted neuroplasticity protocols are conceivable, and evaluating even one of them is lengthy, laborious, and expensive. Computational models can address this problem by triaging numerous candidate protocols rapidly and effectively. Animal and human empirical testing can then concentrate on the most promising ones. Here, we simulate a neural network of corticospinal neurons that control motoneurons eliciting unilateral finger extension. We use this network to (i) study the mechanisms and patterns of cortical reorganization after a stroke; and (ii) identify and parameterize a targeted neuroplasticity protocol that improves recovery of extension torque. After a simulated stroke, standard training produced abnormal bilateral cortical activation and suboptimal torque recovery. To enhance recovery, we interdigitated standard training with trials in which the network was given feedback only from a targeted population of sub-optimized neurons. Targeting neurons in secondary motor areas on ~20% of the total trials restored lateralized cortical activation and improved recovery of extension torque. The results illuminate mechanisms underlying suboptimal cortical activity post-stroke; they enable the identification and parameterization of the most promising targeted neuroplasticity protocols. By providing initial guidance, computational models could facilitate and accelerate the realization of new therapies that improve motor recovery.

1 Biology and Biological Engineering, California Institute of Technology, Pasadena, CA 91125, USA
2 Mechanical and Aerospace Engineering, University of California: Irvine, Irvine, CA 92697, USA
3 National Center for Adaptive Neurotechnologies, Stratton VA Medical Center and State University of New York, Albany, NY 12208, USA
4 Mechanical and Aerospace Engineering, Anatomy and Neurobiology, University of California: Irvine, Irvine, CA 92697, USA

Correspondence to: Sumner L. Norman, Ph.D.
1200 E California Blvd, MC 216-76
Pasadena, CA 91125, USA
E-mail: sumnern@caltech.edu

Keywords: stroke; motor control; reinforcement learning; activity-dependent plasticity; rehabilitation

Abbreviations: CS = corticospinal; dPM = dorsal premotor cortex; fMRI = functional MRI; M1 = primary motor cortex; MN = motoneuronal; SCI = spinal cord injury; TMS = transcranial magnetic stimulation; TNP = targeted neuroplasticity

Received August 16, 2021. Revised June 18, 2022. Accepted October 19, 2022. Advance access publication October 21, 2022
Published by Oxford University Press on behalf of the Guarantors of Brain 2022.

Graphical Abstract

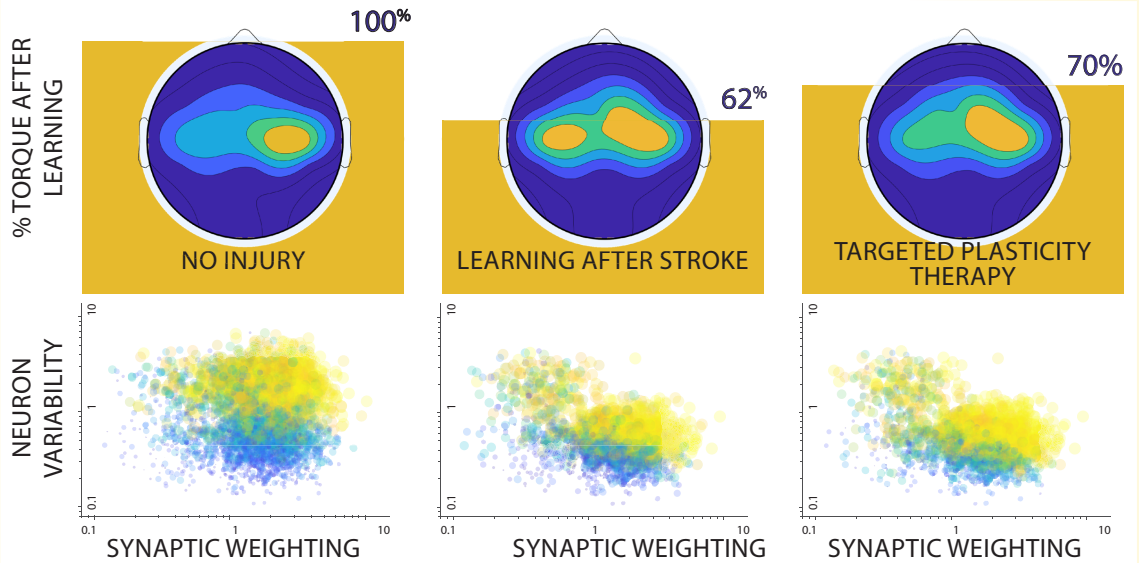

Using a neural network model of corticospinal motor function, *Norman et al.* show non-normative cortical reorganization and decreased motor function after a simulated stroke and how targeted neural feedback could enhance neuroplasticity, restore cortical organization and improve motor recovery.

Introduction

Activity-dependent neuroplasticity occurs throughout life, affecting the CNS from cortex to spinal cord.[1–4] When trauma or disease [e.g. stroke, spinal cord injury (SCI)] impairs motor function, traditional rehabilitation concentrates on intensive practice of the impaired motor skills. Although this usually produces some recovery, significant disability often remains. Thus, the present challenge is to guide CNS plasticity to maximize functional recovery.[3–6] Targeted neuroplasticity (TNP) protocols are an innovative approach to addressing this challenge.[7–13]

A TNP protocol creates a sensorimotor interaction with the CNS that induces activity-dependent plasticity at a key site (e.g. a particular spinal reflex pathway, a specific region in motor cortex).[7–11] This plasticity improves function. By doing so, the targeted plasticity enables activity that produces wider beneficial plasticity at other important sites.[10] For example, after incomplete SCI, a TNP protocol that weakens a hyperactive spinal reflex can reduce the ankle clonus or foot-drop that prevents effective locomotor practice; it can thereby enable more effective practice, which produces wider beneficial plasticity.[14,15]

While TNP protocols are a promising new therapeutic approach, their design and evaluation are formidable tasks. The many kinds of CNS plasticity, the many sites where they occur, the many new biological and technical methods, and the probability that the best therapies will combine several methods, generate an overwhelming number of appealing protocols. Testing even one of them is lengthy, demanding and expensive, especially in humans. Computational models offer a solution. They can provide rapid and efficient screening of many potential protocols; only the most promising ones would then be tested in animals and/or humans.[16,17]

This study develops a neural network model first proposed in Reinkensmeyer *et al.*[18] of motor corticospinal (CS) plasticity before and after a simulated stroke. We used it to predict the therapeutic efficacy of a TNP training protocol as a function of the brain region targeted and the TNP dosage. The results provide insight into the mechanisms of cortical reorganization after stroke and into the design of maximally beneficial TNP protocols. They indicate how judicious use of computational models might shape the development of effective new rehabilitation therapies.

Materials and methods

We simulated the impact of different treatment protocols on recovery of contralateral finger extension after a stroke that damaged motor cortex in one hemisphere. The foundations

of this mathematical model are based on the structure and learning model first presented in Reinkensmeyer *et al.*[18] That model investigated use-dependent recovery of movement strength following a stroke using a network of CS neurons connected to downstream motor neuronal pools. The network learned using a biologically plausible reinforcement learning rule. To investigate the impact of different TNP protocols following a stroke, we now extend the structure of the model to represent CS neurons in both hemispheres; each neuron now has its own connection strength to the motor neuronal pool *and* its own intrinsic firing rate variability. We then simulate three scenarios: (1) the undamaged network underwent standard finger extension training trials; (2) the trained network was damaged by a stroke affecting contralateral motor cortex (i.e. contralateral to the finger) and then underwent standard finger extension training trials; and (3) the trained network was damaged by a stroke affecting contralateral motor cortex and then underwent standard finger extension training trials interspersed with training trials in which trial outcome was determined by the behaviour of a specific population of CS neurons (i.e. TNP trials).

Architecture

The model incorporates a network of n CS neurons that fire with activation levels x_i (assumed to vary between 0 and 1 and to correspond, proportionally, to firing rates). Each CS neuron is connected to a motoneuronal (MN) pool via a scalar connection weight w_i. The MN pool sums the product of the neuron activation x_i and connection weight w_i using a saturation nonlinearity, g_i

$$S_e = \sum g_i(x_i)w_i, \tag{1}$$

where function g_i sets the saturation limit of neuron i. In this presentation, the model has a constant saturation limit of +1 for all neurons

$$g_i(x_i) = \begin{cases} x_i, & x_i \le 1 \\ 1, & x_i > 1 \end{cases} \tag{2}$$

The MN pool activation level S_e is proportional to a unitless finger extension torque T_e. Thus, the finger extension torque generated is proportional to, and determined by, the weighted summed output of the CS network activation pattern (Fig. 1). Varying the saturation limits across neurons did not significantly affect network dynamics. Thus, this presentation uses the constant saturation limit of +1 for all neurons.

Network learning

The goal of the network is to learn the CS neuronal activation pattern that produces the maximum possible finger extension torque, i.e. the best performance. Since the network consists only of neurons that excite extensor motor neuronal pools, the optimal activation pattern is achieved when the activation of every neuron is increased to the neuron's saturation limit of +1.

