Computational and Experimental Advances in Clinically Relevant Neurostimulation of Non-Human Primates

Julien Bloch

A dissertation
submitted in partial fulfilment of the
requirements for the degree of

Doctor of Philosophy

University of Washington
2023

Reading Committee:
Azadeh Yazdan-Shahmorad, Chair
Steve Perlmutter
Amy Orsborn

Program Authorized to Offer Degree:
Bioengineering

University of Washington

Abstract

Computational and Experimental Advances in Clinically Relevant Neurostimulation of Non-Human Primates

Julien Bloch

Chair of the Supervisory Committee:
Azadeh Yazdan-Shahmorad
Department of Bioengineering

Neural engineering capabilities in non-human primates (NHPs) dramatically lag the capabilities in rodents, despite NHPs being the most neurophysiologically similar animal model to humans. This work puts forward advances in hardware, computation, and theoretical grounding to improve neural engineering ability in NHPs, especially with respect to clinically relevant neural stimulation. We begin by investigating how cortical stimulation affects cortical coherence, a form of functional connectivity, in the NHP cortex (Chapter 2). Our findings highlight how various factors, such as the state of the cortex during stimulation, can modulate coherence. Building on this, Chapter 3 examines how network structure impacts functional reorganization resulting from optogenetic stimulation in the NHP sensorimotor cortex. The results, obtained through graph theory-inspired computational analysis, reveal that network structure, more than the stimulation protocol, can predict changes in functional connectivity. To further our understanding of behavioral and neural disruption studies in NHPs, we introduce a novel opto-electric neural interface in Chapter 4. This tool combines the specific modulation abilities of optogenetics with the dual functionality of electrical stimulation and recording. Consequently, this allows for a more precise investigation into the link between brain activity and behavior. Transitioning to computer-based work in Chapter 5, we develop a biophysically accurate neural network simulation to derive structural connectivity from functional connectivity at the local field potentials (LFP) level. This important advancement aids in understanding the relationship between LFP functional connectivity and structural connectivity. Finally, in Chapter 6, we present and validate a computational model predicting lesion sizes in a new photothrombotic stroke model. This key tool contributes to our understanding of cortical physiology and stroke in NHPs. In summary, this thesis provides a multifaceted approach to enhance our understanding of neurostimulation and its potential therapeutic benefits for neural disorders. By improving experimental and computational tools, grounding functional connectivity in a biological framework, and advancing stroke modeling in NHPs, this work helps to bring neurotechnologies closer to their clinical application in humans.

Acknowledgements

I would like to thank Prof. Azadah Yazdan-Shahmorad for her ceaseless championing and mentorship. I would like to thank Prof. Eric Shea-Brown for his enthusiastic support and excitement. I would like to thank my closest collaborators in my work, Dr. Devon Griggs and Dr. Alec Greaves-Tunnell, for the honor of working with and learning from them. I would like to thank Prof. Ali Shojaie and Prof. Zaid Harchaoui for their collaboration and willingness to share their depth of knowledge of statistical machine learning and how it could be applied to neurotech. I would like to thank all students and industry-side research assistants who I have mentored – I hope you have learned nearly as much from me as I from you. I would like to thank all members of the Yazdan lab, who ceaselessly provided me with support, feedback, and friendship. I would like to thank my mentors and friends in Prof. Eb Fetz' lab, especially Irene Rembado, Jon Mishler, and Richy Yun. Finally, I want to thank my family, especially my mother, my father, my brother Roman, and my sister Vaiana – thank you for always supporting and believing in me.

Table of Contents

Chapter 1: Introduction

We begin with a fundamental question: How do we develop a neurotechnology solution for a given neural disorder? The answer involves an understanding of the disorder and its response to various treatments. We can often understand the gross presentation of a neural disorder through human studies via imaging, post-mortem analysis, basic electrophysiology, and neuropsychology. However, a large majority of our mechanistic understanding of neural disorders and their responses to treatments is derived from studies using preclinical animal models (Perretta 2009, Ellenbroek 2016). Given the wide variation in pathological presentations and responses across species, the discoveries and therapeutic advances made in these animal models often fail when translated to humans (van der Worp 2010). To maximize the probability of successfully translating from animal models to humans, it is necessary to establish a knowledge base of disease dynamics and responses in experimental conditions that closely mimic human neurology and pathophysiology. Achieving this requires several key components: 1) an animal model that is more similar to humans in both its symptom presentation and response to treatments, 2) technical tools and hardware to run the needed experiments, 3) computational ability to optimize the experiments and analyze the data, and 4) a fundamental understanding of the measured outcomes to ground and interpret the data within the currently established framework.

To address the first requirement, non-human primates (NHPs) represent the closest animal model to human neurology (Frederick 1988, Perretta 2009, Van Dam 2017). Unlike most other commonly used animal models, NHPs possess large, wrinkled brains with volume ratios of different brain regions similar to those found in humans and feature direct connections from the cortex to motor neurons. Despite their similarity to humans, the complexity and cost of training and housing NHPs render them a less common research model than rodents, with consequently fewer established disease models. Nevertheless, the use of NHPs is essential to maximize clinical translatability of developing neurotechnologies. The Northstar Neuroscience Everest clinical trial is an example of why NHPs are necessary for the development of such therapies. The trial's aim was to use sub-motor-threshold stimulation to enhance sensorimotor rehabilitation post-stroke. Despite some promising results in preclinical models, the study failed to meet its primary endpoint at 4 weeks and was ended (Levy 2016). The stimulation protocol for the trial was fully developed and optimized in rodents (Teskey 2003, Kleim 2003, Adkins 2003, Adkins 2008), with only a proof-of-concept performed in NHPs (Plautz 2003).

While the Northstar trial may have benefitted from stimulation optimization in NHPs, the direct replication of many contemporary rodent-based stimulation experiments in NHPs is not yet feasible. This is partly due to the more advanced genetic techniques available and commonly used in rodent research, one of which is optogenetics. Since its inception in 2010, optogenetics has revolutionized the field of neural stimulation by providing unparalleled spatiotemporal precision and enabling the selective targeting of specific cell types and projections (Boyden 2015). Despite its prevalence in rodent research, the adoption of optogenetics in NHP studies remains limited (El-Shamayleh 2019). This highlights our second requirement: the necessity for advanced technical tools to conduct precise neural manipulations. Therefore, to optimize stimulation interventions in NHPs such that these interventions can more successfully be translated to humans, it is essential to advance the capabilities of optogenetics for NHP research to the same level as for rodent studies.

Beyond behavioral outcomes, there are other pathology-relevant measures that can be optimized in therapies. One promising target is functional connectivity, or the degree to which two different regions of the brain cooperate. A substantial body of research has linked abnormal functional connectivity to a plethora of neural disorders (Garrity et al. 2007, Nakai et al. 2021, Stam et al. 2007). Furthermore, the use of neural stimulation to modify functional connectivity is well-documented, with the University of

Washington's very own Prof. Fetz pioneering its application in NHPs (Jackson 2006). Despite its potential, two factors currently limit its use in rehabilitative-relevant contexts: the historically inconsistent results encountered when trying to modify functional connectivity over a network (Rebesco 2011, Seeman 2017, Song 2013) and the limited biological understanding of the phenomenon itself. Fortunately, new graph-centric computational tools are improving our ability to understand stimulation induced functional reorganization over a network (Khambatini 2019, Muldoon 2016), and computer simulations are improving our understanding of the biological underpinnings of functional connectivity (Vuksanović 2014). These improvements must continue in order to arrive at the requisite computational ability and fundamental understanding needed to develop and optimize functional connectivity-based interventions for successful clinical translation.

My work addresses each of the aforementioned four needs of having a clinically relevant model, technical tools, computational tools, and fundamental understanding of the measured outcomes. On the experimental front, my work advances the ability to utilize optogenetics in NHPs to manipulate both functional connectivity and behavior. My computational work advances our computational ability to predict stimulation-induced functional connectivity reorganization over a network, to induce targeted lesions for studying stroke-like neurophysiology, and to ground functional connectivity in relation to structural connectivity. These advances are detailed over several chapters, each of which focuses on different yet interconnected aspects of clinically-relevant neurostimulation in NHPs.

In Chapter 2, I delve into the domain of cortical stimulation, and discuss how it modulates coherence, a type of functional connectivity, within the NHP cortex. These results explore how factors such as stimulation latency and intrinsic network connectivity control the coherence changes induced by stimulation. This work is complimented by a series of computational analyses carried out in Chapter 3, where I explore the role of network structure in mediating functional reorganization induced by optogenetic stimulation of the NHP sensorimotor cortex. In these studies, I show that the stimulation protocol is only part of the equation. The underlying network architecture plays an even greater role at mediating stimulation-induced changes. This finding shifts our perspective on targeted neural stimulation and holds potential for the development of novel, more effective stimulation-based interventions for neural disorders.

Chapter 4 introduces a novel opto-electric neural interface designed for advanced behavioral and neural disruption studies in NHPs. This new tool merges the cell-type specific modulation capabilities of optogenetics with the dual functionality of electrical stimulation and recording, opening new avenues for precise, detailed studies of how brain activity shapes behavior and vice versa. With it, I demonstrate the optogenetic manipulation of both neural activity as well as reaching behavior.

In Chapter 5, I use a biophysically plausible artificial neural network to study how structural connectivity can be inferred from local field potentials (LFP). This work provides crucial insights for interpreting LFP functional connectivity in the context of structural connectivity, thereby grounding the often-studied yet poorly-understood phenomena of functional connectivity in biological understanding.

Chapter 6 details the development and validation of a computational model for predicting induced lesion sizes in a novel photothrombotic stroke model. The model serves as a crucial tool for studying cortical physiology following stroke and can potentially be utilized to develop stimulation-based interventions in NHPs, thus facilitating the development of clinically translatable stroke therapies.

In summary, this thesis offers a comprehensive view of my efforts to advance neurostimulation science and its clinical applicability. By studying how neural stimulation affects the brain, improving the available experimental and computational tools used for such investigations in NHPs, grounding a good candidate clinical measure in biology, and furthering a replicable NHP stroke model, this research works to bridge the gap between experimental neurotechnologies and their use in humans. Through these endeavors, we move closer towards developing and optimizing stimulation-based treatments of neural disorders.

Chapter 2: Cortical stimulation induces network-wide coherence change in non-human primate somatosensory cortex

Published as: **J. Bloch**, K. Khateeb, D. Silversmith, J. O'Doherty, P. Sabes, A. Yazdan-Shahmorad "Cortical Stimulation Induces Network-Wide Coherence Change Across Non-Human Primate Cortex," IEEE EMBC 2019, Berlin, Germany, Jul. 2019.

Abstract

Stimulation of the cortex can modulate the connectivity between brain regions. Although targeted neuroplasticity has been demonstrated in-vitro, in-vivo models have been inconsistent in their response to stimulation. In this paper, we tested various stimulation protocols to characterize the effect of stimulation on coherence in the non-human primate cortex in-vivo. We found that the stimulation latency, the state of the cortex during stimulation, and the stimulation site all affected the modulation of cortical coherence. We further investigated features of a resting-state network that could predict how a connection is likely to change with stimulation.

Introduction

Neurological disorders affect millions of people world-wide. Many of these disorders, such as epilepsy, schizophrenia, and Alzheimer's disease are linked to aberrant connectivity in the brain (Stam 2014). One possible approach to combat these disorders is targeted reorganization of neural connections through neural stimulation. As the brain has neuroplastic mechanisms, stimulation can leverage neuroplasticity to result in changed neuronal network dynamics.

Neuroplasticity has been shown to follow the framework of spike-timing dependent plasticity (STDP), which holds that when neuron A activates immediately before neuron B, the connection from A to B is strengthened (long-term potentiation), while when B fires before A then the connection from A to B is weakened (long-term depression) (Bi and Poo 1998). The degree of connection change is modulated by the latency between the two activations. In this paper we test different stimulation protocols with varying relations to the STDP latency window. For quantifying connection strength, we use the coherence metric, which has shown to be more robust than other commonly used neural connection metrics (Bastos 2016). Specifically, we analyze coherence of the beta, gamma, and high-gamma bands, which are commonly associated with cognitive and motor processes (Lucas 2013, Yazdan-Shahmorad 2018).

Previous studies have explored the effects of stimulation for changing cortical dynamics. Such studies have varied from electrical to optogenetic stimulation and have investigated the immediate and long-term effects of such stimulation between stimulation pairs, with varying degrees of success (Rebesco 2011, Buzsaki 2015, Seeman 2017). In this study, we electrically stimulate cortical regions and measure how the stimulation changes the dynamics over the whole network.

Results

A. Electric Stimulation Alters Connectivity Dynamics Between Stimulated Sites

We obtained results characterizing the effects of the various stimulation protocols on the coherence between the two stimulated channels (channels A and B), as shown in Fig. 1A. The stimulation block coherences of all stimulation protocols except out-STDP were statistically significantly deviated from the pre-stimulation baseline values (p-values<0.01). We repeated the recordings after stimulation and found the new baseline coherence during this period to still be statistically significantly deviated from the pre-stimulation baseline (p-values<0.001). We additionally performed a sham trial, in which no stimulation was applied, as a control and found no statistical significance in the deviation from its baseline coherence.

B. Stimulation-Induced Network Depends on the State of the Network

We stimulated multiple pairs of electrodes in the same experiment session, with later stimulation sessions beginning before coherence values returned to baseline. On one day, we performed in-STDP stimulation, then an extended recording session, and then out-STDP stimulation. On another day, we reversed the order and performed out-STDP stimulation, an extended recording session, and then in-STDP stimulation. This reversal allowed us to compare differences in the evoked coherence of a recently stimulated brain vs. a brain in a resting unstimulated state.

Our results, displayed in Fig. 1B, demonstrated that stimulating the brain when already in a recently stimulated state can affect the stimulation-induced coherence. The induced coherences differed significantly (p-values <0.05) between the two approaches for all trials except in-STDP for beta band and high-gamma band.

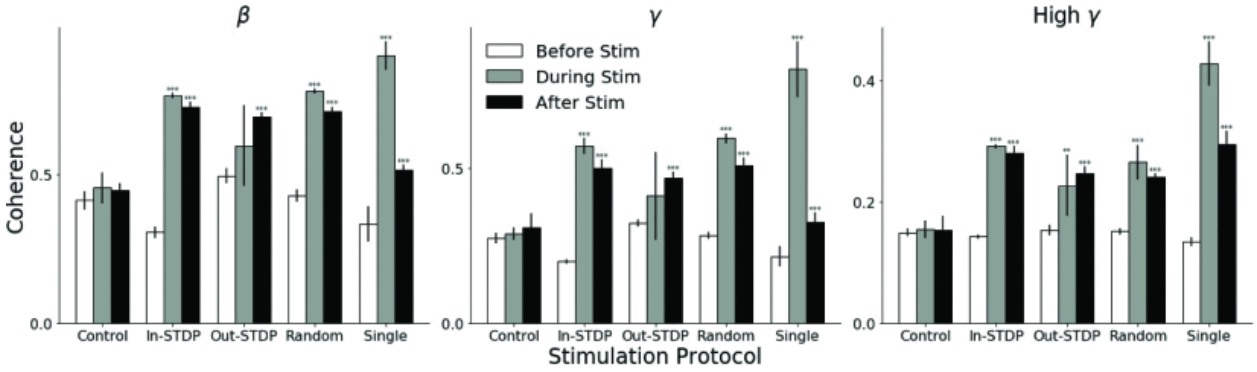

Fig. 1. Coherence between stimulation electrodes before stimulation, during the stimulation recording block, and after stimulation ended. Bar heights indicate mean values, while error bars indicate 95% confidence interval. Asterisks indicate statistical significance, with '*' indicating p<0.05, '**' indicating p<0.01, and '***' indicating p<0.001. Statistical significance was calculated by paired t-test relative to the "Before Stim" data.

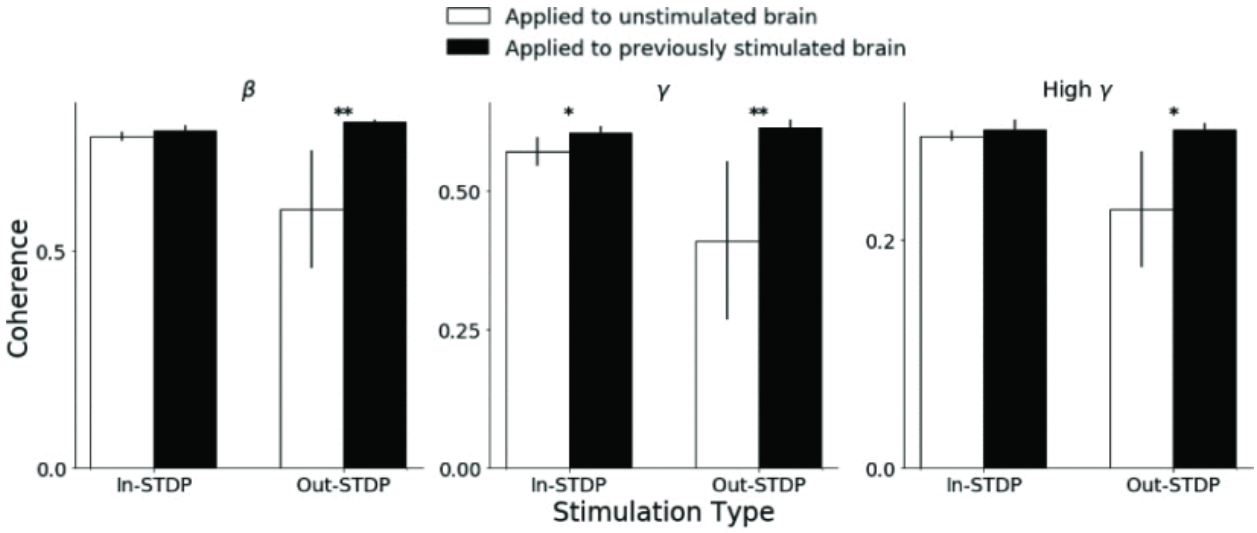

Fig. 2. Comparison of coherence for two types of paired stimulation, when applied as first stimulation of experimental session (white bars) vs. when applied to a brain that had already been stimulated in that experimental session (black bars).

C. Paired Stimulation Can Increase or Decrease Coherence Between Stimulation Sites

All stimulation-driven coherence changes between electrodes A and B corresponded to increases in coherence. We wanted to test whether coherence increase was a truism of neuronal coherence dynamics following stimulation, or a facet of the specific neuronal location of electrodes A and B. We thus tested two other stimulation pairs. One of the pairs, consisting of electrodes C and D, was located far from our original stimulation electrode pair (A and B), while the other pair (E and F) was spatially close to A and B. The intra-pair distance was consistent between all pairs.

The trials for these channels were completed after A and B had already been stimulated, and the coherence had not yet returned to baseline (Fig. 3). The data in this section are tested for significance with respect to the post A-B stimulation excited baselines. Therefore, although we do test these data for significant changes in coherence, we must qualify these results by stating that they are valid within the context of recent previous paired stimulation of sites A and B.

We compared the coherence of the baseline immediately before stimulation with the coherence of the stimulation electrodes during the stimulation block, and with the coherence of the stimulation electrodes during the post-stimulation baseline. C-D pair demonstrated significant increase for all comparisons while E-F pair demonstrated significant decrease for all comparisons (p-values<0.01) except the beta coherence changes of pair E-F.

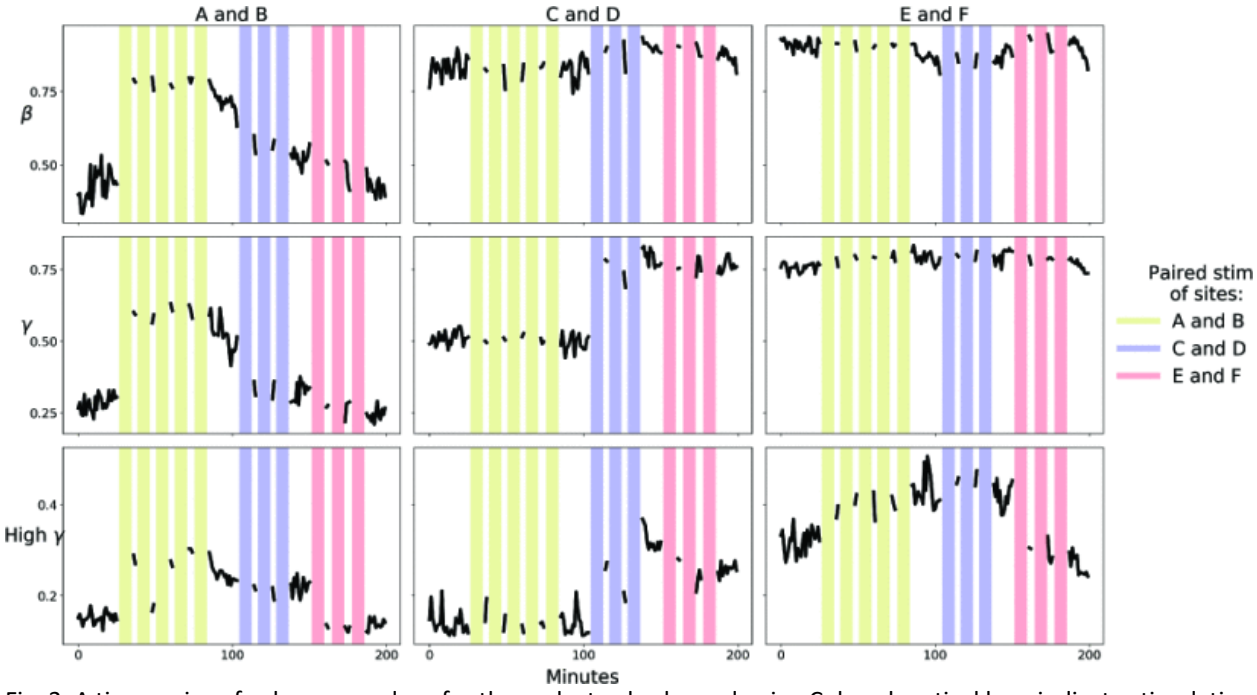

Fig. 3. A time series of coherence values for three electrode channel pairs. Colored vertical bars indicate stimulation sessions. The left, center, and right plots indicate coherence between channels A and B, C and D, and E and F respectively. Yellow, light blue, and pink lines indicate 10 minute periods of paired stimulation between channels A and B, C and D, and E and F respectively.

D. Features of Baseline Recording Are Related to Stimulation Induced Coherence Change Across Entire Network

During recording blocks, we obtained coherence values for all electrode pair combinations of the 96 Utah array electrodes. We quantified the change in coherence of each pair relative to baseline and performed

least squares regression to quantify how baseline coherence, baseline coherence stability, and electrode pair distance are related to coherence change during stimulation. In order to quantify baseline coherence stability, we calculated the standard deviation of the coherences of each electrode pair during baseline recording. As change in coherence includes a term of the baseline coherence, the regression is biased. In order to unbias the regression, we split the coherences up into odd and even minutes, and calculated the baseline with the even minutes, while calculating the change in coherence with the odd minutes. A similar approach has been used in (Rebesco 2011) in order to combat this same regression bias.

Our analysis showed that the baseline coherence, standard deviation of baseline coherence, and electrode pair distance were each significantly (p-value<0.001) linearly related to the change in coherence (Fig. 4). Baseline coherence yielded the highest r squared value for beta and gamma band coherence change, while baseline standard deviation yielded the highest r2 for high gamma. All baseline coherence slopes were negative, and all baseline standard deviation slopes were positive. This indicates that for all frequency bands, both lower baseline coherences and lower connection stabilities (corresponding to a higher coherence standard deviation) were correlated with more positive coherence change.

We then constructed a multiple regression model consisting of the three features in order to evaluate the degree to which each feature could increase the explained variance ratio (r2) of the regression model. The multiple regression model yielded r2 values of 0.236, 0.122, and 0.029 for beta, gamma, and high gamma, respectively. These improved r2 values are appreciably higher than the highest r2 values of the single regression models for gamma (0.072, from baseline coherence) and high-gamma band (0.013, from baseline stability), but not for beta (0.23, from baseline coherence). Each feature continued to be statistically significant (p-values<0.001) in the multiple regression models.

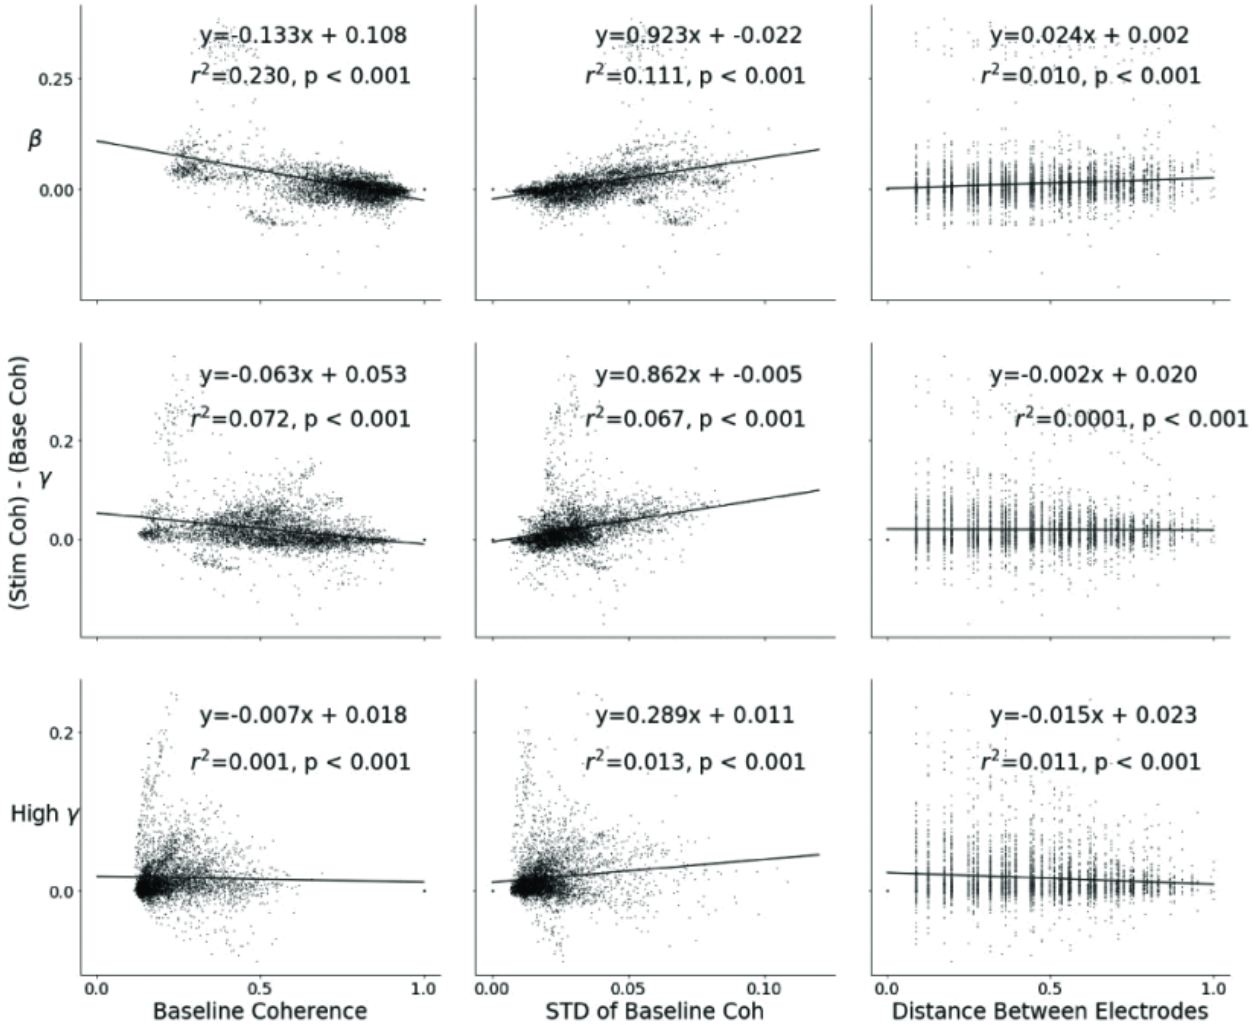

Fig. 4. Scatter plots of stimulation-induced coherence change with respect to (from left to right) mean baseline coherence, standard deviation of baseline coherence, and distance between electrode pairs. The top plots are for beta frequency band, the middle for gamma, and the bottom for high-gamma. The data points represent all 96x96 coherence changes all possible electrode pairs. Each feature is statistically significantly related to the change in baseline coherence.

Discussion

In this study we investigated cortical network dynamics in response to electrical stimulation. We compared stimulation protocols, electrode sites, and predictive baseline features. Our results indicate that coherence between neural sites can be modulated through stimulation, and this modulation is influenced by stimulation latency, site of stimulation, and the state of the network before stimulation.

One hypothesis constructed from a naive extrapolation of STDP excitation rules is that paired stimulation within the long-term potentiation window would result in unilateral increase in coherence between two stimulation sites. This hypothesis would explain the more consistent significant coherence change of in-STDP and random stimulation compared to out-STDP. However, although we did see increases in coherence for pairs A-B and C-D, pair E-F decreased in coherence when stimulated. STDP rules delineate that when A fires before B, the connection from A to B is strengthened but the connection from B to A is weakened. As magnitude of coherence measures the overall connectivity between two sites, and does

not confer any information of directionality, it is possible that there was a larger decrease in connectivity from F to E than there was an increase from E to F, which yielded a net decrease in connectivity. Although this explanation might fit within the STDP framework, it does not take into account how the single-site stimulation of channel B resulted in increased coherence with channel A.

An alternative hypothesis is that connections between neuronal sites are largely unchanged, and that the chief mechanism of electrical stimulation induced coherence modulation is modulation of cortical oscillation rhythms. Since coherence is a measure of the degree to which two signals keep consistent phase with one another, modulation of cortical oscillation synchrony would result in changed coherence. However, similar studies have investigated the change in connection between two sites by using other connectivity metrics such as evoked response amplitudes or Granger Causality (Rebesco 2011, Buzsaki 2015, Seeman 2017) and have reported similar connectivity modulation with these metrics. In addition, coherence change has been directly compared to evoked response change and the two were shown to be significantly correlated (Rebesco 20011).

It is clear from Fig. 3 that there are dynamics determined by the underlying cortical network that dictate the response of the network to stimulation. Fig. 4 is a first attempt at predicting this response from the network. The r squared values and p-values of the regression models indicate that simple features such as inter-electrode distance, baseline coherence, and baseline coherence stability may successfully predict the degree of coherence change during stimulation, but that as the frequency of the coherence rises, so does the difficulty in explaining the change in coherence. This is consistent with past analyses of local-field potential frequencies, in which higher frequencies are hypothesized to underlie more complex local connectivity (Yazdan-Shahmorad 2018). If more data is made available experimentally, more complex modeling tools can be leveraged to gain an understanding of the cortical network structure that underlies connectivity modulation.

Methods

A. Animal Model

One adult male rhesus macaque was used in this study. All experiments were performed under approval of the University of California, San Francisco Institutional Animal Care and Use Committee and were compliant with the Guide for the Care and Use of Laboratory Animals.

B. Data Acquisition

A macaque monkey was implanted with a 96-electrode Utah array in its primary somatosensory cortex (S1). We used a Tucker-Davis Technologies System (FL, USA) to control the recording and stimulation of the electrodes. For the duration of an experiment, the monkey was awake and headfixed, sitting in a primate chair. We monitored the animal during the experiment to ensure it remained awake.

The local field potentials were sampled at a rate of 24kHz, and then downsampled to 3.051kHz. We recorded neural activity before and after stimulation. Additionally, we interleaved 10 minutes of stimulation with neural recordings to obtain data corresponding to stimulation blocks. This interleaving was necessary as stimulation results in saturation of the recording electrodes.

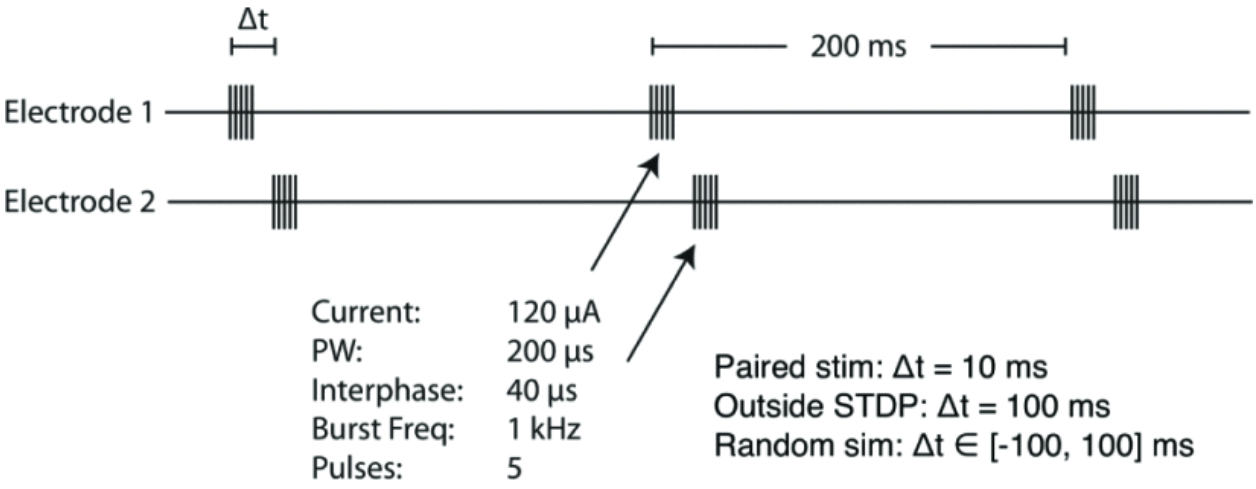

Fig. 5. Paired stimulation protocol. A 5ms 1kHz burst is delivered to electrode 1, and then a 5ms 1kHz burst is delivered to electrode 2. The time delay, Δt, between onsets of the bursts is detailed in the lower-right of the figure. This protocol is applied in 10 minute blocks interleaved with recordings.

C. Stimulation Protocols

Stimulation of an electrode consisted of a 5ms burst of 1kHz stimulation at 120 µA. In this study we used four different stimulation protocols: in-STDP, out-STDP, random, and single-site. In-STDP, out-STDP, and random correspond to protocols of paired stimulations (alternating stimulations of 2 channels) of different time intervals between the stimulations (Fig. 5). In-STDP stimulation consisted of two bursts, one in the first stimulated channel and one in the second, with a separation between the burst onsets of 10ms. Out-STDP had a burst onset separation of 100ms, while random had a burst onset separation randomly sampled from the uniform distribution of -100ms to 100ms. Single-site corresponded to stimulation of a single channel. For each protocol the stimulations were repeated every 200ms for an allotted stimulation time.

The in-STDP and out-STDP stimulation protocols correspond to different relationships to the STDP window. We set the lag between electrode 1 and 2 bursts of in-STDP stimulation at 10ms in order to fall inside the STDP long-term potentiation window, and the out-STDP lag to 100ms in order to fall outside the long-term potentiation window.

D. Quantification and Statistical Analysis

The data analysis was completed in Python. Coherence calculation was completed using the Scipy package. Statistical analysis such as paired t-test and construction of linear regression confidence intervals were completed using the Scipy and Statsmodel packages.

E. Coherence

The raw signals from each recording block were split into time-frames of 1-minute long recordings, and the coherence was calculated for each 1-minute long block. Within each 1-minute block, pairwise coherences were calculated with a Hamming window of 10 seconds and 50% overlap. Coherence is defined as:

$$C_{xy} = \frac{|G_{xy}(f)|^2}{G_{xx}(f)\,G_{yy}(f)} \quad (1)$$

Where Gxx and Gyy correspond to channel power spectral density of channels x and y, respectively, and Gxy corresponds to the cross-spectral power density of the two channels. Coherence values were binned into beta, gamma, and high-gamma frequency ranges corresponding to 12Hz-30Hz, 30Hz-55Hz, and 65Hz-

200Hz, respectively. Significant changes between different recording blocks were detected using paired 2-sided t-tests.

Conclusion

As a means to combat neural disorders stemming from aberrant connectivity, cortical stimulation is a powerful method. The inherent complexity of the brain, however, makes targeted reorganization of brain connectivity using stimulation a complex task. As connections between neural areas are mediated by many intermediary connections, the dynamics of paired-site connectivity modulation can only be understood by gaining an understanding of the connectivity dynamics of the entire network. In this paper we evaluated the effect of different stimulation modalities on coherence between brain regions, both between stimulation sites and of the network. This work can offer direction to researchers choosing a methodology of stimulation in order to affect cortical connections, for either rehabilitative or purely scientific purposes.

Chapter 3: Network structure mediates functional reorganization induced by optogenetic stimulation of non-human primate sensorimotor cortex

Published as**: J. Bloch**, A. Greaves-Tunnell, E. Shea-Brown, A. Harchaoui, A. Shojaie, A. Yazdan-Shahmorad, "Network structure mediates functional reorganization induced by optogenetic stimulation of non-human primate sensorimotor cortex," *iScience*, May 2022.

Abstract

As aberrant network-level functional connectivity underlies a variety of neural disorders, the ability to induce targeted functional reorganization would be a profound development towards therapies for neural disorders. Brain stimulation has been shown to alter large-scale network-wide functional connectivity, but the mapping from stimulation to the modification is unclear. Here, we leverage advances in neural interfaces, interpretable machine learning, and graph theory to arrive at a model which accurately predicts network-wide functional reorganization in response to optogenetic stimulation of non-human primate primary sensorimotor cortex. The model jointly considers the stimulation protocol and the cortical network structure, departing from the standard approach which only considers the stimulation protocol. We observe that the stimulation protocol only predicts a small portion of the induced functional connectivity changes while the network structure predicts much more, indicating that the underlying network is the primary mediator of the response to stimulation. We extract the relationships linking the stimulation and network characteristics to the functional connectivity changes and observe that the mappings diverge over frequency bands and successive stimulations. Finally, we uncover shared processes governing real-time and longer-term effects of stimulation. Our framework represents a paradigm shift for targeted neural stimulation and can be used to interrogate, improve, and develop stimulation-based interventions for neural disorders.

Introduction

From schizophrenia to epilepsy, aberrant network-level functional connectivity underlies a variety of neural disorders (Garrity et al. 2007; Nakai et al. 2021; Stam et al. 2007). As such, the ability to modify network-level functional connectivity would be a profound development towards therapies for neural disorders. Brain stimulation has demonstrated promise as a means for modifying large-scale functional connectivity (Edwardson et al. 2013; M. D. Fox, Halko, et al. 2012; Sehm et al. 2012), but the mapping from stimulation to the resulting modification remains unclear. The prevalent approach for inducing targeted connectivity change is neural stimulation informed by Spike-Timing Dependent Plasticity (STDP). STDP is a phenomenon in which synaptic strengths are modified as a function of the delay between the firing times of the pre- and post-synaptic neurons (Bi and Poo 1998; Bliss, T.V.P. and Collingridge 1993; Markram et al. 1997). While STDP was initially observed in monosynaptic connectivity strength between isolated neuron pairs *in vitro* (Bi and Poo 1998; Bliss, T.V.P. and Collingridge 1993; Markram et al. 1997), its application was eventually translated to larger-scale, less isolated connections *in vivo* (Bloch et al. 2019; Jackson, Mavoori, and Fetz 2006; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013; Yazdan-Shahmorad, Silversmith, et al. 2018). In response to the more challenging *in vivo* setting, this translation coincided with stimulation of a larger set of neurons instead of monosynaptic connections of isolated neuron pairs and the measurement of functional connectivity changes (FCC) instead of direct synaptic strength changes (Bloch et al. 2019; Jackson, Mavoori, and Fetz 2006; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013; Yazdan-Shahmorad, Silversmith, et al. 2018). In line with the original *in vitro* studies, these experiments hypothesized that pairwise stimulation would drive FCC only between the stimulation targets. Some reported successful induction of targeted FCC (Jackson, Mavoori, and Fetz 2006; Rebesco and L. E. Miller 2011; Rebesco,

Stevenson, et al. 2010; W. Song et al. 2013) while others reported inconsistent results across stimulation targets (Bloch et al. 2019; Seeman et al. 2017). Notably, these studies also reported a novel finding: stimulation-induced FCC were found to extend beyond the stimulation sites to other recorded regions (Bloch et al. 2019; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013). Follow up studies investigated the extent of this phenomenon in primates (Huang et al. 2019; Keller, Huang, et al. 2018; Yazdan-Shahmorad, Silversmith, et al. 2018) and found that paired stimulation reliably results in FCC extending over a large-scale surrounding network. These results highlight the need for an updated framework which accurately predicts network-wide stimulation-induced FCC *in vivo*.

Recent research indicates that network-level functional connectivity is not only affected by stimulation but also is involved in shaping the response to stimulation. For example, underlying functional connectivity has been shown to shape stimulation-induced neural activity propagation (K. C. Fox et al. 2020; Huang et al. 2019; Keller, Bickel, et al. 2011; Momi et al. 2021; Solomon et al. 2018) as well as influence the therapeutic outcome of neural stimulation interventions for Parkinson's disease (Horn et al. 2017; Younce et al. 2021) and depression (M. D. Fox, Buckner, et al. 2012). Early results have also indicated that network-level FCC induced by stimulation are correlated to baseline functional connectivity (Bloch et al. 2019; Keller, Huang, et al. 2018; Khambhati et al. 2019). Despite the demonstrated relevance of the underlying network for shaping the response to stimulation, the question of how to effectively parse and analyze a brain network to inform stimulation remains unanswered. The answer to this question could be used to arrive at an updated framework for network-level stimulation induced FCC.

