

NORTHWESTERN UNIVERSITY

Providing proprioceptive information through biomimetic multi-electrode stimulation patterns

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Biomedical Engineering

By

Joseph Tebben Sombeck

EVANSTON, ILLINOIS

September 2022

Abstract

Persons with spinal cord injuries can use state-of-the-art brain-computer interfaces to control robotic arms. Despite this high-tech solution, their movements are slow and imprecise, much like those made by individuals who have lost proprioception, the sense of body position and movement. Intracortical microstimulation (ICMS) used to reactivate neural circuits in the somatosensory cortex is a promising approach for providing artificial proprioceptive feedback.

While tactile interfaces have advanced to the point where ICMS can provide force and contact location feedback to a spinal cord injured patient, proprioceptive interfaces have proven more difficult to develop. Previous proprioceptive interfaces either required months of training to use or evoked unreliable sensations. Part of the difficulty in designing these interfaces is the complicated somatotopy in proprioceptive cortical areas, where even simple limb movements evoke a complex spatial pattern of neural activity. It may be that stimulation patterns that evoke neural activity more nearly like that of limb movements will elicit naturalistic sensations and reduce the training time required to use proprioceptive interfaces.

This dissertation presents my work to develop such biomimetic stimulation patterns. By quantifying the spatiotemporal pattern of neural activity evoked by ICMS in Chapter 2, I show that stimulation through many electrodes with small amplitudes will be needed to recreate the complex spatial pattern of activity evoked by limb movements. In Chapter 3, I show that multi-electrode ICMS (mICMS) can replicate the rapid feedback provided by natural proprioception, something that single electrode stimulation cannot do. By modeling the evoked sensation with an artificially generated cortical map, I find that mICMS can produce effects as large as normal limb movements and in predictable directions in Chapter 4. Together, these results suggest that mICMS will be necessary to provide proprioceptive feedback in future afferent interfaces.

Acknowledgements

Throughout my time at Northwestern and in the Miller lab, I have been fortunate to be surrounded by caring individuals who were around to pick me up when I had setbacks and disappointments. For this, I would like to thank all the former and current members of the Miller lab. In particular, I would like to thank Juliet, Josie and Ege.

Special thanks go to Dr. Lee Miller, my PhD advisor. None of this would have been possible without you. I enjoyed the supportive environment that you cultivated in your lab and felt like your insightful comments always steered me in the correct direction. I hope to be able to replicate your mentorship and support in my future career.

I would also like to thank Drs. Sliman Bensmaia, Mitra Hartmann, and Sandro Mussa-Ivaldi, my dissertation committee. Your helpful comments and keen insight throughout my PhD have greatly improved this final product.

I would like to thank my family. You have supported me throughout my entire PhD even when though you may not have understood exactly what I was doing. To my friends, thank you for your continued support. I always looked forward to Dryhop Mondays and late-night soccer games.

Finally, I would like to thank my partner Cassie, who has been in my corner the entire time. She always listened when I discussed the difficulties I was having with my research and supported me through both the good times and bad times, something that I hope I can reciprocate.

Table of Contents

Abstract.....	2
Acknowledgements.....	3
Table of Contents.....	4
Table of Figures and Tables.....	8
Chapter 1 - Introduction	9
Importance of proprioception for motor control	11
Efferent interfaces to restore movement.....	13
Afferent interfaces to restore sensation.....	13
Neural response to ICMS	17
Multi-electrode stimulation to provide proprioceptive feedback.....	21
Importance of somatotopy for eliciting naturalistic sensations of limb movement	23
Summary	25
Chapter 2 - Characterizing the short-latency evoked response to intracortical microstimulation across a multi-electrode array	27
Foreword.....	27
Abstract.....	28
Introduction	29
Methods.....	31
Animal Subjects.....	31
Implant and data collection	31
Pipeline to record at short latencies after ICMS.....	32
Evaluating the performance of the RRA and acausal, time-reversed filtering	36
Stimulation protocol for characterizing the evoked response	36
Data analysis	39
Statistical Analysis.....	41
Results.....	41
Recording pipeline performance	41
Excitatory and inhibitory response to single pulses of ICMS.....	44
Temporal response to trains of ICMS	48

Spatial pattern of the response to ICMS trains.....	54
Discussion.....	57
Comparison of artifact suppression to previous techniques	57
Mode of activation of recorded spikes on stimulated and non-stimulated channels.....	59
Limitations due to missing directly evoked spikes.....	60
Qualitatively similar evoked responses across different experimental conditions.....	62
Implications for biomimetic stimulation patterns	63
Linking evoked activity to sensation	64
Online recording in the presence of stimulation artifact	65
Supplementary Materials.....	67
Technical description of the rapid-recovery amplifier.....	67
Chapter 3 - Short reaction times in response to multi-electrode intracortical microstimulation may provide a basis for rapid movement-related feedback	74
Foreword.....	74
Abstract.....	75
Introduction	76
Methods.....	78
Monkeys.....	78
Reaction time task	80
Stimulation and data collection	81
Data analysis	83
Statistical analysis	83
Results.....	84
Reaction time in response to natural proprioceptive and visual cues	84
Reaction time in response to single-electrode ICMS.....	86
Reaction time in response to high-amplitude sICMS.....	92
Reaction time in response to multi-electrode stimulation.....	94
Comparison of mICMS to sICMS	98
Effect of inter-electrode distance on mICMS reaction time.....	98
Discussion.....	101

Summary of results	101
Reaction time to single-electrode stimulation	101
Effect of electrode location on sICMS reaction time	103
Reaction time to multi-electrode stimulation	104
Effect of number of stimulation electrodes and total current	105
Effect of distance between electrodes	106
Implications for neuroprosthetics.....	106
Chapter 4 - Modeling proprioceptive feedback evoked by intracortical microstimulation using a topographical variational autoencoder	109
Foreword.....	109
Abstract.....	110
Introduction	111
Methods.....	113
Topographical variational autoencoder.....	113
Arm kinematics dataset	117
Training hyperparameters	117
Stimulation experiments and statistical analysis.....	118
Results.....	119
Artificial stimulation pipeline.....	119
Effect of stimulation at two example locations	122
Effect of increasing amplitude on the magnitude and direction of the effect of stimulation.....	124
Efficacy of stimulation at cluster centers.....	127
Effect of cluster size on the magnitude and direction of the effect of stimulation.....	131
Discussion.....	133
Developing proprioceptive interfaces	133
Parallels to psychophysical work in area MT	137
Assumptions and limitations.....	138
Chapter 5 - Discussion.....	140
Multi-electrode stimulation to provide proprioceptive feedback.....	141
Framework for designing biomimetic stimulation patterns	145

Rapid assessment of the evoked direction of limb movement	146
Targeting other proprioceptive areas	147
Combining biomimetic and learning approaches	148
Statistics of ICMS-evoked activity are inherently unnatural.....	149
Rebound excitation may be problematic for future afferent interfaces	151
Decoding motor intent during multi-electrode ICMS.....	151
Final Conclusions.....	152
References	154

Table of Figures and Tables

Fig. 2.1. Overview of artifact reduction pipeline.	34
Table 2.1. Experimental parameters.....	38
Fig. 2.2. Evaluation of rapid-recovery amplifier (RRA).....	43
Fig. 2.3. Excitatory response on the stimulated channel after single pulses of stimulation.	47
Fig. 2.4. Inhibitory response recorded on the stimulated channel after single pulses of stimulation	49
Fig. 2.5. Evoked response throughout 4-s long trains of stimulation.....	51
Fig. 2.6. Rebound excitation recorded on the stimulated electrode after short trains of stimulation.	53
Fig. 2.7. Evoked response on non-stimulated channels during 4-s long trains of ICMS.....	56
Fig. 2.8. Circuit diagram of the rapid-recovery amplifier.....	70
Fig. 2.9. Connection diagram	71
Fig. 2.10. Pictures of Stim Headstage and Front-end amplifier	72
Fig. 2.11. Diagram of breakout board.....	73
Fig. 3.1. Reaction time task description.....	79
Fig. 3.2. Dependence of reaction time on mechanical cue force magnitude.	85
Fig. 3.3. Dependence of reaction time to single-electrode stimulation on stimulation parameters.	87
Fig. 3.4. Reaction time to single-electrode stimulation for many electrodes.	89
Fig. 3.5. Dependence of reaction time to single-electrode stimulation on electrode position.....	91
Fig. 3.6. Reaction time to single-electrode stimulation at large currents.	93
Fig. 3.7. Reaction time to 16-electrode stimulation.	95
Fig. 3.8. Dependence of multi-electrode reaction time on total current and number of electrodes.	97
Fig. 3.9. Effect of distance between electrodes on multi-electrode reaction time.	100
Fig. 4.1. Architecture of the topoVAE model.....	115
Fig. 4.2. Lateral effect between neurons in the cortical map.....	116
Fig. 4.3. Simulating the effect of ICMS in proprioceptive cortex using a topographic variational autoencoder (topoVAE)	121
Fig. 4.4. Effect of stimulation at two example locations.....	123
Fig. 4.5. Effect of stimulation across locations.	126
Fig. 4.6. Efficacy of stimulation at progressively greater distances from cluster centers.	128
Fig. 4.7. Efficacy of many-location stimulation applied at cluster centers..	130
Fig. 4.8. Effect of cluster size.....	132

Chapter 1 - Introduction

Most people are aware of the five senses: sight, smell, taste, touch, and hearing. Proprioception, the sense of where our body is in space and how it is moving, is sometimes termed the ‘hidden sense’ because it generally remains largely below our conscious perception. Because of proprioception, I can bring my fingers to my nose when my eyes are closed. In fact, without proprioception, I can’t touch my nose accurately even *with* vision. Despite being “hidden”, proprioception is vital for controlling movements, as evidenced by the dramatic loss of body control exhibited by individuals who have lost it but retain the ability to activate their muscles, as depicted in the BBC documentary “The Man Who Lost His Body.”

In addition to losing the ability to generate movements, patients with spinal cord injuries also lose proprioception. While brain-computer interfaces (BCIs) have advanced to the point where patients can reach and grasp objects with a highly anthropomorphic robotic limb or even their own limbs (Collinger et al. 2013; Ajiboye et al. 2017; Flesher et al. 2021; Hochberg et al. 2012; Bouton et al. 2016), the movements of these patients are slow and imprecise, like those of people without proprioception. Intracortical microstimulation (ICMS) in the somatosensory cortex (S1) is a promising approach for providing somatosensory feedback, as it can elicit detectable sensations in rats, monkeys, and humans (Devecioğlu and Güçlü 2017; Fridman et al. 2010; Romo et al. 2000; London et al. 2008; O’Doherty et al. 2012; Flesher et al. 2016; Salas et al. 2018). ICMS has been used to provide tactile force and contact location feedback to a spinal cord injured patient, enabling him to grasp objects faster while controlling a robotic arm with both ICMS and visual feedback than with visual feedback alone (Flesher et al. 2021).

Proprioceptive interfaces have proven more difficult to develop than tactile interfaces. One previous interface evoked unreliable sensations (Tomlinson and Miller 2016), while another required multiple months of training in part because it elicited arbitrary sensations (Dadarlat, O'Doherty, and Sabes 2015). To reduce the training time required to use proprioceptive interfaces, stimulation patterns that evoke naturalistic sensations could be used. In my doctoral work, I aimed to develop stimulation patterns which evoke cortical activity like that evoked by natural limb movements, as I expected these patterns to elicit more nearly naturalistic sensations of proprioception.

Part of the difficulty in developing these biomimetic stimulation patterns is the complicated organization of neurons in proprioceptive cortices. In these areas, limb movements evoke a complex spatial pattern of neural activity. To mimic this spatial pattern, it is important to understand the pattern of activity evoked by ICMS. I measured the spatiotemporal pattern of cortical activity evoked by a wide range of ICMS parameters, as described in Chapter 2. Stimulation through single electrodes at large amplitudes activated neurons far from the site of stimulation, an approach ill-suited to mimic the complex spatial pattern of cortical activity evoked by limb movements. Instead, stimulation would need to be provided through multiple electrodes each with small amplitudes to evoke activity in a spatially restricted population of neurons.

One important aspect of natural proprioception is that it provides feedback much faster than vision, feedback fast enough to update ongoing movements. Surprisingly, in one study, single electrode ICMS provided feedback slower than even visual cues (Godlove, Whaite, and Batista 2014). To determine whether ICMS could provide feedback as rapid as natural proprioception, I compared the reaction times to proprioceptive and ICMS cues in Chapter 3. While single electrode stimulation also evoked long reaction times in my experiments, multi-electrode stimulation could

evoke reaction times as short or shorter than the proprioceptive cue, implying that this approach can provide feedback as fast as natural proprioception.

Without a well-defined somatotopy, stimulation in proprioceptive cortical areas may activate neurons with different encoding properties, resulting in unnatural sensations or perceived hand movements in unpredictable directions. To explore the effect of topography on the resulting sensation, I simulated the evoked sensations in response to multi-electrode stimulation using a computational model of proprioceptive cortex in Chapter 4. In this model, single-site stimulation produced weak effects in unpredictable directions. Multi-site stimulation could produce effects as large as normal limb movements and in predictable directions.

This introduction will review background information relevant to my doctoral work beginning with an overview of the importance of proprioception for motor control and previous afferent interfaces. After this, I will summarize the neural response to ICMS and cover a few benefits of multi-electrode stimulation.

Importance of proprioception for motor control

Proprioception, the sense of where the body is in space and how it is moving, is vital for planning and executing fast and accurate movements, as evidenced by people who have lost this sense but retain the ability to generate movements. To better understand the role of proprioception in controlling movements, the reaches of individuals without this sense were compared to able-bodied control subjects. When performing a gesture similar to slicing bread, control subjects made straight, highly planar movements that were consistent across trials. People without proprioception made curved and nonplanar movements which varied wildly across trials (Sainburg, Poizner, and Ghez 1993). These individuals made their largest errors when the hand needed to reverse directions, likely because this required coordination of the shoulder and elbow joints (Sainburg et

al. 1995). These individuals even struggled to coordinate activation of muscles during simple elbow flexion movements (Forget and Lamarre 1987). After they initially accelerated their arm by activating their biceps, subjects needed to activate their triceps to decelerate the arm. Control subjects activated their triceps near peak velocity and modulated the magnitude of this activation based on the speed of the arm. Patients without proprioception activated their triceps less than did control subjects, with magnitude that did not increase with speed. Furthermore, the timing of the burst was more variable relative to peak velocity than in control subjects, resulting in less accurate braking. Proprioceptive feedback from the moving limb is important for coordinating movements across joints and coordinating activation across muscles.