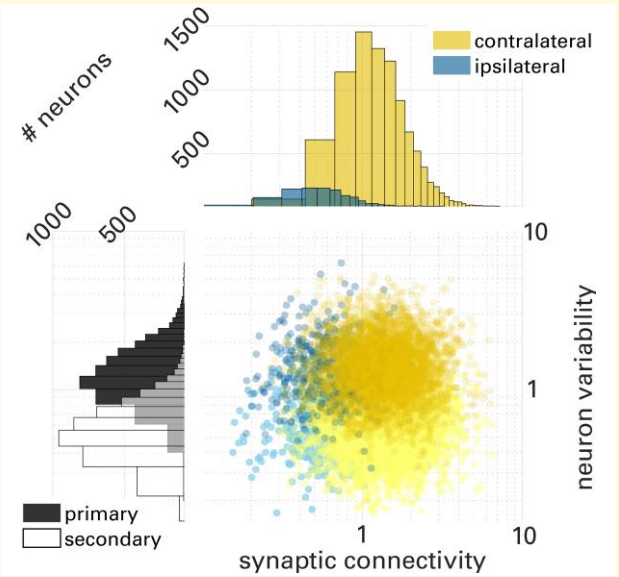

Figure. 1 Parameter distributions for a network of 10 000 corticospinal (CS) neurons. Synaptic connectivity adheres to a bimodal distribution resulting from use of two lognormal probability density functions, one for the contralateral cortex and one for the ipsilateral cortex. The mean of the bimodal distribution was chosen to be one. Most (90%) neurons reside in the cortex contralateral to the finger to be extended and have stronger connectivity (contralateral = yellow, high connectivity). The remaining neurons reside in the ipsilateral cortex and have weaker connectivity (ipsilateral = blue, low connectivity). Neuronal firing variability also adheres to a bimodal distribution arising from two lognormal probability density functions. Neurons in primary motor areas are more task-related and exhibit more trial-to-trial variability during movement attempts (primary = dark, high variability). Neurons in secondary motor areas are less task-related and exhibit less trial-to-trial variability during movement attempts (secondary = light, low variability). The resulting network has four broad types of neurons: 1, high-connectivity/high-variability (dark yellow); 2, high-connectivity/low-variability (light yellow); 3, low-connectivity/high-variability (dark blue); 4, low-connectivity/low-variability (light blue).

To learn this pattern, the network employs an iterative reinforcement learning protocol: after each trial (i.e. each movement attempt), the network adjusts the activation patterns based on a scalar teaching signal, which is finger extension torque T_e. Using a single signal to optimize a large network presents a credit assignment problem: if finger extension torque T_e increases on a given trial, which neurons are responsible for the increase? Reinforcement learning can solve this credit assignment problem, albeit imperfectly.[18–22] As will become clear below, TNP trials are designed to mitigate the credit-assignment problem and to thereby improve the learning outcome.

We implement reinforcement learning with stochastic search. The algorithm uses a noise process to generate a new activation pattern for each trial (i.e. each attempted

movement). If the new activation pattern increases finger torque compared with the previous trial, the algorithm stores (i.e. switches to) the new activation pattern. The stochastic search algorithm is a simplified form of the random search with chemotaxis algorithm:[19]

Given an initial activation pattern X_0 that produces a torque T_0:

1. Activate CS neurons with pattern $X_i = X_0 + v_i$, where v_i is random noise, and measure the torque T_i produced by this pattern.
2. Store (i.e. switch to) the new pattern X_i if the torque T_i it produces is greater than T_0 (i.e. if $T_i > T_0$, then let $X_0 = X_i$ and $T_0 = T_i$).
3. Repeat

We also tested a gradient descent stochastic search method, another biologically plausible solution to the credit assignment problem.[21] It produced comparable results. Here, we present the results from the stochastic search algorithm to allow more direct comparison to Reinkensmeyer *et al.*[18] and to provide conceptual clarity.

Neuron parameters

The CS neurons in the model are described by their current activity level (x_i, proportional to firing rate), and their excitatory synaptic connectivity to the extensor MN pool (w_i). In the present work, we augment the model of Reinkensmeyer[18] to allow different levels of trial-to-trial firing rate variability for different CS neurons (σ_i). This feature is based on evidence[23] that neuronal variability differs across cortical areas. The network updates each neuron's activation level after a successful movement trial (i.e. a trial that produces more torque than the previous trial). Each neuron's connectivity and variability remain constant throughout the simulation. In reality, spared descending pathways are plastic after an injury. Although we do not change weighting (w_i), altering the firing rate is mathematically equivalent in this model; it is a heuristic for the dynamic nature of downstream synaptic plasticity.

Another change from Reinkensmeyer *et al.*[18] is that neuronal parameters, including initial activation level, fixed connectivity and fixed variability, were initialized by sampling from lognormal distributions. Functional and structural parameters in the brain, including synaptic connectivity and firing rates, are typically not normally distributed; they are skewed with a heavy tail. Thus, they are closely approximated by lognormal distributions[24] (see Supplementary Material, Selection of neuronal parameter distributions). To represent different cortical areas, we used different lognormal distributions (Fig. 2).

Synaptic connectivity

Monosynaptic and multisynaptic CS pathways are represented in the model by a single, fixed connectivity from each CS neuron to the finger extensor MN pool (i.e. weights labelled w_i in Fig. 1). We explicitly represent both cortical hemispheres in the model, since the hemisphere ipsilateral to the moving finger is known to be able to activate the requisite MN pools through uncrossed pathways, and these pathways are thought to play a significant functional role after stroke.[25] We designed the distributions of the neuronal connectivities to reflect known physiology: neurons from the hemisphere contralateral to the motor task are, on average, more strongly connected to the MN pools than ipsilateral neurons (Fig. 2).[26] Furthermore, these contralateral neurons outnumber the ipsilateral neurons by a ratio of 9:1. This is consistent with the physiological situation wherein about 90% of axons cross over from the lateral CS tract and about 10% of fibres travel within the uncrossed anterior CS tract.[27]

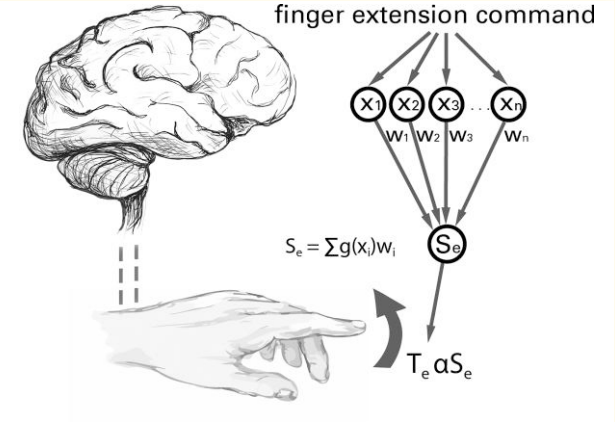

Figure 2 Network architecture. A two-layer feedforward neural network incorporates *n* corticospinal (CS) neurons with activation levels x_i. These activation levels are generated when the neurons are given a command to maximize finger extension torque T_e. A motoneuronal pool S_e sums the weighted activation pattern. A nonlinear function g_i implements the physiological observation that the contribution of any single CS neuron to the excitation of the motoneuronal pool saturates at some activation level. The network optimizes the activation pattern × through reinforcement learning in simulated consecutive movement practice trials in which extension torque T_e is the teaching signal.

Variability

Motor variability is necessary for motor learning;[28,29] higher levels of motor variability accompany motor skill acquisition. Furthermore, such variability is present from the neuronal level to the behavioural level.[23] Thus, the algorithm incorporates trial-to-trial corticomotor variability to drive motor learning in our training scenarios.

The algorithm varies neuronal firing rate by sampling an activation noise level from a normal distribution that is specific to each neuron. A higher-variability neuron has a wider normal distribution; it averages a larger stochastic perturbation to its activation level on each trial compared with a lower-variability neuron. The most task-relevant brain areas exhibit more variability. Thus, the model gives neurons in primary motor areas of both hemispheres higher variability for the finger extension task than neurons in secondary

motor areas. At the same time, high- and low-variability neurons are unlikely to be wholly separated spatially.[30,31] For this reason, the model overlaps high- and low-variability neurons (Fig. 2). A neuron's firing rate is related to its mean firing rate by a power function, the parameters of which vary across cortical areas and behavioural conditions.[32,33] Our model simplifies this relationship to clarify the effects of neuronal variability on the network dynamics: a neuron's trial-to-trial variability differs across cortical areas but remains constant over time and does not depend on its firing rate.

Simulations

We use this model to simulate three scenarios for learning finger extension: (1) learning by the uninjured network; (2) learning by the network following a unilateral stroke; (3) learning by the network following a unilateral stroke with TNP training trials interdigitated. The initial, uninjured network consisted of 10 000 CS neurons. To simulate the stroke, Scenarios 2 and 3 fix the activation and connectivity of CS neurons lost to the stroke to zero. To simulate the disruptive effects of the stroke on the properties of surviving CS neurons,[34] their initial post-stroke activation patterns for Scenarios 2 and 3 are randomized.

Since network learning is driven by movement attempts (i.e. trials), the number (i.e. dosage) of movement trials affects the results. Lang *et al.*[35] found that participants with stroke completed an average of 32 functionally oriented movements/day during upper extremity rehabilitation sessions. However, recent rehabilitation interventions have successfully administered 150–250 trials/session, for up to 23 sessions, for a total of 3450–5750 trials.[7] In the simulations presented here, the network trained in each given scenario for a total of 20 000 total movement attempts. In the TNP scenario, a subset of these attempts was given TNP feedback. For example, in the case where 20% of attempts were dedicated to TNP therapy, this simulates 4000 trials in the clinic.