Here, we leverage advances in neural interfaces and computational modeling to develop a model which jointly considers characteristics of the stimulation protocol and the underlying network-level functional connectivity for prediction of network-wide stimulation-induced FCC. We use optogenetics, a stimulation technology in which neurons are rendered light-sensitive by viral-mediated expression of opsins and thus able to be activated by incident light (Boyden et al. 2005), to activate target neural regions while simultaneously recording without electrical artifact. We measure FCC at the network level by recording neural activity via a micro-electrocorticography (µECoG) array covering 1cm^2 of the primary sensorimotor cortex. We then parse underlying network-level functional connectivity with summarizing metrics borrowed from and inspired by graph theory. Finally, we employ a nonparametric hierarchical additive model (Haris, Shojaie, and Simon 2019) to predict network-level FCC while allowing maximal interpretability of identified feature mappings. We develop and test the model in the sensorimotor cortex of two non-human primates (NHPs) to ensure maximal clinical translation.

We interrogated our model over multiple contexts to identify the relationship between stimulation and FCC. Characteristics of the stimulation protocol alone poorly predicted stimulation induced FCC, whereas the characteristics of the underlying network-level functional connectivity contained information which yielded accurate prediction of the FCC. This trend was true both when protocol and network features were analyzed as groups and when features were analyzed individually. The mappings from input features to FCC were similar between the stimulated-state and resting-state changes. Continuous stimulation modified the mapping from input features to resulting FCC. Our methods represent a promising new framework for parsing network-level FCC to predict the effects of stimulation, while our model and insights can be leveraged towards novel therapies of neural disorders.

Results

Large-scale stimulation and recording of non-human primate sensorimotor cortex

We performed a series of experimental sessions consisting of stimulating and recording rhesus macaque cortex. To stimulate neurons without causing a recording artifact we used optogenetic stimulation. We

chose to express the C1V1 opsin as it is a red-shifted excitatory opsin and thus allows for greater light penetration and subsequent neural activation (Prakash et al. 2012). To record over a large-scale ($1cm^2$) area of the cortex while allowing light to pass through the recording hardware, we used a transparent µECoG array which recorded local-field potentials (LFP) from the surface of the primary sensorimotor cortex (Yazdan-Shahmorad, Diaz-Botia, et al. 2016).

For each experimental session we randomly chose two locations corresponding to two electrodes of the µECoG array to be stimulated. The stimulation consisted of alternating 5 ms laser pulses delivered to the stimulation sites, temporally offset by a session-specific delay (Fig. 1) and repeated every 200 ms for 10 minutes. The session-specific delays were randomly chosen at the beginning of each experimental session to be 10 ms, 30 ms, or 100 ms. These values were chosen because 10 ms and 30 ms have been shown to be in the effective range for connectivity change while 100 ms is outside the effective range according to STDP (Bi and Poo 1998; Jackson, Mavoori, and Fetz 2006). Locations of stimulation, the delay between the stimulation of the first and second sites, and the location and orientation of the µECoG array placement varied between experimental sessions but were consistent over individual sessions.

The set of experimental sessions spanned multiple days and two rhesus macaques. We recorded neural activity during 10-minute stimulation periods, which we refer to as "stimulation blocks," as well as during 5-minute periods before and after each stimulation period, which we refer to as "resting-state blocks." Each experimental session consisted of alternating resting-state blocks and stimulation blocks, with 6 total resting-state blocks and 5 total stimulation blocks.

From the neural data recordings we calculated the FCC for each pairwise connection between recorded electrodes of an experimental session. We quantified FCC with the coherence metric for the Theta (4-7 Hz), Beta (12-30 Hz), Gamma (30-70 Hz), and High-gamma (70-200 Hz) bands.

As the effects of stimulation in both the stimulated-state and in the resting-state are relevant for neural disorders therapies (Levy et al. 2016; Lozano and Lipsman 2013), we estimate and model the FCC in both states. We calculated the change from a resting-state block to the following stimulation block which we refer to as SS-FCC, and calculated the change from a resting-state block to the following resting-state block which we refer to as RS-FCC.

Stimulation delay is a poor predictor of resting-state functional connectivity change between stimulation sites

Studies have shown that change in anatomical connectivity between two monosynaptically connected neurons is a function of activation times, or stimulation delay, between the neurons (Bi and Poo 1998; Bliss, T.V.P. and Collingridge 1993; Markram et al. 1997), a phenomenon termed STDP. As STDP relates delay to structural changes in connectivity but not real-time modulation of neural activity, we evaluated how well stimulation delay could predict RS-FCC between the sites of stimulation.

The standard function relating stimulation delay to corresponding connectivity change according to the STDP hypothesis is an exponential function (S. Song, K. D. Miller, and Abbott 2000; Zhang et al. 1998). However, here we recognize that our experimental setting bears little resemblance to the controlled monosynaptic studies yielding such a curve. As such, we did not enforce an exponential delay function but rather tested whether varying the delay had an effect on RS-FCC by fitting a linear model with delay as categorical input and RS-FCC as continuous output (Fig. 2a). Since STDP is a process governing pairwise connectivity change, we modeled only RS-FCC calculated between stimulation sites, and did not model RS-FCC nor connectivity between unstimulated sites. We observed that delay was a significant predictor

of RS-FCC in Theta and Gamma bands, but not for Beta or High-gamma (p: Theta 0.016, Beta 0.125, Gamma 0.002, High-gamma 0.239; one-sided F-test with 165 observations, 2 degrees of freedom). For all frequencies we observed that delay had low explanatory power (R^2: Theta 0.050, Beta 0.025, Gamma 0.075, High-gamma 0.018). These results indicate that, although a naive extrapolation of the STDP curve to our large-scale setting would hint otherwise, delay between paired stimulation of two neural sites minimally controls RS-FCC.

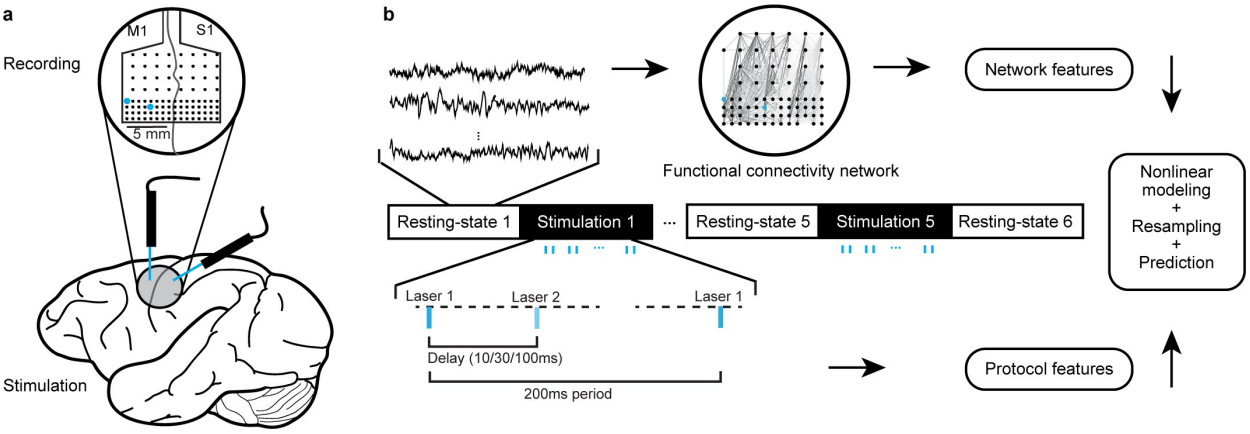

Figure 1: Schematic of the experimental setup and basic modeling pipeline. **a**, Depiction of the transparent 96-electrode μECoG array which we used to record from macaque primary sensorimotor cortex, and the approximate location of the array on the cortex. Optogenetic stimulation was applied via laser light. **b**, We recorded traces of neural LFP during the resting-state and stimulation blocks making up our experimental sessions. During the stimulation blocks, stimulation was applied in a paired pulse protocol whereby the stimulation at the second site followed the first by a session-specific delay. Network coherences were constructed for each individual recording block by calculating pairwise LFP coherences over the network. Network coherence from resting-state recordings were parameterized to generate a set of network features. Details of the protocol was parameterized to generate a set of protocol features. The sets of network and protocol features were fed to a nonparametric model to predict network-wide FCC.

Pairwise stimulation drives functional connectivity changes in the stimulated-state and resting-state across the entire recorded network

In order to determine whether our stimulation induced FCC only between the stimulation sites or over a broader area of the recorded network, we examined the network-wide FCC. We plot an example network of FCC over our recording array in Fig. 2b. In this example, we see that the connectivity change between stimulation sites (bolded) is simply one change among many in the recorded network. To quantify the network-wide FCC over all experiments we calculated the FCC distributions. We observed that stimulation generally increased network-level functional connectivity, as SS-FCC had positive mean over all frequency bands. We additionally compared the SS-FCC of our stimulation experiments to a set of control experiments in which no stimulation was applied, and found that stimulation resulted in significantly higher SS-FCC ($p < 0.0001$ in all frequency bands; Welch's unequal variances t-test) (Fig. S1.)

We then compared the SS-FCC and RS-FCC distributions. We found SS-FCC was highly correlated to RS-FCC, as the Pearson's correlation coefficient was above 0.6 in all frequency bands (Fig. 2d). The means of RS-FCC distributions were less offset from 0 than the means of SS-FCC distributions of the same frequency bands, and variances were similar between RS-FCC and SS-FCC distributions of the same frequency bands. These widespread FCC raise the question of explaining and predicting the stimulation-induced FCC across

the entirety of the recorded network.

Nonlinear modeling to predict stimulation-induced network-wide functional connectivity changes

Motivated by the observation of network-wide FCC, we pursued an accurate and interpretable model to predict FCC over the entire recorded network. While linear modeling is the standard method for statistically rigorous interpretable models, the assumption of linearity is scientifically unreasonable for statistical relationships among complex experimental measurements such as neural data (Friedman, Hastie, and Tibshirani 2001). We relax this assumption through the use of a nonparametric hierarchical additive model which can identify complex nonlinear input-output relationships (Haris, Shojaie, and Simon 2019). As this model retains the structure of summation over individual features familiar from linear modeling, it permits feature-wise investigation of the results (Fig. 3c-f). In addition, careful penalization of the estimation objective (Eq. 18) ensures that the model retains important theoretical guarantees and is encouraged to learn simple, smooth components (Haris, Shojaie, and Simon 2019). In the following sections we use such nonparametric hierarchical models to arrive at our results. These models outperformed linear models uniformly across frequency bands and prediction tasks (Fig. S2).

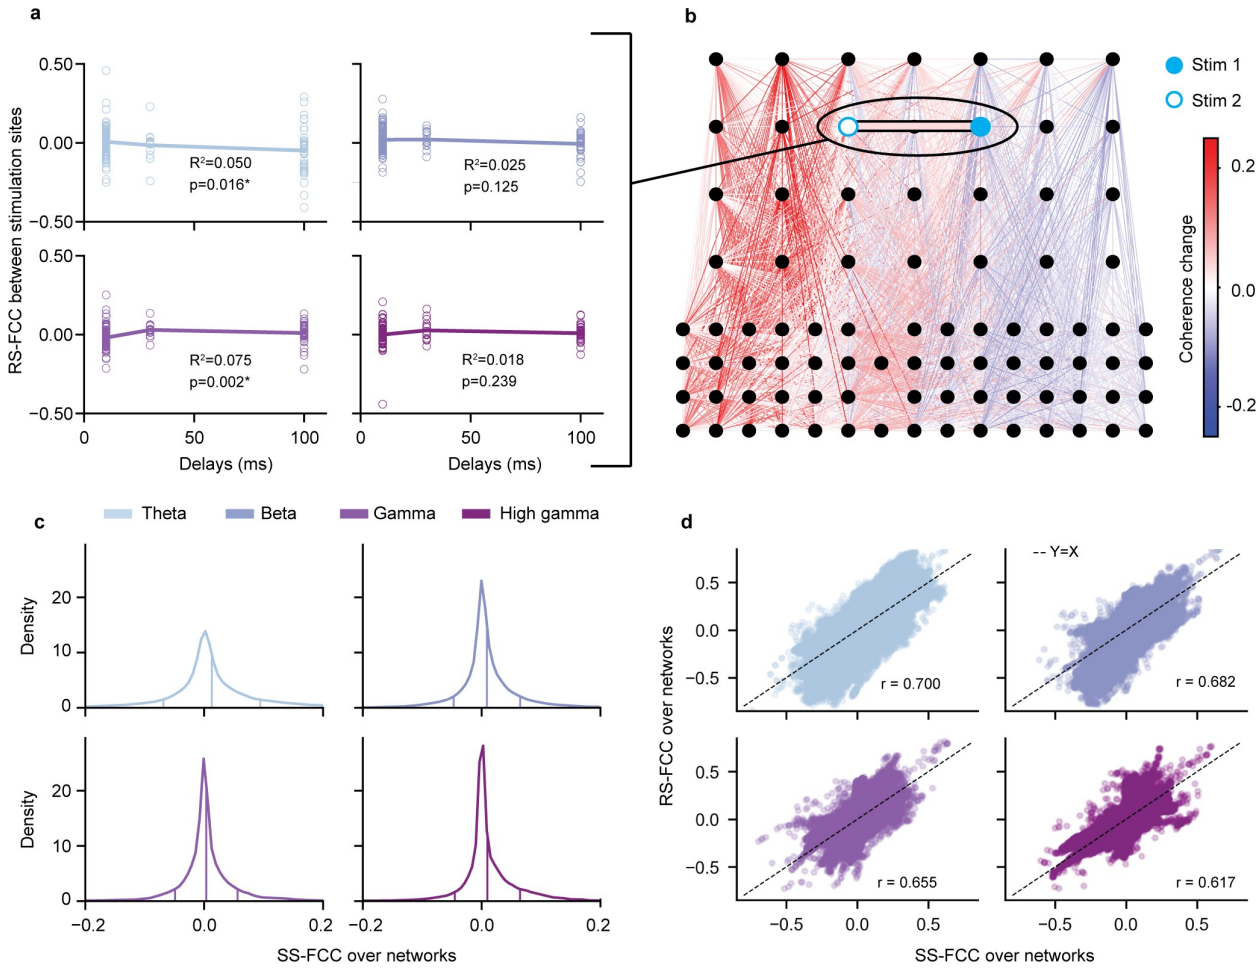

Figure 2: **a**, Delay-based prediction of RS-FCC between stimulation sites has poor explanatory power, and does not achieve statistical significance in the Beta and High-gamma. **b**, Visualization of the RS-FCC of a single example session over the µECoG array. Black circles correspond to electrodes while the edges indicate the change in

coherence relative to the preceding baseline period. Optogenetic stimulation sites are indicated by light blue circles. The bolded line connecting the stimulation sites indicates the coherence change between the two sites. **c**, Distribution of the SS-FCC in each frequency band over all sessions. Internal lines indicate the mean, and one standard deviation below and above the mean. **d**, Scatter plot of SS-FCC versus RS-FCC showing high correlation.

The stimulation protocol alone is a poor predictor of network-wide functional connectivity changes

To test the degree to which the factors of the stimulation protocol controlled the network-wide FCC, we constructed models to predict SS-FCC and RS-FCC from the stimulation protocol. We constructed a set of 8 features consisting of: stimulation delay, a categorical identifier of the time ordering of the stimulation block, the distance between a given electrode pair, the distances from the sites of stimulation to the closer or further electrode of an electrode pair, anatomical region of electrodes (Fig. 3a) (see Methods for more details of how they are calculated). Throughout the paper we refer to this set of predictive features as "protocol features." We also include a categorical animal subject indicator in the model in order to capture differences in mean between the two subjects. With this feature set we trained models with output of pairwise FCC between all recording electrodes of an experiment, over each block of all experiments. We fit individual models to predict SS-FCC or RS-FCC over individual frequency bands. We held out 30% of the data to evaluate the predictive performance of the model on unseen data.

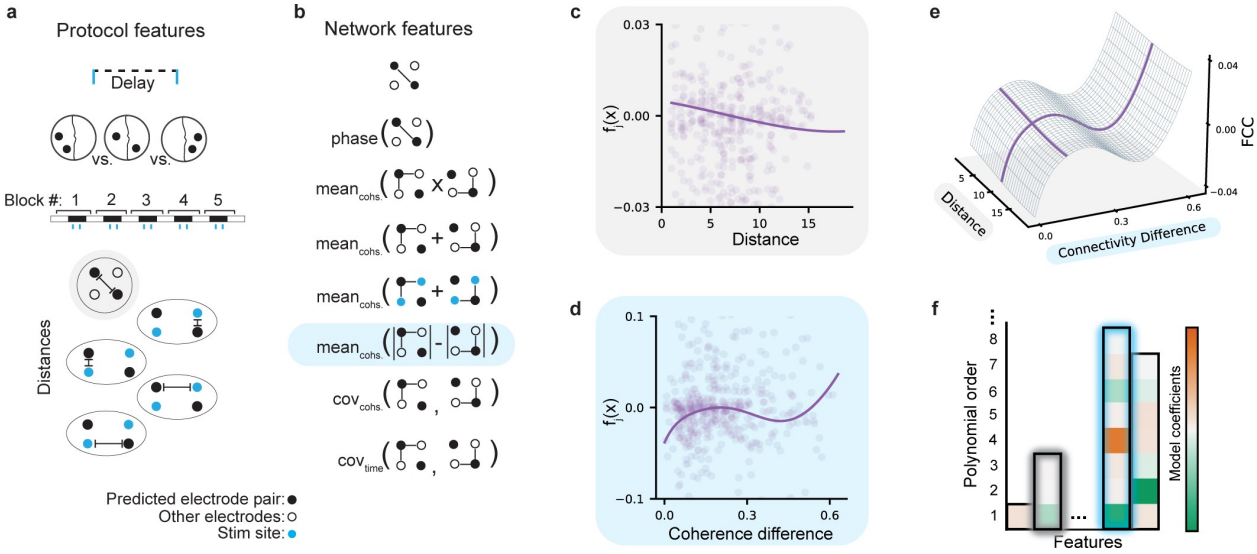

Figure 3: Feature generation and predictive modeling. **a**, Visualization of the protocol features. From top to bottom, the features are: stimulation delay, anatomical region, block number, distance, stim 1 distance to closer, stim 2 distance to closer, stim 1 distance to further, stim 2 distance to further. **b**, Visualization of the network characteristics which the network features quantify. From top to bottom, the visualizations correspond to: initial coherence, phase, length 2 path strength, coherence with network, coherence with stim sites, coherence difference, electrode covariance, and time covariance. **c-d** The nonparametric model can identify feature mappings of a variety of complexities, such as (**c**) near-linear fits such as that found for the distance feature, or (**d**) higher-order polynomials such as that found for the connectivity difference feature. $f_j(x)$ indicates the contribution of a feature to the predicted FCC. Throughout the figure, the gray shadowing indicates the distance feature, while the light blue shadowing indicates the connectivity difference feature. **e**, An example of how the feature mappings are additively combined. **f**, Individual feature mappings are represented in a basis of increasingly complex nonlinear functions. The order of this representation is selected automatically by a penalized estimation procedure, with the y-axis of this figure indicating the polynomial order. The order selection for the distance feature and coherence difference feature are shadowed in gray and light blue, respectively.

We observed that the protocol-based model of FCC had low predictive power for both SS-FCC (R^2: Theta 0.045, Beta 0.026, Gamma 0.057, High-gamma 0.078) and RS-FCC (R^2: Theta 0.028, Beta 0.014, Gamma 0.020, High- gamma 0.014). These results indicate that while factors related to the protocol alone explain some of the variation in the observed coherence change, they are insufficient to generate practically relevant predictions.

Network structure determines stimulated-state and resting-state functional connectivity changes

Our observations that stimulation induces a change in connectivity over an entire network, in tandem with the poor predictive power of the protocol-based model, lead us to hypothesize that the effects of stimulation may be mediated by the existing cortical network. We tested this hypothesis by constructing a set of 8 features extracted from the baseline functional connectivity network. These features are: initial coherence, which is the pairwise coherence between an electrode pair; phase, which is the absolute value of difference between electrode spectral phases; length 2 path strength, which is the average strength of a connection between an electrode pair passing through another recording site; coherence with network, which is the average coherence between a given electrode pair and the rest of the network; coherence with stim sites, which is the average coherence between a given electrode pair and the stimulation sites; coherence difference, which is the average absolute difference in connection strength between each element of the electrode pair and the recorded network; electrode covariance, which is the average covariance of an electrode pair's coherence to the other electrodes in the network over time; and time covariance, which is the average covariance across time series between a given electrode pair's coherence and the others in the network (Fig. 3b) (see Methods for details of how these are calculated). These features are primarily derived from graph theoretic measures (Barrat et al. 2004; R. Milo et al. 2002; Ron Milo et al. 2004; Onnela et al. 2005) which have been used with success in past theoretical neuroscience analyses (Hu et al. 2014; Ocker et al. 2017; Pernice et al. 2011; Zhao et al. 2011). Throughout this paper we refer to this set of features as "network features."

We observed that the model predictions on the test set was dramatically improved compared to the protocol- only model for each frequency band for both SS-FCC (R^2: Theta 0.240, Beta 0.130, Gamma 0.221, High-gamma 0.323) (Fig. 4) and RS-FCC (R^2: Theta 0.218, Beta 0.087, Gamma 0.138, High-gamma 0.218) (Fig. 6a). This improvement in performance across all frequency bands and for both stimulated-state and resting-state contexts suggests that properties of the local network at baseline may be highly informative with respect to prediction of stimulation-induced changes.

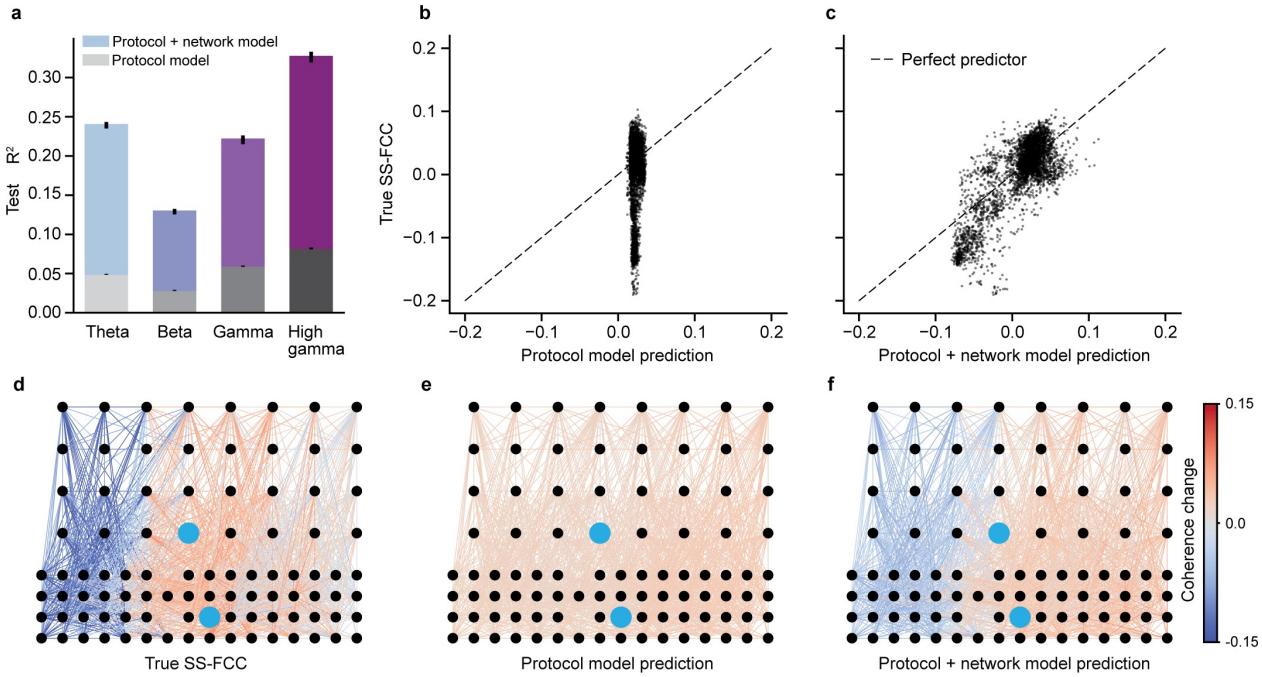

Figure 4: A model trained to predict stimulation-induced FCC over a cortical network based on protocol features alone is outperformed by one which incorporates information about existing network connectivity. Results in this figure correspond to prediction quality for SS-FCC; results are similar but slightly worse for RS-FCC. **a**, Bar graph indicating r-squared accuracy on held-out data. Grayscale bars indicate accuracy of models trained on only protocol features and colored bars indicate accuracy of models trained on both protocol and network features. Error bars indicate standard deviation. **b-f**, Results of prediction of a representative experimental session from the protocol versus the protocol-and-network models. **b**, Scatter plot of the protocol model's predictions of network SS-FCC versus the actual changes observed during that experimental session. **c**, Scatter plot of the protocol-and-network model's predictions of SS-FCC versus the actual changes observed during that experimental session. **d**, Example SS-FCC over a cortical network, visualized over the µECoG array used for recording. Black circles indicate electrodes of the array, while lines indicate coherence change between electrode pairs. The observed coherence changes correspond to the y-axis of **b** and **c**. **e**, The protocol model's SS-FCC prediction, corresponding to the x-axis of **b**. **f**, The protocol-and-network model's SS-FCC prediction, corresponding to the x-axis of **c**.

Individual network features are more important than individual stimulation protocol features for predicting functional connectivity changes

We quantified the relative importance of each feature through the decrease in accuracy of prediction on the test set when that feature was excluded from the data. The statistical uncertainty associated with the importance was quantified through a resampling procedure (Meinshausen and Bühlmann 2010) (see Methods for details).

Using this procedure, we show the importances of each feature per frequency band in Fig. 5a for SS-FCC and in (Fig. S3) for RS-FCC. Four of the five most important features for SS-FCC prediction and each of the five most important features for RS-FCC prediction were network features. Features appearing in the top five of both contexts were time covariance, coherence difference, and initial coherence, which were all network features. Delay and time-ordered block number were the two most important of the protocol features for both SS-FCC and RS-FCC prediction.

Feature mappings display diverging trends across frequency bands

We decomposed the nonlinear model into feature mappings representing the estimated mapping between each individual feature and the change in coherence. The statistical uncertainty associated with

the estimated nonlinear mapping from the feature to the response was quantified through a resampling procedure (Meinshausen and Bühlmann 2010) and representative feature mappings were obtained through point-wise interpolation of the resampled feature mappings (see Methods for details).

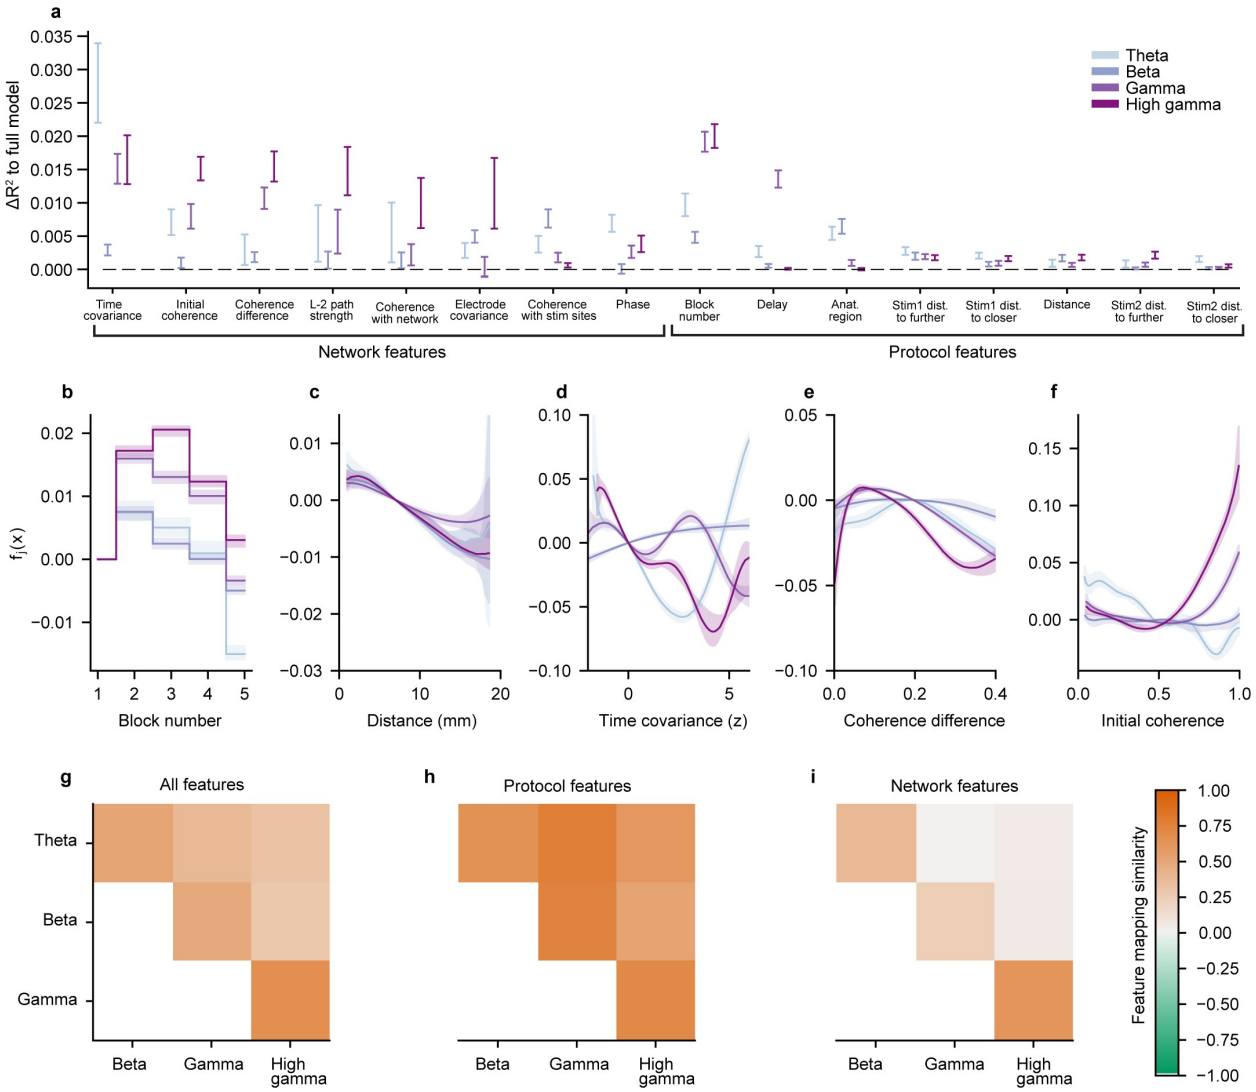

Figure 5: Feature-wise analysis of the nonlinear additive model. Results in this figure correspond to the SS-FCC model; results were similar for the RS-FCC model. **a**, Feature importances as measured by the decrease in model accuracy when the feature was omitted. Features were partitioned into the protocol and network categories, then sorted by average importance over all frequency bands. Line ranges indicate 95% confidence intervals. **b-f**, Feature mappings for the **b** time-ordered block number, **c** electrode pair distance, **d** time covariance (z-scored), **e** coherence difference, and **f** initial coherence features. $f_j(x)$ indicates the contribution of a feature to the predicted FCC. Shading indicates 95% confidence intervals. **g-i**, Average cosine similarity of the component functions between each pair of frequency bands, over **g** all features, **h** protocol features, and **i** graph features.

We plot the SS-FCC feature mappings of some important features as well as others of interest in Fig. 5b-f. We observed that feature mappings were generally more complicated for network features than for protocol features. For example, the time covariance feature varies from a near-linear increase in the Beta band to a high-order polynomial in the High gamma band, while the distance feature mapping exhibits a nearly linear decrease in the high-confidence regions of all frequency band models (Fig. 5c). We also

observed that in the higher frequency bands, electrodes which had a higher initial coherence tended to strongly increase their coherence when stimulated (Fig. 5f). Additionally, we found a small effect of increasing SS-FCC for the first two or three temporal blocks depending on frequency band, followed by decreasing SS-FCC in later blocks (Fig. 5b). We evaluate the temporal characteristics of the FCC in more detail in a later section.

We summarized the similarity of all features, protocol features, and network features by calculating the cosine similarity of the feature mappings (Fig. 5g-i). Generally, we observed that the mappings from protocol features to SS-FCC were fairly consistent across frequency bands, while the mappings from network features to SS-FCC varied sharply between frequency bands. Only the network feature mappings of Gamma and High-gamma were similar.

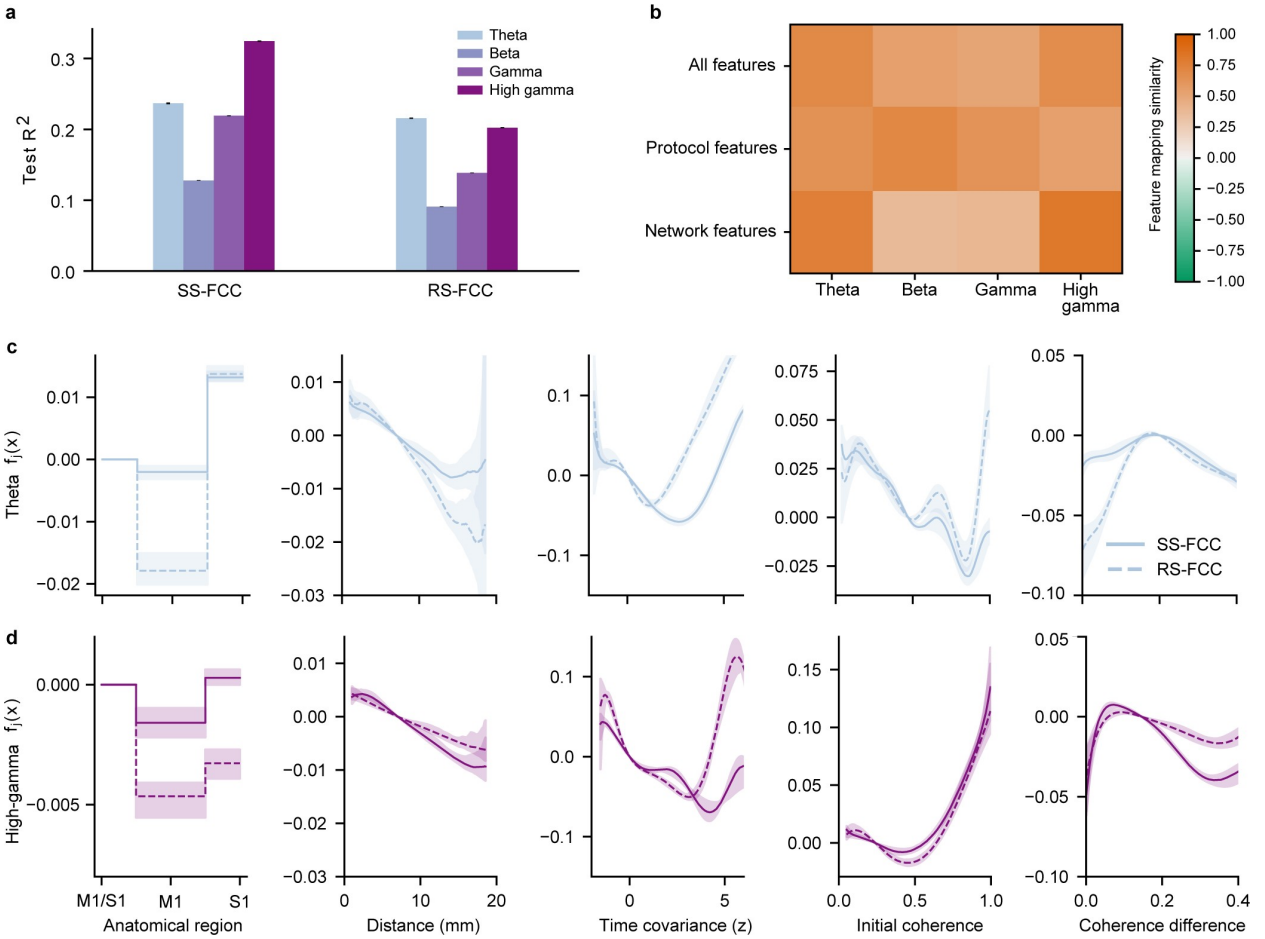

Figure 6: **a**, Comparison of the predictive accuracy for nonlinear additive models of SS-FCC prediction versus RS- FCC prediction. Error bars indicate standard deviation. **b**, Heatmap of average cosine similarity between feature mappings of SS-FCC and RS-FCC over frequencies. **c-d**, Comparison of feature mappings for SS-FCC and RS-FCC models for the **c** Theta band and **d** High-gamma band. The features from left to right are anatomical region of electrodes, electrode pair distance, time covariance (z-scored), initial coherence, and coherence difference. Shading indicates 95% confidence intervals.

Feature mappings are similar for stimulated-state and resting-state functional connectivity changes

We calculated the similarity of feature mappings between SS-FCC and RS-FCC and observed that they exhibited high similarity in all frequency bands (Fig. 6b). Notably, the groups with the highest similarity between SS-FCC and RS-FCC feature mappings were the network features in the Theta and High-gamma

bands. We plot representative SS-FCC and RS-FCC feature mappings of five features from the Theta and High-gamma models in Fig. 6c-d.

Repeated stimulation modifies mappings from features to functional connectivity changes

The modeling approach employed thus far constrained the learned feature mappings to be constant across all temporal blocks of an experimental session. The influence of the block's time location was represented by a learned global shift in the FCC between all electrode pairs measured in that block. We hypothesized that successive stimulation blocks within an experiment had a more complex effect on FCC than a simple average shift, and rather modified the actual feature mappings over the time course of stimulation. We tested this hypothesis by allowing the nonlinear relationship between each feature and the response to vary across the five stimulation blocks by adding an interaction term between the block number feature and all other features. Rather than averaging over potential heterogeneity in the network dynamics resulting from the accumulated effects of stimulation, this approach allowed for the investigation of a time-course of smoothly varying predictors of coherence changes.

In line with the hypothesis that repeated stimulation indeed produces such heterogeneity across blocks within a session, we found a substantial improvement in prediction quality with the time-varying model for both SS-FCC (R^2 Theta 0.426, Beta 0.247, Gamma 0.415, High-gamma 0.527) (Fig. 7a) and RS-FCC (R^2 Theta 0.410, Beta 0.251, Gamma 0.359, High-gamma 0.436) over the static model. We also calculated the cosine similarities of feature mappings for each block of the time-varying SS-FCC model and found that successive stimulation results in the mappings diverging from each other, as their similarity decreased over subsequent blocks (Fig. 7d). These results indicate that the relationship between features and connectivity change evolved dynamically over the course of repeated stimulation.

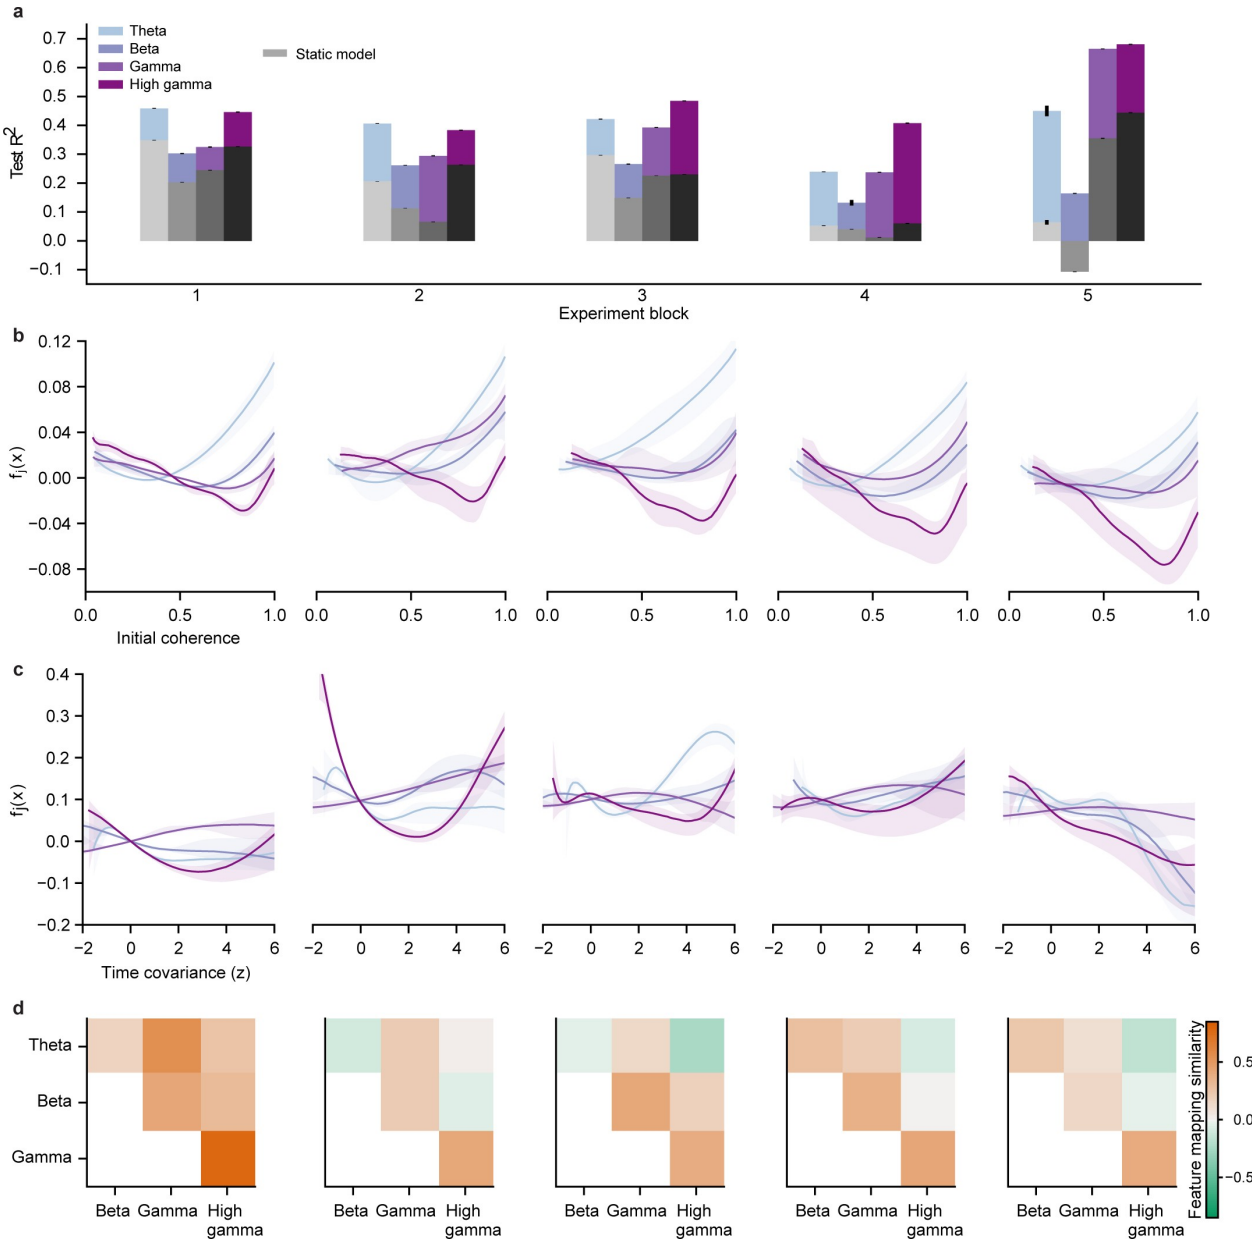

Figure 7: Feature mappings mediating FCC over a network are dependent on the stimulation history. Results in this figure correspond to prediction of SS-FCC; results are similar for prediction of RS-FCC. **a**, Bar graphs indicating r-squared accuracy on held-out data of the time-varying and static model predictions. Greyscale bars indicate the static model while colored bars indicate the time-varying model. The time-varying models outperformed the static model over all frequencies and experiment blocks. Error bars indicate 95% confidence intervals. **b-c**, Example feature mappings of two network features, **b** initial coherence and **c** time covariance (z-scored) used in the time-varying model. We observed that the feature mappings change over successive experiment blocks. Shading indicates 95% confidence intervals. **d**, Heatmaps of similarities of feature mappings between models trained for different frequency bands. Similarity was quantified by feature-averaged cosine similarity of the feature mappings. We observed that while feature mappings were similar at the onset of the stimulation, they diverged over the time-course of stimulation.

