Implementing an Integrate-and-Fire Neural Network on a Bidirectional Brain-Computer Interface

Jonathan Mishler

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Eberhard Fetz, Chair

Rajesh Rao

Steve Perlmutter

Azadeh Yazdan-Shahmorad

Program Authorized to Offer Degree:

Bioengineering

University of Washington

Abstract

Implementing an Integrate-and-Fire Neural Network on a Bidirectional Brain-Computer Interface

Jonathan Mishler

Chair of the Supervisory Committee:

Eberhard Fetz

Department of Physiology and Biophysics

There is a growing interest in using intracortical microstimulation (ICMS) as a means of neurorehabilitation, from using it to rewire synaptic connections in the brain, to providing a means of providing artificial sensations to provide feedback to patients controlling external devices such as robotic limbs via decoded neural activity. Towards the goal of neurorehabilitation, there have been recent efforts to integrate artificial neural networks (ANNs) with the brain and train them to manipulate neuronal activity in a context-specific manner with the goal of developing brain-computer interfaces (BCIs) that can restore function to damaged neural circuitry. In this thesis, I present my work whose aim was to interface artificial spiking neurons with biological neurons in primary motor cortex to create hybrid biological/artificial neural networks that altered the firing dynamics of the biological neurons. In these hybrid networks, spikes that are detected from the biological neurons send artificial postsynaptic potentials (PSPs) to the artificial neurons, whose magnitudes and polarities are defined by weights that characterize the strengths of their connections. When the membrane potentials of the

artificial neurons exceed a predetermined threshold, they spike, and subsequently trigger ICMS to manipulate the activity of the biological neurons.

We first characterize the effects of ICMS on neural activity in primary motor cortex (M1) of pigtail macaques, and show how it elicits a brief excitatory response followed by a longer inhibitory response. We then show how the probability of evoking single action potentials has at least three dependences – the stimulation amplitude, time delay between the neuron's previous spike and stimulation onset, and its firing rate. Finally, we show how repetitive stimulation can increase or decrease the probability over time, likely due to mechanisms of short-term plasticity.

I used these results to explore how various properties of the hybrid biological/artificial neural networks shape the closed-loop dynamics between the biological and artificial neurons. To do so, I measure changes in the auto-, and cross-correlograms of the biological neurons between their spontaneous and closed-loop dynamics, and show how features within the correlograms are related to the size, connectivity, number of hidden layers, magnitude of inhibition, and stimulation delays of the network.

I then show how the closed-loop dynamics can be simulated by recording the spontaneous activity as well as the open-loop responses of the biological neurons to ICMS, and use the model to further characterize how the features of the hybrid network influence the closed-loop dynamics free from experimental constraints such as the non-stationarity of the firing rates and evoked spike probabilities of the biological neurons over time. I also use the model to explore how stimulation artifacts impacted the closed-loop operation of the hybrid neural network by obstructing the detection of spikes. To do so, I characterize the number of obstructed spikes as a function of both the network size and connection strength between the biological and artificial neurons. I found that compared to the size of the network, the strength of the connections

between the biological and artificial neurons was the greater determinant of how many spikes were blanked.

Lastly, I discuss preliminary work of the development of a computational model whose goal is to demonstrate that ANNs can be interfaced with the brain to restore lost motor function. I first train a small recurrent ANN, which simulates a motor circuit within the brain, to perform a 1D motor task with a reinforcement learning algorithm. I then lesion the network by deleting a subset of the neurons and their related connections, and show how even with retraining, the network is incapable of relearning the task. Finally, I discuss future steps in which a second ANN, whose outputs use ICMS to activate neurons within the original ANN, can be used to restore network function and task performance.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would first like to acknowledge the members of my lab for their support and guidance throughout my time at the University of Washington. I am extremely grateful to have been surrounded by such a warm and supportive group of people who helped me succeed. It has been once in a lifetime experience to work on such exciting projects in the Washington National Primate Center. I am incredibly fortunate to have had the leadership of Eb and Raj, who have been two of my biggest cheerleaders throughout my PhD. Through their mentorship, I have developed into the scientist I have always aspired to be. I am particularly thankful for their never-ending encouragement, understanding, and patience which gave me the confidence to push through difficult times. I would also like to also acknowledge Steve Perlmutter. While Steve was never officially listed as my principal investigator, he attended almost every lab meeting of mine and was always happy to give me feedback, encouragement, and his support in every matter.

I would like to give a very special thanks to my friend and colleague, Richy Yun. We both joined the lab around the same time and supported each other the entire way. He is the only person in the world that can truly relate to the experience that I went through. I thank him for all of his enthusiastic support, feedback, and friendship. I would also like to thank Irene Rembado. I cannot count the number of times that she dropped what she were doing in lab to help me out when I needed technical support, advice, or just a person to vent to. I would like to thank her for never making me feel judged, and for always being someone I could rely on. It was largely because of her support that I was able to carry on at times. She has set the standard for what I want to be as a postdoc, and I will do my best to follow in her footsteps and return the favor to graduate students in my lab when I am a postdoc.

I want to also give a warm thanks to Larry Shupe for the countless number of times that he gave me technical assistance, as well as the numerous times that he dropped what he was doing to program features on the Neurochip3 for me.

Another person in the lab that I would like to thank is Samira Moorjani for her surgery assistance, journal club contributions, and friendship. I always looked forward to eating dinner with her family. I would also like to thank Becky Burch for her enthusiastic general lab assistance.

I would also like to thank the primate center support staff, including the machine shop and surgery assistance for being so accommodating.

I also acknowledge my previous undergraduate and postbac research mentors. Dr. Keith Roper gave me the opportunity to work in his lab as a freshman with no prior research experience, which set me up for the amazing career I have now. I later joined Dr. Joseph Herzog's lab who helped me greatly improve my computational, written, and presentation skills. Last, I would like to acknowledge Dr. Dietmar Plenz for mentoring me at the NIH. He gave me a lot of flexibility in my projects which allowed me to not only improve on my programming skills, but also helped my foster my scientific creativity.

Some other people that I would like to thank are my parents for their endless encouragement and support throughout my life. I would like to thank Joseph Strom and Claudia Quittner who were at times like a second set of parents to me. Finally, I would like to thank my grandfather. My grandfather immigrated to the United States from Poland in the late 1930's without a single penny to his name. Through his incredible hard work, self-sacrifice, and love for his family, I was able to pursue my dream to become a first-generation college student. Without him, none of this would have been possible.

Last but not least, I would like to thank and acknowledge my wonderful wife, Lauren Gentles. I went through a lot of extremely difficult times, including both the typical stressors of graduate school, as well as numerous health related issues. She was always willing to drop what she was doing to help me out, and was my rock these last five and a half years. I would not be where I am today without her loving support, and am truly grateful to have her in my life. Also, thanks for finding my postdoc for me.

Chapter 1.

INTRODUCTION

1.1 The use of stimulation for neurorehabilitation

1.1.1 The need to develop novel therapeutics for neurological diseases and disorders

The Global Burden of Disease Study (GBD) released a recent study showing that neurological diseases have been rising throughout the world [1], [2]. According to the study, neurological diseases and disorders are the leading cause of disability and the second leading cause of death worldwide [1]. Increasing global lifespans, unhealthy diets, and other factors such as smoking have all contributed to these increases [1], [2]. Stroke is responsible for the majority of these increases, and accounts for 42.2% of disability-adjusted life years, and 64.4% of deaths [2]. These disabilities are not only physically and emotionally crippling to the afflicted, but also substantially contribute to a loss of economic productivity. Currently, there are limited treatment options for most of these disorders, and the treatments that do exist often rely on early intervention for them to be efficacious [3], [4]. Moreover, diseases such as stroke have a variable scope in both the location and extent of damage [4], [5], which further complicates treatment. Therefore, there is a great need to develop novel treatments for these disorders that can be administered to patients in both early and advanced stages of their diseases, irrespective of disease location and severity.

1.1.2 Stimulation as a means of promoting synaptic plasticity

There have been numerous efforts to use invasive stimulation of various modalities in both open-, and closed-loop manners to promote plasticity in the central nervous system, with the goal of developing protocols to apply in patients with neurological damage [6]–[12]. Due to a smaller risk of serious complications, noninvasive stimulation has also provided researchers with the opportunity to test various protocols for the treatment of patients with neurological disorders such as stroke. Altogether, the invasive and noninvasive methods include transcranial magnetic stimulation (TMS) [10], [11], [13], transcranial

direct current stimulation (TDCS) [10], intracortical microstimulation (ICMS) [7]–[9], epidural [10], and optogenetic stimulation [6], [12]. The mechanisms of rehabilitation have largely been attributed to various mechanisms including increased neural excitability, Hebbian mechanisms of synaptic plasticity, and structural plasticity [10], [11], [14]. Paradigms that leverage increased neural excitability operate by enabling the damaged circuits to function with less synaptic excitation, which ultimately strengthens the spared connections via activity-dependent plasticity to promote lasting functional rehabilitation. Other work has utilized stimulation to strengthen connections in the brain and spinal cord via Hebbian mechanisms of synaptic plasticity [10], [14]. The Fetz lab has previously demonstrated that electrical stimulation in the primate brain can be used to strengthen connections in the brain via both open-, and closed-loop ICMS, as measured via stimulation-evoked movements and cortico-cortical evoked potentials [7]–[9]. Other work has explored how optogenetic stimulation can be used to induce Hebbian synaptic plasticity. Azadeh-Yazdan et al. used μ-ECoG arrays in combination with open-loop optogenetic stimulation in both the primary motor and sensory cortices to facilitate changes in their functional connectivity, which they measured using both the evoked responses throughout the array and coherence between the recorded sites [6]. In addition to invasive means of stimulation, other work has used noninvasive forms of stimulation to promote Hebbian mechanisms of rehabilitation. Revill et al. used TMS of the ipsilesional primary motor cortex of stroke patients paired with motor hand training to promote lasting recovery compared to motor hand training alone [11]. Finally, it has been well-established that structural plasticity plays an important role in stroke rehabilitation. Not only do new synaptic connections form in both the intralesional and contralesional cortices following stroke [10], [14], [15], it has been demonstrated that electrical stimulation of the contralesional cortex can promote the growth of new contralesional corticospinal projections in rats [10]. Altogether, these works demonstrate the capabilities of stimulation to promote rehabilitation in patients following stroke.

1.1.3 Intracortical microstimulation for brain-computer interface control

Stimulation is also being explored for its uses in bidirectional brain-computer interfaces (BCIs) utilizing its ability to elicit artificial cutaneous and proprioceptive sensations in the subject [16]–[19]. In these systems, the neural activity of the user are decoded to control an external device such as a robotic arm, while stimulation is delivered to the somatosensory cortex to provide feedback, without which, the subject's ability to interact with the outside world is limited. Non-human primates (NHPs) have also demonstrated their capabilities of utilizing these artificial sensations to perform BCI tasks [16], [17], [19]. To these ends, there are efforts underway to understand how the frequency and amplitude of ICMS shape the artificial sensations [16]. One limitation of this approach is that the sensations elicited by ICMS are typically artificial and are described as tingling or buzzing, rather than being natural. However, more recent research has demonstrated that ICMS can also elicit natural sensations, and that stimulation amplitude, not frequency, is the determining factor in which sensory modality – cutaneous or proprioceptive – is experienced by the subject [17]. Interestingly, ICMS is not restricted to eliciting a single type of sensory modality, and can often elicit both cutaneous and proprioceptive sensations [17]. Stimulation frequency has also been shown to increase the sensitivity with which the subjects can detect the ICMS [16]. Non-human primates (NHPs) have been shown to be able to discriminate between different frequencies, although performance has been shown to taper at frequencies greater than 200 Hz [16].

1.1.4 Deep brain stimulation

A common means of using electrical stimulation for neurorehabilitation is deep brain stimulation (DBS). The most common clinical application of DBS is in the treatment of Parkinson's disease, although it is also used to treat essential tremor, dystonia, pain, multiple sclerosis, and psychiatric disorders such as depression, anxiety, obsessive-compulsive disorder, and Tourette Syndrome [20]–[23]. DBS electrodes are commonly placed in subthalamic nuclei, the thalamus, or global pallidus, after which high frequency stimulation of approximately 130 Hz is applied to disrupt the abnormal brain signals [20], [22], [23]. It is currently believed that there are numerous mechanisms of action for DBS, including the manipulation of

neuronal activity, release of neurotransmitters and neuromodulators, and synaptic plasticity in the stimulated region [22]. There are also efforts underway to develop closed-loop systems that trigger stimulation from neural activity, rather than applying the stimulation continuously [20]–[23]. The closed-loop DBS systems may have the advantage of increased efficacy, and longer battery life [23].

1.2 The effects of ICMS on neuronal activity

1.2.1 ICMS primarily activates neurons transsynaptically

ICMS is known to activate neurons via two primary methods – direct and indirect activation [24], [25]. Direct activation occurs when ICMS excites the soma or axon of a neuron [24], [25]. Indirect activation occurs when ICMS activates neurons transsynaptically [24], [25]. Although antidromic activation of neurons does occur, it has been shown using antidromic collision tests that the primary means of activation are orthodromic in nature [26]. Transsynaptic activation of neurons allows ICMS to excite axons that project to neurons up to several millimeters away from the stimulation site [24], [27]. It has also been shown that the probability of evoking spikes increases as a function of the current amplitude due the larger spread of current which activates a greater number of axons [26]. Furthermore, the background network activity has also been shown to influence the activation of neurons both *in vitro* [28], [29] and *in vivo* [30]; a greater level of background activity increases the probability of exciting neurons.

1.2.2 The excitatory response following ICMS is followed by long-lasting inhibition

The brief excitatory response elicited by ICMS is followed by a prolonged period of inhibition which can last over 200 ms [26], [31]. The inhibitory response is commonly attributed to GABAergic synapses that are thought to be activated via feedforward and feedback circuits. GABA B receptors have been implicated in this long lasting inhibition in rodents [26], [31]. For this study, Butovas et al. pharmacologically blocked GABA B receptors in mouse neocortex which reduced the duration of inhibition. Furthermore, they showed that repetitive stimulation sublinearly increased the length of inhibition; however, this was eliminated when the GABA B receptors were blocked [31]. They also showed that the inhibition was weaker and longer in

connexin 36 knockout mice – a protein which forms gap junctions in neocortical cells to electrically couple them and synchronize their activity – which also demonstrates that gap junctions contribute to the inhibitory response by synchronizing the activations of the inhibitory neurons [31]. Although the duration of inhibition was different, other work done in NHPs have shown a similar phenomenon in which inhibition occurs immediately after the excitatory response [32]. Interestingly, they found that the inhibitory response tapers in strength more quickly with distance compared to the excitatory response.

1.2.3 Repetitive stimulation modulates the responses of neurons to stimulation

The responses of neurons to ICMS are dynamic and can change with repetitive stimulation. High-frequency stimulation in particular has been shown to decrease both the probability and spatial extent of stimulation over time [33]. Although the authors in this study did not conclude that any particular mechanism was responsible for these changes, short-term decreases in the activation of neurons are commonly attributed to short-term synaptic depression [34]–[36]. Short-term depression is commonly attributed to vesicle depletion, but can also manifest from postsynaptic mechanisms in which the ligand-gated receptors progressively become desensitized to the corresponding neurotransmitter over time [34]. Low-frequency stimulation can also lead to short-term synaptic potentiation, in which the probability of evoking spikes increases over time. Due to the influx of calcium into the presynaptic terminal following action potentials, subsequent presynaptic action potentials can lead to greater concentrations of calcium in the axonal terminal if the second action potential occurs before all the calcium has left [34]–[36]. The increased calcium concentration increases the release probability of synaptic vesicles, thereby increasing the amount of neurotransmitter released, which increases the probability of evoking a spike. Finally, short-term changes in synaptic strength have been shown to alter the latency of the evoked spikes – with short-term depression increasing the latency of the evoked spike and facilitation decreasing the latency [37].

1.3 Artificial neural networks

1.3.1 The evolution of artificial neural networks

The first concept of artificial neurons, known as McCulloch-Pitts (MCP) neurons, was developed by Warren McCulloch and Walter Pitts Jr. in 1943 [38]; these were designed to represent linearly separable Boolean functions. MCP neurons receive binary inputs from a collection of signals. Each connection has a weight of -1, 0, or 1, corresponding to an inhibitory connection, no connection, and an excitatory connection, respectively to the MCP. The MCP neuron has a binary output, which is 1 if its output exceeds a threshold value, and 0 if it doesn't. The primary limitations of the MCP neuron model are that the model must be preprogrammed, and cannot learn or perform nonlinear computations.

The MCP model, colloquially referred to as the first generation of artificial neural networks [39] was improved on by Frank Rosenblatt in 1958 at Cornell Aeronautical Laboratory, who developed what is known as the perceptron [40]. The perceptron improved on the MCP model by allowing real-valued inputs and biases in the artificial neuron. Furthermore, Rosenblatt developed an algorithm to train the weights to map unique stimuli to unique outputs [40]. However, because the perceptron was a simple binary classifier, it was incapable of computing nonlinear functions such as the XOR function, as shown by Marvin Minsky and Seymour Papert in 1969 [41].

The concept of perceptrons was later generalized to what are now known as artificial neural networks (ANNs), which are the second generation of neural networks. ANNs extend the perceptron model by allowing the output of the artificial neurons to take on real values instead of binary values [39]. They also introduce the concept of activation functions, which are nonlinear functions that are applied to the output of each artificial neuron after the weighted inputs are summed together with the biases [42]. These activation functions add an element of nonlinearity to the neural networks which allow them to compute nonlinear functions [42]. Examples of such activation functions include sigmoidal and rectified linear (ReLU) functions. The ANNs can have several hidden layers (deep learning), and multiple outputs. The ANNs can be trained for regression or multiclass classification, and are typically trained via the gradient

descent algorithm [39], [43]. There are several types of ANNs including feedforward, convolutional, and recurrent neural networks, each of which with its own purpose [43]. Convolutional neural networks are frequently used in image processing applications, while recurrent neural networks are used when the neural network needs a memory of past activity. Although ANNs represent a significant step forward from the original perceptron, they have a significant limitation in that they are rate-based neural networks [39]. For instance, experimental evidence suggests that neural computations based on firing rate alone likely cannot account for the rapid visual pattern analysis and classification performed by both NHPs and humans [44]–[46]. Therefore, a neural network that can encode information in spike timing rather than rate would not only be more biomimetic, but would theoretically be able to be much more computationally powerful [39].

To investigate computation mediated by action potentials, spiking neural networks (SNNs) were developed. Like ANNs, neurons within SNNs receive inputs, but rather than having continuous outputs, they communicate with discrete binary events, or "spikes" [39]. The most commonly used type of spiking neuron is the leaky integrate-and-fire neuron (LIF) [47], [48]. LIF neurons receive inputs which raise their membrane potentials. A leaky term is added to the model which slowly decays the membrane potential of each neuron. The neurons spike when they reach a predetermined threshold, after which they reset to a baseline level. There are several adaptions of the LIF model including the integrate-and-fire-or-burst, quadratic integrate-and-fire, resonate-and-fire, and integrate-and-fire with adaptation model, each of which has their own unique advantages [47], [48]. Another prominent spiking neuron model is the Izhikevich neuron, which combines the sophistication of the Hodgkin-Huxley model with the low computational cost, measure in floating point operations (FLOPS), of integrate-and-fire models [47], [49]. The Izhikevich neuron model is capable of simulating several different kinds of neurons found in the brain, which adds a greater degree of complexity to the model [47], [49].

1.3.2 Increasing the complexity of spiking neural networks

Beyond the various types of spiking neuron models, there are several other features that can be incorporated into the models to make them more biomimetic and increase their complexity. Synapses can be modelled

as current-, or conductance-based [50]. Current-based synapses send post-synaptic potentials (PSPs) from the presynaptic neuron to the postsynaptic neuron with a constant amplitude regardless of the membrane potential of the postsynaptic neuron. Conductance-based synapses are more biomimetic in that they incorporate the concept of reversal potential into the model. For instance, a more depolarized postsynaptic neuron will have a lower sodium conductance, thereby lowering the size of the PSPs.

Spiking neurons can also have adaptive thresholds that increase after spiking. As a neuron continues to spike in close succession, its threshold continues to rise and will exponentially decay over time [51]–[53]. The threshold can also be stochastic, rather than deterministic [51], [54]. Neurons with stochastic thresholds are not guaranteed to spike at any given threshold but are progressively more likely to fire as their membrane potentials increase. There is some evidence that stochastic thresholds increase the ability of the SNN to more efficiently separate different classes of inputs, thereby leading to better function [51].

1.3.3 Synaptic plasticity in spiking neural networks

SNNs can also include other biological mechanisms such as short-, and long-term synaptic plasticity [55]–[71]. Some of the most commonly used models for short-term plasticity utilize data from Markram which show how different types of connections in the brain (i.e., excitatory-inhibitory versus inhibitory-inhibitory connections) have unique manifestations of short-term plasticity [72], [73]. The incorporation of short-term synaptic plasticity in SNNs has been shown to greatly increase their information processing capabilities, for instance by helping them maintain working memories[55], [71]. Long-term synaptic plasticity is also frequently incorporated into SNN models. The ability of SNNs to communicate with spike times lends itself to biologically inspired learning algorithms such as reinforcement learning and spike-timing dependent plasticity (STDP). These training have been used to train SNNs for both supervised [67], unsupervised [70], [74], and reinforcement learning tasks [75]–[78]. In addition to the potentiation and depression of existing synaptic connections, new algorithms are being developed to flexibly add or prune existing synaptic connections of the SNN [60], [79]. These training algorithms allow the SNN to optimize their connectivity by retaining the useful connections and eliminating the unnecessary ones. This type of learning

is particularly interesting since it allows the SNN to continuously learn new, while forgetting unnecessary information, just like the brain. Furthermore, because the structure of neural networks is paramount to their computational capabilities, the ability to rewire themselves is a great advantage to maximize their ability to learn. Many of these training algorithms also utilize homeostatic mechanisms of plasticity to prevent overexcitation of the neurons [80], [81].

1.4 Layout of dissertation

The following chapters detail my work which first investigates the responses of individual neurons in primary motor cortex (M1) of pigtail macaques to single-pulse ICMS, then explores how biological neurons can be interfaced with integrate-and-fire neurons (which will henceforth be called integrate-and-fire units to distinguish them from the biological neurons) to manipulate their connectivities and firing dynamics.

In **Chapter 2** I show how ICMS impacts the spikes of individual cortical neurons in M1 of two pigtail macaque monkeys. This characterization includes the excitation and subsequent inhibition of activity following single-pulse ICMS, the dependencies that the neurons have on their evoked spike probabilities, and how their responses change over time.

In **Chapter 3** I explore how various properties of the SNN, including network architecture and stimulation delays impact the closed-loop dynamics between the biological neurons and artificial integrate-and-fire units. I then show how a model consisting of two elements – a Markov model to simulate spontaneous spike trains, and a stimulation response model can reproduce the closed-loop experimental dynamics. I then use the model to further probe how various features of the neural network, including inhibition strength and stimulation delay influence the closed-loop dynamics free from experimental and hardware constraints. Finally, I explore how the stimulation artifacts impact the functioning of the closed-loop operation by blanking spikes as a function of the network size and synaptic weight.

In **Chapter 4** I discuss preliminary work whose goal is demonstrate in a computational model how an SNN could conceivably be used to restore motor function in a damaged biological network via ICMS. For this

project, I first train a recurrent SNN (SNN_1) with a reinforcement learning algorithm to track the movement of a target with a cursor in 1D. I then lesion SNN_1 by deleting a random subset of its neurons and related connections, thereby compromising its ability to perform the task, and show how even with retraining, is incapable of relearning the task. I then discuss future steps towards interfacing a second SNN (SNN_2) with SNN_1 to restore its proper function via ICMS.

Chapter 2.

RESPONSES OF CORTICAL NEURONS TO INTRACORTICAL MICROSTIMULATION IN AWAKE PRIMATES

[Paper by Richy Yun*, Jonathan Mishler*, Steve Perlmutter, Rajesh Rao and Eberhard Fetz]

2.1 Abstract

Intracortical microstimulation (ICMS) is commonly used in many experimental and clinical paradigms; however, its effects on the activation of neurons are still not completely understood. To document the responses of cortical neurons in non-human primates to stimulation, we recorded single unit activity while delivering single-pulse stimulation via Utah arrays implanted in primary motor cortex of three macaque monkeys. Stimuli between 5-50 μA delivered to single channels reliably evoked spikes in neurons recorded throughout the array with delays of up to 12 milliseconds. ICMS pulses also induced a period of inhibition lasting up to 150 ms that typically followed the initial excitatory response. Higher current amplitudes led to a greater probability of evoking a spike and extended the duration of inhibition. The likelihood of evoking a spike in a neuron was dependent on the spontaneous firing rate as well as the delay between its most recent spike time and stimulus onset. Repetitive stimulation often modulated both the probability of evoking spikes and the duration of inhibition, although high frequency stimulation in particular was more likely to change both responses. Comparisons between stimuli that evoked a spike and stimuli that did not revealed that the excitatory and inhibitory responses were independent of one another on a trial-by-trial basis; however, their changes over time were frequently positively or negatively correlated. Our results document the complex dynamics of cortical neural responses to electrical stimulation that need to be considered when utilizing ICMS for scientific and clinical applications.

2.2 Introduction

Intracortical microstimulation (ICMS) is widely used for interfacing with the brain in both basic and clinical research, from inducing plasticity to employing sensory neuroprostheses in various animal models [82]–[85]. The applicability of ICMS arises from the fact that it has the highest spatial and temporal specificity of all clinically applicable cortical stimulation techniques [86]. However, the circuit mechanisms that drive the responses of neurons following ICMS, and the ways in which other factors such as timing and stimulation frequency affect the stimulus responses are not fully understood.

ICMS was originally thought to activate neural elements around the electrode tip. Regions closer to the tip would have higher activation in a sphere with an isotropic gradient, and the volume would grow with increasing current amplitude [25], [87], [88]. However, it is now accepted that ICMS predominantly excites axons near the electrode tip that transsynaptically excite neurons up to several millimeters away, favoring pathways with similar functional properties [27], [89]–[93]. Additionally, the effects of ICMS are not limited to excitation, and includes a long-lasting inhibitory response that is commonly attributed to GABAergic synapses [26], [31], [90], [94].

Single neuron responses to ICMS are dynamic and can be modulated with repeated stimulation. The changes, in part, also depend on stimulus frequency. In particular, the excitation of neurons generally decreases over time and becomes more localized with higher frequencies [95], [96]. However, the reported frequency ranges and timescales are variable, and the driving processes remain unclear. The changes over time are often attributed to short-term synaptic plasticity or intrinsic plasticity of neurons which both depend on the frequency and pattern of stimulation [97]–[99].

Altogether, these studies demonstrate that the effects of ICMS are not restricted to regions proximal to the electrode tip, and that responses consist of interplay between excitation and inhibition [92], [100], [101]. Despite our increasing understanding of how ICMS activates cortical circuits, several significant questions remain. How does the background neuronal activity, including firing rate and previous spike time impact the stimulus response? How do the responses change over time as a function of both the frequency of stimulation and proximity to stimulation site? Is the inhibitory response coupled to the excitatory response, or are they independent?

We addressed the questions above by delivering ICMS and examining responses of single neurons in primary motor cortex (M1) of three macaque monkeys with chronically implanted Utah arrays. Single-pulse ICMS was delivered to one channel for up to 20 minutes while the spikes of single neurons were simultaneously recorded from all other electrode channels. We tracked the probability of evoking spikes as well as the duration of the evoked inhibition and varied both the sites and frequency of the stimulation

between sessions. Our results expand upon previous findings by characterizing the dependencies of the neuronal responses to background neuronal activity, distance, and stimulation frequency, and exploring the interactions between the excitatory and inhibitory responses.

2.3 Materials and Methods

2.3.1 Implants and surgery

Three pigtail macaque monkeys (Macaca nemestrina) were unilaterally (right hemisphere) or bilaterally implanted with 96-channel Utah microelectrode arrays (Blackrock Microsystems; 10x10, 400μm inter-electrode distance, 1.5 mm depth, Iridium oxide) in the hand region of M1. Sterile surgeries were performed under isoflurane anesthesia and aseptic conditions with continuous monitoring of all vitals. Animals received postoperative courses of analgesics and antibiotics following each surgery. All procedures conformed to the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the University of Washington Institutional Animal Care and Use Committee.

Implantation of the arrays was guided via stereotaxic coordinates. A 1.5 cm wide square craniotomy centered at 4 mm lateral of Bregma was performed to expose the dura. Three sides of the exposed dura were cut to expose the cortex, after which a Utah microelectrode array was implanted. Two reference wires were inserted under the dura and two were inserted between the dura and the skull. The dura was then sutured around the implant, and the bone flap from the craniotomy was reattached to the skull with a titanium strap and titanium skull screws. A second, smaller titanium strap was screwed onto the skull to secure the wire bundle that connected the array to a connector "pedestal" that was also secured to the skull with eight skull screws. The skin incision was then sutured around the pedestal base.

To facilitate the chronic recording of neuronal activity, the monkeys were also implanted with halos made with 3/8" aluminum bars in an egg-shaped oval that was 17 cm long and 15.3 cm wide. Four titanium straps were affixed to the skull via titanium bone screws. Two of the straps were implanted bilaterally over the occipital ridge, and two were bilaterally implanted temporally. After the plates integrated with the skull for

6 weeks, the halo was secured to the skull with four pins, each of which were seated in one of the four plates.

2.3.2 Experimental setup

The monkeys were trained to calmly sit in a primate chair while periodically receiving an apple smoothie reward without performing a task (Figure 2.1). Each session included a prestimulus epoch lasting between 5 and 10 minutes and a stimulus epoch lasting between 5 and 20 minutes. During the stimulus epoch we delivered tonic or Poisson distributed single pulse stimulation (cathodal, biphasic, 200µs phase width) to a single channel at rates between 1-20 Hz. For testing the effects of current amplitude, a range of 2-50 µA was used. Current amplitude was fixed at 15 µA for all other experiments and analyses. The stimulation frequency was fixed during the stimulus epoch for each session.

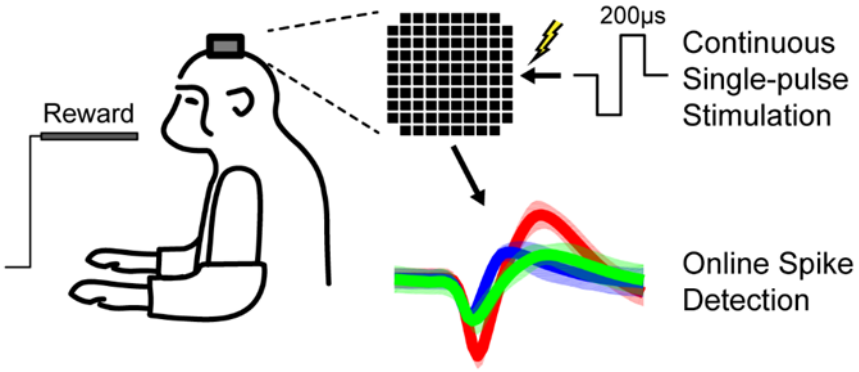

Figure 2.1 Experimental Design

Macaques calmly sat in a chair receiving apple smoothie reward through the experiment. Cathodic, 200us phase width, single-pulse ICMS was delivered to one channel of the Utah array in primary motor cortex while unit responses were recorded across the array. Each session consisted of a prestimulus and stimulus epoch.

2.3.3 Evoked spike acquisition

Spikes were sorted using two time-amplitude windows, initially online and subsequently confirmed offline. Stimulus artifacts lasted around 1.1-1.6 ms. Spikes were frequently detected immediately following the

artifact (Figure 2.2A). The timing of evoked spikes was found by calculating the peristimulus time histogram (PSTH, 0.5 ms bin widths) of spikes in the window from -20 to 20 ms from the time of stimulation. To isolate the evoked spikes from the spontaneous activity, we defined upper and lower thresholds in the PSTH as the histogram mean plus or minus 2 times the standard deviation from -20 to -2 ms. We then found the largest peak in the PSTH from 1 to 15 ms after stimulation that was larger than the upper threshold and tracked adjacent bins in both directions until we reached the lower threshold on both sides. All spikes occurring within this window were denoted as stimulus-evoked spikes (Figure 2.2B). If no peak was greater than the threshold the spike was not considered to have been evoked by stimulation. The probability of evoked spikes was calculated as the number of evoked spikes divided by the number of total stimuli. For any analysis over time, the evoked spike probability was calculated for stimuli within overlapping 30 second bins with 1 second steps.

We tracked a total of 602 units, 17 of which were exclusively used to analyze the effects of stimulus timing on the probability of evoking a spike since the analysis required high firing and/or stimulation rates. For these experiments, stimulation was sometimes delivered to multiple channels in the same experiment. All units were used for all other analyses.