Scenario I: Learning with an undamaged network

To provide a baseline, we simulated learning in the uninjured network. For each trial, finger extension torque was determined using all 10 000 neurons. The teaching signal was the difference from the previous trial in finger extension torque. If the torque was greater than that of the previous trial, the network switched to the new activation pattern.

Scenario II: Learning after stroke

To simulate neuronal death after a stroke, we disconnected (de-weighted) a subpopulation of 3333 CS neurons by permanently setting their connectivity and activation levels to zero. This scenario focuses on the most severe stroke: a stroke that damages contralateral primary motor cortex, where high-connectivity, high-variability CS neurons are concentrated. The teaching signal from each trial is the finger extension torque produced by the remaining intact CS neurons of both hemispheres.

Scenario III: Learning after stroke with TNP trials interdigitated

Finally, we simulated learning the finger extension task after stroke with a portion of the training trials now dedicated to a TNP protocol. In TNP trials, only the targeted neurons determined the teaching signal. That is, in a TNP trial, separate training signals were generated for the total network (finger extension torque, I_A) and the targeted network subpopulation (targeted intervention, I_B). The training algorithm used I_B to determine success or failure (i.e. to determine whether the network switched to the new activation pattern). If I_B indicated success, the parameters of all the surviving CS neurons (targeted and untargeted) were updated. The total number of trials did not change; Standard trials and TNP trials were interspersed in different ratios to determine the optimal dosage. Thus, we systematically evaluated the dependence of torque recovery on two training parameters: (i) which subpopulation of CS neurons was targeted; and (ii the proportion of the total trials that were TNP trials (i.e. dosage).

Statistical analysis

Neuron parameters adhere to a bimodal distribution resulting from two lognormal probability density functions. These parameters were set according to known physiological parameters; no statistical comparisons were made. We ran each simulation ten times and present the mean torque (solid lines) and standard deviation of torque (shaded area) as functions of movement trial. We ran each simulation 20 times to determine how targeting each of eight cortical regions (160 simulations total) affected recovery after simulated stroke. We present the mean recovery for each area and the standard deviation of that value across 20 simulations using bars and whiskers. We assessed the effect of varying TNP dosage and present recovery as a function of dosage as a mean (solid line) and standard deviation (shaded area) across 20 simulations. To determine the behaviour of different neuron groups, we used a Monte Carlo method where we ran the same single trial of the model 10 000 times—once at the beginning of network training, and once after the network had been trained for 20 000 trials. We present the mean change in torque due to neuron activity changes, grouped by their parameters. We also present the probability of a neuron increasing its firing rate, the mean change in firing rate and mean change in torque across all successful trials and across all trials.

Results

We trained the network in three Scenarios: (1) an undamaged network; (2) after stroke; and (3) after stroke with TNP trials interdigitated. We ran each scenario 10 times, with 20 000 trials each time. In Scenario 1, the undamaged network achieved 84.5% of the maximum possible finger extension torque (i.e. the maximum torque possible for all CS neurons) following a stereotypical learning curve (see Supplementary Material, Characterizing the learning curve

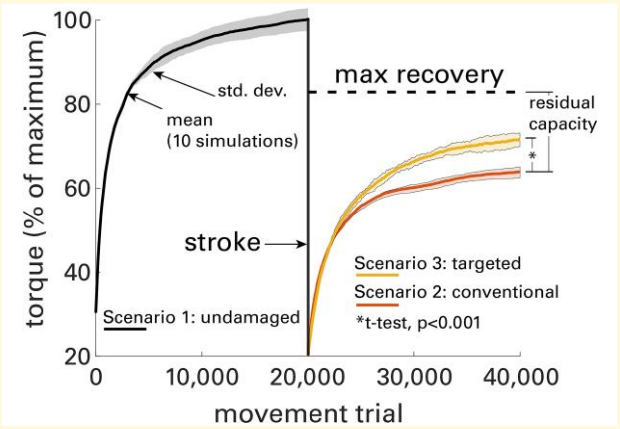

Figure 3 Torque production by the three Training Scenarios. Torque (in % of the maximum torque achieved by the uninjured network after 20 000 trials) as a function of the number of movement trials. After a stroke, max recovery is the theoretical maximum torque possible for the surviving CS neurons. Solid lines represent the mean result and shaded areas represent the standard deviation of ten simulations. Conventional Training (Scenario 2) leaves substantial residual capacity for recovery. Targeted plasticity (TNP) (Scenario 3) recovers much of this capacity; thus, it produces greater torque recovery (student's *t*-test, $t = 9.72$, $P < 0.001$).

with residual capacity). For simplicity, we scale all results to the average maximum torque produced in Scenario 1 (Fig. 3).

In Scenario 2, we simulated neuronal death due to stroke by removing a subpopulation of CS neurons from the network (by fixing their connectivities and activation levels at zero). We randomized the initial firing rates of the surviving CS neurons. The teaching signal was the finger extension torque produced by all the surviving CS neurons. Finger torque increased exponentially and then approached a recovery plateau in which it generated 64.9% (+/− 2.9% SD) of the maximum torque possible before the stroke (Fig. 3, Scenario 2). While this is a substantial recovery, significant capacity to generate torque remained unused (Fig. 3, residual capacity for recovery).

Scenario 3 began in the same way as Scenario 2, i.e. with a stroke and 20 000 trials of subsequent training. In our initial evaluation of Scenario 3, every fifth trial was a TNP trial; in which the teaching signal was provided by the surviving CS neurons in the secondary motor areas of both hemispheres. After this training period, the network reached 72.8% of the maximum torque possible before the stroke, an increase over Scenario 2 (conventional rehab). Scenario 3 enabled the network to use 43.6% of the latent capacity for recovery that Scenario 2 did not capture.

Impact of training on different neuronal types

Scenario 1 (the uninjured network) preferentially optimized activation of high-variability (M1, primary motor cortex), high-connectivity (contralateral) neurons. Thus, it left some residual capacity unachieved.

Scenario 2 (the injured network with standard training) also favoured optimization of high-variability, high-connectivity neurons (Fig. 4C). However, because many such neurons were gone, the training then optimized high-variability neurons with low connectivity; it shifted activity towards the weakly connected, but undamaged ipsilateral hemisphere.

Scenario 3 was the same as Scenario 2 except that every fifth trial based its outcome only on neurons in secondary motor areas of either hemisphere (i.e. generally low-variability neurons not optimized by Scenario 2). These targeted areas had the potential to significantly improve overall network performance. However, without TNP, the network struggled to access them. In standard trials, the stronger impact on trial outcome (i.e. success or failure) of the random changes in high-variability neurons masked the weaker impact of low-variability neurons. The TNP trials of Scenario 3 removed this masking. They promoted activation of low-variability/high-connectivity neurons (Fig. 4D).

Impact of training on the topography of neuronal activation

To visualize the predicted effect of the three scenarios on the topography of the task-related neuronal activation, we mapped neuronal parameters to brain areas. Specifically, we mapped high- and low-connectivity neurons contralateral and ipsilateral to the movement. We mapped high and low-variability neurons to primary motor cortex (M1) and dorsal premotor cortex (dPM; a secondary motor area), respectively (see Table 1). In Fig. 5, we use this mapping to show network reorganization across hemispheres during learning before a stroke (Scenario 1), and after a stroke without (Scenario 2) or with (Scenario 3) TNP training.

Standard learning by the intact network (Scenario 1) optimized contralateral primary motor (i.e. high-variability, high-connectivity) neurons. This resulted in contralateral activity (Fig. 5). Standard learning after stroke (Scenario 2) shifted activation to the undamaged ipsilateral hemisphere (i.e. towards high-variability neurons with relatively low connectivity). This produced bilateral activation for the unilateral task. Standard learning after stroke plus TNP (Scenario 3) also optimized neurons in secondary motor areas of both hemispheres (i.e. low-variability neurons). This produced more normal lateralized activation and improved torque recovery.