Discussion

Here we leverage advances in neural interface technology and interpretable machine learning to induce and model large-scale FCC. We augment the standard framework of stimulation-induced FCC by considering the effects of the pre-existing network structure and by analyzing FCC over a large-scale ($1cm^2$) network. We show that stimulation induces widespread FCC in both the stimulated-state and the resting-state, and that the underlying network structure mediates these effects. Our modeling framework represents a readily translatable method of characterizing the underlying network structure to predict the outcome of stimulation. Our findings explain inconsistent results historically reported *in vivo* (Bloch et al. 2019; Huang et al. 2019; Keller, Huang, et al. 2018; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013; Yazdan-Shahmorad, Silversmith, et al. 2018) and pave the way for increased understanding and development of neural stimulation interventions.

The influence of the stimulation protocol

While brain stimulation can be delivered to a local area, the highly-interconnected nature of the brain ensures that any local region tightly modulates and is modulated by other regions. An analysis of the effects of stimulation which only considers the stimulation protocol thus prioritizes the direct, local effect of stimulation over the mediation of the response by the brain network.

For neural stimulation interventions targeting longer-term FCC *in vivo*, the protocol-focused approach has been found lacking. The dominant framework for stimulation-induced targeted connectivity change was developed *in vitro*, a setting in which the cortical network is largely nonexistent and the stimulation protocol can be extremely precise. Attempts to translate STDP from its native context of monosynaptic connectivity between isolated neuron pairs *in vitro* to functional connectivity between larger-scale neural structures *in vivo* have been largely inconsistent, with some reports of promising results (Jackson, Mavoori, and Fetz 2006) but also of off-target (unstimulated) connectivity changes (Bloch et al. 2019; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013) and a lack of response to the protocols (Bloch et al. 2019; Seeman et al. 2017).

Here, we study real-time network-level FCC as SS-FCC and longer-term FCC as RS-FCC. We found that, in accordance with existing protocol-focused approaches, the stimulation protocol influences both SS-FCC and RS-FCC. However, further analysis indicated that this influence is minor, and that a protocol-focused analysis of stimulation only explains a small fraction of the effects. This is a novel finding, and represents a substantial departure from the standard analysis of the effects of stimulation which prioritizes the stimulation protocol.

Our observation that the stimulation delay was able to predict some RS-FCC between stimulation sites, and that the stimulation delay was one of the top two protocol features for network-wide FCC prediction, is aligned with the delay-focus of STDP-informed connectivity modification approaches (Bi and Poo 1998; Bliss, T.V.P. and Collingridge 1993; Jackson, Mavoori, and Fetz 2006; Markram et al. 1997; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013). However, the small amount of variance explained by delay indicates that delay alone is not sufficient for accurate FCC control.

Network-level analysis of neural stimulation

Here, we model the cortical network as both the mediator and target of the effects of stimulation. Our focus on the network structure is facilitated by our use of novel neurotechnologies for stimulation and recording. We used a large- scale µECoG array to record from a large-scale cortical network. This differed from past studies in which the inclusion of a small set of unstimulated neural sites had primarily been

relegated to that of a statistical control (Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017). In addition, we used convection-enhanced delivery of the optogenetic virus to obtain widespread opsin expression (Khateeb et al. 2019; Yazdan-Shahmorad, Diaz-Botia, et al. 2016), which allowed us to stimulate many distinct sites of the sensorimotor cortex and to record neural activity during stimulation without artifact.

We estimated a complete network between all recording sites by calculating their pairwise coherences, which is highly correlated to other functional connectivity metrics (Huang et al. 2019; Yazdan-Shahmorad, Silversmith, et al. 2018). In addition, unlike metrics such as evoked response, coherence can be calculated in the absence of stimulation. This had the dual advantage of allowing us to estimate connectivity during resting-state blocks and between unstimulated sites. We note that the latter point is especially relevant, as most "network-level" analyses of stimulation only observe a partial network as they do not analyze connectivity between unstimulated regions (Huang et al. 2019; Keller, Bickel, et al. 2011; Keller, Huang, et al. 2018; Yang et al. 2021).

We then incorporated details of the distributed structure of these networks for prediction of FCC. Recent studies have indicated that pairwise coherence between two sites is correlated to their change in coherence following stimulation (Bloch et al. 2019; Keller, Huang, et al. 2018; Khambhati et al. 2019). However, any pairwise analysis treats connections as independent from another and does not capture the distributed structure of the network. In order to incorporate details of the distributed network structure into our predictive model, we turned to advances in graph theory and theoretical neuroscience. Graph theory is a branch of mathematics concerned with the study of graphs, and has developed metrics for summarizing and comparing graphs, such as clustering coefficients, measurements of centrality, and graph distances (Barrat et al. 2004; R. Milo et al. 2002; Ron Milo et al. 2004; Onnela et al. 2005). Theoretical neuroscience has similar measures, termed "motifs," which are used to analyze simulated brain networks (Hu et al. 2014; Ocker et al. 2017; Pernice et al. 2011; Zhao et al. 2011). Just as simple pairwise connectivity between electrodes describes their first-order connectivity, these measures summarize higher-order, distributed connectivity within a network (Benson, Gleich, and Leskovec 2016). The network features we developed and used in our model were inspired by these existing metrics and similarly encode higher order, distributed connectivity information of the network. For example "coherence with network" and "length 2 path strength" quantify general whole-network connectivity strength between two sites, while "electrode covariance" and "coherence difference" quantify the similarity of connectivity structure to the entire network between two site. Indeed, our model confirmed their relevance by identifying "time covariance" and "coherence difference," two higher-order connectivity metrics, as the most important features for FCC prediction.

One alternative framework for inducing targeted network-wide FCC is network control theory (Bassett and Sporns 2017), which attempts to use control-theoretic methods to control network activity in real-time. However, it has thus far has been limited to quantifying theoretical controllability of networks (Khambhati et al. 2019; Muldoon et al. 2016; Yang et al. 2021) rather than inducing targeted FCC. Here, we validate our framework on the maximally clinically-translatable NHP model and demonstrate that it can accurate predict network-level stimulation-induced FCC.

The cortical network as the mediator and target of stimulated-state and resting-state functional connectivity changes

Both real-time and long-term modulation of brain networks are relevant for neural stimulation-based medical interventions. Here, we induced, estimated, and modeled network-level FCC which we observed during stimulation as well as after stimulation. The high correlation in FCC and high similarity in feature mappings between the two contexts hint at shared processes underlying SS-FCC and RS-FCC. Practically, this also means that observing stimulation-induced FCC in real-time may lead to more accurate longer-

term FCC.

Interventions targeting real-time neuromodulation often do not consider effects beyond local activity modulation. However, this neuromodulation has a secondary effect of altering real-time network synchronization, observed here as network-level SS-FCC, and which persist after stimulation as RS-FCC. As such, the complete effects of these interventions can only be fully understood by augmenting their scope of analysis to include network-level FCC.

Our observation that the underlying network mediates RS-FCC demonstrates that interventions directly targeting RS-FCC must consider the underlying network for accurate induction of connectivity changes. The historical neglect of the underlying network structure in mediating RS-FCC could provide an explanation for reports of off-target connectivity changes and inconsistent connectivity changes between stimulation sites (Bloch et al. 2019; Rebesco and L. E. Miller 2011; Rebesco, Stevenson, et al. 2010; Seeman et al. 2017; W. Song et al. 2013; Yazdan-Shahmorad, Silversmith, et al. 2018). We note that, while many of these studies use STDP-informed stimulation which is designed to elicit changes in anatomical connectivity, we cannot pinpoint the nature of our observed FCC as either solely anatomical or activity-based, as functional connectivity is related to both (Waites et al. 2005; Wang et al. 2013).

Finally, our results offer a potential explanation for how underlying network-level functional connectivity shapes the therapeutic response to stimulation. Studies have shown that aberrant functional connectivity underlies many neural disorders (Garrity et al. 2007; Kaiser et al. 2015; Nakai et al. 2021; Stam et al. 2007), and that neural stimulation alters network-level functional connectivity (Huang et al. 2019; Keller, Huang, et al. 2018; Yazdan-Shahmorad, Silversmith, et al. 2018) and can treat neural disorders (Edwardson et al. 2013; Lozano and Lipsman 2013). As such, therapeutic response to stimulation may be a direct result of a stimulation-induced change from aberrant to healthy network- level functional connectivity. This change may in turn be determined by the underlying network-level functional connectivity, as recent work has indicated that underlying functional connectivity shapes the therapeutic response to stimulation (M. D. Fox, Buckner, et al. 2012; Horn et al. 2017; Younce et al. 2021), and here we have also shown that it shapes stimulation-induced network-level FCC. Further work which holistically studies the underlying network, neural stimulation, FCC, and neural disorders is needed.

The evolving response to successive neural stimulation

Our observation that the time-ordered block number was the most important protocol feature indicates that the temporal location of repetitive stimulation is important for consideration. This observation is in accordance with earlier work from our lab (Bloch et al. 2019) in which we reported sharp differences in functional connectivity between an initial stimulation block following stimulation blocks. We further investigated this phenomenon with our time-varying model.

With the time-varying model we observed that feature mappings describing the response to stimulation evolve in response to repeated stimulation. This result indicates that stimulation interventions likely need to be adapted over time, as repeatedly stimulating will not repeatedly induce the same response nor even be governed by the same network rules mediating the response. This is consistent with the increased success of some closed-loop neural stimulation interventions over open-loop alternatives (Edwardson et al. 2013), as it acknowledges that the effects of stimulation can be made more accurate when the stimulation is informed by recent neural activity. However, we additionally show that stimulation should be adapted to characteristics of the whole network, not only to local activity.

Interpretable predictive modeling beyond the linear paradigm

In this study we capitalized on our network-level treatment of FCC by using a data-driven nonlinear

modeling approach. In doing so, we were able to develop an accurate predictive model and estimate feature mappings displaying the relationships between the protocol and network features and FCC. These feature mappings have the dual purpose of elucidating underlying mechanisms of network-level response to stimulation, and informing future stimulation interventions.

Data-driven statistical models often fall into one of two approaches: simplistic linear models or complex "black- box" models. While the simple relationships identified by linear models ensure that the mappings are interpretable, linearity is an assertion, not a property learned from the data, and therefore linear models may not capture the complexity of the true input-output relationships. At the other extreme, prediction-focused black-box models can learn complex input-output mappings at the expense of interpretative ability such as observation of the mappings. The unique modeling approach we employ in our study displays advantages of each as it can identify complex mappings between the features and FCC while allowing the features mappings to be observed. As it is a data-driven and generalized additive model, it retains the ability of identifying linear mappings when such a mapping is supported by the data, while also being able to identify more complex nonlinear mappings when appropriate.

We found that eliminating the artificial and restrictive assumptions of the linear model yielded dramatically improved predictive performance. In particular, as the complexity of our input features increased from simple stimulation parameters to the actual structure of the cortical network, their feature mappings exhibited strong nonlinearities and thus would not be able to be fully represented by a strictly linear approach.

Finally, in light of recent unsuccessful attempts to model the effects of neural stimulation with nonlinear models (Yang et al. 2021), our data-driven method for inferring feature mappings of variable complexity represents a timely contribution for principled and accurate modeling of the effects of brain stimulation.

Neuroscientific and clinical utility of our insights

Here we have shown that the cortical network structure mediates the response to stimulation, both for real-time and longer-term contexts. Our findings provide a possible explanation for how underlying functional connectivity shapes the therapeutic response to stimulation, and for previously reported inconsistent results of neural stimulation. Our framework, only made possible by novel advances in neural interfaces, interpretable machine learning, and graph theory, represents an efficient method for parsing network structure for accurate prediction of stimulation-induced FCC. It can be readily applied to other network-informed approaches of neural stimulation, and can be used to interrogate, improve, and develop novel neural stimulation interventions for neural disorders.

Limitations of the study

While we have shown that FCC induced by stimulation are more positive than FCC observed in the control condition, it is possible that the modeled FCC in the stimulation condition are composed of a combination of stimulation-driven changes and natural fluctuations. Regardless, the model represents an accurate method for predicting functional connectivity changes in the presence of stimulation.

Our neural recordings were limited to a µECoG on the primary sensorimotor cortex surface; further work is needed to verify that our findings are consistent across brain regions, at various cortical depths, and at larger and smaller scales. In addition, as all stimulation experiments consisted of paired optogenetic stimulation, our observations should be verified with other stimulation modalities and protocols to verify their generality. Our method for parsing the network structure, along with the hierarchical additive model, can be easily extended to these cases. If our findings are verified across a broad range of brain regions and stimulation modalities, then this would indicate that existing stimulation-based therapies must factor in the underlying network in order to arrive at a complete understanding of their effects.

Here we have constructed a set of network features which we have shown are relevant for FCC prediction. While some interpretable mappings of interest were described in the results section, others which were more complicated, such as the most-important "time covariance" feature, could not be easily interpreted. Our features broadly encapsulate various spatiotemporal connectivity motifs present in the network connectivity, and as such can be further dissected and parsed for purposes of interpretation. For example, the importance of "time covariance" indicates that the temporal relations of pairwise coherence contains information relevant for FCC prediction; this knowledge can be used to construct more detailed studies with the goal of uncovering the underlying processes governing this relation. Our finding that the network mediates the plastic response to stimulation can guide the development of stimulation treatments, but does not explicitly offer a prescriptive method for specifying an exact stimulation protocol. Instead, our model acts as an encoder which yields a prediction for the effects of stimulation given a specific stimulation protocol and network state. The creation of a corresponding decoder could be used to directly prescribe a stimulation protocol to elicit desired FCC. Such a model could also be deployed in real-time as a network-informed controller for FCC.

Methods

Two adult male rhesus macaques (monkey G: 8 years old, 17.5 Kg; monkey J: 7 years old, 16.5 Kg) were used in this study. All procedures were performed under the approval of the University of California, San Francisco Institutional Animal Care and Use Committee and were compliant with the Guide for the Care and Use of Laboratory Animals.

Large-scale neural interface

Stimulation and recording of the cortex were achieved using our large-scale optogenetic interface consisting of a semi- transparent micro-electrode array, semi-transparent artificial×dura, titanium implant, widespread optogenetic expression, and laser-delivered optical stimulation (Yazdan-Shahmorad, Diaz-Botia, et al. 2016). Neurons in the primary sensorimotor cortex were optogenetically photosensitized via viral-mediated expression of the C1V1 opsin. We injected a viral cocktail of AAV5-CamKIIa-C1V1(E122T/E162T)-TS-eYFP-WPRE-hGH (2.5 10^{12} virus molecules/ml; Penn Vector Core, University of Pennsylvania, PA, USA, Addgene number: 35499) in the primary sensory (S1) and primary motor (M1) cortices of the left hemisphere of the two rhesus macaques using the convection-enhanced delivery technique as described in (Yazdan-Shahmorad, Diaz-Botia, et al. 2016; Yazdan-Shahmorad, Tian, et al. 2018). We infused 200 µL of the cocktail was over four sites in monkey G and 250 µL over five sites in monkey J.

The chronic neural interface was implanted by performing a 25mm craniotomy over the primary sensorimotor cortices of the left hemisphere of two rhesus macaques, replacing the dura mater beneath the craniotomy with a chronic transparent artificial dura, and attaching a titanium cap over the craniotomy. During experiments we removed the artificial dura and placed a custom 92mm^2 micro-electrocorticography array of 96 electrodes on the cortical surface for recording neural activity. The µECoG arrays consisted of platinum-gold-platinum electrodes and traces encapsulated in Parylene-C (Ledochowitsch et al. 2015).

Neural data was recorded by sampling local-field potentials at 24 kHz from the µECoG array using a Tucker-Davis Technologies system (FL, USA). We stimulated the cortex by delivering light via a fiber optic (core/cladding diameter: 62.5/125 µm, Fiber Systems, TX, USA) connected to a 488 nm laser (PhoxX 488-60, Omicron-Laserage, Germany) positioned above the array such that the tip of the fiber-optic cable touched the array. A 488 nm laser was used due to equipment availability and because it is close to the optimal wavelength of C1V1 activation (Prakash et al. 2012). The interface used for the subjects in this study is further described in (Yazdan-Shahmorad, Diaz-Botia, et al. 2016). Of the dataset used for this study, a subset overlaps with the dataset used in our previous publication (Yazdan-Shahmorad,

Silversmith, et al. 2018).

Verification of optogenetic expression and activation

We verified expression of the C1V1 opsin by fluorescent imaging of the eYFP marker (Fig. S4), and by recording optogenetically-evoked neural responses (Fig. S5), as described in (Yazdan-Shahmorad, Diaz-Botia, et al. 2016). We observed widespread fluorescence in the primary sensorimotor cortices (Fig. S4). We also observed neural activity elicited by illumination site, propagating within the stimulated anatomical region and across the central sulcus. An example propagation of neural activity from one experiment is plotted in (Fig. S5), after filtering to the high-gamma which reflects more local activity (Ray et al. 2008). The same effect of activity propagation can be found without filtering, but displays more spatially-distributed low frequency noise. Also, as detailed in the "Signal preprocessing" section, at the beginning of each experimental session we verified that recorded activity was not due to photoelectric artifacts by stimulating at 500 Hz and confirming that we did not see matching LFP traces. We also verified that the evoked neural responses were due to optogenetic activation and not another light-induced factor such heating by illuminating outside the sites expressing fluorescence and recording neural activity. We found that such illumination did not evoke neural activity, confirming that the observed neural activation from our experimental sessions was due to optogenetic activation.

Structure of experimental sessions

Throughout each experimental session, subjects were watching cartoons while seated in a primate chair and headfixed. At the beginning of an experiment the chamber was opened, the artificial dura was removed, and the µECoG array was placed on the cortical surface. Throughout the experiment, the exposed cortex was irrigated with saline to ensure that it remained moist. At the end of each experiment, the µECoG array was removed from the cortical surface, the artificial dura was replaced on the brain, and the chamber was closed.

We performed a total of 36 experimental sessions over multiple days and two subjects. The experimental sessions consisted of recording while no stimulation took place, which we termed a "resting-state block," and recording during stimulation, which we termed a "stimulation block." Each resting-state block was 5 minutes long and each stimulation block was 10 minutes long. Each experimental session consisted of alternating resting-state blocks and stimulation blocks, with 6 total resting-state blocks and 5 total stimulation blocks, as a resting-state block was completed before the first and after the final stimulation block. The µECoG array was replaced on the cortex for each experimental session. Additionally, each recording block of an experimental session had the same stimulation protocol.

We stimulated two locations per experimental session which varied between experimental sessions but remained fixed throughout an experimental session. The locations were randomly chosen across the array. Optogenetic activation at the chosen stimulation site was verified before beginning the experimental session, by stimulating and observing activation in the LFP traces. We ensured that these traces were due to neural activation and not photo- electric artifact by the method described in the "Signal preprocessing" section. We imaged the stimulation locations on the array and subsequently identified the closest recording electrode to each stimulation location. The distance from stimulation site to the nearest electrode was no more than 500 µm.

We applied stimulation using a paired-pulse protocol of 5 ms pulses separated by a session-specific delay (Fig. 1). A pulse width of 5 ms was chosen because it elicited reliable responses in M1 and S1, and has a close duration to previously used targeted plasticity protocols (Jackson, Mavoori, and Fetz 2006). The session-specific delays were randomly chosen at the beginning of each experimental session to be 10 ms, 30 ms, or 100 ms; once chosen they remained fixed for an experimental session. These values were motivated by previous studies reporting that 10 ms and 30 ms delays induced reliable connectivity

changes while 100 ms delays did not induce changes (Bi and Poo 1998; Jackson, Mavoori, and Fetz 2006). The paired-pulses were repeated every 200 ms for the entirety of each stimulation block; this period was chosen to allow for enough time for the chosen delays while maximizing the number of paired-pulse stimulations within a conditioning block.

We performed a total of four additional control sessions over 3 days and two subjects. The structure of the sessions remained identical to the experimental sessions described in this section, except that no stimulation took place.

Signal preprocessing

Before beginning an experimental session, we tested for photoelectric artifact. This was done by stimulating at frequencies faster than C1V1 off-kinetics (eg. 500 Hz) (Prakash et al. 2012) and examining the LFP traces. If this high frequency stimulation elicited response in the LFP traces then this was deemed to be photoelectric artifact, and the illumination site was moved to another location (Ledochowitsch et al. 2015).

Before recording, electrode impedances were measured and those with high impedance were excluded from further analysis. We also examined the broadband recorded surface potentials and excluded those with low signal-to-noise ratio. The total amount of data removed constitutes no more than 15% of the raw time series. We then downsampled the data to 1 kHz after applying a low-pass Chebychev filter for anti-aliasing.

Signal processing to obtain a time-varying coherence network

Within each stimulation block and each resting-state block, the recorded LFP were partitioned into non-overlapping 20-second windows. All quantities computed from the raw LFP time series were computed on a per-window basis, then summarized across all windows in the block.

We computed an estimate of the coherence between all non-faulty electrodes for each 20-second window in every resting-state block and stimulation block. We denote by E the number of non-faulty electrodes and denote the multivariate LFP time series for a given window by the two-dimensional array $x_{1:T} \in R^{E\times T}$, which consists of the time-ordered concatenation of the LFP recordings $x_t = (x_{1t}, \ldots, x_{Et})^T$ for each electrode at time t. We estimated the multivariate spectral density function $S(\lambda) \in C^{E\times E}$ for the time series $x_{1:T}$ by Welch's method (Welch). The coherence between electrodes i and j at frequency λ is denoted as $C_{ij}(\lambda)$ and computed from the estimated spectral density via $C_{ij}(\lambda) = |S_{ij}(\lambda)|^2 / S_{ii}(\lambda)S_{jj}(\lambda)$. We computed the coherence at 403 linearly spaced frequencies in the interval $(0, 200)$ Hz.

Summaries over four frequency bands, 4-7 Hz (Theta), 12-30 Hz (Beta), 30-70 Hz (Gamma), and 70-199 Hz (High Gamma), were computed by averaging the coherence estimates within each band. Thus, for each session, the raw LFP time series was transformed into a sequence of matrix-valued quantities representing the average band-limited coherence in each 20-second interval. Since the coherence is symmetric and the diagonal conveys no information on pairwise behavior of the LFP electrode signals, we retained only values in the upper triangle above the diagonal for modeling and prediction.

We defined SS-FCC as the difference between the mean coherence of an electrode pair calculated during a stimulation block and the preceding resting-state block. We defined RS-FCC as the difference between the mean coherence of an electrode pair calculated during a resting-state block and the preceding resting-state block. We use FCC as a general term to refer to either SS-FCC or RS-FCC. Let $C_{ij}^{(w)}(\lambda_k)$ be the coherence between electrodes i and j at frequency λ_k in the 20-second window w, as estimated by the procedure described above. The band-limited coherence was obtained by averaging over the indices $k \in K_b$ corresponding to frequency band b: $C_{bij}^{(w)} = |K_b|^{-1} \sum_{k\in K_b} C_{bij}^{(w)}(\lambda_k)$. Within an experimental session,

we denote by W_{R_l}, $l \in \{1,2,3,4,5,6\}$ and W_{S_l}, $l \in \{1,2,3,4,5\}$, the collections of 20-second windows belonging to resting-state block l or stimulation block l, respectively. Note that the additional resting-state block corresponds to a final recording block after the 5th stimulation block. The electrode coherences were summarized for each of these blocks by their mean values:

$$\tilde{C}_{bij}^{(R_l)} = \frac{1}{|W_{B_l}|} \sum_{w \in W_{B_l}} C_{bij}^{(w)}$$

$$\tilde{C}_{bij}^{(S_l)} = \frac{1}{|W_{S_l}|} \sum_{w \in W_{S_l}} C_{bij}^{(w)}$$

The SS-FCC was then computed as

$$y_{bijl} = \tilde{C}_{bij}^{(S_l)} - \tilde{C}_{bij}^{(R_l)}$$

and the RS-FCC was computed as

$$y'_{bijl} = \tilde{C}_{bij}^{(R_{l+1})} - \tilde{C}_{bij}^{(R_l)}.$$

Nonlinear modeling of stimulation-induced coherence change

We used a nonlinear additive model to explain the observed FCC in terms of features derived from both the experimental protocol and the information on functional connectivity available in the resting-state block prior to stimulation. In this framework, a single observation was denoted by the pair(y_n, x_n), $n = 1, \ldots, N$, with $y_n \in \mathbb{R}$ an FCC measurement and $x_n \in \mathbb{R}^p$ the corresponding features. In switching to the single index n we indicate that the data was aggregated over all unique electrode pairs and all experimental blocks. The dependence on the band b was subsequently omitted for notational simplicity; the same analysis was repeated for the data in each frequency band. The FCC was modeled as

$$y_n = \beta_0 + \sum_{j=1}^{p} f_j(x_{nj}) + \varepsilon_n. \qquad (1)$$

The model consists of an intercept β_0, nonlinear functions f_j controlling the impact of each feature $j = 1, \ldots, p$ on the response, and an error term ε_n.

The utility of the model derives from the nonlinearity of the feature mappings and the additive procedure through which they are combined to yield a prediction for the FCC. The additive structure of the model allows for the visualization and interpretation of individual feature-response relationships. For each feature $j = 1, \ldots, p$, this is represented by the feature mapping f_j. Moreover, the predictive impact of feature j can be measured by the decrease in prediction error on a test set relative to a model in which the feature is omitted (or equivalently, a full model is estimated while enforcing $f_j = 0$). The same analysis can be applied to groups of features. While additivity is essential for model investigation, the nonlinearity of the feature mappings enables the identification of more complex relationships than can be expressed in a linear model. Linearity is a strong modeling assumption that is both difficult to justify scientifically, and as demonstrated in the results harmful in terms of predictive accuracy.

Feature representation of processed data

A single observation of the FC corresponds to a unique pair of electrodes and a specific recording block. The corresponding features were constructed to satisfy two objectives: first, that the subsequent analysis separates the influence of the protocol parameters from that of the network structure of the functional connectivity; second, that all information in the features was available prior to stimulation.

To satisfy the first objective, we partitioned the features into two groups: protocol features, which summarize aspects of the experimental setting and protocol; and network features, which summarize information in the electrode coherences during baseline recording. The protocol features describe the key parameters of the experimental framework (Yazdan-Shahmorad, Silversmith and Kharazia). The designation "network features" derives from consideration of the estimated band-limited coherence as the adjacency matrix of an undirected, edge-weighted graph. The network features were intended to serve as simple but informative summaries of the connectivity information available in recordings of baseline activity prior to stimulation. They correspond to summary statistics computed over spatial or temporal ranges, basic quantities pertaining to the estimated spectrum (Brockwell, Davis and Fienberg), or network features that have found previous application in graph analysis (Barrat, Barthélemy and Pastor-Satorras; Salton and McGill). Finally, a subject-level indicator was included to adjust for potential global differences in FCC between the two macaque subjects. The adjustment is included in all models, as the subject is neither an aspect of the stimulation protocol nor a property of the baseline connectivity network.

Protocol features

The protocol features were comprised of the categorical measurements *Anatomical region*, *Delay*, and *Block number*, as well as the real-valued features *Distance*, *Stim1 distance to closer*, *Stim1 distance to further*, *Stim2 distance to closer*, and *Stim2 distance to further*. *Anatomical regions* indicates whether the two electrodes corresponding to a given measurement are both in region M1, both in region S1, or if one is in each region. *Delay* encodes three levels of time-delay in the pulse of the paired laser stimulation: 10, 30, or 100 ms. While the delay parameter in the stimulation protocol is real-valued, in the context of our data it is only measured at three distinct settings. We therefore chose to estimate an effect for each setting individually rather than to estimate a nonlinear function of the delay given observations at only three points. *Block number* indicates the time-order position of the experimental block in which the observation was recorded. In the "block analysis" regression design, we removed this feature and instead allowed all feature mappings to vary with the block number, thus investigating the predictive impact of allowing the entire model to evolve dynamically over the discrete time-stages of an experimental session. *Distance* measures the distance between the two electrodes. *Stim1 distance to closer*, *Stim1 distance to further*, *Stim2 distance to closer*, and *Stim2 distance to further*, encode distances between the stimulation sites and the closer and further electrodes of the electrode pair being predicted.

Network features

The network features were computed from the LFP time series recorded during the resting-state block preceding the stimulation for which FCC are being predicted. Except for the *Phase* feature, all of these features were derived from the tensor obtained by time-concatenation of the coherence matrices across all windows in the resting-state block.

We introduce some tensor notation prior to defining these features. The coherence tensor for a given resting-state period is a three-dimensional array, denoted as $C \in \mathbb{R}^{E \times E \times T}$. Two-dimensional slices through the array at time t or electrode i are indicated as $C_{t\cdot\cdot}$ and $C \cdot i \cdot$, respectively. One-dimensional vectors obtained by fixing an electrode i and a time point t are denoted $\overrightarrow{c_{ti\cdot}}$; vectors obtained by fixing

both electrodes i and j are denoted $\overrightarrow{c_{\cdot ij}}$. The scalar coherence value at time t for electrodes i and j is denoted c_{tij}.

The *Initial coherence* indicates the mean coherence between electrodes i and j at baseline:

$$\frac{1}{T}\sum_{t=0}^{T} c_{tij} = \text{mean}\left(\overrightarrow{c_{\cdot ij}}\right)$$

The *Coherence with network* summarizes the average coherence between electrodes i and j and the remaining electrodes in the array. This corresponds to the normalized sum of vertex strengths of nodes i and j (Barrat, Barthélemy and Pastor-Satorras):

$$\frac{1}{T}\sum_{t=0}^{T}\frac{1}{E}\sum_{k=0}^{E}\left(c_{tik} + c_{tjk}\right) = \text{mean}\left(C_i + C_j\right)$$

The *Coherence difference* measures the mean absolute difference in coherence between electrodes i and j to other electrodes in the network:

$$\frac{1}{T}\sum_{t=0}^{T}\frac{1}{E}\sum_{k=0}^{E}\left|c_{tik} - c_{tjk}\right|$$

The *Length 2 path strength* represents the average strength of length-2 paths connecting electrodes i and j. This feature is similar to standard graph clustering coefficient metrics quantifying the total edge-weight of triangle motifs that include a given node (Barrat, Barthélemy and Pastor-Satorras; Onnela, Saramäki and Kertész). However, it excludes the direct edge between electrodes i and j, which reduces correlation with the initial coherence. Alternatively, the *L=2 path strength* can be considered the unnormalized cosine similarity between the network-coherence vectors of electrodes i and j (Salton and McGill):

$$\frac{1}{T}\sum_{t=0}^{T}\frac{1}{E}\sum_{k=0}^{E} c_{tik}c_{tjk} = \frac{1}{T}\sum_{t=0}^{T}\frac{1}{E}\left(\overrightarrow{c_{ti\cdot}}\right)\cdot\left(\overrightarrow{c_{tj\cdot}}\right)$$

The *Coherence with stim sites* is the average coherence between electrodes i and j to the two optogenetic stimulation sites. Below, the indices a and b refer to the electrode sites corresponding to the laser locations for a given experiment:

$$\frac{1}{4T}\sum_{t=0}^{T}\left(c_{tia} + c_{tib} + c_{tja} + c_{tjb}\right)$$

The *Electrode covariance* and *Time covariance* capture the variability of the coherence measurements over the electrode array and over the total number of 20-second time windows in a resting-state block, respectively.

The *Electrode covariance* represents a time-average of the covariance between the vectors representing the connectivity of electrodes i and j to the rest of the network:

$$\frac{1}{T}\sum_{t=0}^{T} Cov\left(\overrightarrow{c_{ti\cdot}}, \overrightarrow{c_{tj\cdot}}\right).$$

The *Time covariance* is the average across all other electrodes k of the covariance between the time series of coherence values to electrodes i and j:

$$\frac{1}{E}\sum_{k=0}^{E}\mathrm{Cov}\left(\overrightarrow{c_{\cdot ik}},\overrightarrow{c_{\cdot jk}}\right).$$

Finally, the *Phase* was the only network feature not derived from the coherence tensor, since the coherence only contains magnitude information from the estimated spectral density. Writing the $(i,j)^{th}$ cross-spectral component computed within window w as

$$S_{ij}^{(w)}(\lambda_k) = \left|S_{ij}^{(w)}(\lambda_k)\right| e^{i\theta_{ijk}^{(w)}},$$

and denoting by K_b the set of indices k such that λ_k is in frequency band b, the phase feature θ_{ij} in block l is given by

$$\tilde{\theta}_{ij}^{(B_l)} = \frac{1}{|W_{B_1}|}\sum_{w\in W_{B_1}}\frac{1}{|K_b|}\sum_{k\in K_b}\theta_{ijk}^{(w)}$$

Outlier removal

Nonlinear regression methods are sensitive to extreme outliers, which can exert highly disproportionate influence on the model estimate. We wished to achieve robustness against such outliers while minimizing perturbation of the data and subsequent analysis. Therefore, we adopted a highly permissive definition, by which an observation was considered an outlier if its absolute deviation from the mean exceeded 20 times the interquartile range along any dimension (i.e. feature). Across all frequency bands and models, the maximum amount of data excluded by this procedure was less than 0.1% of the total number of observations.

Regression designs

We investigated two different configurations of the regression design matrix within our nonlinear modeling framework. For each design, the data were aggregated across all the experimental sessions for each subject. While significant session heterogeneity suggested that even stronger prediction results were possible for models fit to individual sessions, our objective was to estimate a model that generalized well in the sense of accurate prediction on data from many sessions.

Full data

First, we investigated the performance of the nonlinear model on the full data, which included all experimental blocks from all sessions. This yielded $N = 481{,}505$ observations of $p = 16$ features, which expanded to dimension 21 after dummy-coding of the categorical variables *Delay*, *Region*, and *Block number*, and to dimension 22 after inclusion of the categorical animal subject indicator.

Block interactions

Under the full data design, the time evolution of the response was limited to global shifts corresponding to the categorical feature *Block number*. We subsequently investigated the impact of allowing the shape of the feature mappings to vary with the block number, thus expanding the model's capacity to capture time variation in feature-response relationships over the course of repeated stimulation. This constituted an interaction design, whereby the categorical feature *Block number* was removed, and the remaining features were each augmented with interaction terms for binary indicators that denoted whether the measurement was made in each of blocks 2, 3, 4, and 5. The number of observations remained N = (481, 505), while the number of features expanded to p = 94.

Estimation of the nonlinear additive model

We used a train-test paradigm to assess the predictive performance of the nonlinear model on unseen data. For each frequency band and regression design, the data were split uniformly at random such that

70% of the observations were assigned to the training set and the remaining 30% were assigned to the test set. All model parameters and hyperparameters were selected using data from the training set only. Prior to estimation of the model parameters, the real-valued features were standardized such that they have mean zero and unit standard deviation. Whenever the data were split between train and test sets, the parameters of the linear transformation corresponding to standardization were computed from the training data and applied to the test data.

The nonlinear feature mappings f_j, $j = 1, \ldots, p$, in Eq. 1 were represented in the basis of polynomial functions

$$f_j(x) = \beta_{j1}x + \beta_{j2}x^2 + \cdots + \beta_{jK}x^K.$$

The complete set of parameters to be estimated was thus$\beta = (\beta_0, \beta_1, \ldots, \beta_p)$, where $\beta_j = (\beta_{j1}, \ldots, \beta_{jK})^T \in R^K$. The order of the polynomial representation for f_j is defined as the maximum integer value $k \in 1, \ldots, K$ for which β_{jk} is nonzero. The smoothness of f_j is determined by the order k of its polynomial representation and the magnitude of the coefficients $\beta_{j1}, \ldots, \beta_{jk}$. Numerical implementation requires the specification of an upper bound K, corresponding to truncation of the infinite basis at some maximum order. We used $K = 10$ in all experiments; this decision was justified by the fact that the parameter estimates decay either to negligible values or to exactly zero before reaching order $K \leq 10$, as a result of the hierarchical penalization of the model coefficients.

All categorical features were dummy-coded such that the C categories for each feature are represented by $C - 1$ indicator variables. Polynomial expansion of these variables generates identical columns and thus rank-deficiency in the design matrix, which leads to instability in the optimization routine. This problem was avoided by truncating the polynomial expansion of all categorical variables to order 1 rather than order K. There is no resulting loss of generality from this truncation. Rather, this expanded the framework to allow for simultaneous estimation of smooth nonlinear functions of continuous features and discrete functions (that can be visualized for example, as a step function) of categorical features.

Automatic order selection via hierarchically penalized estimation

Given an observation of the features x, the model prediction of the SS-FCC or RS-FCC is given by

$$\hat{f}(x) = \beta_0 + \sum_{j=1}^{p} f_j(x_j).$$

We used the standard square loss on the training set D^{train}

$$\mathcal{L}(\beta) = \frac{1}{|D^{\text{train}}|} \sum_{(x,y) \in D^{\text{train}}} \left(y - \hat{f}(x)\right)^2$$

as the model goodness-of-fit criterion. Direct optimization of model fit on a training set over a sufficiently complex model class is well-known to yield estimates that perform poorly on unseen data, an issue commonly known as overfitting. We avoided this problem by adoption of a penalized estimation framework specifically designed to control the magnitude of the parameter estimates and automatically select the order of each feature mapping in the model. This approach augments the standard square loss with a penalty term $\Omega_j(\beta_j; \alpha, \lambda)$ for each feature j, defined as

$$\Omega_j(\beta_j;\alpha,\lambda) = \alpha\lambda\sum_{k=1}^{K} w_k\left(\sum_{i=1}^{N}\left(\beta_{jk}x_{ij}^{k}\right)^2 + \cdots + \left(\beta_{jK}x_{ij}^{K}\right)^2\right)^{\frac{1}{2}} + (1-\alpha)\lambda\left(\sum_{i=1}^{N}\left(\beta_{j1}x_{ij}\right)^2 + \cdots + \left(\beta_{jK}x_{ij}^{K}\right)^2\right)^{\frac{1}{2}}, \quad (2)$$

with hyperparameters $\lambda \in R_+$ and $\alpha \in [0,1]$. The weights $w_k = k^3 - (k-1)^3$ were chosen to satisfy theoretical criteria that guarantee statistical convergence of the estimator over a class of smooth functions (Haris, Shojaie and Simon).

The model parameters were estimated by solving

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}}\, \mathcal{L}(\beta) + \sum_{j=1}^{p}\Omega_j(\beta_j;\alpha,\lambda).$$

We used a block coordinate descent algorithm as implemented in the R package *HierBasis* (Haris, Shojaie and Simon). The first term in the penalty $\Omega_j(\beta_j;\lambda,\alpha)$ induces hierarchical sparsity in the estimate of $\beta_j = \left(\beta_{j1},\ldots,\beta_{jK}\right)^T$, while the second term in Eq. 2 shrinks the magnitudes of the estimated coefficients towards zero. Hierarchical sparsity guarantees that if an estimated coefficient $\widehat{\beta_{jk}} = 0$, then all higher-order coefficients for that feature $\widehat{\beta_{jk'}} = 0$, with $k < k' \leq K$; this is equivalent to selecting an order $k-1$ representation for the feature mapping f_j. The selection procedure is thus automatic and data-driven in that it emerges as a direct mathematical consequence of the penalized estimation framework, which seeks the best model fit to the data under the structural constraints imposed by $\Omega_j(\beta)$.