Individuals without proprioception struggle to move even when they have full visual feedback. To better understand the role of vision, these individuals were provided various amounts of visual feedback as they made reaches to targets (Ghez et al. 1990). Without any feedback, they could not accurately reach to the targets. When visual feedback was provided only before the reach, allowing the individuals to plan their movement, accuracy improved slightly, but was still much worse than reaches made by control subjects. Even with visual feedback both prior to and during the reach, these individuals still made less accurate reaches than control subjects. Furthermore, individuals without proprioception made smaller than expected forward/backwards reaches and larger than expected reaches laterally, indicating that they were unable to compensate for anisotropies in the inertial properties of the arm (Ghez et al. 1990). While visual feedback can slightly improve the accuracy of reaches made by individuals without proprioception, it cannot fully replace proprioception in either the planning or execution of movements.

Efferent interfaces to restore movement

Patients with spinal cord injuries or amputations lose both somatosensation and the ability to generate movements. To restore movement, researchers have implanted arrays of electrodes into the motor cortex of affected individuals to record neural signals, from which decoders can be used to predict motor intent (Serruya et al. 2002; Taylor, Tillery, and Schwartz 2002; Ethier et al. 2012). With such a decoder, participants were able to control a robotic limb to grasp and move objects (Collinger et al. 2013; Hochberg et al. 2012; Wodlinger et al. 2014). Instead of robotic limbs, therapeutic devices could use functional electrical stimulation to activate muscles in the participant's own arm (Ajiboye et al. 2017; Ethier et al. 2012).

Afferent interfaces to restore sensation

Devices which electrically stimulate the nervous system have previously been used to provide artificial sensory feedback. Perhaps the most successful sensory neural prosthesis is the cochlear implant, which has been implanted in more than 500,000 individuals, including ~100,000 children who were born deaf (Kral, Dorman, and Wilson 2019; Wilson and Dorman 2008). Cochlear implants bypass damaged cells in the cochlea by directly stimulating the auditory nerve, which then transmits information to the brain. Part of the success of the cochlear implant is due to the powerful processing and highly adaptable nature of the brain. Because of this, subjects with cochlear implants improved their speech understanding from 40% of "everyday" sentences to 90% within 12 months (Helms et al. 1997). Additionally, children born deaf who receive cochlear implants before 2 years of age, when their brains are still developing, can reach near-normal language skills by 8 years of age (Kral and O'Donoghue 2010).

To restore vision to blind individuals, electrical stimulation could be applied through electrodes implanted into visual cortex, an approach called intracortical microstimulation (ICMS).

Early experiments showed that stimulation in visual cortex evoked small phosphenes of various colors (Brindley 1973; Troyk et al. 2003). In monkeys, these results have been extended to the generation of simple shapes, motions, and letters (Chen et al. 2020).

Feedback about the direction of motion could be provided to blind individuals by stimulating area MT of visual cortex, an area involved in processing the direction of motion of a visual stimulus (Albright, Desimone, and Gross 1984; Van Essen, Maunsell, and Bixby 1981). Area MT neurons respond preferentially to specific directions of movement. Neurons with similar direction preferences are organized into columns, with columns further organized into pinwheel-like structures (Malonek, Tootell, and Grinvald 1994). Low amplitude stimulation applied in the middle of these columns can bias a monkey's perception of a moving dot field (Murasugi, Salzman, and Newsome 1993; Salzman et al. 1992). The direction of this bias was well predicted by the preferred motion direction of neurons recorded on the stimulated electrode.

To restore the sense of touch, ICMS could be applied to early tactile areas in S1, as this can evoke sensations of pressure and vibration (Romo et al. 1998; Flesher et al. 2016; Fifer et al. 2020). One approach to providing artificial tactile feedback is to elicit sensations that mimic those provided by the natural sense. ICMS in early tactile areas can approximate this, as it evokes localized sensations on the body, such as on a single finger, and sensation intensity can be controlled by modulating the stimulation amplitude (Kim, Callier, Tabot, Gaunt, et al. 2015; Tabot et al. 2013; Flesher et al. 2016). With this biomimetic approach, tactile interfaces provided force and contact location feedback to a spinal cord injured patient. When an experimenter touched fingers on a robotic hand, pressure sensors triggered stimulation on corresponding electrodes in early tactile areas, allowing the participant to correctly identify which finger was touched without any training (Flesher et al. 2016). The participant could also identify which of multiple fingers

were touched, though with less accuracy than when only a single finger was contacted. Expanding this study to an activity of daily living, the participant grasped and moved objects by controlling a robotic limb using signals recorded in motor cortex (Flesher et al. 2021). With both visual and ICMS feedback about the force on the index finger and middle fingers, the participant could grasp and move objects faster than with visual feedback alone. Much of the improvement (88%) came during the grasping phase.

Proprioceptive interfaces have proven more difficult to develop than tactile interfaces. In one example, monkeys needed to learn an arbitrary mapping between stimulation and feedback. Monkeys were trained to reach to invisible targets using feedback about the error vector between the position of their hand and the target (Dadarlat, O'Doherty, and Sabes 2015). Initially, monkeys used visual feedback in the form of a random moving-dot field. Once monkeys could use this feedback to accurately reach to the targets, visual feedback was paired with ICMS feedback through eight electrodes implanted in S1. The frequency of stimulation through the eight electrodes scaled with the distance between the hand and the target. To provide direction feedback, each of the electrodes was arbitrarily assigned a different movement direction, on which the error vector between hand and target position was projected. This projection further scaled the frequency of stimulation for each electrode. After months of training with increasingly noisy visual feedback, monkeys learned to use ICMS feedback alone to reach to the invisible targets. Even with all this training, monkeys took twice as long to reach the target with ICMS feedback than with visual feedback, implying that ICMS feedback was difficult to use.

It may be that stimulation patterns which evoke naturalistic sensations of limb movement would be easier to learn, thus reducing the time required to learn to use current proprioceptive interfaces (Bensmaia and Miller 2014). To develop biomimetic patterns for proprioception, it is

important to first understand the responses of neurons in proprioceptive cortical areas, such as area 2 of S1, to limb movements. The firing rates of neurons in area 2 modulate strongly during limb movements and during perturbations applied to the hand (Prud'homme and Kalaska 1994; London and Miller 2012; Chowdhury, Glaser, and Miller 2020). The velocity dependence of neural firing is classically analyzed using sinusoidal tuning functions (Georgopoulos et al. 1982), and neurons in area 2 are typically roughly sinusoidally tuned to both the velocity and load applied to the hand (Prud'homme and Kalaska 1994; London and Miller 2012). The direction of movement (or load) that evokes maximal firing is termed the 'preferred direction' (PD). The distribution of PDs in area 2 is bimodal, with a preference towards forward and backwards movements, an effect that is likely due to the biomechanics of the arm (Versteeg, Chowdhury, and Miller 2021). Within area 2, there is no discernable long-range organization of PDs, in part because receptors from many muscles, some spanning multiple joints, need to be mapped onto the two-dimensional cortical surface (Iwamura, Iriki, and Tanaka 1994; Pons et al. 1985). Even simple limb movements result in a complex spatial pattern of cortical activity across proprioceptive cortex. There is some local structure in area 2, though, as neighboring neurons are slightly more likely to have the same PD than non-neighboring neurons (Weber et al. 2011).

We previously tried to exploit the tendency of neurons with similar PDs to be near each other to elicit naturalistic sensations of proprioception. To infer the effect of stimulation in area 2, we trained monkeys to report which of two opposing directions was closest to the direction of a mechanical perturbation applied to the hand (Tomlinson and Miller 2016). On some trials, we applied ICMS during the mechanical perturbation to bias the monkey's perception of the direction of the perturbation. We stimulated on 4 electrodes each with similar PDs. Since neighboring neurons are likely to have the same PD in area 2 (Weber et al. 2011), we hypothesized that ICMS

would bias the monkey's perception in the direction of the recorded PDs. Stimulation through 6 out of 7 sets of electrodes predictably biased one monkey's perception of the mechanical perturbation. We could not reproduce this result in three subsequent monkeys. In these animals, stimulation produced biases, but their direction could not be predicted based on the PDs recorded on the stimulated electrodes. The different results across monkeys may have been due to slightly different array placements across monkeys or the specific population of neurons activated by stimulation. If we could have monitored the evoked activity after stimulation, perhaps we could better understand the differences across monkeys and design better stimulation patterns.

Neural response to ICMS

ICMS activates neurons directly and transsynaptically (Tehovnik et al. 2006). Stimulation changes the membrane potentials of nearby cells, sometimes causing them to fire. These directly evoked action potentials are typically initiated in axons, which have a higher density of sodium channels and thus have lower activation thresholds than cell bodies or dendrites (Nowak and Bullier 1998a, 1998b). Directly evoked action potentials propagate antidromically to cell bodies and orthodromically to presynaptic terminals, where they may cause further activity.

To activate neurons with homogeneous encoding properties within the complex somatotopy in proprioceptive cortical areas, it is important to understand the spatial pattern of directly and transsynaptically activity evoked by ICMS. It is difficult to monitor this activity, as ICMS causes a large shock artifact that corrupts electrical recordings (Hao, Riehle, and Brochier 2016; Butovas and Schwarz 2003; Weiss et al. 2018). Because of this, most previous studies recorded activity on electrodes distant from the site of stimulation and at long latencies after stimulation, likely missing short latency activity evoked near the site of stimulation (Hao, Riehle, and Brochier 2016; Voigt, Yusuf, and Kral 2018; Butovas and Schwarz 2003). Based on the

latency of recorded spikes, these studies likely recorded transsynaptically evoked activity and missed directly evoked activity.

Recording directly evoked activity is difficult because it typically occurs during the stimulus pulse, where the shock artifact is largest. Instead of recording this activity, researchers inferred the presence of directly evoked spikes in nearby neurons by blocking antidromically activated spikes (Stoney, Thompson, and Asanuma 1968). Stimulation applied to medullary pyramid evoked action potentials in axons that traveled antidromically to cell bodies in the pericruciate cortex. A recording electrode was inserted into the pericruciate cortex until it recorded such antidromic activity from a single neuron. The position of this electrode was then fixed. Stimulation was applied simultaneously in the medullary pyramid and near the recording electrode in cortex. If both stimuli evoked a spike, then these two spikes would collide, resulting in no recorded activity. Instead, if cortical stimulation did not evoke a spike, then the antidromic spike would appear on the recording electrode. Thus, the presence of the antidromic spike indicated that cortical stimulation did not evoke a spike. Since the timing of the antidromic spike was highly consistent across trials, researchers could determine precisely when the orthodromic spike was initiated, allowing them to infer activity evoked directly by the stimulating pulse instead of activity evoked through synapses.

Researchers used this setup to analyze the spatial spread of directly evoked activity by systematically moving the site of cortical stimulation. At each site, they measured the activation threshold, the minimum current required for 50% of pulses to evoke a spike (Stoney, Thompson, and Asanuma 1968). The activation threshold was proportional to the square of the distance from the neuron to the stimulating electrode:

$$I = kr^2$$

where I is the activation threshold, r is the distance from the neuron to the site of stimulation, and k is a constant. Thus, increasing amplitude resulted in a growing sphere of neural activation around the stimulation electrode.

While short-latency electrical recordings are difficult due to the shock artifact caused by ICMS, optical recording methods are unaffected. Using calcium imaging, the spatial pattern of directly evoked activation was sparser and more widely distributed than predicted by Stoney's simple spherical model (Histed, Bonin, and Reid 2009). Stimulation at low amplitudes activated cell bodies hundreds of micrometers from the stimulus site, a distance that did not increase with amplitude. Instead, increasing amplitude increased the activation density within the same activation volume.

We recently investigated this apparent conflict between Stoney's spherical model and Histed's recordings using a biophysical model of cortical stimulation (Kumaravelu et al. 2022). In the model, stimulation evoked action potentials within axons, consistent with their lower activation thresholds measured in physiological experiments (Nowak and Bullier 1998a, 1998b). The volume of axonal activation grew with stimulation amplitude in a manner well predicted by Stoney's model. Activation in axons propagated antidromically to somas, causing activation in a sparse and widely distributed population of cell bodies, similar to Histed's observations. In the model, increasing amplitude increased both the extent of somatic activation and activation density.

Evoked activity spreads farther from the stimulated electrode via synaptic activation. The amount of transsynaptic activity is largest near the stimulated electrode and decreases exponentially with distance (Logothetis et al. 2010; Butovas and Schwarz 2003). Increasing

amplitude increased both the amount and spread of transsynaptically evoked activity (Hao, Riehle, and Brochier 2016), likely due to the increased spread of directly evoked activity. At sufficiently high amplitudes, neurons millimeters from the site of stimulation can be activated by stimulation. The spatial pattern of transsynaptic activation depends on the connections made between neurons in cortex. In visual cortex, neurons with similar orientation preferences are organized into cortical columns, and columns of similar preferences are connected to each other (Roe et al. 2012). Because of this, stimulation in one column selectively activated neurons in distant columns with similar orientation preference (Hu et al. 2020).

The widespread transsynaptic activation implies that ICMS will not evoke spatially precise percepts. In contrast to this conclusion, ICMS in early tactile areas evoked localized sensations on the body, at locations that typically matched the receptive fields recorded on the stimulated electrodes (Tabot et al. 2013). Human participants have reported sensations on individual fingers, including the fingertips, and on the palm, in a manner that was consistent with the expected somatotopy of S1 (Flesher et al. 2016; Salas et al. 2018; Fifer et al. 2020). When stimulation evoked sensations of movement, the evoked sensations typically corresponded to only a small part of the arm (Salas et al. 2018). These results imply that the widespread transsynaptic activation is not important for determining the perceptual effects of ICMS. Instead, the perceptual effect of ICMS depends more on the local directly evoked activity.

These results suggest that ICMS through a single electrode will be insufficient to provide feedback about the whole arm and/or hand because the projection fields are very localized (Tabot et al. 2013; Flesher et al. 2016; Salas et al. 2018). Single electrode stimulation is also ill-suited for mimicking the complex spatial pattern of neural activity evoked by limb movements, as evoked activity is centered around the stimulated electrode. Instead, stimulation will need to be applied

through many electrodes simultaneously to activate clusters of neurons across proprioceptive cortex with similar encoding properties. To restrict the spatial extent of activation, stimulation amplitude through each electrode will need to be small. Multi-electrode stimulation can also sum localized projection fields to provide feedback about the whole arm/hand.

In addition to mimicking the spatial pattern of naturally evoked activity, ICMS should aim to match the temporal dynamics of naturally occurring activity. ICMS evokes activity in a large population of neurons immediately after each pulse, resulting in more synchronous activation across neurons than that evoked by natural stimuli. It has been hypothesized that this synchrony is part of the reason why stimulation at peripheral nerves induces unnatural sensations of tingling or pricking (paresthesia) (Tan et al. 2014). One approach to lower synchrony across the activated population is to deliver stimulation through multiple electrodes asynchronously, a process which does not affect detection thresholds (Kim, Callier, Tabot, Tenore, et al. 2015). Synchrony can even be reduced during single-electrode stimulation by replacing each pulse in a stimulus train with a high-frequency (> 1 kHz) burst of pulses (Formento et al. 2020). By increasing the amplitude of pulses throughout each burst, neurons with different activation thresholds are activated at different times during the burst, resulting in less synchronous activation.