Optimizing target population and dose

One of the model's advantages is its ability to predict which populations to target to maximize therapeutic effect. To this end, we targeted different cortical areas and gave TNP on every fifth trial. We ran each simulation twenty times; each one simulated a stroke that affected a random subset of 75% of the CS neurons in contralateral primary motor

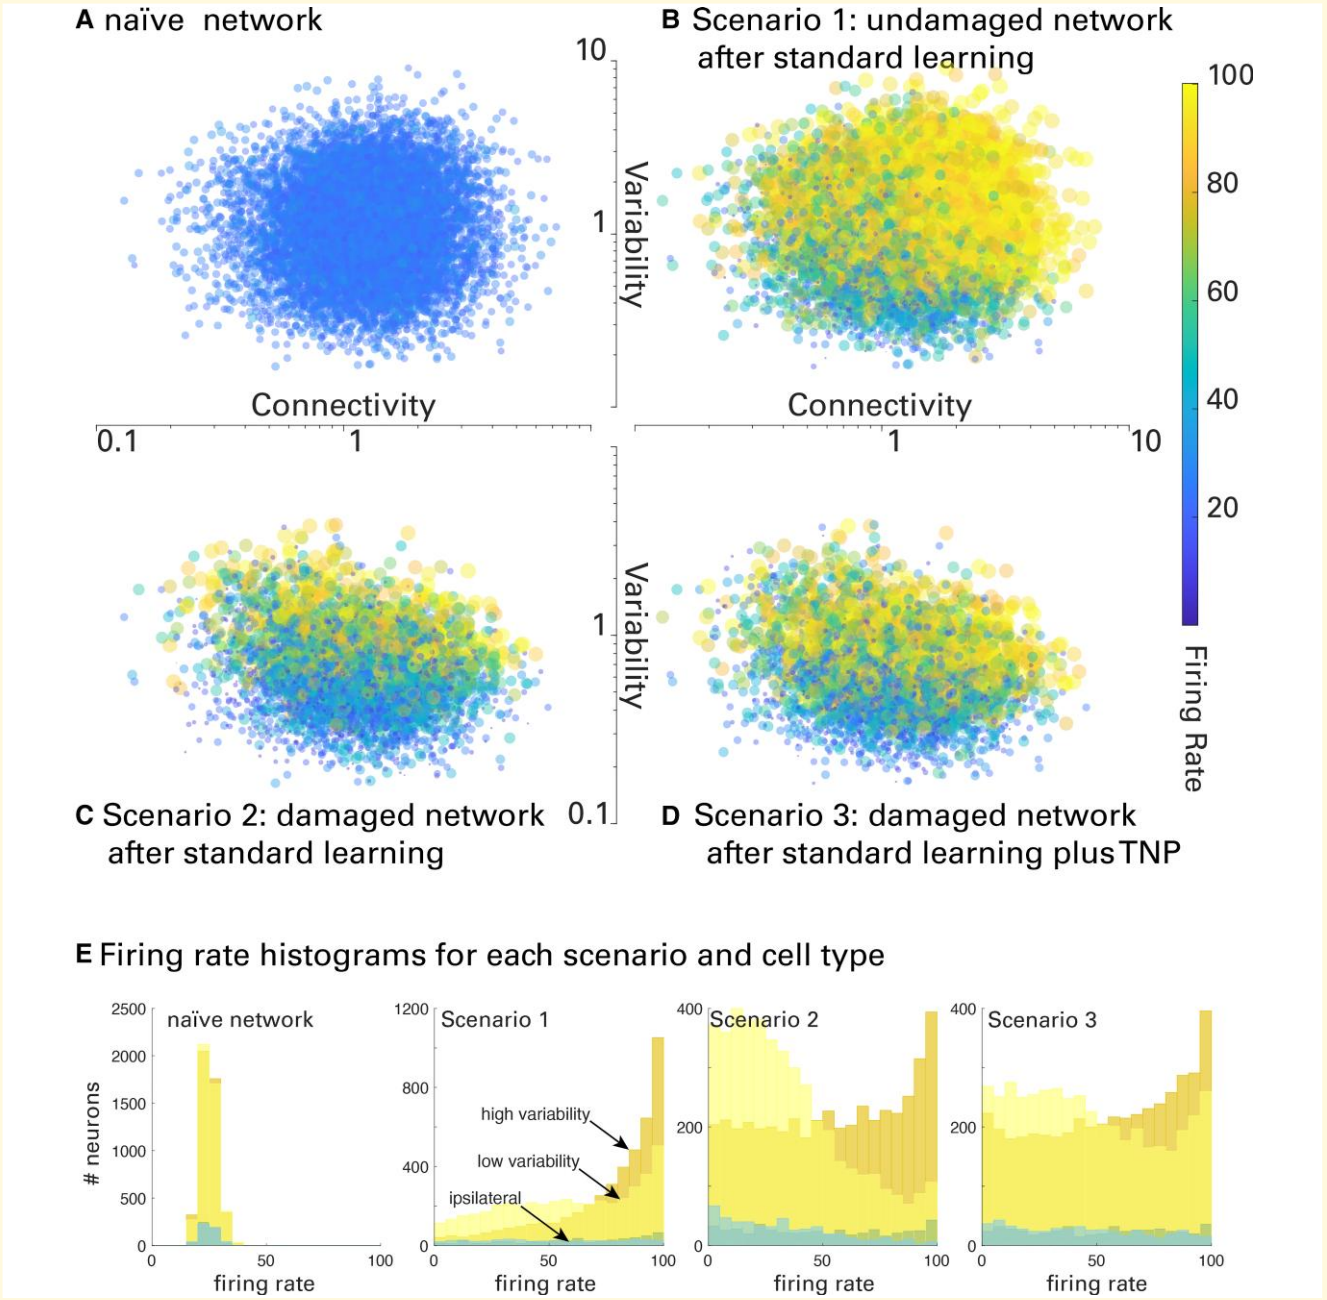

Figure 4 Neuronal activation patterns for the naive network and for the network at the end of each of the three scenarios (shown for a 10 000-neuron network). (**A**) Normal network before learning. Network activation reflects a pseudorandom sample of activation levels. (**B**) Normal network after standard learning and before stroke, with more optimized neurons in yellow/lighter shade (high firing rate). High-connectivity/high-variability neurons are optimized. (**C**) Injured network (i.e. 8125 remaining neurons) after standard learning. Optimization is focused on the remaining high-variability neurons; low-variability neurons are generally not optimized. (**D**) Injured network after standard learning plus targeted neural plasticity (TNP). Optimization of low-variability neurons is increased (i.e. d is substantially more yellow/lighter shade than c). The result is that torque recovery improves (Fig. 3). (**E**) Firing rate histograms for each scenario and neuron type. Low-variability neurons are optimized by TNP training in Scenario 3. versus Scenario 2.

cortex (i.e. mostly high-connectivity, high-variability neurons). We then determined for each area the improvement in residual torque capacity recovered over that provided by standard training alone (Scenario 2) (Fig. 6A). Delivering TNP therapy to secondary motor areas (e.g. dPM) gave the best results, restoring 30.5% of the residual capacity left by Scenario 2. Targeting dPM only in the damaged hemisphere had similar effects. Targeting dPM in the undamaged hemisphere was ineffective. Targeting primary motor cortex (M1) either bilaterally or only in the damagehemisphere

Table 1 Cortical mapping of neuronal parameters (M1: primary motor cortex; dPM: dorsal premotor cortex)

Variability	Connectivity	
	Low	High
High	Ipsilateral-M1	Contralateral-M1
Low	Ipsilateral-dPM	Contralateral dPM

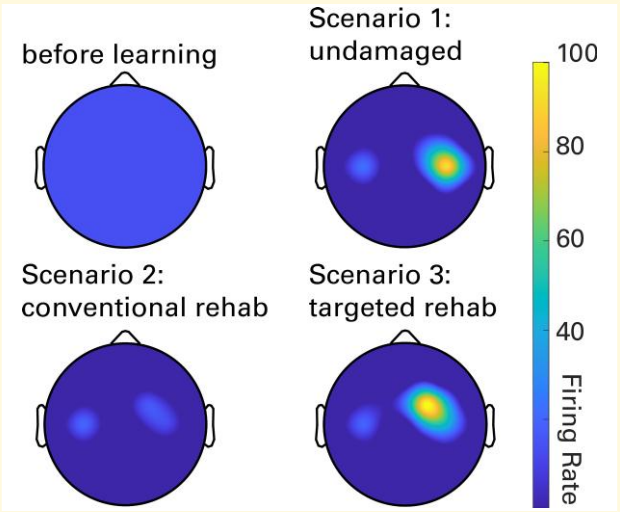

Figure 5 Topography of task-related neuronal activation before training and at the end of training for each of the three training Scenarios. 1, Standard training of the uninjured network; 2, standard training of the injured network; and 3, standard training of the injured network plus TNP. Scenario 1 primarily optimizes neurons in contralateral primary motor cortex (high-connectivity/high-variability neurons). Scenario 2 also optimizes neurons in ipsilateral (undamaged) cortex (high-variability/low-connectivity neurons). This produces abnormal bilateral activation and diffuse activation in the remaining contralateral (damaged) primary motor cortex. Scenario 3 also optimizes neurons in secondary motor areas (e.g. dPM); and it restores more normal laterality.

caused maladaptive plasticity that reduced torque production compared with Scenario 2. Targeting primary motor cortex only in the undamaged hemisphere had no significant effect.

We then focused on the most beneficial targeting (i.e. dPM bilaterally) and assessed the effect of varying TNP dosage from 1/2000 trials to every trial. We ran each simulation twenty times; each one simulated a stroke that affected a random subset of 75% of the CS neurons in contralateral primary motor cortex (i.e. mostly high-connectivity, high-variability neurons). We then determined for each TNP dosage the residual capacity recovered compared with conventional rehabilitation (Fig. 6B). Doses <1% were ineffective. As dose increased from zero, training efficacy increased; it reached a maximum when 20% of the trials were TNP trials. Efficacy declined at doses >50%. The results were similar when other cortical areas were targeted. To our knowledge, these results provide the first substantive insight into the most effective dosage of TNP as a fraction of total training. The consistency of the result across different targeted areas suggests that the result may apply across different lesion types and training protocols.

Mechanisms: neuron-specific optimization rates and 'blocking'

Why did activity lateralize in the undamaged brain, but become bilateral after the simulated stroke? And why did TNP help remediate this situation, improving torque recovery? Two mechanisms account for these results. First, high-variability neurons optimize faster than low-variability neurons. Second, once optimized, high-variability neurons prevent (block) optimization of low-variability neurons.