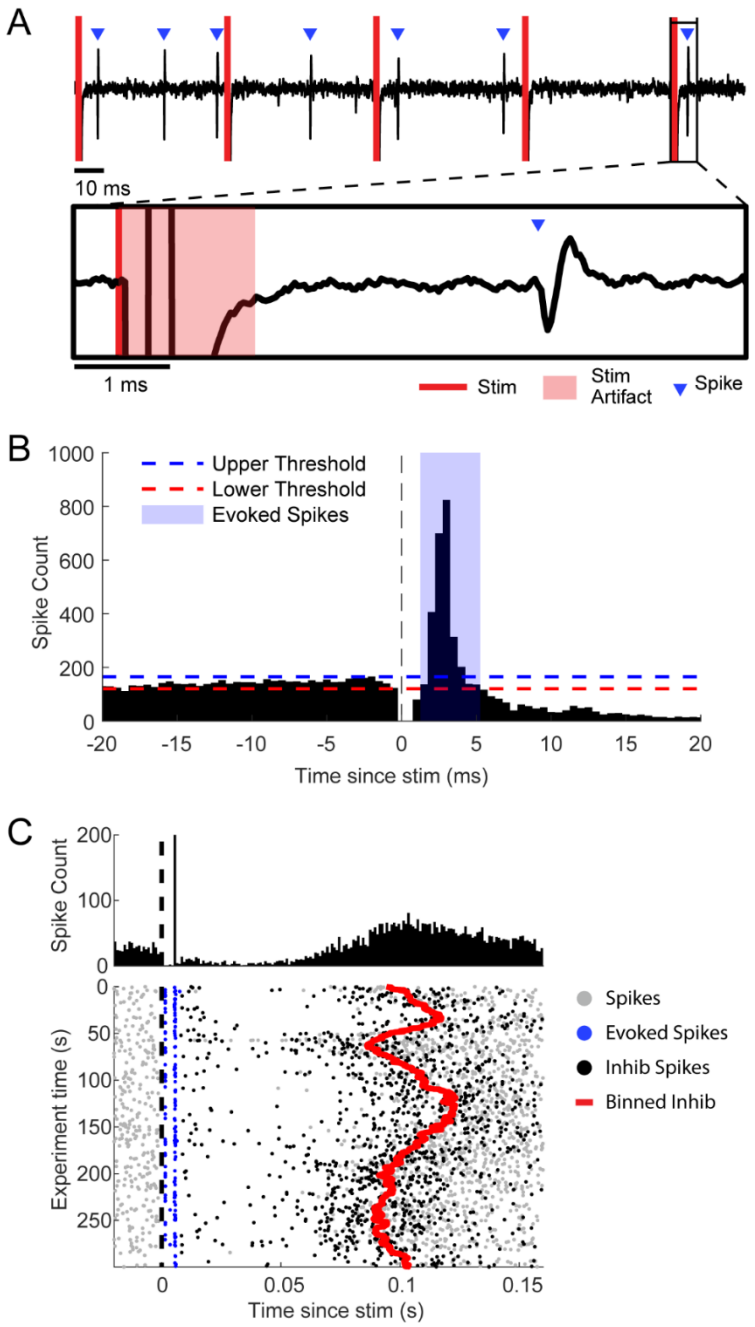

Figure 2.2 Detection of Evoked Spikes and Inhibition

A. Example of filtered data trace. The inset shows a stimulus followed by an evoked spike after 3.5 ms. **B.** Defining evoked spikes. A PSTH with 0.5 ms bins was generated. Peaks after the time of stimulation greater than the upper threshold (mean+2 standard deviations of -20 to -2ms in the PSTH) down to the lower threshold (mean-2 standard deviations) were called evoked spikes. **C.** Characterizing the duration of inhibition. Top: PSTH. Bottom: corresponding raster plot. The timing of the first spike that is not an evoked spike was recorded following each stimulus. The red line shows the median inhibition averaged over 30 second bins with 1 second steps.

2.3.4 Inhibitory response acquisition

The inhibitory response was measured using the PSTH in previous studies [26], [32]. While evoked spike timing can easily be determined with the PSTH due to their high probability and narrow time window, the inhibitory response depends on a broad window with sparse activity, particularly for units with low firing rate. As a result, a large number of stimuli is required to reliably detect inhibition in the PSTH. Therefore, rather than using the PSTH, we measured the duration of inhibition by removing all evoked spikes and calculating the time between the onset of stimulation and the next spontaneous spike (Figure 2.2C). Inhibition was deemed to be stronger when the delay from stimulation onset to the next spontaneous spike was longer, giving us a measure directly comparable to the PSTH but with much higher temporal resolution. We discarded any stimuli for which the subsequent stimuli occurred before the next spontaneous spike. For any analysis over time, we used the median inhibition duration of stimuli within overlapping 30 second bins with 1 second steps.

2.3.5 Evoked spike probability dependencies

To calculate the spontaneous firing rate for each unit over time, we disregarded the times from each stimulus onset to the next spontaneous spike. This effectively removed the stimulus response from the firing rate calculation, providing us with an independent measure of spontaneous activity.

The autocorrelation histograms for the 17 units in which we characterized the evoked spike probability as a function of the time delay between their previous spike times and stimulation onsets were calculated from their respective prestimulus epochs. The histograms were binned between 0 to 50 ms in non-overlapping 1 ms bins.

The dependencies of evoked spike probabilities on the timing of stimulation were calculated by first measuring the delays between each stimulus and the preceding spike. We separated this into two groups corresponding to whether the preceding spike was spontaneous or evoked. Stimuli that did not have a spike following the previous stimulus were discarded from the analysis. To account for the transmission delay

from the stimulated site to the evoked spike, we added the average evoked spike latency for each unit to the delays. We then calculated the probability of evoking a spike for stimuli with delays from 0 to 50 ms with 1 ms time bins, then applied a moving average with a 5 ms window. Since the average spike latencies were added to each delay, time bins that were less than the average evoked spike latency were not included in correlation calculations.

2.3.6 Changes over time

Pairwise correlations and their statistical significance between firing rate, evoked spike probability, and inhibition duration over time were calculated using the Pearson correlation coefficient r:

$$r = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \tag{2.1}$$

where $\bar{x}$ is the mean of the x-variable, and $\bar{y}$ is the mean of the y-variable. Firing rate was calculated by removing all times between each stimulus onset to the first spontaneous spike (the inhibitory period) to remove any confounding affects between firing rate and the inhibitory response.

Linear and exponential fits were performed on binned evoked spike probabilities and inhibition durations to determine changes over time due to repetitive stimulation:

Linear $$y = a * t + b \tag{2.2}$$

Exponential $$y = a + b * (1 - c)^t \tag{2.3}$$

where y is either the evoked spike probability or inhibition duration, t is time and a, b, and c are the fitted variables. Changes were denoted to occur if the analysis of variance (ANOVA) F-statistic resulted in $p<0.05$. The sign of the linear fit slope (a in Equation 2.2) or the sign of the exponent base (b in Equation 2.3) of the exponential fit determined whether the changes were classified as increasing or decreasing. Neuron-stimulation site pairs with less than 3% average probability of evoking spikes were disregarded for analyses over time due to their inconsistency. The changes over time in spikes and inhibition were designated to be correlated if their correlation had a p-value less than 0.05.

2.3.7 Statistical analysis

Student's t-test was used for statistical comparisons between groups when the data were determined to be normally distributed. Two-sided Wilcoxon signed-rank test (*signrank*, MATLAB) was used to confirm normality and for nonparametric data. Paired t-test or two-sided Wilcoxon rank-sum (*ranksum*, MATLAB) tests were used where appropriate. Fisher's exact test (*fishertest*, MATLAB) or two-way analysis of variance (ANOVA) (*anova2*, MATLAB) was used to compare categorical data. The Pearson correlation was used for all correlation tests. P-values for significance and tests used are reported in individual analyses.

2.4 Results

2.4.1 Evoked spike and inhibitory response

We found that ICMS elicited a brief excitatory response followed by a longer inhibition. Electrodes on the Utah array typically showed evoked spikes occurring 1.5 to 10 ms after single-pulse stimulation. The inhibition typically followed the excitatory response and was observed as the suppression of firing in the PSTH for 5 - 100 ms, although in some instances it lasted up to 200 ms. We also observed that stimulus-evoked inhibition could occur in the absence of the excitatory response.

To ensure that stimulation arriving during inhibition was not affecting the stimulus response, we delivered trains of two or three pulses with each subsequent pulse timed to occur during the inhibitory response of the previous stimulus. Single, double, and triple pulse stimuli were delivered to a single channel, randomly interleaved at 2 Hz. Our results across 5 different sessions show that stimuli delivered during the inhibitory response were able to reliably evoke spikes comparable to when stimuli was delivered at other times, as previously reported [26] (Figure 2.3, left). Furthermore, each stimulus pulse "reset" the inhibitory response such that the duration of inhibition was the same following each pulse train (Figure 2.3, right).

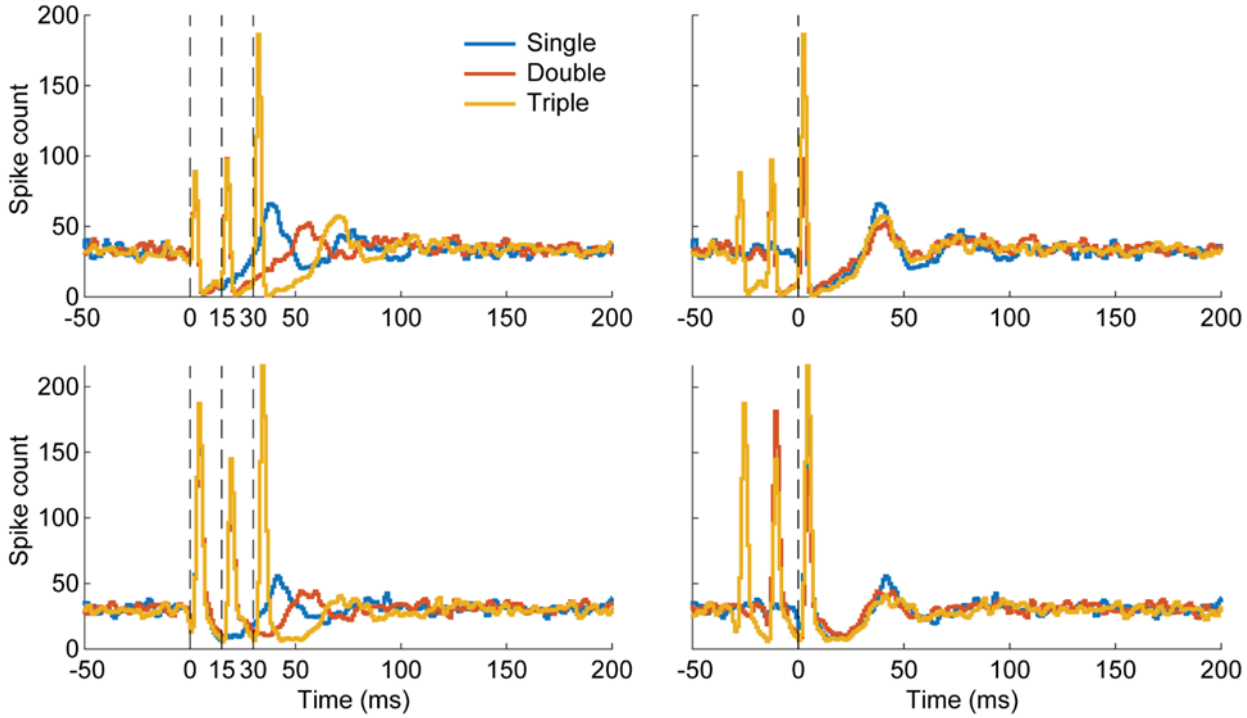

Figure 2.3 Stimulation During Inhibition

Two PSTH examples with 1 ms bins of double- and triple-pulse stimulation in which subsequent pulses arrive during the inhibitory response of the previous pulse (left). Spikes were readily evoked even when stimulating during the inhibitory response. Aligning the PSTHs to the final stimulus pulse (right) shows that the inhibition restarts at each stimulus pulse. Each condition consisted of 1500 stimuli.

2.4.2 Effects of distance from stimulation site and stimulus amplitude

Evoked spikes occurred with greater probability and less variable latencies for units close to the stimulus electrode than for more distant units. The probability of evoking a spike in electrodes <1 mm from the stimulus site was significantly greater than for further sites (p=1.3e-14). In addition, closer sites on average had evoked spikes that occurred at shorter, and less variable latencies (p=8.5e-29), suggesting the presence of mono- and polysynaptic activation (Figure 2.4A).

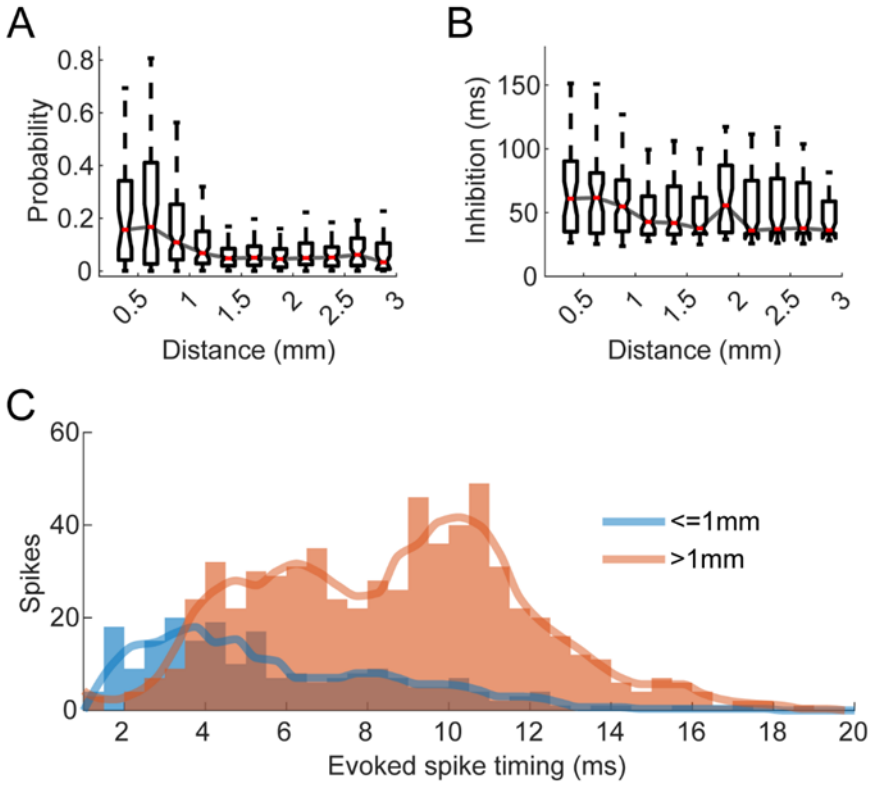

Figure 2.4 Effect of Distance From Stimulus Site on Evoked Spike Probability

A. Evoked spike probability with respect to distance from the stimulated site. **B.** inhibition duration with respect to distance from the stimulated site. **C.** Histogram of evoked spike timings split into sites close (<1mm) to the stimulated site and all other sites. The line shows the cubic interpolated moving average over 3 bins.

The duration of inhibition had a similar trend: recording sites <1 mm from the stimulus site tended to show stronger inhibition compared to further sites (p=0.003) (Figure 2.4B). The probability of evoking spikes and the duration of inhibition increased sigmoidally with the stimulus amplitude for all responsive units (Figure 2.5). The sigmoid curves were always steep: a change of 10 to 20 μA in stimulus intensity generated the difference between 5% and 95% of the maximum value for both evoked spike probability and inhibition duration.

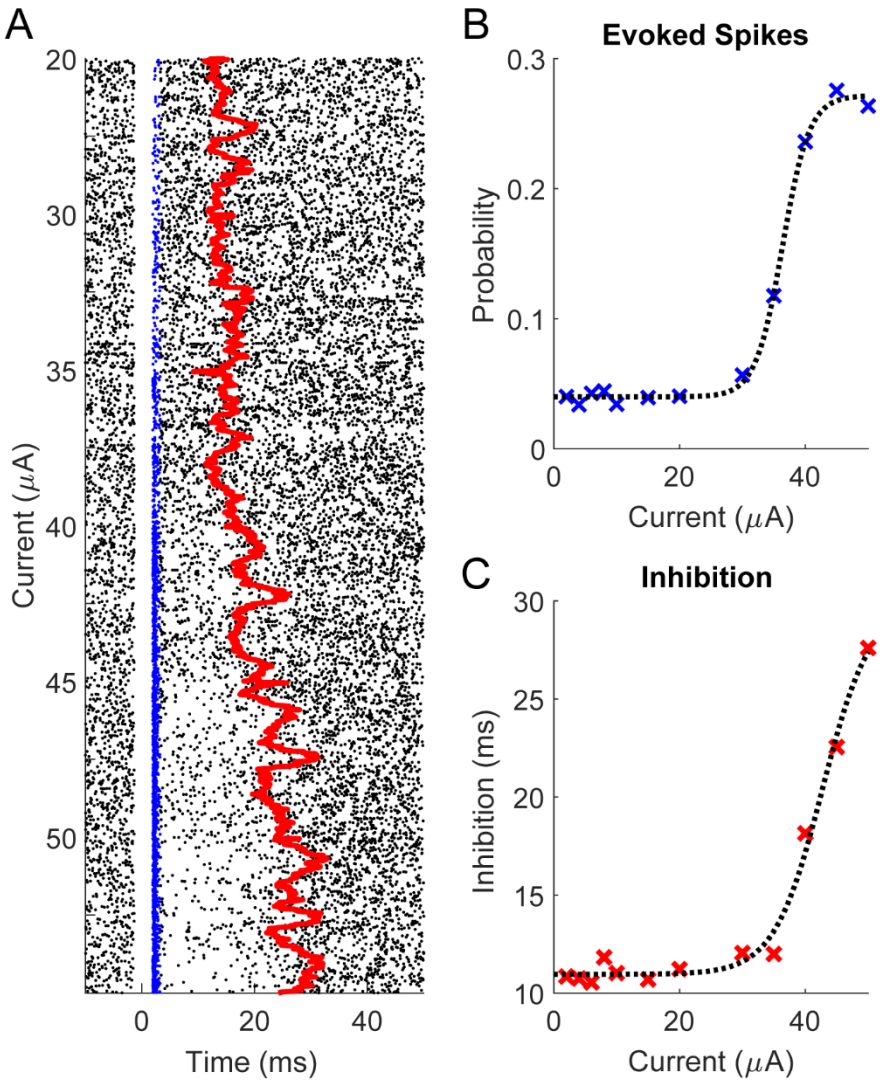

Figure 2.5 Effect of Stimulus Amplitude on Evoked Spike Probability and Inhibition Duration

A. An example raster plot of a spike over different stimulus current amplitudes delivered for five minutes each at 2 Hz. Blue dots represent evoked spikes and the red line shows the median of the inhibition duration binned every 30 seconds with 1 second steps. **B.** An example evoked spike probability as a function of amplitude and **c)** inhibition duration as a function of stimulus amplitude. The dashed lines are fitted sigmoidal curves using least-squares regression.

2.4.3 Spontaneous activity affects stimulus response

In addition to stimulation current and separation of the recording and stimulation sites, we found two other dependencies that affect the response. One is the spontaneous firing rate of the recorded units. Both the evoked spike probability and inhibition duration had statistically significant correlations with spontaneous firing rate over time. Table 2.1 documents the number of units with uncorrelated, positively correlated, or negatively correlated evoked spike probability

Table 2.1 Evoked Spike Probability and Inhibition Duration of a Neuron Are Correlated with its Spontaneous Firing Rate

		Inhibition duration and spontaneous firing rate			
		Uncorrelated	Positive Correlation	Negative Correlation	Total
Evoked spike probability and spontaneous firing rate	Uncorrelated	10	18	37	65
	Positive Correlation	86	77	300	463
	Negative Correlation	12	13	32	57
	Total	108	108	369	585

and inhibition duration with spontaneous firing rate. A slight majority of units tested (300/585, 51%) had evoked spike probabilities that were positively correlated with firing rate (Figure 2.6A, top) and inhibition duration that were negatively correlated with firing rate (Figure 2.6, bottom). We additionally performed a 2-way ANOVA to determine whether the two relationships were dependent on one another but found no significant relationship ($p=0.11$). Units with positively correlated evoked spike probabilities and spontaneous firing rate were typically farther from the stimulated site than units with a negatively correlated relationship (Figure 2.6B, left). In contrast, units with positively correlated inhibition duration and spontaneous firing rate were typically closer to the stimulated site (Figure 2.6, right).

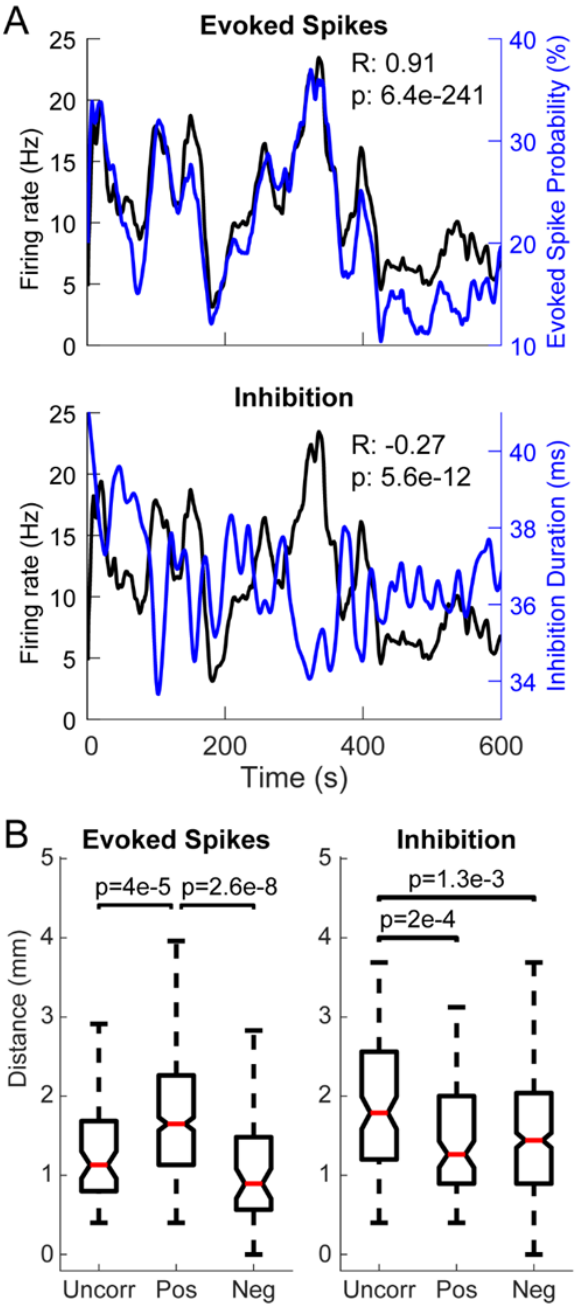

Figure 2.6 Probability of Evoking a Spike and Inhibition Duration are Related to Spontaneous Firing Rate

A. An example of a neuron with positively correlated firing rate (black) and evoked spike probability (blue) over 10 minutes. The rate and probabilities are averaged over 30 second bins with 1 second steps. **B.** The same unit had a negatively correlated firing rate (black) and inhibition duration (blue). The inhibition duration is also averaged over 30 second bins with 1 second steps. **C.** Distance from the stimulated site for units with uncorrelated, positively correlated, and negatively correlated evoked spike probabilities (left) or inhibition duration (right) with firing rate.

The second dependency was the timing of stimuli relative to the most recent spike. We analyzed 17 units (2 from monkey S, 3 from monkey K, and 12 from monkey J). In 14 of the 17 units (2 from monkey S, 3 from monkey K, and 9 from monkey J), the probability of evoking a spike varied as a function of the time between the onset of stimulation and the most recent spontaneous spike. For these 14 units the probability was significantly positively correlated with the unit's autocorrelogram in the absence of stimulation (Figure 2.7, top and middle). The probability distributions were the same even if the most recent spike was evoked by stimulation rather than being spontaneous. The inhibitory response did not depend on the timing of pre-stimulus spikes.

In 3 of the 17 tested units the probability of evoking a spike did not depend on the timing between the previous spike and stimulus onset. In all 3 of these units the probability of evoking a spike was consistently high (Figure 2.7, bottom), which suggests that the stimulation amplitude was large enough to evoke spikes regardless of other properties.

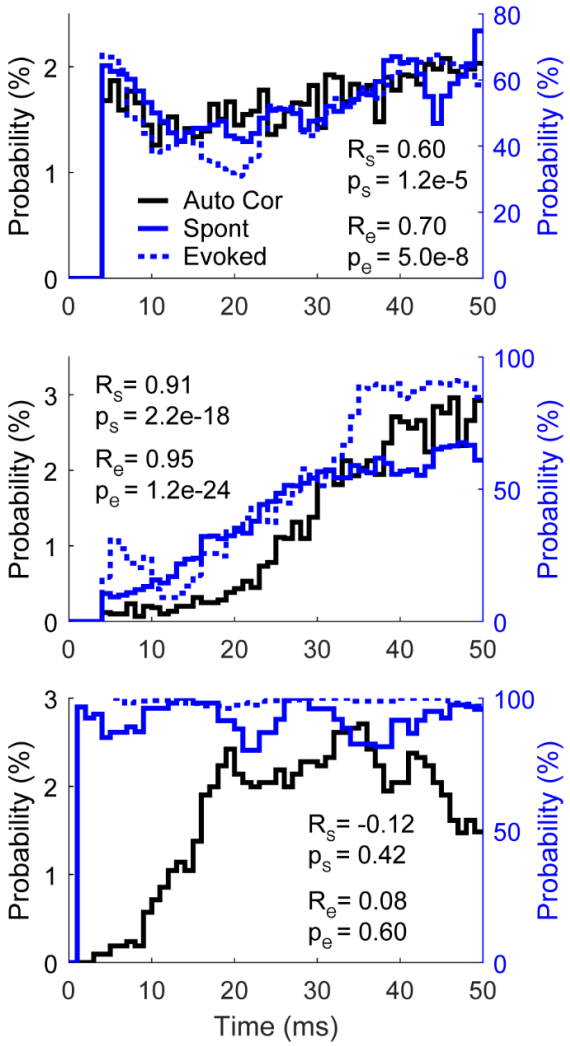

Figure 2.7 Probability of Evoking a Spike is Dependent on the Timing of Stimulus

Three examples of spike autocorrelations (Auto Cor), and probability of a stimulus evoking a spike relative to timing from the most recent spontaneous (Spont) and evoked (Evoked) spike. Top and middle show two different autocorrelation waveforms with correlated evoked spike probability. They are not aligned in the bottom example; this typically occurs when the probability of evoking a spike is high. All traces show a moving average using 5 ms bins with a 1 ms step size. R_s is the correlation coefficient between Auto Cor and Spont, R_e the correlation coefficient between Auto Cor and Evoked, and p_s and p_e are the corresponding p-values of correlation.

2.4.4 Repetitive stimulation changes stimulus response over time

Repetitive microsimulation has been shown to modulate the responses of units to ICMS [33], [34], [102]. We documented the effects of repetitive ICMS over the stimulus period ranging from 5 to 20 minutes on the evoked spike probability and the duration of inhibition by delivering stimuli at between 1 and 20 Hz. The probability of evoking a spike often increased or decreased over time, following a linear or exponential trend over the course of the session to an asymptote (Figure 2.8A).

The evoked spike probability was significantly more likely to change with higher frequencies of stimulation, with a clear trend over successive 5 Hz bins (Figure 2.8B). Of the changes, high frequencies (15-20 Hz) were likely to cause more decreases in evoke spike probability compared to lower frequencies. Changes in inhibition duration over time were also more likely to occur at higher frequencies, but there was no clear frequency dependence in the direction of change. (Figure 2.8B)

We also observed changes depending on the distance from the recording site to the stimulated site (Figure 2.8C). The probability of evoking spikes was significantly more likely to change the closer the units were to the stimulated site. Sites further away tended to be more likely to have decreased probability over time, but this was not statistically significant ($p>0.05$ for all pairwise comparisons). Similarly, inhibition duration also had a higher likelihood of changing over time in sites closer to the stimulated site. The changes tended to be more likely to decrease further away, but this was also not statistically significant ($p>0.05$ for all pairwise comparisons).

Finally, in 72/585 (12%) of the tested units, we also observed changes over time in the latency of the evoked spikes (Figure 2.9A). Latency changes typically occurred in units that were recorded on electrodes closer to the stimulation site (Figure 2.9B), and more commonly occurred with high frequency stimulation (Figure 2.9C). Of the units with changes in their evoked spike latency, 51/72 (71%) units had increasing latencies and 21/72 (29%) had decreasing latencies.

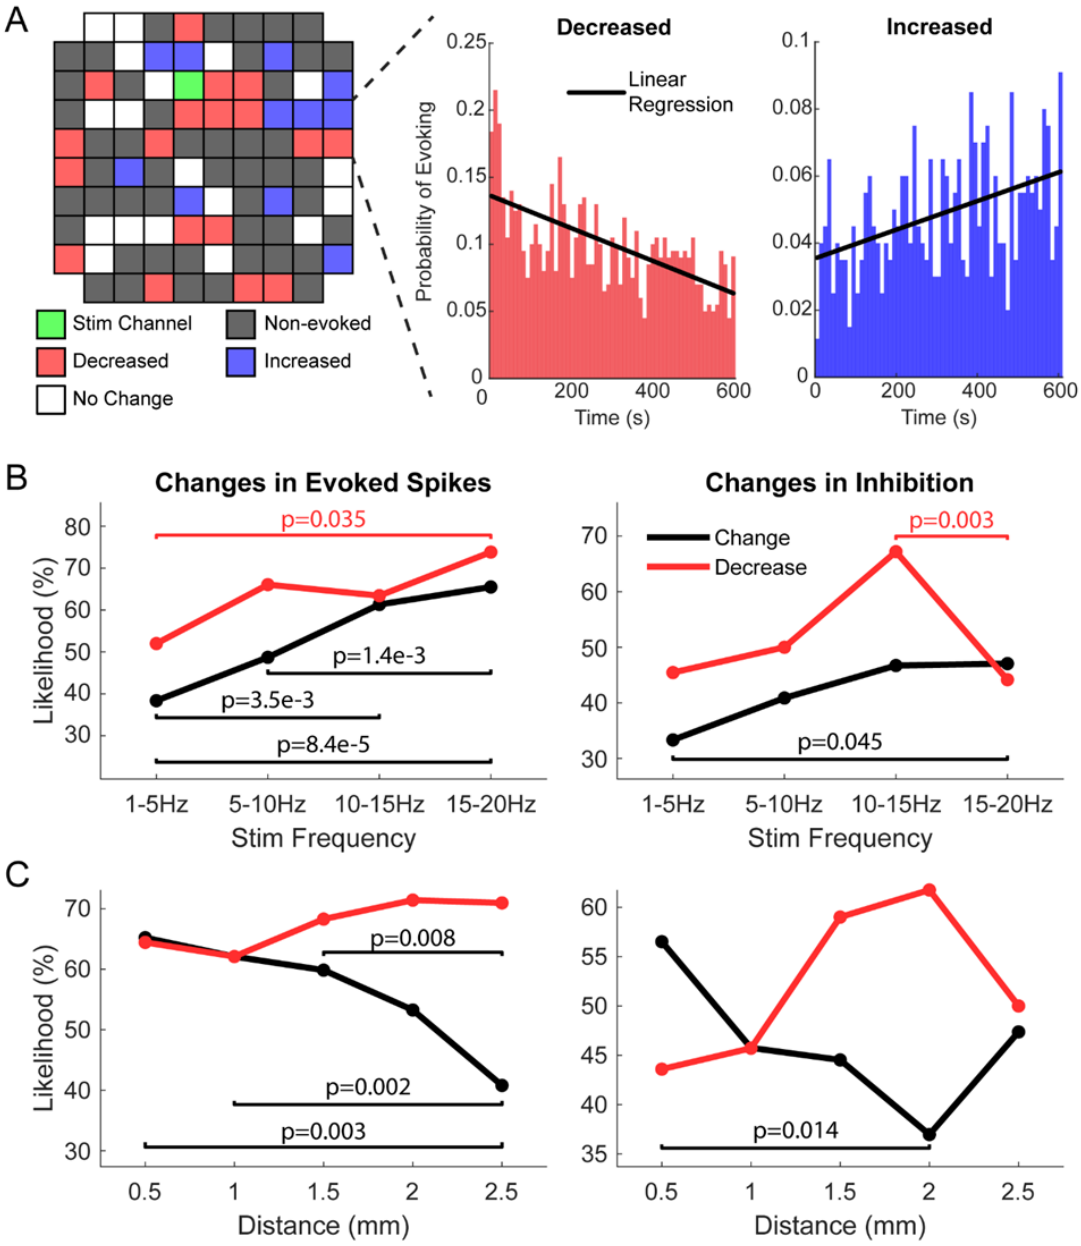

Figure 2.8 Repetitive Stimulation Changes Evoked Spike Probability and Inhibition Duration

A. An example of changes in the likelihood of evoking spikes at different sites across the array with 10 Hz repetitive stimulation at one site. Insets show examples of a decrease and increase of evoked spike probability over time. **B.** The likelihood of evoked spike probability (left) and inhibition duration (right) changing over time (Change), and the likelihood that the changes are decreasing (Decrease) dependent on stimulation frequency. Labeled p-values show all statistically significant differences. Evoked spike probability is more likely to change over time with higher stimulation frequencies, and the changes are more likely to be decreasing. Inhibition is also more likely to change with higher stimulation frequency. **C.** Likelihoods of changes dependent on the distance of the recorded site from the stimulated site. Labeled p-values show all statistically significant differences. It is less likely for the probability of evoked spikes at sites closer to the stimulated site to change over time.

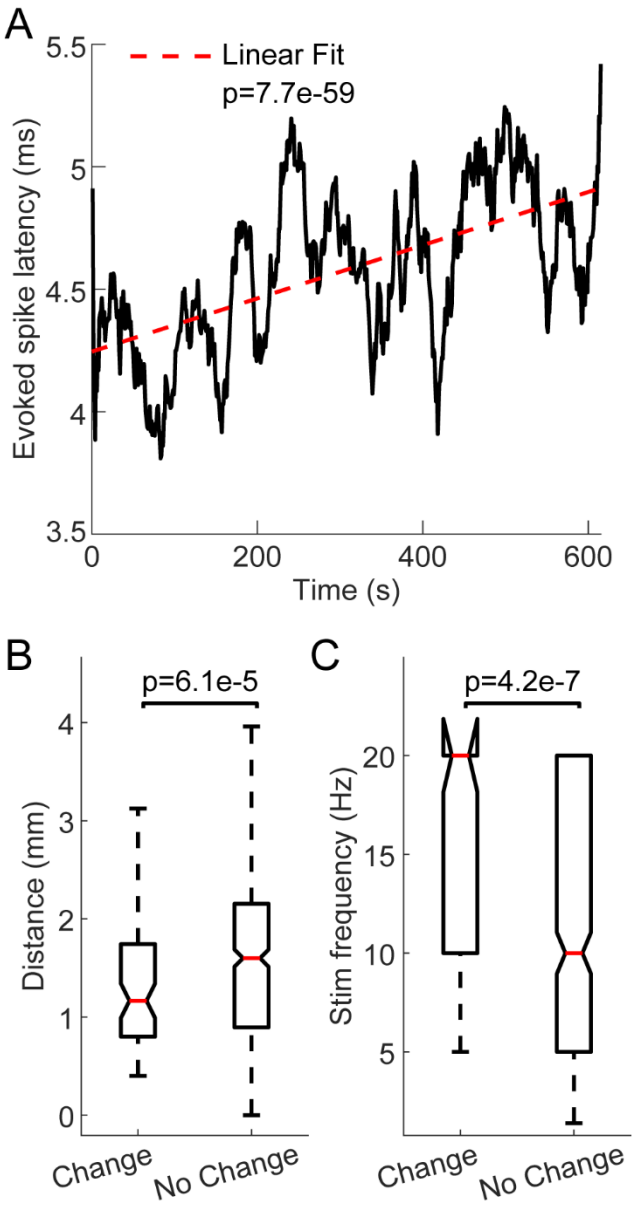

Figure 2.9 Changes in Evoked Spike Latency with Repetitive Stimulation

A. An example of evoked spike latency changing over time. Evoked spikes that had changes in their latency over time were **(B)** typically at sites closer to the stimulated site and **(C)** occurred when delivering stimulation at a higher frequency.