Multi-electrode stimulation to provide proprioceptive feedback

Multi-electrode stimulation seems well suited for recreating the complex spatial pattern of neural activity evoked by limb movements and for reducing synchronous activation across proprioceptive cortex. Stimulation through many electrodes with small amplitudes, though, may not result in detectable sensations, as the detection threshold to single electrode stimulation is ~ 10 - $30 \mu\text{A}$ (Kim, Callier, Tabot, Gaunt, et al. 2015; Zaaïmi et al. 2013).

Two studies explored how stimulation sums across multiple electrodes by measuring the detectability of stimulus trains. In one study, multi-electrode stimulus trains were more detectable than predicted by combining the detection threshold of each individual electrode linearly (Zaaimi et al. 2013). Because of this, stimuli through single electrodes that normally was subthreshold summed to generate readily detectable sensations. This supralinear summation may reduce the amplitude required through each electrode, and thus improve the specificity of activation. However, these results conflicted with a separate study, where the detection thresholds during multi-electrode stimulation matched predictions of an independence model, implying linear summation across electrodes (Kim, Callier, Tabot, Tenore, et al. 2015). Understanding how sensations combine during multi-electrode stimulation will be important when designing stimulation patterns to provide feedback about the whole arm and hand.

These experiments measured the effect of multi-electrode stimulation near the detection threshold, implying that the evoked sensations were weak. For ICMS to be useful, it will need to evoke strong sensations that are easily detectable. For strong sensations, the detectability of a stimulus is no longer a good measure of intensity, as most stimuli are detected on every trial. Instead, sensation intensity can be measured by measuring the reaction time to that stimulus, as increasing intensity decreases reaction times (Pins and Bonnet 1996). Surprisingly, monkeys responded slower to high amplitude stimulation through a single electrode than to even visual cues (Godlove, Whaite, and Batista 2014). This is surprising, as stimulation is applied directly to the brain, and problematic because visual feedback is too slow to replace the rapid feedback of natural proprioception. It remains to be seen whether small amplitude stimulation through many electrodes can evoke intense sensations and provide rapid feedback.

Multi-electrode stimulation may also increase the number of distinguishable feedback levels provided by ICMS. The just noticeable difference (JND) is the minimum change in stimulus intensity that can be detected. For most natural stimuli, JND increases with increased sensation intensity due to corresponding increases in both the neural response and the noise of that response (Ekman 1959). In contrast, the JND for ICMS amplitude was constant over a wide range of amplitudes (Kim, Callier, Tabot, Gaunt, et al. 2015), implying a different relationship between stimulation amplitude and the mean response rate and variance. Across ICMS amplitudes, the JND is $\sim 15\text{-}30\ \mu\text{A}$ in monkeys and humans (Flesher et al. 2016; Kim, Callier, Tabot, Gaunt, et al. 2015). With detection thresholds of $\sim 20\ \mu\text{A}$, and a safety limit of $\sim 100\ \mu\text{A}$ through a single electrode (Chen et al. 2014; Rajan et al. 2015), implanted participants would be able to distinguish 3-6 levels of feedback. It may be possible to increase the number of distinguishable levels by delivering charge through multiple electrodes simultaneously, thus increasing the maximum deliverable charge.

Importance of somatotopy for eliciting naturalistic sensations of limb movement

We previously stimulated in area 2 during a mechanical perturbation applied to the hand, finding that this stimulation could bias a monkey's perception of the direction of the perturbation (Tomlinson and Miller 2016). The direction of the bias, though, could not be predicted by the mean PD of the stimulated electrodes in three out of four monkeys. This experiment was heavily inspired by similar experiments performed in area MT of visual cortex, where stimulation could predictably bias the perception of the direction of a moving dot field across multiple monkeys (Murasugi, Salzman, and Newsome 1993; Salzman et al. 1992). The different outcomes of these experiments were surprising, as the organization of neurons within area 2 is like that of area MT. In area MT, neurons with preference for similar directions of visual motion are found in the same cortical

column, and cortical columns with different preferences form pinwheels (Malonek, Tootell, and Grinvald 1994). In area 2, there is similar local structure, as neurons with similar PDs tend to form clusters (Weber et al. 2011). While there is no discernable long-range organization of neurons in area 2 (Pons et al. 1985), our recent computational model of proprioceptive cortex predicted that clusters of neurons with similar PDs would form pinwheel-like structures (Blum, Grogan, et al. 2021).

One important difference between these studies, apart from modality, is that the stimulating electrode was placed in the center of a group of neurons with similar direction preferences in area MT (Murasugi, Salzman, and Newsome 1993). To find the center of a group of neurons, researchers moved the electrode forwards and backwards until 200 μm of electrode travel recorded neurons with the same preferred motion direction. The electrode was then placed in the center of this group. At low amplitudes, where activated neurons were likely primarily located within the same group, stimulation in area MT could bias the monkey's perception in a direction well predicted by the direction preference of neurons recorded on the stimulated electrode. At higher amplitudes (80 μA), the monkeys could no longer report the direction of the visual stimulus, likely because high amplitude stimulation activated neurons far from the site of stimulation, neurons with heterogeneous direction preferences.

The importance of the stimulus location was further highlighted by systematically moving the stimulation electrode away from the center of a group of neurons with similar direction preferences. As the electrode moved away from the center, the magnitude of the bias caused by stimulation decreased greatly (Murasugi, Salzman, and Newsome 1993). Perhaps part of the reason our attempts to bias the perception of a proprioceptive stimuli failed was that we were unable to place our electrodes within the center of groups of neurons with similar PDs. These results show

that the placement of electrodes, in addition to the underlying organization of neurons, is important for determining the efficacy of stimulation.

Summary

This chapter has given a brief introduction to the importance of proprioception for controlling movements, previous efferent and afferent interfaces, the neural response to ICMS, and multi-electrode stimulation. The following chapters will detail my doctoral work toward developing multi-electrode stimulation patterns to provide proprioceptive feedback. In Chapter 2, I develop and evaluate tools to record neural activity at short latencies after stimulation on every electrode across a multi-electrode array, including the stimulated one. I quantify the spatiotemporal pattern of activity evoked by ICMS, confirming that the widespread activation due to stimulation through single electrodes is unlikely to recreate the complex spatial pattern of activity evoked by limb movements.

In one experiment reported several years ago, stimulation through a single electrode did not replicate the fast feedback provided by natural proprioception (Godlove, Whaite, and Batista 2014). In Chapter 3, I report reaction times to multi-electrode stimulation, to proprioceptive, and to visual cues. Those results show that multi-electrode stimulation can evoke reaction times as short or shorter than a proprioceptive cue, implying that this approach can recreate the feedback speed of natural proprioception.

Chapter 4 details a computational study that models the sensation evoked by stimulation in proprioceptive cortex. This study uses a model of proprioceptive cortex that regenerates arm kinematics via a cortical layer, in which neurons with similar PDs form clusters. I artificially stimulate this model to explore the effects of stimulus parameters and local topography on the evoked sensation. The simulations show that single electrode stimulation with adequate current to

evoke large magnitude sensations also activates a large population of neurons with dissimilar encoding properties, resulting in sensations in unpredictable directions. Instead, multi-electrode stimulation with small amplitudes activates a more homogeneous population of neurons, resulting in strong sensations more-nearly aligned with the mean PD of the stimulation sites. Finally, Chapter 5 presents the implications of the results in this thesis for the development of proprioceptive interfaces, along with future directions to extend this work.

Chapter 2 - Characterizing the short-latency evoked response to intracortical microstimulation across a multi-electrode array

Joseph T. Sombeck, Juliet Heye, Karthik Kumaravelu, Stefan M. Goetz, Angel V. Peterchev, Warren M. Grill, Sliman Bensmaia, Lee E. Miller

Foreword

The following chapter has been adapted from a manuscript published in the *Journal of Neural Engineering* in April 2022. The purpose of this project was to quantify the spatiotemporal pattern of neural activity evoked by intracortical microstimulation (ICMS) to aid in the design of biomimetic stimulation patterns. Recording electrical signals after stimulation is not trivial because stimulation causes a large shock artifact that lasts multiple milliseconds. In this project, I developed and evaluated hardware and software techniques to record short-latency neural activity after ICMS. With these techniques, I was able to monitor neural activity ~ 0.7 ms after stimulation offset on every channel across a multi-electrode array, including on the stimulated channel. I quantified the neural response to a wide range of stimulation parameters, a description that can both help to interpret results that probe the function of cortical areas and contribute to the design of stimulation patterns to improve afferent interfaces.

Abstract

Objective. Persons with tetraplegia can use brain-machine interfaces to make visually guided reaches with robotic arms. Without somatosensory feedback, these movements will likely be slow and imprecise, like those of persons who retain movement but have lost proprioception. Intracortical microstimulation (ICMS) has promise for providing artificial somatosensory feedback. If ICMS can mimic naturally occurring neural activity, afferent interfaces may be more informative and easier to learn than interfaces that evoke unnaturalistic activity. To develop such biomimetic stimulation patterns, it is important to characterize the responses of neurons to ICMS.

Approach. Using a Utah multi-electrode array, we recorded activity evoked by single pulses and trains of ICMS at a wide range of amplitudes and frequencies in two rhesus macaques. As the electrical artifact caused by ICMS typically prevents recording for many milliseconds, we deployed a custom rapid-recovery amplifier with nonlinear gain to limit signal saturation on the stimulated electrode. Across all electrodes after stimulation, we removed the remaining slow return to baseline with acausal high-pass filtering of time-reversed recordings.

Main results. After single pulses of stimulation, we recorded what was likely transsynaptically-evoked activity even on the stimulated electrode as early as ~ 0.7 ms. This was immediately followed by suppressed neural activity lasting 10–150 ms. After trains, this long-lasting inhibition was replaced by increased firing rates for ~ 100 ms. During long trains, the evoked response on the stimulated electrode decayed rapidly while the response was maintained on non-stimulated channels.

Significance. The detailed description of the spatial and temporal response to ICMS can be used to better interpret results from experiments that probe circuit connectivity or function of cortical areas. These results can also contribute to the design of stimulation patterns to improve afferent interfaces for artificial sensory feedback.

Introduction

Efferent brain-machine interfaces (BMIs) have advanced to the point where a spinal-cord injured patient can move a robotic arm using signals recorded from motor cortex (Collinger et al. 2013; Wodlinger et al. 2014; Hochberg et al. 2012). Without somatosensory feedback, the effectiveness of the movements generated through these interfaces will be limited, perhaps like those of people who have lost somatosensation (Ghez et al. 1990; Sainburg et al. 1995). Intracortical microstimulation (ICMS), which has been shown to elicit percepts in rats, monkeys, and humans (Devecioğlu and Güçlü 2017; Fridman et al. 2010; London et al. 2008; Romo et al. 2000), is a promising approach for providing artificial somatosensory feedback via an afferent interface (Tabot et al. 2013; Flesher et al. 2016). In the first such bidirectional BMI, monkeys could move a virtual arm to explore the “texture” of different virtual objects, a property conveyed by two different temporal patterns of ICMS (O’Doherty et al. 2011b). The monkeys moved the arm sequentially to the objects to find the one with the rewarded texture. More advanced methods have been used to supply a spinal cord injured patient with information about object contact location and force (Flesher et al. 2016; Flesher et al. 2021). Using a robotic arm, the patient was able to pick up, move, and place objects faster using vision combined with ICMS feedback than with visual feedback alone, primarily because they spent less time attempting to grasp the object (Flesher et al. 2021).

While some sensations of limb movement have been elicited with ICMS (Salas et al. 2018), achieving usable feedback about the position and movement of the arm has proven more difficult than providing the analogous artificial sense of touch. In one approach, monkeys learned to reach to invisible targets using ICMS feedback through eight arbitrarily chosen electrodes which provided information about the error vector between hand and target position (Dadarlat,

O'Doherty, and Sabes 2015). Monkeys only learned to use this feedback after a few months of training. To shorten this long learning period, it may be possible for ICMS to provide more naturalistic feedback (Bensmaia and Miller 2014). In a second approach, researchers attempted to evoke perceptions of hand movement by stimulating on sets of electrodes in somatosensory cortical area 2, that had similar preferred directions (Tomlinson and Miller 2016). This biomimetic approach was successful for six of seven sets of electrodes in one monkey but failed in three other monkeys. The difference across monkeys may have been due to the array placement across monkeys or the particular population of neurons activated by stimulation. Had it been possible to monitor the homogeneity of preferred directions of activated neurons, the explanation may have been clearer.

To better interpret experiments which use ICMS and to achieve more successful mimicry of naturally occurring activity, it will likely be important to quantify the evoked response of neurons to a range of stimulus parameters. However, recording at short latency after stimulation is difficult due to the large shock artifact it causes (Hao, Riehle, and Brochier 2016; Weiss et al. 2018). Many experiments have been limited to recordings made on electrodes hundreds of microns away or even on a separate array (Hao, Riehle, and Brochier 2016; Butovas and Schwarz 2003; Chen et al. 2020; Allison-Walker et al. 2021), thereby missing evoked activity near the stimulated electrode. Further, previous studies have typically characterized the evoked response to only single pulses of stimulation, whereas future afferent interfaces will need to employ trains of stimulation throughout a grasp and/or movement (Flesher et al. 2021).

We developed novel hardware and software techniques allowing us to record ~ 0.7 ms after stimulation offset on every electrode in an implanted microelectrode array, including even the stimulated one. We first used single pulses across a wide range of amplitudes to characterize the

short-latency excitatory and long-lasting inhibitory responses of neurons recorded on the stimulated electrode, as was done previously for non-stimulated electrodes (Hao, Riehle, and Brochier 2016; Butovas and Schwarz 2003). In preliminary experiments, we noticed that the evoked response on the stimulated electrode decreased rapidly throughout ~0.2-s, high-frequency trains, and that neurons greatly increased their firing rates for ~0.1 s after the end of the train. We extended the train length to 4 seconds, more akin to the prolonged stimulation provided by an afferent interface (Flesher et al. 2021). During these longer trains, the excitatory response recorded on the stimulated electrode decayed, while the response on non-stimulated electrodes was typically maintained throughout the train. The results in this paper can inform the interpretation and design of stimulation patterns for providing somatosensory feedback.

Methods

Animal Subjects

We performed experiments using two male rhesus macaques. Monkey H was 12.0 kg and monkey D was 10.0 kg when we performed the experiments. We performed all procedures in this study in accordance with the Guide for the Care and Use of Laboratory Animals. The institutional animal care and use committee of Northwestern University approved all procedures in this study under protocol #IS00000367.