To illustrate these mechanisms, we used a Monte Carlo method wherein we ran the same, single trial of the model 10 000 times—once at the beginning of network training, and once after the network had been trained for 20 000 trials, when the network had learned to generate more torque (Fig. 7A) by increasing activation across neuron groups (Fig. 7B). Using the results from the Monte Carlo simulations, we estimated the probability that each neuron group would, through their summed activity that resulted from the random perturbation on that trial, contribute towards a positive change in torque. Specifically, we calculated the mean change in activation $\delta\ X = \frac{1}{n}\sum_{i=1}^{n} \delta x_i$ and the mean change in torque $\delta\ T = \frac{1}{n}\sum_{i=1}^{n} g(\delta x_i w_i)$ across all neurons $x_i = 1,2, \ldots, n$ in each group. We report these values for successful trials that advance the optimization (Fig. 7C–E) and for all trials (Fig. 7F–H).

Highly variable neurons optimize first. Before training, no neurons are optimized (Fig. 7B). Thus, on learning trial 1, all neurons increase or decrease their activity with equal (i.e. 0.5) probability on each trial (Fig. 7F). Recall that we ran this first learning trial many times. Across all these first learning trials that are successful (i.e. produce an increase in torque), high-variability (i.e. fast) neurons change activation (Fig. 7D), and thus their contributions to total torque (Fig. 7), by relatively large amounts. The contribution to increase in torque is particularly high for the fast, high-connectivity (i.e. strong) neurons, which will cause them to optimize quickly. These neurons will saturate as training proceeds (they cannot exceed the maximum firing rate defined by the *g* function). At trial 20 000, they do not contribute any more to increasing the total torque (Fig. 7E, after training), or it is extremely unlikely that they will do so (Fig. 7C, after training). The network then favours at this later time stage (i.e. following the optimization of fast/strong neurons) the optimization of high- variability/low-connectivity (fast/weak) neurons (Fig. 7D, after training).

The network learns at increasingly slow rates, failing to optimize some neurons. Saturated neurons cannot contribute to further increases in torque production. Thus, for saturated neurons, the probability that δx is positive approaches zero

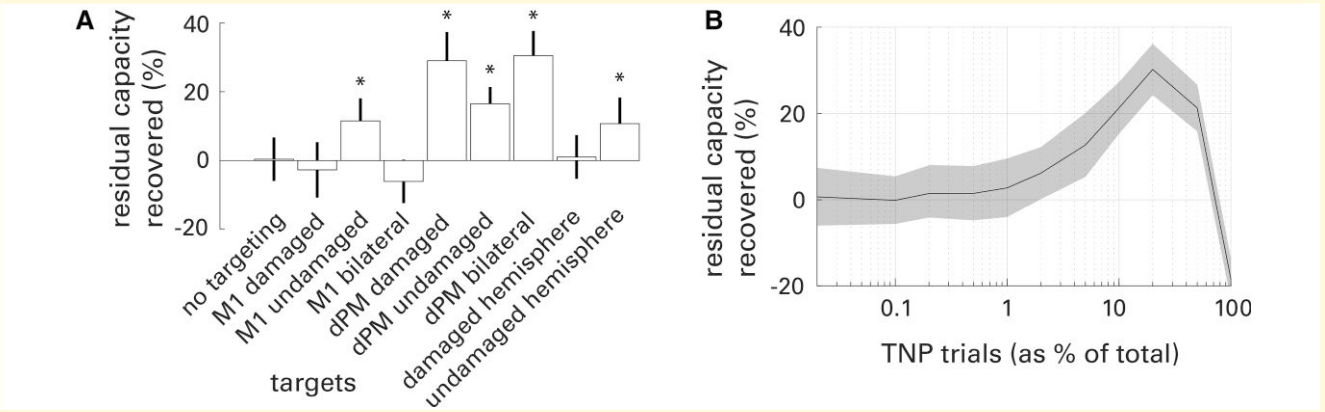

Figure 6 The benefits of targeted neuroplasticity (TNP): the impact of the area targeted and the TNP dosage. (**A**) Additional residual torque recovered by standard learning plus TNP (given on 20% of trials) over that recovered with standard learning alone, for each of eight different cortical regions (M1: primary motor cortex; dPM: dorsal premotor cortex). Results are given as a percentage of the residual capacity for recovery left uncaptured by Scenario 2 (i.e. no TNP trials). Positive values indicate a better outcome with TNP. Error bars are standard deviation of 20 simulations. (**B**) Residual torque recovery as a function of the percentage of trials that were TNP trials (results shown for targeting dPM). The solid line represents the mean and shaded area indicates the standard deviation of 20 simulations. Recovery improved with increased dosage of targeted feedback. Recovery reached a maximum when 20% of trials were given targeted feedback and declined thereafter.

(Fig. 7F). As observed above, more variable and strongly connected (fast/strong) neurons optimize more quickly, and thus, on average, saturate first. For the network to switch to a new activation pattern, yet-to-be optimized neurons must not only produce their own net positive effect on torque but must also overcome any negative effect of the more-variable neurons that, once saturated, tend to decrease torque on any given trial (Fig. 7E). This blocking phenomenon, created by the nonlinear saturation function, g, causes the latent residual capacity shown in Fig. 3.

Fast but weak neurons optimize before slow but strong neurons. Because the training signal is total torque, the network preferentially optimizes neurons by their torque change, not their activation change. Late in learning, fast/weak and slow/strong neurons are similar in mean torque change per successful trial (Fig. 7E, after training), but fast/weak neurons increase activity more on successful trials, thus optimizing more quickly (Fig. 7d, after training). As these neurons saturate, they also block optimization of the slow/strong neurons, and thus they limit total torque recovery.

Discussion

After a stroke, EEG and functional magnetic resonance imaging (fMRI) studies reveal substantial reorganization of movement-related cortical activity.[1,36–40] For example, unilateral movements that are contralateral to the affected hemisphere can elicit bilateral activity.[25,39,41,42] This loss of hemispheric laterality correlates with decreased motor function. It may reflect a suboptimal compensatory strategy that limits motor recovery.[36,38] The present study uses a computational model to explain these shifts in laterality as arising from the motor system's stochastic search for neurons that can help drive MN pools following injury. The model further demonstrates the ability of a new therapeutic strategy, TNP, to modify abnormal movement-related cortical activation and to thereby improve motor recovery.[43]

In traditional rehabilitation, patients simply practice the skills that have been impaired by stroke (e.g. locomotion, reach and grasp, speech). In the model, practicing finger extension causes the neural search process to settle on a suboptimal pattern of activation. In contrast, the TNP protocol simulated here uses operant conditioning to modify the task-related activation of specific subpopulations of CS neurons and to thereby enhance the functional recovery of the entire population. In both animals and humans, TNP protocols can target beneficial plasticity to a specific CNS site (e.g. a spinal reflex pathway or a cortical area) by operantly conditioning EMG or EEG features that reflect activity in that site.[11,44] This plasticity improves function and enables wider plasticity that further improves function.[14,45,46] For example, a TNP protocol that operantly conditions a spinal reflex pathway can improve walking in rats[11] or people with incomplete spinal cord injuries.[14] A TNP protocol that operantly conditions specific EEG features can enhance functional recovery after stroke.[7,8,10,46]

To date, only a handful of studies have attempted to model the mechanisms underlying sensorimotor rehabilitation;[16,17] none have modelled a TNP protocol. In this paper, we do this by building on an approach that employed a simplified CS neural network with inherent stochastic noise to simulate finger movement recovery after stroke.[18] The model used a biologically plausible reinforcement learning (operant conditioning) algorithm to optimize CS activation patterns over repeated motor practice. This network reproduced major clinically observed features of motor recovery after stroke, including exponential recovery