Cross-validation for model hyperparameters

The penalized loss function requires specification of two regularization hyperparameters $\lambda \in R_+$ and $\alpha \in [0,1]$, which control the overall regularization strength and tradeoff between terms inducing hierarchical and feature-wise sparsity, respectively. The parameter α was selected over the range from 0 to 1, inclusive, in increments of 0.1. The parameter λ was selected over 100 log-linearly-spaced values between λ_{max} and $\lambda_{max} \times 10^{-5}$.

The value λ_{max} was selected such that all estimated feature mappings are equal to zero for every value of investigated α; it was obtained by a backtracking line search algorithm (Boyd and Vandenberghe). This algorithm requires a gross upper bound λ_0, defined as any value of $\lambda > 0$ such that $\hat{\beta} = 0$ for all grid values of α. It was found manually; we found that $\lambda_0 = 0.1$ sufficed for all designs and frequency bands.

We selected the pair (α^*,λ^*) by first finding the (α,λ) pair minimizing the average R^2 on the validation set over 5-fold cross-validation on the training data. We selected α^* as the α-coordinate of this pair, and selected λ^* as the largest value of λ such that the mean validation R^2 of (α^*,λ) was within one standard error of the mean validation R^2 at (α^*,λ^*). This follows the "one standard error" strategy for one-dimensional cross-validation (Friedman, Hastie and Tibshirani) and represents a conservative approach to regularization corresponding to our preference for smoother feature mapping estimates.

Performance measure

We quantified model performance by predictive accuracy on the test set, as measured by the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{i=1}^{n_{\text{test}}}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n_{\text{test}}}(y_i - \bar{y}_i)^2},$$

where $\hat{y}_i$ was the model prediction after cross-validation and estimation on the training set and $\bar{y}_i$ was the mean of the regression target on the test set.

Resampling for feature stability

We took a resampling approach to assess the stability of the estimated feature mappings of the nonlinear model. Following the subsampling heuristic of (Meinshausen and Bühlmann), we drew a subsample of size $\lfloor N/2 \rfloor$ from the data without replacement and ran the full estimation procedure, including hyperparameter selection, on the subsampled data. The procedure was repeated for 100 independent trials, resulting in 100 estimated functions for each feature. The feature-wise variability in the nonlinear additive feature mappings was then assessed pointwise, by computing upper and lower quantiles for the range of function values at each point in a fine grid.

A predictive measure of feature importance

Traditional tests of statistical significance are not available for the coefficients of the nonlinear additive model as the asymptotic distribution of the penalized estimator is not known. Instead, we assessed importance of a feature or group of features by evaluating the difference in predictive performance on the test set between the full model and a model fit without the parameter or parameter group of interest. Removing the feature or feature group is equivalent to estimating the full model under the constraint that their corresponding parameters are all equal to zero. The comparison can thus be viewed as between the full model and one in which a hypothesis of null response is enforced for the feature or group of interest. When estimating the reduced model, we used the regularization parameters (α^*, λ^*) obtained by cross-validation on the full data.

Quantifying feature similarity across frequency bands

Let $\widehat{f_j^a}$ and $\widehat{f_j^b}$ be the estimated feature mappings for feature j on the data corresponding to frequency bands a and b, respectively. We computed a quantitative measure of their similarity,

$$s_j^{ab} = \frac{\langle f_j^a, f_j^b \rangle}{\|f_j^a\| \|f_j^b\|}.$$

The quantity s_j^{ab} is the cosine similarity of the estimated feature mappings, considered as elements of a common Hilbert space of functions. By definition, $s_j^{ab} \in [-1,1]$. Similarity increases as s_j^{ab} approaches 1 or -1 (which indicates f_j^a and $-f_j^b$ are perfectly similar) and is minimized at 0.

The inner product $\langle f_j^a, f_j^b \rangle = \int_{\mathcal{X}_j} f_j^a(x) f_j^b(x) dx$ and norm $\|f_j\| = \langle f_j, f_j \rangle$ are given by integrals whose domain $\mathcal{X}_j$ depends on the feature j. For real-valued features, we took $\mathcal{X}_j$ to be the interval $[-5,5]$, which after standardization corresponds to the range of all observed measurements within 5 standard deviations of the mean. Due to the polynomial representation of the nonlinear feature mappings, these could be computed exactly.

Chapter 4: Behavioral and neural disruption in non-human primates via optogenetic inhibition through chronic opto-electric interface

Julien Bloch* and Devon Griggs*, Azadeh Yazdan-Shahmorad

Abstract

We introduce a novel optogenetic-electric neural interface designed for advanced behavioral and neural disruption studies in non-human primates (NHPs). The interface combines the spatiotemporal control of optogenetics with the functionality of electrical recording, allowing simultaneous stimulation and readout of neural activity at a fine scale. We employed this interface to disrupt instructed delay center-out reach tasks, effectively showcasing its potential to modulate voluntary actions in NHPs. Simultaneously, we monitored ongoing electrocorticographic activity and documented the disrupted patterns underlying the behavioral changes. The system exhibits long-term stability and functionality, with implant stability and epifluorescence lasting at least two years, and clear optical access lasting three months between window tissue growth resections, thereby facilitating large-scale electrocorticography, optical imaging, and optogenetics over extended periods. This characteristic is crucial for studying neurological diseases and disorders with translational relevance. The interface is designed for maximum ease of implementation across various research specialties, supporting large-scale, long-term multi-modal experiments in this highly translational animal model.

Introduction

Optogenetics is unique among neuromodulation technologies because it allows for cell-type specific manipulation of neural networks with high spatial and temporal precision. This ability enables the precise manipulations needed for research in fundamental neuroscience (Carter and de Lecea, 2011; Chen *et al.*, 2013; Adhikari *et al.*, 2015; Barnett *et al.*, 2018; Camporeze *et al.*, 2018), including learning and memory (Spellman *et al.*, 2015; Barnett *et al.*, 2018; Yang *et al.*, 2018), the relationships between behavior and neural circuits in the brain (Gradinaru *et al.*, 2009; Chen *et al.*, 2013; Adhikari *et al.*, 2015; Jiang *et al.*, 2018; Yang *et al.*, 2018), the changes in neural circuitry over time (Jiang *et al.*, 2018; Yazdan-Shahmorad *et al.*, 2018), and therapies for neurological disorders (Chen *et al.*, 2013; Tønnesen, 2013; Ruggiero *et al.*, 2017; Barnett *et al.*, 2018; Camporeze *et al.*, 2018).

The study of optogenetics in non-human primates (NHPs) is especially vital for translational medical research, especially in the context of neurological diseases and disorders. This is due to the evolutionary proximity between the neurophysiology of NHPs and humans (Roelfsema and Treue, 2014; Lear et al., 2022). Studying neurological disease models in NHPs is challenging because human-like pathology and anomalous neural activity are often multi-regional chronic phenomena. These spatiotemporal characteristics necessitate the development of large-scale, chronically-stable optogenetic interfaces for NHPs. In response to these needs, we introduce optical cranial windows designed to facilitate multi-modal experimentation for researchers working with NHPs, and bridge the transition of multi-modal studies from rodents to monkeys.

Optogenetics alone provides a potent interface for neuromodulation, but the integration of optogenetics with an electrical interface can greatly enhance capabilities. A combined interface allows for simultaneous electrical stimulation and recording with optogenetic manipulation, enabling precise control over specific cell types and the concurrent measurement of neural activity and functional connectivity with high spatial and temporal resolution. Furthermore, optogenetic and electrical stimulation have been shown to work synergistically, for example by enhancing the efficacy of evoking saccadic eye movements in macaques compared to using electrical stimulation alone (Ohayon et al., 2013). The development of multimodal

neural probes and closed-loop brain-computer interfaces (BCIs) that integrate optogenetics and electrophysiology is opening new avenues for neuroscience, furthering the transition of multi-modal studies from rodents to primates (Tian et al., 2022).

While strides are being made in optogenetic and electrical interfaces, of similar importance is the advancement of *in vivo* optical imaging capabilities. *In vivo* optical imaging of NHP brains through cranial windows has been a small but vibrant field with decades of development (Arieli *et al.*, 2002; Slovin *et al.*, 2002; Ruiz *et al.*, 2013; A Yazdan-Shahmorad *et al.*, 2016; Yazdan-Shahmorad *et al.*, 2018; Macknik *et al.*, 2019). More recently, the technology has been combined with optical stimulation methods such as infrared neural stimulation (Chernov and Roe, 2014) and optogenetics (Ruiz *et al.*, 2013; Zaraza *et al.*, 2020), and electrical recording methods such as electrocorticography (A Yazdan-Shahmorad *et al.*, 2016; Yazdan-Shahmorad *et al.*, 2018). However, achieving long-term stability in these multi-modal approaches remains a challenge (Belloir *et al.*, 2022). In this work, we demonstrate large-scale and chronically-stable integration of optogenetics and electrocorticography in an imaging window on NHP cortex. To our knowledge, this is the first demonstration of such a synergistic approach in any animal model.

We demonstrated the versatility of our interface by impacting behavior in a center-out reach task through the optogenetic deactivation of the posterior parietal cortex (PPC). This showcases the applicability of our interface to higher order sensory and motor behavioral experiments. We demonstrate the stability of our system for over 3 months, which is critical for studying models of neurological diseases and disorders at timescales with translational relevance. The ability to link millisecond-resolution neural activity with behavioral effects and cell-type specific neuronal manipulation, as provided by optogenetics, is a valuable addition to the NHP researcher's toolbox. We took measures to ensure accessibility to our methods for a diverse range of specialties, enabling large-scale and long-term multi-modal experiments in a highly translational animal model.

Results

Hardware and surgical techniques

To plan a 25 mm diameter craniotomy over the left PPC in two healthy adult male rhesus macaques (*Macaca mulatta*, monkeys H and L), we used our previously described surgical planning method (Ojemann *et al.*, 2020) (Fig. 1A), making use of magnetic resonance (MR) data of the skulls of the NHPs.

This method allowed us to design a chamber with a curvature precisely fitting the surface of the skull] (Fig. 1C). We performed the craniotomy and implanted the chamber along with our multi-modal artificial dura (MMAD, (Griggs *et al.*, 2021)), which provided electrophysiological and optical access to the PPC (Fig. 1B, D).

During the surgery, we used convection enhanced delivery (CED) (Yazdan-Shahmorad et al., 2016; Yazdan-Shahmorad et al., 2018; Karam Khateeb et al., 2019; Griggs et al., 2022b) to infuse an inhibitory optogenetic viral vector (AAV8-hSyn-Jaws-GPF) at several locations across the craniotomy to maximize the area of viral coverage and minimize the time required for infusion.

Two years after the surgery, the implant remains implanted and useable in both macaques.

We began experiments after recovery from surgery. At the start of an experiment, we attached custom electrophysiological recording hardware (Fig. 1E-F) and optogenetic stimulation hardware (Fig. 1G) to the chamber. The lasers and fibers were commercially sourced, while our LED arrays and driving circuitry were custom made (Fig. 2). We designed our driving circuitry to be controlled by analog signals to avoid noise

complications in our electrophysiology, and we powered our circuitry with two rechargeable 9V batteries to mitigate line noise. We designed and fabricated two LED array versions for patterned stimulation, a 4 × 4 array with individually drivable columns and a 3 × 5 array with each LED independently drivable with identical driving circuitry (Fig. 2A). The LED array was protected from moisture by a cover glass between the LED array and the MMAD (Fig. 1E-G, 2B). This arrangement created an air gap between the LEDs and the cover glass, which mitigated conductive heating of the brain (Fig. 2B).

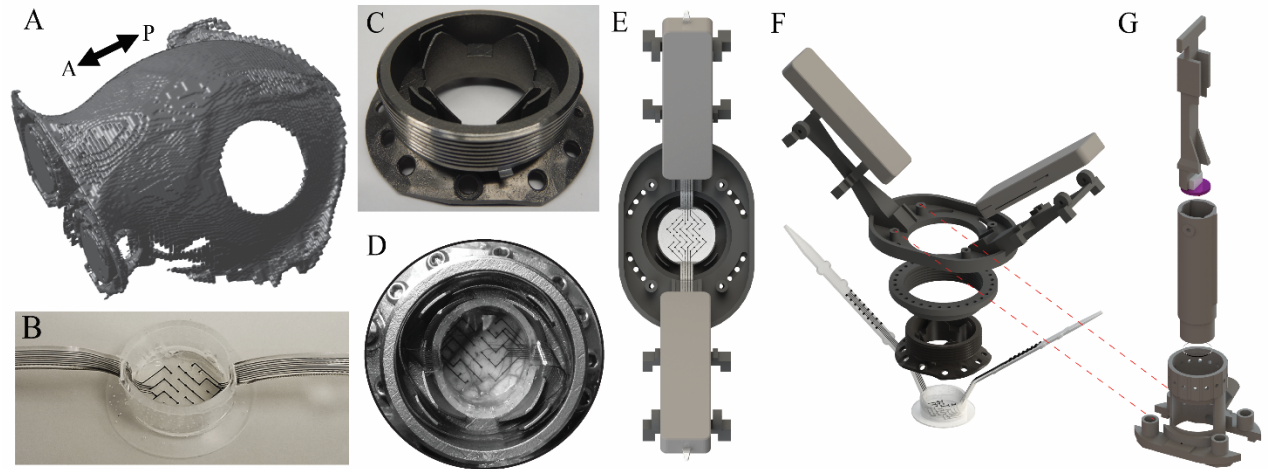

Figure 1: Hardware. (A) Digital representation of the skull with craniotomy. (B) Molded MMAD. (C) Chamber implanted on skull. (D) Implanted chamber with MMAD. (E) Model of electrophysiological recording hardware, top view and (F) trimetric view. The chamber contains the MMAD, and custom made mounting equipment is secured above, which allows the electrode arms of the MMAD to be securely connected to clamp connectors, and allows the stimulation hardware to be connected to the mounting equipment. From bottom up: MMAD, chamber, ring, clamp connector tray, clamp connectors. (G) LED-based optical stimulation hardware. From bottom up: Stimulation holder, cover glass, outer tube, LED array, inner LED holding rod.

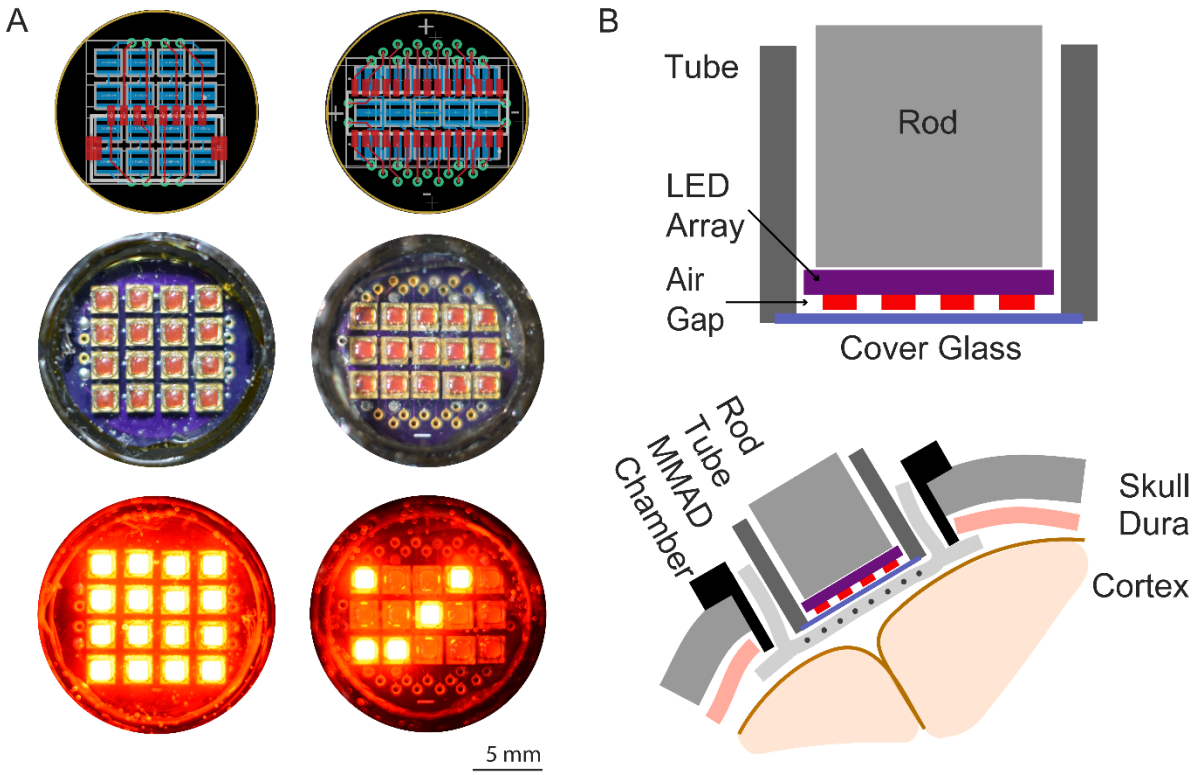

Figure 2: LED array. (A) Row-drivable 4 × 4 LED array (left column) and individually drivable 3 × 5 LED array (right column). PCB schematic (top), photograph of LED array in tube and behind cover glass with LEDs off (middle) and on (bottom) in example configurations. (B) Schematic of LED array and related hardware for optical stimulation through MMAD, including an air gap to mitigate tissue heating.

Optical window stability

For Monkey H, we imaged the brain weekly through the MMAD and observed only modest tissue growth over the first 14 weeks after implantation (Fig. 3). After 14 weeks, an acceleration in tissue growth was observed, and after it matured, we resected the tissue to perform additional experiments. For Monkey L, there were surgical complications with the initial surgery that hindered optical access, but we reimplanted an MMAD 20 weeks after the initial surgery with clear optical access that lasted for 10 weeks; at this point the MMAD was explanted and the tissue began to grow over the cortex. As mentioned before, our implants have been stable for multiple years, proving their applicability for long-term NHP experiments. For example, for monkey H, following a tissue resection approximately two years after first implantation, the MMAD was replaced and clear optical access was maintained for 7 weeks. At week 7 while optical access was still clear, the MMAD was removed, then approximately one week later tissue began to grow over the cortex.

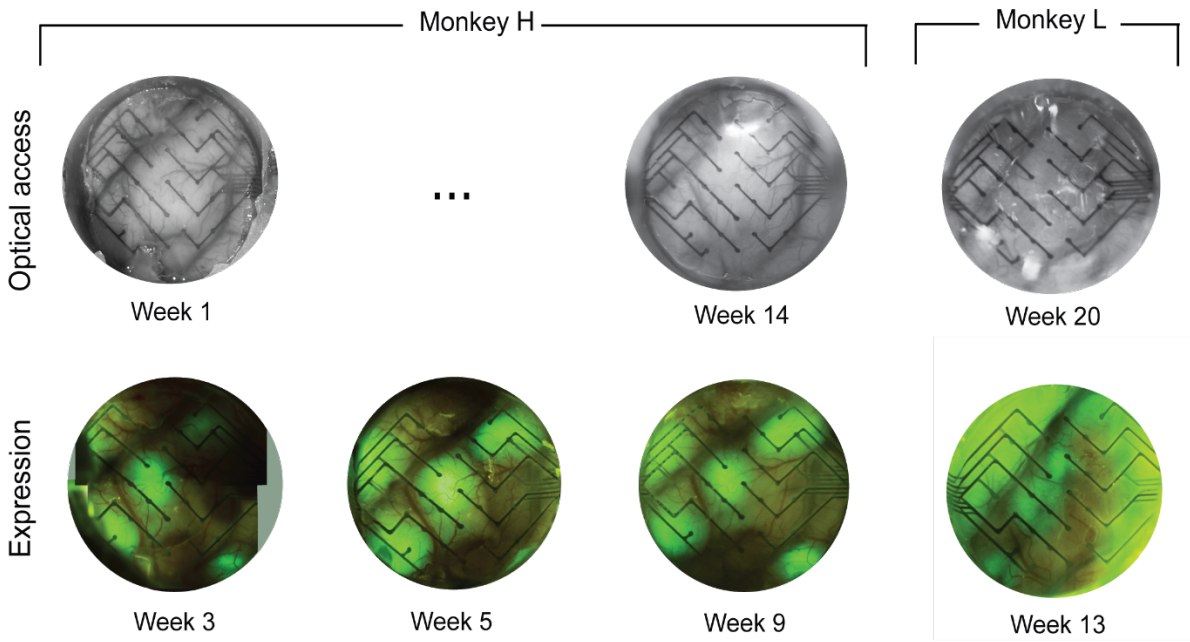

Figure 3: Optical access (top) and epifluorescence (bottom) over time, for both NHPs. Expression patterns differ between the two subjects, with fluorescent areas corresponding to the CED viral injection sites. Optical access and epifluorescence are maintained over time.

Epifluorescence

Optogenetic expression was monitored over time for both subjects. Subject H underwent initial testing for epifluorescence two weeks after viral infusion. Expression was evident during the first imaging session and increased in coverage and intensity until the fifth week after infusion, when expression plateaued at tens of mm^2 coverage (Fig. 3). We performed epifluorescent imaging four and five weeks post-infusion in subject L and also identified large-scale expression, despite aforementioned surgical complications. Thirteen weeks post-infusion, we performed epifluorescent imaging again, this time with little to no tissue growth impeding optical access, and we confirmed tens of mm^2 expression coverage (Fig. 3). We repeated epifluorescent imaging for subject H approximately two years after viral infusion and observed continued expression (data not shown). These large areas of coverage and long time scales of expression are comparable with previous work (Yazdan-Shahmorad *et al.*, 2016) and are important proofs of concept for our method of optogenetic CED in NHP cortex without live MRI guidance (A Yazdan-Shahmorad *et al.*, 2016).

Behavioral task

Subjects were trained to perform an instructed delay center-out reach task to targets on a video monitor in front of the animal using the hand contralateral to the implant (Fig. 4A). The finger was held in the center at a start target and then the end target was shown randomly in four locations: left, up, right, down.

We applied optical stimulation for 900 ms with all LEDs of the full LED array (634 nm, ~12 mW/mm2) for a randomly sampled 50% of the trials, with the onset of the stimulation aligned with the go-tone. Subjects performed several hundred trials each day for 15 total days.

Each day began with a few trials with easier reach settings such as a shorter interval of the wait time before the go tone, and/or larger circle sizes. For analysis, these early warmup trials were excluded.

Widespread optogenetic stimulation disrupts reaching behavior

We found that optogenetic inhibition aligned to the onset of the go tone resulted in significantly longer reach times for both subjects. This was true if we grouped the data from the subjects together (Fig. 4B) and if we analyzed each subject individually (Fig. 4C-D). For subject H, there were 3904 analyzed reach trials split over 7 sessions, while for subject L there were 8588 trials split over 12 sessions.

For the combined data, the reach times for trials with stimulation (stim condition) were significantly greater than the reach times for the control trials in which stimulation was not applied (nostim condition) when reaching to the left (median difference = 24.0, $p = 1.54 \times 10^{-6}$) and down (median difference = 28.0, $p = 1.35 \times 10^{-10}$). When all directions were combined, the reach times in the stim condition were significantly greater than in the nostim condition (median difference = 19.0, $p = 4.59 \times 10^{-8}$).

For subject H, the stim condition reach was significantly longer when reaching to the left (median difference = 22.0, p = 0.0015) and to the right (median difference = 10.5, p = 0.0316). When all directions were combined, the stim condition reaches were significantly larger (median difference = 8.0, p = 0.024).

For subject L, the stim condition reach was significantly longer when reaching to the left (median difference = 23.5, p = 7.02e-06), down (median difference = 32.0, p = 9.96e-10), and combined directions (median difference = 13.0, p = 1.04e-06).

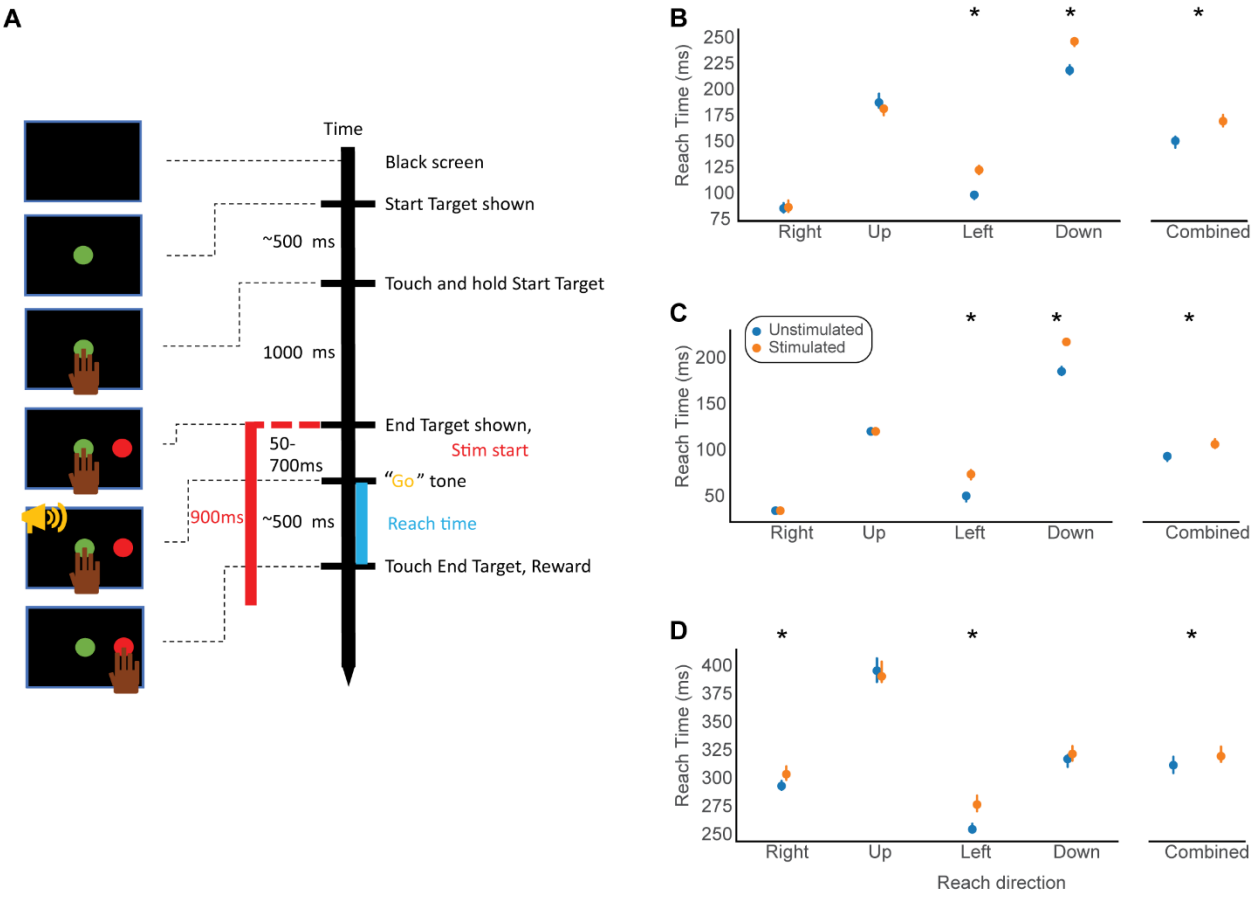

Figure 4: A) Schematic of behavioral task and stimulation intervention timeline. B-C) Behavioral disruption results for B) combined data, C) subject L, and D) subject H. Blue circles indicate unstimulated reach times and orange circles indicate stimulated reach times. Error bars show 95% confidence interval of the median as determined by a bootstrapping procedure. Asterisks indicate reach times significantly slowed by stimulation.

Widespread optogenetic stimulation disrupts neural activity during reaching

We analyzed the neural activity via electrocorticography recorded during reaches with and without optogenetic inhibition to quantify the change in neural activity underlying the disrupted reach.

First, the recorded data was cleaned by the photoelectric-artifact rejection procedure detailed in methods. Briefly, it involved measuring a signal in saline when light was shined onto the array such that only artifact is present, then fitting an artifact basis to the response using PCA. Then, the projection of recorded neural data onto this basis was subtracted before analysis. We validated the approach by showing that it effectively removes nearly all signal in saline, and that it negligibly and insignificantly alters recorded neural activity in the unstimulated condition while visually removing the non-neuronal artifact from the stimulated condition.

After artifact rejection, we quantified the effect of stimulation on reaching neural activity by analyzing the time-frequency content of the data (Fig. 5). The ECoG data associated with each successful reach were grouped by whether stimulation had been applied, and by subject. Data from each trial were analyzed in the time-frequency domain via multitaper spectrograms. The average spectrogram for a reach in the presence of stimulation is shown in Figs. 5A and 5D, and the average activity for a reach in the absence of a stimulation is shown in Figs. 5B and 5E. In these figures, the values are normalized such that the average pre-stimulation activity is unity for each individual frequency. Grey bars indicate the start and termination of stimulation, as the 50ms immediately following the start and termination are excluded from analysis due to difficulty of removing the artifact in comparison to the rest of the signal recorded during the 900ms stimulation. From these figures it is apparent that during stimulation, there is generally increased power throughout the time course and over most frequencies. Patterns of neural activation during reaching, as in Fig. 5E, are present in the stimulation condition as well but seem to be modulated.

To quantify this modulation, we calculated the power ratios of the spectrograms from the stimulated to the unstimulated conditions, individually for both subjects (Fig. 5C and 5F). Each time-frequency bin of the stimulated spectrogram is divided by the bin value of the spectrogram corresponding to the unstimulated condition to obtain a power ratio. In addition, the distribution of powers of that time frequency bin over all trials of all channels of all sessions are compared between the two conditions to generate statistical significance values, which are then Bonferroni transformed for multiple comparison correction. This approach clearly shows that stimulation during reaching strongly increases the brain activity during reaching over the entire course of stimulation, with nearly all time and frequency bins exhibiting significant ($p < 0.05$) modulation from stimulation.

We repeat this analysis in the frequency domain, allowing the use of longer time windows to improve the frequency resolution (Fig 5G-J). This is done for the neural activity recorded during stimulation (Fig. 5G and I) and immediately after stimulation (Fig. 5H and J). Significances are similarly calculated here and indicated by black lines and dots above the plots of the power ratios between the stimulated and unstimulated conditions. Nearly all frequencies are modulated during stimulation for both monkeys, whereas only a few frequencies are modulated after stimulation.

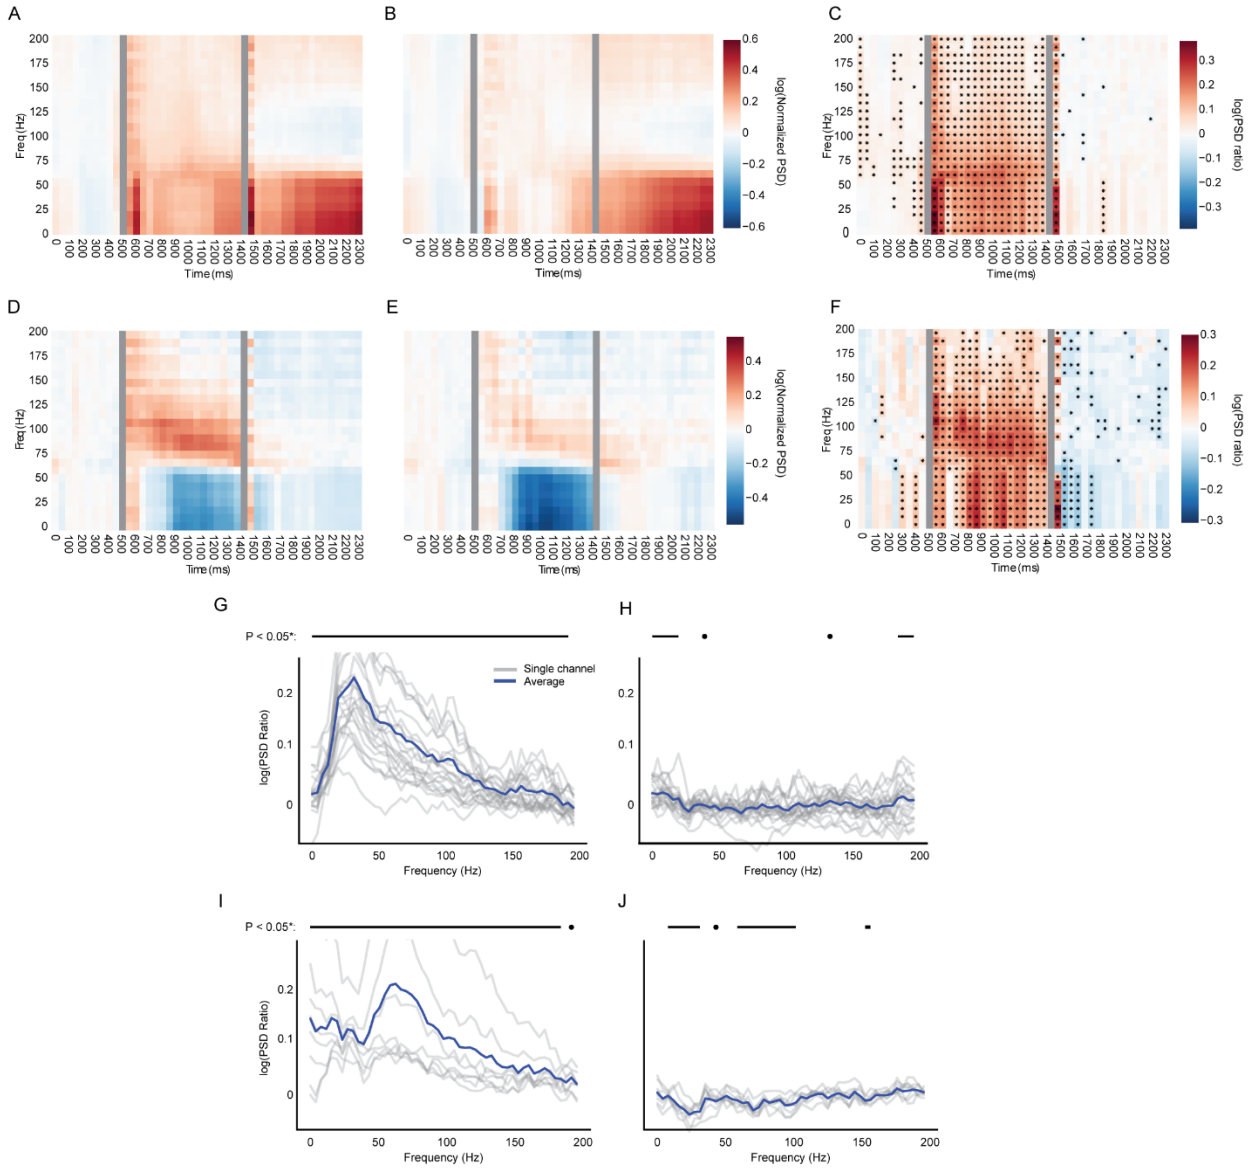

Figure 5: The effect of stimulation on neural activity in the reach condition. A-B) Spectrograms corresponding to the neural data of the reaching monkey L in the A) stimulated condition and in the B) unstimulated condition. Data values are averaged over all trials, sessions, and channels. The area between the two grey vertical lines indicates the analyzed data from the go tone until 900ms after, corresponding to the entire time course of optical stimulation. Grey bars indicate bins during the first 50ms of stim onset to stim offset and vice versa; data during this time iare greyed out due to difficulty rejecting the artifact during this period. For ease of visualization, values are normalized to the average pre-stimulation power, and the log base 10 of this number is shown. C) Ratio of average stimulated condition data to unstimulated condition data, per time-frequency bin. Asterisks indicate significant ($p < 0.05$) differences between the conditions. D-F) same as A-C) but for monkey H. G-H) PSD ratios between the stimulated and unstimulated condition for monkey L G) during stim and H) after stimulation. Power values are log base 10 transformed. Grey lines indicate values averaged over all trials for an individual channel of an individual session; blue line indicated ratio of means of all trials of all channels of all sessions. Black lines and circles above plots indicate frequencies which are significantly ($p < 0.05$) different between the stimulated and unstimulated condition. I-J) same as G-H) but for monkey H.

Effect of Stimulation in the Absence of Reach

We applied optogenetic stimulation in the absence of reach to quantify the result of the optogenetic inhibition on resting-state neural activity. This was completed in subject L, in the absence of any behavioral task or reaching (Fig. 6A). As before, statistical results are generated by comparison to the unstimulated condition. We observe a significant increase in power during stimulation, particularly above 60Hz. These findings differ from those obtained during reaching behavior for this subject, where the increase in power was strongest from 20Hz – 70Hz. This result indicates that the behavioral context and ongoing neural activity affect the frequency characteristics of the stimulation induced modulation.

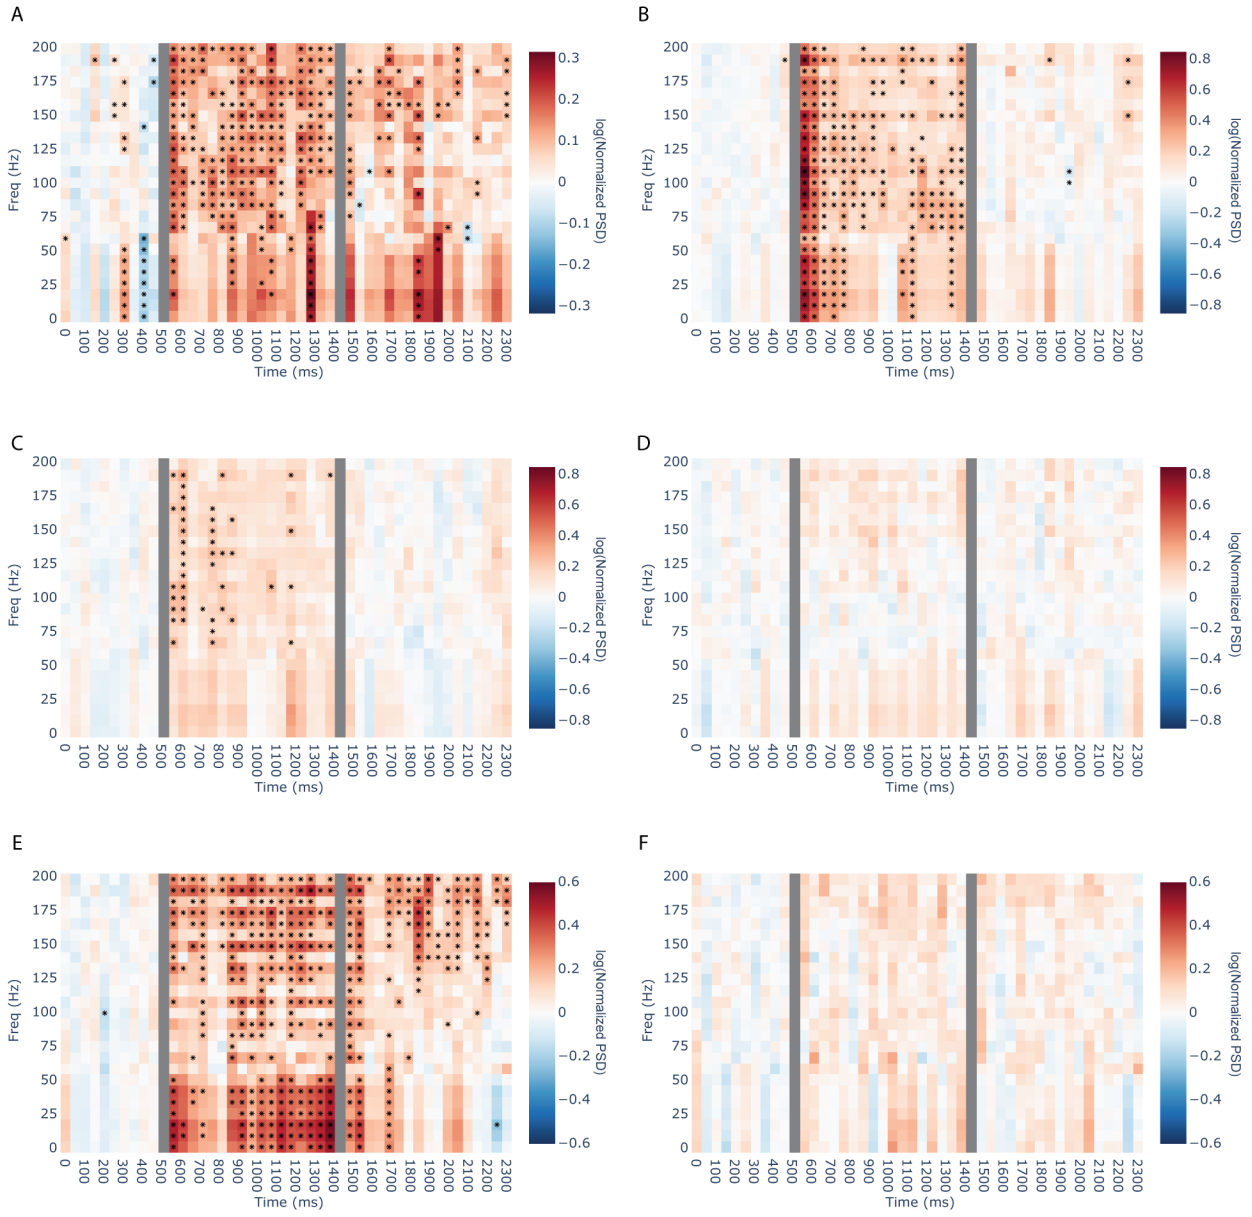

Figure 6: Neural response to various stimulations in the no-reach condition. A) Normalized PSD for stimulation applied when no behavioral task of reach is occurring. B-D) Normalized PSD for channels while a single-LED illumination is taking place, recorded from B) the primary targeted region C) a secondarily electrode region and D) a nontargeted region. E-F Normalized PSD for channels while a single laser guided illumination is taking place, from E)

the targeted region and F) a nontargeted region. For all plots, power values are normalized on a frequency-wise manner by the average power of that frequency measured before stimulation is applied. Asterisks indicate significant ($p < 0.05$) differences between distributions of powers in a specific time-frequency bin and the distribution of powers of pre-stim time-bins for that same frequency.