Implant and data collection

Each monkey was implanted with a 96-electrode sputtered iridium-oxide multi-electrode array with 1.0 mm electrodes (Blackrock Neurotech, Salt Lake City, UT) in the proximal arm area of somatosensory cortical area 2. In addition to surface landmarks, we recorded intraoperatively from the cortical surface while manipulating the arm and hand to find the arm representation (for

more details, see Weber et al. 2011). We performed sensory mappings after implantation to confirm that recorded neurons had receptive fields corresponding to the proximal arm.

We used the Blackrock Stim Headstage, Front-End amplifier, and Neural Signal Processor (Blackrock Neurotech, Salt Lake City, UT) to record signals at 30 kHz. We delivered ICMS from the Blackrock CereStim R96. Unless otherwise noted, electrodes were stimulated with biphasic pulses, each phase lasting 200 μ s and separated by 53 μ s. We used the sync line from the CereStim R96 to determine stimulation onset, accounting for the 60 μ s delay between sync line going high and stimulation.

During all experiments, monkeys performed a center-out reaching task while holding the handle of a robotic manipulandum (for more details, see (London and Miller 2012)) or sat idly in the chair. Stimulation was delivered independently of the monkey's behavior.

Pipeline to record at short latencies after ICMS

Typically, ICMS causes large electrical artifacts which prevent neural recordings for an extended period after stimulation. When using the Blackrock Stim Headstage and Front-End amplifier to record on the stimulated electrode, the recorded signal saturated the amplifier for several milliseconds (Fig. 2.1a, dashed lines), after which the signal slowly recovered to baseline. To enable recording at shorter latencies, we developed a rapid-recovery amplifier (RRA, see Supplementary Materials) and used it instead of the Blackrock Stim Headstage and Front-end Amplifier. We used a custom-made breakout board to pass a single channel from the Cereport to the RRA, bypassing the Stim Headstage (see Supplementary Materials). The RRA has several features that allow it to operate on the same electrode as the stimulator, yet still recover rapidly after stimulation. The wide input range (± 15 V) of the first stage of the RRA prevents input voltage

clamping and current shunting as well as output saturation during the stimulus pulse. To prevent saturation of subsequent stages, the gain of the RRA declines rapidly from a maximum of ~1000 to a minimum of 1 during large dynamic swings of the front-end voltage. The output of the RRA, which was limited to ± 5 V, was connected to an analog input on the Blackrock Neural Signal Processor (Fig. 2.1b).

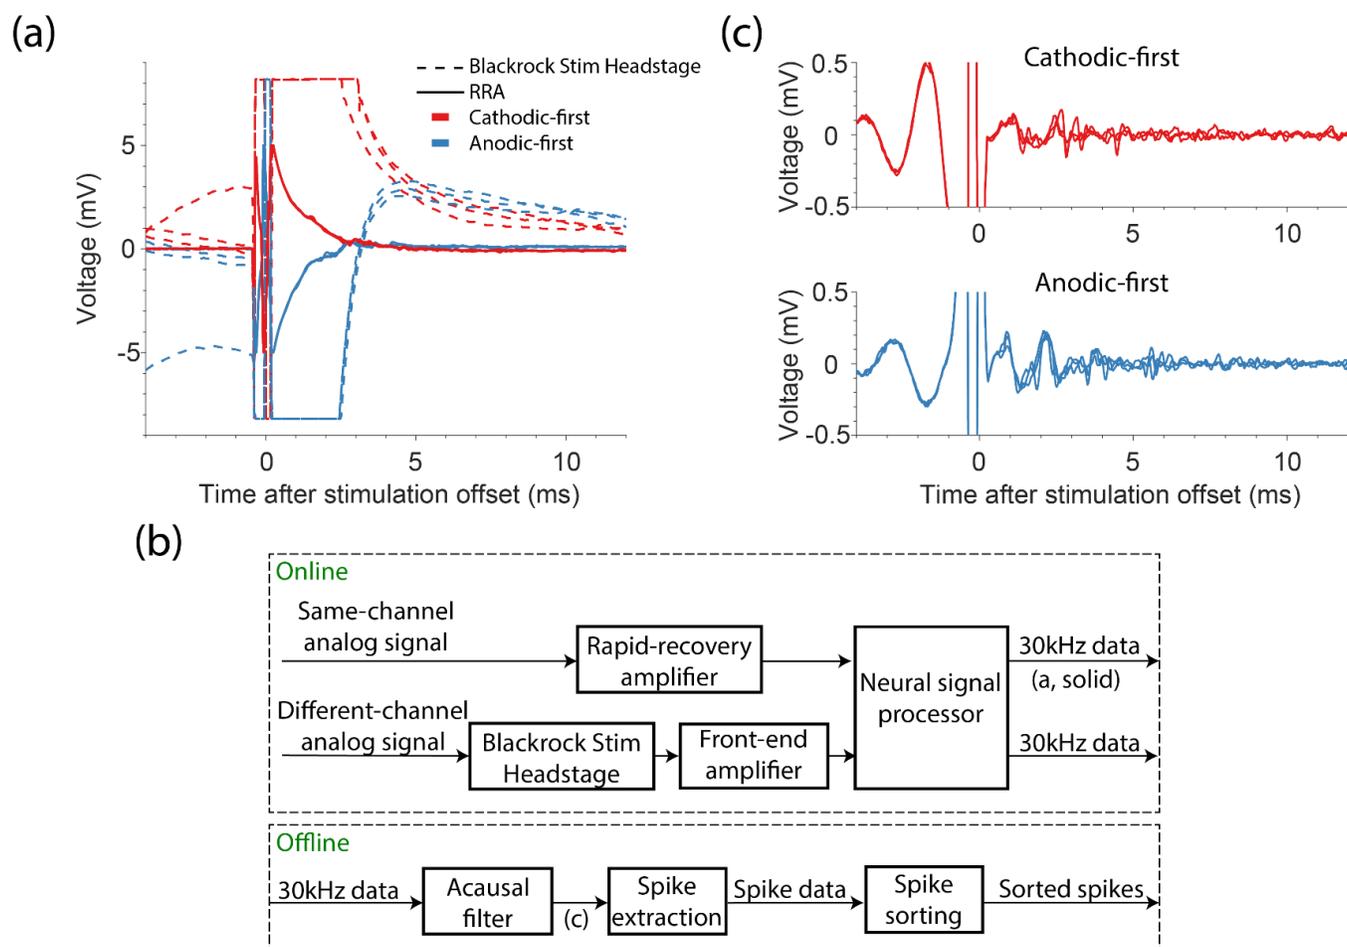


Fig. 2.1. Overview of artifact reduction pipeline. (a) Example recordings from the stimulated channel are shown when recording with the Blackrock Stim Headstage and Front-end amplifier (dashed lines) and the rapid recovery amplifier (RRA; solid lines). We stimulated with anodic-first (blue) or cathodic-first (red) biphasic pulses with phase duration of $200 \mu\text{s}$, phases separated by $53 \mu\text{s}$, and with an amplitude of $50 \mu\text{A}$. (b) Block diagram depicting the artifact reduction pipeline. The rapid-recovery amplifier receives signals and passes them to the Blackrock Neural Signal Processor. Signals from channels that were not stimulated were amplified by the Blackrock Stim Headstage and Front-End amplifier. All signals were sampled at 30 kHz and filtered offline. After filtering, we extracted spikes via threshold crossings and then sorted the spike data. (c) Voltages recorded using the RRA after acausal time-reversed high-pass filtering. Traces in (c) correspond to those in (a).

To measure the progressive gain recovery of the RRA after stimulation when stimulating and recording on the same electrode, we monitored the size of the artifact evoked by much lower current stimulation on a remote electrode. We tested gain recovery following alternating cathodic- and anodic-first biphasic pulses at 10 Hz, with amplitudes of 5–30 μA in 5 μA steps and 40–100 μA in 10 μA steps. We tested 25 stimulation electrodes across the two monkeys and delivered 32 ± 2 (mean \pm sd) pulses per condition. The remote channel was stimulated at 3000 Hz for 4.5 ms, with cathodic-first biphasic pulses (53 μs pulse length with 53 μs between phases). We used 1 μA to monitor gain recovery on four stimulation electrodes in one session, and 5 μA on the remote channel in later sessions.

Even with the RRA, full recovery to baseline took ~ 3 ms (Fig. 2.1a, solid lines). While a high-pass filter removed this drift, ringing caused by filtering the large artifact prevented neural recording for ~ 10 ms. Instead, we applied a 500 Hz high-pass Butterworth filter acausally, backwards in time, thereby preventing the introduction of a ringing artifact (Fig. 2.1c). We adjusted the timestamps of recorded spikes to account for the ~ 100 μs phase shift caused by filtering. Even with this acausal filtering we avoided filtering through the artifact, which would have obscured the pre-stimulus data (as seen in Fig. 2.1c). To account for the changing gain of the RRA, we divided the recorded signal by the measured gain recovery. After filtering, we extracted neural activity by finding threshold crossings and then sorting single units using OfflineSorter (Plexon Inc., Dallas, TX).

Recordings on non-stimulated electrodes using the Stim Headstage and Front-end Amplifier were saturated for ~ 0.7 ms after stimulation offset. In our testing, the RRA did not shorten the recording latency on non-stimulated electrodes. Because of this, we did not use the

RRA when recording on non-stimulated channels. Nevertheless, we filtered acausally before extracting neural activity as we did for recordings made on the stimulated electrode.

Evaluating the performance of the RRA and acausal, time-reversed filtering

To evaluate the performance of our pipeline for recording neural activity on the stimulated electrode, we tested how well we could recover simulated activity, added artificially at different latencies after stimulation. For this simulation, we recorded in the presence of artifacts with either the RRA or the Stim Headstage and Front-End amplifier on 10 representative stimulation electrodes. To simulate neural activity, we recorded naturally occurring spike waveforms during a period without stimulation, which we added to the recordings at random times after stimulation. We added spike waveforms at random latencies between 0.2–7 ms following 50% of the stimuli for each of the 10 electrodes. For each electrode, we generated 200,000 stimulation artifacts, half from recordings made with the RRA and half with the standard Blackrock hardware. We tested the same amplitudes described above for measuring RRA gain but used only cathodic-first pulses since our subsequent experiments used this polarity. We computed the percentage of spikes recovered by comparing the time stamps of recovered spikes to the artificial ones, tolerating ± 0.33 ms of error.

Stimulation protocol for characterizing the evoked response

After evaluating our recording capability, we characterized the response evoked on the stimulated channel by single pulses or pulse trains. Table 2.1 shows the numbers of sessions for each monkey, electrodes tested, neurons recorded, and trains per condition for all experiments. The final column (inter-train period) indicates the time between successive stimulation conditions. We slightly jittered the inter-train period for each condition to prevent synchronizing stimulation

with any physiological process by adding 0–100 ms sampled from a uniform distribution. We measured the progressive gain recovery of the RRA at each stimulation amplitude for each stimulation electrode.

In initial experiments we measured the response evoked by single pulses at a range of amplitudes typically used in BMIs (Rajan et al. 2015; Flesher et al. 2019; Salas et al. 2018). In four experiments, we tested 10–60 μA in 10 μA steps, 80 and 100 μA . Later, we probed the lower stimulation amplitudes more thoroughly using 10–30 μA in 5 μA steps and 40, 50, and 100 μA for another 4 electrodes, then added a 5 μA condition for the final 21 electrodes.

We next characterized the responses to short (~ 0.2 s) trains. We stimulated at 50 μA and at 20, 49, or 94 Hz for 7 channels, from which we recorded 7 neurons. After noticing a modest decay in the neuronal responses throughout the 0.2 s train at 94 Hz, we added a 179 Hz condition for the remaining 12 channels, from which we recorded the remaining 12 neurons. After short, high frequency trains, we observed rebound excitation, which we analyzed with this data.

We then characterized the evoked response to longer (~ 4 s) trains of stimulation, a duration that approximates that required for a BMI user to grasp an object (Flesher et al. 2021). Because the recorded neural response decayed rapidly with 179 Hz stimulation, we used a maximum of 131 Hz when stimulating with 4-s long trains. We stimulated with all combinations of 51, 80, 104, 131 Hz and 20, 40, 60 μA , amplitudes, these ranges chosen to evoke robust responses. While delivering continuously varying sensory feedback will require trains with time-varying amplitude or frequency, testing this full parameter space would not be feasible. Thus, we used trains with constant stimulation parameters to partially sample the space. Data were collected simultaneously on the stimulated and non-stimulated channels during this experiment. The results for non-stimulated channels may include a given neuron activated by different stimulation electrodes.

Table 2.1. Experimental parameters are shown for the single pulse and short train experiments, and the continuous (Cont.) long train experiments when recording on either the stimulated channel (stim) or non-stimulated channels (nonstim). The numbers of sessions for monkey H and monkey D are denoted with ‘H’ and ‘D’ respectively.

	# sessions	# stimulation electrodes	# neurons	# pulses or trains (mean \pm std.)	Inter-train period (s)
Single Pulse	4H; 4D	29	30	82.6 \pm 14.4	0.5
Short Train	7H; 5D	19	19	253 \pm 17.2	0.5
Cont. Long Train (stim)	6H; 4D	24	25	8 \pm 0	20
Cont. Long Train (nonstim)	6H; 4D	24	437	8 \pm 0	20

Data analysis

All data analysis was performed using MATLAB (MathWorks Inc., Natick, MA). To quantify the amount of activity evoked by each pulse, we counted spikes between 0.5 and 5.0 ms after the offset of each pulse and averaged across pulses. To account for different baseline firing rates across neurons, we subtracted the expected number of spontaneous spikes based on the baseline firing rate measured 10 ms to 80 ms before onset of single pulses or 0.2 to 2 s before train onset.

We computed an activation threshold for each neuron in response to single pulses. To do so, we measured the proportion of stimulation pulses with at least one spike occurring 0.5–5 ms after stimulation offset for each condition and neuron. We defined the activation threshold as the smallest amplitude at which the proportion of trials with a spike was significantly larger than that expected based on the baseline firing rate (Chi-Square test, $\alpha < 0.05$). We determined if a neuron was responsive to long trains of stimulation in a similar manner. Since the evoked response decayed throughout long trains, we considered only the first 20 pulses in each train. For each condition, neurons with significantly more spikes than chance (Chi-Square test, $\alpha < 0.05$) were considered responsive.

Multiple spikes were typically evoked at consistent latencies by single stimulus pulses. We grouped spikes based on their response latency across trials for each neuron and condition. To do so, we computed a firing rate for the spikes evoked after single pulses by convolving them with a non-causal Gaussian kernel of width equal to a standard deviation of 0.2 ms, which we then averaged across pulses. We found peaks in this average with MATLAB's *findpeaks* algorithm. This algorithm uses “prominence”, the height of a peak and its location relative to other peaks, to measure how much a peak stands out. We required peaks to have a minimum prominence of 1.0

and to be separated by at least 0.7 ms. This algorithm also computes the width at half maximum of each peak. Spikes that occurred within the width of each peak were included in the corresponding group. We measured the latency of each peak and computed the standard deviation of the spike times within each group. Our results were only slightly affected by small changes to the smoothing kernel width, minimum peak spacing, and minimum prominence.