2.4.5 Relationship between evoked spikes and inhibition

Since evoked spikes and inhibition often both exhibited changes over time, we sought to determine whether they were directly related on a trial-by-trial basis. Individual stimuli in each experiment were divided into

two categories: those that evoked spikes and those that did not. For the unit in Figure 2.10A the PSTHs of the two classifications of stimuli are very similar except for the evoked spike peak. We found no statistical pairwise difference between the inhibition induced by stimuli that evoked spikes compared to the stimuli that did not evoke spikes (585 units, Figure 2.10B).

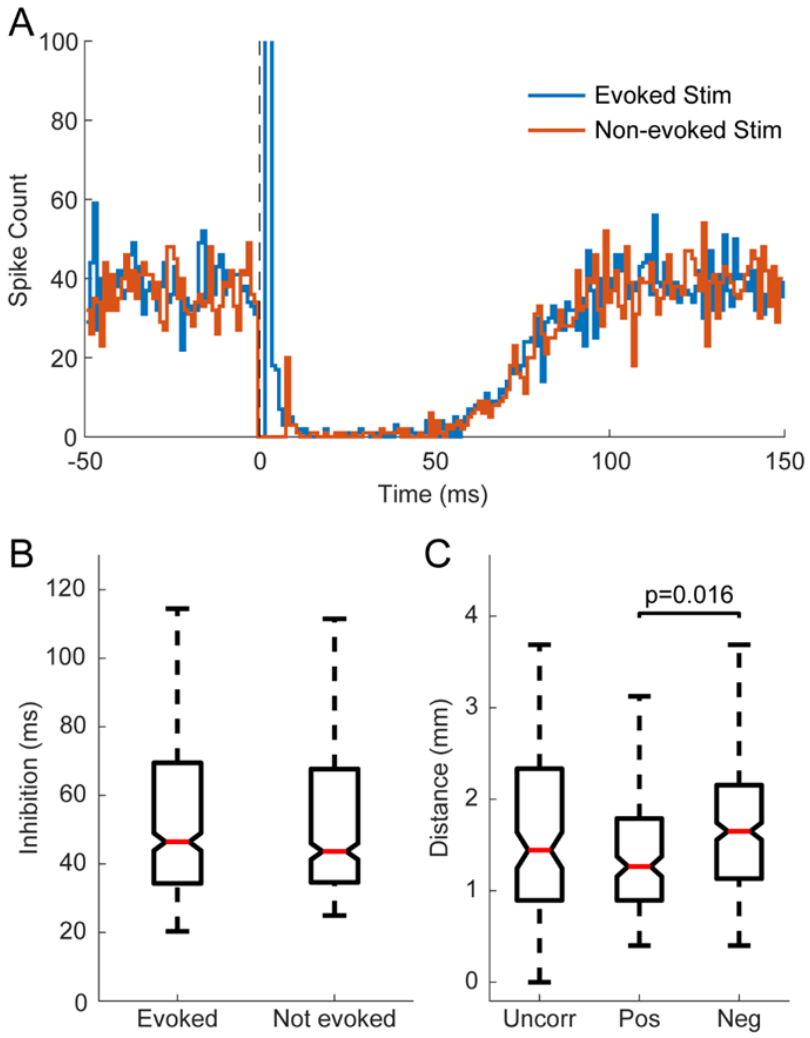

Figure 2.10 Relationship Between Evoked Spikes and Inhibition

A. Example PSTH with 1 ms bins following stimuli that evoked spikes and those that did not demonstrating similar inhibitory response. Each condition consisted of 1500 stimuli. Smoothed (2 ms wide gaussian moving window) PSTH of the two different stimulation classifications. Note the inhibition is extremely similar for both. **B.** A comparison of the inhibition strength in the two different classifications for 470 units. There was no statistically significant pairwise difference between the two groups. **C.** Comparisons of distance from the stimulus site of units with uncorrelated, positively correlated, and negatively correlated evoked spike probability and inhibition duration.

Although evoked spikes and inhibition seemed to be independently activated on a trial-by-trial basis, the mechanisms underlying the changes over time could nevertheless be related. We analyzed the correlation between the probability of evoking a spike and duration of inhibition over time to determine whether they were positively or negatively correlated. Correlations with $p<0.05$ were considered significant, and all other instances were denoted to be uncorrelated. Of the 585 units tested, 27% (161 units) had positively correlated changes in the probability of evoking spikes and the duration of inhibition, 50% (292 units) had negatively correlated changes, and 23% (132 units) were uncorrelated. Units with positively correlated evoked spike probabilities and inhibition duration tended to be closer to the stimulated site (Figure 2.10C).

2.4.6 Spike type does not correlate with stimulus response properties

To determine whether the cell type of the recorded units influenced their responses to ICMS, we used both the spike waveform and firing rate to classify each neuron as fast-spiking (FS) or regular-spiking (RS). Although previous studies have used the spike width and the inter-spike interval to differentiate between the two [103], [104], we used a proxy measure of the area under the action potential curve divided by the trough magnitude, which resulted in a more continuous distribution.

Figure 2.11A and Figure 2.11B show an example of the two different spike waveforms, as well as the distribution of their action potential areas and firing rates, respectively. We separated the units into two groups based on their defined spike width shown at the dotted line in Figure 2.11B. Around 10% of all units fell to the left of the line and were denoted to be FS; the rest were denoted to be RS. We found that the putative cell type did not correlate with the distribution of distances from the stimulated site, whether a spike could be evoked, the probability of evoking a spike, the duration of inhibition, or how any measure changed over time (Figure 2.11C-G). Thus, all results reported herein are independent of the type of neuron recorded.

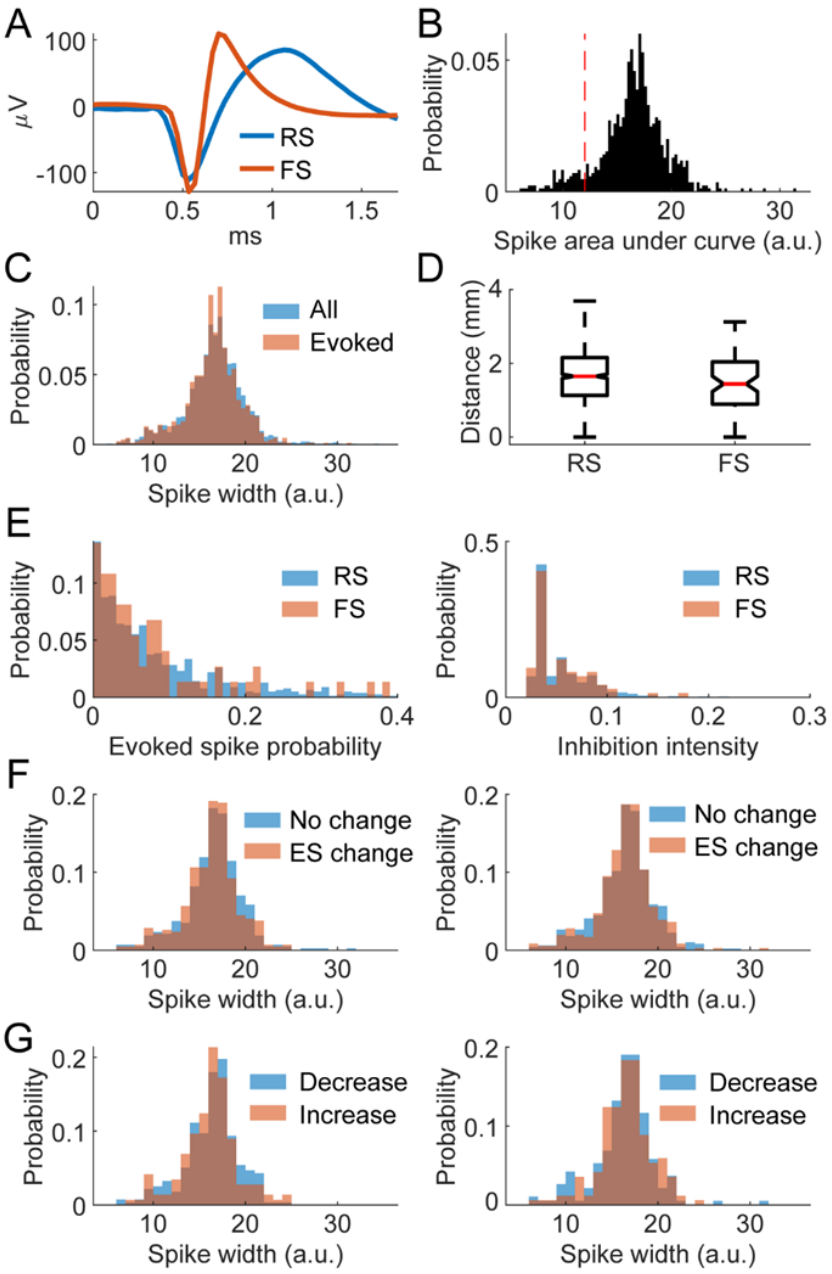

Figure 2.11 Comparisons Between Regular Spiking and Fast Spiking Neurons

A. Example of regular spiking (RS) and fast spiking (FS) waveforms. **B.** Distribution of measured area under the curve of all spikes. Vertical dotted line indicates the classification boundary – fast spiking falls to the left and regular spiking to the right. **C.** Spike width distribution of all recorded spikes and spikes that were evoked by stimuli. **D.** Distance from the stimulated site of evoked spikes grouped by spike width ($p=0.03$). **E.** Evoked spike probability distribution (left) and inhibition duration distribution (right) of RS and FS neurons. **F.** Spike width distributions of evoked spike probability change vs no change over time (left) and inhibition duration change vs no change over time (right). **G.** Spike width distributions of evoked spike probability decrease vs increase over time (left) and inhibition duration decrease vs increase over time (right).

2.5 Discussion

2.5.1 Comparisons to previous studies

ICMS has been shown to predominantly activate neurons transsynaptically [26], [105], [106]. We observed in our study that spike latencies fluctuated more than would occur with antidromic activation, and stimuli that were delivered within 1 millisecond after a spike were still able to evoke spikes. Since antidromic activation would result in collision and an absence of a spike at such short latencies, this suggests that the recorded units were predominantly activated orthodromically. Also consistent with previous experiments, we saw evoked spikes in units recorded up to 4.5 mm away from the stimulation site, which suggests that ICMS activates a distributed population rather than only a concentrated sphere of neurons around the electrode tip [26], [27], [32].

Although the excitatory response was directly measurable in our experiments, the subsequent inhibitory response manifested as a lack of spikes. Butovas et al. concluded that a similar inhibitory response to ICMS was likely caused by $GABA_B$ receptors, which they confirmed with a follow-up study with pharmacological blocks [26], [31]. GABAergic inhibition would also explain the similarity in the sigmoidal curves of evoked spikes and inhibition with stimulus amplitude in our experiments – if more excitatory neurons are excited by the higher intensity stimulation, more inhibitory neurons will be activated via feedforward and feedback circuits, thereby increasing the amount of inhibition (Figure 2.12). A lack of relationship between inhibition and distance to stimulation site as well as the presence of inhibition in units without evoked spikes is likely due to the high connectivity of interneurons compared to principal cells [107]–[110]. Interestingly, Hao et al. found that inhibition decreased as a function of distance, which we did not find [32]. Potential reasons for this discrepancy could be the differences in how we measured inhibition and the stimulus amplitude.

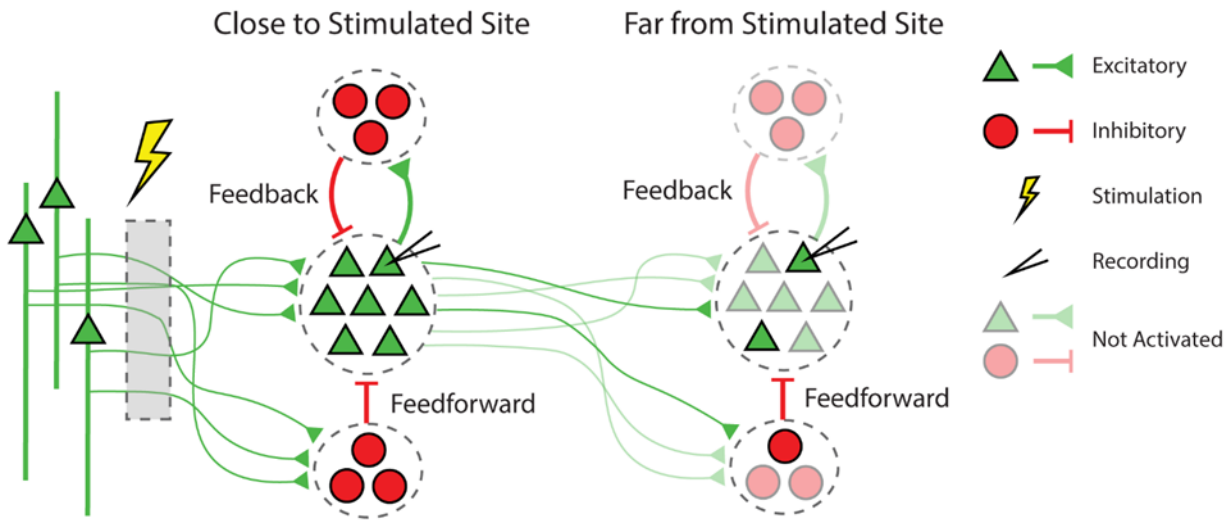

Figure 2.12 Stimulation Response Circuitry

Model of ICMS activated circuitry that generates excitatory and inhibitory responses through feedforward and feedback mechanisms. Stimulation activates axons projecting to the recording site. Sites closer to the stimulated site have more complete activation compared to sites further from the stimulated site. Possible direct connections from the stimulated site to the far site would also be sparsely activated.

However, we found that inhibition typically lasted between 5-100 ms and rarely over 100 ms, which is significantly shorter than the average time constants of $GABA_B$ inhibitory postsynaptic potentials of 150-200 ms [111], [112]. Therefore, $GABA_A$ mediated inhibition is a more likely candidate to explain our results; previous studies have shown that $GABA_A$ is involved in recurrent polysynaptic inhibition, which we are likely activating via ICMS [113], [114]. The different animal models and recorded cortical region in these studies may account for these discrepancies.

ICMS activates long horizontal fibers to both feedforward inhibitory networks and the recorded excitatory neurons (Figure 2.12). These afferent fibers have stronger excitatory connections to the inhibitory interneurons than the principal cells, particularly in layer 2/3, which may explain why we sometimes observe an inhibitory response without any excitatory response [115]–[117]. This, coupled with the fact that inhibitory neurons often target somatic or perisomatic compartments [107]–[109], suggests that the observed inhibition may be initially activated via feedforward circuitry, and is subsequently followed by

the excitatory response. The excited principal cells may then activate feedback circuitry that contributes to the inhibitory response (Figure 2.12).

2.5.2 Stimulus response depends on network activity and intrinsic membrane properties

Previous research has shown that stimulus responses depend on network activity both *in vitro* [118], [119] and *in vivo* [120]. Our findings were consistent with these results; the probability of evoking a spike was often positively correlated with spontaneous firing rate, whereas the duration of inhibition was often negatively correlated with firing rate.

English et al. demonstrated that the transmission probability for postsynaptic spikes of inhibitory neurons in the hippocampus *in vivo* is a function of the timing between the previous postsynaptic spike and presynaptic spike [121]. Moreover, that study found that this dependency was independent of whether the previous postsynaptic spike was spontaneous or evoked, which suggested that intrinsic properties of the postsynaptic membrane were responsible for the dependence. Our results extend these outcomes by demonstrating that the timing of the previous spike affects not only the transmission probability for spontaneous presynaptic spikes, but also the stimulus-evoked spikes in cortical neurons. Furthermore, we found that the probability distributions for evoking a spike as a function of the timing between the previous spike time and stimulation onset and the spike train autocorrelation were often significantly positively correlated, which further reinforces the notion that this dependency reflects the intrinsic properties of the recorded units.

Altogether, our results reveal that the stimulus response has at least two dependencies other than stimulus amplitude and distance from stimulus site: the intrinsic membrane properties of the recorded neurons and the activity of the network. However, a sufficiently large stimulus current may saturate the responses and overcome these dependencies.

2.5.3 Repetitive stimulation modulates stimulus responses

Michelson et al. showed that the number of neurons activated by electrical stimulation diminished over time with higher frequencies measured with calcium imaging in a 407 μm x 407 μm window [33]. Their study showed that changes began at regions distal from the stimulated site when delivering stimuli with frequencies greater than 10 Hz. These results are consistent with our study across the larger spatial field (4x4 mm) of the Utah array, as evoked spikes were more likely to be diminished with higher frequency of stimulation. The changes we observed occurred over much longer timescales and continuously for up to 20 minutes. However, we also observed that in some units the probability of evoking a spike increased over time even at longer distances and higher frequencies, which cannot be fully explained by a diminishing region of activation.

Although we did not explicitly measure the duration of changes induced by repetitive stimulation, we observed that they typically lasted less than 2 minutes. Due to the short-lived nature of the induced changes, various mechanisms of short-term synaptic plasticity such as vesicle depletion and facilitation by calcium influx [34], [122] may best explain our results. Evidence suggests that different forms of short-term plasticity exist for synaptic connections between different cell-types [123], [124]. Beyond these differences, previous *in vitro* work by Markram showed that synaptic connections between pyramidal neurons of the same morphological class and interneurons had similar facilitating and depressing characteristics, but with different time courses [125].

Together, these results may explain why the frequency-dependent changes that we measured were different for each unit. Although we did not discern any differences between regular and fast spiking neurons for any measure, there are limitations in such cell type classifications with extracellular recordings. Furthermore, we also observed changes in the latency of evoked spikes due to repetitive stimulation, which has previously been shown to occur in the presence of short-term plasticity [37]. Future studies with specific differentiation between cell and synapse types may shed more light on whether cell-type specific differences account for the variability across spikes.

2.5.4 Excitation and inhibition are independently activated but modulated together within an interconnected network

The balance between excitation and inhibition within the cortex is a much-studied topic and is highly relevant to neural computation. Though the examined network size, location, and synaptic connections vary greatly, the strong consensus is that excitation and inhibition are generally comodulated [107], [126]–[130]. We found that the probability of evoking a spike and the duration of inhibition were frequently positively or negatively correlated over time, though they were independent on a trial-by-trial basis. Units close to the stimulated site typically had positively correlated evoked spike probability and inhibition duration, both of which were negatively correlated with firing rate. Units far from the stimulation channel, however, had positively correlated evoked spike probabilities and firing rates, which were both negatively correlated with inhibition duration.

The effect of distance can be explained by the fact that sites closer to stimulation are more likely to be activated by ICMS [26], [32]. Due to the feedforward and feedback inhibitory circuitry, if the total excitation increased or decreased due to short-term plasticity, the inhibition should change in a positively correlated manner. Sites far from stimulation, however, are not activated as comprehensively and are thus less likely to be susceptible to short-term plasticity. Similarly, the negative correlations in evoked spike probability and inhibition at these far sites are likely due to network dynamics, whereas the positive correlations in closer sites are likely due changes in short-term plasticity caused by direct activation via ICMS.

Author Contributions

Richy Yun and Jonathan Mishler performed the experiments and prepared the manuscript which forms the basis for this chapter, including data analysis and interpretation. Steve Perlmutter, Rajesh Rao, and Eberhard Fetz helped edit the manuscript and guide discussions.

Chapter 3.

MANIPULATION OF NEURAL ACTIVITY BY AN ARTIFICIAL SPIKING NEURAL NETWORK IMPLEMENTED ON A BIDIRECTIONAL BRAIN-COMPUTER INTERFACE IN MACAQUE MONKEYS

[Paper by Jonathan Mishler, Richy Yun, Steve Perlmutter, Rajesh Rao, and Eberhard Fetz]

3.1 Abstract

Objective. The absence of effective therapies for patients with neurological damage necessitates the development of new technologies that facilitate recovery. One promising new approach is to use artificial neural networks (ANN) to supplement deficient neural circuits. Such ANNs can be trained to deliver context-specific intracortical microstimulation (ICMS) that manipulates circuit activity with the intent to restore lost function, and can be described as co-processors to the brain. Here we demonstrate that small artificial spiking neural networks (SNNs) can be bidirectionally interfaced with single neurons in the primate brain to manipulate their activity. *Approach.* The activity of single neurons was recorded from primary motor cortex of two awake non-human primates (NHPs) during rest. The neurons were then interfaced with spiking neural networks comprised of integrate-and-fire (IF) units that were programmed on a closed-loop brain-computer interface. When spikes were detected from the biological neurons, simulated postsynaptic potentials (PSPs) were sent to the artificial IF neurons, which spiked when they reached a set threshold. The spikes of the IF output neurons triggered intracortical microstimulation (ICMS) to evoke spikes in the biological neurons, thereby creating artificial connections that altered their spiking dynamics. *Main results.* Interfacing biological neurons with small networks of IF neurons created artificial connections that altered the spiking dynamics of the biological neurons in predictable ways. We show that these altered dynamics depend on the timing of stimulation, and that they can be modeled from open-loop measurements. We then used the model to compare and explore how PSP shape, inhibition strength, and stimulation delay impacted the closed-loop dynamics free from experimental constraints such as non-stationarities in the brain, experimental run-time, and hardware limitations. *Significance.* Our results represent one of the first demonstrations of a neural co-processor that interfaces artificial units with biological neurons in the brain to shape neural activity *in vivo* in ways predictable from open-loop measures. Our model can be used to probe how different characteristics of an interfaced SNN modulate the closed-loop neural activity and can potentially provide a means of developing training algorithms that derive SNNs designed to shape neuronal activity for neurorehabilitation.

3.2 Introduction

The mammalian neocortex utilizes complex interconnected neural networks to integrate and process multimodal sensory information to perform complex tasks. Damage to these networks can disrupt function, which necessitates the development of novel technologies to rehabilitate the brain. Brain-Computer Interfaces (BCIs) [131] are increasingly being considered for neurorehabilitative applications, including stroke [132]–[134], blindness [135], and neurodegenerative diseases such as Parkinson's [21], [134]. For some of these diseases, one approach is to use stimulation to promote synaptic plasticity to rewire the damaged circuits in the central nervous system [7], [136]; however, the ability of the brain to rewire itself is limited, particularly with increasing damage and time after stroke [4], [137]. Other approaches have utilized bidirectional BCIs, such as those that can both record and stimulate the brain to decode neural activity to control an external device while simultaneously using ICMS to create artificial sensations that help guide movement [19], [138]. However, traditional decoding algorithms such as the Kalman filter are linear, which limits their ability to model the complex movements that are normally produced by the brain [18].

One potential strategy to restore damaged neural circuits in the brain could be to build upon previous approaches by using bidirectional BCIs to artificially reconnect or rewire individual neurons within the brain. Such a BCI could serve as a "co-processor" to the brain [139], [140], in which the activity of a group of neurons is decoded while ICMS simultaneously activates neurons based on the decoded activity to restore network activity and neurological function. For clinical applications, an ideal neural co-processor would have at least three characteristics. First, the co-processor could create artificial connections between neurons to manipulate their activity in a context-, or task-specific manner. Second, in addition to manipulating the neuronal activity, it would be able to co-learn and co-adapt with the brain to account for non-stationarities in the brain. Third, accumulating evidence supports the notion that the brain may utilize temporal coding – that is, information is contained within the timing of action potentials, rather than their rate [44]–[46], [141], however this has not been conclusively demonstrated [142]. Nevertheless, the ability to utilize both

decoding strategies would be a beneficial feature of the co-processor. Furthermore, we have previously shown that the likelihood of evoking a spike in most neurons depends on the relative timing between its previous spike time and stimulation onset [143], which highlights the potential importance of temporal encoding as well. Fourth, it would be able to both isolate authentic neuronal activity from stimulation artifacts, as well as limit the number of signal lost due to such artifacts so that its performance is not compromised. Given these demands, spiking neural networks (SNNs) are excellent candidates for such a co-processor. SNNs can be trained to deliver context-specific stimulation, support both rate and temporal coding [144], and can continuously adjust their synaptic connections over time, which could provide a means for them to co-learn and co-adapt with the brain [60], [62]. An additional clinical benefit of SNNs is that they can also be implemented on devices with low power consumptions [145].

Towards some of these goals, there have been recent efforts to interface small numbers of biological neurons with artificial units. Tafazoli et al. [146] trained deep convolutional neural networks with an evolutionary algorithm called adaptive, closed-loop stimulation (ACLS) that triggered multi-site ICMS to shape the activity of multiunits recorded in primary visual cortex to mimic the responses of those multiunits to visual stimuli. Their system was also able to account for noise as well as drift in the responses of the multiunits to stimulation.

Xu et al. [147] developed a neuroprosthetic system in which a biomimetic cerebellar neural circuit was programmed on a field programmable gate array (FPGA) and trained to substitute the biological circuitry responsible for learning delay eyeblink conditioning. In their system, multiunit activity was recorded from both the Inferior Olive nucleus and the Pontine nucleus in an anesthetized rat. The neurons recorded from the inferior olive nucleus were shown to respond to an unconditioned stimulus of blowing an air-puff into the rat's eye, while the neurons recorded from the Pontine nucleus responded to a conditioned stimulus of an audio tone. The artificial cerebellar circuitry was trained to stimulate the facial nucleus of a second anesthetized rat to trigger a delayed eyeblink following the presentation of the conditioned stimulus in the

first rat. Their experiment demonstrated that a spiking neural network can be used to recreate biological circuitry.

In addition to *in vivo* experiments, other work by Buccelli et al. [148] has demonstrated in cell cultures that spiking neural networks can be used to bridge two lesioned populations of neurons, or replace one of the populations to restore their pre-lesioned dynamics. Altogether, these works show the potential for combining artificial neural networks with biological neurons for the purposes of manipulating their activity and/or repairing damaged neural circuitry.

Here we present a novel implementation of hybrid biological/artificial integrate-and-fire neural networks (IFNNs) in awake NHPs. The experiments used a bidirectional brain-computer interface called the Neurochip3 [149] that recorded single neuron activity from multiple channels and delivered stimulation at multiple sites. We programmed artificial IF neurons on a FPGA that received inputs from the set of extracellularly recorded single neurons. The spikes of the IF neurons triggered ICMS to evoke single action potentials in the recorded neurons, thus artificially interconnecting them and changing their relative spiking dynamics. We built a model of the closed-loop dynamics by building a library of responses for each recorded neuron to open-loop ICMS, and found that the model closely predicted the altered dynamics. We then used the model to explore how various parameters of the IFNNs such as strength of inhibition and stimulation delays impact the closed-loop spike dynamics of the biological neurons. Finally, we investigated the effects of stimulation artifacts on the functioning of the IFNN as a function of both network size and average connection strength.

Our system has four crucial differences with previous *in vivo* studies. First, we did not train the IFNNs to shape the activity of the biological neurons, but rather explored how various features of the IFNNs shaped the closed-loop dynamics. In particular, we analyzed how different circuit properties such as recurrent excitatory connections, hidden layers and stimulation delays impacted the closed-loop dynamics. We also implemented IFNNs that connected groups of neurons via artificial biomimetic circuitry and investigated the extent to which the circuits were capable of producing the desired neuronal firing dynamics. Second,

we did not input an external signal into the IFNN to trigger the ICMS, but instead used the activity from the recorded neurons to provide all the inputs to the artificial IF neurons, thereby creating artificial connections between the biological neurons. Third, each output IF neuron triggered a single stimulation pulse which evoked single spikes, or in a few instances, doublets, in the recorded neurons. We were able to isolate the single evoked spikes from the stimulation artifacts, enabling our closed-loop system to function continuously with minimal interference from the stimulation artifacts. Finally, we used IFNNs that utilized the timing of the spikes relative to the background network activity, which we have shown to influence the activation of the recorded neurons [143], thereby serving as a powerful tool to shape the closed-loop dynamics. Altogether, this work represents a first step towards developing neural co-processors [139], [140] that can manipulate neural activity and deliver context-specific stimulation to the brain, with the potential to restore the function of damaged neural networks within the brain.

3.3 Materials and Methods

3.3.1 Implants and Surgery

Two *Macaca Nemestrina* NHPs were chronically implanted unilaterally in the hand region of primary motor cortex with 96-channel Utah microelectrode arrays (Figure 3.1A) (10x10 grid, 400 μm inter-electrode spacing, 1.5 mm depth, iridium oxide, Blackrock Microsystems). Sterile surgeries were performed with isoflurane anesthesia under aseptic conditions, and continuous monitoring of all vitals. Stereotaxic coordinates were used to guide the implantation. A skin incision was first made to expose the skull. A 1.5 cm square craniotomy was then performed, after which the dura was cut along three sides to expose the cortex. The arrays were then implanted, and the dura was sutured around the implant. Two reference wires were placed between the cortex and the dura, and two were placed between the dura and the skull. The bone flap that was removed from the craniotomy was then reattached to the skull with a titanium strap. A second titanium strap was connected to the skull with skull screws to secure the wire bundle that connected the array to the recording pedestal, which itself was secured to the skull with skull screws. The skin incision was then sutured around the base of the pedestal. Following surgery, the animals received postoperative

courses of antibiotics and analgesics. All procedures conformed to the National Institutes of Health *Guide for the Care and Use of Laboratory Animals* and were approved by the University of Washington Institutional Animal Care and Use Committee.

3.3.2 Electrophysiology

Neural activity was recorded on one of two systems – the Neurochip3 (32 channels, 20 kHz sampling rate) [149], or the Neural Interface Processor (Ripple Neuro, 96 channels, 30 kHz sampling rate). Although closed-loop experiments were exclusively performed on the Neurochip3, the Ripple system was occasionally used to record the spontaneous activity of the units to first determine which units would be used on a given day, since it was capable of recording all 96 channels simultaneously. Single neuron activity was isolated with two-window discrimination; stimuli consisted of single cathodal pulses (biphasic; 200 µs phase widths; 15-25 µA).

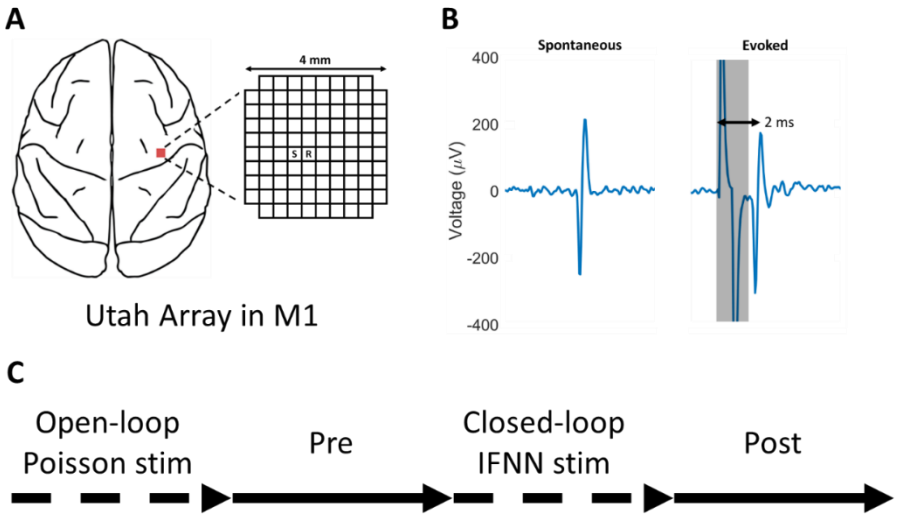

Figure 3.1 Experimental Design

(A) Diagram showing placement of Utah array in primary motor cortex. The stimulation channels (S) were always adjacent to the recording channels (R) to reliably evoke spikes in the biological neurons. (B) Example of a spontaneous and evoked spike. (C) Typical experiment started with open-loop stimulation with Poisson distributions to build a library of responses to ICMS, followed by the recording of spontaneous activity (Pre). The biological neurons were then interfaced with the IF neurons during closed-loop stimulation, which was usually followed by another recording of spontaneous activity (Post) to measure any lasting effects from the stimulation.

3.3.3 Spiking Neural Network Model

We used IF neurons for our spiking neural network model [150] with current-based synapses [50]. The network was programmed on a Cyclone V FPGA on the Neurochip3. The sampling frequency of the *in vivo* neural recordings was 20 kHz, and the IFNN operated at 10 kHz. Each of the input neurons could have an excitatory, inhibitory, or no connection to each of the IF neurons. Due to the small number of combined biological neurons and IF neurons that we were limited to, we did not require each tested network to have separate excitatory and inhibitory neurons (Dale's principle). The postsynaptic response kernel was modelled as the difference between two exponential functions which produced waveforms that resembled PSPs found in the brain:

$$E_S = k \times C_{Slow\ Decay} \times E_S \tag{3.1A}$$

$$E_F = k \times C_{Fast\ Decay} \times E_F \tag{3.1B}$$

$$PSP = E_S - E_F \tag{3.1C}$$

where E_S is a slowly decaying recursively defined exponential, E_F is a quickly decaying recursively defined exponential, $C_{slow\ decay}$ is a constant that determines the rate of decay for E_S, and $C_{fast\ decay}$ is a constant that determines the rate of decay for E_F. We calibrated the values of $C_{slow\ decay}$ and $C_{fast\ decay}$ to achieve the desired rise time and fall time constant, which were typically 1.5 and 4 ms, respectively.

A weight matrix, **W**, stored each of the weights between the input neurons and the IF neurons. The values of both E_S and E_F for any given connection were increased by a factor of k x $w_{i,j}$ every time an input neuron fired, where $w_{i,j}$ is the weight of the connection between input neuron i, and IF neuron, j, and k is a scaling factor that was used to adjust the PSPs such that their amplitudes were equal to their corresponding weights. When the membrane potential reached a predefined threshold, T, the IF neuron spiked, and the membrane potential was reset to 0 (Figure 3.2A, right panel). Although the model was capable of assigning each IF neuron a refractory period, we set it to 0 for each of the experiments since for the small networks that we used the effective refractory period was already greater than 2 ms. Furthermore, each of the IF neurons had the same threshold, as well as rise and fall time constants. Finally, each output neuron had a delay t_{stim}

between when it spiked and when stimulation was delivered. The delays for each IF neuron were typically chosen to be at least 4 ms apart from one another to enable the detection of the spikes evoked by stimulation with minimal interference from stimulation artifacts.