Single-LED illumination disrupts local activity

We then tested the ability of the modular LED array to shine localized illumination on a small region of the optical window by only illuminating through a single LED. The illumination primarily targeted one region corresponding to a channel of the electrode array (Fig. 6B), and partially targeted another which was in proximity (Fig. 6C). All other regions did not show stimulation-induced modulation (Fig. 6D). With these figures we clearly observe the spatial control possible with a more localized stimulation.

Laser-guided illumination disrupts highly localized activity

Using our laser setup eight weeks after infusion of Monkey H, we illuminated neural tissue with fiberoptic cable near selected electrodes and evoked spatially specific changes in our ECoG signal confirmed with various illumination parameters. Results are shown at the site of stimulation (Fig. 6E) and at an untargeted site (Fig. 6F). Unlike the activation profile from the single LED illumination, the laser delivered activity was only visible on a single channel.

Discussion

We present a chronic interface for large-scale electrocorticography, optical imaging, and optogenetics in NHP cortex by which behavior can be modulated. Our work culminates decades of effort across the field largely hindered by the challenge of maintaining optical access, especially at the large scales necessary to study complex network interactions across multiple brain regions *in vivo* in NHPs. As with past work, we resected both the skull and the opaque native dura and implanted transparent artificial duras that are biocompatible and can provide optical access to the brain for 3-9 months (Arieli *et al.*, 2002; Chen *et al.*, 2002; Azadeh Yazdan-Shahmorad *et al.*, 2016) before tissue growth stifles optical access (Azadeh Yazdan-Shahmorad *et al.*, 2016). This tissue is a neo-membrane and cannot be easily removed during early stages of growth because it is nourished through the brain, but the tissue can be surgically removed once the neo-membrane matures and separates from the brain. This growth process forces the researchers to wait for about one to two months before resuming experimentation, and thus highlights the importance of maintaining optical windows for an order of months, as we have done.

While others have experimented with window-less LED-based NHP optogenetics (Rajalingham *et al.*, 2021), this work facilitates routine epifluorescent monitoring of large-scale optogenetic expression, revealing experimental insights otherwise challenging to ascertain. For example, in our experiment the optical window empowered us to confirm that convection-enhanced delivery performed without the complexity of live-MRI (unlike (A Yazdan-Shahmorad *et al.*, 2016)) resulted in a progression of large-scale expression up to a plateau in expression at 5 weeks, and which was still observed at two years.

Early attempts to integrate electrode arrays into a large-scale optical window did not result in long-term optical access, however. In previous work, a Parylene-C-embedded micro-electrode array was implanted under a transparent silicone artificial dura to achieve simultaneous electrical recording and optical stimulation (Yazdan-Shahmorad *et al.*, 2015; Azadeh Yazdan-Shahmorad *et al.*, 2016). Surprisingly, the brain exhibited such a high biocompatibility with Parylene-C that the tissue quickly grew over the top of the Parylene-C during the experiments and produced only 13-18 days of optical access (Yazdan-Shahmorad *et al.*, 2015). In this work, we achieve over 3 months of optical and electrophysiological access by embedding electrodes into a transparent polymer and then embedding the product into an artificial dura to create a single MMAD. The brain's response is similar to traditional (electrode-free) silicone artificial

duras, which facilitates long-term optogenetic experiments with electrophysiological recordings. Also, the chronic MMAD eliminates an alternative, tedious approach of daily implantations and explanations of micro-electrode arrays that accelerated tissue growth and thus maintained optical clarity for only about a month (A Yazdan-Shahmorad *et al.*, 2016). We are among the first to embed a miniature electrocorticography array into an artificial dura (Orsborn *et al.*, 2015; Kleinbart *et al.*, 2018; Chiang *et al.*, 2020), and we are the first to use this embedded approach chronically in NHPs for large-scale imaging and optogenetics in addition to electrophysiology; this suite of techniques has not been demonstrated in rodents to our knowledge (Chen *et al.*, 2020; Renz *et al.*, 2020; Zátonyi *et al.*, 2020), and yet our work brings this novel integration of techniques for long-term use to the highly translational NHP model.

We designed our interface to be modular and replicable in a number of ways. Our MMAD molding process does not require any special facilities or tools aside from the custom mold pieces, and the array itself is commercially available (Ripple Neuro Inc.). Furthermore, we predict our demonstration will spur the development of MMADs with higher channel counts of smaller and transparent electrodes to increase both electrophysiological resolution and optical access. We also expect LEDs will be incorporated into future MMAD designs, much like past work with electrodes arrays (e.g., (Ji *et al.*, 2018; Reddy *et al.*, 2019)). We have refined surgical planning and chamber design processes to accelerate the field (Ojemann *et al.*, 2020). The ring and clamp connector tray facilitate experimental flexibility for imaging, electrophysiological recording, and both fiberoptic- and LED-based optical stimulation. Most of the optical stimulation components (bases, stimulation ring, rod, and tube) can be 3D printed in-house with common filament and printers, which facilitates quick and inexpensive revisions for unique experimental needs. Parts can be soldered onto the LED-related PCBs in-house with minimal experience, and minimalistic battery-powered circuitry mitigates line-noise. The LEDs are commercially available in many colors and are driven continuously (as opposed to pulse-width modulation), thus mitigating complexity of photo-artifacts in electrophysiology data. The air-gap in our LED setup mitigates tissue heating, which is a significant concern in window-less LED-based NHP optogenetic methods (Rajalingham *et al.*, 2021). The stimulation setup of previous work allowed for the optical stimulation of only two or three point-locations at a time due to the physical interference of fiberoptic wires within the enclosure (Azadeh Yazdan-Shahmorad *et al.*, 2016), and in contrast, both of our LED arrays provide broad, non-invasive patterned optical stimulation and eliminate concerns regarding fiber fragility. We anticipate development of higher-count and multicolor LED arrays to further complement our large-scale optical interface, which would be especially powerful with multi-opsin and step-function opsin experiments (Diester *et al.*, 2011a). The behavioral task was based off previous work (Griggs *et al.*, 2020), and our code was written in MATLAB, which is a common software language among researchers. Our behavioral task code and the design files of our electronic and hardware parts are publicly available.

Past studies have attempted to optogenetically evoke musculoskeletal activity in NHPs at rest. Early trials were unsuccessful (Diester *et al.*, 2011b), but more recently, a proof-of-concept was demonstrated (Watanabe *et al.*, 2020). Building on this work, recent work (a preprint at the time of writing this) has modulated sensory-guided reaching behavior using excitatory optogenetics (O'Shea 2022), but there have been no reports of modulating this behavior using inhibitory optogenetics. Here we present the first report of successfully modulating sensory-guided reaching behavior using inhibitory optogenetics. We used a center-out reach task to test for optogenetic behavioral modulation, which we achieved in both animals. Specifically, we confirmed our hypothesis that inhibition of the PPC, a sensory integration area of the brain, would disrupt multi-faceted sensory processing and thus slow reaches informed by audial and visual cues. The variation of reach time effects in direction between animals is reasonable, given the variance in expression maps between animals. Of note, motor dynamics of reaches were qualitatively similar with and without stimulation (data not shown), demonstrating optogenetics can modulate behavior while preserving natural and voluntary characteristics of the behavior. Our results indicate that our interface is

suitable for studying neural mechanisms of behavior with spatially-specific, millisecond-precision and reversable modulation of a variety of brain areas.

We hypothesized that our inhibitory opsin would reduce electrophysiological signal power when stimulated, but we were surprised to find an increase in electrophysiological signal power. We suspect upper layers of the cortex, which are known to have higher ratios of inhibitory neurons than other layers (Meyer *et al.*, 2011; Kooijmans *et al.*, 2020), were more strongly inhibited due to greater proximity to the light source, thus netting higher electrophysiological signal power from lower cortical layers. In addition, as we were recording LFPs, the recorded signals represented a summation of neural activity which adds more complexity to the interpretation of the result. Also, our inhibition length of 900ms is longer than commonly used and the opsin kinetics of inhibition at this time length are not well studied.

Collectively, our methods make cutting-edge neuroengineering techniques approachable to monkey researchers with diverse expertise. Importantly, our work has demonstrated the success of large-scale optogenetic stimulation by three key measures: epifluorescence imaging, electrophysiology, and behavior. Together, these data validate our chronic setup and indicate our method is suitable for fostering large-scale and long-term optogenetic experimentation with electrophysiology in behaving NHPs.

Our interface opens the door for investigation into a variety of questions which require flexible control of neural activation with high temporal and spatial precision. For example, this interface will better enable the study of plasticity and functional connectivity between different cortical areas (Yazdan-Shahmorad *et al.*, 2018; Bloch *et al.*, 2019, 2022) and the study of neurological disorders such as stroke, which is a particularly prevalent area of research in the cortex (Choi *et al.*, 2018; Frías *et al.*, 2018; Obi *et al.*, 2018; Yazdan-Shahmorad *et al.*, 2018; K Khateeb *et al.*, 2019; Khateeb *et al.*, 2022; Zhou *et al.*, 2022). Notably, the optical access afforded by our interface is compatible with photothrombotic stroke model (Yao and Yazdan-Shahmorad, 2018; K Khateeb *et al.*, 2019; Devon J. Griggs *et al.*, 2021; Khateeb *et al.*, 2022; Zhou *et al.*, 2022) which allows for repeated induction and quantification of precise spatially localized lesions. Importantly, our interface supports simultaneous optogenetic and electrical stimulation, which is useful for experiments comparing differences in their effects. Research in this area is key because optogenetic stimulation is still largely in a preclinical stage, yet may reveal diagnostic or therapeutic opportunities. In such a scenario, an interface supporting parallel development of both optogenetic and electrical stimulation protocols could speed the translation of the new diagnostic or therapeutic techniques from animals to humans. We also anticipate that energy storage technologies and wireless data transfer technologies will be integrated into designs based on our interface to make our work practical for experimentation on freely moving animals. This would further open the door for the advancement of translational optogenetic recovery models for neurodegenerative diseases, stroke, and traumatic brain injury.

Methods

MMAD

To address tissue growth concerns, we chose to embed a multi-modal artificial dura (MMAD, model “58ot58ri ecog vX”, Ripple Neuro) electrocorticography (EcoG) array, similar to past work (Griggs *et al.*, 2021), into an artificial dura (Fig. 1B). Briefly, the MMAD consisted of four conductive layers sandwiched between five transparent layers of medical-grade copolymer, where the conductive layers were aligned with each other to maximize optical access. Conductive layers were composed of platinum particles dispersed in a matrix within the polymer for mechanical flexibility, which was complimented by the absence of printed circuit boards (PCBs). This design choice also side-steps the challenge of protecting PCBs from

moisture and fluids. The MMAD was designed in close collaboration with Ripple Neuro (Salt Lake City, UT, USA), who printed the arrays.

A traditional artificial dura is generally described as “top-hat” shaped: a single silicone piece is composed of a cylindrical wall with an optical window and a flange extending under the native dura to reduce regrowth of the native dura (Fig. 1B), (Yazdan-Shahmorad *et al.*, 2016). Here, we embedded the MMAD into the artificial dura with a previously described molding process (Griggs_SPIE_2019). Disassembly of the mold resulted in an artificial dura (skirt diameter: 31.2 mm; skirt thickness: 0.4 mm; wall outer diameter: 21.0 mm; wall inner diameter: 18.6 mm; height: 11.9 mm) with a fully embedded MMAD (Fig. 1B).

Detachable MMAD cables, chamber, and cap

The detachable MMAD cables, chamber, and cap are detailed in (Griggs 2022a) and are reprinted here for completion.

We designed the entire MMAD, including the cables, as a corrosion resistant single piece. When not in use, the cables of the MMAD were housed in the chamber, and during experiments were attached to the electrophysiological recording and stimulation hardware.

The chamber is a cylinder (internal diameter 23 mm) 3D-printed and milled from titanium (Ti-6Al-4V, Hybex Innovations, Anjou, Quebec, Canada) that provides a sturdy base for experimental equipment and protects the brain when the animal is freely moving between experiments, similar to previous work (Yazdan-Shahmorad *et al.*, 2016). To design the chamber to fit the skull of each animal, we generated a custom software pipeline (MATLAB, MathWorks, Natick, MA, USA; Solidworks, Waltham, MA, USA) based on past work (Ojemann *et al.*, 2020) to extract the skull curvature from an MRI file and then used the curvature to design a chamber with a unique skull-fitting skirt (Fig. 1A, C). We included an intracranial rim under the skirt to aid in positioning the chamber on the skull during surgery and to reduce tissue growth around the wall of the MMAD, and we included 12 holes in the skirt for skull tap screws (Christ Instruments, Hagerstown, MD, USA). We incorporated concentric tracks within the chamber to store the MMAD cables between experiments (Fig. 1C-D). We threaded the chamber with a high thread pitch (M38 $\times$ 1.5) to allow either affixation of experimental equipment or a cap to close the chamber between experiments, and tabs on the outside of the chamber prevented users from screwing equipment too far down the chamber (Fig. 1C).

We designed a cap (stainless steel 304, machined by Hybex Innovations) to close off the chamber between experiments. We applied polytetrafluoroethylene (PTFE) tape to the threads before each closure to prevent binding of the threads. We used a set screw (HSS11400187HD, #4-40 X 3/16) to fasten the cap shut against the top rim of the chamber.

Ring, Tray, and PCB Clamp Connector

Our experimental equipment included a ring, a clamp connector tray, and PCB clamp connectors (Fig. 1E-G). To prepare the interface for recording and stimulation, we screwed the ring onto the chamber and secured the clamp connector tray onto the ring with 2–4 screws (M2.5 × 6). The ring was machined from stainless steel (304; Hybex Innovations) and the clamp connector tray was 3D printed with titanium (TI-64; i.materialise, Leuven, Belgium). The ring had 32 screw holes (M2.5) around its rim to offer optimal flexibility in the orientation of attachments. The upper face of the ring included a slight recess just above the threads for additional hardware to be easily aligned with the ring for fastening. The bottom of the clamp connector tray included an intra-ring rim around the window to fit into the recess of the ring. The base of the tray included a right-angled outer rim for grounding electronics as well as four elevated screw holes (M2.5 or 4–40 thread) designed for optical stimulation equipment to be secured to the tray. We designed the tray to be compatible with both this work’s LED-based optical stimulation setup and our lab’s previously described fiberoptic-based setup (Yazdan 2016). We designed the arms of the tray to be low

weight and provide support for the MMAD cables as they extended from the artificial dura walls upwards and at an angle (45°). The tip of the tray arm was approximately 9.5 cm away from the cranial window. The arms of the tray supported PCB clamp connectors (Fig. 1E-F; Ripple Neuro) held in place by rubber bands (not shown), which secured the MMAD cables approximately 4.5 cm away from the brain and provided an electrical connection between the cables and commercial neurophysiology equipment (Grapevine Nomad, Ripple Neuro). Two of these setups may be simultaneously deployed bilaterally, although this work is exclusive to the left hemisphere.

Optical Stimulation Setup

To optically stimulate across large cortical areas (~1 cm^2), we incorporated either a 4 × 4 or 3 × 5 LED array (Fig. 2). To prevent the LED array from moving when in use, we designed a system involving a rod, tube, coverslip, and stimulation holder (Fig. 1G). We custom designed and 3D printed the rod, tube, and stimulation holder in-house with FDM PLA. We glued (Better Ultimate Adhesive, slow dry, Cemedine Co. Ltd., Shinagawa City, Tokyo, Japan) the coverslip (16 mm diameter) to one end of the tube. The base, which we affixed to the clamp connector tray, secured the coverslip against the top side of the MMAD and a set screw (HSS11400187HD, #4–40 × 3/16) secured the rod inside the tube. The rod positioned the LED array over the brain and above the coverslip. Our setup created an air gap between the LEDs and the coverslip, which protected the brain from conductive heating. Four long screws in the base secure the tube in place.

The LED array consisted of a circular custom rigid printed circuit board (PCB; 15.5 mm diameter; OSH Park, Portland, OR, USA). The LEDs (L1C1-RED1 or L1C1-BLU1, Lumileds, Schipol, Netherlands) of the 3 × 5 array were individually addressable, while the 4 × 4 array was wired to drive four individually addressable columns of LEDs. Both PCB versions included an internal copper plane for heat dissipation. We soldered the LEDs on one side of the PCB and a connector (501331-0807, Molex Inc., Lisle, IL, USA) on the other side. We drove the LEDs with a custom external PCB (OSH Park) connected to the LED array by a cable assembly (151330806, Molex Inc.). The external PCB consisted of voltage-controlled current sources. Each driving circuit consisted of a high output current opamp (TLE2301, Texas Instruments, Dallas, TX, USA) with a capacitor between the compensation network terminals (not shown) and a resister (15 Ω, 5 W, SMW515RJT, TE Connectivity, Schaffhausen, Switzerland), where the resister value controls the ratio of input voltage to output current. The input voltage is provided by our neurophysiology system's analog voltage output (Ripple Neuro) and is prevented from floating high (e.g., when the neurophysiology system is turned off) by a pull-down resistor (10 kΩ). We powered our PCBs with two in-series batteries (9V Li-ion rechargeable, Keenstone, Industry, CA, USA) to reduce line noise in our electrophysiological recordings. We depict our circuit and arrays in Fig. 2 and Supp. Fig. 1.

We also designed our clamp connector tray to be compatible with our previous laser stimulation setup (Yazdan-Shahmorad *et al.*, 2015), and its slightly modified version is briefly described here. Our setup was composed of a base (similar to the base for the LED array setup) and stimulation ring, both FDM 3D-printed in-house with black PLA. The base was screwed to the clamp connector tray. The stimulation ring guided telescoping cannulas (HTX-13R, HTX-16T, and HTX-18T, Component Supply Co., Sparta, TN, USA) to different stimulation locations, and optical fibers were slid through the cannulas. We coupled lasers (520 nm or 638 nm, LDFLS_520_520_638_638, Doric Lenses, Quebec, Canada) to the optical fibers.

Animals

We recruited two healthy, socially housed male rhesus macaques (*Macaca mulatta*; Monkey H: 9 y, 13 kg; Monkey L: 9 y, 11 kg) for this study. Water was available *ad libitum* and no feeding restriction or scheduling protocols were used. All animal care and experiments were approved by the University of Washington's

Office of Animal Welfare, the Internal Animal Care and Use Committee, and the Washington National Primate Research Center.

Surgical procedures

Surgical procedures similar to those in other works (Yazdan-Shahmorad et al., 2016) were conducted on both animals. We sedated the animal and shaved its head before anesthetizing and head-fixing him in a stereotactic frame. The respiratory rate, heart rate, and body temperature of the animal were monitored throughout the procedure. We created an incision several cm long to the left of the head's midline and used elevators to peel back the skin, musculature, and muscle fascia. Then we test-fitted a plastic version of the chamber to the skull and marked the center of the craniotomy. We confirmed the location of the center of the craniotomy with stereotactic coordinates. Using a 25-mm diameter trephine, we resected the skull. Then we resected the dura by lifting the dura with a curved needle and trimming with ophthalmic scissors, revealing the posterior parietal cortex. After test-fitting the titanium chamber to the skull, we placed a flat transparent silicone artificial dura on the brain to keep the brain moist.

With the craniotomy complete, we performed convection-enhanced delivery (CED) of Jaws (rAAV8/hSyn-Jaws-KGC-GFP-ER2, 5.4 x 10^{12} virus molecules/mL, University of North Carolina Vector Core), a red-shifted inhibitory viral vector, which has previously been used to evoke behavioral effects in NHPs (Chuong *et al.*, 2014; Acker *et al.*, 2016). Prior to surgery, we had manufactured a 1-mm stepped-tip silica cannula (inner cannula: inner diameter 320 µm, outer diameter 435 µm; outer cannula: inner diameter 450 µm, outer diameter 673 µm; Polymicro Technologies, Phoenix, AZ, USA) for CED similar to our previous work (A Yazdan-Shahmorad *et al.*, 2016; Griggs *et al.*, 2022b). We removed the artificial dura from the brain, punched a hole through the artificial dura, laid the artificial dura back on the brain, and then inserted the cannula tip through the hole in the artificial dura and approximately 2 mm into the cortex and began the infusion protocol similar to previous work (Yazdan-Shahmorad *et al.*, 2016). We infused up to 50 µL at each of several locations across the cortex. Unlike previous work where live MRI was performed during infusions (Yazdan-Shahmorad *et al.*, 2016; Karam Khateeb *et al.*, 2019), we elected to monitor the infusions by eye to simplify the process. Once complete, we placed a new artificial dura on the brain.

Following CED, we placed the chamber on the skull. We used a piezo-drill to drill screw holes and we used skull tap screws (Crist Instruments) to secure the chamber to the skull. We removed the artificial dura and placed the MMAD on the brain and coiled the cables in the slots of the chamber (Fig. 1D). Finally, we applied PTFE tape to the threads, closed the cap, secured the set screw, and applied bone wax to the set screw. Having secured and closed the chamber, we closed the wound using standard surgical techniques. Following the surgery, the animal was placed on oral antibiotics for 4–5 weeks.

We had implanted a titanium headpost (Crist Instruments) on each animal prior to this study (Ojemann *et al.*, 2020).

Chamber cleaning and hardware assembly

We developed a process of cleaning the chamber at least three times each week, starting one week after the surgery, and assembling hardware on the chamber for experimentation. We briefly describe the processes here. After head-fixing the animal, we cleaned the cap, chamber margin, surrounding scalp and outside of the ear. Once clean, we removed the cap and cleaned the margin. We used cotton-tipped applicators dipped in saline to clean the threads and the rim of the chamber. To clean the inside of the chamber, we used a sterile syringe and blunt-tipped needle to spray warm saline on and in the MMAD, its cables, and the storage area for the cables. We removed the saline during rinses with a sterile suction tip.

To prepare for electrophysiological recording, we screwed a sterile ring onto the chamber and screwed a sterile clamp connector tray to the ring with sterile screws. Then we removed the cables from the chamber and placed them into the clamp connectors, and rubber banded the clamp connectors to the clamp connector tray arms.

To prepare for optical stimulation, we attached the stimulation holder for either the LED array or for the lasers. In the case of laser stimulation, we slid the stimulation ring down around the stimulation holder and we inserted the fiberoptic cable through the cannula and cannula through the simulation ring. In the case of LED stimulation, we assembled the LED array, rod, and tube with coverslip, and secured them together with a set screw. Then we lowered the assembly down through the stimulation holder and into the MMAD, and secured the tube into place with set screws. Black foil was wrapped around the stimulation equipment to prevent the animal from seeing any lighting changes in the experimental rig due to stimulation.

When the experiment was complete, we removed the stimulation equipment and recording equipment, cleaned the MMAD, and used forceps to return the cables back into the chamber as shown in Fig. 5. Then we cleaned the chamber using saline irrigation and applied antibiotic to the center of the MMAD. To avoid building antibiotic resistance, we rotated through polymyxin B sulfate, gentamicin, and amikacin sulfate antibiotics approximately every 4-5 weeks. To prevent the cap from binding to the chamber between cleanings, we wrapped a strip of sterile PTFE around a portion of the threads. Then we screwed on a sterile cap and tightened by hand, and then we tightened the set screw. Finally, we returned the animal to his cage.

Behavioral task

Subjects were trained to perform a center-out reach task (Fig. 4. A reflective sticker, placed on the middle finger of the animal's right hand (contralateral to the implant), was tracked by a motion capture system (Motive, OptiTrack, NaturalPoint Inc., Corvallis, OR, USA) to monitor the movement of the animal's finger. This tracking data was live-processed in MATLAB. A small white circle, corresponding to the finger location, appeared on the screen when the animal's finger came within a specified distance of it.

The task initiated with a green circle at the screen's center. Once the animal placed its finger into this initial circle, a red target circle would manifest in one of four reach directions (up, down, left, or right), marking the onset of stimulation. A go-tone would sound after a random delay of several hundred milliseconds, during which the animal had to keep its finger within the initial circle. The animal would then touch the target circle to earn a juice reward (5-RLD-D1 D.A.R.I.S., Crist Instruments). The animal remained in a head-fixed state, situated within a dark experimental rig, throughout the task. We designed and operated the task in MATLAB using Psychophysics Toolbox version 3 (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007), as done in previous studies (Griggs et al., 2020).

Behavioral analysis

Reach times were analyzed from the behavioral trials. A reach time was calculated as the time between the onset of the go tone and the monkey's finger being placed in the target circle. All successful trials from all sessions were combined and labeled as stim or nonstim trial depending on if stimulation had occurred that trial, and also labeled with the direction of the reach. Depending on the analysis, trials were either combined or kept separate between the two monkeys. Relevant groups were compared statistically by paired t-test of the distribution of reach times.

Neural data analysis

We analyzed the neural activity recorded during optogenetic inhibition during a reach, to quantify the change in neural activity underlying the disrupted reach. All neural data analysis was conducted separately for the two monkeys.

The process involved segregating all reach trials into stimulation and non-stimulation trials. Each trial was then split into three epochs: before stimulation, during stimulation, and after stimulation. During the "before stimulation" period, we analyzed the 500 milliseconds before stimulation began. The "during stimulation" period corresponded to the LFP data recorded during the 900ms stimulation, excluding the first 50ms of the stimulation. The "after stimulation" period corresponded to the 900 milliseconds following the end of the stimulation, excluding the first 50ms. This decision was informed by the challenges associated with removing artifacts during this specific period. For each of these epochs, artifact rejection was applied.

Time-frequency analysis

Each epoch was analyzed in the time-frequency domain by multitaper based analysis of 50ms non-overlapping windows. Five tapers were used, with a 128 sample length Fast Fourier Tranform (FFT). For a given trial, the three epochs of "before stimulation", "during stimulation", and "after stimulation" were realigned and rejoined along the time axis, with a 50ms empty placeholder corresponding to the removed 50ms immediately after stim onset and after stim offset. Frequencies above 200Hz were excluded from analysis.

For spectrograms depicting the average time-frequency activity of a specific condition (e.g., stimulated reach of monkey H), the spectrograms were normalized to the "before stim" epoch activity, on a trial-by-trial basis. For every channel of every trial, the average power in the "before stim" epoch for a specific frequency was calculated and then each time-frequency bin of that frequency for all bins in the three epochs were divided by that value. The values were then log base 10 transformed.

For spectrograms depicting the time-frequency ratios of stimulated to unstimulated condition, no within-trial normalization was applied. Instead, the mean spectrograms for the stimulated and unstimulated conditions were calculated, averaged over all trials and channels and sessions. Then each time-frequency bin of the stimulated condition spectrogram was divided by that of the unstimulated condition spectrogram, and the data were log base 10 transformed.

Statistics were calculated between distributions of stimulated and unstimulated conditions. For the reaching condition data, the unstimulated condition was defined as a reach undertaken in the absence of stimulation. For the data recorded in the absence of behavior such as for the single LED and laser data, the unstimulated condition was defined as the data recorded before stimulation. Time-frequency bins of spectrograms were grouped over all trials of all channels of all sessions, log base 10 transformed, and then the distributions were compared via paired t test. To account for multiple comparisons and mitigate the risk of Type I errors, a Bonferroni correction was applied post-hoc.

Frequency analysis

For analysis of the frequency-domain-only content of the signal, visualized by power spectral density (PSD) plots in the Results section, similar analyses were performed. The epochs were identical, except that the during stim epoch had not only its first 50ms of data excluded from analysis but also the last 50ms (bringing the total length from 850ms to 800ms). The neural data was transformed to the time-frequency domain by multitaper, but with parameters selected to prioritize information along the frequency axis. Windows were 200ms long, with 50% overlap. 6 tapers, and a 256 length FFT were used. From the

resulting time-frequency data, the data were averaged within an epoch over the time axis, yielding one power value per frequency.

The statistical calculations mirrored that of the time-frequency analysis, with p-values being derived by considering all trials, channels, and sessions concurrently, and by Bonferroni correction applied post-hoc.

Artifact rejection

To isolate photoelectric artifacts from the recorded data, acquired during simultaneous stimulation in non-human primates, we generated a dataset of these artifacts, absent any neural signals, by stimulating an electrode array submerged in a saline bath. A model, fit to this saline data, was used to compute the photoelectric artifact signals in the non-human primate data. These were then subtracted from the signals. The procedure involved fitting an artifact basis to the saline data with the following steps:

1. The time surrounding each stimulation in the saline data was isolated, from 100 milliseconds before the onset of stimulation to one second after the end of stimulation. Each of the 30 channels were stacked with their 100 corresponding pulses to form a (30 by 100 by length of time) matrix.
2. The matrix was split along the time axis into three matrices corresponding to before stimulation during stimulation, and after stimulation.
2. The saline ECoG data at the time of stimulation onset and offset (data point time=0 for matrices “during stimulation” and “after stimulation”) were set to zero millivolts and the whole traces were shifted accordingly.
3. Principal component analyses (PCA) were performed individually on matrices “during stimulation” and “after stimulation” and the three top components were extracted from each.

Then the NHP ECoG data were handled with the following steps:

4. The NHP data was isolated and stacked in the same way as the saline data, such that a data matrix of shape (channels by pulses by length of time), with identical time length, was formed.
5. The matrix was split along the time axis into three matrices corresponding to before stimulation during stimulation, and after stimulation.
6. The NHP ECoG data at the time of stimulation onset and offset (data point time=0 for matrices “during stimulation” and “after stimulation”) were set to zero millivolts and the whole traces were shifted accordingly.
7. The NHP data of the “during stimulation” and “after stimulation” matrices were projected onto the corresponding PCA saline basis and inverse transformed to obtain the reconstruction of the non-human primate data from the saline data bases.
8. This reconstruction was subtracted from the raw NHP recorded data.
9. The three matrices were rejoined along the time axis, adding an offset to the later matrices such that the average LFP value of the last 50ms of the preceding matrix and the first 50ms of the later matrix had the same average value.

This procedure enabled the separation of the recording data in non-human primates into neural activity and photoelectric artifact, and the removal of the photoelectric artifact, thus retaining only the neural activity in the recorded signal.

Acknowledgments

We thank WaNPRC staff for their help with animal care and surgeries. We thank Ripple Neuro, especially Alex Johnson, Jose Ortega, and Jessi Michel, for their help with designing and manufacturing the MMADs.

Chapter 5: How to infer structural connectivity from LFP-level functional connectivity: a simulation study

Julien Bloch, Larry Shupe, Patrick Zhang, David Lewis, Eric Shea-Brown, Eberhard Fetz, Azadeh Yazdan-Shahmorad

Abstract

Local field potentials (LFPs) are stable and readily accessible measures of neural activity that have been employed to calculate functional connectivity between various brain regions. However, the relationship between structural connectivity and LFP-level functional connectivity remains poorly understood. To address this gap, we developed a large-scale, biophysically plausible neural network simulation in the VERTEX simulation platform that generates LFPs using validated conduction equations. This simulation is based on connectivity statistics derived from structural studies and features hundreds of thousands of virtual neurons that produce both spiking activity and LFPs. Our analysis of the simulation reveals that undirected functional connectivity in low-frequency ranges exhibits the strongest correlation with structural connectivity strength, while higher frequencies demonstrate weaker correlations and directed metrics show no correlation at all. These findings enhance our understanding of the origins and interpretation of LFP functional connectivity metrics and may have implications for future studies investigating brain function and dysfunction.

Introduction

The human brain is a complex network comprised of approximately 100 billion neurons and 100 trillion synapses, with an intricate organization of connections between neurons (Binzegger et al., 2004). Mapping this extensive network poses significant challenges due to the large number of structural nodes and connections. Therefore, experiments designed to quantify connectivity often focus on "functional connectivity," which represents large-scale statistical correlations in activity between brain regions (Friston, 2011; Bastos, 2016). Despite the widespread use of functional connectivity metrics in neuroscience, the relationship between functional and structural connectivity remains poorly understood.

Various modalities, such as fMRI, EEG, ECOG, and depth electrode recording of LFP and spikes, have been employed to study functional connectivity in brain activity (Buzsáki and Draguhn, 2004; Friston, 2011; Horwitz, 2003; Nolte et al., 2004). Although there has been some work investigating the relationship functional connectivity with structural connectivity (Friston, 2011; O'Reilly, 2013; Fukushima, 2017) there is still much to be explored. For example, functional relationships may emerge naturally as a result of strong structural connections between brain areas, but they may also emerge via the state-based modulation of brain activity coordination without impacting structural connectivity (Friston, 2011; O'Reilly, 2013; Fukushima, 2017). The degree and prevalence of these two sources of functional connectivity is not well understood.

Functional connectivity has been used as a proxy to infer the extent of physical connections between brain areas, particularly in rehabilitation contexts, where forming new or strengthening existing connections can be beneficial (Beliveau et al., 2017). However, there is limited direct evidence demonstrating that these types of metrics are correlated. Functional connectivity metrics can be categorized into three classes: symmetric connectivity, directed connectivity, and conditional connectivity (Horwitz, 2003). Symmetric connectivity measures the strength of connections between two brain regions in both directions, while directed connectivity measures connection strength in a specific direction. Conditional

connectivity measures the strength of connectivity between two brain regions based on the activity of another brain region (Horwitz, 2003). Furthermore, some metrics require stimulation for inference, whereas others can be inferred from recordings alone.

The relationship between functional and structural connectivity has been studied at large scales, particularly in the context of white matter tract-defined structural connectivity between discrete brain regions and functional connectivity computed from fMRI signals and, to a lesser extent, EEG signals. This approach provides valuable information about whole-brain statistics but lacks high-resolution insights. As a result, findings from large-scale studies cannot be directly applied to small scale experiments. As neural engineering approaches which aim to alter connectivity often operate at these small scales and rely on functional connectivity to infer their effects, there is a need to understand relationships at this level.

Our platform features a large-scale (>100,000 neuron) simulation of virtual neurons with a connectivity structure that mirrors the cerebral cortex (Tomsett et al., 2015). This simulation also generates readouts of voltages at surface and intracortical locations, analogous to local-field potential (LFP) recordings sampled from the cortex using electrodes. The model represents a section of primate cortex measuring 1.5mm x 1.5mm x 2.6mm. This size is comparable and relevant to a plethora of experiments quantifying functional connectivity between electrodes on neural implants such as Utah array electrodes (Zanos et al., 2011; Stevenson et al., 2012; Snyder et al., 2015; Bloch et al., 2019).

This innovative platform enables us to calculate functional connectivity metrics from the recorded LFP signals and compare them to the virtual structural connectivity defined by the model. We first compute various measures of structural connectivity strength between distinct virtual electrode locations. Next, we calculate standard functional connectivity metrics, including coherence and Granger causality, from the LFP signals collected from virtual electrodes and employ correlational analysis to investigate their relationship.

In this paper, we aim to investigate the relationship between functional connectivity and LFP-based metrics recorded from electrodes in a simulated brain at a spatial scale comparable to common electrode arrays like the Utah array. By examining the relationships between functional connectivity and LFP-based metrics, we aim to contribute to the growing body of knowledge on this subject and provide new insights into the complex interactions between brain regions.

We analyze the relationships between various functional connectivity metrics, such as coherence, Granger causality, and others, and compare them to the underlying structural connectivity defined in the simulation. We also examine the impact of experimental constraints, such as the placement of electrodes at constant depths in the brain, on the correlations between functional and structural connectivity metrics.

Furthermore, we explore the potential of combining different functional connectivity metrics in a data-driven manner to improve the prediction of structural connectivity using stepwise regression models. This approach will be tested on both excitatory and inhibitory structural connectivity and compared to the performance of individual metrics.

Finally, we investigate the ability of functional connectivity metrics to accurately estimate changes in structural connectivity induced by stimulation in a model with plastic synapses. By correlating changes in

functional connectivity metrics with changes in structural connectivity before and after stimulation, we aim to identify the most promising metrics for predicting connectivity changes induced by stimulation.

Overall, our study aims to shed light on the complex relationship between functional and structural connectivity at a small scale, which is crucial for understanding brain function and for advancing neural engineering approaches that aim to manipulate connectivity. By leveraging the power of large-scale simulations and a range of functional connectivity metrics, we hope to provide valuable insights that can inform future experimental designs and computational models.

Results

Biorealistic neural tissue and activity

To ensure the relevance of our simulation findings to real brains, we aimed for single-neuron resolution of neural activity, sub-neuron level of LFP creation and propagation, and biorealistic cell type statistics and connections, all while encompassing a broad spatial scale. Utilizing the Vertex 2.0 neuron simulation framework (Thornton 2019), which employs multicompartment models of neurons for precise local field potential (LFP) calculations and leverages real voltage propagation equations for discrete spatial point calculations, our platform features a large-scale simulation of over 100,000 virtual neurons. This simulation has a connectivity structure mirroring the cerebral cortex and generates readouts of voltages at both surface and intracortical locations, analogous to LFP recordings from the cortex using electrodes. Specifically, our study focused on a spatial scale of 1.5mm x 1.5mm x 2.6mm of cortex, which is representative of a section of the primate cortex and maintains applicability to research employing neural implants, such as the Utah arrays. Additionally, the foundation of our model is based on a published model from the Vertex 1.0 framework (Tomsett 2015) with anatomically derived neural cell type and connectivity statistics that closely match neuron statistics observed in the mammalian brain. The chosen depth of 2.6mm was selected to simulate primate brains more closely.

We modified the simulation by changing synapses from a set conductance value for specific neuron groups to a normally distributed conductance centered around the previously established value. This modification led to a smoother power spectral density (PSD), exhibiting primarily the characteristic $1/freq^{\alpha}$ observed in the brain with a small bump in the Gamma power (Fig. 1H-I). The time series of the LFP looked qualitatively similar to LFP from in vivo electrophysiology (Fig. 1G), unlike the original model for which the LFP was dominated by a clear Gamma rhythm.

In our model, we conducted an analysis of firing rate histograms and inter-spike interval (ISI) distributions within neuron groups to ensure the display of characteristic spiking statistics. We found that the firing rate distributions adhered to a log-normal pattern (Fig. 1D), a finding that aligns with empirical observations in both excitatory and inhibitory neuronal populations (Antolik et al., 2018; Kajiwara et al., 2020). Furthermore, the ISI histograms also exhibited a log-normal distribution (Fig. 1E-F), a characteristic feature of neuronal firing dynamics as demonstrated in both experimental and theoretical studies (van Vreeswijk and Sompolinsky, 1996; Matthews and Clements, 2014; Hasegawa, 2009). These findings underscore the physiological relevance of our model, as it successfully replicates key statistical properties of real neurons.

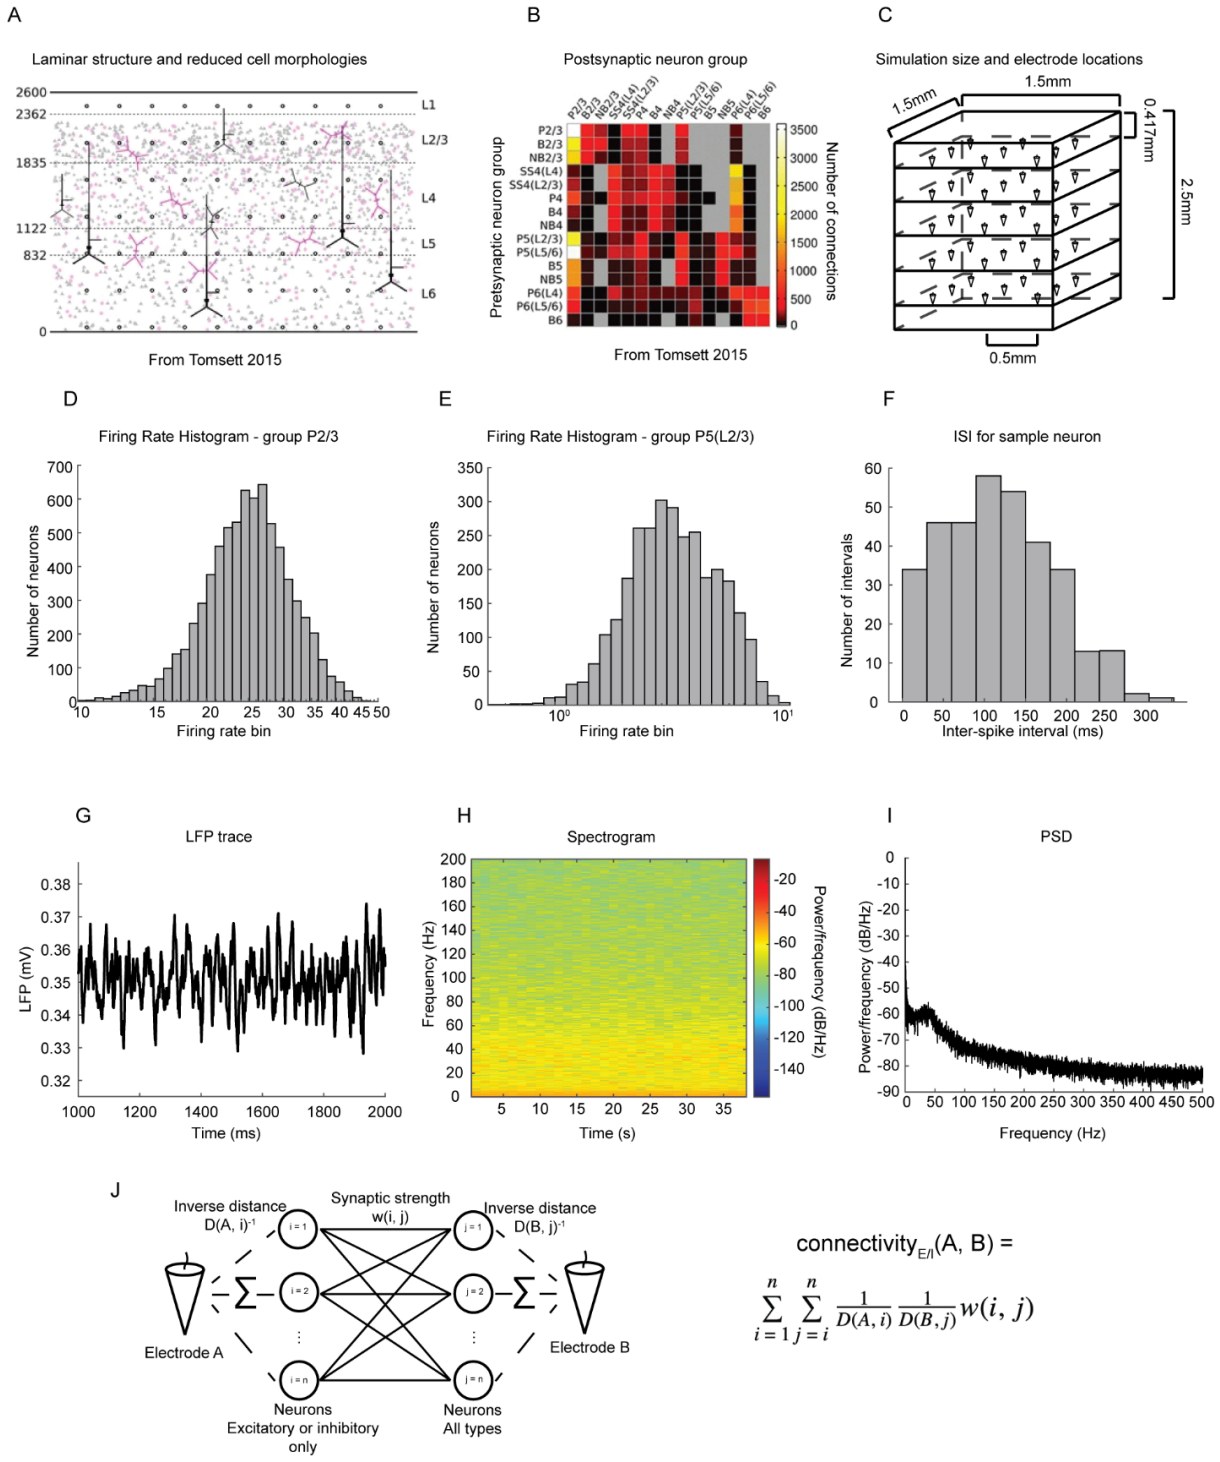

Figure 6: Model depiction and validation. A) Schematic showing laminar structure of model and examples of reduced neuron morphologies at different depths. B) Cell connectivity statistics, between all neuron types in the model. A) and B) from Tomsett 2015. C) Simulation scale and placement of recording electrodes, depicted as small cones. D) Firing rate histogram of pyramidal neurons in layers 2 and 3. E) Firing rate histogram of pyramidal neurons in layer 5 projecting to layers 2 and 3. F) Inter-spike interval for a randomly sampled. G) Example LFP trace from top-most central recording electrode. H) Spectrogram of signal from same electrode. I) Power spectral density from same electrode. J) Visualization and equation of the structural connectivity calculation between two electrodes.