After an evoked response, many neurons underwent either long-lasting inhibition or rebound excitation, which we quantified by computing the average firing rate across trials using a two-bin running average across 5 ms bins. We defined an inhibitory response as firing rates below three-quarters of the mean baseline firing rate for two consecutive bins (a similar threshold as (Butovas and Schwarz 2003)) and measured the time the firing rate remained below this threshold. We defined a rebound excitatory response if two consecutive bins exceeded twice the mean baseline firing rate and the corresponding duration.

For many neurons, the evoked response decreased throughout long trains of stimulation. We measured the decay rate for each responsive neuron. To do so, we measured the mean firing rate in 50 ms bins from 0.0 to 3.9 s after train onset, excluding the initial 1 ms after each pulse, which was obscured by the stimulus artifact. We then fit the firing rate with an exponential decaying function,

$$a * e^{(-b*x)},$$

with a as the intercept and b as the decay rate. A large decay rate indicates the response decayed rapidly, while a decay rate near zero indicates the response was maintained throughout the train.

Statistical Analysis

Statistical analyses were performed using MATLAB (MathWorks Inc., Natick, MA). We used linear and logistic models to analyze many of our results. We included two interaction terms in the model when analyzing the effect of amplitude, time, and amplifier on the proportion of simulated spikes recovered: one between amplitude and amplifier, to test whether the effect of amplitude was reduced with the RRA, and a second between time and amplifier, to see if the rate of spike recovery increased with the RRA. When analyzing the effect of amplitude on the latency of evoked spikes, we included an interaction term between amplitude and spike group number. Finally, we included an interaction term between amplitude and frequency when analyzing the decay rate throughout long trains of stimulation. After fitting the models, we performed F-tests on the resulting parameters from the linear models and t-tests on the resulting parameters from the logistic models.

We performed Wilcoxon rank-sum tests to compare the magnitude of evoked activity recorded on non-stimulated channels at 20 μ A and 60 μ A for each neuron and stimulation electrode pair. Here, we aggregated data across stimulation frequencies.

Results

Recording pipeline performance

To evaluate the performance of the RRA, we first measured its dynamic gain recovery after stimulation at different amplitudes. We delivered a single biphasic pulse through the electrode to which the RRA was connected and simultaneously injected a known signal to a remote electrode (Fig. 2.2a). After acausal, time-reversed filtering, we determined the gain of the amplifier by dividing the amplitude of each pulse in the known signal by the mean amplitude of the final three

pulses, which were well after full gain recovery. The mean gain recovery curves, aggregated for 25 stimulation electrodes across two monkeys, are shown for both cathodic- and anodic-first pulses at several stimulation amplitudes in Fig. 2.2b. We compared the gain of the amplifier at 1 ms across stimulation amplitudes and polarities using a repeated measures ANOVA ($F(26,481) = 40.6$, $p = 6.58E-104$). The gain of the amplifier recovered more slowly as amplitude increased ($F(1,481) = 762.65$, $p \cong 0$) and roughly 140 μs faster for cathodic-first pulses than for anodic-first pulses ($F(1,481) = 142.2$, $p \cong 0$). Subsequently, when measuring actual neural signals, we accounted for the changing gain by dividing the recorded signal by the gain function (Fig. 2.2a, bottom).

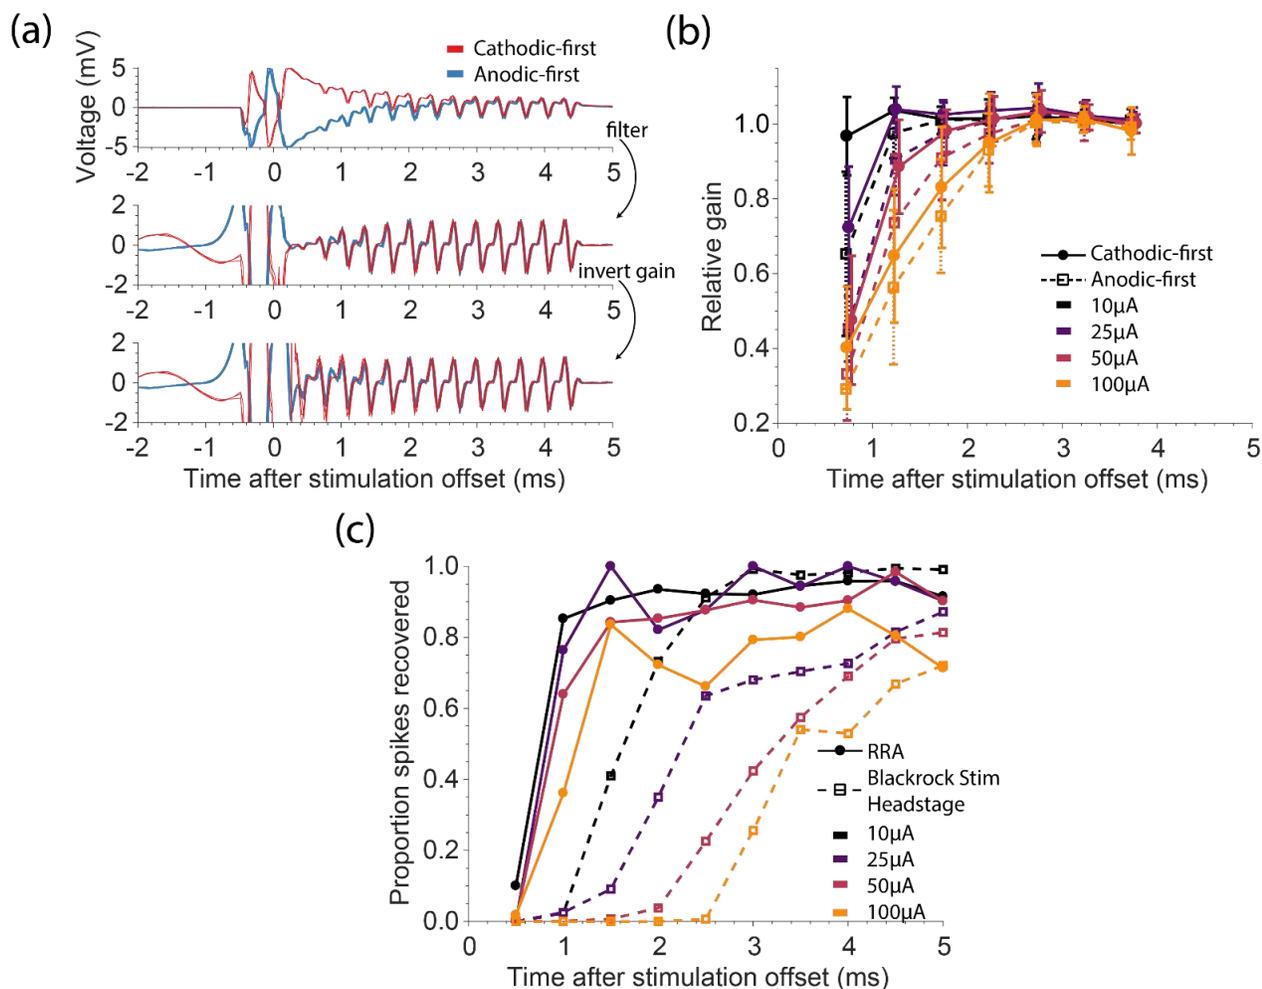


Fig. 2.2. Evaluation of rapid-recovery amplifier (RRA). (a) Example recordings on the stimulated channel when evaluating the gain of the RRA are shown both before (top) and after (middle) acausal, time-reversed filtering, and after accounting for the changing gain (bottom). We stimulated with biphasic anodic-first (blue) or cathodic-first (red) pulses with phase duration of 200 μ s, phases separated by 53 μ s, and an amplitude of 50 μ A. Pulses were simultaneously delivered on a remote channel to inject a ‘known’ signal. (b) The relative gain of the RRA for stimulation at different amplitudes. The gain was determined by measuring the peak-to-peak voltage of the injected signal. Error bars denote standard deviation across electrodes ($n = 25$). (c) Spikes were artificially added to artifact traces recorded on the stimulated channel. The proportion of simulated spikes recovered using our pipeline for both the RRA (solid lines) and the Blackrock Stim Headstage (dashed lines) across stimulation amplitudes ($n = 10$).

We tested the ability to recover spikes following stimulation by adding representative, naturally occurring spike waveforms to actual recordings of stimulation artifacts to establish a ground-truth reference. The proportion of these spikes that could be recovered with the Blackrock Headstage and with the RRA are shown in Fig. 2.2c. We used logistic regression to predict the proportion of spikes recovered based on the stimulation amplitude and time after stimulation, (overall model $\chi^2 = 1.97 \times 10^3$, $p \cong 0$). Not surprisingly, spike recovery worsened with increasing stimulation amplitude regardless of amplifier ($p \cong 0$, t-test), but spikes were recovered at much shorter latencies with the RRA than with the Blackrock Stim Headstage ($p \cong 0$, t-test). The RRA also reduced the effect of amplitude ($p = 0.0015$, t-test) and increased the recovery rate ($p = 0.0090$, t-test).

Excitatory and inhibitory response to single pulses of ICMS

After evaluating the performance of the RRA, we used it for a series of experiments to quantify the neural responses evoked on the stimulated electrode. We first characterized the excitatory and inhibitory responses to single stimulus pulses across a wide range of current amplitudes (5–100 μA). Example raw and acausal filtered spikes for action potentials recorded at least 100 ms (top) and 1–3 ms after (bottom) stimulation offset are shown in Fig. 2.3a. The shape of filtered spikes recorded shortly after stimulation was similar to those recorded long after stimulation offset. Responses for this example neuron are shown in Fig. 2.3b. While it was not possible to record throughout stimulation (red shading indicates region obscured by the artifact), using the RRA allowed us to record many spikes that we could not have seen if we had used the Blackrock Headstage (grey shading). To quantify the amount of evoked activity, we measured the number of spikes evoked for each amplitude and subtracted the expected number of spikes due to

baseline firing. The number of spikes evoked above baseline firing across amplitudes is shown in Fig. 2.3c. The number increased significantly as amplitude increased (overall model $F(30,223) = 4.88$, $p \cong 0$; amplitude factor $F(1,223) = 12.029$, $p = 6.3 \times 10^{-4}$). Among the 29 out of 30 neurons that were activated with the range of currents tested, the median activation threshold was $10 \mu\text{A}$ (Fig. 2.3c).

Sufficiently high stimulation amplitude evoked multiple spikes within 10 ms of stimulation offset. These spikes occurred at consistent latencies across trials, with later spikes having more varied timing than earlier ones. To quantify this, we grouped evoked spikes based on their latency (Fig. 2.3b and Supplementary Materials show example groups). Fig. 2.3e shows the standard deviation of spike times within a group compared to the latency of that group for multiple stimulation amplitudes. This standard deviation increased significantly as group latency increased (overall model $F(32,302) = 103$, $p \cong 0$; latency factor $F(1,302) = 574.13$, $p \cong 0$). We also noticed that latencies decreased as stimulation amplitude increased, seen as a leftward shift in Fig. 2.3a as current increased to $100 \mu\text{A}$, at which point the artifact likely obscured the first recorded group of evoked spikes. Using a linear model across all neurons, we determined that the latency of groups decreased by $3.6 \pm 0.7 \mu\text{s}/\mu\text{A}$ as amplitude increased (overall model $F(31,303) = 55.1$, $p \cong 0$; amplitude factor $F(1,303) = 25.497$, $p = 7.7 \times 10^{-7}$).

The percentage of neurons that responded with different numbers of spike groups for stimulation at various amplitudes is shown in Fig. 2.3f. For our example neuron, stimulation at $100 \mu\text{A}$ appeared to evoke 4 groups of spikes, as did 7% of neurons we recorded. However, at $100 \mu\text{A}$, the extended artifact and decreased latency likely obscured the entire initial group of spikes, as can be seen in panel b. When we determined that this occurred, we increased the number of

groups for the corresponding neuron by one (increasing the example neuron's group count from 4 to 5 in Fig. 2.3f). Even without this compensation, the number of groups increased significantly with stimulation amplitude (overall model: $F(20,223) = 5.65$, $p \cong 0$; amplitude factor: $F(1,223) = 85.5$, $p \cong 0$).