3.3.4 Experimental Design

The NHPs were trained to calmly sit in a primate chair while receiving a periodic smoothie reward to maintain their attention. Our experiments began by selecting the set of input neurons (and potentially output neurons), stimulation sites, IFNN weights, and the delays between when each IF neuron fired and the corresponding electrical stimulus (Figure 3.2). We tested a variety of circuit features to explore their effects on the closed-loop operation, including recurrent connections, hidden layers, and strong reciprocal inhibition. Due to the small number of neurons that we interfaced, we set the excitatory weights to be strong enough to evoke spikes in connected IF neurons if two input spikes in excitatory connections were detected within 1.1-12.6 ms or less, depending on the strength of the connections. Due to the small network sizes, we set the inhibitory connections to be weaker than the excitatory connections to prevent the IF neurons from being inhibited too strongly. The biological neurons were selected based on a few criteria. The first was that the neurons had to be recorded on electrodes as far away from each other as possible to prevent each stimulation site from activating more than one interfaced cortical neuron at a time, although this was frequently unavoidable. If more than one neuron was recorded on a channel, one of the neurons could be selected for the experiment if its waveform was different enough for two-window discrimination to reliably sort them. The stimulation electrodes were always adjacent to a recording electrode that was detecting the spikes of one of the interfaced neurons (Figure 3.1A) to maximize the probability of evoking a spike in the recorded neurons. Altogether, we performed a total of 6 experiments, 5 in monkey J and 1 in monkey K.

After selecting the neurons, stimulation sites, and IFNN weights, we started data collection, which began with 10-15 minutes of open-loop ICMS at 2-7 Hz (Poisson distributed pulses) to document each neuron's responses (open loop or 'OL' epoch) (Figure 3.1C), and 5-10 minutes of recording spontaneous activity ('Pre' epoch). In most of the experiments, stimulation was randomly interleaved to all channels during the

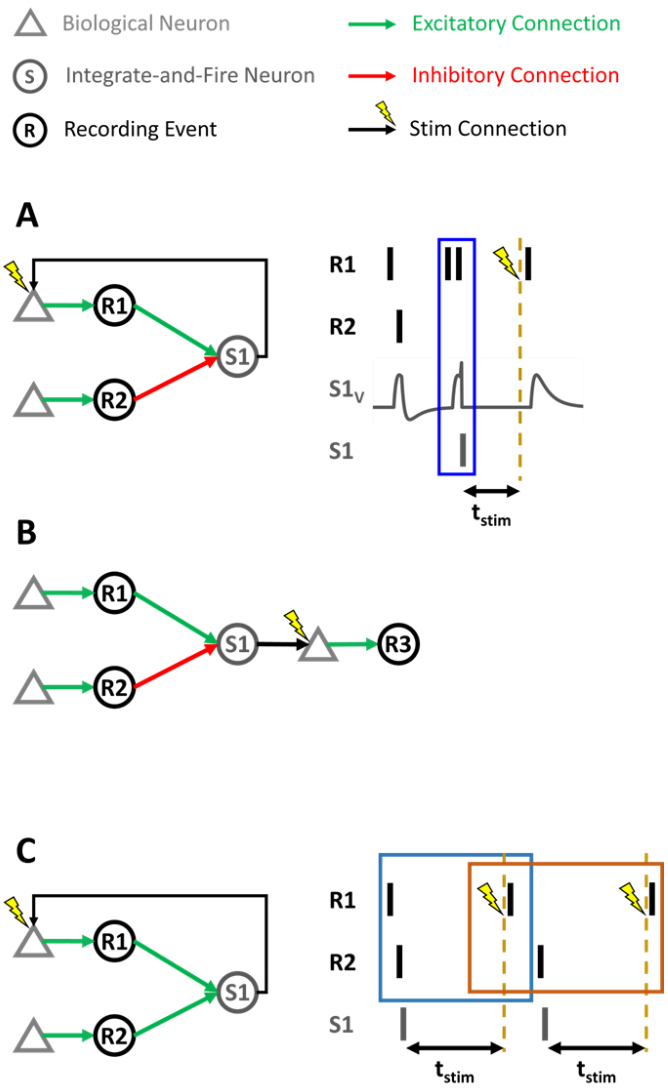

Figure 3.2 Interfacing an IFNN with Single Neurons in Primary Motor Cortex

(A) Spikes are recorded from two single neurons (triangles) producing pulses in recording event units R1 and R2. When R1 or R2 detects a spike from its input neuron, it evokes an EPSP or IPSP in S1, which spikes when its membrane potential reaches a fixed threshold. When S1 spikes, it sends a trigger to stimulate one of the two input neurons. (B) Similar to A, but when S1 spikes, it sends a trigger to stimulate a separate output neuron. (C) Example of a feedback loop generated from the IFNN. Spontaneous spikes can trigger stimulation to evoke spikes. The spikes that trigger the stimulation are grouped together with the evoked spikes to create elementary patterns (blue and orange boxes). Elementary patterns with overlapping spikes are grouped together to isolate the complete feedback loop. The size of the feedback loop shown is 5 spikes.

OL epoch to minimize the length of the epoch. Due to experimental time constraints, we did not include an OL epoch for the experiment in monkey K. We then interfaced the selected neurons with the IFNN for 5-15 minutes to record the closed-loop dynamics ('CL' epoch). In four of the experiments for which time permitted, we collected 5 minutes of data following the closed-loop stimulation ('Post' epoch) to determine the extent to which the closed-loop stimulation changed the natural firing characteristics of the interfaced neurons.

Once the experiments were concluded, the spikes were then re-sorted offline with a two-window discriminator and compared with the online sorted spikes to identify any false positives or negatives. False positives were defined to be spikes that were detected online, but not offline, while false negatives were defined to be spikes that were detected offline, but not online. Noise (mostly from movement artifacts) was removed from both the online and offline sorted spikes. To simulate the CL dynamics free from noise and false positives/negatives, we used the offline sorted spikes to create the models of the CL dynamics. The simulated dynamics were compared to the spike dynamics of the online sorted spikes. Table 3.1 shows the lengths of each experiment and the percentage of false positives and negatives.

Table 3.1 Experimental Times and Statistics

Figure	OL (min)	Pre (min)	CL (min)	Post (min)	Noise (%)	False Negatives (%)	False Positives (%)	Stim Rate (Hz)
4A	0	5	5	5	0.3	R1, R2, R3 3.1, 7.2, 13.0	R1, R2, R3 1.2, 0.4, 5.8	S1, S2, S3 3.2, 3.6, 3.7
4B	10	5	5	5	0.1	R1, R2, R3, R4 0.5, 3.4, 2.9, 9.2	R1, R2, R3, R4 0.5, 2.0, 2.2, 7.7	S1, S2, S3, S4 5.3, 4.9, 5.7, 5.0
4C	15	10	10	0	8.6	R1, R2, R3, R4 5.2, 6.5, 3.2, 10.4	R1, R2, R3, R4 1.0, 4.3, 0.7, 9.1	S1, S2, S3, S4 7.3, 5.6, 3.1, 5.2
4D	10	5	5	5	0.1	R1, R2, R3 0.8, 0.4, 4.6	R1, R2, R3 5.0, 0.8, 0.2	S4, S5 2.3, 2.2
5B (top)*	5 / 5	5	15	5	2.5	R1, R2 0.4, 1.5	R1, R2 0.2, 0.6	S1, S2 9.8, 15.2
5B (bot)*	5 / 5	5	15	5	13.2	R1, R2 0.1, 0.1	R1, R2 1.9, 5.5	S1, S2 20.9, 7.2

3.3.5 Simulating the Closed-Loop Dynamics

To facilitate analysis of the operation of the IFNN, we developed a model to simulate the closed-loop dynamics between a set of single cortical neurons interfaced with IF neurons. The model was updated at the same 10 kHz rate that was used experimentally and consisted of two elements – (1) a Markov model that generated neuronal spike trains with the same statistics as recorded neurons, and (2) a stimulation response model that simulated a response to ICMS. Although this was a simplified model of the neurons and their responses to ICMS, it allowed us to explore how various parameters of the IFNNs such as the weights and stimulation delays influenced the closed-loop dynamics free from experimental confounds such as the non-stationarities in neuronal firing rates and evoked spike probabilities (Supplementary Figure 3.6.1). Furthermore, the model allowed us to overcome hardware constraints such as the limitation of requiring each PSP shape to be the same for each connection.

The Markov model was generated from the neuronal firing data collected during the Pre epoch. The interspike intervals of each recorded neuron were calculated and rounded to the nearest millisecond. For each unique interspike interval, ISI_1, every subsequent interspike interval, ISI_2, was measured and stored in an array (Supplementary Figure 3.6.2A). To generate a spike train, a random ISI_1 was selected from the

data to be a seed. To generate the next spike, one of the possible ISI_2's was randomly selected from the array (Figure 3.3A). This process was repeated until a spike train of an arbitrary length was created.

The second component of the simulation was the stimulation response model, which simulated each neuron's response to ICMS. The stimulation response model (Figure 3.3B) involved two elements: evoked spikes and inhibition. The first step in generating the evoked spike response to ICMS was to isolate the evoked spikes from the spontaneous spikes in the OL epoch. To do so, we plotted the peristimulus time histograms of all the spike times for each neuron for each stimulation channel. We then found all of the peaks after stimulation that were followed by a period of inhibition and counted each spike in that peak as an evoked spike (Figure 3.3C). Although evoked doublets were infrequent, we counted all spikes that occurred within 3 ms of an evoked spike as part of the evoked spike response. As shown in [143], the probability of evoking a spike was dependent on the time interval between the last spike and electrical stimulus (Figure 3.3E). To determine the probability that any given stimulus in the simulation would evoke a spike, the time interval Δt_{prev} between the preceding spike (spontaneous or evoked) and the stimulus was calculated (Figure 3.3B). The probability that a spike was evoked at any Δt_{prev} was obtained from the experimental data measured in the open loop (OL) epoch. The probability distributions were smoothed using a 10 ms window moving average (Figure 3.3E). Due to the typically high rate of stimulation and high spike rate of the biological neurons during the OL epoch, the probability distribution of evoking spikes as a function of Δt_{prev} had a low sampling density at larger values of Δt_{prev} and could vary greatly. Therefore, in the simulations, if the value of Δt_{prev} was greater than 100 ms, or was undefined, we set the probability to be the average probability of evoking a spike for the unit to eliminate this variability. Finally, if an evoked spike was generated, the timing of the evoked spike following the stimulus was randomly chosen from a distribution of timings Δt_{evoked} obtained from the OL epoch data (Supplementary Figure 3.6.2B).

To model the inhibition, we extracted the collection of first spontaneous spikes (so-called "rebound spikes") that occurred after stimulation from the OL experimental data to create a distribution of rebound spikes. After each stimulus in the simulation, a random $\Delta t_{rebound}$ from the array of timings of rebound spikes was

selected (Figure 3.3D). The distribution for rebound spikes was the same independent of whether a spike was evoked or not [143]. Next, we smoothed the rebound spike distributions by first taking 10000 random samples from each distribution. We then added noise that was uniformly distributed between [-10,10] ms to each sample, and combined them with the original distributions. Any rebound spikes in the modified distributions that occurred within the window of time for evoked spikes were discarded. Finally, because multiple channels were stimulated during the OL epoch, there were two cases in which the average rebound spike occurred after the next stimulus which was corrected for by adding Gaussian noise to the distributions (Figure 3.3F) to increase the average value of $\Delta t_{rebound}$. The neurons whose rebound distributions were modified are noted in their respective figures and represent a total of 2/14 neurons that we modelled throughout all our experiments.

To simulate the CL epoch dynamics, we combined the Markov and stimulation response models. To start the simulation, a random sample, ISI_1, from the Markov model was first selected for each neuron. The spike trains for each neuron were then generated by randomly selecting an ISI_2 from the previously obtained distributions (Supplementary Figure 3.6.2A). When the membrane potential of any given IF neuron reached threshold, a stimulus was generated after the corresponding delay. The probability that an evoked spike occurred was calculated based on the value of Δt_{prev} for that neuron. A random number between [0,1] was then drawn from a uniform distribution. If the random number was less than or equal to the evoked spike probability, an evoked spike was generated, the timing of which was drawn from the distribution of Δt_{evoked} for that neuron (Supplementary Figure 3.6.2B). Independent of whether a spike was evoked or not, a rebound spike was then drawn from the neuron's distribution of $\Delta t_{rebound}$ (Figure 3.3D). To continue the spike train after a rebound spike, a random ISI sample from the Markov model was then drawn to continue the spike trains. This process was repeated for each neuron that was evoked from any single stimulus.

To increase the accuracy of the simulations, we also incorporated stimulation artifacts into the model. Each artifact lasted 1 ms and blanked every spike that occurred within that interval. If another stimulus was set to occur within this 1 ms period, it was delayed until the end of the 1 ms, as in the *in vivo* experiments.

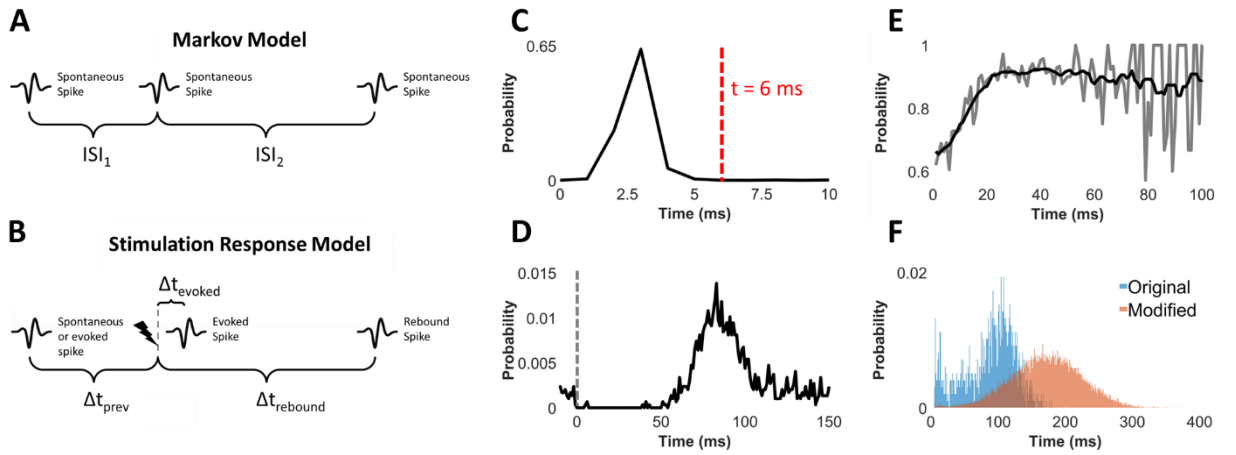

Figure 3.3 Simulating the Closed-Loop Dynamics Between Single Units and an IFNN

(A) Markov model to generate neuronal spike trains. For every interspike interval, ISI_1, all possible subsequent interspike intervals, ISI_2 are stored in an array. To generate a spike train, a seed is randomly selected from the ISI_1's, after which an ISI_2 is randomly chosen from the array. (B) Stimulation response model. The probability of evoking a spike is dependent on the Δt_{prev} between the previous spike time and stimulation onset, and is measured from the data obtained in the OL epoch. Similarly, the rebound spikes are randomly selected from the distribution of responses measured experimentally. (C) Example PSTH of R1 spike times following stimulation (S4 → R1 from Figure 3.4D) collected from the OL epoch. Dashed red line indicates evoked spike cutoff time. (D) PSTH of rebound spikes (S4 → R1 from Figure 3.4D) following stimulation. Dashed line shows stimulation onset. (E) Example probability distribution of evoking a spike (R3 from Figure 3.4A) as a function of Δt_{prev}. Grey curve shows unsmoothed distribution, black curve shows the distribution smoothed with a moving average of 10 ms. (F) Example distribution of $\Delta t_{rebound}$ spike times with and without correction (R2 from Figure 3.4A). Larger time values sampled from a gaussian distribution were added to the $\Delta t_{rebound}$ distribution to correct for the large number of short interstimulus intervals in the OL epoch.

3.3.6 Measuring Feedback Loops

Beyond providing a means of artificially connecting the biological neurons, IFNNs with recurrent connections between the IF neurons and biological neurons were also capable of generating sustained patterns of evoked activity, henceforth referred to as feedback loops, whose characteristics depended on several factors including the size, connectivity, and stimulation delays of the IFNN, as well as the evoked spike probabilities of the biological neurons. These feedback loops were similar in nature to the sustained activity generated by other types of artificial recurrent neural networks [151]. To generate such feedback loops, spontaneous spikes could excite one or more IF neurons, which would then trigger stimulation to evoke spikes in the biological neurons. The evoked spikes could then themselves, or with other spontaneous spikes, excite more IF neurons, which in turn could evoke more spikes. Different combinations of

spontaneous input neurons generated distinct feedback loops via the artificial connections established by the IFNN.

To isolate the feedback loops from the CL neuronal spike data, we first identified the times of each evoked spike. For each evoked spike, we then found the artificial spike times of the respective IF neuron that triggered the corresponding stimulus. We then found each of the biological spikes that caused the artificial spikes in the IF neurons (only biological spikes that sent excitatory PSPs to the IF neurons were included). The biological spikes that excited the IF neurons, and subsequent biological evoked spikes were then grouped together to create "so-called" elementary patterns (Figure 3.2C). Finally, each of the elementary patterns with overlapping biological spikes were grouped together to isolate the complete feedback loops (Figure 3.2C). The size of each feedback loop was defined to be the total number of grouped biological spikes. It was common for elementary patterns to have no overlapping biological spikes with other elementary patterns. Although these could not be strictly described as feedback loops, the extent to which they occurred was directly related to the degree to which the IFNNs could store the memory of previous inputs, and as such were included in all the feedback loop-related analyses.

3.4 Results

3.4.1 Recurrent IFNNs Create Simple Artificial Connections Between Biological Neurons

Rather than train the IFNNs for a specific function, the purpose of these experiments was to explore how various features in the recurrent IFNNs such as their connectivities and stimulation delays influenced the closed-loop dynamics of the neuronal spike times. A greater understanding of these relationships could be used in future studies to design training algorithms for applications such as neuronal control. To this end, we interfaced single neurons with recurrent IFNNs in 4 experiments – 3 in monkey J, and 1 in monkey K. Each of the IFNNs had unique architectures and stimulation delays which allowed us to document their effects on closed-loop function via the artificial connections they created.

We first tested an IFNN with three input neurons and three output IF neurons to understand how the biological-artificial-biological connections established by an IFNN with no hidden layers altered the spike dynamics of the biological neurons. The connectivity was designed so that each of the output IF neurons served as a coincidence detector for the spikes of the input neurons that had excitatory connections to them. An output IF neuron would fire if two of its excitatory input connections fired within a small window of time that depended on the magnitude of their excitatory weights. For an input neuron to prevent an output IF neuron from spiking via an inhibitory connection, it had to fire coincidentally with the two excitatory inputs. Figure 3.4A shows the connectivity of the tested IFNN (left), in which each input neuron had 2 excitatory connections and 1 inhibitory connection to one of the IF neurons. Because each input neuron had excitatory connections to two output IF neurons, if one input neuron spiked in close succession, the two corresponding output IF neurons would fire. However, if two separate input neurons spiked, only one output IF neurons with the shared excitatory connection would fire. The numbers to the right of the network diagram are the delays between the spikes of the three output IF neurons, S1-S3, and the subsequently triggered ICMS. We measured the changes in dynamics with the auto- and cross-correlograms of the neurons, examples of which are shown in Figure 3.4A (right). In contrast to the Pre epoch, the CL epoch shows large and small peaks, which correspond to the evoked spikes generated by the artificial connections between the biological neurons. Due to the small size of the IFNNs, the average size of the generated feedback loop was only 3.4 spikes (Supplementary Figure 3.6.3). The small average feedback loop size indicated that the IFNN created basic artificial connections between the biological neurons with a minimal ability to store memories of past inputs. Supplementary Figure 3.6.4 shows the complete set of auto-, and cross-correlograms of the neuronal spike times for this experiment, as well as the IFNN weights/parameters, and the positions of the recording and stimulation channels on the Utah array.

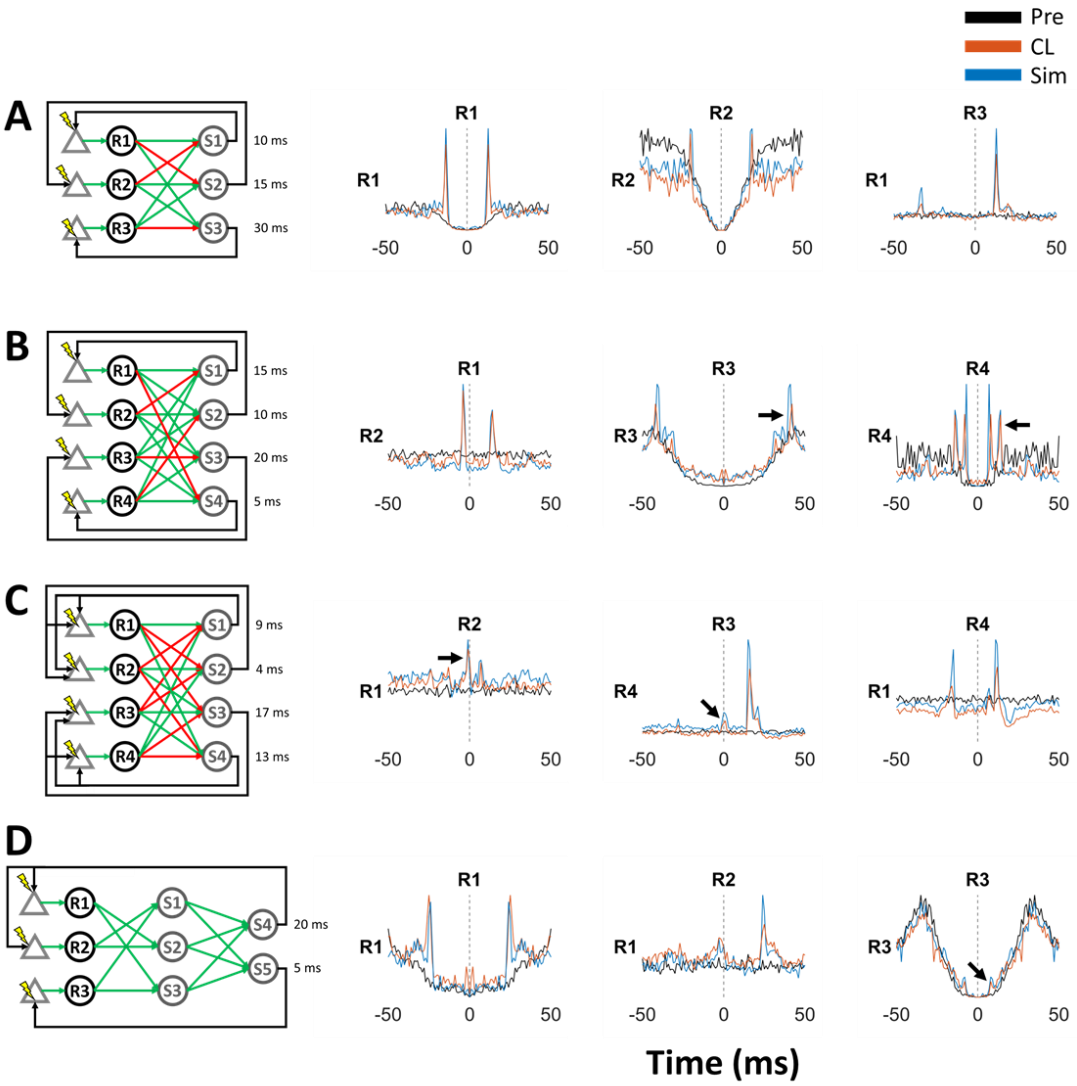

Figure 3.4 Comparing Spike Dynamics Between Pre Epoch, Closed-Loop Epoch, and Simulations for recurrent IFNNs

(Left) Network architectures (green and red arrows show excitatory and inhibitory connections, respectively) of the tested IFNNs. (Right) Comparison of a selection of auto-, and cross-correlograms (1 ms bins) of spike times during Pre epoch, CL epoch, and simulation. 0 ms corresponds to the time at which the biological neurons listed at the top of the plots fire. As shown, the simulation captures the millisecond precision spatiotemporal firing dynamics during the CL epoch, albeit with some deviations. (A) IFNN with 3 input neurons and 3 output IF neurons implemented in monkey K. Peaks in the plots show the artificial connections established by the IFNN. (B) IFNN with 4 input neurons and 4 output IF neurons implemented in monkey J. Beyond creating artificial connections between the neurons, the IFNN created feedback loops in the spike times (black arrows). (C) Same as B, but with a different number of excitatory/inhibitory connections, as well as two evoked spikes per stimulation output. Black arrows show examples of the correlated activity created from stimulation exciting more than one neuron. (D) IFNN with a hidden layer implemented in monkey J. Because S4 evoked spikes in R1 and R2, there is a broad peak in their auto-correlogram. The low probability of evoking a spike in R3 for small values of Δt_{prev} can be seen in the auto-correlogram of R3 (black arrow).

3.4.2 Recurrent IFNNs Create Feedback Loops in Spike Times of Biological Neurons

To understand how IFNNs with more input neurons and output IF neurons affected the closed-loop dynamics, we tested two IFNNs, each of which had four input neurons and four output IF neurons. Figure 3.4B shows the first of these tested networks, in which each input neuron had three excitatory and one inhibitory connection to one of the four output IF neurons. Because each input neuron had excitatory connections to three distinct output IF neurons, if one input neuron spiked twice in close succession, the three corresponding output IF neurons would fire. However, if two separate input neurons spiked coincidentally, only two output IF neurons with shared excitatory connections would fire. In contrast to the IFNN with just three input neurons and three output IF neurons, this IFNN was capable of generating clear feedback loops that were reflected in the auto-, and cross-correlograms, as can be seen in the auto-correlograms of R3 and R4 (black arrows). Examples of how both feedback loops could be generated are shown in Figure 3.5D (see below). Importantly, as with most of the generated feedback loops throughout each of the experiments, these feedback loops occurred when at least two stimuli were triggered in close succession to one another and evoked spikes within a short time interval.

To measure the effects of having different numbers of excitatory and inhibitory connections on the closed-loop operation, we tested a second IFNN with four input neurons and four output IF neurons (Figure 3.4C). For this IFNN, each of the four input neurons had two excitatory connections and two inhibitory connections to one of the four IF neurons. Whereas the coincident firing of any two separate input neurons for the IFNN in Figure 3.4B guaranteed that two output IF neurons fired, any two input neurons firing for this IFNN would drive between 0-1 output IF neurons to fire, depending on the number of shared excitatory connections. If any single input neuron fired twice in close succession, two output IF neurons would fire. However, if a second input neuron with an inhibitory connection to either/both of those output IF neurons fired coincidentally with the two spikes from the input neuron, the outputs were suppressed. An unintended feature of this IFNN was that each stimulation output evoked spikes in two separate input neurons, examples of which can be seen in the cross-correlogram of R1 and R2, as well as R3 and R4 (black arrows). However,

because each of the pairs of input neurons that were simultaneously activated shared no common excitatory connections with the output IF neurons, their simultaneous activations did not themselves generate feedback loops. Despite each stimulus evoking spikes in two channels, the average size of the feedback loops was smaller for this IFNN than the one shown in Figure 3.4B due to the smaller number of excitatory connections (Supplementary Figure 3.6.3). Nevertheless, both of these networks displayed much larger feedback loops than the IFNN with three inputs and outputs. Both of these experiments illustrate the importance of the IFNN connectivity on shaping the characteristics of the closed-loop neuronal spike dynamics. Supplementary Figure 3.6.5 shows the complete set of auto-, and cross-correlograms of the neuronal spike times for these experiments, as well as the IFNN weights/parameters, and the positions of the recording and stimulation channels on the Utah array.

3.4.3 Hidden Layers Gate Patterns of Activation

The addition of hidden layers to neural networks is known to increase their computational capabilities [152]. Furthermore, it has been well-established that both conventional rate-based, and spike-based feedforward artificial neural networks with just one hidden layer are capable of approximating any continuous function [144], [153]. We therefore sought to document the effects of adding a hidden layer to an IFNN on the closed-loop dynamics. Since the IFNNs were restricted to a combined 8 biological neurons and IF neurons, we were limited to using three biological neurons for the IFNN, which is shown in Figure 3.4D. Furthermore, because the output IF neurons were sparsely activated, even with large weights, we did not add any inhibitory connections to the network. The weights between the hidden layer and the output IF layer were set so that if any two hidden units fired, both output IF neurons would typically co-fire, although due to differences in weights this did not always occur. Although each of the input neurons were in some way connected to each of the output IF neurons, in contrast to the feedback networks with no hidden layers, the hidden layer gated which patterns of input spikes could activate the output IF neurons. For instance, any input neuron could activate the output IF neurons by spiking twice within a small enough window, but two spikes from two separate neurons was insufficient to activate the output IF neurons. Three spikes from any

of the three input neurons, however, was sufficient to evoke the output IF neurons. The hidden layer also extended the window of time for which the input neurons could spike to evoke outputs in the output IF neurons – R1, R2, and R3 could fire sequentially in any order over a period of time that outlasted the duration of the first PSP and still trigger the output IF neurons.

Figure 3.4D (right) shows three examples of the auto-, and cross-correlograms of neuronal spike times during the Pre and CL epochs. One feature of this IFNN was that due to the low probability of evoking a spike in R3 at small values of Δt_{prev} (Figure 3.5F), the auto-correlogram of R3 shows a very small peak near 7 ms (black arrow) corresponding to the evoked spikes. This is in contrast to the large peaks in the auto-correlogram of R1. Another interesting feature is that because R2 was evoked from the same stimuli that evoked R1, there was a broad peak in their cross-correlogram near 0 ms. However, due to the large stimulation delay and evoked spike latency for R2 (Supplementary Figure 3.6.2B, top), its CL epoch auto-correlogram was nearly identical to its Pre epoch auto-correlogram (Supplementary Figure 3.6.6). These results suggest that it is possible to create artificial correlations in activity, while enabling individual neurons to retain their natural firing properties. Finally, although the stimulation delays were 15 ms apart, this IFNN was capable of generating feedback loops, an example of which is shown in Figure 3.5G. To generate this feedback loop, a spontaneous spike from R1 could fire just prior to a stimulus pulse that evoked a spike in R1. The two spikes from R1 would then trigger ICMS, which could then evoke more spikes. Supplementary Figure 3.6.6 shows the complete set of auto-, and cross-correlograms of the neuronal spike times for this experiments, as well as the IFNN weights/parameters, and the positions of the recording and stimulation channels on the Utah array.

3.4.4 The Closed-Loop Dynamics Can be Predicted from Open and Closed-Loop Measures

Once we performed the initial set of experiments *in vivo*, we sought to model the closed-loop dynamics between the biological and artificial neurons to study their interactions free from experimental constraints such as non-stationarities in the brain (Supplementary Figure 3.6.1), experimental run-time, and hardware

limitations. Figure 3.4 compares the simulated dynamics from each of the recurrent networks to the Pre and CL dynamics. As shown, the auto-, and cross-correlograms generated by the simulation match the experimental results, albeit with minor deviations. Importantly, the simulation was able to predict the dynamics during the CL epoch from data collected from the Pre and OL epochs, despite non-stationarities in the brain (Supplementary Figure 3.6.1). For the network shown in Figure 3.4A in which we did not collect the open-loop responses of the neurons to stimulation, we used the CL epoch to build the stimulation response model, but still used their Pre epochs to generate the Markov model. We also found that the average feedback loop size, and distributions were similar in the CL and simulation epochs, albeit with some deviations (Supplementary Figure 3.6.3). In particular, we found that the average feedback loop sizes for the simulated dynamics were larger, which could arise from false negatives or rebound spike distributions that were skewed to smaller values due to the high rate of stimulation.

3.4.5 Inhibitory Connection Strength and Stimulation Delays Impact Feedback Loop Characteristics

We used the model of the closed-loop dynamics to probe how inhibitory connections and stimulation delays affect the average sizes of the feedback loops. For this first analysis, we used the network architecture from Figure 3.4B to probe how inhibition influenced the size of the feedback loops. We simulated the closed-loop dynamics with the same connections and stimulation delays but varied the magnitude and fall time constants of the IPSPs. The magnitude of all the inhibitory connections was varied between amplitudes of 0 – 200% of the spike threshold in increments of 25%, and was equal for all inhibitory connections. The IPSP fall time constants between 4 - 20 ms in increments of 2 ms, while keeping the rise times the same at 1.5 ms. Each of the simulations was 15 minutes long. Both the fall time constant and magnitude of the IPSPs similarly influenced the size of the average feedback loop (Figure 3.5A), which monotonically decreased as either the PSP time constant or inhibition strength increased. We found that the larger feedback loops (black arrows) were decreased to a much greater extent than smaller feedback loops, as shown in the

smoothed auto-correlograms of R3 and R4 shown in Figure 3.5B. Examples of each of these feedback loops are shown in Figure 3.5C.