Calculation of intrinsic functional connectivity metrics

To elucidate the relationship between structural and functional connectivity metrics, we calculated structural and functional connectivity between all virtual electrode pairs in the simulation. As detailed in Methods, the structural connectivity from electrode A to B is computed as the sum of all synaptic strengths between pairs of neurons, where each strength is inversely weighted by the distance from electrode A to the presynaptic neuron and by the distance of electrode B to the postsynaptic neuron (Fig. 1J). This measure represents a gross synaptic connectivity strength between two electrode locations. As our simulation had distinct excitatory and inhibitory neuron populations, we calculated structural connectivity independently for excitatory and inhibitory synapses, resulting in two distinct measures, which we term excitatory structural connectivity and inhibitory structural connectivity, respectively.

We employed four functional connectivity metrics, each corresponding to a unique class of functional connectivity metrics: pairwise symmetric, partial symmetric, pairwise directional, and conditional directional. The corresponding metrics were coherence, partial spectral coherence, Granger causality, and multivariate Granger causality. Each of these metrics was calculated for the following standard frequency bands: Delta (0.1-4Hz), Theta (4-8Hz), Alpha (8-12Hz), Beta (12-30Hz), Gamma (30-60Hz), High-Gamma (60-200Hz). Pairwise connectivity metrics assess the strength of statistical dependency between two signals, while conditional metrics determine the partial contribution of one signal to another, conditioned on all other recorded signals. Undirected metrics yield a similarity score for two signals, generating a symmetric connectivity matrix. In contrast, directed measures distinguish the influence of two signals on each other, producing a non-symmetric matrix. Each metric type offers its own advantages and disadvantages (Bastos 2016); thus, we evaluated one standard metric from each class.

Metric correlations

Functional connectivity metrics were assessed for each frequency band, and their relationships with each other, structural connectivity, and electrode distance were analyzed (Fig. 2). A heatmap displaying Spearman's correlation between all metrics is presented in Fig. 2A. First, the correlations between excitatory and inhibitory structural connectivity are highlighted within the black rectangular bounding box in the heatmap. These two distinct forms of structural connectivity exhibit a high correlation (Spearman correlation = 0.9) with each other, as they are implemented similarly in the model. As such, most of the later subfigures in Fig. 2 show correlations between functional connectivity metrics and the excitatory structural connectivity only, as the relationship with inhibitory structural connectivity tends to be similar.

The heatmap is arranged in blocks of metrics grouped by frequency band. Examining the heatmap through this organized block structure reveals high within-metric correlations across frequency bands for coherence, Granger causality, and multivariate Granger causality, while a relatively low correlation is observed between different metrics.

Fig. 2B-G features scatter plots that illustrate the relationships between metrics. Coherence displayed the highest correlation with excitatory and inhibitory structural connectivity, followed by multivariate Granger causality. Granger causality and partial coherence showed much lower correlations with structural connectivity in comparison to the other two metrics. These observations generally hold true for both excitatory and inhibitory structural connectivity, except for Granger causality, which demonstrated a strong correlation with inhibitory structural connectivity and an almost zero correlation with excitatory structural connectivity.

All functional connectivity metrics, except for multivariate Granger causality, exhibited weaker correlations with structural connectivity in the Gamma band.

The distance between electrodes exhibited a negative correlation with all functional connectivity metrics, indicating that a larger distance corresponds to weaker connections. Importantly, the negative correlation with distance is more pronounced for metrics that display a strong correlation with structural connectivity.

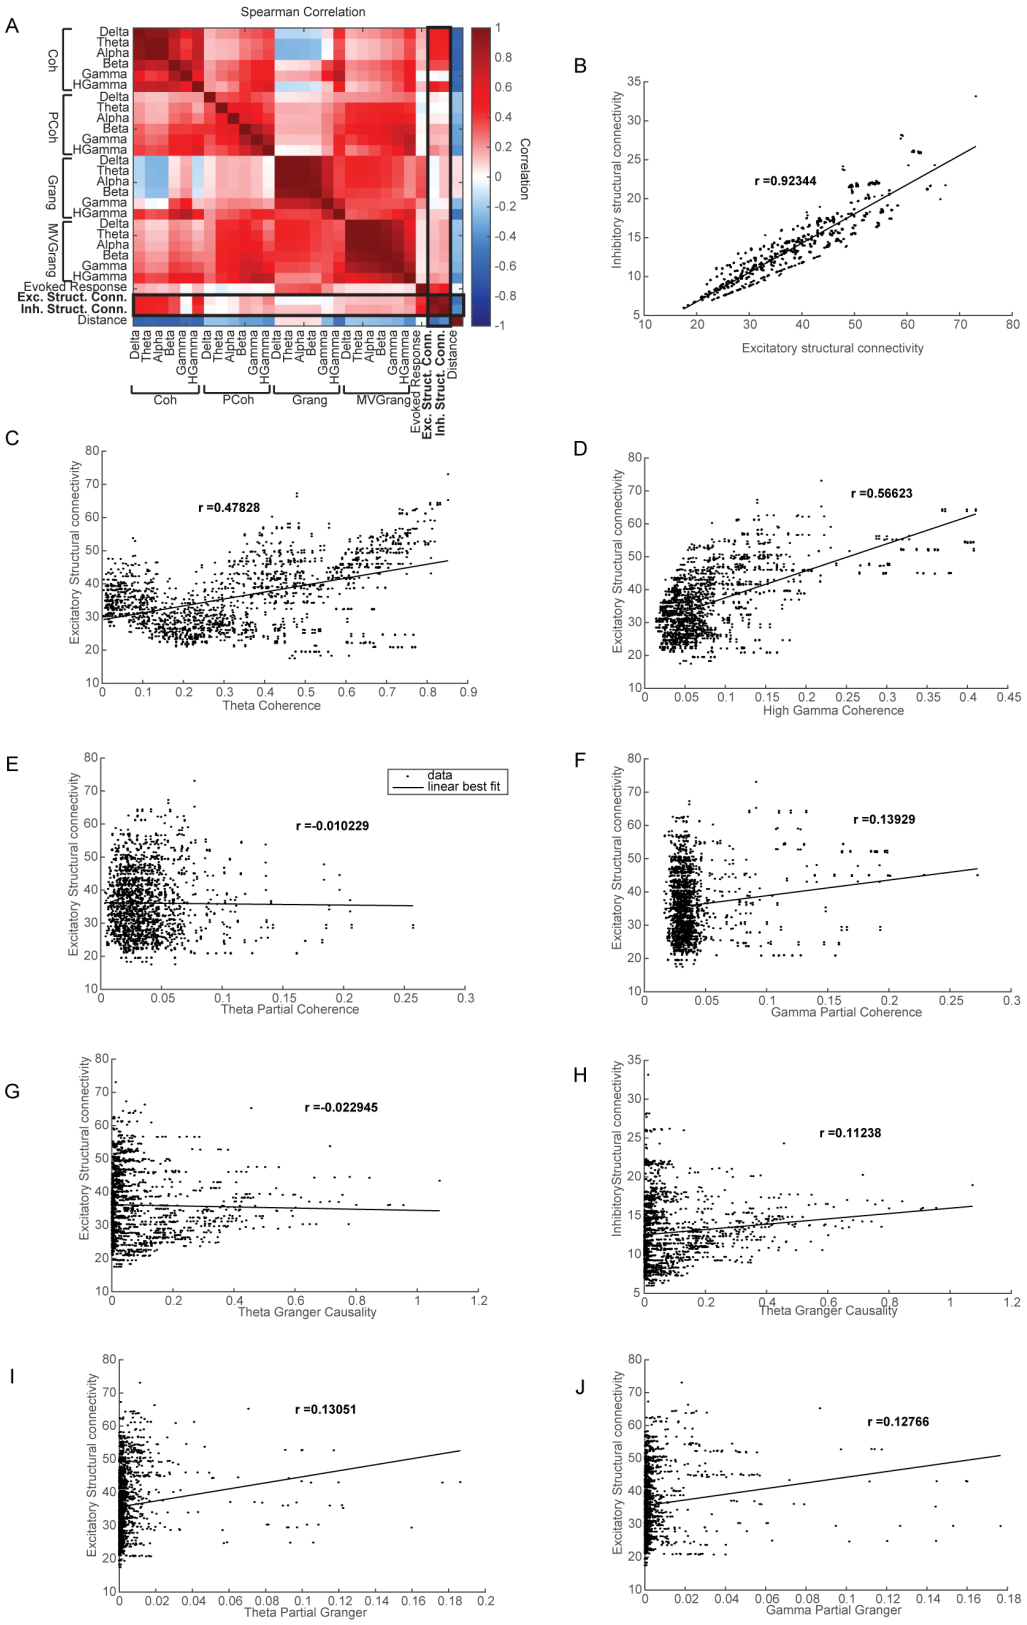

Figure 7: Correlations and scatter plots between functional connectivity metrics and structural connectivity. A) Heatmap of all correlations. Scatter plots of B) excitatory vs inhibitory structural connectivity, C) Theta coherence

and D) High Gamma coherence vs excitatory structural connectivity, E) Theta partial coherence and (F) High Gamma partial coherence vs excitatory structural connectivity, G) Theta Granger causality and H) High Gamma Granger causality vs excitatory structural connectivity, I) Theta multivariate Granger causality and J) High Gamma multivariate Granger causality vs excitatory structural connectivity.

Nonlinear Scaling of Granger-Based Metrics with Structural Connectivity

In the previous section, we presented only Spearman correlation. However, we computed both Pearson (Fig. 3A) and Spearman (Fig. 2A) correlation coefficients for all metrics. While Pearson correlation quantifies the degree of linear correlation, Spearman correlation captures rank-ordering agreement, thus detecting nonlinear correlations. For most metrics, the Pearson and Spearman correlation coefficients exhibited minimal differences. However, for Granger-based measures, particularly multivariate Granger causality, the Spearman correlation coefficients were substantially higher, indicating a nonlinear correlation with structural connectivity values.

Granger causality and especially multivariate Granger causality values were very strongly right-skewed, suggesting that a log transformation of their values would make their values Gaussian. Moreover, a log transformation significantly improved their linear correlation to structural connectivity (Fig. 3B-E). Note that the Spearman correlation is not affected by this transformation, as it computes similarity of rank ordering, which is preserved by transformation. As such, and since the use of raw Granger causality values is standard, in subsequent sections we chose to use the untransformed Granger causality values and report Spearman correlations.

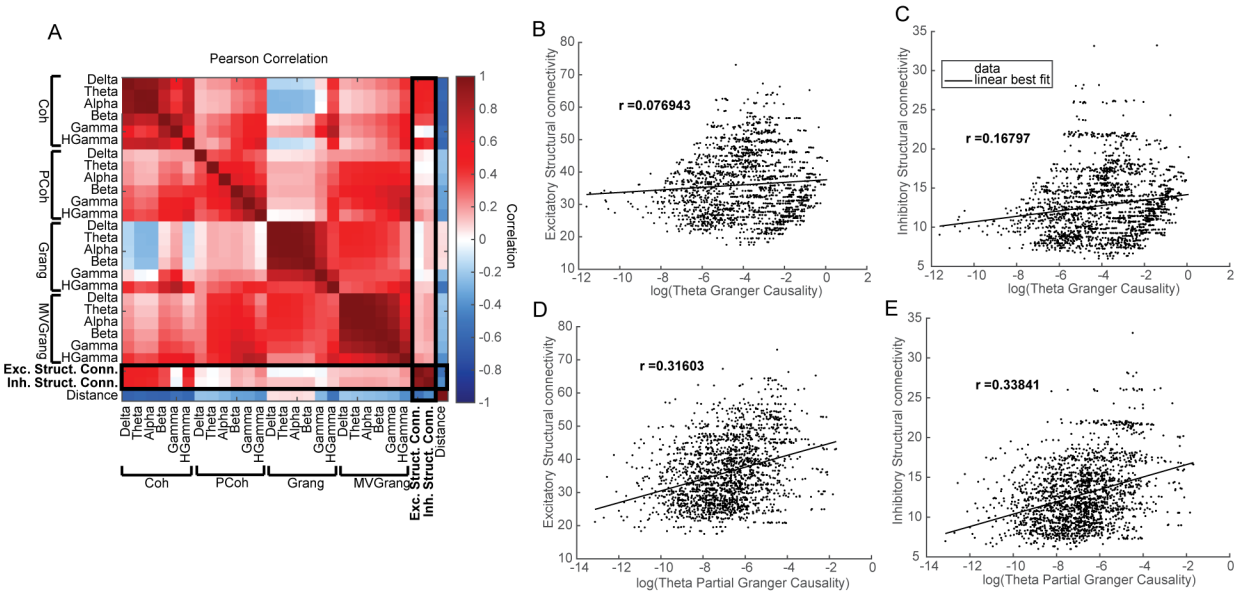

Figure 8: A) Pearson (linear) correlation between all tested metrics. B-C) log transformed Theta Granger causality vs (B) excitatory structural connectivity and vs. (C) inhibitory structural connectivity. D-E) log transformed Theta Partial Granger Causality vs. (D) excitatory structural connectivity and vs. (E) inhibitory structural connectivity.

Distance Information in Functional and Structural Connectivity Relationships

To further examine the relationship between functional and structural connectivity, we removed the distance information from the structural connectivity and assessed the correlation of functional connectivity metrics with this modified value (Fig. 4). We did this by fitting a linear model from distance to structural connectivity (independently for excitatory and inhibitory structural connectivity), and

subtracting the structural connectivity estimates generated by the model from the true structural connectivity values. Then the residual values were termed the “distance removed structural connectivity” and compared to the functional connectivity values.

In general, we observed similar trends to our previous findings: coherence exhibited the highest correlation, followed by multivariate Granger Causality, and pairwise Granger Causality showed selectivity for inhibitory structural connectivity (Fig. 4C and 4E). However, these correlations were reduced to about 0.2, suggesting that the distance information contributes to the predictive capability of functional connectivity metrics.

We also examined the residuals of the distance-based predictions of structural connectivity and found high variance for predictions of stronger structural connectivity (Fig. 4D). In other words, there was less information contained in distance for electrodes that were closer together, while distance could predict near-perfectly the structural connectivity for electrodes that were further apart.

With distance information removed, multivariate Granger causality continued to display a nonlinear scaling with structural connectivity, as evidenced by the substantially higher Spearman correlation coefficient compared to the Pearson correlation coefficient. In contrast, pairwise Granger causality did not exhibit a pronounced nonlinear trend, unlike our previous observations. Interestingly, coherence in the Alpha band increased from about 0 to nearly 0.2 for Spearman correlation.

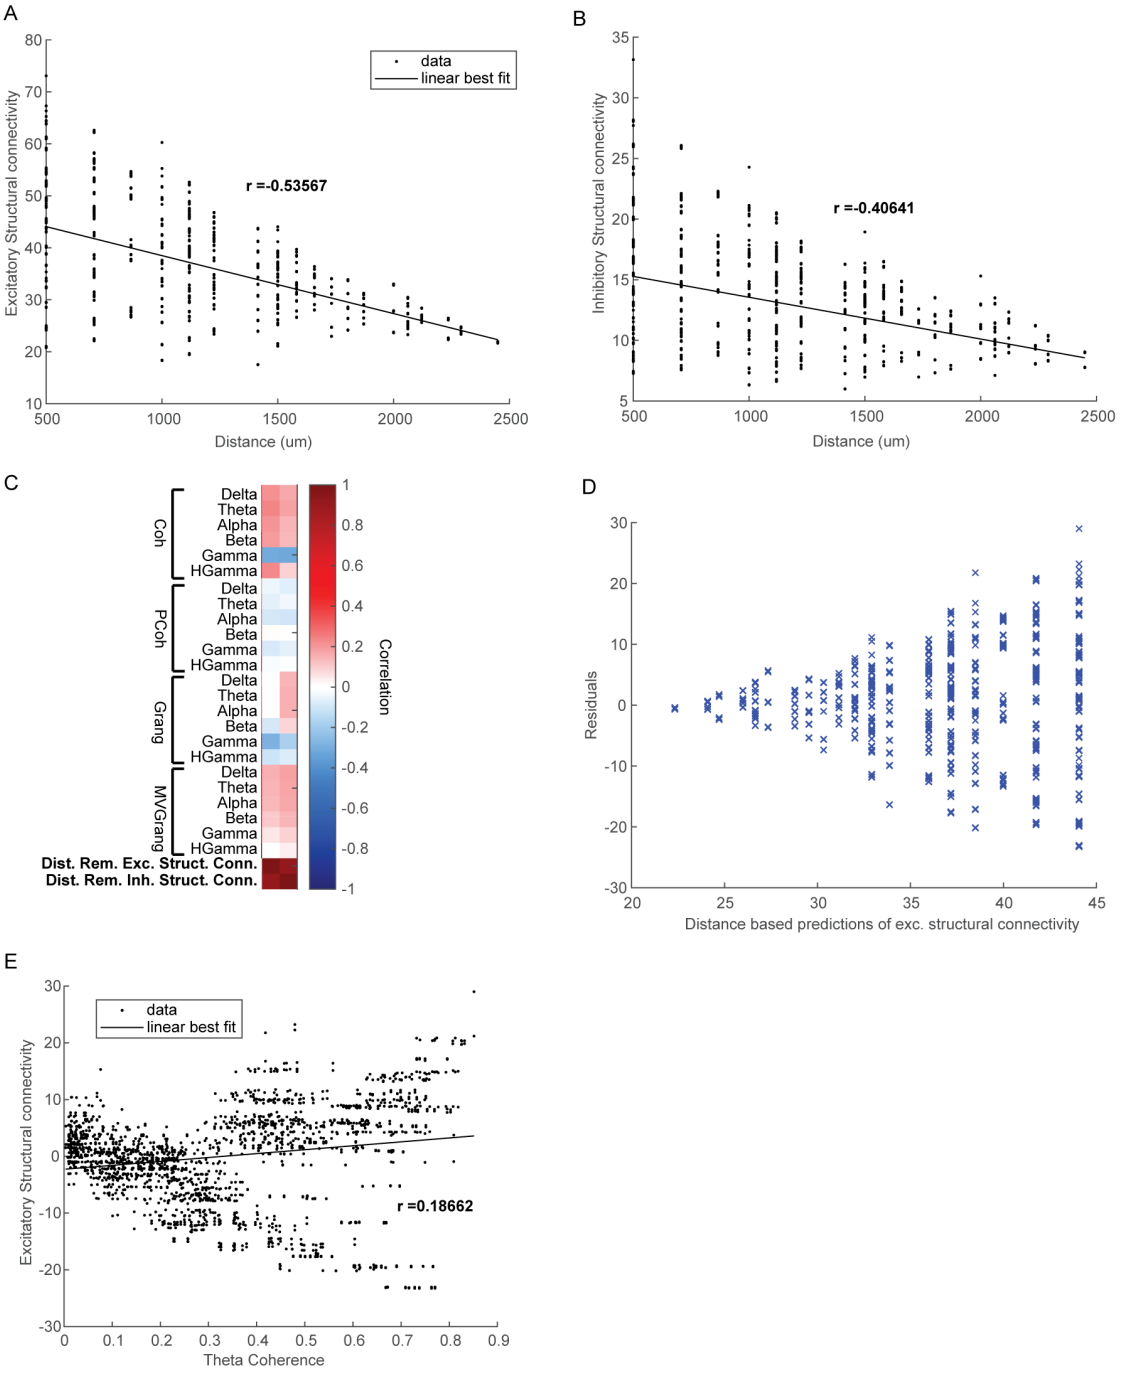

Figure 9: Distance effect on relationship between functional connectivity metrics and structural connectivity. A) Distance between electrodes vs excitatory structural connectivity and B) inhibitory structural connectivity. C) Correlations between functional connectivity metrics and structural connectivity with distance information removed. D) Residual plot of distance based prediction of excitatory structural connectivity. E) Theta coherence vs excitatory structural connectivity when distance information is removed.

Combining metrics via stepwise regression

We tested whether the metrics could be combined in a data-driven way to yield a better predictor of structural connectivity using a stepwise regression model (Fig. 5). The model was fit on randomly sampled

training data and optimized by Bayesian information criterion, with R-squared on the held-out test set of around 0.7 for both excitatory and inhibitory activities. The fitting and selection process for the excitatory and inhibitory structural connectivity are shown in Figs 5A and 5B respectively.

The results of the fit models tested on held-out test data are shown in Fig. 5A-D. We observe that a model trained to predict excitatory structural connectivity achieved an explained variance score (r2) of 0.65 (Fig. 5A), while the model trained to predict inhibitory structural connectivity achieved an r2 of 0.67 (Fig. 5B). We also tested the ability of a model to predict the distance removed structural connectivity described in the previous section and found no drop in performance. Specifically, we observed an r2 of 0.74 for the test set of distance removed excitatory structural connectivity (Fig. 5C), and r2=0.67 for distance removed inhibitory structural connectivity (Fig. 5D).

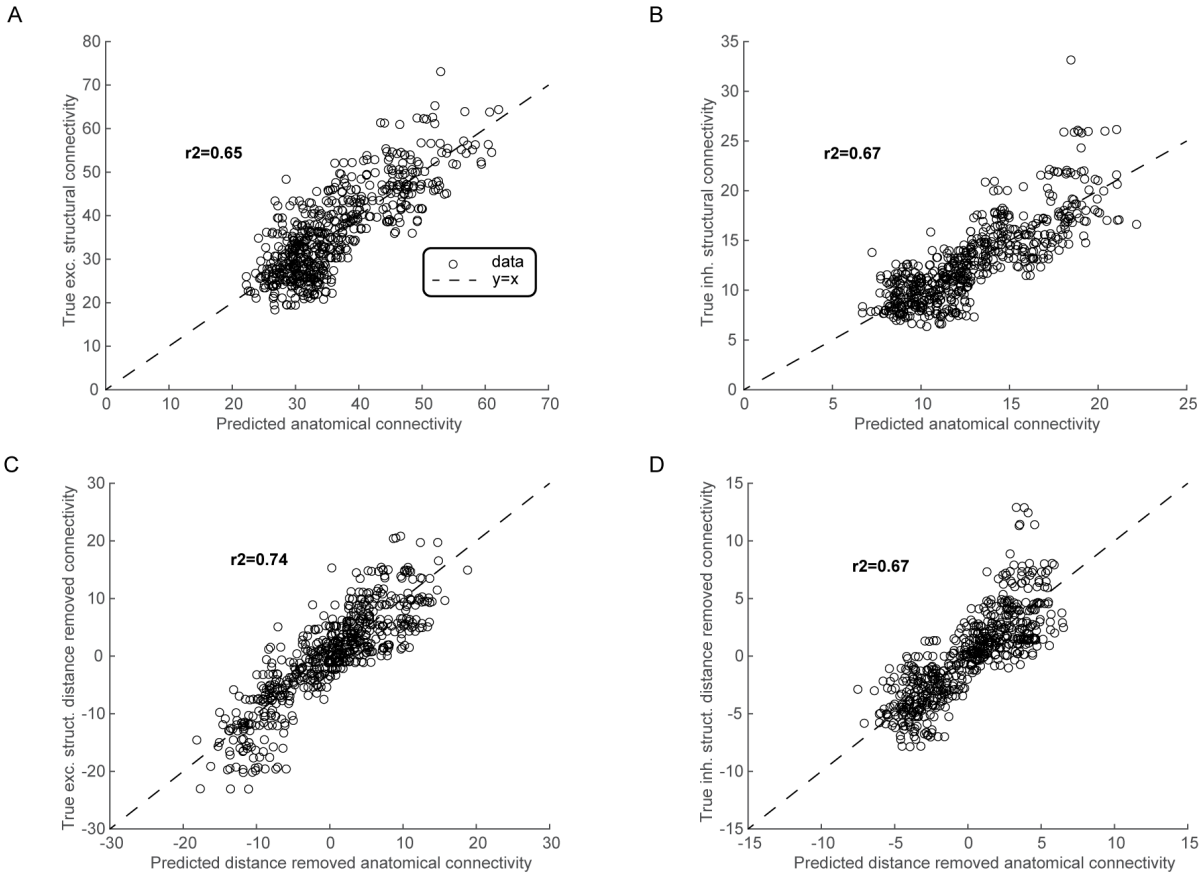

Figure 10: Data driven combination of functional connectivity metrics by stepwise regression. A) Predicted excitatory anatomical connectivity on held-out test data. B) same as A but for inhibitory. C-D) same as A-B but for predicting distance-removed structural connectivity.

Constant depth recordings

We further explored the possibility of obtaining higher correlations between functional and structural connectivity metrics by restricting the depth, or z-location, of electrode placements in the brain, as typically done in physical experiments. Specifically, we calculated the correlation of functional connectivity metrics to structural connectivity using only data recorded at the same depth, independently for all depths used in previous sections. The results revealed a substantial degree of correlation for all tested metrics in this constant-depth setting (Fig. 6). Then, by averaging across various layers and depths,

the average correlations were obtained, which are depicted in the heatmap in Figure 6A. These findings suggest that the majority of the examined connectivity metrics demonstrate a strong association with structural connectivity when experimental constraints, such as confining electrodes to specific depths, are employed.

Upon removing distance as a factor and rerunning the analysis, the functional connectivity metrics were less correlated with structural connectivity, with much variation over depths (Fig. 6F-H). Generally, the pairwise metrics exhibited superior performance. The results were similar for the Pearson correlation, with the note that Granger was selective for inhibitory connectivity only in the Pearson case.

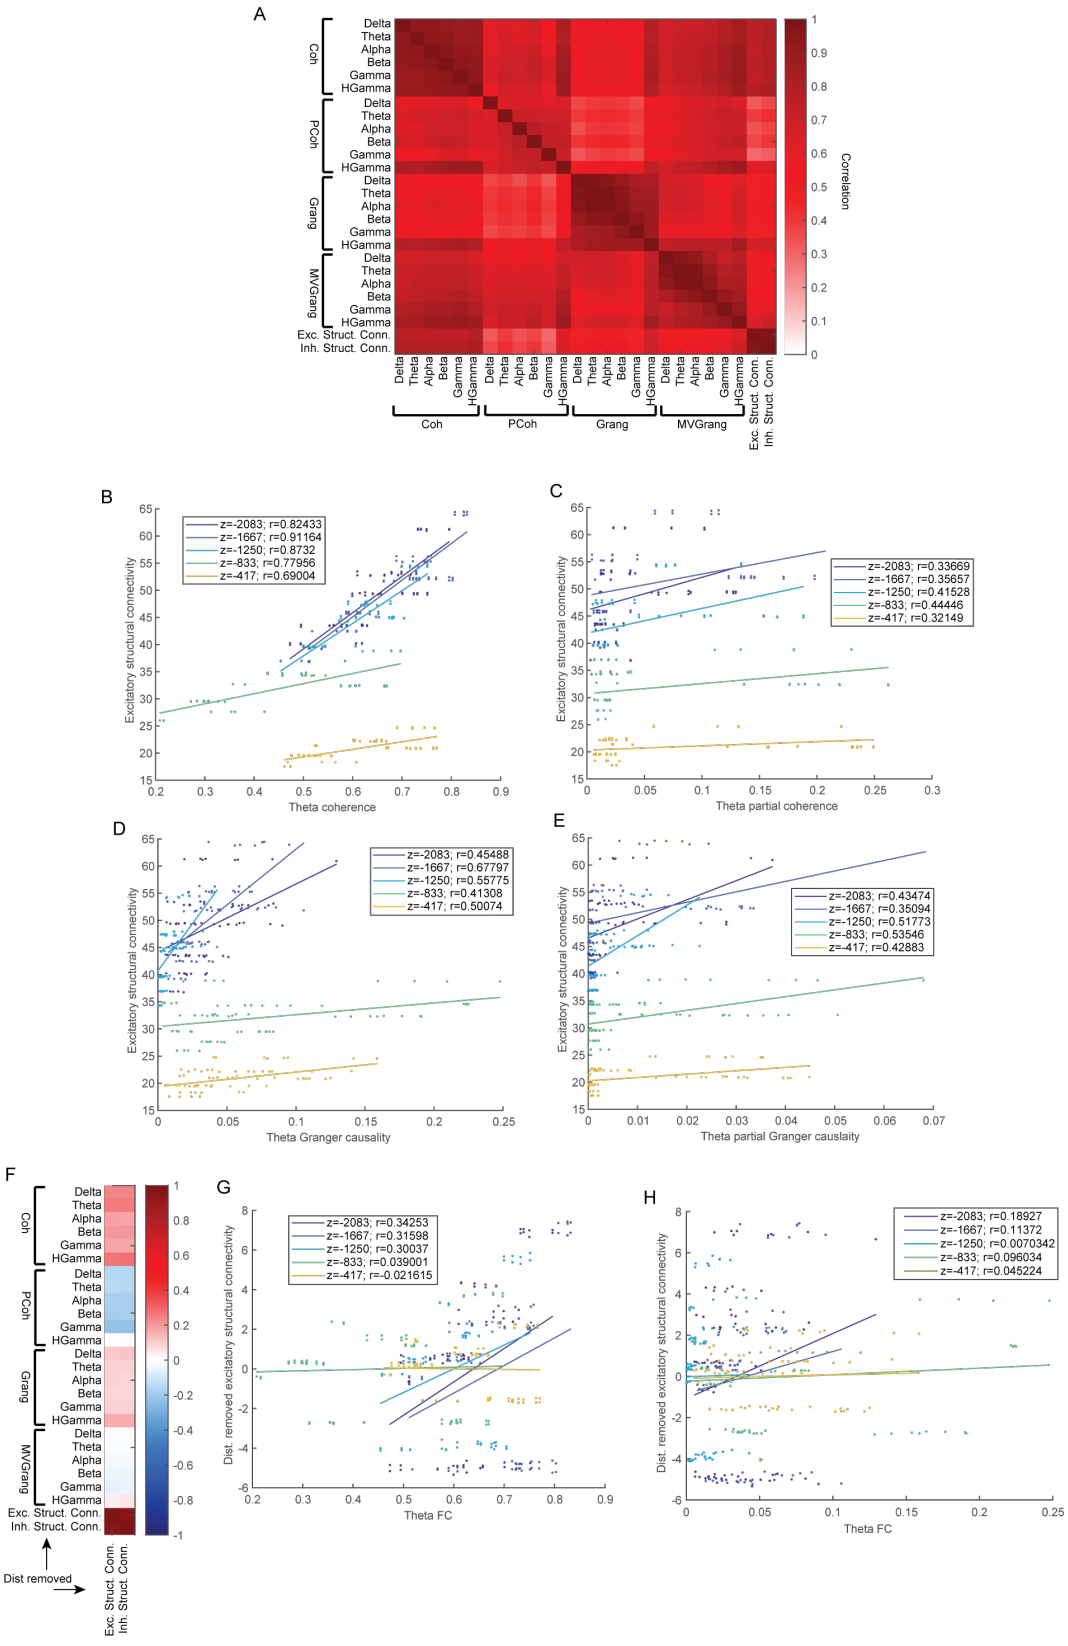

Figure 11: Correlations when electrodes are placed at consistent depths. A) heatmap of correlations calculated on a depth specific basis, averaged over all depths. B-E) Theta functional connectivity metric values vs. excitatory ground

truth on a depth-specific basis for B) coherence C) partial coherence D) Granger causality and E) multivariate Granger causality. Colors of lines and data points correspond to connection strengths between electrodes at specific depths beneath the surface of the simulated cortex, given in um in the legend, alongside the Pearson correlation values.. F) Heatmap of correlations on depth basis when layer-specific distance information is removed. G) Theta coherence vs distance removed excitatory structural connectivity and H) vs. distance removed inhibitory structural connectivity.

Discussion

This study aimed to illuminate the relationship between local field potential (LFP)-level functional connectivity and structural connectivity in the brain. Our approach involved a large-scale neural network simulation that generates LFPs based on a biophysically realistic model. We leveraged this model to evaluate a set of functional connectivity metrics and determine their ability to infer structural connectivity. Our results offer valuable insights for neuroscience research and neural engineering applications, underscoring the potential of coherence as a promising metric and the advantages of combining multiple functional connectivity metrics.

Coherence as a Key Metric for Inferring Structural Connectivity

Our findings underscored coherence, especially in the lower frequency bands, as a particularly promising metric for estimating structural connectivity from functional connectivity measurements. This is consistent with previous research that has highlighted the utility of coherence as a proxy for structural connectivity between brain regions. Moreover, the coherence metric demonstrated robust performance even when distance information was removed and when structural connectivity changes were induced by stimulation, further validating its effectiveness in capturing structural connectivity.

Coherence is the simplest measure tested in this work, as it is neither directed nor conditional. As such, it is somewhat surprising that it tended to perform best, at least for the purpose of inferring structural connectivity. However, as coherence is the simplest measure it also makes the fewest assumptions about the data, and as such has the least potential for errors (Bastos 2016). For example, while pairwise metrics will incorrectly identify a connection between two regions when their only connection is through another region, this capacity for bias is theoretically equivalent for all pairs of regions. Contrast this with conditional metrics, which seek to exclude indirect connections, but will only do so accurately if all indirect connections are observed, and thus have dependently biased errors as a function of observed regions. In other words, conditional metrics assume that all contributing signals are recorded and included for analysis, while coherence makes no such assumption. Similarly, directed measures assume that the dynamics of two signals can be simply (even just additively, in the case of Granger causality) split into influence from region A to region B, and region B to region A. Symmetric measures make no such assumptions, and thus do not suffer if the influences cannot be naively parsed in this way; this may be even more relevant in this study as we analyze LFPs which represent distributed neural activity, as opposed to spikes which are easier to separate between recording sites.

Combining Functional Connectivity Metrics: A Data-Driven Approach

We employed a stepwise regression model to combine multiple functional connectivity metrics in a data-driven manner, aiming to improve the predictive capability for structural connectivity. This approach yielded a model with a high degree of explained variance (around 0.7) for both excitatory and inhibitory activities on the held-out test set. This finding suggests that integrating multiple functional connectivity metrics can yield more accurate inferences of structural connectivity. Such an approach should be developed and validated in vivo and tested for generalizability, as the potential for overfitting, and thus

wildly incorrect inference, is higher with a model with multiple variables in comparison to a single variable. Future work could explore nonlinear transformations of these metrics or utilize models that can identify nonlinear relationships to further improve accuracy.

Nonlinear Relationships Between Functional Connectivity Metrics and Structural Connectivity

Our results indicated the presence of nonlinear relationships between some functional connectivity metrics, particularly those based on Granger causality, and structural connectivity. These nonlinear relationships were captured by the higher Spearman correlation coefficients compared to the Pearson correlation coefficients. This finding suggests that future studies should consider these nonlinear relationships when using Granger-based measures to infer structural connectivity, and that simple transformations of the data such as log transforms may be relevant before use of the functional connectivity values.

Challenges in Defining Connectivity

Establishing a clear definition for structural and functional connectivity can be challenging, as there is no universally agreed-upon measure to characterize the strength and nature of connections between neural elements. Different metrics and approaches may capture different aspects of connectivity, and as a result, the choice of metrics can have a significant impact on the observed relationships between structural and functional connectivity (Bastos & Schoffelen, 2016). Our study aimed to address this challenge by evaluating a diverse set of functional connectivity metrics, each representing a unique class of connectivity measures, to provide a comprehensive view of the relationship between functional and structural connectivity.

Impact of strong rhythms on functional connectivity relations

Our study also highlighted the impact of strong rhythms, such as gamma oscillations, on the estimation of functional connectivity relations. We modified the initial model to reduce the overly strong gamma frequency signal and generate a smoother power spectral density (PSD) that closely resembled the characteristic $1/freq^a$ curve observed in the brain. However, we also noted that as the Gamma power of our simulation was higher than the 1/f curve would predict, it coincided with a lower correlation between the functional connectivity metrics within the Gamma band and the structural connectivity. This finding suggests that investigating the PSD curve can be used to make pragmatic decisions about how to use functional connectivity metrics to infer structural connectivity, on an individual experiment basis.

Relevance of Findings for Future Studies

Our findings have important implications for future studies investigating the relationship between functional and structural connectivity. By identifying the most promising metrics, such as coherence, and demonstrating the value of combining metrics, our study provides a foundation for further refining the methods used to infer structural connectivity from functional connectivity data. Additionally, our results highlight the impact of experimental constraints, such as the placement of electrodes at constant depths, on the observed connectivity relationships. This information may guide researchers in designing more effective experimental paradigms and analysis techniques to better understand the complex interplay between functional and structural connectivity in the brain.

In addition, several new technologies are emerging as status quo for recording neural activity at with better resolution and higher channel count than established technologies. For example, Neuropixels are being commonly used now to record spikes and LFPs (Paulk 2022), from which functional connectivity can

be calculated. In this era of larger neural datasets, the validation of the use of simple functional connectivity metrics such as coherence is a pragmatic finding to guide analysis of these data.

Other Factors Contributing to Functional Connectivity

Research suggests that several factors can influence functional connectivity in the brain. It is well established that structural connectivity has a causal role in the generation of functional connectivity, through observational (Huang and Ding 2016, Straathof 2018) as well as interventional studies (O'Reilly 2013). For example, O'Reilly et al. (2013) investigated this relationship by examining interhemispheric functional connectivity in rhesus monkeys before and after sectioning the corpus callosum—a thick band of nerve fibers that connects the left and right cerebral hemispheres. Their findings revealed that while sectioning the forebrain commissure significantly reduced interhemispheric functional connectivity, this effect was considerably mitigated if the anterior commissure—a fiber bundle that connects the two hemispheres—was left intact. In other words, the specific structural connections between two brain regions are causally and strongly linked to the formation of functional connectivity.

However, there are also other factors that can influence functional connectivity, such as brain state (Fukushima et al., 2017; Chen et al., 2012; Straathof et al., 2018). For instance, Fukushima et al. (2017) reported a higher similarity between structural and functional connectivity during integrated network states, indicating that brain states may influence the degree to which structural connectivity predicts functional connectivity. Furthermore, the study linked the fluctuations between segregated and integrated network states with human behavior, shedding light on the potential behavioral implications of functional connectivity. Chen et al. (2012) also highlighted the dynamic nature of functional connectivity, showing that connectivity patterns derived from electroencephalography (EEG) data can change with different brain states. This finding underscores the importance of the contribution of brain states to the complex dynamics of functional connectivity. Future research should consider these factors to develop a more comprehensive understanding of functional connectivity in the brain.

Assumptions and Limitations of the Modeling Approach

Our study employed a biophysically plausible neural tissue model with a focus on a spatial scale of a few millimeters, which is relevant to research employing Utah arrays and similar mammalian brain implants. While this approach offers the advantage of high accuracy in neuronal signal measurements and the ability to closely match neuron statistics observed in the brain, it also comes with certain assumptions and limitations.