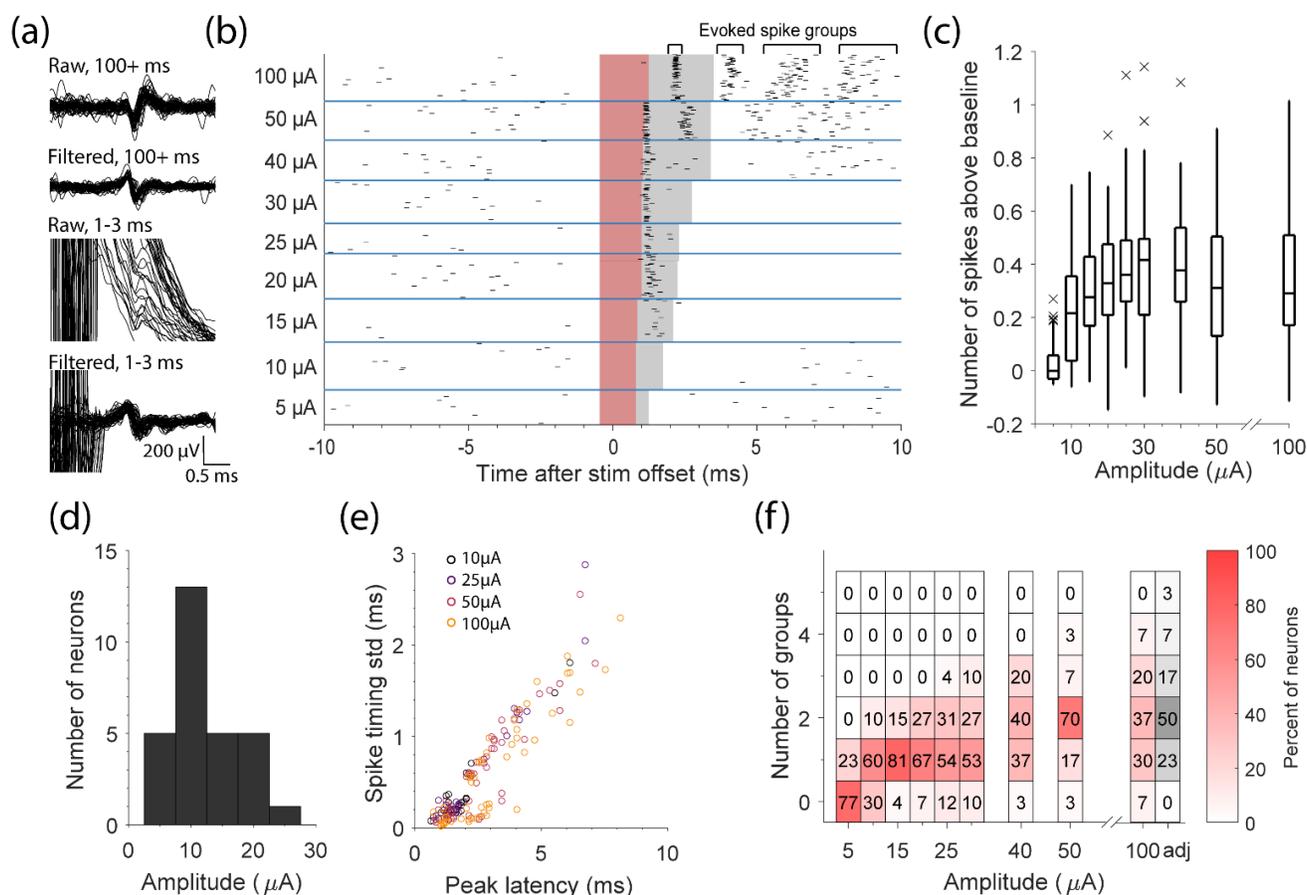


Fig. 2.3. Excitatory response on the stimulated channel after single pulses of stimulation. (a) Example spikes recorded on the stimulated electrode at least 100 ms after stimulation offset before and after acausal filtering. Spikes from the same neuron recorded 1-3 ms after stimulation offset before and after acausal filtering. (b) Response of the neuron in (a) to single cathodic-first pulses at different amplitudes. Each row is a different stimulation trial (728 total), and each tick represents an action potential from this neuron. Blue, horizontal lines separate stimulation trials at different amplitudes. Red shading depicts the time interval in which we were unable to record neural signal with the RRA. Grey shading depicts the corresponding time had we used the Blackrock Stim Headstage and Front-End amplifier. (c) The number of evoked spikes above baseline is shown across neurons ($n = 30$) for each stimulation amplitude. The box represents the 25th, 50th and 75th percentiles, whiskers represent the extent of the data, and X's mark outliers. (d) Distribution of activation thresholds across neurons. (e) The standard deviation of spike times within an evoked spike group is shown against the latency of that group for different stimulation amplitudes. (f) The number of groups evoked for different stimulation amplitudes. The number within each box and the shading of each box indicates the percentage of neurons. The prolonged artifact that occurred when stimulating with 100 μA likely obscured the entire initial group of spikes. When we determined that this occurred, we increased the group number by one (displayed in grey; 'adj').

Neuronal activity was typically suppressed anywhere from 10 to 150 ms after single pulses, depending on the stimulation amplitude (Fig. 2.4a). We fit a linear model to predict inhibition duration by amplitude across neurons (overall model $F(30,180) = 1.9$, $p=0.0057$). We found that increasing stimulation amplitude significantly increased the inhibition duration (amplitude factor $F(1,180) = 32.43$, $p = 5.0 \times 10^{-8}$) and increased the fraction of cells undergoing inhibition (Fig. 2.4b). Stimulation amplitudes $\geq 40 \mu\text{A}$ caused inhibition in $\sim 90\%$ of neurons.

Temporal response to trains of ICMS

We hypothesized that the activity evoked by ICMS would decrease throughout long stimulus trains as a consequence of the long-lasting inhibition on stimulated electrodes following single pulses (Fig. 2.4). To test this, we stimulated on single electrodes with 4-s long trains at several amplitudes (20, 40, 60 μA) and frequencies (51, 80, 104, 131 Hz). Example spikes recorded from a neuron on the stimulated electrode are shown in Fig. 2.5a (grey) using the same format as Fig. 2.3a. The mean responses across eight trains for nine of the 12 stimulation conditions are shown as grey traces in Fig. 2.5b for this example neuron. For this neuron, the evoked response rapidly decayed throughout the train, particularly for the larger amplitudes and frequencies. For the 21.5 ± 2.0 neurons that were activated significantly for each condition (Chi-Square test, $\alpha < 0.05$), we computed a decay rate by fitting the firing rate during stimulation with an exponential (Fig. 2.5c). Using a linear model ($F(26,231) = 14.7$, $p \cong 0$), we determined that the evoked response decayed significantly faster with greater stimulation amplitude or frequency (amplitude: $F(1,231) = 119$, $p \cong 0$; frequency: $F(1,231) = 134$, $p \cong 0$). Increased frequency (amplitude) had a larger effect at higher amplitudes (frequencies) (interaction term: $F(1,231) = 71.4$, $p \cong 0$).

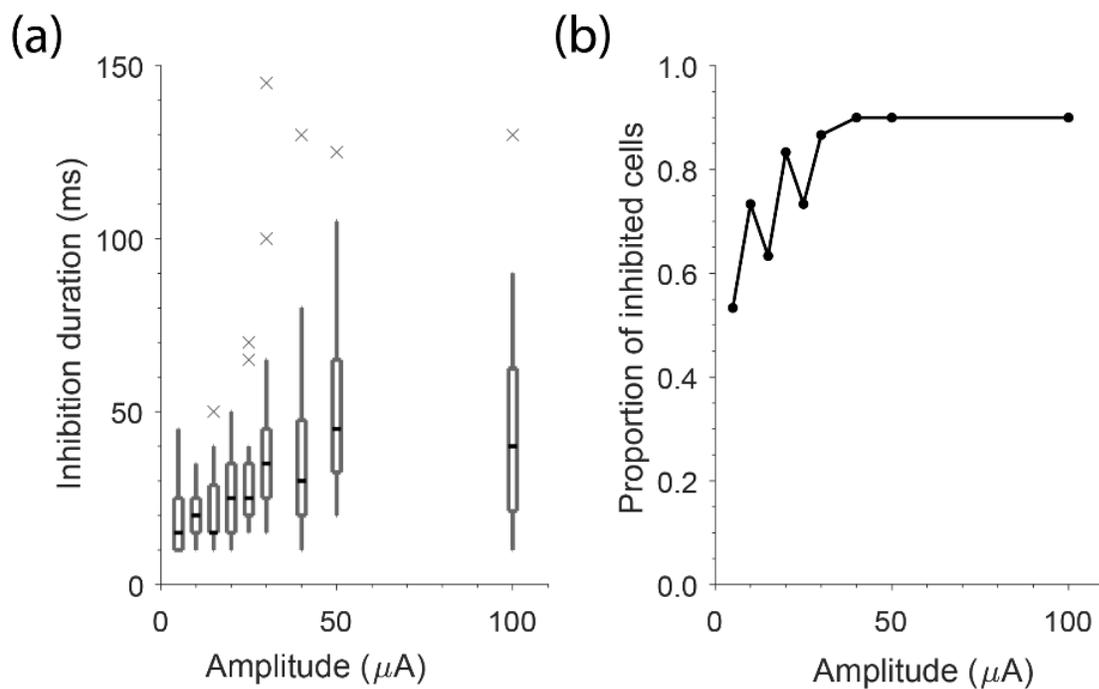


Fig. 2.4. Inhibitory response recorded on the stimulated channel after single pulses of stimulation. (a) The inhibition duration across neurons ($n = 30$) recorded on the stimulated channel after single cathodic-first pulses of stimulation across stimulation amplitudes. (b) The fraction of cells with an inhibitory response is shown for each stimulation amplitude.

If neurons recorded on non-stimulated electrodes were driven transsynaptically by neurons activated near the stimulated electrode, then we would expect to see a similar rapid decay in the evoked activity for neurons on non-stimulated electrodes. If, on the other hand, neurons even on distant electrodes are driven directly, their decay rate may differ from that of neurons recorded on the stimulated electrode. To determine this, we examined the neuronal activity evoked on non-stimulated electrodes. Example spikes are shown for an example neuron in response to same- (gray traces) and different- (black traces) channel stimulation (Fig. 2.5a). In contrast to its response on the stimulated electrode, the activity of this neuron did not decay appreciably when a different electrode was stimulated (Fig. 2.5b, black traces). The 260 ± 90 neurons recorded on non-stimulated electrodes that were activated by stimulation (Chi-Square test, $\alpha < 0.5$) all had maintained responses throughout the stimulation train, as summarized in Fig. 2.5d. Using a linear model with data aggregated across amplitudes and frequencies (overall model: $F(81,3269) = 51.6$, $p \cong 0$), we determined that the evoked response decayed significantly faster for neurons recorded on the stimulated channel than on non-stimulated channels (stimulated channel factor: $F(1,3269) = 980.3$, $p \cong 0$). These results imply that the response on non-stimulated electrodes is driven directly, or by evoked activity that occurs before we can record it.

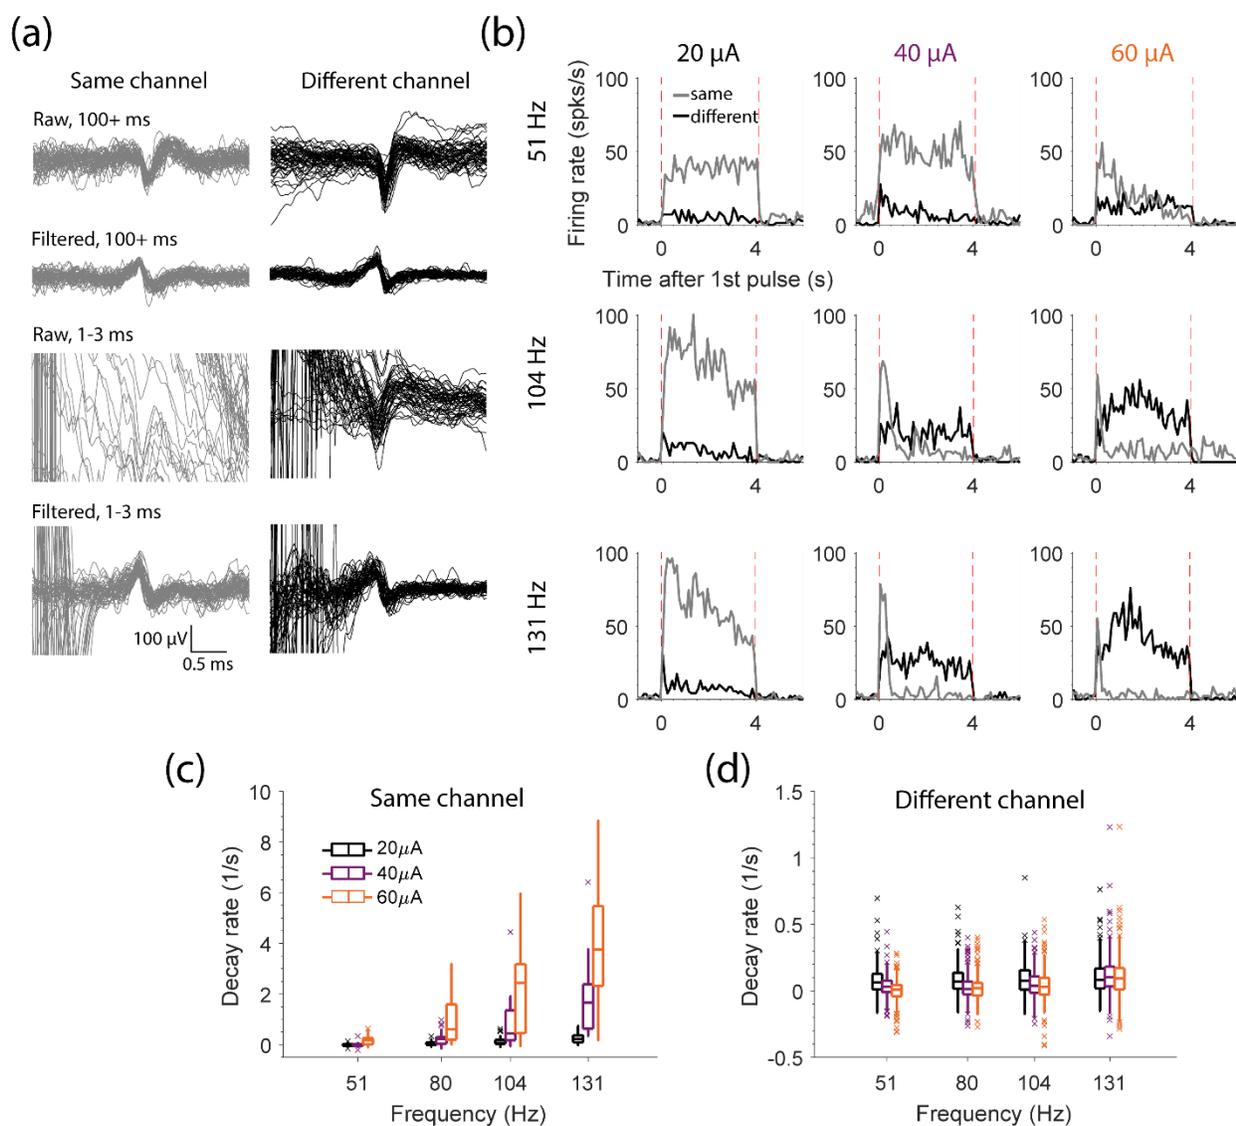


Fig. 2.5. Evoked response throughout 4-s long trains of stimulation. (a) Example spikes for a neuron when the channel it was recorded on was stimulated (grey) and when a different channel was stimulated (black) in the same format as Fig. 2.3a. (b) The mean firing rate across stimulation trials for the same neuron when the channel it was recorded on was stimulated (grey) and when a different channel was stimulated (black) for different stimulation amplitudes (columns) and frequencies (rows). Amplitudes and frequencies are noted above and to the right of the panels, respectively. Vertical, red dashed lines indicate train onset and offset. (c) The decay rates across neurons recorded on the stimulated channel for each amplitude and frequency ($n = 21.5 \pm 2.0$ across conditions). (d) The decay rates for each neuron recorded on non-stimulated channels (for each amplitude and frequency ($n = 258 \pm 86$ across conditions). Note the smaller y-limits in (d) compared to (c).

After the end of an ICMS train, we expected neurons on the stimulated electrode to be inhibited for many milliseconds, as we observed with single pulses (Fig. 2.4). Indeed, low-frequency, 50 μ A trains delivered for ~ 0.2 s caused inhibition (see example in Fig. 2.6a) in about 50% of neurons, lasting from 10-250 ms (Fig. 2.6b). Faster stimulus frequency increased inhibition duration (Model: $F(19,21) = 6.38$, $p = 5.9 \times 10^{-5}$, frequency factor $F(1,21) = 36.0$, $p = 6.0 \times 10^{-6}$) but this effect was not observed in all 16 tested neurons. At 179 Hz, the highest frequency we tested, the fraction of cells with an inhibitory response was only $\sim 8\%$. Instead of inhibition in these cases, we observed a large burst of activity immediately after the stimulation train. This rebound excitation occurred for 75% of cells following stimulation at 179 Hz and lasted from ~ 25 -240 ms (Fig. 2.6c). If a neuron exhibited rebound excitation for multiple stimulation frequencies, higher frequencies almost always resulted in longer lasting rebound. During the longer 4-s trains, we observed rebound excitation very infrequently (2/25 cells) potentially because of the longer train duration.