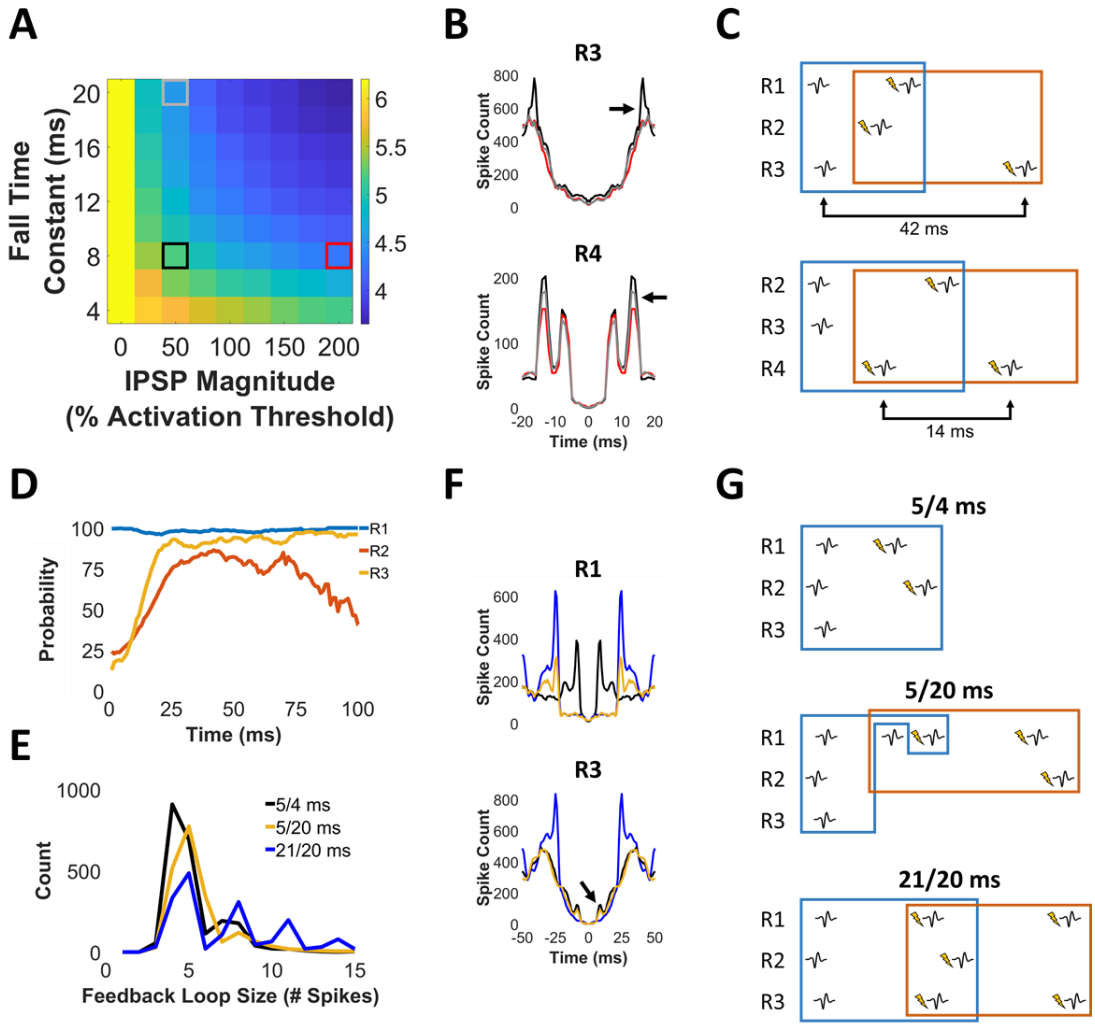

Figure 3.5 Measuring the Effects of Inhibition Strength and Stimulation Delays on the Closed-Loop Spike Dynamics

(A) Color plot of average feedback loop size showing monotonic effects of IPSP magnitude and fall time constant. Simulations were 15 minutes each and were performed using the IFNN in Figure 3.4B. (B) Smoothed auto-correlograms (3 ms moving average) of R3 and R4 showing changes in feedback loops with increased IPSP magnitude (red) and increased IPSP fall time constant (gray). As shown, larger feedback loops are much more strongly diminished than smaller feedback loops. (C) Examples of how each corresponding feedback loop is generated. (D) Probability of evoking a spike as a function of Δt_{prev} for the IFNN shown in Figure 3.4D. (E) Number of feedback loops between 0-15 spikes for each of the three 15 minute long simulations of the IFNN in Figure 3.4D with three different combinations of delays. Feedback loops in 21/20 ms IFNN have the largest feedback loops due to R2 and R3 having much larger probabilities of being evoked at larger values of Δt_{prev}. (F) Auto-correlograms of R1 and R3 for each of the simulations. (G) Examples of feedback loops for each of the three simulations.

We explored how stimulation delay impacted the size of the feedback loops using the IFNN shown in Figure 3.4D, which showed a strong dependence of evoked spike probability on the value of Δt_{prev} for R2 and R3 (Figure 3.5D). For this analysis, we simulated 15 minutes of the closed-loop operation for three different combinations of delays. The first was the original set of 5 and 20 ms (5/20 network), the second had stimulation delays of 5 and 4 ms, (5/4 network), and the third had stimulation delays of 21 and 20 ms (21/20 network), respectively. Because the evoked spike timing of R3 occurred earlier than R1 and R2, we set its stimulation time to be 1 ms after R1 and R2 to prevent the blanked evoked spikes from the simulated stimulation artifacts. Figure 3.5E shows the number of feedback loops of various sizes between 3 - 15 spikes for each of the three networks tested. The size of the feedback loops was the greatest for the 21/20 network, which was when the probability of evoking a spike was greater for R2 and R3. The larger feedback loops also manifested as an increase in the features in the autocorrelogram of R1 and R3 shown in Figure 3.5F. Interestingly, despite having delays of 15 ms apart, therefore making it impossible for the evoked spikes to trigger feedback loops without additional spontaneous spikes, the feedback loops tended to be larger in the 5/20 network than the 5/4 network. Although the stimulation delays were much closer in the 5/4 network, the evoked spike probabilities for R2 and R3 at these delays were too low to generate large feedback loops. However, for the 5/20 ms network, the 20 ms delay was long enough to allow spontaneous spikes from R1 to occur just before stimulation. Since R1 had a uniformly high probability of being evoked at all values of Δt_{prev}, the spontaneous spike of R1 could pair with an evoked spike of R1 to trigger more stimulation and generate a feedback loop. Example feedback loops from each of these networks are shown in Figure 3.5G. Supplementary Figure 3.6.7 shows the complete set of auto-, and cross-correlograms of the neuronal spike times for these simulations.

Altogether, these results highlight how both the strength of inhibition, as well as the stimulation delays can impact the ability of the IFNNs to generate feedback loops, and therefore maintain memories of their previous inputs.

3.4.6 IFNNs with Reciprocal Inhibition Generate Rhythmic Firing in Biological Neurons

In addition to implementing the 4 IFNNs with recurrent connections, we implemented 2 other IFNNs in monkey J to demonstrate the capability of IFNNs to recreate biomimetic circuitry. These IFNNs had an architecture that is known to generate rhythmic activity in the central nervous system (Figure 3.6A) [154]. Because ICMS generated both an excitatory and inhibitory response in the recorded neurons [143], an advantage of testing this specific circuit was that it could function without the need to exclusively produce an excitatory or inhibitory response. Figure 3.6B shows examples of auto-, and cross-correlograms from both experiments. The cross-correlogram of R3 and R4 in both experiments show the generated rhythms in the CL epoch compared to the spontaneous dynamics in the Pre epoch. Since the properties of the rhythms were dependent on the weights of the network connections, for both these networks, we input the spike trains of the input neurons from the Pre epoch directly into the network and symmetrically tuned the weights and PSP shape until rhythmic stimulation outputs were generated, after which we implemented the networks *in vivo*. Supplementary Figure 3.6.8 shows a general set of rules that we used to change the characteristics of the rhythms. Note that we did not include the responses of either the input or output neurons to ICMS for this tuning to minimize the amount of time necessary to adjust the weights. Figure 3.6C show the cross-correlograms of the predicted and experimental rhythmic outputs of the stimulation outputs S1 and S2, as well as the spike times of the output neurons R3 and R4 during the CL epoch. There were also recurrent connections onto R1 and R2 in the experiments which slightly altered the properties of the rhythms. We observed that the spikes of the output neurons R3 and R4 in the experiment shown at the bottom of Figures 6B and 6C were not as locked to the rhythms as the output neurons in the experiment shown at the top of Figures 6B and 6C. This could in principle arise from less stimulation, a lower probability of evoking a spike, or because the spontaneous firing rates of the output neurons were higher. We did not find support for either the first or second choice – both the rate of stimulation (25 and 28 Hz for networks top and bottom of Figures 6B and 6C, respectively), as well as the evoked spike probabilities of both sets of output neurons

were similar (>90% evoked for both outputs in both experiments). However, the firing rates of the output neurons for the experiment shown at the bottom of Figures 6B and 6C were larger (Table 3.2)

Table 3.2 Comparing Firing Rates of the Biological Neurons During Pre, Closed-Loop, and Post Epochs

Figure	Pre (Hz)	CL (Hz)	Post (Hz)
4B	R1, R2, R3, R4 20.5, 24.4, 18.1, 1.4	R1, R2, R3, R4 19.2, 17.5, 18.3, 5.9	R1, R2, R3, R4 22.7, 23.6, 14.6, 1.5
4D	R1, R2, R3 11.5, 19.9, 16.5	R1, R2, R3 14.1, 22.3, 14.7	R1, R2, R3 14.3, 20.9, 11.0
5B (top)*	R1, R2, R3, R4 25.0, 12.5, 7.5, 1.1	R1, R2, R3, R4 17.6, 10.0, 12.8, 7.4	R1, R2, R3, R4 22.3, 12.3, 9.0, 1.2
5B (bot)*	R1, R2, R3, R4 19.7, 22.3, 7.4, 10.4	R1, R2, R3, R4 19.9, 9.1, 14.2, 13.1	R1, R2, R3, R4 18.9, 6.5, 6.7, 12.5

The rhythmic outputs were also observed in the spike raster plots of the output neurons R3 and R4. Figure 3.6D shows a 10 second snippet of the spike raster during the Pre and CL epochs for the experiment shown at the top of Figures 6B and 6C. The firing rates of R3 and R4 over the 10 second period are also shown above, and were calculated by first measuring the firing rate in 1 ms bins and smoothing them with a 500 ms Gaussian filter. The correlation coefficients in these firing rates show how the activity of the neurons, which were very slightly positively correlated during the Pre epoch, became much more negatively correlated during the CL epoch. Supplementary Figure 3.6.9 shows the complete set of auto-, and cross-correlograms of the neuronal spike times for these experiments, as well as the IFNN weights/parameters, and the positions of the recording and stimulation channels on the Utah array.

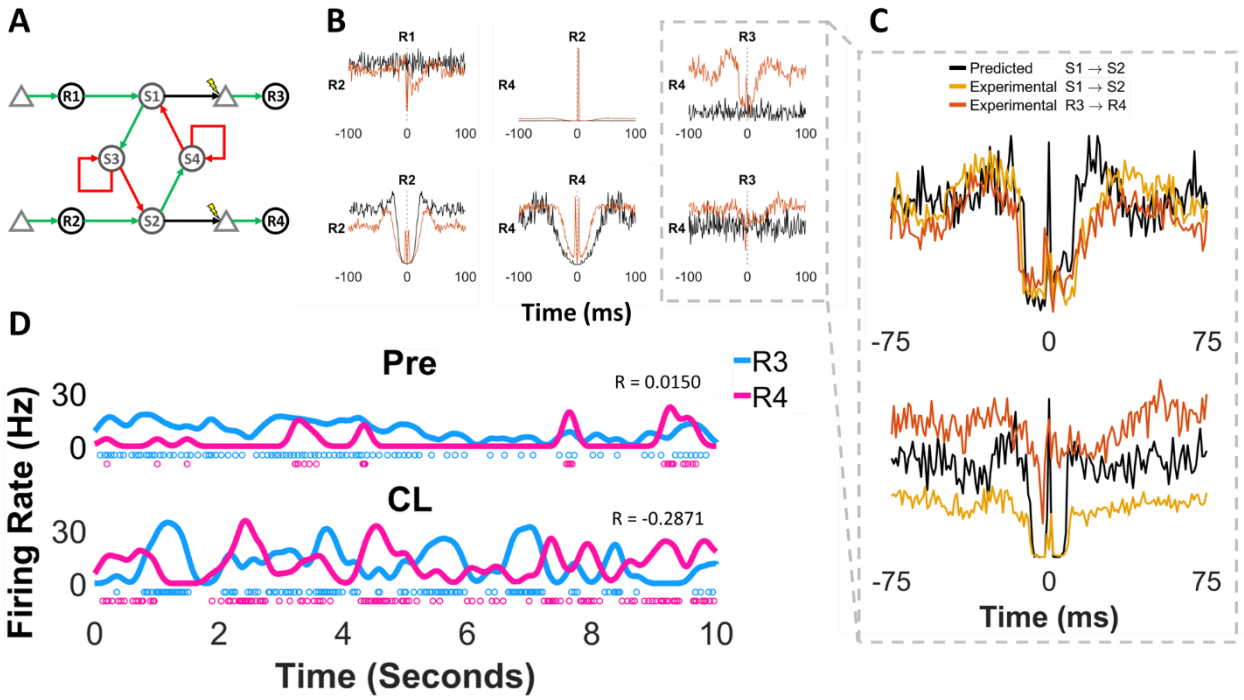

Figure 3.6 Generating Rhythmic Activity in Neurons

(A) IFNN architecture tested in monkey J for two experiments. (B) Examples of auto-, and cross-correlograms for a selection of the neurons. Cross-correlograms between R3 and R4 show generated rhythmic activity during the CL epoch. Cross-correlogram between R1 and R2 (top), and auto-correlogram of R2 (bottom) show small unintentional feedback connections from stimulation outputs that slightly altered CL dynamics. (C) Comparisons between predicted and experimental cross-correlograms of stimulation outputs S1 and S2, as well as the cross-correlograms of the output neurons R3 and R4 during the CL epoch for both experiments. Top experiment shows that R3 and R4 are more locked to the stimulation outputs due to their lower spontaneous firing rates. (D) 10 second snippet of a spike raster and firing rates of R3 and R4 during the Pre and CL epochs. R3 and R4 have a small positive correlation in their firing rates during the Pre epoch, which turns into a more negative correlation during the CL epoch.

3.4.7 The Spontaneous Firing Properties of the Cortical Neurons are Unchanged After Closed-Loop Stimulation

In the 4/6 experiments in which we recorded the activity of the cortical neurons following stimulation (durations of stimulation shown in Table 3.1), we found that the spike characteristics including firing rate (Table 3.2) and auto-, and cross-correlogram shape (Supplementary Figure 3.6.10) were typically unchanged after the CL operation. The typically small changes were in line with the changes observed throughout the Pre epoch (Supplementary Figure 3.6.1A).

3.4.8 Stimulation Artifacts Obscure Spikes as a Function of IFNN Size and Synaptic Weights

Although the IFNNs were able to continuously function in the presence of stimulation artifacts, some spikes were obscured, or blanked due to their coincident timing with the artifacts. In general, increasing the number of stimuli will lead to an increase in the amount of signal lost, and therefore the number of blanked spikes. We therefore sought to determine the extent to which stimulation artifacts from ICMS degraded the functioning of the IFNNs due to loss of spikes. In particular, we probed how the number of inputs and outputs of the IFNN, as well as the magnitude of the weights were related to the number of blanked spikes. To eliminate several experimental constraints, we performed this analysis via computational modelling. We performed two sets of simulations – In the first set of simulations, the stimuli did not evoke any spikes or inhibition, but simply blanked the spikes for 1 ms following each stimulus. In the second set, each stimulus could evoke a spike and inhibit one of the input neurons, while blanking the spikes of the remaining neurons. We selected a total of 13/14 of the simulated neurons from Figure 3.4 for this analysis (the putative polysynaptic connection between S4 and R2 in Figure 3.4D was excluded due to the very high variability in evoked spike latency). As with the *in vivo* experiments, any two stimuli that were set to occur within the same interval were staggered in time. For each of the sets of simulations, we tested IFNNs with 4-10 neurons/units each in size increments of 2 neurons/units. Each output IF neuron triggered stimulation at a randomly assigned delay of n*4, for n = 1,2…10, which was done to keep each of the delays at least 4 ms apart to minimize the amount of staggered stimulation artifacts. To isolate the effects of increased network size, rather than the total number of synaptic connections, we restricted each output IF neuron to only having 4 random presynaptic connections from the biological neurons. Furthermore, no inhibitory connections or hidden layers were used to simplify the analysis and gain a basic understanding of the relationships between network size, weights, and blanked spikes. For each of these IFNNs, we varied the magnitude of the weights between 50-87.5% of the activation threshold of the output IF neurons in increments of 12.5%. We repeated each of these simulations 10 times each with random connectivities, neurons, and stimulation delays to obtain an average number of blanked spikes.

Figure 3.7A shows two color plots of each set of simulations. For the simulations in which stimulation did not evoke any spikes or inhibit activity, the percentage of blanked spikes had a clear monotonic relationship with the magnitude of the weights and network size. However, for the simulations in which evoked spikes were included, this relationship was less clear. While the magnitude of the weights had a strong monotonic relationship with the percentage of blanked spikes, the size of the network was less important. This occurred because the outputs of the IF neurons were more synchronized in smaller IFNNs due to more overlapping connections between the input biological neurons and output IF neurons. This synchrony caused a greater fraction of the blanked spikes to be evoked spikes rather than spontaneous spikes in smaller IFNNs (Figure 3.7B), which manifested as a near uniform fraction of total blanked spikes for IFNNs of all sizes. However, larger IFNNs still displayed some level of synchronized outputs, which in turn synchronized some of the evoked spikes and increased the total number of blanked spikes in simulations with evoked spikes compared to those without (Figures 7A and 7C).

Finally, we combined each set of simulations together to measure the average number of spikes missed for each of the neurons as a function of their spontaneous firing rate, obtained from the Pre epoch (Figure 3.7C). When the stimulation did not evoke any spikes, there was an almost perfect positive correlation between the firing rate of the neuron and the number of spikes that were missed. Although the total number of spikes that were missed was greater when stimulation was allowed to evoke spikes, the relationship between firing rate and number of evoked spikes was significantly less correlated due to the varying evoked spike probabilities and magnitudes of inhibition.

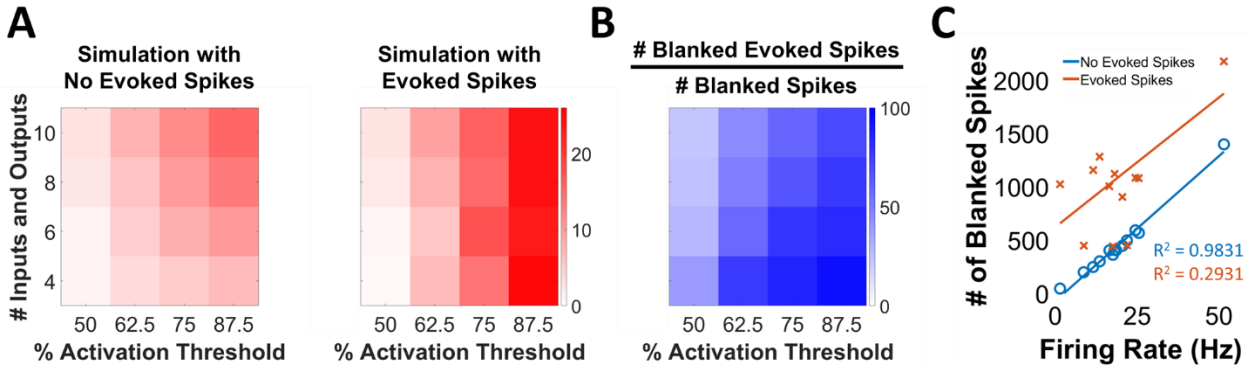

Figure 3.7 Quantifying the Effects of Stimulation Artifacts

(A) Color plots showing the percentage of blanked spikes as a function of the magnitude of the weights and number of input biological neurons and output IF neurons for the simulations with no evoked spikes (left) and with evoked spikes (right). (B) Color plot showing the percentage of blanked spikes that were evoked spikes compared to spontaneous spikes for the simulations with evoked spikes. (C) Relationship between the firing rates of the biological neurons and the number of spikes that were blanked due to stimulation artifacts.

3.5 Discussion

We defined a neural co-processor as a neuromorphic bidirectional BCI that could artificially connect neurons together to provide neurorehabilitation or augmentation [139], [140], and listed a set of at least three characteristics it should have. Towards the goal of developing such a co-processor, we interfaced collections of biological and IF neurons together with a variety of network architectures with the goal of understanding how different features such as recurrent connections, hidden layers, and strong reciprocal inhibition shape neuronal activity.

The first listed criterion was that it would be able to create artificial connections between neurons and manipulate their activity in a context-specific manner. Although we did not completely satisfy this criterion, the IFNNs did create clear artificial connections, as reflected in the auto-, and cross-correlograms, between the biological neurons during the closed-loop stimulation. Interestingly, we found that beyond creating simple connections between the biological neurons, the recurrent IFNNs created positive feedback loops in the spike times of the biological neurons through the biological-artificial-biological pathways. The characteristics of these feedback loops were dependent on the IFNN architectures, stimulation delays, as

well as the specific inputs that generated them (Figure 3.5C). The uniqueness of the feedback loops based on these parameters could potentially enable the IFNN to generate input-specific short-term memories in the spike activity of the biological neurons, similarly to other recurrent neural networks such as echo state networks and liquid state machines [151].

Because ICMS can generate both an excitatory and inhibitory response in the biological neurons whose relative strengths depend on both the simulation timing and amplitude [143], future work could explore how the weights, as well as stimulation delays and stimulation amplitudes of the IFNNs could be trained to selectively increase or decrease the firing rate of the biological neurons in a context-specific manner. One of our most significant findings was that the combination of two simple models – a Markov model to simulate spike trains, and a stimulation response model to simulate the responses of neurons to ICMS – was capable of accurately reproducing the closed-loop dynamics from open-loop measures. This computational model provides a means of developing such a training algorithm free from the constraints of developing them *in vivo* due to challenges such as non-stationarities in firing rates and responses to ICMS (Supplementary Figure 3.6.1). Due to the necessity of changing the weights, as well as the stimulation delays and amplitudes, a reinforcement or evolutionary learning algorithm that could efficiently explore this state-space would be an ideal learning algorithm.

The ability of the co-processor to co-learn and co-adapt with the brain to account for non-stationarities in the brain was the second criteria for the co-processor. However, if the IFNN is trained offline, these non-stationarities will eventually compromise its ability to properly function. To overcome this limitation, the IFNN could first be trained offline, while smaller adjustments to its parameters are made online over time to change with the brain. The approach of first training the IFNN offline with the model would likely have the added benefit of decreasing training time compared to a completely model-free online learning algorithm since the state-space can be explored much more rapidly offline. Together, this learning approach would satisfy both the first and second criteria for the co-processor. Another limitation with using ICMS to shape neuronal activity is that due to the mechanisms by which ICMS activates neurons, the excitatory and

inhibitory effects cannot be fully decoupled, which constrains its ability to selectively control neuronal activity. The coupling of excitatory and inhibitory responses likely arises from the complex recruitment of both excitatory and inhibitory neurons with ICMS [143], [155], [156]. This limitation might be overcome by optogenetic or juxtacellular stimulation, both of which could activate neurons more locally and selectively.

The third criterion was that the co-processor would utilize temporal coding. Because the IF neurons that we used required two excitatory inputs to arrive within a predetermined coincident interval, two sets of inputs could have had the same average firing rates, but activated the IF neurons in completely different ways due to differences in the relative timings of the inputs. We also showed how the stimulation delays of the IF neurons can profoundly influence the closed-loop dynamics by controlling the probability of evoking spikes (Figure 3.5D-F and Supplementary Figure 3.6.7). Together, these show how the IFNN utilized a highly simplified form of temporal coding for both decoding and encoding. Although we did not explicitly test this, the characteristics closed-loop dynamics could also be controlled by allowing the synaptic connections within the IFNN to have different delays, which could be used to control the timing of the individual output IF neurons for different patterns of input spikes. Finally, while the IF neurons utilized a simple temporal coding scheme of detecting coincident inputs, more sophisticated SNN models such as the Izhikevich model could be used to increase the complexity of the temporal code by simulating the membrane characteristics of various cell-types [47], [49].

The fourth criterion that we listed was that the co-processor should be able to isolate authentic neural signal from stimulation artifacts, and that the signal lost due to the stimulation artifacts should limited and not compromise the functioning of the IFNN. To eliminate the confounds of performing this analysis experimentally due to non-stationarities, we used the model that we developed to systematically probe how the size and weights of the IFNN obstructed the detection of spikes due to stimulation artifacts. We found that when evoked spikes were included in the model, a greater number of spikes were blanked due to the synchronization of IF outputs, which in turn synchronized the evoked spikes. For this reason, it may be

possible to reduce the number of blanked spikes by more carefully tuning the stimulation delays for each IF neuron rather than choosing them randomly to account for both the evoked spike probabilities and latencies. This would enable the IFNN to control the relative magnitudes of excitation and inhibition in each neuron while minimizing the number of blanked spikes. One challenge with performing this analysis is that the IFNNs were not designed to perform a task, which limited our ability to determine the extent to which the functioning of the IFNN was compromised. Future studies that train the IFNN for such goals can perform this analysis with this question in mind. Nevertheless, our system enabled us to continuously record and stimulate, while allowing us to separate spikes from the stimulation artifacts. With all of this in mind, there are numerous efforts underway to develop systems capable of subtracting the artifacts out of the recordings [157], [158], which will mitigate this limitation.

Finally, if a co-processor is to be used to manipulate neuronal activity to provide functional neurorehabilitation or augmentation, it is important to document how ICMS changes the activity of neurons due to a combination of factors such as stimulation-induced short-term plasticity [96], changes in neuronal excitability [159], and homeostatic plasticity [160]. Interestingly, we did not observe any substantial changes to the firing rates (Table 3.2), or features in the auto-, and cross-correlograms for the single neurons between the Pre and Post epochs (Supplementary Figure 3.6.10). Most changes were inconsistent between neurons and were almost always small enough to be accounted for by the neurons' non-stationarity (Supplementary Figure 3.6.1A). Whether the brain adapts to the continuous closed-loop interface, or extended periods of stimulation were not investigated due to the challenges in reliably recording extracellular action potentials from the same neurons continuously for several days. It will be essential to explore this question further to develop an understanding of how the BCI will need to adapt to any such changes.

Author Contributions

Jonathan Mishler prepared the manuscript, including data analysis and interpretation. Richy Yun assisted with coding and performing experiments. Steve Perlmutter, Rajesh Rao, and Eberhard Fetz helped edit the manuscript and guide discussions.

Chapter 4.

COMPUTATIONAL MODELING OF REPAIRING A BIOLOGICAL NEURAL NETWORK WITH AN ARTIFICIAL NEURAL NETWORK TO RESTORE LOST MOTOR FUNCTION

Jonathan H. Mishler, Steve I. Perlmutter, Rajesh P. N. Rao and Eberhard Fetz

4.1 Abstract

As the technology to interface with large collections of neurons from various regions of the brain develops, it is important to develop algorithms that can decode and encode information into the brain for neurorehabilitative and neural augmentation applications. Spiking neural networks (SNNs) are an excellent tool for such applications due to their ability to detect patterns of activity that are encoded in both spike rate and spike timing and deliver context-specific stimulation at specific times relative to the background neuronal activity. Computational models present a great opportunity to test these learning algorithms free from experimental constraints. In this chapter, I discuss preliminary work in which I trained a recurrently connected simulated "biological" SNN (BN) comprised of artificial integrate-and-fire (IF) neurons to perform a 1D tracking motor task using a reinforcement learning algorithm. I then lesioned the neural network by randomly deleting a subset of the neurons which compromised its ability to perform the task, and showed how even with retraining, the SNN was incapable of relearning the task. Finally, I discuss future work in which a second "artificial" SNN (AN) could be interfaced with the BN to supplement the damaged network by simulating ICMS to shape the dynamics in the BN and restore lost function.

4.2 Introduction

While technology that enables the chronic recording and manipulation of neurons throughout the brain is developing, it is important to develop tools that can decode specific patterns of activity and use stimulation to shape neuronal activity in a context-specific manner. There are numerous efforts underway to develop various algorithms for neuronal control both electrically and optogenetically, using various tools such as neural networks and state-feedback controllers [146], [161]. One limitation of these approaches, however, is that due to technological limitations such as the relatively small number of neurons that can be interfaced with, these control algorithms are incapable of being utilized to restore lost function arising from neurological damage. While the technology to interface with more neurons is being developed, there has been an interest in using computational models of the brain to develop the control algorithms. Towards these goals, the Lytton lab developed a model of the sensorimotor system which consisted of several

populations of spiking neurons to simulate proprioceptive, sensory, and motor neurons that were connected to a simulated musculoskeletal arm [162]. The SNN was trained to perform a reaching task using a reinforcement learning algorithm, after which it was lesioned, thereby compromising its performance. They developed an inverse model of their neuronal networks by probing the neuronal responses to stimulation, which they then used to find the patterns of stimulation that were necessary to restore the pre-lesioned network activity, thereby restoring function. However, there are several limitations with such an approach that limit its therapeutic utility. First, as we have previously shown in our lab, the timing of stimulation relative to the background network activity is important in determining the responses of the neurons to stimulation [143]. Second, the inverse controller used a single burst of stimulation for 200 ms between 100 – 300 ms post trial onset, which prevents the model from flexibly adjusting the trajectory the movement of the arm during the task. While this may be sufficient for the simple and short task that was simulated, it cannot be utilized for more naturalistic tasks. Third, the inverse model required the relationship between the stimulation of the neurons and subsequent motor movements to be explored, which is not feasible for *in vivo* applications.

Toward addressing these limitations, I developed a recurrently connected BN comprised of integrate-and-fire neurons to perform a 1D tracking motor task. I trained the BN using a reinforcement learning algorithm, after which the network was lesioned by deleting a subset of neurons. I showed how even with new retraining, the lesioned BN was incapable of relearning the task. I then discuss future work in which a second AN could be interfaced with the BN to restore lost function via stimulation. This work would address each of the limitations of the Lytton study. First, the probability of evoking spikes in the BN neurons would depend on the time delay between the previous spike time and stimulation onset. Second, stimulation would occur continuously throughout each trial, which could be used to continuously update the position of the cursor to track the target. Third, I use a reinforcement learning algorithm to train the networks which does not require probing the relationship between stimulation and motor response.

4.3 Materials and Methods

4.3.1 1D Motor Task

The task goal was to control a cursor to track a target in 1D. The positions of the target and cursor were randomly initialized to be on a line whose length was 21 arbitrary units (AU). Thus, their relative positions, P_r, spanned [-10, 10] AU. The target was randomly placed on a location on the line and held its position for 10 seconds while allowing the cursor to track it, after which it repeated this process an arbitrary number of times.

4.3.2 Biological Neural Network Architecture

The BN was comprised of a total of 37 IF neurons, and included 11 input neurons, 11 excitatory neurons, 11 inhibitory neurons, and 4 output neurons that controlled the cursor position (Figure 4.1).

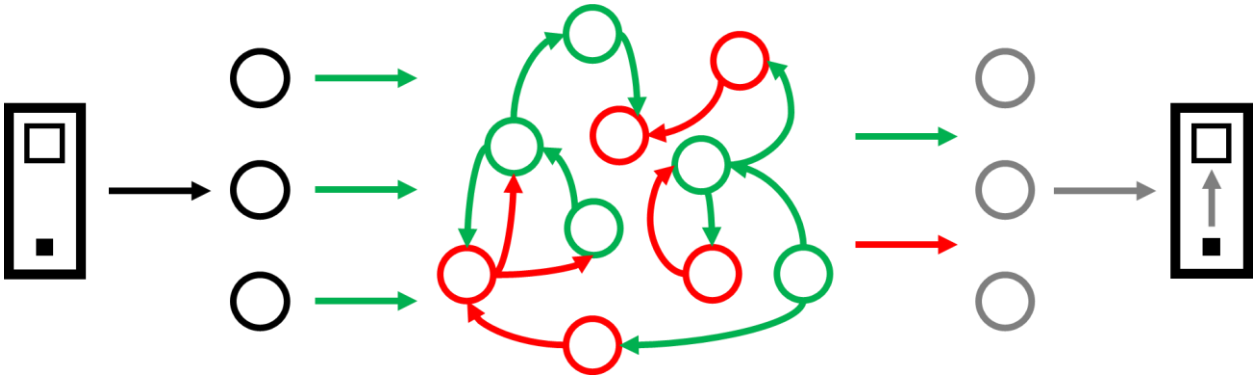

Figure 4.1 Schematic of Biological Network

Based on the relative positions of the cursor and target, the 11 input neurons (black) send excitatory inputs to the 11 recurrent excitatory neurons. Both the recurrent excitatory (green) and inhibitory (red) neurons send outputs to the output (gray) neurons which control the movement of the cursor.

Table 4.1 shows the probability with which each type of neuron was connected to each other, and their randomly initialized weights which was drawn from a normal distribution.

Table 4.1 SNN Connectivity and Initial Weights

Connection Type	Connection Probability	Connection Weight (μ +/- σ)
Input → Exc	40%	100 +/- 40
Exc → Exc	40%	25 +/- 10
Exc → Inh	40%	150 +/- 60
Inh → Inh	40%	25 +/- 10
Exc → Out	40%	75 +/- 30
Inh → Out	40%	25 +/- 10

When each of the excitatory or inhibitory neurons fired, they sent PSPs to each of their connected neurons. The shapes of the PSPs were governed by Equations 3.1A-C. For these experiments, the rise times of the PSPs was approximately 3 ms, while the fall time constants were approximately 20 ms. The simulations were updated every 1 ms. The neurons spiked when they reached a predetermined threshold. Each of the input neurons sent suprathreshold Poisson inputs to their connected excitatory IF neurons, whose firing rates were dependent on the value of P_r. The inputs did not change once the absolute values of P_r exceeded 3 AU. The input neurons were designed so that 1/11 of them would send inputs to the BN at a rate of 10 Hz, 2/11 would send them at a rate of 91 Hz, and another 2 would send them at a rate of 83 Hz for each value of P_r between [-3, 3].

4.3.3 Cursor Control

Each of the 4 output neurons were assigned random "spontaneous firing rates" between 1-52 Hz (thresholds for cursor control). The firing rates were chosen to match the firing rates of biological neurons that were used in the *in vivo* experiments outlined in Chapter 3. The cursor was controlled by the differences in firing rates between the output neurons during the task and their randomly assigned thresholds. For 2 of the output neurons, if their input-driven firing rates during the task exceeded their randomly assigned thresholds, the cursor would move up, whereas lower firing rates caused the cursor to move down. For the other 2 output neurons, the cursor response was reversed. The spikes of the output neurons were binned every 200 ms,

after which their firing rates were collected. The net change in the cursor position, ΔC, was driven by changes in firing rates of the output neurons, and was calculated using Equation 4.1. First, the differences between each output neuron's firing rate and its randomly assigned thresholds were calculated. Next, the differences in firing rates, ΔFR_n for each output neuron was multiplied by $\vec{v}$ which was set to 1 or -1 depending on which direction it controlled the cursor. The 4 values were then added together and multiplied by a scaling factor, set to 0.05.

$$\Delta C = 0.05 \times \sum_{n=1}^{N} \Delta FR_n \times \vec{v} \quad 4.1$$

Since the selected input neurons were selected based on discrete values of P_r, they were chosen based on the rounded value of the updated cursor positions.

4.3.4 Training the Biological Neural Network

I developed a heuristic reinforcement learning algorithm to train the BN to control a cursor in 1D and have it track a moving target. To start training, I selected 4 unique, integer values of P_r between [-3, 3], and performed 4 corresponding simulations for each of these values. In each simulation, the unique value of P_r activated a unique set of inputs to control the cursor. At the end of each simulation, I computed the baseline error between the final position of the cursor and the target. The goal of the training algorithm was to minimize this error for each value of P_r. The simulation time was unique to each value of P_r and was calculated via Equation 4.2:

$$t_{sim} = min\left[500, abs\left(round\left(1000 \times \frac{3}{25} \times P_r\right)\right)\right] \quad 4.2$$

where t_{sim} is the total time in samples. The purpose of using different simulation times for different values of P_r was to train the cursor to move towards the target with a constant velocity. If the simulation time was as large for small values of P_r as it was for larger values, the velocity of the cursor would progressively slow down as the cursor neared the target. In addition to measuring the baseline error, I measured the "activation" of each synapse, which measured the extent to which each synapse was utilized over the course of each trial. Synapses that were more active during trials with lower error were strengthened, while those that were

less active were weakened. To calculate the activation for each excitatory synapse, I first measured how closely a postsynaptic neuron would fire after the presynaptic neuron fired. For these calculations, only postsynaptic spikes that occurred within 25 ms of a presynaptic spike were counted. I then measured the times between the pre and postsynaptic spikes, and subtracted those values from 25 ms. Therefore, these values were larger if the postsynaptic neuron spiked immediately following a presynaptic spike. To compute the final activation for the synapse, I then summed each of these values from the entirety of the trial. The calculation for the activations of inhibitory synapses was similar, but rather than requiring postsynaptic neurons to fire within 25 ms, they could be inhibited for up to 50 ms. For each inhibitory synapse, I measured the length of time between the pre and postsynaptic spikes. I then summed each of those values together from the entirety of the simulations to get the activation for each inhibitory synapse.