One, it is unclear if observations made at this spatial scale extend to larger scales – that is, if the same relationships observed between functional and structural connectivity at this small scale extends to larger scales such as between different brain areas. Despite this, there is a long history of new functional connectivity metrics being developed and despite only being validated on toy models of just a few neurons (Baccala 2001, Eichler 2006), drastically smaller than what we use here, deployed for use in larger scales in real brains such as for EEG data.

Second, our results are generated in a simulation, and while the simulation is made to be as biorealistic as possible, it is just an approximation of a real brain. As such, while our analyses and findings are useful, and we have pragmatically chosen the use of a simulation in order to conduct analyses not possible in vivo, they should be validated in animal models before any claims of generalizability of our findings.

Despite these limitations, our study provides valuable insights into the relationship between functional and structural connectivity and offers a foundation for further research in this area. Future studies should consider exploring alternative models, incorporating additional cell types and connectivity patterns, and examining the relationship between functional and structural connectivity across different spatial scales and recording modalities.

Our study revealed that coherence is a promising metric for inferring structural connectivity from functional connectivity data, and that combining functional connectivity metrics using a stepwise regression approach can further improve the accuracy of these inferences. Our findings also highlighted the impact of various factors, such as the choice of metrics, the presence of strong rhythms, and the spatial scale of analysis, on the observed relationships between functional and structural connectivity. These insights have important implications for the design and interpretation of future studies investigating the complex interplay between functional and structural connectivity in the brain.

Methods

Simulation of neural activity

For our simulation we use a model from (Tomsett 2015) which we have tweaked. The model is composed of hundreds of thousands of virtual neurons, with scale and number corresponding to a brain section of size 1.5mm x 1.5mm x 2.6mm. The neurons are modeled as adaptive exponential integrate and fire neurons (Brette 2005) connected to each other via conductance-based synapses and with background currents given by Ornstein-Uhlenbeck processes. The model was tweaked by varying the neuron group connectivity characteristics with each other and by varying the background current mean and standard deviation to neurons, such that the power spectral density exhibited maximally close to $1/freq^a$ power and the neuron firing rates were as similar to in vivo experimental values as possible, as detailed below. The model was run for 40 seconds of virtual neural activity. Both excitatory and inhibitory neurons are in the simulation. Each virtual neuron is associated to a specific spatial X Y and Z location.

Neurons are divided into 16 different cell types, in accordance with (Binzegger 2004). The laminar statistics of neuron spatial locations, as well as their connectivity statistics to each other and their arborization radii are also taken from structural connectivity information (Binzegger 2004).

Simulation LFP generation

Within the spatial boundaries of the simulation, we define 45 points corresponding to the locations of virtual electrodes. For each neuron in the simulation, a reduced compartmental spatial model (Tomsett 2015) generates resulting LFP at each recording site. The voltages of each neuron are superimposed to yield LFPs at each recording electrode at each time point, as described in (Tomsett 2015).

The electrodes were placed in 5 z-slice planes as 3x3 electrodes regular distance from each other. They were located at z slices 500um intervals below the surface of the simulated cortex, with x and y locations of 250, 500, 750, 1000, and 1250um. We used these locations to avoid spurious dynamics and LFP at the boundaries of the simulation volume.

Computational tools

All analysis was conducted in Matlab. The simulation and calculation of the structural connectivity was completed on the Neuroscience Gateway Supercomputer (Sivagnanam 2013), and all analyses were conducted on PCs.

Model validation

We validated the model by ensuring its adherence to certain empirical and theoretical standards established within neurophysiology. First, we confirmed that the firing rate histograms of neurons within a population exhibited normal or log-normal distributions. This was grounded in the pioneering work of van Vreeswijk and Sompolinsky in 1996, which highlighted irregular firing patterns in cortical neurons, evinced by their Poisson-like histograms of interspike intervals. This pattern of irregularity, characteristic of neuronal firing, suggests the presence of an underlying log-normal distribution. Further empirical substantiation comes from Antolik et al. 2018, and Kajiwara et al. 2020.

The second point of validation involved ensuring that interspike intervals also showed a log-normal distribution. This was supported by the successful modelling of neuronal interspike interval durations using a log-normal distribution by Matthews and Clements in 2014.

Finally, the model was evaluated based on whether the power spectral density at the recording electrodes demonstrated the canonical 1/frequency^α decrease, reflecting the model's capability to simulate natural neuronal dynamics accurately.

LFP intrinsic functional connectivity metric calculation

Recorded time series of LFP from each virtual electrode were extracted over **t**he 40 second run from 241 to 39990 milliseconds.

Coherence was computed between pairs of electrodes in the local field potential (LFP) data. For each pair of electrodes, the coherence function mscohere from MATLAB was used. This function calculates the magnitude-squared coherence estimate, which is a measure of the linear correlation between the two signals in the frequency domain. For each pair of electrode signals, a hamming window of 1000 ms was applied. This signal was then zero padded to length 1024 via the nfft variable. The overlap between successive windows was set to half the window size.

Partial Spectral Coherence (PSC) was calculated for pairs of electrodes in the local field potential (LFP) data. The windowing parameters were identical to that of coherence. In order to calculate PSC, the Cross Power Spectral Density (CPSD) was computed for all pairs of columns in the LFP data using a Hamming window with 50% overlap. The matrix of cross power spectral densities (CPSDs) for each frequency point was then inverted. This inversion was performed using Singular Value Decomposition (SVD) for numerical stability. The SVD of each frequency matrix was computed, and then the inverse was calculated by reciprocating the singular values and multiplying the resulting diagonal matrix by the transpose of the right singular vectors and the left singular vectors. Finally, the PSC was computed as the square of the absolute value of the normalized inverted matrix of CPSDs. This normalization was done by dividing each element of the inverted CPSD matrix by the square root of the product of the corresponding diagonal elements.

Granger causality and Multivariate Granger causality were calculated with the MVGC Toolbox (Barnett 2014). A 16 order vector autoregressive (VAR) model was fit to the time series data using Ordinary Least Squares Regression. The autocovariance was calculated from the VAR model, and then the spectral granger causality was calculated from the autocovariances. For pairwise Granger causality this process took place individually for each pair of electrodes, while for Multivariate Granger causality this process happened over the full LFP time series matrix of all electrodes.

Structural connectivity calculation

At model initialization, connections between neurons are created according to structurally-derived probability distributions. These connections are stored in a weighted adjacency matrix, where the weights correspond to the conductance of a synapse connecting neurons i and j.

First, we separate the matrix into 2 separate matrices one which holds the synapses from excitatory presynaptic neurons, and one with synapses from inhibitory pre-synaptic neurons. Then for each matrix, we calculate a summary-level "structural connectivity structural" between individual pairs of virtual recording electrodes. We do so by weighting these the connectivity weights in the adjacency matrix by the distances from each neuron to the recording electrode, as in the equation and in the figure below.

$$connectivity(A,B) = \sum_{i=1}^{n}\sum_{j=i}^{n}\frac{1}{D(A,i)}\frac{1}{D(B,j)}w(i,j)$$

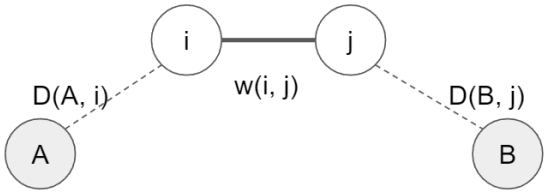

Figure 12: Depiction of the variables used in structural connectivity calculation.

This yields two measures of structural connectivity for each electrode pair, an excitatory structural connectivity and an inhibitory structural connectivity.

Distance-based Analysis

To determine the distances between electrodes, we calculated the Euclidean distance in the x, y, and z coordinates of each electrode. The distance values were based on raw untransformed distance calculations. An intercept term was fitted to the data to obtain the structural connectivity information that was not present within distance. Subsequently, other features were regressed against these values.

Partial Regression

For the partial regression results, a linear model was utilized to predict structural connectivity as a function of distance. The predicted value was then subtracted from the actual structural results for each data point to obtain the structural connectivity information that was not present within distance.

Stepwise Regression

To generate a model that integrated multiple functional connectivity metrics to estimate structural connectivity, we employed stepwise regression. First, the data was divided into a training and test set, with the training set comprising 70% of the data and the test set comprising the remaining 30%. Then the default MATLAB implementation of stepwise regression was used, with a feature selection criterion for maximizing the Bayesian information criterion (BIC). The process consisted of taking the model implementation at iteration i (where the model at iteration 0 consists of only an intercept term), testing all possible iteration i+1 models by testing every possible combination of iteration i model plus or minus one feature in the feature set, and greedily setting the model with the maximum BIC as the iteration i+1 model. This was repeated iteratively until the BIC would no longer be improved by subtracting or removing a feature. The final model was fit to the training set, and the R squared was calculated using the holdout

data of the test set. This procedure was carried out for both excitatory and inhibitory structural connectivity, as well as for distance-removed structural connectivity.

Chapter 6: Computational model of cortical light penetration for a novel photothrombosis-based lesion model in non-human primates

Full version published as: K. Khateeb, **J. Bloch**, J. Zhou, M. Rahimi, D. Griggs, V. Kharazia, M. Le, R. Wang, A. Yazdan-Shahmorad, "A Toolbox for Studying Cortical Physiology in Primates," *Cell Reports Methods*, Mar. 2022.

Abstract

Lesioning and neurophysiological studies have facilitated the elucidation of cortical functions and mechanisms of functional recovery following injury. Clinical translation of such studies is contingent on their employment in non-human primates (NHPs), yet tools for monitoring and modulating cortical physiology are incompatible with conventional lesioning techniques. To address these challenges, we developed a toolbox validated in seven macaques. Our lab introduces the photothrombotic method for inducing focal cortical lesions, a quantitative model for designing experiment-specific lesion profiles, and optical coherence tomography angiography (OCTA) for large-scale (~5 cm2) monitoring of vascular dynamics. We integrate these tools with our electrocorticographic array for large-scale monitoring of neural dynamics and testing stimulation-based interventions. Advantageously, this versatile toolbox can be incorporated into established chronic cranial windows. By combining optical and electrophysiological techniques in NHP cortex, we can enhance our understanding of cortical functions, investigate functional recovery mechanisms, integrate physiological and behavioral findings, and develop neurorehabilitative treatments.

As my primary role in the project was the development and validation of the computational model, I have selected the most relevant passages from the published work which highlight my contributions while still communicating the motivation and context of the work.

Introduction

The primate neocortex is responsible for a variety of complex tasks and behaviors, including long-term memory storage, sensory processing, and movement. Historically, both lesioning and neurophysiological studies have been critical for elucidating functions of specific cortical regions (Ferrier, 1876) such as somatosensory (Borich et al., 2015; Brinkman et al., 1985; Gerlai et al., 2000), visual (Humphrey, 1974; Wurtz and Goldberg, 1972), auditory (Heffner and Heffner, 1986), and posterior parietal (Murphy et al., 2016; Vallar et al., 1994) cortices. Recently, such strategies have also been employed to investigate mechanisms of plasticity and recovery following injury (Friel et al., 2007; Guo et al., 2021; Harrison et al., 2013; Kaeser et al., 2010; Khanna et al., 2021; Liu and Rouiller, 1999; Nudo and Milliken, 1996; Nudo et al., 1996; Padberg et al., 2010; Pons et al., 1988; Xerri et al., 1998). The study of these phenomena in non-human primate (NHP) models with strong evolutional and physiological relevance to the human cortex is critical for understanding fundamentals of cortical physiology and designing novel clinical treatments for cortical injury. This was particularly highlighted following the lack of success of clinical trials for novel stroke therapies in recent decades. One major contributing factor to this lack of success has been attributed to a heavy reliance on rodent studies and a lack of pre-clinical studies in NHPs (Kotak et al., 2005; Levy et al., 2016). Thus, there is a strong need to expand the tools available for studying NHP cortical physiology for clinical translation.

Typical strategies for investigating in vivo NHP cortical physiology include either monitoring or perturbing the native cortical activity then correlating neural activity with behavior. Neural activity is frequently monitored through electrical recording, or calcium imaging, while perturbations include lesioning, electrical stimulation, or optogenetic manipulation (Acker et al., 2016; Tremblay et al., 2020). Previously, we developed a large-scale interface enabling optogenetic neuromodulation in concert with simultaneous

electrical recording (Ledochowitsch et al., 2015a; Yazdan-Shahmorad et al., 2015, 2016, 2018c, 2018a, 2018b). However, to the best of our knowledge, there is no single unifying paradigm through which the full spectrum of strategies can be combined for the unhindered investigation of cortical physiology in NHPs. Conventional lesioning techniques lack compatibility with tools for monitoring and modulating cortical physiology and lack flexibility in controlling lesion location and extent. Here, we adapt and integrate in NHPs a versatile focal ischemic lesioning technique with large-scale monitoring of cortical vascular dynamics, and electrophysiological recording and stimulation. The mutual compatibility of the integrated tools within this versatile toolbox enables for the inclusion or exclusion of any of the tools depending on experimental needs. Importantly, with large-scale stable optical access, this toolbox can be combined with established chronic cranial windows for optical stimulation and imaging techniques such as optogenetics and calcium imaging, respectively.

Commonly utilized NHP cortical lesioning techniques are challenging to employ. One method of lesioning cortex in NHPs is through middle cerebral artery occlusion (Maeda et al., 2005; Virley et al., 2004). Because middle cerebral artery occlusion is used to mimic ischemic stroke as observed in the clinic, the resulting lesions are broad and are constrained only to regions downstream of the middle cerebral artery. Moreover, obtaining access to the middle cerebral artery for occlusion requires complex surgical intervention regardless of the occlusion technique. Similarly, common focal cortical lesioning techniques such as endothelin-1 (Dai et al., 2017; Teo and Bourne, 2014), electrocoagulation (Nudo et al., 2003; Xerri et al., 1998), ibotenic acid (Kaeser et al., 2010; Liu and Rouiller, 1999), cooling (Brinkman et al., 1985), and aspiration (Heffner and Heffner, 1986; Padberg et al., 2010; Pons et al., 1988) also involve technically challenging surgical procedures that are susceptible to variability across animals. Importantly, these techniques lack compatibility with tools for monitoring cortical physiological dynamics during lesion formation and recovery.

In this study, we addressed the technical shortcomings of studying cortical physiology in NHPs. We successfully demonstrate the combination of a photochemical lesion induction technique with a controlled spatial profile, a computational model designed to predict the induced lesion sizes and shapes, an in vivo lesion validation method with high spatial resolution, and the ability to simultaneously monitor the underlying neural activity and blood flow as lesions form at a large scale (~5 cm2). Moreover, the sizes of focal lesions induced with our toolbox are comparable to previously reported focal lesions capable of eliciting behavioral deficits (Murata et al., 2008; Padberg et al., 2010). Importantly, the tools presented here are compatible with previously established interfaces offering large-scale access for optical and electrical stimulation as well as imaging (Griggs et al., 2021b; LM et al., 2002; Macknik et al., 2019; Trautmann et al., 2021; Yazdan-Shahmorad et al., 2016), providing unparalleled access to the primate cortex. Although not demonstrated here, the tools described in this study can easily be adapted for use in smaller model organisms, expanding its utility beyond research groups with access to macaques. Through this integrative approach, the large-scale monitoring of cortical network activity and reorganization concurrently with vascular dynamics expands the questions and interventions we can explore, profoundly impacting the development of neurorehabilitative therapies.

Relevant results

Prediction of Lesion Size by Simulation of Light Propagation through Cortical Tissue

We developed a computational model to predict the lesion sizes and shapes induced by illumination parameters tested here and beyond. We incorporated both light intensity and aperture diameter into a biophysically inspired model to jointly predict both lesion diameter and depth. Photothrombotic lesions arise due to thrombi formation and subsequent cell death mediated by the topography of the underlying microvasculature. As such, we developed a two-stage modeling process in which we simulated photons penetrating brain tissue (Figure 2A) to generate a profile of the spatial fluence distribution (Figure 2B) and

transformed fluence contours to recreate the lesion shapes observed through histology (Figure 2C,D, see Methods for details). Using this method, we identified the best optical properties as a gray matter absorption coefficient of 0.395 mm^{-1}, gray matter scattering coefficient of 53.6 mm^{-1}, white matter absorption coefficient of 0.09 mm^{-1}, white matter scattering coefficient of 54.066 mm^{-1}, and a spatial fluence distribution threshold for lesion induction of 19.9 $\mu W/mm^2$.

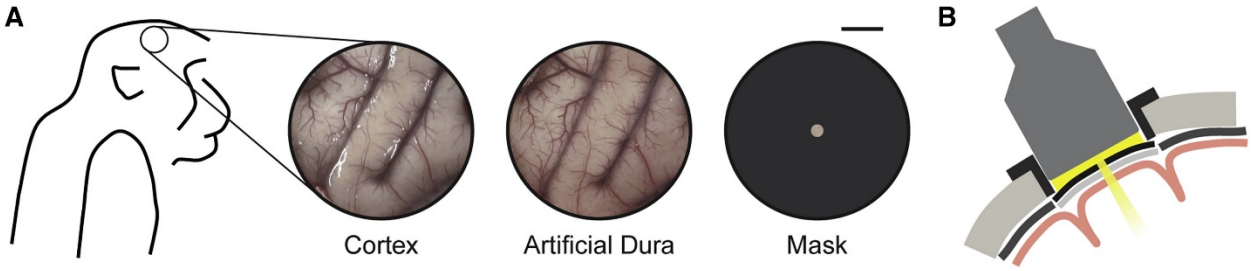

Fig. 1. Figure 1. Schematic of photothrombotic technique application to induce focal cortical ischemic lesions. (A) Following a 25-mm-diameter circular craniotomy, a thin, transparent artificial dura was placed over the exposed cortical surface. Next, an apertured mask was placed on top of the artificial dura. Scale bar is 5 mm. (B) Coronal schematic of light illumination through the apertured mask following intravenous Rose Bengal infusion. Portions of this figure have been adapted for inclusion in this manuscript from Khateeb et al. (2019b) with permission.

From identified best-matching light simulations, we quantified the degree to which they predict lesion shapes and sizes. We observed that while the simulation results qualitatively align with the lesion profiles, the light simulation alone was not adequate for explaining the observed lesions. By using the maximal depth and average diameter of individual fluence threshold contours of the simulation as predictions of their corresponding lesion profiles, we observed that the contours resulted in an r-squared value of -0.04 for lesion depth and 0.32 for diameter. These results indicate that the extent of induced photothrombotic ischemic lesions *in vivo* is not simply a function of the fluence profile in the cortex but is likely also governed by biological factors such as vascular topography.

We therefore incorporated a scaling process to transform the light intensity contours to the biological lesion profiles observed through histology (Figure 2C,D). The depth did not scale linearly well, however, as the best linear scaling of depth yielded an r-squared value of 0.13. From examining the residuals, we observed that square root transforming the depth before linear scaling would yield a more accurate prediction. Indeed, by square root transforming the depth and scaling by a factor of 1.6969 we obtained an r-squared value of 0.41 (Figure 2E). We observed that the diameters from the simulation scaled linearly to match the diameters observed from histology, with a scaling factor of 0.726 yielding an r-squared value of 0.82 (Figure 2F). Interestingly, although our model was developed based on histologically derived lesion contours, the resulting model-derived lesion diameters were also highly predictive of OCTA-measured lesion diameters with an r-squared value of 0.60 (Figure 2G). Thus, our quantitative model can accurately predict lesion shape, diameter, and depth.

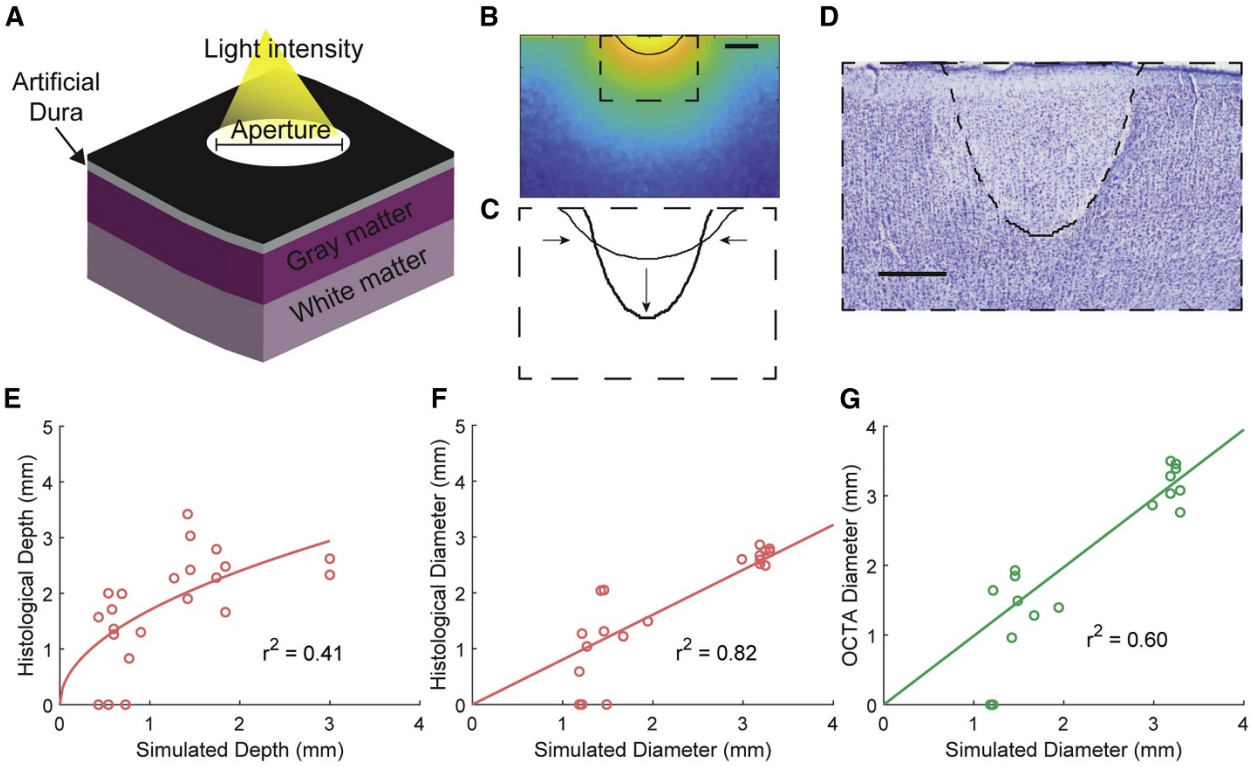

Fig. 2: Prediction of lesion size by simulation of light propagation through cortical tissue. (A) Schematic of simulated cortical volume with our experimental setup. An uncollimated light beam passes through an aperture and a transparent artificial dura (0.5 mm thick) into gray and white matter of a virtual cortical medium. Gray-matter thickness is 2.5 mm. (B) A contour is identified from the light profile matching the light intensity threshold (19.9 µW/mm2) most closely matching the lesions. Scale bar is 100 µm. (C) The light-intensity contour is scaled to generate a biological lesion contour. Scaling factors were obtained through regression on our dataset of simulated lesion dimensions and corresponding histologically measured lesions. (D) Predicted lesion contour overlayed on a coronal Nissl-stained slice of a corresponding lesion from monkey B. Scale bar is 50 µm. (E–G) Simulated lesions accurately predict histologically measured lesion depths (0.41 r-squared) (E), diameters (0.82 r-squared) (F), and OCTA-measured diameters (0.60 r-squared) (G).

Relevant Discussion

The presented toolbox offers significant advantages for investigating critical questions regarding cortical physiology. Here, we demonstrate the photothrombotic technique for inducing focal ischemic lesions in NHP cortex. The extent of subsequent neuronal cell loss was later confirmed histologically. Moreover, we developed a quantitative model for predicting photothrombotic lesion sizes based on illumination parameters. Through this model, lesions can be designed in accordance with experimental needs. We also present OCTA imaging as a tool for observing vascular dynamics and validating lesion induction *in vivo*. Additionally, our ECoG array allows for the investigation of neural dynamics before, during, and after lesion formation. Critical for developing therapeutic interventions, we also establish the ability to test stimulation-based strategies with our ECoG array. Although not demonstrated here, the toolbox allows for experimentation at clinically relevant time scales when employed in NHPs in contrast to rodent studies.

The photothrombotic technique allows for the induction of targeted focal ischemic lesions in any cortical region of interest. The spatial control of photothrombotic lesion induction is an improvement over other commonly employed models such as middle cerebral artery occlusion models (Sommer, 2017). In addition, the ability to induce consistent smaller, targeted lesions can allow for the induction of lesions to

impact specific cortical functions without risking serious injury or mortality compared to middle cerebral artery occlusion (Wu et al., 2016). This capability has been validated in rodents as the photothrombotic technique has enabled the study of local cortical remapping in sensorimotor cortex (Harrison et al., 2013). Larger lesions can also be induced with areas as large as the applied cranial window. Moreover, middle cerebral artery occlusion (Maeda et al., 2005), electrocoagulation (Friel et al., 2007; Nudo et al., 1996; Xerri et al., 1998), and aspiration methods (Padberg et al., 2010) of cortical lesioning require complex surgical intervention and are performed in an invasive manner. In contrast, photothrombotic lesioning only requires optical access to cortex. With the aid of our predictive computational model, lesions can be designed based on experimental needs. However, as is typical with optical techniques, photothrombotic lesion induction with minimal invasiveness is limited to optically accessible regions. To access subcortical areas, more invasive approaches can be used, such as the use of penetrating optical fibers.

We also developed a computational model to predict the shape and scale of photothrombotic lesions. Our regression results and Monte Carlo simulation-based results together allow for planning of precise lesions with a variety of illumination parameters, informed by previous literature and validated against our experimental results. Importantly, through our grid search, we identified the optimal optical properties of cortical tissue to accurately predict lesion sizes. It is important to note, in our experiments we did not vary illumination time across different lesions, with an illumination time based on previous photothrombotic studies (Gulati et al., 2015). Therefore, our computational model was constructed based on lesions induced with consistent illumination times. However, our model may be modified in future studies to predict lesions induced with different illumination times to account for differences in total energy that would be delivered to induce Rose Bengal photoactivation. There are also notable assumptions which are applied to our model. First, we assume a consistent gray matter depth at 2.5 mm, whereas in reality, the depth of gray matter is variable across cerebral cortex, which can affect lesion depth predictions. Moreover, the final step of our model involves scaling the profiles from probability maps to better reconstruct the profiles of our histologically detected lesions. This was done to account for the topography of the cortical vasculature (Figure 5B-D). The depth-wise upscaling of our model contours is likely due to the perpendicular orientation of penetrating arterioles and venules relative to the cortical surface (Gould et al., 2017). The diameter-wise downscaling can be ascribed to the redistribution of blood flow in vessels downstream of an occluded vessel to preserve perfusion in these downstream regions (Schaffer et al., 2006), thus potentially reducing the extent of ischemic damage.

Together, these tools can be used to answer critical questions in neuroscience and drive studies of cortical physiology and related disorders such as stroke and traumatic brain injury. Importantly, our toolbox enables the perturbation of the brain and simultaneous monitoring of physiological dynamics over time. The combination of the presented tools can be used to evaluate vascular dynamics, neural dynamics, and neurovascular coupling in the context of either cortical lesioning, stimulation, or both. Like any toolbox, the mutual compatibility of the tools allows for the exclusion or inclusion of any of the tools described here, depending on the research questions. One main advantage of this toolbox lies in the adaptability and mutual compatibility of this family of tools with each other and other conventional tools in the field. Although our toolbox is demonstrated here in an acute setting, the same tools are capable of implementation in chronic experiments. Previous studies have demonstrated the long-term safety and efficacy of optical windows utilizing an artificial dura for optogenetics, calcium imaging, and other optical tools across various cortical regions (Bollimunta et al., 2021; Ju et al., 2018; LM et al., 2002; Trautmann et al., 2021; Yazdan-Shahmorad et al., 2016, 2018a). The increasingly widespread use of these optical interfaces make the presented tools easily integrable with established cranial windows. Such a capability would allow for the monitoring of the behavioral and physiological effects of lesioning and subsequent recovery over time, in combination with testing stimulation-based interventions. Furthermore, the compatibility of our toolbox with optical tools such as optogenetics can enable the comparison of

temporary and permanent cortical lesioning techniques. The tools described here can also be adapted for use in smaller organisms such as marmosets and rodents, expanding the utility of this toolbox beyond research groups with access to large NHP animal models. Importantly, long-term studies employing our toolbox can address key questions regarding cortical functions and drive the development of future rehabilitative therapies for stroke, traumatic brain injury, and other relevant neurological disorders.

Relevant Methods

Light Simulation

Our computational model of light penetration and subsequent lesion in the brain consisted of a Monte Carlo simulation of photons propagating through brain tissue, extraction of a light intensity contour, and scaling of the contour. A virtual volume was constructed which mimicked the experimental setup and had the relevant optical properties of brain tissue. Photons were virtually propagated through this volume to obtain a probability distribution of photon fluence, or light energy passing through a given area. Optical parameters were informed by published literature regarding primate cortical optical properties and then refined to maximally match our experimental results obtained through histology. The Monte Carlo simulations were designed using the Monte Carlo eXtreme (MCX) software (Fang and Boas, 2009) and run on discrete graphics cards (Nvidia GTX 1080Ti and Nvidia GTX 2060).

The virtual volume consisted of a light source and mask, an artificial dura, and gray and white matter (Figure 5A). The entire volume was 8 mm x 8 mm x 5.6 mm (width x width x height), with voxel resolution of 0.01 mm^3. We specified the light source as a disk parallel to the surface of the brain with collimated light exiting the side of the disk facing the brain, and then passing through a thin highly scattering disk which uncollimated the light. We set the diameters of both disks to be equal to the aperture diameter of the mask of the corresponding *in vivo* experiment. We specified the volume peripheral to the disks as a highly absorbing and highly scattering medium to mimic the opacity of the non-aperture part of the mask (Supplementary Table 1). The diameter used in our simulations were 0.5 mm, 1 mm, 1.5 mm, and 2 mm, corresponding with the apertures tested in our *in vivo* experiments. After specifying the volumes, we ran simulations of 1 million photons through the volume. The simulations yielded voxel-wise fluence values, or values of radiant energy received per unit area. The fluence results were imported into MATLAB and analyzed with custom code.

experimental sessions consisted of recording while no stimulation took place, which we termed a “resting-state block,” and recording during stimulation, which we termed a “stimulation block.” Each resting-state block was 5 minutes long and each stimulation block was 10 minutes long. Each experimental session consisted of alternating resting-state blocks and stimulation blocks, with 6 total resting-state blocks.

Research on the optical properties of gray matter is often contradictory, and values of index of refraction, anisotropy coefficient, and absorption and scattering coefficients vary by orders of magnitude in published research. We defined a grid search boundary of optical properties by using the values reported in previous literature (Gottschalk, 1992; Yaroslavsky et al., 2002) as upper and lower bounds of the properties. As the absorption spectrum for Rose Bengal peaks at 559 nm, we used the optical properties reported for this wavelength by interpolating from nearest reported wavelengths. The range of the grid was [0.03, 0.76] mm^{-1} for gray matter absorption, [9.9, 53.6] mm^{-1} for gray matter scattering, [0.09, 0.36] mm^{-1} for white matter absorption, and [41.9, 78.4] mm^{-1} for white matter scattering. For anisotropy coefficients the two publications reported similar values, therefore we averaged the values between the publications to obtain an anisotropy coefficient of 0.92 for gray matter and 0.8 for white matter. The refractive index was set as 1.36 for gray matter and 1.38 for white matter (Yaroslavsky et al., 2002). From the optical property ranges we constructed a 5 x 5 x 4 x 4 size grid with 5 values for the gray matter absorption and scattering coefficients and 4 values for the white matter absorption and scattering

coefficients. For each of the 400 combinations of parameters, four simulations were run, corresponding with the four aperture sizes tested in our experiments. Each simulation yielded a volumetric fluence distribution. We further obtained representative 2D central slices to compare with histologically obtained lesions.

For each photon intensity distribution corresponding to an optical property combination, we identified the light intensity threshold which yielded contours most closely matching those obtained from central slices of lesions obtained from histology. By treating the maximum depth and average width of an individual contour as the prediction of maximum depth and average width of the corresponding lesion, we quantified the predictive power of the simulation. Denote by SSR_d the residual sum of squares of prediction of maximum depth of a lesion, SST_d the total sum of squares of maximum depth of a lesion, SSR_w the residual sum of squares of prediction of average width of a lesion, and SST_w the total sum of squares of average width of a lesion. The objective function used to find the best light intensity threshold was:

$$\min_{\text{thresh}} \left(\frac{{SSR_d}^2}{{SST_d}^2} + \frac{{SSR_w}^2}{{SST_w}^2}\right). \qquad (1)$$

This is equivalent to maximizing the (squared or unsquared) Euclidean norm of the r-squared values of depth and width predictions less 1. Denote by r_d^2 the r-squared value of the simulation-based prediction of maximum depth of a lesion, and r_w^2 the r-squared value of the simulation-based prediction of average width of a lesion. Then the equation above can be rewritten as:

$$\max_{\text{thresh}} \left((r_d^2 - 1)^2 + (r_w^2 - 1)^2\right). \qquad (2)$$

We solved this objective function individually for each optical property combination using Bayesian optimization. We then identified the simulation yielding the minimum value of this objective function as that most closely matching our histology data and used the results of that simulation for all later parts of the analysis. The optical properties of this best simulation were gray matter absorption coefficient of 0.395 mm^{-1}, gray matter scattering coefficient of 53.6 mm^{-1}, white matter absorption coefficient of 0.09 mm^{-1}, white matter scattering coefficient of 54.066 mm^{-1}, and a light intensity threshold for lesion induction of 19.9 $\mu W/mm^2$.

We then identified a transformation to accurately scale the profiles of the light simulation contours to that of the lesions obtained from histology (Figure 5C,D). We opted to do this by independently scaling the width and depth of the contours. Our first approach was to fit a univariate linear regression from the simulation depths and widths to the lesion depths and widths, respectively. In this linear regression we did not allow an intercept term. Based on the results of our linear regressions, we further evaluated the need for non-linear transformations. Once these scaling factors were set, we treated them as the final stage of our modeling platform, whereby light intensity contours from the Monte Carlo would be stretched and shrunk accordingly to transform from the light profile to the lesion profile.

Chapter 7: Conclusion

This thesis set out broaden the toolkit of neurostimulation technology, focusing on making the transition from preclinical models to human applications more seamless. At the start of this journey, four critical needs were identified: a neurological model that mirrors humans, technical tools for conducting precise experiments, advanced computational capabilities for optimizing experiments and analyzing data, and a comprehensive and grounded understanding of the outcomes we measure.

The use of NHPs throughout this work and the stroke model presented in Chapter 6 addressed the first need. By improving the capabilities surrounding NHP neurostimulation research we helped establish NHPs as an effective model for human neurological studies. Additionally, by developing an easily replicable and focal stroke model we provided a valuable framework for studying stroke—a prevalent neurological disorder. The stroke model offered a direct method for studying cortical functions and recovery mechanisms after injury, thus expanding our scope for developing neurorehabilitation treatments.

Chapter 4 addressed the second need by introducing a novel optogenetic-electric neural interface. This interface combined the specificity of optogenetics with the dual functionality of electrical stimulation and recording, showcasing its potential in future neurostimulation research and interventions. The experiment also validated the optogenetic perturbation by showing disruption of neural activity and behavior, two fundamental targets of neural stimulation.

For the third need, Chapters 2 and 3 significantly improved our computational capabilities. These chapters offered insights into how various parameters influence cortical coherence modulation and developed a model predicting network-wide functional reorganization, setting a course for more efficient experiment optimization and more accurate data interpretation.

Chapter 5 focused on the fourth need: a fundamental understanding of the outcomes we measure. By developing a large-scale, biophysically plausible neural network simulation that generates LFPs, we deepened our understanding of functional connectivity metrics and their relationship to structural connectivity.

In conclusion, this work has made strides towards addressing the four identified needs at the onset of this thesis. By advancing the NHP model, creating new technical tools, enhancing our computational capabilities, and deepening our understanding of the outcomes we measure, we move closer to the effective translation of neurotechnologies for human use.

References

Acker, L., Pino, E.N., Boyden, E.S., and Desimone, R. (2016). FEF inactivation with improved optogenetic methods. Proc. Natl. Acad. Sci. U. S. A. 113, E7297–E7306.

Adkins, D.L., Hsu, J.E., & Jones, T.A. (2008). Motor cortical stimulation promotes synaptic plasticity and behavioral improvements following sensorimotor cortex lesions. Exp Neurol, 212(1), 14-28. doi: 10.1016/j.expneurol.2008.01.031.

Aoki, F., Fetz, E.E., Shupe, L., Lettich, E., and Ojemann, G.A. Increased gamma-range activity in human sensorimotor cortex during performance of visuomotor tasks.

Arieli, A., Grinvald, A. and Slovin, H. (2002) 'Dural substitute for long-term imaging of cortical activity in behaving monkeys and its clinical implications', Journal of Neuroscience Methods, 114(2), pp. 119–133. doi: 10.1016/S0165-0270(01)00507-6.

Baccalá, L. & Sameshima, K. (2001). Partial directed coherence: a new concept in neural structure determination. Biol Cybern, 84, 463–474.

Barnett, S. C., Perry, B. A. L., Dalrymple-Alford, J. C. and Parr-Brownlie, L. C. (2018) 'Optogenetic stimulation: Understanding memory and treating deficits', Hippocampus, 28(7), pp. 457–470. doi: 10.1002/hipo.22960.

Barrat, A. et al. (2004). "The architecture of complex weighted networks". Proceedings of the National Academy of Sciences of the United States of America 101.11, pp. 3747–3752.

Bassett, Danielle S. and Olaf Sporns (2017). "Network neuroscience". Nature Neuroscience 20.3, pp. 353–364. issn: 15461726. doi: 10.1038/nn.4502.

Bastos, A.M. & Schoffelen, J.-M. (2016). A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. Front. Syst. Neurosci., 9.

Belloir, T., Belloir, T., Vargo, S. M., Ahmed, Z., Griggs, D., Fisher, S., Brown, T., Chamanzar, M. and Yazdan-shamorad, A. (2022) 'iScience Large-scale multimodal surface neural interfaces for non-human primates', ISCIENCE.

Beliveau, V., Ganz, M., Feng, L., Ozenne, B., Højgaard, L., Fisher, P.M., & Greve, D.N. (2017). A high-resolution in vivo atlas of the human brain's serotonin system. Journal of Neuroscience, 37(1), 120-128.

Benson, Austin R., David F. Gleich, and Jure Leskovec (2016). "Higher-order organization of complex networks". Science 353.6295, pp. 163–166. issn: 00107514. doi: 10.1080/00107510412331283531. arXiv: 0405123.

Bi, Guo-qiang and Mu-ming Poo (1998). "Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type". J. Neurosci 18.24, pp. 1–9. issn: 0270-6474. doi: 10.1038/25665.

Binzegger, T., Douglas, R. J., & Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. Journal of Neuroscience, 24(39), 8441–8453. doi:10.1523/JNEUROSCI.1400-04.2004

Bliss, T.V.P. and Collingridge, G.L. (1993). "A synaptic model of memory: LTP in the hippocampus". Nature 361, pp. 31–39.

Bloch, J. A., Khateeb, K., Silversmith, D. B., O'Doherty, J. E., Sabes, P. N. and Yazdan-Shahmorad, A. (2019) 'Cortical Stimulation Induces Network-Wide Coherence Change in Non-Human Primate Somatosensory Cortex∗',

Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS, 6446–6449. doi: 10.1109/EMBC.2019.8856633

Bloch, Julien A. et al. (2019). "Cortical Stimulation Induces Network-Wide Coherence Change in Non-Human Primate Somatosensory Cortex". Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS, pp. 6446–6449. issn: 1557170X. doi: 10.1109/EMBC.2019.8856633.

Bollimunta, A., Santacruz, S.R., Eaton, R.W., Moxon, K.A., Carmena, J.M., Correspondence, J.J.N., Xu, P.S., Morrison, J.H., and Nassi, J.J. (2021). Head-mounted microendoscopic calcium imaging in dorsal premotor cortex of behaving rhesus macaque. CellReports 35, 109239.

Borich, M.R., Brodie, S.M., Gray, W.A., Ionta, S., and Boyd, L.A. (2015). Understanding the role of the primary somatosensory cortex: Opportunities for rehabilitation. Neuropsychologia 79, 246–255.

Boyden, Edward S. et al. (2005). "Millisecond-timescale, genetically targeted optical control of neural activity". Nature Neuroscience 8.9, pp. 1263–1268. issn: 10976256. doi: 10.1038/nn1525.

Brainard, D. H. (1997) 'The Psychophysics Toolbox', Spatial Vision, 10(4), pp. 433–436.

Brinkman, J., Colebatch, J.G., Porter, R., and York, D.H. (1985). Responses of Precentral Cells During Cooling of Post-Central Cortex in Conscious Monkeys.

Brockwell, Peter J, Richard A Davis, and Stephen E Fienberg (1991). Time Series: Theory and Methods. Springer Science & Business Media.

Brette, R., & Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. Journal of Neurophysiology, 94(5), 3637–3642. doi:10.1152/jn.00686.2005

Buzsáki, G., & Anastassiou, C. A. (2011). Neurophysiology. Brain states and brain oscillations. Trends in cognitive sciences, 15(7), 329-336.

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science, 304(5679), 1926-1929.

Buzsáki, G., & Mizuseki, K. (2014). The log-dynamic brain: how skewed distributions affect network operations. Nature Reviews Neuroscience, 15, 264–278. doi:10.1038/nrn3687

Camporeze, B., Manica, B. A., Bonafé, G. A., Ferreira, J. J. C., Diniz, A. L., de Oliveira, C. T. P., Mathias Junior, L. R., de Aguiar, P. H. P. and Ortega, M. M. (2018) 'Optogenetics: the new molecular approach to control functions of neural cells in epilepsy, depression and tumors of the central nervous system.', American journal of cancer research, 8(10), pp. 1900–1918.