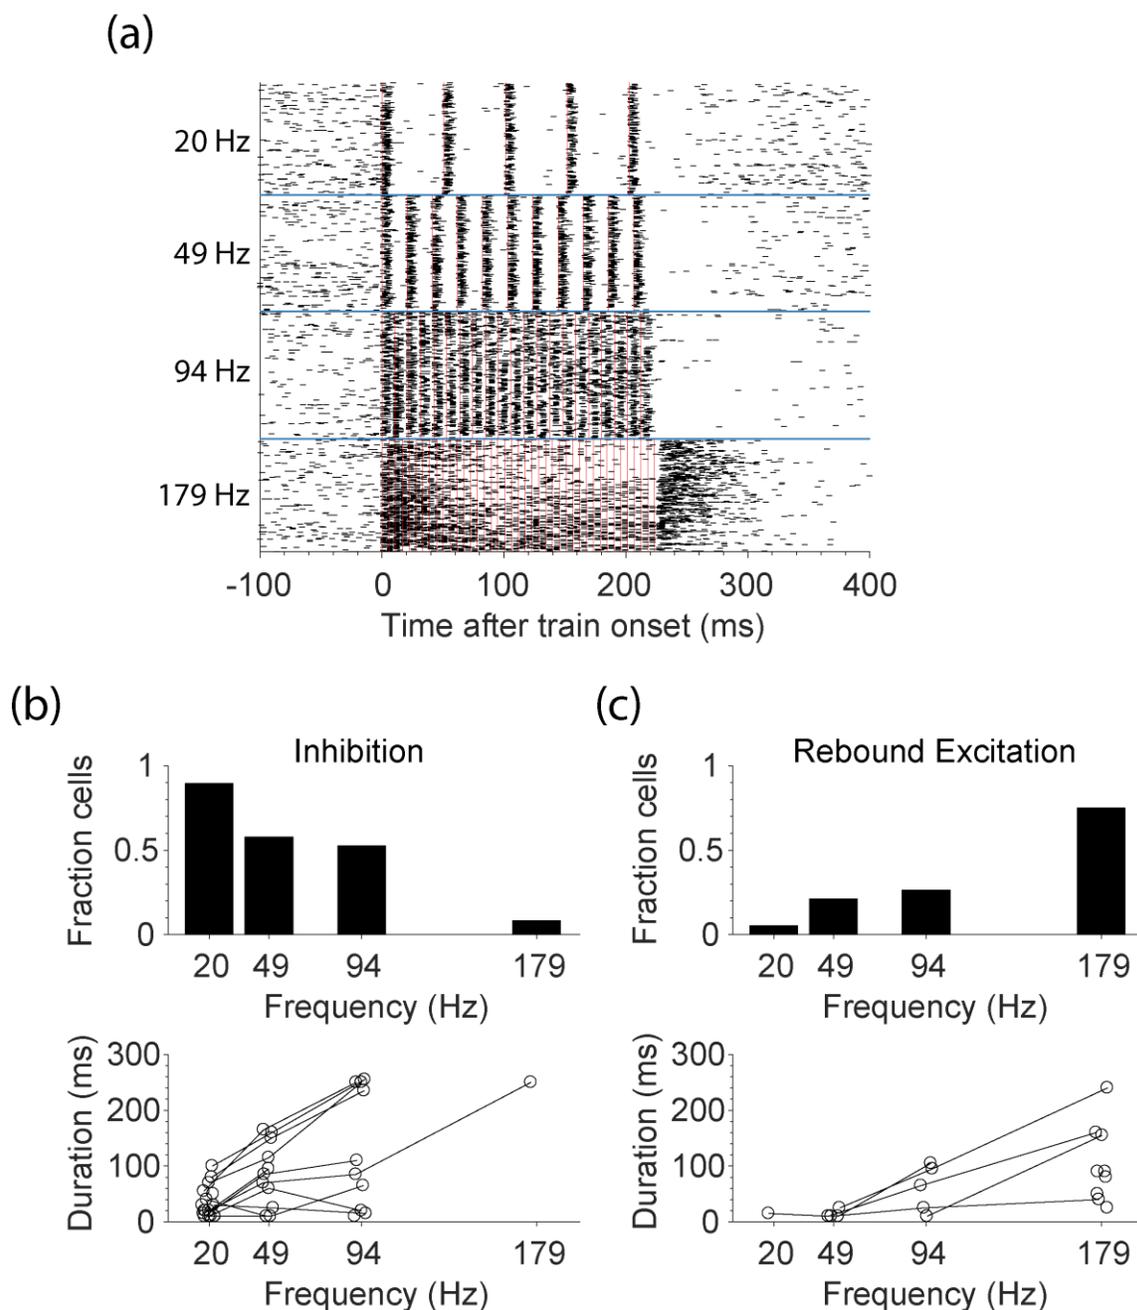


Fig. 2.6. Rebound excitation recorded on the stimulated electrode after short trains of stimulation. (a) The response of an example neuron recorded on the stimulated electrode during ~ 200 ms trains at different frequencies. Red lines indicate stimulation pulses. Stimulation frequencies are shown on the left of the figure for $50 \mu\text{A}$ stimulation. (b) The fraction of cells ($n = 19$ for 20, 49, and 94 Hz; $n = 12$ for 179 Hz) that displayed an inhibitory response after the end of the short trains (top) and the duration of the inhibitory responses (bottom) for each frequency. (c) The fraction of cells that displayed rebound excitation (top) and the duration of the rebound excitation (bottom) for each frequency. Lines connecting points represent data from the same neuron.

Spatial pattern of the response to ICMS trains

Both increased amplitude and frequency typically increase ICMS detectability, perhaps because of increased charge delivery (Kim, Callier, Tabot, Gaunt, et al. 2015; Otto, Rousche, and Kipke 2005; Sombeck and Miller 2019). Increasing amplitude leads both to more activity near the stimulated electrode (Fig. 2.3) as well as a wider spread of activity recorded across a multi-electrode array (Hao, Riehle, and Brochier 2016; Stoney, Thompson, and Asanuma 1968; Kumaravelu et al. 2021), likely because increased amplitude results in more charge delivered per pulse. Greater frequency, though, does not change the charge per pulse and thus may not lead to equivalent effects. To study these effects, we measured activity on non-stimulated electrodes throughout 4-s trains of continuous stimulation. We computed the mean firing rate above baseline for each neuron and amplitude / frequency combination across 8 trains. Fig. 2.7a shows the mean firing rate above baseline for each stimulation electrode aggregated across two monkeys against distance from the stimulated electrode. For each stimulation electrode, we only analyzed neurons that had activation thresholds at or below 20 μA when stimulating at 51 Hz. The evoked activity per pulse at 60 μA was significantly larger than that at 20 μA for 290 out of 437 neurons ($p < 0.001$, Wilcoxon rank-sum test) (Fig. 2.7b). Using a linear model (overall model: $F(125,1362) = 19.4$, $p \cong 0$), we determined that increasing amplitude and frequency increased the evoked firing rate (amplitude factor: $F(1,1362) = 674$, $p \cong 0$, frequency factor: $F(1,1362) = 472$, $p \cong 0$).

We wondered whether the effect of frequency was simply due to the different number of stimulation pulses in the train. To analyze this, we normalized firing rates by the number of pulses and repeated our statistics (overall model: $F(125,1362) = 20.9$, $p \cong 0$). Increasing frequency no

longer significantly increased the evoked activity per pulse (frequency factor: $F(1,1362) = 0.81$, $p = 0.37$). The apparent effect of frequency was only due to the greater number of pulses in the train.

Neural discharge typically has a fixed FANO factor, meaning that the variance of the firing rate increases with the mean rate (Softky and Koch 1993; Tolhurst, Movshon, and Dean 1983). We wondered whether the variance of ICMS-induced neural activity also increases with amplitude and frequency, along with increases in mean rate (Fig. 2.7b). To test this, we measured the variance in firing rate across trains for each neuron and each condition (Fig. 2.7c). We observed no appreciable change in the variance with increasing amplitude and only a slight increase with frequency. To quantify these effects, we used a linear model to determine the effect of amplitude and frequency on the variance in firing rate (overall model: $F(126,1362) = 7.21$, $p \cong 0$). While increasing amplitude significantly decreased variance (amplitude factor: $F(1,1362) = 34.0$, $p = 6.8 \times 10^{-9}$) and increasing frequency increased variance ($F(1,1362) = 34.4$, $p = 5.7 \times 10^{-9}$), the effect sizes for both effects were tiny. This implies that ICMS-evoked activity does not have a fixed FANO factor; increasing either amplitude or frequency increases the mean rate without an equivalent effect on variance.

We also hypothesized that increased stimulation amplitude would increase the distance at which neurons are activated while increased frequency would not. Data for a subset of stimulation conditions are shown in Fig. 2.7d. We used logistic regression to determine the effect of amplitude, frequency, and distance on the proportion of activated neurons (overall model $\chi^2(104) = 1.69 \times 10^3$, $p \cong 0$). While increasing either amplitude or frequency increased the proportion of activated neurons (amplitude: $p \cong 0$; frequency: $p = 1.6 \times 10^{-9}$), the effect of frequency was an order of magnitude smaller.

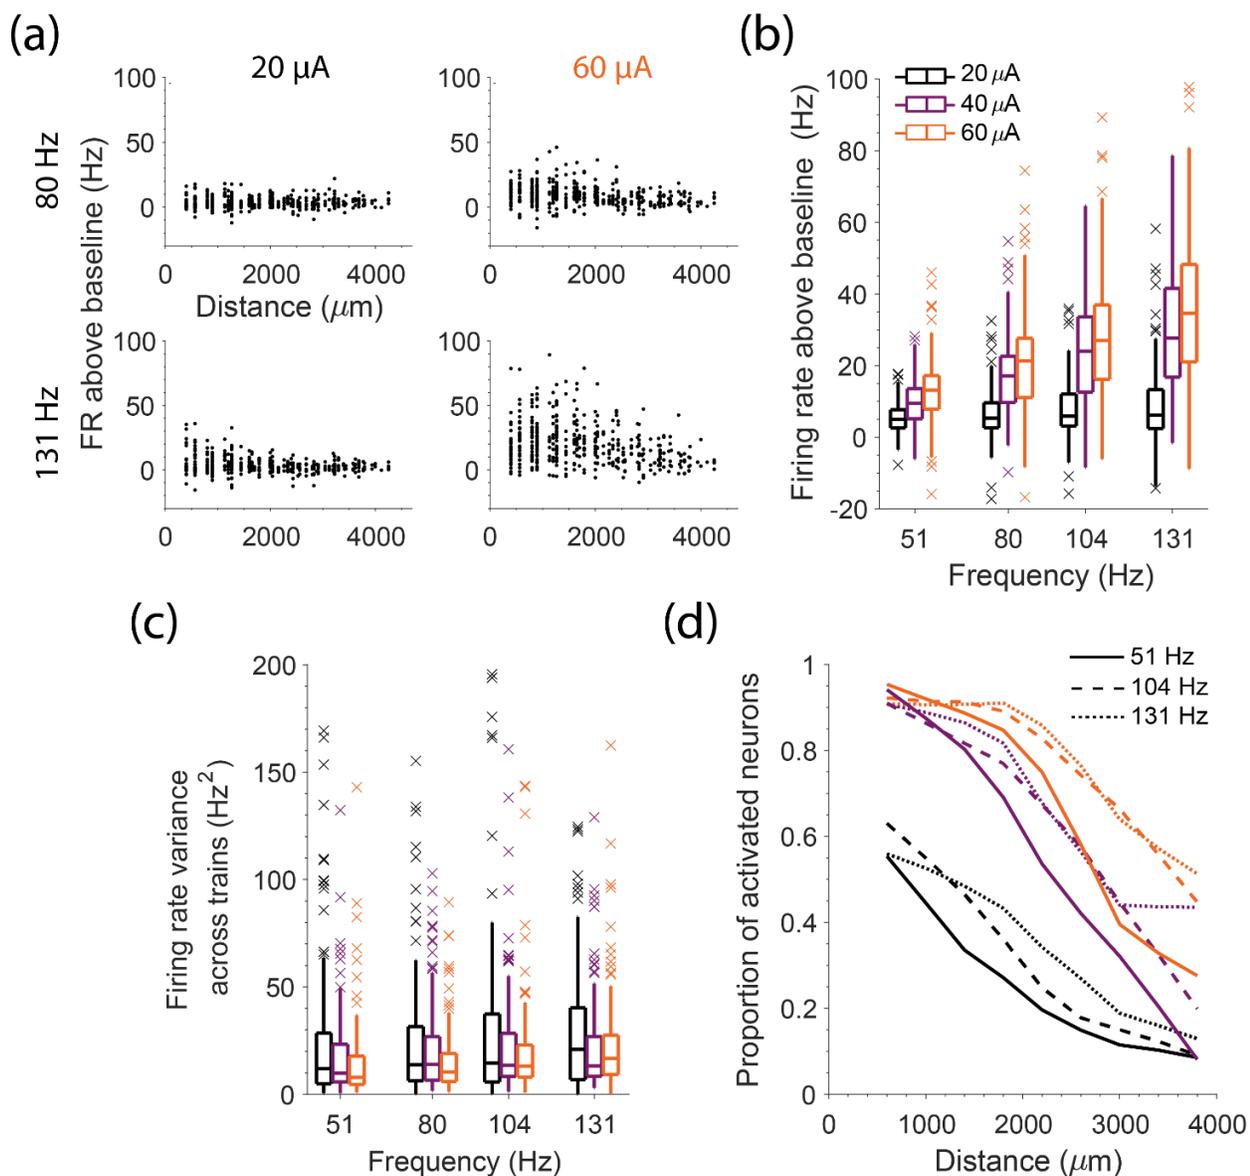


Fig. 2.7. Evoked response on non-stimulated channels during 4-s long trains of ICMS. (a) The firing rate above baseline against distance from the stimulated electrode for different amplitudes (columns) and frequencies (rows). Each point represents a neuron and stimulated electrode pair ($n = 258.2 \pm 86.1$ across conditions). (b) The firing rate above baseline for each frequency and amplitude condition for responsive neurons. (c) The variance of the firing rate across trains for each condition. (d) The proportion of neurons activated at different distances is shown for a subset of amplitudes (color) and frequencies (line-style) (437 total neurons).

Discussion

We developed hardware and software tools to enable recording at short latency after ICMS. With these tools, we were able to record roughly 0.7 ms after the end of stimulation, even from the stimulated channel. We investigated the evoked response to single pulses, short trains, and long trains of ICMS of varying amplitude and frequency to better understand the neural response to stimulation. Here, we compare our methods and results to those of previous studies, discuss the mode of activation for the spikes we recorded, and how our results may impact the design of biomimetic stimulation patterns in afferent interfaces.