Once the initial error and activations were calculated for the 4 values of P_r, the same 4 simulations were performed another 4 times each, but in each of these 16 simulations the network was perturbed by stochastically injecting noise into it [see below]. The purpose of the noise was to perturb the network dynamics enough to change the final cursor position to explore how changes in the synaptic activity were related to successful task performance. If the difference in the final cursor position was closer to the target than the baseline performance, the more active synapses were strengthened, while the less active ones were weakened. To inject the noise, Poisson inputs between 20 – 100 Hz were sent to a random set of 3 – 6 recurrent excitatory and input neurons. These inputs were randomly chosen to be excitatory or inhibitory for each neuron. The magnitudes of the Poisson input PSPs were suprathreshold. Furthermore, only input neurons that fired for the particular value of P_r were capable of having noise injected into them. Noise was also added by randomly changing the thresholds of 3 – 6 recurrent excitatory and recurrent inhibitory connections by up to +/- 37.5% of the neurons' typical activation thresholds. The change in error as well as the activations of each of the synapses was calculated for each of the 16 simulations. For each of the 16 simulations, the difference in activation was calculated for each synapse via:

$$Diff_{exc} = \tanh\left(\frac{Activation_{exc,noise} - Activation_{exc,baseline}}{1000}\right) \quad 4.3A$$

$$Diff_{inh} = \tanh\left(\frac{Activation_{inh,noise} - Activation_{inh,baseline}}{1000}\right). \quad 4.3B$$

where $Diff_{pos}$ and $Diff_{neg}$ correspond to the differences in activation for the excitatory and inhibitory synapses, respectively. Rather than update the weights each iteration, the weights were updated in batches after performing the 16 simulations via:

$$Weights_{batch} = Weights_{batch} + 50 \times \Delta Error * \left(-Diff_{pos} + Diff_{neg}\right) \quad 4.4$$

where ΔError corresponds to the change in error between the noisy and the baseline simulations. Importantly, the batch weights were only updated when the noise lowered the error of task. After the 16 noisy simulations, the batch weights were added to the original weights, and a new baseline with 4 new values of P_r was obtained, after which the process repeated until the error became asymptotic. Excitatory and inhibitory synapses that decreased or increased to values past 0 were clipped to 0.

4.4 Results

4.4.1 The Biological Neural Network can be Trained to Perform a 1D Motor Task

I used the training algorithm to tune the BN weights to allow it to track a target in 1D. Figure 4.2A shows the total error (gray), and its moving average (black) over the training iterations. The total error for each iteration was obtained by summing the error of the 4 baseline error measurements. The total error reached an asymptote at just under 4 AU. Figure 4.2B shows the performance of the cursor tracking the target over a period of 100 seconds.

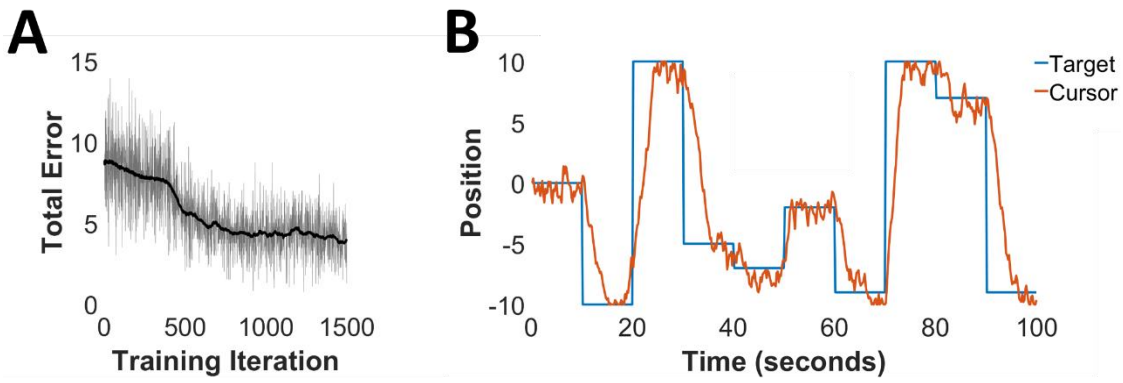

Figure 4.2 Training the Biological Network to Perform a 1D Motor Tracking Task

(A) Total error of the training algorithm as a function of the training iteration, which was asymptotic near 1500 training iterations. (B) Performance of the 1D target tracking task. The target stays at a random position on a line of length 21 AU for 10 seconds while the cursor follows. The position of the target then changes to a new random position.

4.4.2 Lesioning the Biological Neural Network Impairs Task Performance

Once the BN was trained to track a moving target, I lesioned the network by deleting a subset of the recurrent excitatory and recurrent inhibitory neurons, thereby compromising its ability to perform the task (Figure 4.3).

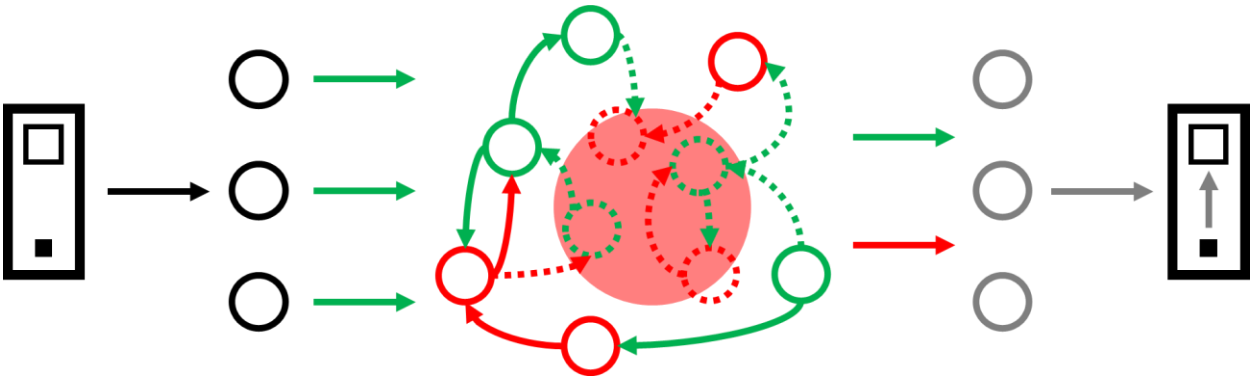

Figure 4.3 Lesioning the Biological Network

A subset of the recurrent excitatory and inhibitory neurons, including all their connections are deleted (dashed lines) to simulate a lesion, thereby compromising task performance.

For this analysis, 40% of these neurons and their connections were deleted. I then tried to retrain the network using the same reinforcement learning algorithm to restore its ability to perform the task. Figure 4.4A shows the total error over as a function of the training iteration. In contrast to the pre-lesioned BN, the training error did not decrease over time. Figure 4.4B shows the same target tracking task as the pre-lesioned BN performed. However, unlike the pre-lesioned BN, the lesioned BN could not perform the task at all.

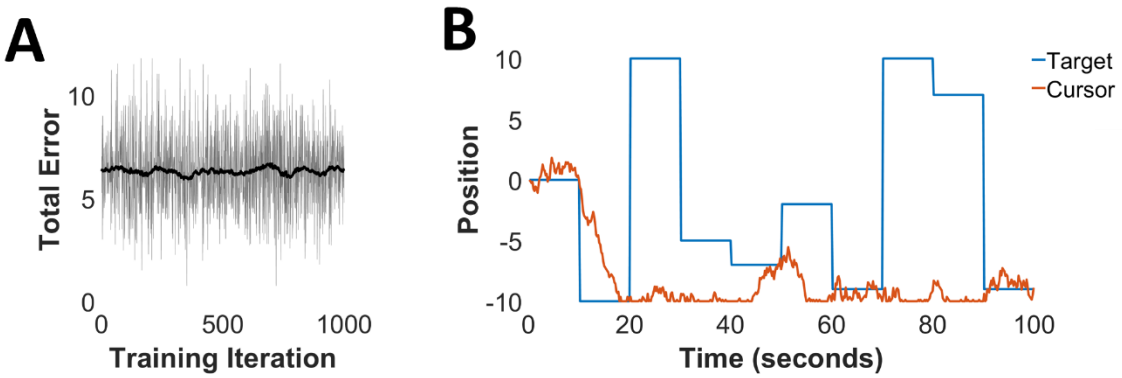

Figure 4.4 Training the Lesioned Biological Network to Perform a 1D Motor Tracking Task

(A) Total error of the training algorithm as a function of the training iteration, which did not decrease during the first 1000 training iterations. (B) Performance of the 1D target tracking task. Unlike the pre-lesioned network, the lesioned network was incapable of performing the task at all.

4.5 Discussion

In these sets of preliminary experiments, I trained a recurrently connected SNN to perform a 1D tracking motor task. The reinforcement algorithm utilized reward alone to strengthen synaptic connections that were beneficial to task performance, and weaken those that were unimportant. Lesioning the network by deleting a random 40% of the recurrent excitatory and inhibitory neurons destroyed task performance and prevented relearning the task. From these results, there are two immediate questions that will need to be addressed in future work. First, what is the extent to which the network can be retrained as a function of the lesion size? In the simulation presented above, the lesion compromised 40% of the total BN, which is a significant fraction of its original size. Furthermore, is the ability of the BN to relearn the task more hampered when excitatory or inhibitory neurons are preferentially lesioned? The second question concerns the ability of the network to relearn the task if more sophisticated training algorithms are utilized. In these models, a heuristic

reinforcement learning algorithm was used to train the BN. However, the inefficiencies of this learning algorithm may have itself contributed to the inability of the network to learn the task. Perhaps the use of state-of-the-art reinforcement learning algorithms would help the BN relearn the task.

A next step for this project would be to interface a second SNN with the BN, referred to as AN, that would utilize stimulation to restore lost function in the network (Figure 4.5).

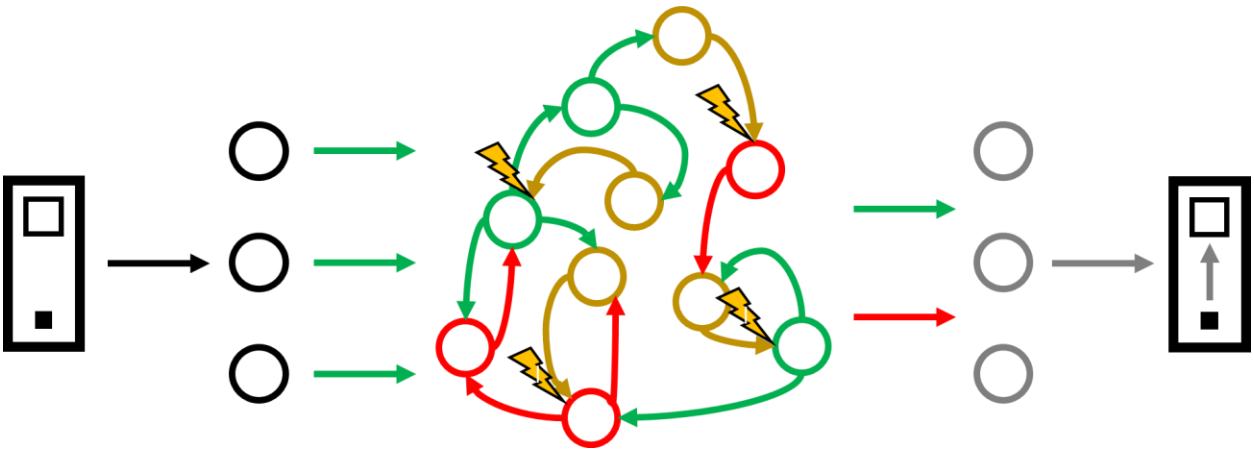

Figure 4.5 Repairing the Lesioned Biological Network with a Second Artificial Network
The biological network is repaired with a second artificial network (yellow) that uses ICMS to restore proper network function and subsequent task performance.

As with the *in vivo* experiments outlined in Chapters 2 and 3, the stimulation would both evoke spikes and inhibit the neuronal activity. A substantial advantage of using SNNs to restore network function is that they can be trained to utilize the timing of stimulation relative to the background network activity, which we have shown to influence the probability of evoking a spike. To train the BN/AN hybrid network, the connections between the BN, and between the BN and AN would be tuned with a similar reinforcement learning algorithm to when the BN was trained alone – noise would be injected into the neurons to change the performance of the network, and connections that increased the network's performance would be strengthened, while the other connections would be weakened. Whereas these connections can be changed in a continuous manner, the connections between the AN and BN, however, use ICMS which excites the neurons within the BN in a probabilistic manner. Therefore, to train these connections, both the ICMS

amplitude, which affects both the probability of evoking a spike and duration of inhibition, as well as the stimulation delay, which affects the probability of evoking a spike, could both be independently and randomly altered on a subset of the trials. Changes that improve the performance of the hybrid BN/AN network would be kept, while others would be discarded. This computational model would provide a platform to test the hypothesis that an artificial SNN that modulated neuronal activity is capable of repairing a lesioned biological network. The model could then be used to explore more sophisticated co-adaptation learning algorithms to eventually apply *in vivo*.

Chapter 5.

CONCLUSIONS AND FUTURE DIRECTIONS

With the increasing need to develop novel technologies for neurorehabilitation, ICMS is increasingly being considered for its neurorehabilitative applications. One challenge of utilizing ICMS for such applications is the limited understanding of how it excites neurons in the cortex, both directly and indirectly. Although previous studies have probed this question, we addressed important gaps in these studies by analyzing the trial-by-trial dynamics of neural responsiveness. I extended this work to probe how biological neurons can be interfaced with artificial integrate-and-fire units to manipulate the spike dynamics of the biological neurons. I explored and characterized how the closed-loop dynamics were highly dependent on the size, connectivity, number of hidden layers, magnitude of inhibition, and stimulation delays of the IFNNs. I also developed a computational model to simulate the closed-loop dynamics which captured the millisecond resolution spatiotemporal dynamics and used the model to further probe the relationship between inhibition and stimulation delays on the closed-loop dynamics. Below I will outline this work in greater detail and describe ways to move this work forward by training this co-processor to shape the spike activity of a group of neurons in task-specific manner to serve as a functional co-processor to the NHPs.

5.1 Analyzing the Responses of Neurons to ICMS on a Trial-By-Trial Basis Reveals Novel Evoked Spike Dependencies

ICMS is commonly used tool to probe cortical circuitry in a variety of animal models [26], [90], [163]. Previous studies have characterized the responses of single neurons to ICMS with a variety of parameters [26], [27], [90]. However, these studies typically explored the averaged effects of ICMS throughout each experimental session, rather than by a trial-by-trial measure for each stimulation pulse. By shifting the approach, we explored other dependencies of their activations. Of particular interest, we found that the probability of evoking a spike was dependent on the time delay between the previous spike time and

stimulation onset, irrespective of whether the previous spike was evoked or spontaneous. Furthermore, this probability distribution tended to be correlated with the autocorrelation of the neuron's spontaneous spike times; this suggests that probability of responding may reflect intrinsic membrane properties of the neurons. We found no such effect for the magnitude of inhibition. However, we did find that while the duration of inhibition was independent of whether a spike was evoked or not on a trial-by-trial basis, that their relative changes over time were typically correlated, which suggests that while different individual neurons contributed to each trial-by-trial response, that they are part of an interconnected network whose neurons are modulated together.

While our findings revealed these novel effects of ICMS on the activation of neurons, several more studies need to be performed to further elucidate the effects of ICMS. First, although we frequently observed that high frequency stimulation lowered the probability of evoking spikes for a short period of time, we did not explicitly measure the duration of these changes, but noted that they typically lasted less than a few minutes. Second, while we observed increases and decreases in the probability of evoking a spike over time due to short-term plasticity, we did not measure how the corresponding probability changed as a function of the time delay between the previous spike time and stimulation onset. For instance, are there broad changes in the probability distribution, or do certain time delays show preferential increases or decreases in the evoked spike probability, and can these results shed further light on the mechanisms by which ICMS activates the neurons? Third, and in a similar vein, how do these evoked spike probabilities and magnitude of inhibition change during different brain states such as NREM sleep?

5.2 The Closed-Loop Dynamics Between Biological Neurons and Artificial Integrate-and-Fire Units can be Simulated

I first explored how various features of the IFNNs, including connectivity and stimulation delays influenced the closed-loop dynamics of the IFNN. I found that the recurrent IFNNs frequently generated sustained patterns of activity, referred to as feedback loops, that maintained a memory of previous inputs similar to other kinds of recurrent neural networks. The properties of these feedback loops depended on the size of

the IFNNs, the number of excitatory connections, the stimulation delays, strength of inhibition, and the number of hidden layers in the IFNN. I also connected groups of biological neurons with IFNNs that had biologically inspired circuitry known to generate rhythmic activity [154]. Of particular interest, I found that the extent to which the IFNNs could manipulate the activity of the neurons was highly dependent on the spontaneous firing rates of the interfaced neurons.

I then designed a model that simulated the closed-loop dynamics of the recurrent IFNNs by combining a Markov model to simulate the spontaneous spike trains of the neurons, and a stimulation-response model to simulate the responses of the biological neurons to ICMS. I then used the model to further probe how features of the IFNN such as magnitude of inhibition and stimulation delays influence the closed-loop operation free from experimental and hardware constraints. I found that greater amounts of inhibition decreased the average size of the feedback loops, and that the stimulation delays also controlled the feedback loop sizes because they influenced the probability of evoking a spike. Future studies can elaborate on this model by incorporating a learning algorithm that could tune the weights, synaptic/stimulation delays, and ICMS amplitude to control the activity of the biological neurons for different tasks. Furthermore, the model could also incorporate other features that we have shown to exist such as short-term plasticity. The model that I developed to simulate the closed-loop dynamics between biological neurons and artificial IF units could be expanded and used for this purpose.

One of the most common training algorithm for neural networks is backpropagation [164]. In this training paradigm, the weights and biases of an artificial neural network are randomly initialized, after which training data is fed through the network to evaluate its performance. The weights are iteratively changed to reduce the difference between output and target output via gradient descent in weight space. However, backpropagation has limitations that prevent its applicability for the previously described co-processor. First, it employs units with continuous activation functions rather than spiking activity. Second, it cannot be used to adjust the synaptic and stimulation delays to account for the timing of stimulation relative to the background network activity, which strongly impacts the probability of evoking the spikes [39]. Third, it relies on large, labeled training sets, which would prevent the patient from organically learning new tasks

and would require them to have the co-processor retrained and calibrated for each task. Finally, large collections of labeled training data would have to be generated for each new task which would be impossible to do in a damaged brain since the correct patterns of activity for successful task performance would be unknown. For these reasons, a reinforcement learning algorithm would be a better approach. Reinforcement learning algorithms can explore a variety of parameters of the SNN such as the synaptic weights and delays that promote reward. Furthermore, reinforcement learning algorithms can enable the SNN to manipulate its connectivity over time to learn new information [60], [79]. Finally, future implementations of the co-processor could also conceivably record the activity of dopaminergic neurons that encode for reward [165] to instruct the co-processor to strengthen or weaken the artificial connections between the brain and SNN which would eliminate the need to generate new training sets for each novel task.

Towards these goals, the project outlined in Chapter 4 could be elaborated on to first characterize the ability of the biological network to relearn the task as a function of lesion size. Once the size of the lesion required to compromise the ability of the network to relearn the task was found, a collection of simulated artificial excitatory and inhibitory neurons could then be connected to the original set of simulated biological neurons. While the original set of simulated biological neurons would communicate with one another via PSPs, the artificial neurons would trigger ICMS to evoke spikes and inhibit the biological neurons. The probability of evoking spikes would also depend on the value of Δt_{prev}. The reinforcement learning algorithm would randomly explore various connectivities, weights, and synaptic/stimulation delays to search for parameters and connectivities that maximize reward. Dopaminergic neurons could also be added to the model to encode for reward and punishment, which could then be interpreted to guide these changes.

5.3 Implement Reinforcement Learning Algorithms *In Vivo* to Train Neural Co-Processor

One inherent limitation of implementing a neural co-processor *in vivo* is the challenge of manipulating enough neurons to shape behavior. Although novel recording interfaces are being developed that can record many hundreds to thousands of neurons [166], there is no way to extracellularly control neurons in a cell-type specific manner. Furthermore, not only does ICMS activate a broad collection of neurons

transsynaptically [25]–[27], which limits the spatial resolution of neural activation, the stimulation artifacts associated with the targeted manipulation of thousands or more neurons would prevent accurate functioning of the co-processor. However, while new technologies are developed to address these limitations [157], [158], it is important to develop algorithms that are capable of training neural co-processors to manipulate neuronal activity in a task-specific manner to assist the user with the performance of a task.

To this end, neuroprosthetic tasks would be an ideal platform to test such algorithms. Neuroprosthetic tasks enable the researcher to assign a causal neuron-behavior relationship to a group of neurons, such that increases or decreases in the firing rates of the selected neurons can control an external device such as a cursor or robotic limb [167]. The small number of neurons that can be arbitrarily chosen to perform these tasks provides an excellent platform to test co-adaptation algorithms *in vivo*. However, rather than repair damage from a lesion, which would necessitate the use of a co-processor that could interface with a substantial number of neurons, the goal of this co-processor would be neural augmentation, in which the rate at which the NHP learned the neuroprosthetic task was expedited. For this system, task-specific inputs would be input into a multilayer SNN reflecting the relative positions of the cursor and target. The outputs of the SNN would trigger ICMS at specific delays to evoke spikes and/or inhibit the neurons involved in the task. It has been well-established that when an animal learns a novel task, including neuroprosthetic tasks, they increase the firing rate variability of the neurons to search for the correct patterns of neuronal activity that lead to reward [167]–[169]. Therefore, the co-processor would utilize the previously developed reinforcement learning algorithm to tune the weights, synaptic/stimulation delays, and ICMS amplitude to help the animal find the correct patterns of activity for successful task performance. For the first implementation of this project, the outcome of each trial (reward or punishment) would be given to the co-processor, which could then implement the modifications to its weights, synaptic/stimulation delays, and ICMS amplitude. By assisting the NHPs with learning the task, both the NHPs and the co-processor would be required to continuously adapt with one another to jointly maximize reward. As the NHPs improved their task proficiencies, the co-processor could begin to minimize its contributions by lowering the amount of stimulation it delivers, until the NHPs completely learn the task independently of the co-processor. \

Chapter 6. APPENDIX A:

Supplementary Material for Chapter 3

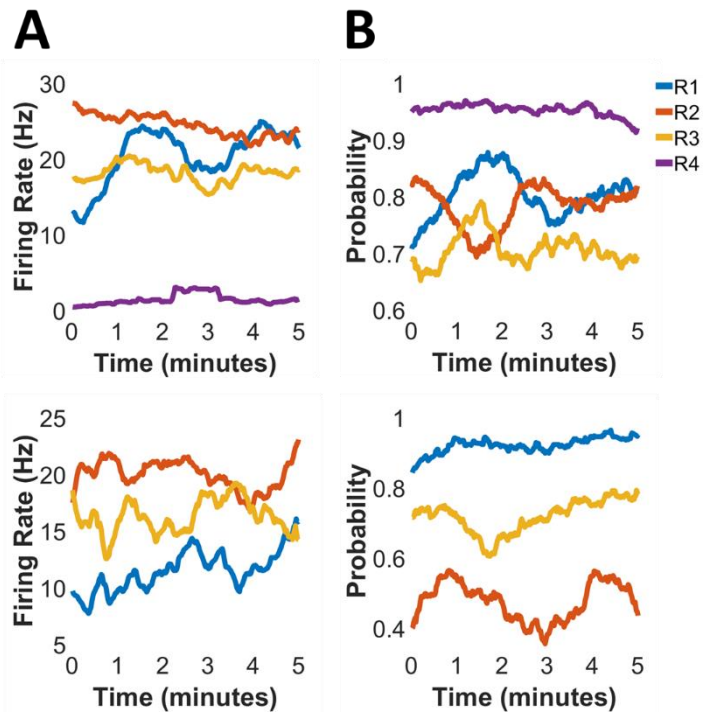

Supplementary Figure 3.6.1 Non-Stationarity in Firing Rate and Stimulation Responses

(A) Changes in firing rates for R1 – R4 throughout the Pre epoch (for networks in Figure 3.4B and Figure 3.4D). Calculated by measuring the firing rate in 1 second bins and performing a moving average of 60 seconds. (B) Probability of evoking spikes in R1 – R4 through the OL epoch (for networks in Figure 3.4B and Figure 3.4D). Calculated by measuring the probability in 1 second bins and performing a moving average of 60 seconds.

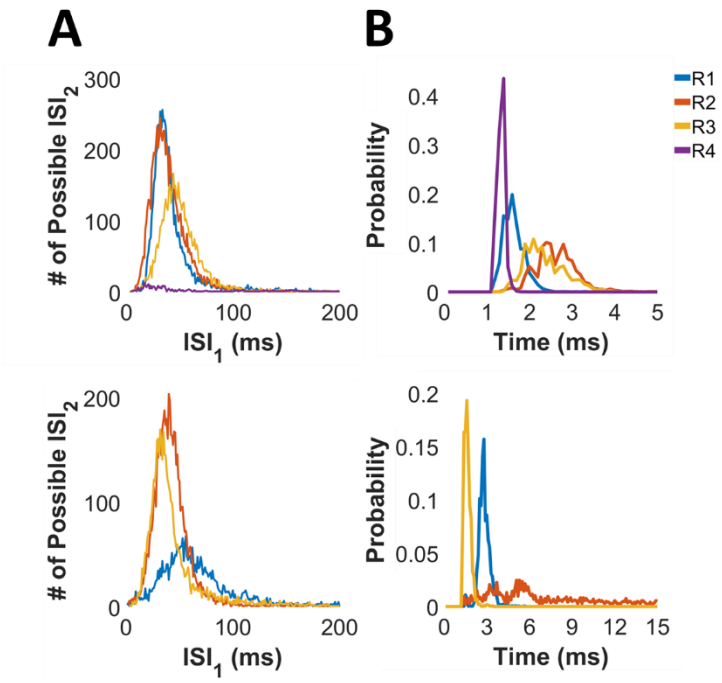

Supplementary Figure 3.6.2 Training the Markov Model and Stimulation Response Model

(A) Number of possible ISI_2's for every possible ISI_1 is a unimodal distribution for R1 – R4 (for networks in Figure 3.4B and Figure 3.4D). (B) Evoked spike timing for R1 – R4 measured during OL epoch (for networks in Figure 3.4B and Figure 3.4D).

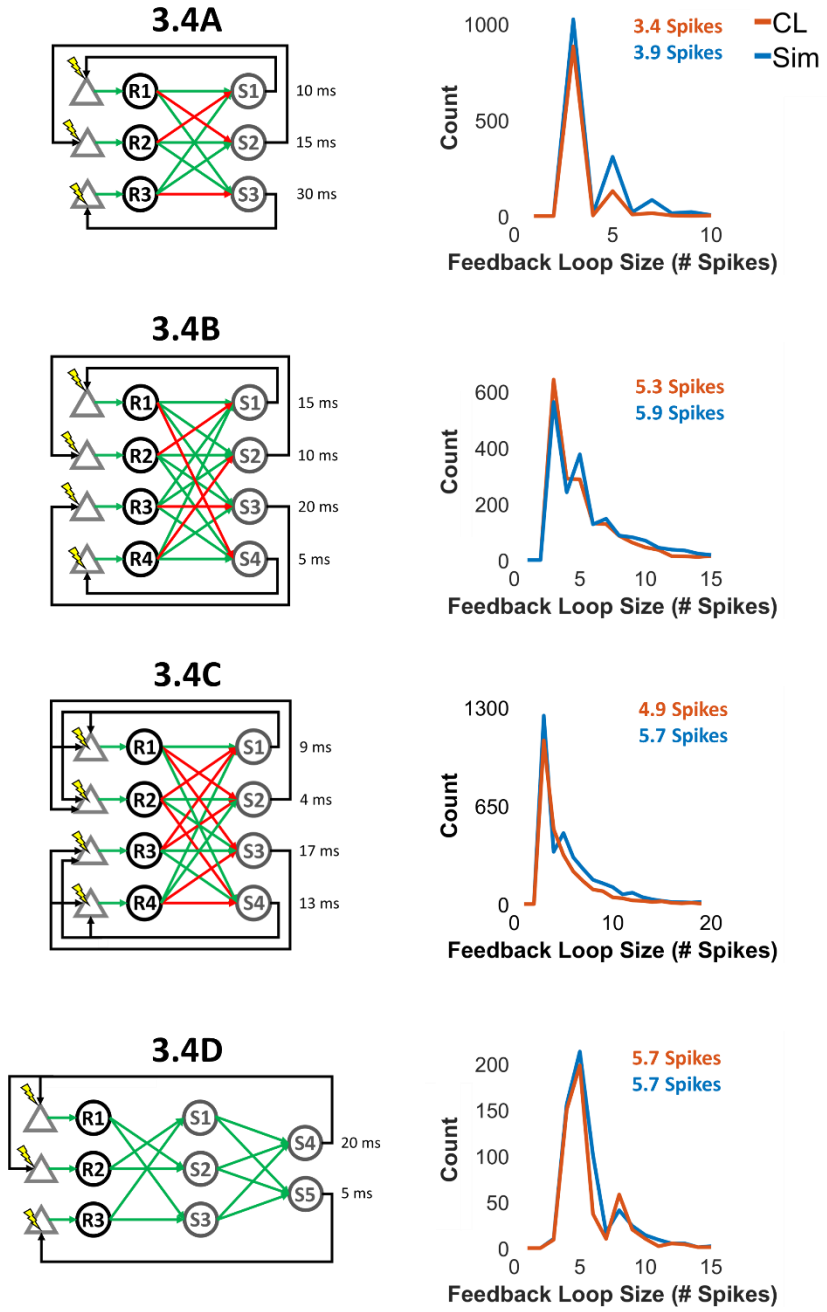

Supplementary Figure 3.6.3 Comparison of feedback loop sizes between the CL epoch and simulations

Number on top of each IFNN show the corresponding figure. Numbers shown in top right of plots are the average feedback loop size for the CL epoch and simulation. In general, the simulations had larger feedback loop sizes, likely due to false negatives in CL epoch, as well as imperfect rebound spike distributions from the OL epoch.

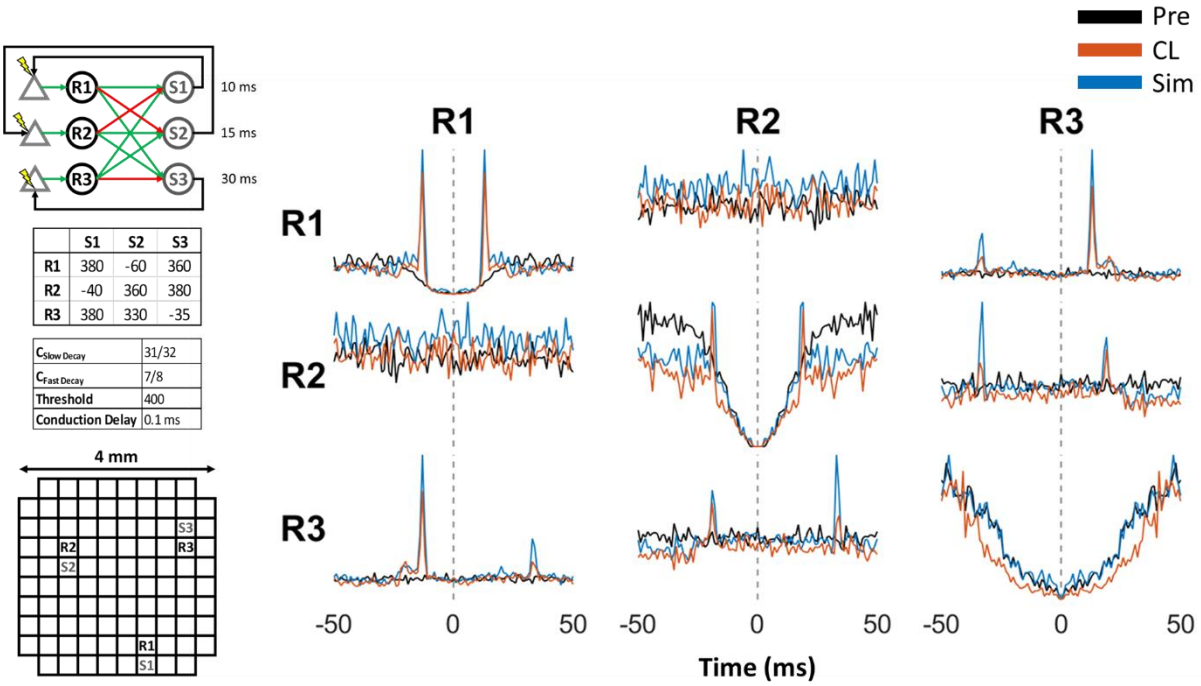

Supplementary Figure 3.6.4 Comparing Spike Dynamics Between Pre Epoch, Closed-Loop Epoch, and Simulations for IFNNs with three Input Neurons and three Output IF Units tested in monkey K

(Left) Network architecture (green and red arrows show excitatory and inhibitory connections, respectively), weights and IFNN parameters, and location of recording/stimulation channels on Utah Array. (Right) Comparison of auto-, and cross-correlograms of spike times during Pre epoch, CL epoch, and simulation.