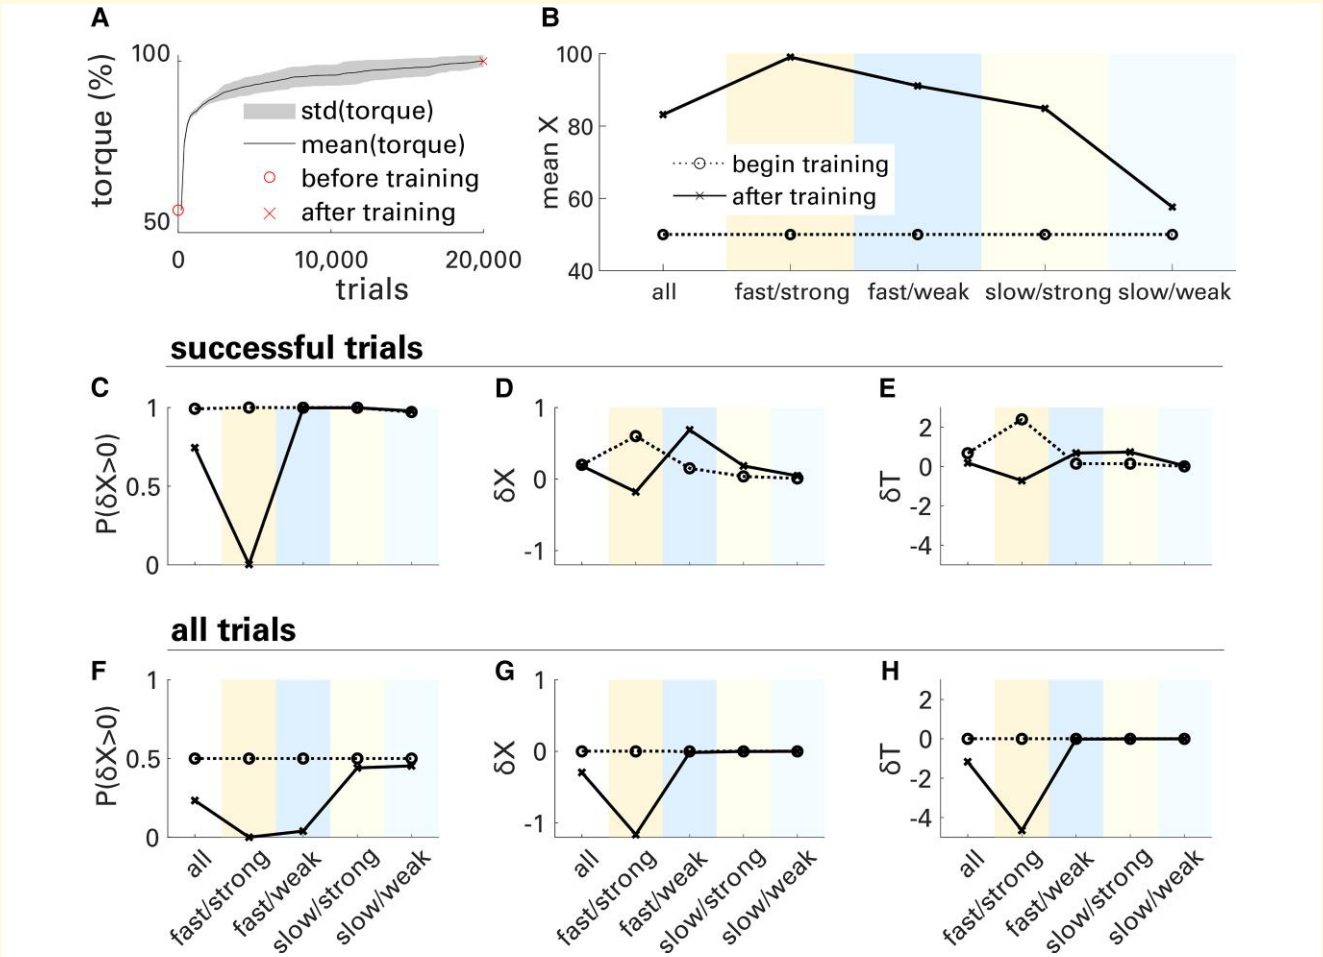

Figure 7 Network effect statistics. **(A)** Total torque generation (as percentage of maximum) over 20 000 training trials. **(B–H)** Neuronal population measures for trials at the beginning of training (dotted lines) and after 20 000 trials (solid lines), calculated by a Monte Carlo simulation. **(B)** Mean activation levels for each population. **(C)** Probability that each population's mean change in activation was positive on a successful trial (i.e. a trial that increased the total torque). **(D)** Mean change in each population's activation level for successful trials. This is effectively the rate of optimization. **(E)** Each population's mean contribution to the torque change on a successful trial. This indicates which population the model is currently optimizing. **(F)** Probability that a population's mean change in activation was positive on any given trial. **(G)** Mean change in a population's activation level for all trials. **(H)** A population's mean contribution to the torque change on all trials.

and latent residual capacity. A subsequent paper extended this model to simulate multiple limbs and explore the effects of strength and coordination training after neurological injury.[47] Here, we extend the model using biologically plausible neuronal population parameters and employ it to predict the results of combining a standard rehabilitation protocol with a TNP protocol.

The computational model presented here suggests that, after a stroke damages cortex, traditional rehabilitation methods may fail to optimize cortical reorganization, limiting motor recovery. Furthermore, this simulation supports the hypothesis that appropriate TNP interventions can improve cortical reorganization and enhance motor recovery. We found that Scenario 2, which simulated traditional rehabilitation, left the network with residual capacity for motor recovery. Scenario 3, which incorporated TNP training, was able to access much of this residual capacity.

While mounting evidence suggests TNP interventions affect cortical reorganization, it is not clear how to maximize their beneficial effects on functional recovery.[4] Our reductionist model suggests that targeting neurons that are not easily accessed by regular motor practice is particularly beneficial. With conventional training, these neurons were blocked from full optimization. The simulated stroke primarily affected high variability neurons but never completely wiped them out. This presents an interesting comparison to the real-world physiological situation wherein peri-infarct neurons are often hyperexcitable. In fact, the results of this model suggest that such peri-infarct neurons would contribute to immediate learning but would eventually block less excitable populations in other motor areas that might otherwise contribute to motor recovery. TNP circumvents this blocking and thereby increases total torque. Furthermore, the model results predict the optimal dosage

and target of TNP training. The following sections compare the model predictions regarding network reorganization with the results of imaging studies, summarize TNP principles derived from the model, and consider the study's limitations and the most promising and important directions for further research.

Cortical organization before and after cortical injury

Extensive EEG and fMRI data indicate that activation of contralateral primary motor areas normally precedes and accompanies motor actions.[48–50] In accord with these experimental data, our model, applied prior to stroke, optimized mainly the high-variability, high-connectivity (i.e. fast/strong) neurons concentrated in contralateral primary motor cortex. Thus, learning a right-hand task optimized activation of neurons in the primary motor cortex of the left hemisphere.

After the simulated stroke destroyed primary motor cortical neurons, the model exhibited a profound but suboptimal reorganization of network recruitment. This phenomenon is prominent in clinical data.[40,48,51] Injuries to sensorimotor regions often result in extension of motor representation into perilesional regions.[38,52–54] Another effect is increased activation in the uninjured hemisphere;[55,56] this is especially prominent for unilateral movements of the affected limb.[25,39,48] In accord with these clinical data, after the stroke, the model produced diffuse perilesional activation in the lesioned hemisphere and a prominent increase in activation of the uninjured hemisphere.

The model provides insight into the possible mechanisms of this suboptimal reorganization. For a network that uses stochastic search to optimize neuronal activation, the trial-to-trial variability of each neuron, and, secondarily, its connectivity to the MN pool, determines the rate at which the network is able to recruit that neuron (i.e. to increase its contribution to the total finger extension torque). Once the most quickly optimized neurons (e.g. those in contralateral primary motor areas) are saturated, they can only remain saturated or incrementally decrease their activity as they experience trial-to-trial variability. After stroke, these incremental decreases in their activity block optimization of other neurons. While the shift towards bilateral activation does increase torque, it is suboptimal because it blocks recruitment of high-connectivity (but more slowly recruited) neurons in other areas that could better enhance total torque. This leaves a substantial residual (i.e. unused) capacity for further recovery of torque.

Simulating TNP training partially re-lateralized hemispheric activity and improved torque recovery. Such re-lateralized activity has also been found to accompany better functional outcome in the stroke recovery literature.[57] The model provides insight into how TNP can improve motor recovery. As discussed above, neurons in secondary motor areas with the potential to contribute to torque production may be blocked from optimization by neurons in primary motor cortical areas with higher trial-to-trial variability (i.e. neurons that are rapidly recruited). The model predicts that TNP therapies may remove the blocking effect by modifying the roles of specific cortical areas in determining the teaching signal (i.e. in determining whether a trial is successful and thus updates the activation levels of all the CS neurons). Simulated TNP therapy removed the block by ignoring the high-connectivity/high variability neurons concentrated in primary motor areas and deriving the teaching signal solely from the low-variability neurons concentrated in secondary motor areas.

In summary, the model replicated the patterns of network organization that are found in people before and after a unilateral cortical injury (e.g. a stroke). Furthermore, it identified a mechanism that can account for both the normal (pre-stroke) and abnormal (post-stroke) patterns: the dynamics of a stochastic search. With a stochastic search, the final topography of activation reflects the topographies of neuronal variability and connectivity across the two hemispheres. Stroke changes those topographies, and thereby changes the results of the stochastic search. If the blocking and reorganization effects predicted by the model do occur *in vivo*, appropriate TNP training could improve functional recovery.

Computational principles of TNP

These results provide a rationale for using TNP training to enhance neuronal recruitment after injury and thereby improve recovery of motor function. The first principle is that interspersing TNP trials that enable recruitment of under-used neuronal populations with standard trials can access residual torque capacity inaccessible to standard training alone. This principle is consistent with animal and human evidence that appropriately designed TNP therapy can induce widespread adaptive plasticity leading to network reorganization and enhanced motor function.[3,4,7,8,10,11,13,14,46,58] In these TNP studies, real-time neuroimaging (i.e. EEG or fMRI) or measurement of a key physiological parameter (e.g. an H-reflex) provides the teaching signal. Stated in behavioural terminology, this teaching signal operantly conditions the person to modify key aspects of CNS activity (e.g. the activation levels of neurons in a specific cortical area).

A second principle based on the model is that, as TNP dose (i.e. TNP trials as % of all trials) increases, overall motor recovery increases up to a maximum and then declines. For the model presented here, maximum recovery occurs with a dose of ~20%. This optimum exists because standard trials alone leave a large latent residual capacity; and TNP trials alone leave most neurons untrained. A proper balance between them is essential. This suggests that studies such as,[7] which provided only TNP trials, might achieve still better results by interspersing TNP trials with standard, non-targeted trials.