Carter, M. E. and de Lecea, L. (2011) 'Optogenetic investigation of neural circuits in vivo', Trends in Molecular Medicine, 17(4), pp. 197–206. doi: 10.1016/j.molmed.2010.12.005.

Chen, B. T., Yau, H. J., Hatch, C., Kusumoto-Yoshida, I., Cho, S. L., Hopf, F. W. and Bonci, A. (2013) 'Rescuing cocaine-induced prefrontal cortex hypoactivity prevents compulsive cocaine seeking', Nature, 496(7445), pp. 359–362. doi: 10.1038/nature12024.

Chen, L. M., Heider, B., Williams, G. V., Healy, F. L., Ramsden, B. M. and Roe, A. W. (2002) 'A chamber and artificial dura method for long-term optical imaging in the monkey', Journal of Neuroscience Methods, 113(1), pp. 41–49. doi: 10.1016/S0165-0270(01)00475-7.

Chen, R. C.-H., Atry, F., Richner, T., Brodnick, S. K., Pisaniello, J., Ness, J., Suminski, A., Williams, J. C. and Pashaie, R. (2020) 'A system identification analysis of optogenetically evoked electrocorticography and cerebral blood flow responses', Journal of Neural Engineering.

Chernov, M. and Roe, A. W. (2014) 'Infrared neural stimulation: a new stimulation tool for central nervous system applications', Neurophotonics, 1(1), p. 011011. doi: 10.1117/1.nph.1.1.011011.

Chiang, C.-H., Won, S. M., Orsborn, A. L., Yu, K. J., Trumpis, M., Bent, B., Wang, C., Xue, Y., Min, S., Woods, V., Yu, C., Kim, B. H., Kim, S. B., Huq, R., Li, J., Seo, K. J., Vitale, F., Richardson, A., Fang, H., et al. (2020) 'Development of a neural interface for high-definition, long-term recording in rodents and nonhuman primates', Science Translational Medicine, 12(538), p. eaay4682. doi: 10.1126/scitranslmed.aay4682.

Choi, I. A., Lee, C. S., Kim, H. Y., Choi, D. H. and Lee, J. (2018) 'Effect of inhibition of DNA methylation combined with task-specific training on chronic stroke recovery', International Journal of Molecular Sciences, 19(7). doi: 10.3390/ijms19072019.

Chuong, A. S., Miri, M. L., Busskamp, V., Matthews, G. A. C., Acker, L. C., Sørensen, A. T., Young, A., Klapoetke, N. C., Henninger, M. A., Kodandaramaiah, S. B., Ogawa, M., Ramanlal, S. B., Bandler, R. C., Allen, B. D., Forest, C. R., Chow, B. Y., Han, X., Lin, Y., Tye, K. M., et al. (2014) 'Noninvasive optical inhibition with a red-shifted microbial rhodopsin', Nature Neuroscience, 17(8), pp. 1123–1129. doi: 10.1038/nn.3752.

Dai, P.M., Huang, H., Zhang, L., He, J., Zhao, X.D., Yang, F.H., Zhao, N., Yang, J.Z., Ge, L.J., Lin, Y., et al. (2017). A pilot study on transient ischemic stroke induced with endothelin-1 in the rhesus monkeys. Sci. Rep. 7, 1–12.

Diester, I., Kaufman, M. T., Mogri, M., Pashaie, R., Goo, W., Yizhar, O., Ramakrishnan, C., Deisseroth, K. and Shenoy, K. V. (2011a) 'An optogenetic toolbox designed for primates', Nature Neuroscience, 14(3), pp. 387–397. doi: 10.1038/nn.2749.

Diester, I., Kaufman, M. T., Mogri, M., Pashaie, R., Goo, W., Yizhar, O., Ramakrishnan, C., Deisseroth, K. and Shenoy, K. V. (2011b) 'An optogenetic toolbox designed for primates', Nature Neuroscience, 14(3), pp. 387–397. doi: 10.1038/nn.2749.

Ding, M. (2016). Linking Functional Connectivity and Structural Connectivity Quantitatively: A Comparison of Methods. Brain Connectivity.

Edwardson, M A et al. (2013). "New modalities of brain stimulation for stroke rehabilitation". Experimental Brain Research 224.3, pp. 335–358. doi: 10.1007/s00221-012-3315-1.

Eichler, M. (2006). On the Evaluation of Information Flow in Multivariate Systems by the Directed Transfer Function. Biol Cybern, 94, 469–482.

El-Shamayleh, Y., & Horwitz, G. D. (2019). Primate optogenetics: progress and prognosis. Proc. Natl. Acad. Sci. U.S.A. 116. doi: 10.1073/pnas.1902284116

Ellenbroek, B., & Youn, J. (2016). Rodent models in neuroscience research: is it a rat race? Dis Model Mech, 9(10), 1079-1087. doi: 10.1242/dmm.026120.

Fang, Q., and Boas, D.A. (2009). Monte Carlo Simulation of Photon Migration in 3D Turbid Media Accelerated by Graphics Processing Units. Opt. Express 17, 20178.

Fox, Kieran C.R. et al. (2020). "Intrinsic network architecture predicts the effects elicited by intracranial electrical stimulation of the human brain". Nature Human Behaviour 4.10, pp. 1039–1052. issn: 23973374. doi: 10. 1038/s41562-020-0910-1.

Fox, Michael D., Randy L. Buckner, et al. (2012). “Efficacy of transcranial magnetic stimulation targets for depression is related to intrinsic functional connectivity with the subgenual cingulate”. Biological Psychiatry 72.7, pp. 595– 603. issn: 00063223. doi: 10.1016/j.biopsych.2012.04.028.

Fox, Michael D., Mark A. Halko, et al. (2012). “Measuring and manipulating brain connectivity with resting state functional connectivity magnetic resonance imaging (fcMRI) and transcranial magnetic stimulation (TMS)”. NeuroImage 62.4, pp. 2232–2243. issn: 10538119. doi: 10 . 1016/j.neuroimage. 2012 . 03 . 035.

Frederick A. King et al., Primates. Science 240, 1475-1482 (1988). DOI:10.1126/science.3287624

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani (2001). The Elements of Statistical Learning. Vol. 1. 10. Springer Series in Statistics New York.

Frías, I., Starrs, F., Gisiger, T., Minuk, J., Thiel, A. and Paquette, C. (2018) ‘Interhemispheric connectivity of primary sensory cortex is associated with motor impairment after stroke’, Scientific Reports, 8(1), pp. 1–10. doi: 10.1038/s41598-018-29751-6.

Friston, K. J. (2011). Functional and effective connectivity: a review. Brain connectivity, 1(1), 13-36.

Fukushima, M., Betzel, R. F., He, Y., et al. (2017). Structure–function relationships during segregated and integrated network states of human brain functional connectivity. Brain Struct Funct.

Garrity, A. G. et al. (2007). “Aberrant “default mode” functional connectivity in schizophrenia”. American Journal of Psychiatry. doi: 10.1176/ajp.2007.164.3.450.

Gerlai, R., Thibodeaux, H., Palmer, J.T., Van Lookeren Campagne, M., and Van Bruggen, N. (2000). Transient focal cerebral ischemia induces sensorimotor deficits in mice. Behav. Brain Res. 108, 63–71.

Gibbs, R. A., Rogers, J., Katze, M. G., Bumgarner, R. E., Weinstock, G. M., Mardis, E. R., ... & Zwieg, A. S. (2007). Evolutionary and Biomedical Insights From The Rhesus Macaque Genome. Science, 5822(316), 222-234. doi.org/10.1126/science.1139247

Gradinaru, V., Mogri, M., Thompson, K. R., Henderson, J. M. and Deisseroth, K. (2009) ‘Optical deconstruction of parkinsonian neural circuitry’, Science, 324(5925), pp. 354–359. doi: 10.1126/science.1167093.

Griggs, D.J., Khateeb, K., Philips, S., Chan, J.W., Ojemann, W., and Yazdan-Shahmorad, A. (2019). Optimized large-scale optogenetic interface for non-human primates. In Optogenetics and Optical Manipulation 2019, S.K. Mohanty, and E.D. Jansen, eds. (SPIE), p. 3.

Griggs, D., Khateeb, K., Zhou, J., Liu, T., Wang, R., and Yazdan-Shahmorad, A. (2021a). Multi-modal artificial dura for simultaneous large-scale optical access and large-scale electrophysiology in non-human primate cortex. J. Neural Eng. 18, 055006.

Griggs, D.J., Belloir, T., and Yazdan-Shahmorad, A. (2021b). Large-scale neural interfaces for optogenetic actuators and sensors in non-human primates. 11663, 1166305.

Griggs, D.J. (2022a). Accessible Optogenetic Technologies for Non-human Primate Research. Doctoral dissertation, University of Washington, Electrical and Computer Engineering.

Griggs, D.J., Garcia, A.D., Au, W.Y., Ojemann, W.K.S., Johnson, A.G., Ting, J.T., Buffalo, E.A. and Yazdan-Shahmorad, A. (2022b). Improving the efficacy and accessibility of intracranial viral vector delivery in non-human primates. Pharmaceutics, p. 2022.06.06.494543. doi: 10.1101/2022.06.06.494543.

Griggs, Devon J, Khateeb, K., Zhou, J., Liu, T., Wang, R. and Yazdan-Shahmorad, A. (2021). Multi-modal artificial dura for simultaneous large-scale optical access and large-scale electrophysiology in non-human primate cortex. Journal of Neural Engineering, Special Issue 'Neuroelectronic Interfaces', 18(5), p. 2021.02.03.429596. doi: 10.1101/2021.02.03.429596.

Gulati, T., Won, S.J., Ramanathan, D.S., Wong, C.C., Bodepudi, A., Swanson, R.A., and Ganguly, K. (2015). Robust neuroprosthetic control from the stroke perilesional cortex. J. Neurosci. 35, 8653–8661.

Guzman, S.J., Schlögl, A., Frotscher, M., and Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. Science 353, 1117–1123.

Harari, A., Shojaie, A., and Simon, N. (2019). Nonparametric regression with adaptive truncation via a convex hierarchical penalty. Biometrika 106.1, 87–107.

Harrison, T.C., Silasi, G., Boyd, J.D., and Murphy, T.H. (2013). Displacement of sensory maps and disorganization of motor cortex after targeted stroke in mice. Stroke 44, 2300–2306.

Heffner, H.E., and Heffner, R.S. (1986). Hearing loss in Japanese macaques following bilateral auditory cortex lesions. J. Neurophysiol. 55, 256–271.

Horn, Andreas et al. (2017). "Connectivity Predicts deep brain stimulation outcome in Parkinson disease". Annals of Neurology 82.1, pp. 67–78. issn: 15318249. doi: 10.1002/ana.24974.

Horwitz, B. (2003). The elusive concept of brain connectivity. Neuroimage, 19(2), 466-470.

Hu, Yu et al. (2014). "Local paths to global coherence: Cutting networks down to size". Physical Review E - Statistical, Nonlinear, and Soft Matter Physics. issn: 15502376. doi: 10.1103/PhysRevE.89.032802.

Huang, H., Ding, M. (2016). Linking Functional Connectivity and Structural Connectivity Quantitatively: A Comparison of Methods. Brain Connectivity.

Huang, Yuhao et al. (2019). "Intracortical Dynamics Underlying Repetitive Stimulation Predicts Changes in Network Connectivity". The Journal of neuroscience 39.31, pp. 6122–6135. issn: 15292401. doi: 10.1523/JNEUROSCI.0535-19.2019.

Humphrey, N. (1974). Vision in a monkey wihtout striate cortex: a case study. Perception 3, 241–255.

Jackson, A., Mavoori, J. & Fetz, E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. Nature, 444, 56–60. doi:10.1038/nature05226

Jackson, Andrew, Jaideep Mavoori, and Eberhard E. Fetz (2006). "Long-term motor cortex plasticity induced by an electronic neural implant". Nature 444.7115, pp. 56–60. issn: 14764687. doi: 10.1038/nature05226.

Ji, B., Guo, Z., Wang, M., Yang, B., Wang, X., Li, W. and Liu, J. (2018) 'Flexible polyimide-based hybrid opto-electric neural interface with 16 channels of micro-LEDs and electrodes', Microsystems and Nanoengineering, 4(1). doi: 10.1038/s41378-018-0027-0.

Jiang, W., Tremblay, F. and Elaine Chapman, C. (2018) 'Context-dependent tactile texture-sensitivity in monkey M1 and S1 cortex', Journal of Neurophysiology, 120(5), pp. 2334–2350. doi: 10.1152/jn.00081.2018.

Ju, N., Jiang, R., Macknik, S.L., Martinez-Conde, S., and Tang, S. (2018). Long-term all-optical interrogation of cortical neurons in awake-behaving nonhuman primates. PLoS Biol. 16.

Kaeser, M., Wyss, A.F., Bashir, S., Hamadjida, A., Liu, Y., Bloch, J., Brunet, J.F., Belhaj-Saif, A., and Rouiller, E.M. (2010). Effects of unilateral motor cortex lesion on ipsilesional hand's reach and grasp performance in monkeys: Relationship with recovery in the contralesional hand. J. Neurophysiol. 103, 1630–1645.

Kaiser, Roselinde H. et al. (2015). "Large-scale network dysfunction in major depressive disorder: A meta-analysis of resting-state functional connectivity". JAMA Psychiatry 72.6, pp. 603–611. issn: 2168622X. doi: 10.1001/jamapsychiatry.2015.0071.

Keller, Corey J., Stephan Bickel, et al. (2011). "Intrinsic functional architecture predicts electrically evoked responses in the human brain". Proceedings of the National Academy of Sciences of the United States of America 108.41, p. 17234. issn: 10916490. doi: 10.1073/pnas.1114425108.

Keller, Corey J., Yuhao Huang, et al. (2018). "Induction and quantification of excitability changes in human cortical networks". Journal of Neuroscience 38.23, pp. 5384–5398. issn: 15292401. doi: 10.1523/JNEUROSCI.1088-17.2018.

Khambhati, Ankit N. et al. (2019). "Functional control of electrophysiological network architecture using direct neurostimulation in humans". Network Neuroscience 3.3, pp. 848–877. doi: 10 . 1162 / NETN.

Khanna, P., Totten, D., Novik, L., Roberts, J., Morecraft, R.J., and Ganguly, K. (2021). Low-frequency stimulation enhances ensemble co-firing and dexterity after stroke. Cell 184, 912-930.e20.

Khateeb, K, Yao, Z., Kharazia, V. N., Burunova, E. P., Song, S., Wang, R. and Yazdan-Shahmorad, A. (2019b). A Practical Method for Creating Targeted Focal Ischemic Stroke in the Cortex of Nonhuman Primates∗. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS 3515–3518.

Khateeb, K, Zhou, J., Gala, A., Rahimi, M., Griggs, D. J., Ip, Z. and Yazdan-shahmorad, A. (2022). A versatile toolbox for studying cortical physiology in primates. Cell Reports Methods, p. 100183. doi: 10.1016/j.crmeth.2022.100183.

Khateeb, Karam, Yao, Z., Kharazia, V. N., Burunova, E. P., Song, S., Wang, R. and Yazdan-Shahmorad, A. (2019b). A Practical Method for Creating Targeted Focal Ischemic Stroke in the Cortex of Nonhuman Primates∗. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. EMBS 3515–3518.

Kleim, J., Bruneau, R., VandenBerg, P., MacDonald, E., Mulrooney, R., & Pocock, D. (2003). Motor cortex stimulation enhances motor recovery and reduces peri-infarct dysfunction following ischemic insult. Neurol Res, 25, 789–793.

Kleiner, M., Brainard, D. H. and Pelli, D. G. (2007) 'What's new in Psychtoolbox-3?', in Perception 36 ECVP Abstract Supplement.

Kooijmans, R. N., Sierhuis, W., Self, M. W. and Roelfsema, P. R. (2020) 'A Quantitative Comparison of Inhibitory Interneuron Size and Distribution between Mouse and Macaque V1, Using Calcium-Binding Proteins', Cerebral Cortex Communications, 1(1), pp. 1–14. doi: 10.1093/texcom/tgaa068.

Koulakov, A. A., Hromádka, T., & Zador, A. M. (2009). Correlated Connectivity and the Distribution of Firing Rates in the Neocortex. The Journal of Neuroscience, 29(12), 3685-3694. www.jneurosci.org/content/29/12/3685

Le Friec, A., Desmoulin, F., Demain, B., Davoust, C., Robert, L., Duval, T., Rémy, F., Cirillo, C., and Loubinoux, I. (2020). A Reproducible New Model of Focal Ischemic Injury in the Marmoset Monkey: MRI and Behavioural Follow-Up. Transl. Stroke Res. 1–14.

Ledochowitsch, P., Yazdan-Shahmorad, A., Bouchard, K.E., Diaz-Botia, C., Hanson, T.L., He, J.W., Seybold, B.A., Olivero, E., Phillips, E.A.K., Blanche, T.J., et al. (2015a). Strategies for optical control and simultaneous electrical readout of extended cortical circuits. J. Neurosci. Methods 256, 220–231.

Ledochowitsch, P., Yazdan-Shahmorad, A., Bouchard, K.E., Diaz-Botia, C., Hanson, T.L., He, J.W., Seybold, B.A., Olivero, E., Phillips, E.A.K., Blanche, T.J., et al. (2015b). Strategies for optical control and simultaneous electrical readout of extended cortical circuits. J. Neurosci. Methods 256, 220–231.

Levy, R.M., Harvey, R.L., Kissela, B.M., Winstein, C.J., Lutsep, H.L., Parrish, T.B., Cramer, S.C., and Venkatesan, L. (2016). Epidural Electrical Stimulation for Stroke Rehabilitation. Neurorehabil. Neural Repair.

Levy, Robert M. et al. (2016). "Epidural Electrical Stimulation for Stroke Rehabilitation". Neurorehabilitation and Neural Repair 30.2, pp. 107–119. issn: 15526844. doi: 10.1177/1545968315575613.

Liu, Y., and Rouiller, E.M. (1999). Mechanisms of recovery of dexterity following unilateral lesion of the sensorimotor cortex in adult monkeys. In Experimental Brain Research, (Springer Verlag), pp. 149–159.

LM, Chen, B., Heider, G. V., Williams, F. L., H., R. B. M. and Roe, A. W. (2002) 'A chamber and artificial dura method for long-term optical imaging in the monkey', Journal of Neuroscience Methods, 113(1), pp. 41–49. doi: 10.1016/S0165-0270(01)00475-7.

Lozano, Andres M. and Nir Lipsman (2013). "Probing and Regulating Dysfunctional Circuits Using Deep Brain Stimulation". Neuron 77.3, pp. 406–424. issn: 08966273. doi: 10.1016/j.neuron.2013.01.020.

Macknik, S.L., Alexander, R.G., Caballero, O., Chanovas, J., Nielsen, K.J., Nishimura, N., Schaffer, C.B., Slovin, H., Babayoff, A., Barak, R., et al. (2019). Advanced Circuit and Cellular Imaging Methods in Nonhuman Primates. J. Neurosci. 39, 8267–8274.

Maeda, M., Takamatsu, H., Furuichi, Y., Noda, A., Awaga, Y., Tatsumi, M., Yamamoto, M., Ichise, R., Nishimura, S., and Matsuoka, N. (2005). Characterization of a novel thrombotic middle cerebral artery occlusion model in monkeys that exhibits progressive hypoperfusion and robust cortical infarction. J. Neurosci. Methods 146, 106–115.

Markram, Henry et al. (1997). "Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs". Science 275.5297, pp. 213–215. issn: 00368075. doi: 10.1126/science.275.5297.213.

Meinshausen, Nicolai and Peter Bu¨hlmann (2010). "Stability selection". Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72.4, pp. 417–473.

Meyer, H. S., Schwarz, D., Wimmer, V. C., Schmitt, A. C., Kerr, J. N. D., Sakmann, B. and Helmstaedter, M. (2011) 'Inhibitory interneurons in a cortical column form hot zones of inhibition in layers 2 and 5A', Proceedings of the National Academy of Sciences of the United States of America, 108(40), pp. 16807–16812. doi: 10.1073/pnas.1113648108.

Milo, R. et al. (2002). "Network motifs: Simple building blocks of complex networks". Science. issn: 00368075. doi: 10.1126/science.298.5594.824.

Milo, Ron et al. (2004). "Superfamilies of Evolved and Designed Networks". Science. issn: 00368075. doi: 10.1126/science.1089167.

Momi, Davide et al. (2021). "Network-level macroscale structural connectivity predicts propagation of transcranial magnetic stimulation". NeuroImage 229.July 2020, p. 117698. issn: 10959572. doi: 10.1016/j.neuroimage.2020.117698.

Mouse Genome Sequencing Consortium. Initial sequencing and comparative analysis of the mouse genome. Nature 420, 520–562 (2002). doi.org/10.1038/nature01262

Muldoon, Sarah Feldt et al. (2016). “Stimulation-Based Control of Dynamic Brain Networks.” PLoS computational biology 12.9, e1005076. issn: 1553-7358. doi: 10.1371/journal.pcbi.1005076.

Murata, Y., Higo, N., Oishi, T., Yamashita, A., Matsuda, K., Hayashi, M., and Yamane, S. (2008). Effects of motor training on the recovery of manual dexterity after primary motor cortex lesion in macaque monkeys. J. Neurophysiol. 99, 773–786.

Murphy, A.P., Leopold, D.A., Humphreys, G.W., and Welchman, A.E. (2016). Lesions to right posterior parietal cortex impair visual depth perception from disparity but not motion cues. Philos. Trans. R. Soc. B Biol. Sci. 371.

Nakai, Yasuo et al. (2021). “Regional abnormality of functional connectivity is associated with clinical manifestations in individuals with intractable focal epilepsy”. Scientific Reports 11.1, pp. 1–10. doi: 10.1038/s41598-021-81207-6.

Nudo, R.J., and Milliken, G.W. (1996). Reorganization of movement representations in primary motor cortex following focal ischemic infarcts in adult squirrel monkeys. J. Neurophysiol. 75, 2144–2149.

Nudo, R.J., Wise, B.M., SiFuentes, F., and Milliken, G.W. (1996). Neural substrates for the effects of rehabilitative training on motor recovery after ischemic infarct. Science (80-.). 272, 1791–1794.

Nudo, R.J., Larson, D., Plautz, E.J., Friel, K.M., Barbay, S., and Frost, S.B. (2003). A squirrel monkey model of poststroke motor recovery. ILAR J. 44, 161–174.

Obi, K., Amano, I. and Takatsuru, Y. (2018) ‘Role of dopamine on functional recovery in the contralateral hemisphere after focal stroke in the somatosensory cortex’, Brain Research, 1678, pp. 146–152. doi: 10.1016/j.brainres.2017.10.022.

Ocker, Gabriel Koch et al. (2017). “From the statistics of connectivity to the statistics of spike times in neuronal networks”. Current Opinion in Neurobiology. issn: 18736882.

Ohayon, S., Grimaldi, P., Schweers, N., Tsao, D. (2013). 'Saccade Modulation By Optical and Electrical Stimulation In The Macaque Frontal Eye Field', J. Neurosci., 42(33), 16684-16697. doi.org/10.1523/jneurosci.2675-13.2013

Ojemann, Jeffrey G., Robert Turner, Xiaoping Zhang, Brett L. Foster, Kristen L. Lundy, Wes Barker, and Kai J. Miller. (2020). “Human brain networks for novelty detection and attention orienting: A direct cortical stimulation study.” Journal of Neurophysiology 123, no. 6 (2020): 2544-2560.

Ojemann, W. K. S., Griggs, D. J., Ip, Z., Caballero, O., Jahanian, H., Martinez-Conde, S., Macknik, S. and Yazdan-Shahmorad, A. (2020) ‘A mri-based toolbox for neurosurgical planning in nonhuman primates’, Journal of Visualized Experiments, 2020(161), pp. 1–16. doi: 10.3791/61098.

Onnela, Jukka Pekka et al. (2005). “Intensity and coherence of motifs in weighted complex networks”. Physical Review E - Statistical, Nonlinear, and Soft Matter Physics 71.6.

Orsborn, A. L., Wang, C., Chiang, K., Maharbiz, M. M., Viventi, J. and Pesaran, B. (2015) ‘Semi-chronic chamber system for simultaneous subdural electrocorticography, local field potentials, and spike recordings’, International IEEE/EMBS Conference on Neural Engineering, NER, 2015-July, pp. 398–401. doi: 10.1109/NER.2015.7146643.

O'Shea, D.J., Duncker, L., Goo, W., Sun, X., Vyas, S., Trautmann, E.M., Diester, I., Ramakrishnan, C., Deisseroth, K., Sahani, M., & Shenoy, K.V. (2022). Direct neural perturbations reveal a dynamical mechanism for robust computation. *bioRxiv* [Preprint].

Padberg, J., Recanzone, G., Engle, J., Cooke, D., Goldring, A., and Krubitzer, L. (2010). Lesions in posterior parietal area 5 in monkeys result in rapid behavioral and cortical plasticity. J. Neurosci. 30, 12918–12935.

Paulk, A.C., Kfir, Y., Khanna, A.R. et al. (2022). Large-scale neural recordings with single neuron resolution using Neuropixels probes in human cortex. Nat Neurosci, 25, 252–263.

Pelli, D. G. (1997) 'The VideoToolbox software for visual psychophysics: Transforming numbers into movies', Spatial Vision, 10(4), pp. 437–442.

Perretta, G. "Non-human primate models in neuroscience research." Scandinavian Journal of Laboratory Animal Science 36.1 (2009): 77-85.

Pernice, Volker et al. (2011). "How structure determines correlations in neuronal networks". PLoS Computational Biology 7.5. issn: 1553734X. doi: 10.1371/journal.pcbi.1002059.

Plautz, E., Barbay, S., Frost, S., et al. (2003). Post-infarct cortical plasticity and behavioral recovery using concurrent cortical stimulation and rehabilitative training: a feasibility study in primates. Neurol Res, 25, 801–810.

Pons, T.P., Garraghtyt, P.E., and Mishkin, M. (1988). "Lesion-induced plasticity in the second somatosensory cortex of adult macaques". Proceedings of the National Academy of Sciences 85.14, 5279–5281.

Prakash, Rohit et al. (2012). "Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation". Nature Methods 9.12, pp. 1171–1179. issn: 15487091. doi: 10.1038/nmeth.2215.

Rajalingam, R., Sorenson, M., Azadi, R., Bohn, S., DiCarlo, J. J. and Afraz, A. (2021) 'Chronically implantable LED arrays for behavioral optogenetics in primates', Nature Methods, 18(9), pp. 1112–1116. doi: 10.1038/s41592-021-01238-9.

Ray, Supratim et al. (2008). "Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography". Journal of Neuroscience 28.45, pp. 11526–11536. issn: 02706474.

Reddy, J. W., Kimukin, I., Towe, E. and Chamanzar, M. (2019) 'Flexible, Monolithic, High-Density µlED Neural Probes for Simultaneous Optogenetics Stimulation and Recording', International IEEE/EMBS Conference on Neural Engineering, NER, 2019-March, pp. 831–834. doi: 10.1109/NER.2019.8717116.

Rebesco, James M. and Lee E. Miller (2011). "Enhanced detection threshold for in vivo cortical stimulation produced by Hebbian conditioning". Journal of Neural Engineering 8.1. issn: 17412560. doi: 10.1088/1741-2560/8/1/016011.

Rebesco, James M., Ian H. Stevenson, et al. (2010). "Rewiring neural interactions by micro-stimulation". Frontiers in Systems Neuroscience 4.August, pp. 1–15. issn: 16625137. doi: 10.3389/fnsys.2010.00039.

Renz, A. F., Lee, Jihyun, Tybrandt, K., Brzezinski, M., Lorenzo, D. A., Cerra Cheraka, M., Lee, Jaehong, Helmchen, F., Vörös, J. and Lewis, C. M. (2020) 'Opto-E-Dura: A Soft, Stretchable ECoG Array for Multimodal, Multiscale Neuroscience', Advanced Healthcare Materials, 9(17), pp. 1–26. doi: 10.1002/adhm.202000814.

Roelfsema, P. R. and Treue, S. (2014) 'Basic neuroscience research with nonhuman primates: A small but indispensable component of biomedical research', Neuron. doi: 10.1016/j.neuron.2014.06.003.

Ruggiero, R. N., Rossignoli, M. T., De Ross, J. B., Hallak, J. E. C., Leite, J. P. and Bueno-Junior, L. S. (2017) 'Cannabinoids and vanilloids in schizophrenia: Neurophysiological evidence and directions for basic research', Frontiers in Pharmacology, 8(JUN), pp. 1–27. doi: 10.3389/fphar.2017.00399.

Ruiz, O., Lustig, B. R., Nassi, J. J., Cetin, A., Reynolds, J. H., Albright, T. D., Callaway, E. M., Stoner, G. R. and Roe, A. W. (2013) 'Optogenetic interrogation of neural circuits in vivo', J Neurophysiol. 2013/06/14, 110(6), pp. 1455–1467. doi: 10.1152/jn.00153.2013.

Salton, Gerard and Michael McGill (1983). Introduction to Modern Information Retrieval. Auckland: McGraw-Hill, Inc.

Schaffer, C.B., Friedman, B., Nishimura, N., Schroeder, L.F., Tsai, P.S., Ebner, F.F., Lyden, P.D., and Kleinfeld, D. (2006). Two-photon imaging of cortical surface microvessels reveals a robust redistribution in blood flow after vascular occlusion. PLoS Biol. 4, 258–270.

Seeman, S. C., B. J. Mogen, E. E. Fetz, and S. I. Perlmutter, "Paired Stimulation for Spike-Timing-Dependent Plasticity in Primate Sensori- motor Cortex," The Journal of Neuroscience, 2017.

Sehm, Bernhard et al. (2012). "Dynamic modulation of intrinsic functional connectivity by transcranial direct current stimulation". Journal of Neurophysiology 108.12, pp. 3253–3263. issn: 00223077. doi: 10.1152/jn.00606.2012.

Seong, E., Seo, J.K., Kabir, A.M.R., Andrei, A., Wozny, C., Bieri, O., Yoon, G., and Schuermann, F. (2017). Cellular resolution optical access to brain regions in fissures: Imaging medial prefrontal cortex and grid cells in entorhinal cortex. J. Neurosci. Methods 288, 31–42.

Sivagnanam, S., Majumdar, A., Yoshimoto, K., Astakhov, V., Bandrowski, A., Martone, M., & Carnevale, N. T. (2013). Introducing the neuroscience gateway. CEUR Workshop Proceedings, 993.

Slovin, H., Arieli, A., Hildesheim, R. and Grinvald, A. (2002) 'Long-term voltage-sensitive dye imaging reveals cortical dynamics in behaving monkeys', Journal of Neurophysiology, 88(6), pp. 3421–3438. doi: 10.1152/jn.00194.2002.

Snyder, A., Morais, M., Willis, C. et al. (2015). Global network influences on local functional connectivity. Nat Neurosci, 18, 736–743.

Solomon, E. A. et al. (2018). "Medial temporal lobe functional connectivity predicts stimulation-induced theta power". Nature Communications 9.1, pp. 1–13. issn: 20411723. doi: 10.1038/s41467- 018- 06876- w.

Song, Sen, Kenneth D. Miller, and L. F. Abbott (2000). "Competitive Hebbian learning through spike-timing-dependent synaptic plasticity". Nature Neuroscience. issn: 10976256. doi: 10.1038/78829.

Song, Weiguo et al. (2013). "Cortical Plasticity Induced by Spike-Triggered Microstimulation in Primate Somatosensory Cortex". PLoS ONE 8.3. issn: 19326203. doi: 10.1371/journal.pone.0057453.

Sommer, C.J. (2017). Ischemic stroke: experimental models and reality. Acta Neuropathol.

Stam, C. J. et al. (2007). "Small-world networks and functional connectivity in Alzheimer's disease". Cerebral Cortex 17.1, pp. 92–99. issn: 10473211. doi: 10.1093/cercor/bhj127.

Stevenson, I.H., London, B.M., Oby, E.R., Sachs, N.A., Reimer, J., Englitz, B., et al. (2012). Functional Connectivity and Tuning Curves in Populations of Simultaneously Recorded Neurons. PLoS Comput Biol, 8(11): e1002775.

Straathof, M., Sinke, M. R. T., Dijkhuizen, R. M., et al. (2018). A systematic review on the quantitative relationship between structural and functional network connectivity strength in mammalian brains. J Cereb Blood Flow Metab.

Teo, L., and Bourne, J.A. (2014). A reproducible and translatable model of focal ischemic injury in the marmoset monkey. Brain Pathol. 24, 459–474.

Teskey, G., Flynn, C., Goertzen, C. D., Monfils, M. H., & Young, N. A. (2003). Cortical stimulation improves skilled forelimb use following a focal ischemic infarct in the rat. Neurological Research, 25(8), 794-800. doi: 10.1179/016164103771953871

Thornton, C., Hutchings, F., & Kaiser, M. (2019). The Virtual Electrode Recording Tool for EXtracellular Potentials (VERTEX) Version 2.0: Modelling in vitro electrical stimulation of brain tissue. Wellcome Open Research, 4, 20. doi.org/10.12688/wellcomeopenres.15058.1

Tian, H., Xu, K., Zou, L., Fang, Y. (2022) 'Multimodal Neural Probes For Combined Optogenetics and Electrophysiology', iScience, 1(25), 103612. doi.org/10.1016/j.isci.2021.103612

Tomsett, R. J. (2014). A novel simulation framework for modelling extracellular recordings in cortical tissue: implementation, validation and application to gamma oscillations in mammals (Doctoral dissertation). Newcastle University. theses.ncl.ac.uk/jspui/handle/10443/2436

Tomsett, R. J., Ainsworth, M., Thiele, A., Sanayei, M., Chen, X., Gieselmann, M. A., Whittington, M. A., Cunningham, M. O., & Kaiser, M. (2015). Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX): comparing multi-electrode recordings from simulated and biological mammalian cortical tissue. Brain Structure and Function, 220(4), 2333–2353. doi.org/10.1007/s00429-014-0793-x

Trautmann, E.M., O'Shea, D.J., Sun, X., Marshel, J.H., Crow, A., Hsueh, B., Vesuna, S., Cofer, L., Bohner, G., Allen, W., et al. (2021). Dendritic calcium signals in rhesus macaque motor cortex drive an optical brain-computer interface. Nat. Commun. 2021 121 12, 1–20.

Tremblay, S., Acker, L., Afraz, A., Albaugh, D.L., Amita, H., Andrei, A.R., Angelucci, A., Aschner, A., Balan, P.F., Basso, M.A., et al. (2020). An Open Resource for Non-human Primate Optogenetics. Neuron 108, 1075-1090.e6.

Tønnesen, J. (2013) 'Optogenetic cell control in experimental models of neurological disorders', Behavioural Brain Research, 255, pp. 35–43. doi: 10.1016/j.bbr.2013.07.007.

Vallar, G., Rusconi, M.L., Bignamini, L., Geminiani, G., and Perani, D. (1994). Anatomical correlates of visual and tactile extinction in humans: A clinical CT scan study. J. Neurol. Neurosurg. Psychiatry 57, 464–470.

Van Dam, D., & De Deyn, P. P. (2017). Non-human primate models for Alzheimer's disease-related research and drug discovery. Expert Opin Drug Discov, 12(2), 187-200. doi: 10.1080/17460441.2017.1271320.

Virley, D., Hadingham, S.J., Roberts, J.C., Farnfield, B., Elliott, H., Whelan, G., Golder, J., David, C., Parsons, A.A., and Hunter, A.J. (2004). A New Primate Model of Focal Stroke: Endothelin-1-Induced Middle Cerebral Artery Occlusion and Reperfusion in the Common Marmoset. J. Cereb. Blood Flow Metab. 24, 24–41.

Vuksanović, V., & Hövel, P. (2014). Functional connectivity of distant cortical regions: role of remote synchronization and symmetry in interactions. Neuroimage, 97, 1-8. doi: 10.1016/j.neuroimage.2014.04.039.

Waites, Anthony B. et al. (2005). "Effect of prior cognitive state on resting state networks measured with functional connectivity". Human Brain Mapping 24.1, pp. 59–68. issn: 10659471. doi: 10.1002/hbm.20069.

Wang, Zheng et al. (2013). “The Relationship of Anatomical and Functional Connectivity to Resting-State Connectivity in Primate Somatosensory Cortex”. Neuron 78.6, pp. 1116–1126. issn: 08966273. doi: 10.1016/j.neuron.2013.04.023.

Watson, B. O., Maclean, J. N. and Yuste, R. (2008) ‘Imaging action potentials with calcium indicators’, Cold Spring Harbor Protocols, 2008(11), p. pdb.top43. doi: 10.1101/pdb.top43.

Welch, Peter (1967). “The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms”. IEEE Transactions on audio and electroacoustics 15.2, pp. 70–73.

Wu, D., Chen, J., Wang, B., Zhang, M., Shi, J., Ma, Y., Zhu, Z., Yan, F., He, X., Li, S., et al. (2016). Endovascular ischemic stroke models of adult rhesus monkeys: A comparison of two endovascular methods. Sci. Rep. 6, 1–10.

Wurtz, R.H., and Goldberg, M.E. (1972). Activity of superior colliculus in behaving monkey. IV. Effects of lesions on eye movements. J. Neurophysiol. 35, 587–596.

Xerri, C., Merzenich, M.M., Peterson, B.E., and Jenkins, W. (1998). Plasticity of primary somatosensory cortex paralleling sensorimotor skill recovery from stroke in adult monkeys. J. Neurophysiol. 79, 2119–2148.

Yaroslavsky, A.N., Schulze, P.C., Yaroslavsky, I. V., Schober, R., Ulrich, F., and Schwarzmaier, H.J. (2002). Optical properties of selected native and coagulated human brain tissues in vitro in the visible and near infrared spectral range. Phys. Med. Biol. 47, 2059–2073.

Yazdan-Shahmorad, A, Diaz-Botia, C., Hanson, T., Ledochowitsch, P., Maharabiz, M. M., and Sabes, P. N. (2015) ‘Demonstration of a setup for chronic optogenetic stimulation and recording across cortical areas in non-human primates’, SPIE BiOS, 9305(March 2015), p. 93052K. doi: 10.1117/12.2080405.

Yazdan-Shahmorad, A., Diaz-Botia, C., Hanson, T.L., Kharazia, V., Ledochowitsch, P., Maharbiz, M.M., and Sabes, P.N. (2016). A Large-Scale Interface for Optogenetic Stimulation and Recording in Nonhuman Primates. Neuron 89, 927–939.

Yazdan-Shahmorad, A., Silversmith, D.B., Kharazia, V., and Sabes, P.N. (2018a). Targeted cortical reorganization using optogenetics in non-human primates. Elife.

Yazdan-Shahmorad, A., Tian, N., Kharazia, V., Samaranch, L., Kells, A., Bringas, J., He, J., Bankiewicz, K., and Sabes, P.N. (2018b). Widespread optogenetic expression in macaque cortex obtained with MR-guided, convection enhanced delivery (CED) of AAV vector to the thalamus. J. Neurosci. Methods.

Yazdan-Shahmorad, A., Silversmith, D.B., and Sabes, P.N. (2018c). Novel techniques for large-scale manipulations of cortical networks in non-human primates. Conf. Proc. ... Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf. 2018, 5479–5482.

Yang, X., Yao, C., Tian, T., Li, X., Yan, H., Wu, J., Li, H., Pei, L., Liu, D., Tian, Q., Zhu, L. Q. and Lu, Y. (2018) ‘A novel mechanism of memory loss in Alzheimer’s disease mice via the degeneration of entorhinal-CA1 synapses’, Molecular Psychiatry, 23(2), pp. 199–210. doi: 10.1038/mp.2016.151.

Yang, Yuxiao et al. (2021). “Modelling and prediction of the dynamic responses of large-scale brain networks during direct electrical stimulation”. Nature Biomedical Engineering. doi: 10.1038/s41551-020-00666-w.

Yazdan-Shahmorad, A, Diaz-Botia, C., Hanson, T., Ledochowitsch, P., Maharabiz, M. M., and Sabes, P. N. (2015) ‘Demonstration of a setup for chronic optogenetic stimulation and recording across cortical areas in non-human primates’, SPIE BiOS, 9305(March 2015), p. 93052K. doi: 10.1117/12.2080405.

Younce, John R. et al. (2021). “Resting-State Functional Connectivity Predicts STN DBS Clinical Response”. Movement Disorders 36.3, pp. 662–671. issn: 15318257. doi: 10.1002/mds.28376.

Zanos, T.P., Mineault, P.J., Monteon, J.A., & Pack, C.C. (2011). Functional connectivity during surround suppression in macaque area V4. 2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Boston, MA, USA, pp. 3342-3345.

Zátonyi, A., Madarász, M., Szabó, Á., Lőrincz, T., Hodován, R., Rózsa, B. and Fekete, Z. (2020) ‘Transparent, low-autofluorescence microECoG device for simultaneous Ca2+ imaging and cortical electrophysiology in vivo’, Journal of Neural Engineering, 17(1), p. 16062.

Zhang, Li I. et al. (1998). “A critical window for cooperation and competition among developing retinotectal synapses”. Nature. issn: 00280836. doi: 10.1038/25665.

Zhao, Liqiong et al. (2011). “Synchronization from second order network connectivity statistics”. Frontiers in Computational Neuroscience 5. July, pp. 1–16. issn: 16625188. doi: 10.3389/fncom.2011.00028.

Zhou, J., Khateeb, K., Gala, A., Rahimi, M., Griggs, D. J., Ip, Z. and Yazdan-shahmorad, A. (2022) ‘Neuroprotective Effects of Electrical Stimulation Following Ischemic Stroke in Non-Human Primates *’, IEEE EMBC, 119395, pp. 1–4.