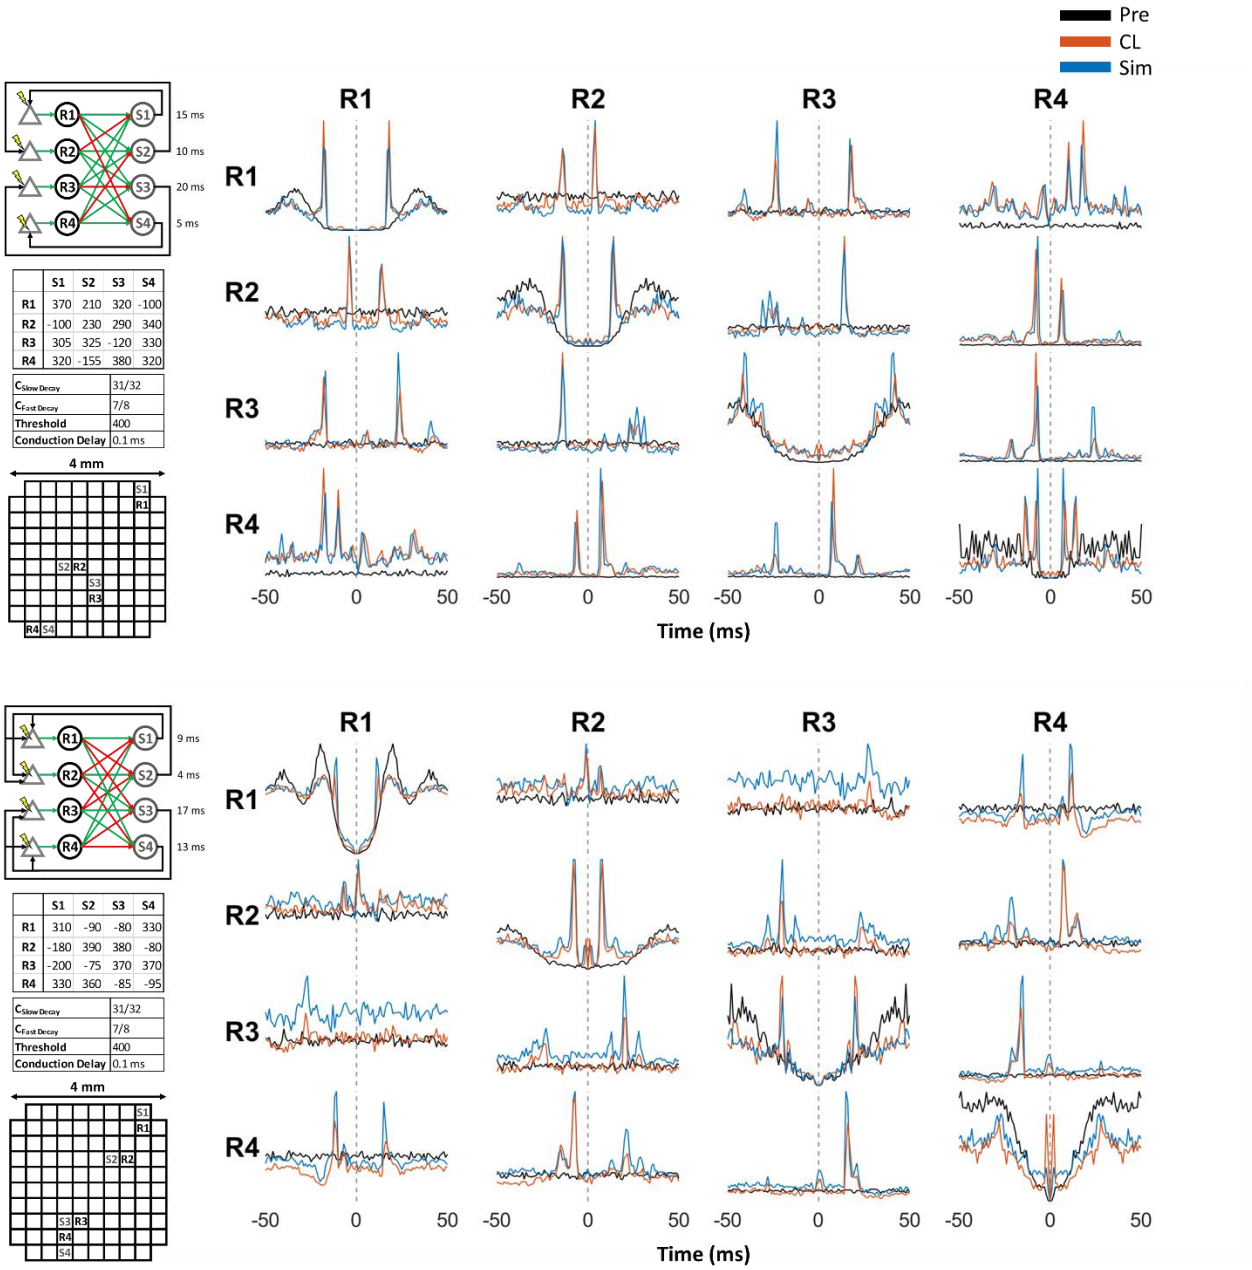

	S1	S2	S3	S4
R1	370	210	320	-100
R2	-100	230	290	340
R3	305	325	-120	330
R4	320	-155	380	320

$C_{Slow\ Decay}$	31/32
$C_{Fast\ Decay}$	7/8
Threshold	400
Conduction Delay	0.1 ms

	S1	S2	S3	S4
R1	310	-90	-80	330
R2	-180	390	380	-80
R3	-200	-75	370	370
R4	330	360	-85	-95

$C_{Slow\ Decay}$	31/32
$C_{Fast\ Decay}$	7/8
Threshold	400
Conduction Delay	0.1 ms

Supplementary Figure 3.6.5 Comparing Spike Dynamics Between Pre Epoch, Closed-Loop Epoch, and Simulations for IFNNs with Four Input Neurons and Four Output IF Units

(Left) Network architectures (green and red arrows show excitatory and inhibitory connections, respectively), weights and IFNN parameters, and location of recording/stimulation channels on Utah Array. (Right) Comparison of auto-, and cross-correlograms of spike times during Pre epoch, CL epoch, and simulation.

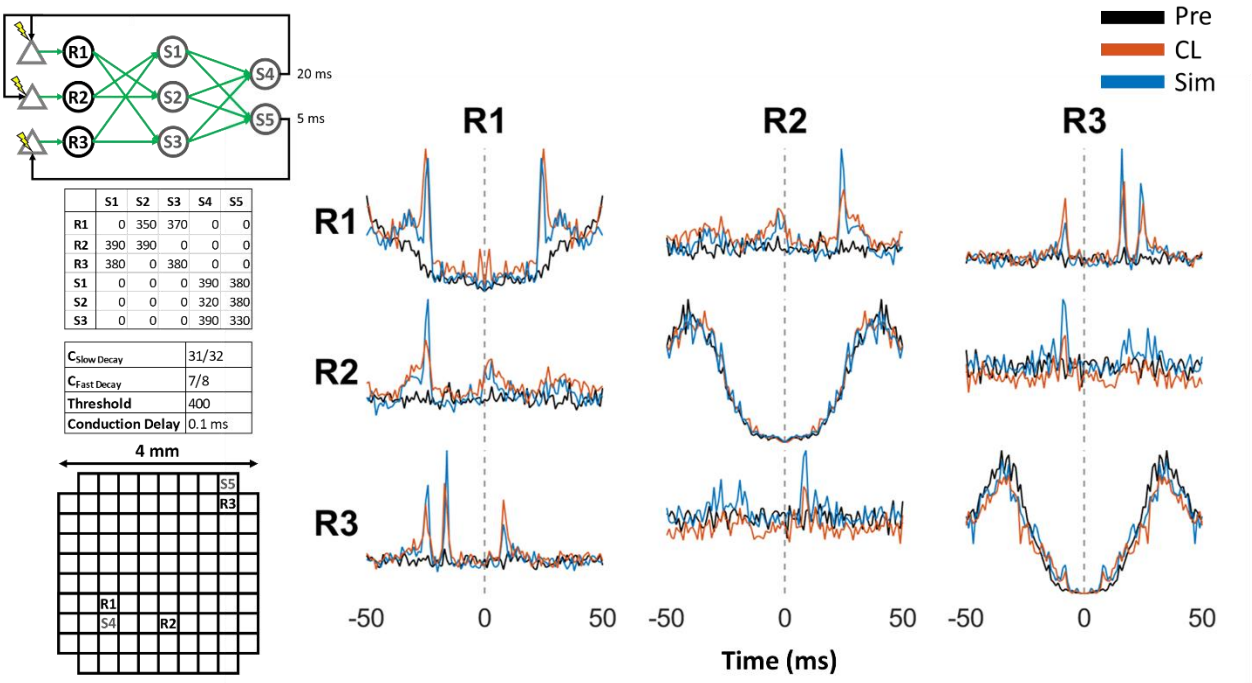

Supplementary Figure 3.6.6 Comparing Spike Dynamics Between Pre Epoch, Closed-Loop Epoch, and Simulations for IFNN with a Hidden Layer

(Left) Network architecture (green arrows show excitatory connections), weights and IFNN parameters, and location of recording/stimulation channels on Utah Array. (Right) Comparison of auto-, and cross-correlograms of spike times during Pre epoch, CL epoch, and simulation.

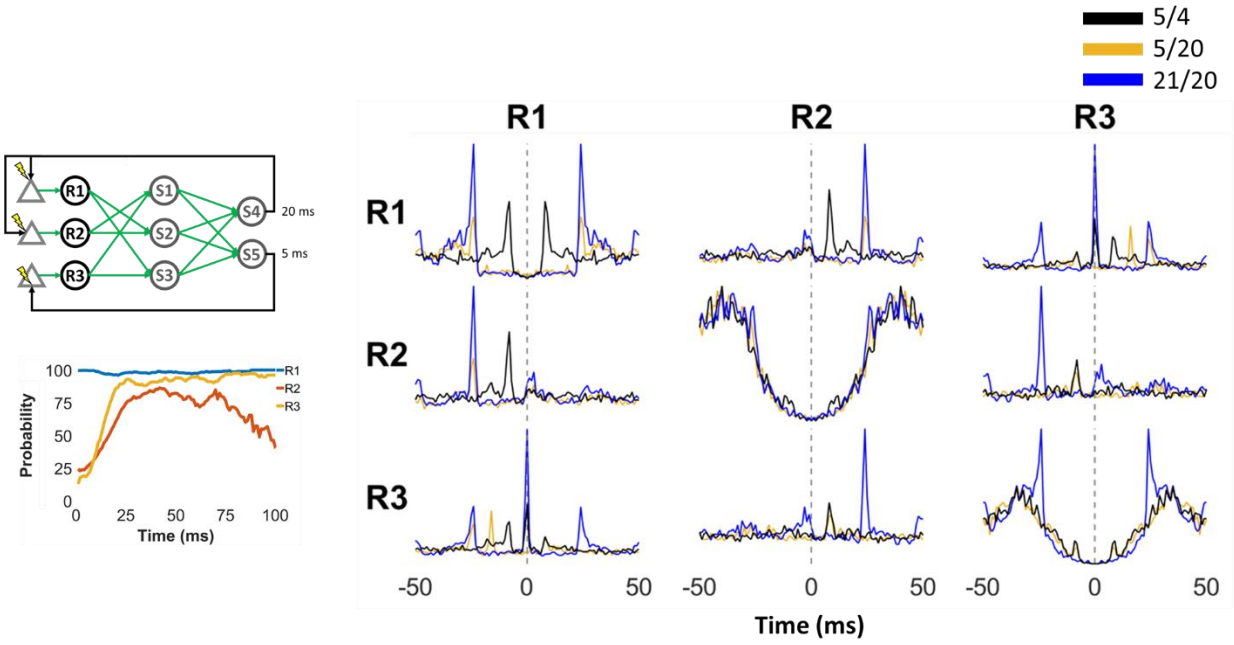

Supplementary Figure 3.6.7 Comparing Simulated Spike Dynamics Between IFNNs with Different Stimulation Delays

(Left) IFNN architecture and probability of evoking a spike in all three input neurons as a function of the delay between their previous spike times and stimulation onsets. (Right) Auto-, and cross-correlograms of R1-R3 for each of the three simulations.

Increase in Magnitude	Frequency	Trough / Peak	# Outputs / # Inputs
R1 → S1	↑		↑
S1 → S3		↓	
S3 → S3		↑	
S3 → S2	↓		
PSP Fall Time Constant	↓		↑

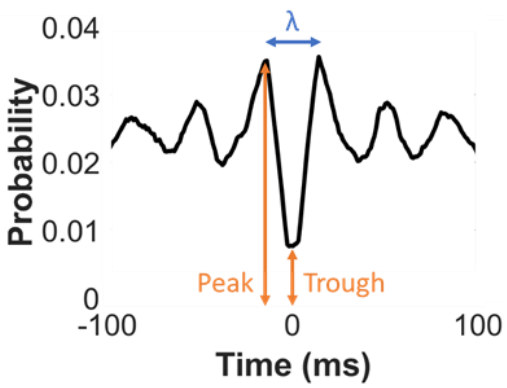

Supplementary Figure 3.6.8 Controlling the Rhythmic Outputs

(left) Relationship between increases in magnitude of weights and PSP fall time constant, and the frequency of the rhythmic outputs (1/λ), ratio of trough / peak height, and ratio of the number of output IF spikes (S1 and S2) / input spikes (R1 and R2). Frequency and ratio of trough / peak height were measured using the cross-correlograms of S1 and S2. Spikes from the Pre epoch were used to change the weights and PSP fall time constant to generate the rhythms.

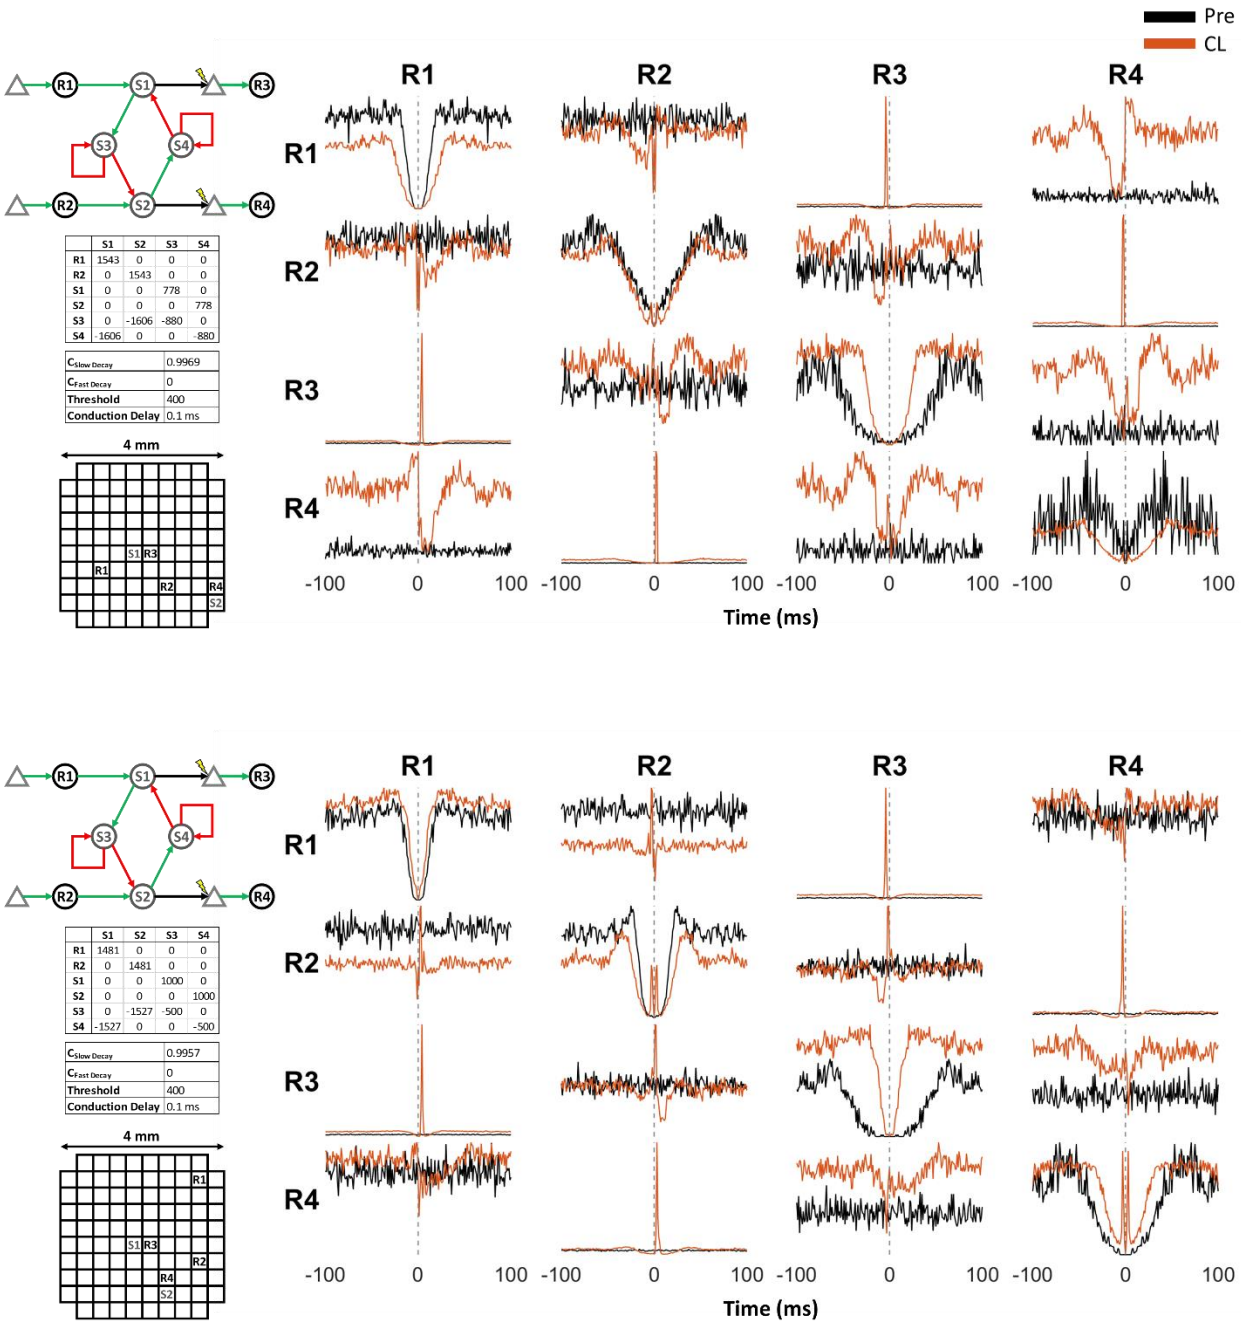

Supplementary Figure 3.6.9 Generating Rhythmic Activity in Biological Neurons. IFNN with reciprocal inhibition in monkey J

(Left) Network architectures (green and red arrows show excitatory and inhibitory connections, respectively), weights and IFNN parameters, and location of recording/stimulation channels on Utah Array. (Right) Comparison of auto-, and cross-correlograms of spike times during Pre epoch and CL epoch. Cross-correlogram between R3 and R4 show generated rhythms between the output neurons.

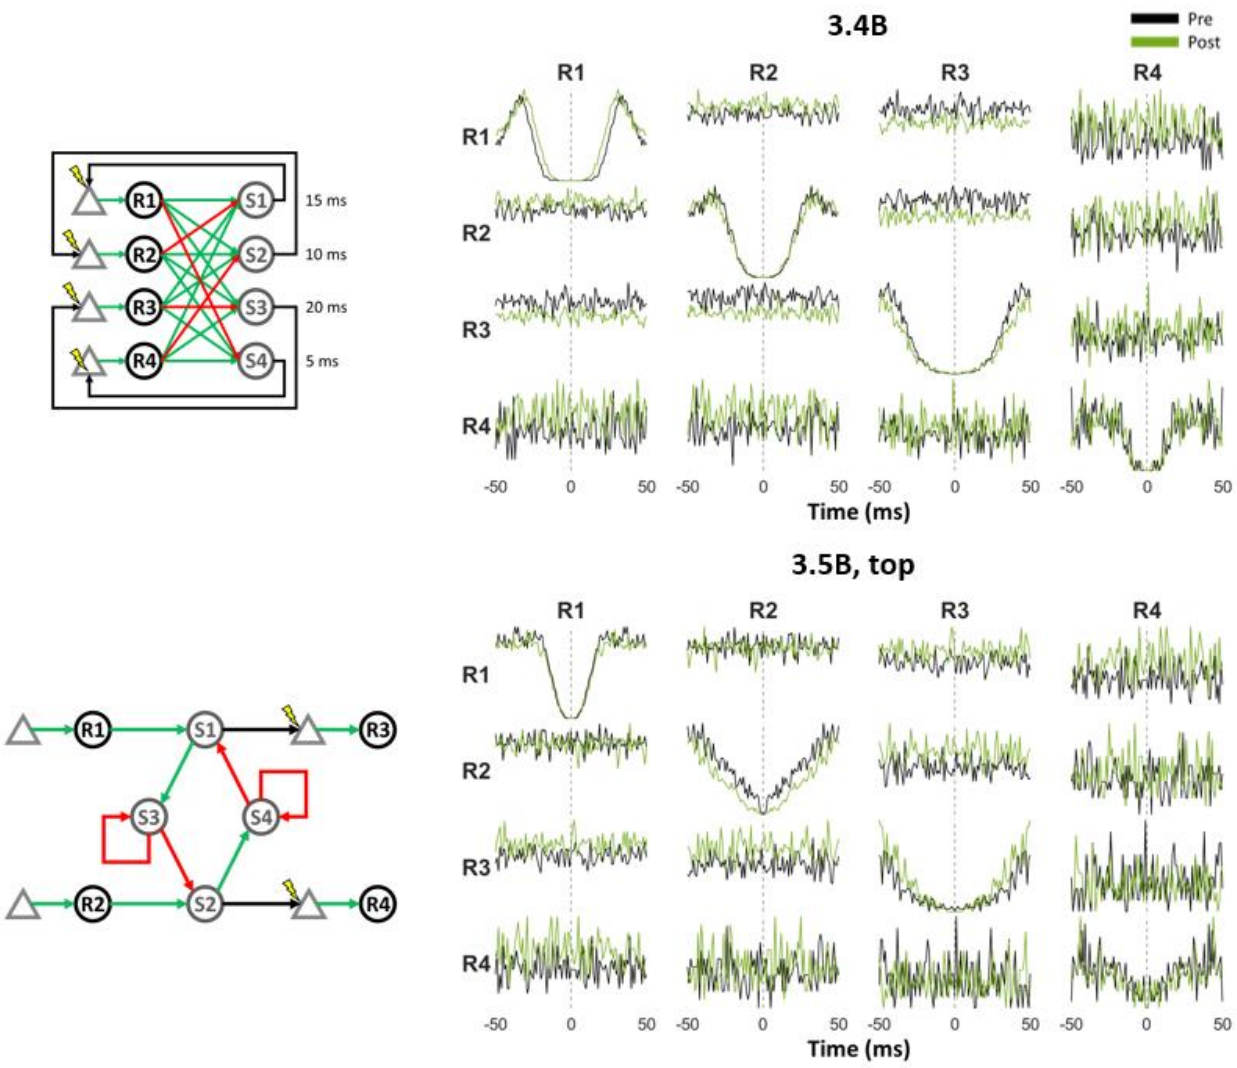

Supplementary Figure 3.6.10 Firing Characteristics During Pre and Post Epochs

(Left) IFNN architectures of both experiments shown. (Right) Auto-, and cross-correlograms of neurons during Pre and Post epochs for IFNNs shown in Figure 3.4B and Figure 3.5B, top.

Chapter 7. BIBLIOGRAPHY

[1] V. L. Feigin *et al.*, "Global, regional, and national burden of neurological disorders, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016," *The Lancet Neurology*, vol. 18, no. 5, pp. 459–480, May 2019, doi: 10.1016/S1474-4422(18)30499-X.

[2] V. L. Feigin *et al.*, "The global burden of neurological disorders: translating evidence into policy," *The Lancet Neurology*, vol. 19, no. 3, pp. 255–265, Mar. 2020, doi: 10.1016/S1474-4422(19)30411-9.

[3] J. Emberson *et al.*, "Effect of treatment delay, age, and stroke severity on the effects of intravenous thrombolysis with alteplase for acute ischaemic stroke: a meta-analysis of individual patient data from randomised trials," *The Lancet*, vol. 384, no. 9958, pp. 1929–1935, Nov. 2014, doi: 10.1016/S0140-6736(14)60584-5.

[4] R. A. Patel and C. J. White, "Acute ischemic stroke treatment: State of the art," *Vasc Med*, vol. 16, no. 1, pp. 19–28, Feb. 2011, doi: 10.1177/1358863X10382945.

[5] C. Laredo *et al.*, "Prognostic Significance of Infarct Size and Location: The Case of Insular Stroke," *Sci Rep*, vol. 8, p. 9498, Jun. 2018, doi: 10.1038/s41598-018-27883-3.

[6] A. Yazdan-Shahmorad, D. B. Silversmith, V. Kharazia, and P. N. Sabes, "Targeted cortical reorganization using optogenetics in non-human primates," *eLife*, vol. 7, p. e31034, doi: 10.7554/eLife.31034.

[7] A. Jackson, J. Mavoori, and E. E. Fetz, "Long-term motor cortex plasticity induced by an electronic neural implant," *Nature*, vol. 444, no. 7115, pp. 56–60, Nov. 2006, doi: 10.1038/nature05226.

[8] S. C. Seeman, B. J. Mogen, E. E. Fetz, and S. I. Perlmutter, "Paired Stimulation for Spike-Timing-Dependent Plasticity in Primate Sensorimotor Cortex," *J. Neurosci.*, vol. 37, no. 7, pp. 1935–1949, Feb. 2017, doi: 10.1523/JNEUROSCI.2046-16.2017.

[9] I. Rembado, S. Zanos, and E. E. Fetz, "Cycle-Triggered Cortical Stimulation during Slow Wave Sleep Facilitates Learning a BMI Task: A Case Report in a Non-Human Primate," *Front. Behav. Neurosci.*, vol. 11, Apr. 2017, doi: 10.3389/fnbeh.2017.00059.

[10] T. A. Jones and D. L. Adkins, "Motor System Reorganization After Stroke: Stimulating and Training Toward Perfection," *Physiology*, vol. 30, no. 5, pp. 358–370, Sep. 2015, doi: 10.1152/physiol.00014.2015.

[11] K. P. Revill, M. W. Haut, S. R. Belagaje, F. Nahab, D. Drake, and C. M. Buetefisch, "Hebbian-Type Primary Motor Cortex Stimulation: A Potential Treatment of Impaired Hand Function in Chronic Stroke Patients," *Neurorehabil Neural Repair*, vol. 34, no. 2, pp. 159–171, Feb. 2020, doi: 10.1177/1545968319899911.

[12] K. A. Tennant, S. L. Taylor, E. R. White, and C. E. Brown, "Optogenetic rewiring of thalamocortical circuits to restore function in the stroke injured brain," *Nat Comms*, vol. 8, p. 15879, Jun. 2017, doi: 10.1038/ncomms15879.

[13] C. Grefkes and N. S. Ward, "Cortical Reorganization After Stroke: How Much and How Functional?," *Neuroscientist*, vol. 20, no. 1, pp. 56–70, Feb. 2014, doi: 10.1177/1073858413491147.

[14] T. H. Murphy and D. Corbett, "Plasticity during stroke recovery: from synapse to behaviour," *Nat Rev Neurosci*, vol. 10, no. 12, pp. 861–872, Dec. 2009, doi: 10.1038/nrn2735.

[15] S. C. Cramer, "Repairing the human brain after stroke: I. Mechanisms of spontaneous recovery," *Ann Neurol.*, vol. 63, no. 3, pp. 272–287, Mar. 2008, doi: 10.1002/ana.21393.

[16] T. Callier, N. W. Brantly, A. Caravelli, and S. J. Bensmaia, "The frequency of cortical microstimulation shapes artificial touch," *Proc Natl Acad Sci USA*, vol. 117, no. 2, pp. 1191–1200, Jan. 2020, doi: 10.1073/pnas.1916453117.

[17] M. Armenta Salas *et al.*, "Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation," *eLife*, vol. 7, p. e32904, Apr. 2018, doi: 10.7554/eLife.32904.

[18] S. N. Flesher *et al.*, “A brain-computer interface that evokes tactile sensations improves robotic arm control,” *Science*, vol. 372, no. 6544, pp. 831–836, May 2021, doi: 10.1126/science.abd0380.
[19] J. L. Collinger, R. A. Gaunt, and A. B. Schwartz, “Progress towards restoring upper limb movement and sensation through intracortical brain-computer interfaces,” *Current Opinion in Biomedical Engineering*, vol. 8, pp. 84–92, Dec. 2018, doi: 10.1016/j.cobme.2018.11.005.
[20] C. J. Hartmann, S. Fliegen, S. J. Groiss, L. Wojtecki, and A. Schnitzler, “An update on best practice of deep brain stimulation in Parkinson’s disease,” *Ther Adv Neurol Disord*, vol. 12, p. 1756286419838096, Jan. 2019, doi: 10.1177/1756286419838096.
[21] A. M. Lozano *et al.*, “Deep brain stimulation: current challenges and future directions,” *Nat Rev Neurol*, vol. 15, no. 3, pp. 148–160, Mar. 2019, doi: 10.1038/s41582-018-0128-2.
[22] T. M. Herrington, J. J. Cheng, and E. N. Eskandar, “Mechanisms of deep brain stimulation,” *Journal of Neurophysiology*, vol. 115, no. 1, pp. 19–38, Jan. 2016, doi: 10.1152/jn.00281.2015.
[23] M. Parastarfeizabadi and A. Z. Kouzani, “Advances in closed-loop deep brain stimulation devices,” *J NeuroEngineering Rehabil*, vol. 14, no. 1, p. 79, Dec. 2017, doi: 10.1186/s12984-017-0295-1.
[24] M. H. Histed, A. M. Ni, and J. H. R. Maunsell, “Insights into cortical mechanisms of behavior from microstimulation experiments,” *Prog Neurobiol*, vol. 103, pp. 115–130, Apr. 2013, doi: 10.1016/j.pneurobio.2012.01.006.
[25] E. J. Tehovnik, A. S. Tolias, F. Sultan, W. M. Slocum, and N. K. Logothetis, “Direct and Indirect Activation of Cortical Neurons by Electrical Microstimulation,” *Journal of Neurophysiology*, vol. 96, no. 2, pp. 512–521, Aug. 2006, doi: 10.1152/jn.00126.2006.
[26] S. Butovas and C. Schwarz, “Spatiotemporal rffects of microstimulation in rat neocortex: A parametric study using multielectrode recordings,” *Journal of Neurophysiology*, vol. 90, no. 5, pp. 3024–3039, 2003, doi: 10.1152/jn.00245.2003.
[27] M. H. Histed, V. Bonin, and R. C. Reid, “Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation,” *Neuron*, vol. 63, no. 4, pp. 508–522, 2009, doi: 10.1016/j.neuron.2009.07.016.
[28] S. S. Kumar, J. Wülfing, S. Okujeni, J. Boedecker, M. Riedmiller, and U. Egert, “Autonomous Optimization of Targeted Stimulation of Neuronal Networks,” *PLoS Comput Biol*, vol. 12, no. 8, p. e1005054, Aug. 2016, doi: 10.1371/journal.pcbi.1005054.
[29] O. Weihberger, S. Okujeni, J. E. Mikkonen, and U. Egert, “Quantitative examination of stimulus-response relations in cortical networks in vitro,” *Journal of Neurophysiology*, vol. 109, no. 7, pp. 1764–1774, Apr. 2013, doi: 10.1152/jn.00481.2012.
[30] P. Kara, J. S. Pezaris, S. Yurgenson, and R. C. Reid, “The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation,” *Proceedings of the National Academy of Sciences*, vol. 99, no. 25, pp. 16261–16266, Dec. 2002, doi: 10.1073/pnas.242625499.
[31] S. Butovas, S. G. Hormuzdi, H. Monyer, and C. Schwarz, “Effects of electrically coupled inhibitory networks on local neuronal responses to intracortical microstimulation,” *Journal of Neurophysiology*, vol. 96, no. 3, pp. 1227–1236, 2006, doi: 10.1152/jn.01170.2005.
[32] Y. Hao, A. Riehle, and T. G. Brochier, “Mapping horizontal spread of activity in monkey motor cortex using single pulse microstimulation,” *Frontiers in Neural Circuits*, vol. 10, no. December, pp. 1–16, 2016, doi: 10.3389/fncir.2016.00104.
[33] N. J. Michelson, J. R. Eles, A. L. Vazquez, K. A. Ludwig, and T. D. Y. Kozai, “Calcium activation of cortical neurons by continuous electrical stimulation: Frequency dependence, temporal fidelity, and activation density,” *Journal of Neuroscience Research*, vol. 97, no. 5, pp. 620–638, 2019, doi: 10.1002/jnr.24370.
[34] R. S. Zucker and W. G. Regehr, “Short-term synaptic plasticity,” *Annual Review of Physiology*, vol. 64, no. 4, pp. 355–405, 2002, doi: 10.1146/annurev.physiol.64.092501.114547.
[35] W. G. Regehr, “Short-Term Presynaptic Plasticity,” *Cold Spring Harbor Perspectives in Biology*, vol. 4, no. 7, pp. a005702–a005702, Jul. 2012, doi: 10.1101/cshperspect.a005702.