The prevalence of stroke is highly correlated with age[59] and there is increasing evidence for the alteration of firing rates throughout the ageing brain.[60] It is not clear whether these changes affect aging, or share a common underlying

age-related mechanism. We found that altering the maximum firing rate of the network did not change the optimal TNP target or dosage result (Supplementary Fig. 1). In other words, the model suggests that age may not be a critical variable when for optimizing TNP target and dosage.

Mechanisms

Why did activity lateralize in the undamaged brain? The network favours optimization of high-variability/strong-connectivity (fast/strong) neurons, which are located mainly in M1 contralateral to the finger movement.

Why did activity become bilateral after the simulated stroke? After injury deprived the network of fast/strong neurons, it preferentially optimized fast/weak neurons, which are located mainly in M1 ipsilateral to the finger movement (Fig. 5B).

Why did targeted plasticity improve torque recovery and relateralize hemispheric activity? Standard training blocked optimization of slow neurons. By focusing on secondary motor areas (where slow neurons are mainly located), TNP training overcame the blocking and enabled optimization of the slow neurons, which are stronger contralaterally.

Model limitations and future directions

Our model simplifies cortical control of movement and the potential consequences of cortical stroke. It reduces the many pathways that interconnect the cortex and other brain areas with the spinal cord and the motoneurons to a group of weighted connections. After injury, white matter hyperintensities and glial cell perturbations may alter system activity, and thus affect function. Effects like these, i.e. outside of direct CS-MN connection, would not be captured by this model; they may be important for understanding the phenomenology of stroke injury and recovery with respect to plasticity and TNP therapy. It also ignores the numerous and poorly understood effects of a stroke (e.g. the impact of impairment of the somatosensory input that guides and maintains motor performance). Furthermore, human CS neurons reside in multiple cortical areas (i.e. primary motor, supplementary motor, premotor, somatosensory, cingulate and parietal). The model simplifies this complex reality into high and low-variability neurons that represent primary and secondary motor areas, respectively. Nevertheless, the model's results are consistent with clinical data. It displays, and helps to explain, phenomena that underlie both normal motor learning and rehabilitation after stroke.[16]

The model's value might increase with additional complexity. For example, a particularly valuable addition could be capacity for changes in neuronal connectivities to the motoneuron pools; in addition to training neural activation patterns, the network would also train the strengths of individual neuronal connections to the motoneurons. This addition would be particularly relevant given the ongoing discussion within the motor learning literature concerning the multiple learning mechanisms that drive short-term adaptation and long-term learning.[15,61,62]

Clinical TNP protocols that are designed based on the results of the simulation presented here could use recent advances in neural recording technology. For example, it may be possible to use high resolution EEG,[63] implanted recording arrays,[64] or ultrasound,[65] to generate teaching signals from a targeted population of sub-optimized neurons. The teaching signal could then be presented to the patient during rehabilitative movement practice to help optimize involvement of the targeted population. Alternately, it may be possible to inhibit populations of already optimized neurons using brain stimulation technologies such as TMS.

The modelling results might also shape the task that a clinical protocol presents to the patient. For example, the simulation presented here shows that incorporating trials that remove the influence of fast/strong CS neurons can access recovery capacity that is inaccessible to standard training alone. As with standard training in our model, traditional rehabilitation for finger extension may improve the performance of stronger, but not weaker, motor units. Stronger motor units are easily fatigued, while weaker units are not.[66] Thus, a protocol that requires prolonged maintenance of position against a steady-state torque could eliminate the contributions of stronger motor units and focus training on the weaker units; it might thereby access otherwise inaccessible capacity for torque recovery. Alternatively, a protocol that requires precise maintenance of a specific force level could discourage activation of strong motor units, which cannot provide such precise control, and encourage activation of weak motor units, which can provide precise control. A combination protocol might incorporate both methods for focusing on weak units.

Although this model applies well to CS control of a muscle, it might not apply as well to other rehabilitation problems. To be useful for a given problem, a model must simulate the motor behaviour to be restored and the CNS mechanisms that underlie the behaviour and its impairment by injury or disease. A variety of existing or conceivable TNP protocols might target plasticity related to specific temporal, kinematic, physiological or anatomical components of a variety of important sensorimotor behaviours. For example, a TNP protocol could target plasticity: in individual joint control during complex arm movements;[67] in responses to perturbations during movement;[68] in physiological measures of activity in key neuronal pathways;[14] or by using precisely paired stimuli to change a key CNS site.[69] As CNS imaging and stimulation technologies continue to improve, the variety and precision of TNP protocols should increase. Hopefully, parallel advances in computational modelling will enable models to facilitate and enhance these advances.

In summary, computational models can accelerate the creation and development of novel neurorehabilitation therapies. In contrast to clinical studies, which are generally demanding and time-consuming for both patients and investigators, modelling allows rapid assessment of many different designs and parameter selections. Properly applied, modelling could guide selection of the most promising

protocols for actual clinical study and could thereby enable efficient and effective realization of new therapies.

Acknowledgements

We thank Dr Peter Brunner for valuable comments on an earlier version of this manuscript.

The National Center for Adaptive Neurotechnologies (NCAN) is a Biomedical Technology Resource Center (BTRC) of the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health (NIH).

Funding

Dr Reinkensmeyer's research is supported by the National Institutes of Health/Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) grant R01HD062744. Dr Wolpaw's research is supported by National Institutes of Health (NIH)/National Institute of Biomedical Imaging and Bioengineering (NIBIB) grant P41 EB018783, National Institutes of Health (NIH)/National Institute of Neurological Disorders and Stroke (NINDS) grant 1R01NS110577, The Veterans Affairs (VA) Merit Award 5I01CX001812, and New York State Spinal Cord Injury Research Board (SCIRB) grant C32236GG.

Competing interests

The authors report no competing interests.

Supplementary material

Supplementary material is available at *Brain Communications* online.

Data availability

The data and code necessary to reproduce key results presented here are available at https://github.com/sumner15/simple_corticospinal_network

References

1. Wolpaw JR, Carp JS. Plasticity from muscle to brain. *Prog Neurobiol*. 2006;78(3):233-263.
2. Wolpaw JR. What can the spinal cord teach us about learning and memory? *Neuroscientist*. 2010;16(5):532-549.
3. Wolpaw JR, Kamesar A. Heksor: the central nervous system substrate of an adaptive behaviour. *J Physiol*. 2022;600(15):3423-3452.
4. Cramer SC, Sur M, Dobkin BH, *et al*. Harnessing neuroplasticity for clinical applications. *Brain*. 2011; 134(Pt 6):1591-1609.
5. Dimyan MA, Cohen LG. Neuroplasticity in the context of motor rehabilitation after stroke. *Nat Rev Neurol*. 2011;7(2):76.
6. Wolpaw J, Wolpaw EW, eds. *Brain-computer interfaces: Principles and practice*. OUP; 2012.
7. Buch E, Weber C, Cohen LG, *et al*. Think to move: A neuromagnetic brain-computer interface (BCI) system for chronic stroke. *Stroke*. 2008;39(3):910-917.
8. Ramos-Murguialday A, Broetz D, Rea M, *et al*. Brain-machine interface in chronic stroke rehabilitation: A controlled study. *Ann Neurol*. 2013;74(1):100-108.
9. Sitaram R, Veit R, Stevens B, *et al*. Acquired control of ventral premotor cortex activity by feedback training an exploratory real-time fMRI and TMS study. *Neurorehabil Neural Repair*. 2012;26(3):256-265.
10. Norman SL, McFarland DJ, Miner A, *et al*. Controlling premovement sensorimotor rhythm can improve finger extension after stroke. *J Neural Eng*. 2018;15(5):056026.
11. Chen Y, Chen XY, Jakeman LB, Chen L, Stokes BT, Wolpaw JR. Operant conditioning of H-reflex can correct a locomotor abnormality after spinal cord injury in rats. *J Neurosci*. 2006;26(48):12537-12543.
12. Thompson AK, Wolpaw JR. Restoring walking after spinal cord injury: Operant conditioning of spinal reflexes can help. *Neuroscientist*. 2015;21(2):203-215.
13. Thompson AK, Wolpaw JR. Targeted neuroplasticity for rehabilitation. *Prog Brain Res* 2015;218:157-172.
14. Thompson AK, Pomerantz FR, Wolpaw JR. Operant conditioning of a spinal reflex can improve locomotion after spinal cord injury in humans. *J Neurosci*. 2013;33(6):2365-2375.
15. Thompson AK, Wolpaw JR. H-reflex conditioning during locomotion in people with spinal cord injury. *J Physiol*. 2019;599:2453-2469.
16. Reinkensmeyer DJ, Burdet E, Casadio M, *et al*. Computational neurorehabilitation: Modeling plasticity and learning to predict recovery. *J NeuroEngineering Rehabil*. 2016;13(1):42.
17. Sedda G, Summa S, Sanguineti V. Chapter 9 - computational models of the recovery process in robot-assisted training. In: Colombo R, Sanguineti V, eds. *Rehabilitation robotics*. Academic Press; 2018:117-135.
18. Reinkensmeyer DJ, Guigon E, Maier MA. A computational model of use-dependent motor recovery following a stroke: Optimizing corticospinal activations via reinforcement learning can explain residual capacity and other strength recovery dynamics. *Neural Netw*. 2012;29:60-69.
19. Anderson RW, Omidvar O, Dayhoff J. Random-walk learning: A neurobiological correlate to trial-and-error. *Prog Neural Netw Ablex Norwood NJ*. 1997.
20. Mazzoni P, Andersen RA, Jordan MI. A more biologically plausible learning rule for neural networks. *Proc Natl Acad Sci USA*. 1991;88(10):4433-4437.
21. Werfel J, Xie X, Seung HS. Learning curves for stochastic gradient descent in linear feedforward networks. *Neural Comput*. 2005;17-(12):2699-2718.
22. Williams RJ. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach Learn*. 1992;8(3–4):229-256.
23. Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008;9(4):292-303.
24. Buzsáki G, Mizuseki K. The log-dynamic brain: How skewed distributions affect network operations. *Nat Rev Neurosci*. 2014;15(4):264-278.
25. Cramer SC, Nelles G, Benson RR, *et al*. A functional MRI study of subjects recovered from hemiparetic stroke. *Stroke*. 1997;28(12):2518-2527.
26. Nudo RJ, Jenkins W, Merzenich MM, Prejean T, Grenda R. Neurophysiological correlates of hand preference in primary motor cortex of adult squirrel monkeys. *J Neurosci*. 1992;12(8):2918-2947.
27. Martin J, Jessell T. *Neuroanatomy: Text and atlas*. Stamford Conneticut Appleton Lange; 1996.

28. Wu HG, Miyamoto YR, Castro LNG, Ölveczky BP, Smith MA. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci*. 2014;17(2):312-321.
29. Herzfeld DJ, Shadmehr R. Motor variability is not noise, but grist for the learning mill. *Nat Neurosci*. 2014;17(2):149.
30. de Ryuter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. *Science*. 1997;275(5307):1805-1808.
31. Warzecha AK, Egelhaaf M. Variability in spike trains during constant and dynamic stimulation. *Science*. 1999;283(5409):1927-1930.
32. Dean A. The variability of discharge of simple cells in the cat striate cortex. *Exp Brain Res*. 1981;44(4):437-440.
33. Lee D, Port NL, Kruse W, Georgopoulos AP. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci*. 1998;18(3):1161-1170.
34. Nudo RJ, Milliken GW. Reorganization of movement representations in primary motor cortex following focal ischemic infarcts in adult squirrel monkeys. *J Neurophysiol*. 1996;75(5):2144-2149.
35. Lang CE, MacDonald JR, Reisman DS, *et al.* Observation of amounts of movement practice provided during stroke rehabilitation. *Arch Phys Med Rehabil*. 2009;90(10):1692-1698.
36. Calautti C, Jones PS, Naccarato M, *et al.* The relationship between motor deficit and primary motor cortex hemispheric activation balance after stroke: Longitudinal fMRI study. *J Neurol Neurosurg Psychiatry*. 2010;81:788-792.
37. Calautti C, Naccarato M, Jones PS, *et al.* The relationship between motor deficit and hemisphere activation balance after stroke: A 3 T fMRI study. *Neuroimage*. 2007;34(1):322-331.
38. Cramer SC, Crafton KR. Somatotopy and movement representation sites following cortical stroke. *Exp Brain Res*. 2006;168(1–2):25-32.
39. Wu CY, Hsieh YW, Lin KC, *et al.* Brain reorganization after bilateral arm training and distributed constraint-induced therapy in stroke patients: A preliminary functional magnetic resonance imaging study. *Chang Gung Med J*. 2010;33(6):628-638.
40. Yozbatiran N, Cramer SC. Imaging motor recovery after stroke. *NeuroRx*. 2006;3(4):482-488.
41. Fu MJ, Daly JJ, Cavusoglu MC. *Assessment of EEG event-related desynchronization in stroke survivors performing shoulder-elbow movements*. In: Proceedings 2006 IEEE international conference on robotics and automation. IEEE; 2006:3158-3164.
42. Rossiter HE, Boudrias MH, Ward NS. Do movement-related beta oscillations change after stroke? *J Neurophysiol*. 2014;112(9):2053-2058.
43. Levin MF, Kleim JA, Wolf SL. What do motor "recovery" and "compensation" mean in patients following stroke? *Neurorehabil Neural Repair*. 2008;23:313-319.
44. Sitaram R, Ros T, Stoeckel L, *et al.* Closed-loop brain training: The science of neurofeedback. *Nat Rev Neurosci*. 2016;18:86-100.
45. McFarland DJ, Sarnacki WA, Wolpaw JR. Effects of training pre-movement sensorimotor rhythms on behavioral performance. *J Neural Eng*. 2015;12(6):066021.
46. Pichiorri F, Morone G, Petti M, *et al.* Brain-computer interface boosts motor imagery practice during stroke recovery. *Ann Neurol*. 2015;77(5):851-865.
47. Norman SL, Lobo-Prat J, Reinkensmeyer DJ. *How do strength and coordination recovery interact after stroke? A computational model for informing robotic training*. In: 2017 International conference on rehabilitation robotics (ICORR). IEEE; 2017:181-186.
48. Grefkes C, Fink GR. Reorganization of cerebral networks after stroke: New insights from neuroimaging with connectivity approaches. *Brain*. 2011; 134:1264-1276 awr033.
49. Kim SG, Ashe J, Georgopoulos AP, *et al.* Functional imaging of human motor cortex at high magnetic field. *J Neurophysiol*. 1993;69(1):297-302.
50. Pfurtscheller G, da Silva FH L. Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clin Neurophysiol*. 1999;110(11):1842-1857.
51. Zemke AC, Heagerty PJ, Lee C, Cramer SC. Motor cortex organization after stroke is related to side of stroke and level of recovery. *Stroke*. 2003;34(5):e23-e26.
52. Muellbacher W, Richards C, Ziemann U, *et al.* Improving hand function in chronic stroke. *Arch Neurol*. 2002;59(8):1278-1282.
53. Nudo RJ, Wise BM, SiFuentes F, Milliken GW. Neural substrates for the effects of rehabilitative training on motor recovery after ischemic infarct. *Science*. 1996:272(5269):1791-1794.
54. Weiller C, Ramsay S, Wise R, Friston K, Frackowiak R. Individual patterns of functional reorganization in the human cerebral cortex after capsular infarction. *Ann Neurol*. 1993;33(2):181-189.
55. Chollet F, DiPiero V, Wise R, Brooks D, Dolan RJ, Frackowiak R. The functional anatomy of motor recovery after stroke in humans: A study with positron emission tomography. *Ann Neurol*. 1991;29(1):63-71.
56. Murase N, Duque J, Mazzocchio R, Cohen LG. Influence of interhemispheric interactions on motor function in chronic stroke. *Ann Neurol*. 2004;55(3):400-409.
57. Dong Y, Dobkin BH, Cen SY, Wu AD, Winstein CJ. Motor cortex activation during treatment may predict therapeutic gains in paretic hand function after stroke. *Stroke*. 2006;37(6):1552-1555.
58. Christopher deCharms R. Applications of real-time fMRI. *Nat Rev Neurosci*. 2008;9(9):720-729.
59. Mozaffarian D, Benjamin EJ, Go AS, *et al.* Heart disease and stroke statistics-2015 update: A report from the American heart association. *Circulation*. 2015;131(4):e29.
60. Rizzo V, Richman J, Puthanveettil SV. Dissecting mechanisms of brain aging by studying the intrinsic excitability of neurons. *Front Aging Neurosci*. 2015;6:337.
61. Wolpaw JR, O'Keefe JA. Adaptive plasticity in the primate spinal stretch reflex: Evidence for a two-phase process. *J Neurosci*. 1984;4(11):2718-2724.
62. Zhou X, Tien RN, Ravikumar S, Chase SM. Distinct types of neural reorganization during long-term learning. *J Neurophysiol*. 2019;121:1329-1341.
63. McFarland DJ, Norman SL, Sarnacki WA, Wolbrecht ET, Reinkensmeyer DJ, Wolpaw JR. BCI-based sensorimotor rhythm training can affect individuated finger movements. *Brain-Comput Interfaces*. 2020;7:38-46.
64. Aflalo T, Kellis S, Klaes C, *et al.* Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science*. 2015;348-(6237):906-910.
65. Norman SL, Maresca D, Christopoulos VN, *et al.* Single-trial decoding of movement intentions using functional ultrasound neuroimaging. *Neuron*. 2021;109(9):1554-1566.e4.
66. Hudspeth AJ, Jessell TM, Kandel ER, Schwartz JH, Siegelbaum SA. *Principles of neural science*. McGraw-Hill; 2013. Health Professions Division.
67. Klein J, Spencer SJ, Reinkensmeyer DJ. Breaking it down is better: Haptic decomposition of complex movements aids in robot-assisted motor learning. *IEEE Trans Neural Syst Rehabil Eng*. 2012;20(3):268-275.
68. Krebs HI, Hogan N, Aisen ML, Volpe BT. Robot-aided neurorehabilitation. *IEEE Trans Rehabil Eng*. 1998;6(1):75-87.
69. Bunday KL, Urbin MA, Perez MA. Potentiating paired corticospinal-motoneuronal plasticity after spinal cord injury. *Brain Stimulat*. 2018;11(5):1083-1092.