[36] J. Xu, L. He, and L.-G. Wu, “Role of Ca2+ channels in short-term synaptic plasticity,” *Current Opinion in Neurobiology*, vol. 17, no. 3, pp. 352–359, Jun. 2007, doi: 10.1016/j.conb.2007.04.005.
[37] S. Boudkkazi *et al.*, “Release-dependent bariations in synaptic latency: A putative code for short- and long-term synaptic dynamics,” *Neuron*, vol. 56, no. 6, pp. 1048–1060, 2007, doi: 10.1016/j.neuron.2007.10.037.
[38] W. McCulloch and W. Pitts, “A logical calculus of the ideas immanent in nervous activity,” *The bulletin of mathematical biophysics*, p. 19, Dec. 1943.
[39] W. Maass, “Networks of spiking neurons: The third generation of neural network models,” *Neural Networks*, vol. 10, no. 9, pp. 1659–1671, Dec. 1997, doi: 10.1016/S0893-6080(97)00011-7.
[40] F. Rosenblatt, “The perceptron: A probabilistic model for information storage and organization in the brain.,” *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958, doi: 10.1037/h0042519.
[41] M. Minsky and S. A. Papert, *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, 1969. doi: 10.7551/mitpress/11301.001.0001.
[42] T. Szandała, “Review and Comparison of Commonly Used Activation Functions for Deep Neural Networks,” in *Bio-inspired Neurocomputing*, vol. 903, A. K. Bhoi, P. K. Mallick, C.-M. Liu, and V. E. Balas, Eds. Singapore: Springer Singapore, 2021, pp. 203–224. doi: 10.1007/978-981-15-5495-7_11.
[43] O. I. Abiodun *et al.*, “Comprehensive Review of Artificial Neural Network Applications to Pattern Recognition,” *IEEE Access*, vol. 7, pp. 158820–158846, 2019, doi: 10.1109/ACCESS.2019.2945545.
[44] D. I. Perrett, E. T. Rolls, and W. Caan, “Visual Neurones Responsive to Faces in the Monkey Temporal Cortex,” *Experimental Brain Research*.
[45] S. J. Thorpe and M. Imbert, “Biological constraints on connectionist modelling,” in *Connectionism in Perspective*, 1989, pp. 63–92.
[46] E. Rolls and M. Tovee, “Processing speed in the cerebral cortex and the neurophysiology of visual masking,” *Proc. R. Soc. Lond. B*, vol. 257, no. 1348, pp. 9–15, Jul. 1994, doi: 10.1098/rspb.1994.0087.
[47] E. M. Izhikevich, “Which Model to Use for Cortical Spiking Neurons?,” *IEEE Trans. Neural Netw.*, vol. 15, no. 5, pp. 1063–1070, Sep. 2004, doi: 10.1109/TNN.2004.832719.
[48] A. N. Burkitt, “A Review of the Integrate-and-fire Neuron Model: I. Homogeneous Synaptic Input,” *Biol Cybern*, vol. 95, no. 1, pp. 1–19, Jul. 2006, doi: 10.1007/s00422-006-0068-6.
[49] E. M. Izhikevich, “Simple model of spiking neurons,” *IEEE Trans. Neural Netw.*, vol. 14, no. 6, pp. 1569–1572, Nov. 2003, doi: 10.1109/TNN.2003.820440.
[50] S. Cavallari, S. Panzeri, and A. Mazzoni, “Comparison of the dynamics of neural interactions between current-based and conductance-based integrate-and-fire recurrent networks,” *Front. Neural Circuits*, vol. 8, Mar. 2014, doi: 10.3389/fncir.2014.00012.
[51] R. Jolivet, A. Rauch, H.-R. Lüscher, and W. Gerstner, “Predicting spike timing of neocortical pyramidal neurons by simple threshold models,” *J Comput Neurosci*, vol. 21, no. 1, pp. 35–49, Aug. 2006, doi: 10.1007/s10827-006-7074-5.
[52] G. Bellec *et al.*, “A solution to the learning dilemma for recurrent networks of spiking neurons,” *Nat Commun*, vol. 11, no. 1, p. 3625, Dec. 2020, doi: 10.1038/s41467-020-17236-y.
[53] C. Huang, A. Resnik, T. Celikel, and B. Englitz, “Adaptive Spike Threshold Enables Robust and Temporally Precise Neuronal Encoding,” *PLoS Comput Biol*, vol. 12, no. 6, p. e1004984, Jun. 2016, doi: 10.1371/journal.pcbi.1004984.
[54] H. Jang, O. Simeone, B. Gardner, and A. Grüning, “An Introduction to Probabilistic Spiking Neural Networks: Probabilistic Models, Learning Rules, and Applications,” *IEEE Signal Process. Mag.*, vol. 36, no. 6, pp. 64–77, Nov. 2019, doi: 10.1109/MSP.2019.2935234.
[55] L. Kozachkov, J. Tauber, M. Lundqvist, S. L. Brincat, J.-J. Slotine, and E. K. Miller, “Robust Working Memory through Short-Term Synaptic Plasticity,” Neuroscience, preprint, Jan. 2022. doi: 10.1101/2022.01.09.475558.

[56] E. L. Lameu *et al.*, “Short-term and spike-timing-dependent plasticity facilitate the formation of modular neural networks,” *Communications in Nonlinear Science and Numerical Simulation*, vol. 96, p. 105689, May 2021, doi: 10.1016/j.cnsns.2020.105689.
[57] J. Schemmel, A. Grübl, K. Meier, and E. Mueller, “Implementing synaptic plasticity in a VLSI spiking neural network model,” 2006.
[58] L. Tauffer and A. Kumar, “Short-Term Synaptic Plasticity Makes Neurons Sensitive to the Distribution of Presynaptic Population Firing Rates,” *eNeuro*, vol. 8, no. 2, p. ENEURO.0297-20.2021, Mar. 2021, doi: 10.1523/ENEURO.0297-20.2021.
[59] L. Leng *et al.*, “Spiking neurons with short-term synaptic plasticity form superior generative networks,” *Sci Rep*, vol. 8, no. 1, p. 10651, Dec. 2018, doi: 10.1038/s41598-018-28999-2.
[60] D. Kappel, R. Legenstein, S. Habenschuss, M. Hsieh, and W. Maass, “A Dynamic Connectome Supports the Emergence of Stable Computational Function of Neural Circuits through Reward-Based Learning,” *eNeuro*, vol. 5, no. 2, p. ENEURO.0301-17.2018, Mar. 2018, doi: 10.1523/ENEURO.0301-17.2018.
[61] Z. Bing, C. Meschede, F. Röhrbein, K. Huang, and A. C. Knoll, “A Survey of Robotics Control Based on Learning-Inspired Spiking Neural Networks,” *Front Neurorobot*, vol. 12, Jul. 2018, doi: 10.3389/fnbot.2018.00035.
[62] L. N. Long, “An adaptive spiking neural network with Hebbian learning,” in *2011 IEEE Workshop on Evolving and Adaptive Intelligent Systems (EAIS)*, Paris, France, Apr. 2011, pp. 17–23. doi: 10.1109/EAIS.2011.5945923.
[63] R. A. Koene and M. E. Hasselmo, “An Integrate-and-fire Model of Prefrontal Cortex Neuronal Activity during Performance of Goal-directed Decision Making,” *Cerebral Cortex*, vol. 15, no. 12, pp. 1964–1981, Dec. 2005, doi: 10.1093/cercor/bhi072.
[64] A. Tavanaei, M. Ghodrati, S. R. Kheradpisheh, T. Masquelier, and A. S. Maida, “Deep Learning in Spiking Neural Networks,” *Neural Networks*, vol. 111, pp. 47–63, Mar. 2019, doi: 10.1016/j.neunet.2018.12.002.
[65] G. Tang, N. Kumar, R. Yoo, and K. P. Michmizos, “Deep Reinforcement Learning with Population-Coded Spiking Neural Network for Continuous Control,” *arXiv:2010.09635 [cs]*, Oct. 2020, Accessed: Jan. 14, 2021. [Online]. Available: http://arxiv.org/abs/2010.09635
[66] C. Pokorny, M. J. Ison, A. Rao, R. Legenstein, C. Papadimitriou, and W. Maass, “STDP Forms Associations between Memory Traces in Networks of Spiking Neurons,” *Cerebral Cortex*, vol. 30, no. 3, pp. 952–968, Mar. 2020, doi: 10.1093/cercor/bhz140.
[67] Z. Bing, I. Baumann, Z. Jiang, K. Huang, C. Cai, and A. Knoll, “Supervised Learning in SNN via Reward-Modulated Spike-Timing-Dependent Plasticity for a Target Reaching Vehicle,” *Front. Neurorobot.*, vol. 13, p. 18, May 2019, doi: 10.3389/fnbot.2019.00018.
[68] A. Bouganis and M. Shanahan, “Training a spiking neural network to control a 4-DoF robotic arm based on Spike Timing-Dependent Plasticity,” in *The 2010 International Joint Conference on Neural Networks (IJCNN)*, Barcelona, Spain, Jul. 2010, pp. 1–8. doi: 10.1109/IJCNN.2010.5596525.
[69] C. Lee, P. Panda, G. Srinivasan, and K. Roy, “Training Deep Spiking Convolutional Neural Networks With STDP-Based Unsupervised Pre-training Followed by Supervised Fine-Tuning,” *Front. Neurosci.*, vol. 12, p. 435, Aug. 2018, doi: 10.3389/fnins.2018.00435.
[70] P. U. Diehl and M. Cook, “Unsupervised learning of digit recognition using spike-timing-dependent plasticity,” *Front. Comput. Neurosci.*, vol. 9, Aug. 2015, doi: 10.3389/fncom.2015.00099.
[71] V. Itskov, D. Hansel, and M. Tsodyks, “Short-Term Facilitation may Stabilize Parametric Working Memory Trace,” *Front Comput Neurosci*, vol. 5, p. 40, Oct. 2011, doi: 10.3389/fncom.2011.00040.
[72] A. Gupta, “Organizing Principles for a Diversity of GABAergic Interneurons and Synapses in the Neocortex,” *Science*, vol. 287, no. 5451, pp. 273–278, Jan. 2000, doi: 10.1126/science.287.5451.273.
[73] H. Markram *et al.*, “Reconstruction and Simulation of Neocortical Microcircuitry,” *Cell*, vol. 163, no. 2, pp. 456–492, Oct. 2015, doi: 10.1016/j.cell.2015.09.029.

[74] P. Panda and K. Roy, “Unsupervised regenerative learning of hierarchical features in Spiking Deep Networks for object recognition,” in *2016 International Joint Conference on Neural Networks (IJCNN)*, Vancouver, BC, Canada, Jul. 2016, pp. 299–306. doi: 10.1109/IJCNN.2016.7727212.
[75] W. Ponghiran, G. Srinivasan, and K. Roy, “Reinforcement Learning with Low-Complexity Liquid State Machines,” *arXiv:1906.01695 [cs, stat]*, Jun. 2019, Accessed: Apr. 26, 2021. [Online]. Available: http://arxiv.org/abs/1906.01695
[76] J. Kaiser *et al.*, “Embodied Synaptic Plasticity With Online Reinforcement Learning,” *Front. Neurorobot.*, vol. 13, p. 81, Oct. 2019, doi: 10.3389/fnbot.2019.00081.
[77] K. El-Laithy and M. Bogdan, “A Reinforcement Learning Framework for Spiking Networks with Dynamic Synapses,” *Computational Intelligence and Neuroscience*, vol. 2011, pp. 1–12, 2011, doi: 10.1155/2011/869348.
[78] S. Aenugu, “Training spiking neural networks for reinforcement learning,” p. 33.
[79] Y. Chen, Z. Yu, W. Fang, T. Huang, and Y. Tian, “Pruning of Deep Spiking Neural Networks through Gradient Rewiring,” in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, Montreal, Canada, Aug. 2021, pp. 1713–1721. doi: 10.24963/ijcai.2021/236.
[80] K. D. Carlson, M. Richert, N. Dutt, and J. L. Krichmar, “Biologically plausible models of homeostasis and STDP: Stability and learning in spiking neural networks,” in *The 2013 International Joint Conference on Neural Networks (IJCNN)*, Aug. 2013, pp. 1–8. doi: 10.1109/IJCNN.2013.6706961.
[81] J. Zierenberg, J. Wilting, and V. Priesemann, “Homeostatic Plasticity and External Input Shape Neural Network Dynamics,” *Phys. Rev. X*, vol. 8, no. 3, p. 031018, Jul. 2018, doi: 10.1103/PhysRevX.8.031018.
[82] S. N. Flesher *et al.*, “Intracortical microstimulation of human somatosensory cortex,” *Sci. Transl. Med.*, vol. 8, no. 361, Oct. 2016, doi: 10.1126/scitranslmed.aaf8083.
[83] K. Hartmann *et al.*, “Embedding a Panoramic Representation of Infrared Light in the Adult Rat Somatosensory Cortex through a Sensory Neuroprosthesis,” *J. Neurosci.*, vol. 36, no. 8, pp. 2406–2424, Feb. 2016, doi: 10.1523/JNEUROSCI.3285-15.2016.
[84] A. Jackson and E. E. Fetz, “Interfacing with the Computational Brain,” *IEEE Trans Neural Syst Rehabil Eng*, vol. 19, no. 5, pp. 534–541, Oct. 2011, doi: 10.1109/TNSRE.2011.2158586.
[85] M. A. Lebedev and M. A. L. Nicolelis, “Brain-Machine Interfaces: From Basic Science to Neuroprostheses and Neurorehabilitation,” *Physiological Reviews*, vol. 97, no. 2, pp. 767–837, Apr. 2017, doi: 10.1152/physrev.00027.2016.
[86] T. J. Sejnowski, P. S. Churchland, and J. A. Movshon, “Putting big data to good use in neuroscience,” *Nat Neurosci*, vol. 17, no. 11, pp. 1440–1441, Nov. 2014, doi: 10.1038/nn.3839.
[87] J. B. Ranck, “Which elements are excited in electrical stimulation of mammalian central nervous system: A review,” *Brain Research*, vol. 98, no. 3, pp. 417–440, Nov. 1975, doi: 10.1016/0006-8993(75)90364-9.
[88] S. D. Stoney, W. D. Thompson, and H. Asanuma, “Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current.,” *Journal of Neurophysiology*, vol. 31, no. 5, pp. 659–669, Sep. 1968, doi: 10.1152/jn.1968.31.5.659.
[89] B. Gustafsson and E. Jankowska, “Direct and indirect activation of nerve cells by electrical pulses applied extracellularly.,” *J Physiol*, vol. 258, no. 1, pp. 33–61, Jun. 1976.
[90] Y. Hao, A. Riehle, and T. G. Brochier, “Mapping Horizontal Spread of Activity in Monkey Motor Cortex Using Single Pulse Microstimulation,” *Front. Neural Circuits*, vol. 10, Dec. 2016, doi: 10.3389/fncir.2016.00104.
[91] R. P. Lesser, H. W. Lee, W. R. S. Webber, B. Prince, N. E. Crone, and D. L. Miglioretti, “Short-term variations in response distribution to cortical stimulation,” *Brain*, vol. 131, no. 6, pp. 1528–1539, Jun. 2008, doi: 10.1093/brain/awn044.
[92] N. K. Logothetis *et al.*, “The effects of electrical microstimulation on cortical signal propagation,” *Nat Neurosci*, vol. 13, no. 10, pp. 1283–1291, Oct. 2010, doi: 10.1038/nn.2631.

[93] C. C. McIntyre and W. M. Grill, “Selective Microstimulation of Central Nervous System Neurons,” *Annals of Biomedical Engineering*, vol. 28, no. 3, pp. 219–233, Mar. 2000, doi: 10.1114/1.262.
[94] N. J. Berman, R. J. Douglas, K. A. Martin, and D. Whitteridge, “Mechanisms of inhibition in cat visual cortex.,” *J Physiol*, vol. 440, pp. 697–722, 1991.
[95] M. C. Dadarlat, Y. Sun, and M. P. Stryker, “Widespread activation of awake mouse cortex by electrical stimulation,” *Int IEEE EMBS Conf Neural Eng*, vol. 2019, pp. 1113–1117, Mar. 2019, doi: 10.1109/NER.2019.8716956.
[96] N. J. Michelson, J. R. Eles, A. L. Vazquez, K. A. Ludwig, and T. D. Kozai, “Calcium activation of cortical neurons by continuous electrical stimulation: Frequency-dependence, temporal fidelity and activation density,” *J Neurosci Res*, vol. 97, no. 5, pp. 620–638, May 2019, doi: 10.1002/jnr.24370.
[97] L. F. Abbott and W. G. Regehr, “Synaptic computation,” *Nature*, vol. 431, no. 7010, pp. 796–803, Oct. 2004, doi: 10.1038/nature03010.
[98] A. Citri and R. C. Malenka, “Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms,” *Neuropsychopharmacol*, vol. 33, no. 1, pp. 18–41, Jan. 2008, doi: 10.1038/sj.npp.1301559.
[99] R. S. Zucker and W. G. Regehr, “Short-Term Synaptic Plasticity,” *Annual Review of Physiology*, vol. 64, no. 1, pp. 355–405, 2002, doi: 10.1146/annurev.physiol.64.092501.114547.
[100] S. Borchers, M. Himmelbach, N. Logothetis, and H.-O. Karnath, “Direct electrical stimulation of human cortex — the gold standard for mapping brain functions?,” *Nat Rev Neurosci*, vol. 13, no. 1, pp. 63–70, Jan. 2012, doi: 10.1038/nrn3140.
[101] D. M. Griffin, H. M. Hudson, A. Belhaj-Saïf, and P. D. Cheney, “Hijacking Cortical Motor Output with Repetitive Microstimulation,” *J Neurosci*, vol. 31, no. 37, pp. 13088–13096, Sep. 2011, doi: 10.1523/JNEUROSCI.6322-10.2011.
[102] M. C. Dadarlat, Y. Sun, and M. P. Stryker, “Widespread activation of awake mouse cortex by electrical stimulation,” *International IEEE/EMBS Conference on Neural Engineering, NER*, vol. 2019-March, pp. 1113–1117, 2019, doi: 10.1109/NER.2019.8716956.
[103] B. W. Connors and M. J. Gutnick, “Intrinsic firing patterns of diverse neocortical neurons,” *Trends in Neurosciences*, vol. 13, no. 3, pp. 99–104, 1990.
[104] D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince, “Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex,” *Journal of Neurophysiology*, vol. 54, no. 4, pp. 782–806, 1985, doi: 10.1152/jn.1985.54.4.782.
[105] A. T. Hussin, J. A. Boychuk, A. R. Brown, Q. J. Pittman, and G. Campbell Teskey, “Intracortical microstimulation (ICMS) activates motor cortex layer 5 pyramidal neurons mainly transsynaptically,” *Brain Stimulation*, vol. 8, no. 4, pp. 742–750, 2015, doi: 10.1016/j.brs.2015.03.003.
[106] P. C. Klink, B. Dagnino, M. A. Gariel-Mathis, and P. R. Roelfsema, “Distinct feedforward and feedback effects of microstimulation in visual cortex reveal neural mechanisms of texture segregation,” *Neuron*, vol. 95, no. 1, pp. 209-220.e3, 2017, doi: 10.1016/j.neuron.2017.05.033.
[107] J. S. Isaacson and M. Scanziani, “How inhibition shapes cortical activity,” *Neuron*, vol. 72, no. 2, pp. 231–243, 2011, doi: 10.1016/j.neuron.2011.09.027.
[108] B. Rudy, G. Fishell, S. Lee, and J. Hjerling-Leffler, “Three groups of interneurons account for nearly 100 % of Neocortical GABAergic Neurons,” *Dev Neurobiol*, vol. 71, no. 1, pp. 45–61, 2013, doi: 10.1002/dneu.20853.Three.
[109] R. Tremblay, S. Lee, and B. Rudy, “GABAergic interneurons in the neocortex: From cellular properties to circuits,” *Neuron*, vol. 91, no. 2, pp. 260–292, 2016, doi: 10.1016/j.neuron.2016.06.033.GABAergic.
[110] M. Matsumura, D. Chen, T. Sawaguchi, K. Kubota, and E. E. Fetz, “Synaptic interactions between primate precentral cortex neurons revealed by spike-triggered averaging of intracellular membrane potentials in vivo,” *Journal of Neuroscience*, vol. 16, no. 23, pp. 1–11, 1996.
[111] B. W. Connors, R. C. Malenka, and L. R. Silva, “Two inhibitory postsynaptic potentials, and GABAA and GABAB receptor-mediated responses in neocortex of rat and cat.,” *The Journal of Physiology*, vol. 406, no. 1, pp. 443–468, 1988, doi: 10.1113/jphysiol.1988.sp017390.

[112] B. Bettler, K. Kaupmann, J. Mosbacher, and M. Gassmann, “Molecular structure and physiological functions of GABAB receptors,” *Physiological Reviews*, vol. 84, no. 3, pp. 835–867, 2004, doi: 10.1152/physrev.00036.2003.
[113] G. Silberberg and H. Markram, “Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells,” *Neuron*, vol. 53, no. 5, pp. 735–746, 2007, doi: 10.1016/j.neuron.2007.02.012.
[114] J. Zhu, M. Jiang, M. Yang, H. Hou, and Y. Shu, “Membrane potential-dependent modulation of recurrent inhibition in rat neocortex,” *PLoS Biology*, vol. 9, no. 3, 2011, doi: 10.1371/journal.pbio.1001032.
[115] S. J. Cruikshank, T. J. Lewis, and B. W. Connors, “Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex,” *Nature Neuroscience*, vol. 10, no. 4, pp. 462–468, 2007, doi: 10.1038/nn1861.
[116] H. Adesnik and M. Scanziani, “Lateral competition for cortical space by layer-specific horizontal circuits,” *Nature*, vol. 464, no. 7292, pp. 1155–1160, 2010, doi: 10.1038/nature08935.
[117] M. Helmstaedter, J. F. Staiger, B. Sakmann, and D. Feldmeyer, “Efficient recruitment of layer 2/3 interneurons by layer 4 input in single columns of rat somatosensory cortex,” *Journal of Neuroscience*, vol. 28, no. 33, pp. 8273–8284, 2008, doi: 10.1523/JNEUROSCI.5701-07.2008.
[118] S. S. Kumar, J. Wülfing, S. Okujeni, J. Boedecker, M. Riedmiller, and U. Egert, “Autonomous optimization of targeted stimulation of neuronal networks,” *PLoS Computational Biology*, vol. 12, no. 8, pp. 1–22, 2016, doi: 10.1371/journal.pcbi.1005054.
[119] O. Weihberger, S. Okujeni, J. E. Mikkonen, and U. Egert, “Quantitative examination of stimulus-response relations in cortical networks in vitro,” *Journal of Neurophysiology*, vol. 109, no. 7, pp. 1764–1774, 2013, doi: 10.1152/jn.00481.2012.
[120] P. Kara, J. S. Pezaris, S. Yurgenson, and R. C. Reid, “The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation,” vol. 99, no. 25, pp. 1–6, 2002.
[121] D. F. English, S. McKenzie, T. Evans, K. Kim, E. Yoon, and G. Buzsáki, “Pyramidal cell-interneuron circuit architecture and dynamics in hippocampal networks,” *Neuron*, vol. 96, no. 2, pp. 505-520.e7, 2017, doi: 10.1016/j.neuron.2017.09.033.
[122] A. Citri and R. C. Malenka, “Synaptic plasticity: Multiple forms, functions, and mechanisms,” *Neuropsychopharmacology*, vol. 33, no. 1, pp. 18–41, 2008, doi: 10.1038/sj.npp.1301559.
[123] A. Losonczy, L. Zhang, R. Shigemoto, P. Somogyi, and Z. Nusser, “Cell type dependence and variability in the short-term plasticity of EPSCs in identified mouse hippocampal interneurones,” *Journal of Physiology*, vol. 542, no. 1, pp. 193–210, 2002, doi: 10.1113/jphysiol.2002.020024.
[124] A. V. Blackman, T. Abrahamsson, R. P. Costa, T. Lalanne, and P. J. Sjöström, “Target-cell-specific short-term plasticity in local circuits,” *Frontiers in Synaptic Neuroscience*, vol. 5, no. DEC, pp. 1–13, 2013, doi: 10.3389/fnsyn.2013.00011.
[125] H. Markram, Y. Wang, and M. Tsodyks, “Differential signaling via the same axon of neocortical pyramidal neurons,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 95, no. 9, pp. 5323–5328, 1998, doi: 10.1073/pnas.95.9.5323.
[126] R. Chen, “Interactions between inhibitory and excitatory circuits in the human motor cortex,” *Experimental Brain Research*, vol. 154, no. 1, pp. 1–10, Jan. 2004, doi: 10.1007/s00221-003-1684-1.
[127] B. Haider, “Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition,” *Journal of Neuroscience*, vol. 26, no. 17, pp. 4535–4545, 2006, doi: 10.1523/jneurosci.5297-05.2006.
[128] G. Turrigiano, “Homeostatic synaptic plasticity: Local and global mechanisms for stabilizing neuronal function,” *Cold Spring Harbor Perspectives in Biology*, vol. 4, no. 1, p. a005736, 2012, doi: 10.1101/cshperspect.a005736.
[129] M. Xue, B. V. Atallah, and M. Scanziani, “Equalizing excitation-inhibition ratios across visual cortical neurons,” *Nature*, vol. 511, no. 7511, pp. 596–600, 2014, doi: 10.1038/nature13321.

[130] R. Rubin, L. F. Abbott, and H. Sompolinsky, “Balanced excitation and inhibition are required for high-capacity, noise-robust neuronal selectivity,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 114, no. 44, pp. E9366–E9375, 2017, doi: 10.1073/pnas.1705841114.
[131] R. P. N. Rao, *Brain-Computer Interfacing: An Introduction*. Cambridge: Cambridge University Press, 2013. doi: 10.1017/CBO9781139032803.
[132] S. R. Soekadar, N. Birbaumer, M. W. Slutzky, and L. G. Cohen, “Brain–machine interfaces in neurorehabilitation of stroke,” *Neurobiology of Disease*, vol. 83, pp. 172–179, Nov. 2015, doi: 10.1016/j.nbd.2014.11.025.
[133] I. Lazarou, S. Nikolopoulos, P. C. Petrantonakis, I. Kompatsiaris, and M. Tsolaki, “EEG-Based Brain–Computer Interfaces for Communication and Rehabilitation of People with Motor Impairment: A Novel Approach of the 21st Century,” *Front. Hum. Neurosci.*, vol. 12, Jan. 2018, doi: 10.3389/fnhum.2018.00014.
[134] D. J. McFarland, J. Daly, C. Boulay, and M. A. Parvaz, “Therapeutic applications of BCI technologies,” *Brain-Computer Interfaces*, vol. 4, no. 1–2, pp. 37–52, Apr. 2017, doi: 10.1080/2326263X.2017.1307625.
[135] A. Lozano *et al.*, “Neurolight: A Deep Learning Neural Interface for Cortical Visual Prostheses,” *Int. J. Neur. Syst.*, vol. 30, no. 09, p. 2050045, Sep. 2020, doi: 10.1142/S0129065720500458.
[136] N. Mrachacz-Kersting *et al.*, “Efficient neuroplasticity induction in chronic stroke patients by an associative brain-computer interface,” *Journal of Neurophysiology*, vol. 115, no. 3, pp. 1410–1421, Mar. 2016, doi: 10.1152/jn.00918.2015.
[137] G. J. del Zoppo, J. L. Saver, E. C. Jauch, and H. P. Adams, “Expansion of the Time Window for Treatment of Acute Ischemic Stroke With Intravenous Tissue Plasminogen Activator: A Science Advisory From the American Heart Association/American Stroke Association,” *Stroke*, vol. 40, no. 8, pp. 2945–2948, Aug. 2009, doi: 10.1161/STROKEAHA.109.192535.
[138] A. G. Rouse and M. H. Schieber, “Advancing brain-machine interfaces: moving beyond linear state space models,” *Front. Syst. Neurosci.*, vol. 9, Jul. 2015, doi: 10.3389/fnsys.2015.00108.
[139] R. P. Rao, “Towards neural co-processors for the brain: combining decoding and encoding in brain–computer interfaces,” *Current Opinion in Neurobiology*, vol. 55, pp. 142–151, Apr. 2019, doi: 10.1016/j.conb.2019.03.008.
[140] R. P. N. Rao, “Brain Co-Processors: Using AI to Restore and Augment Brain Function,” p. 38, 2021.
[141] R. Brette, “Philosophy of the Spike: Rate-Based vs. Spike-Based Theories of the Brain,” *Front. Syst. Neurosci.*, vol. 9, Nov. 2015, doi: 10.3389/fnsys.2015.00151.
[142] E. E. Fetz, “Temporal Coding in Neural Populations?,” *Science*, vol. 278, no. 5345, pp. 1901–1902, Dec. 1997, doi: 10.1126/science.278.5345.1901.
[143] R. Yun, J. H. Mishler, S. I. Perlmutter, R. P. N. Rao, and E. Fetz, “Responses of cortical neurons to intracortical microstimulation in 2 awake primates,” p. 49.
[144] W. Maass, “Computing with Spiking Neurons,” in *Computing with Spiking Neurons*, Pulsed Neural Networks, 1999, pp. 55–85.
[145] M. Pfeiffer and T. Pfeil, “Deep Learning With Spiking Neurons: Opportunities and Challenges,” *Front. Neurosci.*, vol. 12, p. 774, Oct. 2018, doi: 10.3389/fnins.2018.00774.
[146] S. Tafazoli, C. J. MacDowell, Z. Che, K. C. Letai, C. R. Steinhardt, and T. J. Buschman, “Learning to control the brain through adaptive closed-loop patterned stimulation,” *J. Neural Eng.*, vol. 17, no. 5, p. 056007, Oct. 2020, doi: 10.1088/1741-2552/abb860.
[147] T. Xu, N. Xiao, X. Zhai, P. Kwan Chan, and C. Tin, “Real-time cerebellar neuroprosthetic system based on a spiking neural network model of motor learning,” *J. Neural Eng.*, vol. 15, no. 1, p. 016021, Feb. 2018, doi: 10.1088/1741-2552/aa98e9.
[148] S. Buccelli *et al.*, “A neuromorphic prosthesis to restore communication in neuronal networks,” *iScience*, p. S2589004219302731, Aug. 2019, doi: 10.1016/j.isci.2019.07.046.

[149] L. E. Shupe *et al.*, “Neurochip3: An Autonomous Multichannel Bidirectional Brain-Computer Interface for Closed-Loop Activity-Dependent Stimulation,” *Front. Neurosci.*, vol. 15, Aug. 2021, doi: 10.3389/fnins.2021.718465.
[150] L. Shupe and E. Fetz, “An Integrate-and-Fire Spiking Neural Network Model Simulating Artificially Induced Cortical Plasticity,” *eNeuro*, vol. 8, no. 2, p. ENEURO.0333-20.2021, Mar. 2021, doi: 10.1523/ENEURO.0333-20.2021.
[151] J. J. M. Reynolds, J. S. Plank, and C. D. Schuman, “Intelligent Reservoir Generation for Liquid State Machines using Evolutionary Optimization,” in *2019 International Joint Conference on Neural Networks (IJCNN)*, Budapest, Hungary, Jul. 2019, pp. 1–8. doi: 10.1109/IJCNN.2019.8852472.
[152] M. Uzair and N. Jamil, “Effects of Hidden Layers on the Efficiency of Neural networks,” in *2020 IEEE 23rd International Multitopic Conference (INMIC)*, Nov. 2020, pp. 1–6. doi: 10.1109/INMIC50486.2020.9318195.
[153] G. Cybenkot, “Approximation by superpositions of a sigmoidal function,” *Mathematics of Control, Signals and Systems*, pp. 303–314, 1989.
[154] M. Bear, B. Connors, and M. Paradiso, *Neuroscience Exploring the Brain*, 4th ed. Accessed: Sep. 28, 2021. [Online]. Available: https://neurophysics.ucsd.edu/courses/physics_171/Neuroscience%20Exploring%20the%20Brain%20-%20Bear,%20Mark%20F.%20[SRG].pdf
[155] S. Butovas and C. Schwarz, “Spatiotemporal Effects of Microstimulation in Rat Neocortex: A Parametric Study Using Multielectrode Recordings,” *Journal of Neurophysiology*, vol. 90, no. 5, pp. 3024–3039, Nov. 2003, doi: 10.1152/jn.00245.2003.
[156] S. Butovas, S. G. Hormuzdi, H. Monyer, and C. Schwarz, “Effects of Electrically Coupled Inhibitory Networks on Local Neuronal Responses to Intracortical Microstimulation,” *Journal of Neurophysiology*, vol. 96, no. 3, pp. 1227–1236, Sep. 2006, doi: 10.1152/jn.01170.2005.
[157] G. E. Mena *et al.*, “Electrical stimulus artifact cancellation and neural spike detection on large multi-electrode arrays,” *PLoS Comput Biol*, vol. 13, no. 11, p. e1005842, Nov. 2017, doi: 10.1371/journal.pcbi.1005842.
[158] D. J. O’Shea and K. V. Shenoy, “ERAASR: an algorithm for removing electrical stimulation artifacts from multielectrode array recordings,” *J. Neural Eng.*, vol. 15, no. 2, p. 026020, Apr. 2018, doi: 10.1088/1741-2552/aaa365.
[159] T. Alejandre-García, S. Kim, J. Pérez-Ortega, and R. Yuste, “Intrinsic excitability mechanisms of neuronal ensemble formation,” *eLife*, vol. 11, p. e77470, May 2022, doi: 10.7554/eLife.77470.
[160] H.-K. Lee and A. Kirkwood, “Mechanisms of Homeostatic Synaptic Plasticity in vivo,” *Front. Cell. Neurosci.*, vol. 13, p. 520, Dec. 2019, doi: 10.3389/fncel.2019.00520.
[161] M. F. Bolus, A. A. Willats, C. J. Rozell, and G. B. Stanley, “State-space optimal feedback control of optogenetically driven neural activity,” Neuroscience, preprint, Jun. 2020. doi: 10.1101/2020.06.25.171785.
[162] S. Dura-Bernal, K. Li, S. A. Neymotin, J. T. Francis, J. C. Principe, and W. W. Lytton, “Restoring Behavior via Inverse Neurocontroller in a Lesioned Cortical Spiking Model Driving a Virtual Arm,” *Front Neurosci*, vol. 10, Feb. 2016, doi: 10.3389/fnins.2016.00028.
[163] M. H. Histed, V. Bonin, and R. C. Reid, “Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation,” *Neuron*, vol. 63, no. 4, pp. 508–522, Aug. 2009, doi: 10.1016/j.neuron.2009.07.016.
[164] S. Ruder, “An overview of gradient descent optimization algorithms,” *arXiv:1609.04747 [cs]*, Jun. 2017, Accessed: Apr. 03, 2022. [Online]. Available: http://arxiv.org/abs/1609.04747
[165] A. Ramakrishnan, Y. W. Byun, K. Rand, C. E. Pedersen, M. A. Lebedev, and M. A. L. Nicolelis, “Cortical neurons multiplex reward-related signals along with sensory and motor information,” *Proc. Natl. Acad. Sci. U.S.A.*, vol. 114, no. 24, Jun. 2017, doi: 10.1073/pnas.1703668114.
[166] E. Musk and Neuralink, “An Integrated Brain-Machine Interface Platform With Thousands of Channels,” *J Med Internet Res*, vol. 21, no. 10, p. e16194, Oct. 2019, doi: 10.2196/16194.

[167] K. Ganguly and J. M. Carmena, “Emergence of a Stable Cortical Map for Neuroprosthetic Control,” *PLoS Biol*, vol. 7, no. 7, p. e1000153, Jul. 2009, doi: 10.1371/journal.pbio.1000153.
[168] T. Gulati, L. Guo, D. S. Ramanathan, A. Bodepudi, and K. Ganguly, “Neural reactivations during sleep determine network credit assignment,” *Nat Neurosci*, vol. 20, no. 9, pp. 1277–1284, Sep. 2017, doi: 10.1038/nn.4601.
[169] A. J. Peters, S. X. Chen, and T. Komiyama, “Emergence of reproducible spatiotemporal activity during motor learning,” *Nature*, vol. 510, no. 7504, pp. 263–267, Jun. 2014, doi: 10.1038/nature13235.