Comparison of artifact suppression to previous techniques

Recording neurophysiological potentials immediately after passing current through an electrode is difficult; the large shock artifact typically prevents recordings for many milliseconds. We developed and evaluated a rapid-recovery amplifier (RRA) to enable short latency recordings, particularly on the stimulated electrode. The RRA clamps the voltage below that which would saturate downstream electronics by reducing gain as the magnitude of the input voltage increases (Fig. 2.2b). An alternative approach to shorten the duration of the artifact is to electrically disconnect the recording system during stimulation (Zhou, Johnson, and Muller 2018). While this approach is effective on non-stimulated electrodes (Hao, Riehle, and Brochier 2016), it cannot remove artifact on the stimulated electrode, which is caused by residual polarization of the electrode itself (Venkatraman et al. 2008). Our approach is similar to clamping the slew rate (first derivative) of a signal, as has been done previously (Epstein 1995). By reducing the gain, we reduced the size of the artifact and prevented saturation, thereby allowing us to record at earlier

latencies. Another important advantage of the RRA is the wide input voltage range (± 15 V) that avoids input clamping and stimulus current shunting of the relatively high voltage (< 10 V) stimulus pulses. A benefit of our approach is that the RRA can be placed in front of pre-existing recording systems, in our case, the Cerebus system from Blackrock Neurotech. Saturation can also be prevented by using an amplifier with a lower gain and/or an amplifier with a higher maximum input voltage (Jung, Kim, and Nam 2018; Rolston, Gross, and Potter 2009).

While the RRA prevents amplifier saturation that would otherwise be caused by the large shock artifact, the recorded signal still returns slowly to baseline after stimulation (Fig. 2.1a). This slow return is likely caused by slow dissipation of the residual charge on the electrode (Zhou, Johnson, and Muller 2018). To remove excess charge more quickly, custom electronics could be designed to actively discharge the electrode to a pre-stimulus voltage (Brown et al. 2008; DeMichele and Troyk 2003; Freeman 1971), although this may introduce switching artifacts that diminish the effectiveness of this approach.

The slow return to baseline can also be removed offline. When done with a high-pass filter, it is important not to filter through the shock artifact, as this can cause ringing and obscure the neural signal. Some solutions include filtering the data beginning a fixed time after the end of stimulation (Hao, Riehle, and Brochier 2016) or blanking the signal and using a low-order filter to limit ringing (Weiss et al. 2018). Instead, we filtered acausally, backwards in time so that any ringing would be pushed before the stimulation, leaving the post-stimulus data clean (such acausally displaced ringing can be seen before 0 ms in Fig. 2.1c). This approach does not require defining a time at which the shock artifact has ended, though it does push neural signal back in time ~ 100 μ s. We compensated for this time shift by adjusting the time stamps of recorded spikes by 100 μ s. With the RRA and acausal, time-reversed filtering, we were able to record ~ 0.7 ms

after stimulation offset, even on the stimulated electrode (Fig. 2.2c), revealing spikes that we could not have recorded with the Blackrock Stim Headstage (grey shading in Fig. 2.3b).

Mode of activation of recorded spikes on stimulated and non-stimulated channels

ICMS can evoke action potentials both directly and transsynaptically (Tehovnik et al. 2006). Directly evoked spikes occur because stimulation changes the membrane potential of cells near the electrode, causing them to fire. Action potentials are typically initiated in axons, which have a higher density of sodium channels than do somas, resulting in lower activation thresholds (Nowak and Bullier 1998a, 1998b; Tehovnik et al. 2006). Action potentials then propagate antidromically to the cell bodies and orthodromically to presynaptic terminals, where they may elicit further activity transsynaptically.

We wondered whether the spikes we recorded on the stimulated electrode were evoked directly, at either the axon or soma, or transsynaptically. Since we have no direct way of testing this, we inferred the mode of activation from the latency of evoked spikes. When we calculated the latest these spikes could occur, we assumed that directly evoked spikes were generated at the end of the cathodic phase (Gustafsson and Jankowska 1976; Jankowska, Padel, and Tanaka 1975; Stoney, Thompson, and Asanuma 1968), although spikes may occur earlier in the stimulus pulse at higher amplitudes (such a shift is evident in Fig. 2.3b). We may actually observe these spikes somewhat later since they need to propagate from the site of initiation back to the soma. We estimated this potential antidromic distance and latency by first estimating how far spike initiation could have occurred from the stimulated electrode. To do so, we used Stoney's square-root relationship (Stoney, Thompson, and Asanuma 1968):

$$I = kr^2,$$

With $k = 1292 \mu\text{A}/\text{mm}^2$ and $I = 10 \mu\text{A}$, the median activation threshold of neurons in our study (Fig. 2.3c) the maximum spike initiation distance is $\sim 100 \mu\text{m}$. Since somas can be recorded up to $\sim 150 \mu\text{m}$ from the recording electrode (Maynard, Nordhausen, and Normann 1997), the maximum distance an action potential could travel before being recorded is $\sim 250 \mu\text{m}$. With a propagation speed of $1 \mu\text{m}/\mu\text{s}$ (Swadlow 1990), the maximum latency at which we expect to see a directly evoked spike is 0.25 ms after the end of the cathodic phase, (coincident with the end of the biphasic pulses). Hence, the earliest spikes we were able to see on the stimulated electrode, (0.7 ms after the end of the biphasic pulses; Fig. 2.3d), could not have been directly evoked.

Since the shortest synaptic delay is $\sim 0.4 \text{ ms}$ (Gustafsson and Jankowska 1976), we estimate transsynaptic spikes could occur at a latency as short as 0.4 ms after the end of a biphasic pulse, similar to earlier estimates (Gustafsson and Jankowska 1976). This implies that the spikes we observed on the stimulated electrode were evoked transsynaptically.

We asked the same questions about the spikes recorded on non-stimulated electrodes. Due to the increased distance that evoked spikes could propagate, the latency at which we could record directly evoked spikes would also increase. For electrodes within $700 \mu\text{m}$ of the stimulated electrode, the maximum distance an action potential could travel is $\sim 950 \mu\text{m}$. This makes the longest theoretical latency of directly evoked spikes $\sim 0.7 \text{ ms}$ after the end of the biphasic pulses, very close to our observation. For these nearby electrodes, it remains likely we are recording transsynaptic activation.

Limitations due to missing directly evoked spikes

The major limitation of this study is that we were unable to record activity until $\sim 0.7 \text{ ms}$ after stimulation offset, causing us to miss the initial, directly evoked spikes. To assess this impact,

we estimated the proportion of spikes they represented. Since we can begin recording ~ 0.7 ms after stimulation offset, we might miss at most one spike per pulse. In the worst case, where the neuron is directly activated by each pulse, 1.2 – 1.4 spikes are evoked per neuron across amplitudes (Fig. 2.3c). Thus the average of 0.2 – 0.4 transsynaptically evoked spikes we recorded accounts for 17 – 30% of evoked spikes. Even though distance and amplitude affect the proportion of pulses which directly evoke a spike (Stoney, Thompson, and Asanuma 1968), we likely miss a large proportion of evoked spikes near the stimulated electrode.

While we cannot record directly evoked activity, we can infer something about the temporal pattern of directly evoked activity using activity recorded on non-stimulated channels, which must either be driven directly or transsynaptically by directly evoked spikes. Hence, we would expect the firing rate dynamics on non-stimulated channels to be like that of directly evoked spikes. Responses on non-stimulated channels were maintained throughout long trains of stimulation (Fig. 2.5), presumably driven by the maintained responses of at least some directly evoked neurons. In contrast, the likely-transsynaptic response recorded on the stimulated electrode decayed rapidly. It could be that directly evoked activity near the stimulated electrode decayed rapidly, but researchers using calcium imaging found that neurons closer to the stimulated electrode actually maintained their responses longer than those farther away (Michelson et al. 2018). Instead, the decayed response we observed is likely caused by direct activation of local inhibitory neurons which competes with the excitatory effect (Overstreet, Klein, and Helms Tillery 2013).

If the high temporal resolution that electrical recordings provide is not necessary, calcium imaging or voltage-sensitive dye imaging can be used to record activity during the stimulus pulse as these methods are not affected by the shock artifact (Tanaka et al. 2019; Histed, Bonin, and

Reid 2009; Michelson et al. 2018). To isolate directly evoked spikes, pharmacological agents have been used to block synaptic transmission, enabling researchers to study the spatial pattern of directly evoked spikes (Histed, Bonin, and Reid 2009). Both this study and one using biophysical models (Kumaravelu et al. 2022) concluded that ICMS activates a sparse and distributed population of neurons, likely due to local activation in axons propagating antidromically to somas. Combined with our results, we can describe the spatial and temporal pattern of directly and transsynaptically evoked activity.

Qualitatively similar evoked responses across different experimental conditions

Across many studies using different levels of anesthesia, animal models, and recording techniques, the evoked response to ICMS is qualitatively similar. After stimulation, neurons exhibit short-latency excitation due to direct or transsynaptic activation (Margalit and Slovin 2018; Tehovnik et al. 2006). We observed an increase in the amount of evoked activity from activated neurons and an increase in the spread of evoked activity with increasing amplitude (Fig. 2.3), consistent with previous observations (Hao, Riehle, and Brochier 2016; Butovas and Schwarz 2003). With increased frequency, we observed a small increase in the amount of evoked activity per pulse, an effect that is further amplified by the increased number of pulses (Fig. 2.7). After short-latency excitation, neural activity is typically suppressed for long periods, an effect likely mediated by GABA_B receptors (Butovas et al. 2006). The duration of this long-lasting inhibition increased with amplitude in our study (Fig. 2.4), in contrast to previous observations of neurons recorded farther from the stimulated electrode (Butovas and Schwarz 2003). After inhibition, we often saw a large increase in firing rate (Fig. 2.6) (Butovas and Schwarz 2003). This rebound excitation may be due to recurrent excitation within cortical circuits, mediated by calcium channels

(Molineux et al. 2006; McElvain et al. 2010). Throughout trains of stimulation, we observed a rapid decay of the transynaptically evoked activity recorded near the stimulated electrode (Fig. 2.5), similar to previous observations (Michelson et al. 2018).

Implications for biomimetic stimulation patterns

In monkeys, stimulation in tactile cortices evokes sensations at locations corresponding to the receptive field of neurons recorded on the stimulated electrode (Tabot et al. 2013). Different temporal patterns of stimulation can be distinguished and used to convey useful information (Callier et al. 2020; Berg et al. 2013; London et al. 2008; Romo et al. 1998; Dadarlat, O'Doherty, and Sabes 2015). Similar observations have been made in humans with tetraplegia and neuropathy (Salas et al. 2018; Chandrasekaran et al. 2021; Fifer et al. 2020; Hughes et al. 2021), including the ability of one person to identify which of multiple fingers of a robotic hand, linked to somatosensory cortex (S1) stimulation, were touched (Flesher et al. 2016). More recently, somatosensory ICMS was used to provide contact and pressure-related feedback, which improved their ability to control a robotic arm to reach and grasp (Flesher et al. 2021). The stimulus parameters in this most recent example were quite simple, a linear mapping from index and middle finger joint torques to appropriate electrodes, and evoked sensations that were judged to be possibly natural (Flesher et al. 2016). Biomimetic stimulation patterns, those that aim to evoke activity that mimics the spatial and temporal properties of naturally occurring activity, may be necessary to evoke more naturalistic sensations (Bensmaia and Miller 2014).

To develop biomimetic stimulus patterns, it may be useful to compare the spatial and temporal dynamics of naturally occurring activity to the activity evoked by stimulation. In tactile areas, neurons respond to skin indentation with a large transient response and smaller sustained

response (Callier, Suresh, and Bensmaia 2019). Our data suggest that neurons entrain their responses to each pulse in the train. Thus, to recreate the temporal dynamics of this response, the frequency of stimulation needs to be modified throughout the train. In proprioceptive areas, limb movements evoke a complicated spatiotemporal pattern of activity across cortex that is dependent on the direction and speed of reaching movements, as well as interaction forces (London and Miller 2012; Prud'homme and Kalaska 1994). Recreating these patterns may require small amplitudes on many electrodes, in order to target groups of neurons with similar encoding properties (Weber et al. 2011). Even with small amplitudes, though, some directly activated neurons may be located far from the stimulated electrode due to the local activation of axons (Histed, Bonin, and Reid 2009; Kumaravelu et al. 2022). Monitoring the locations of activated neurons may enable researchers to design stimulation patterns that more closely mimic the naturalistic spatial response more closely.

The activity evoked by ICMS is also unnaturally synchronous across neurons. There are several stimulus protocols that may serve to reduce synchrony, including stimulating with amplitudes nearer the activation threshold, where spikes are not evoked reliably (Fig. 2.3b, d). Alternatively, single pulses within a train can be replaced with kilohertz bursts of pulses, with amplitude increasing throughout the burst (Formento et al. 2020). Neurons with different activation thresholds will be activated at different times during the burst. Finally, multi-electrode stimulation could be delivered asynchronously across electrodes. The ability to record the evoked activity would allow the efficacy of any combination of these approaches to be evaluated.

Linking evoked activity to sensation

Most sensory modalities obey Weber's law: The Just Noticeable Difference (JND) in stimulus intensity increases with increased amplitude (Ekman 1959). This log-like relation likely

occurs because the firing rate variance increases with mean firing rate (Johnson 1980b, 1980a). As a consequence, it is possible to detect very small differences within natural low-intensity stimuli. In contrast, for ICMS stimulation, the JND remains constant across amplitudes and frequencies (Kim, Callier, Tabot, Gaunt, et al. 2015), potentially limiting the number of discrete intensities that can be discriminated. While we recorded a linear increase in mean firing rate with increased amplitude and frequency, there was little change in the variance as these parameters changed (Fig. 2.7), likely the source of the JND that is unchanging with increased stimulus intensity.

In our experiments, we typically recorded a maintained response on non-stimulated channels throughout 4-s long trains (Fig. 2.5), a response which likely reflects the temporal dynamics of directly evoked spikes. Because of this, we would predict the perceived intensity due to longer trains to be constant for at least this length of time at the frequencies and amplitudes we tested. Our observation is consistent with the observations of a human participant in a recent study, who reported constant perceptual intensity for ~ 7 s (Hughes, Flesher, and Gaunt 2021).

However, after the end of long stimulus trains, the evoked sensation described above did not disappear immediately, but persisted for a couple of seconds. We often observed a large burst of rebound excitation after the end of high frequency trains (Fig. 2.6), which could potentially lead to persistent sensations. Since rebound excitation primarily occurred at high stimulation frequencies, it may be that there is a maximum frequency that future afferent interfaces can use to avoid the effect.

Online recording in the presence of stimulation artifact

For most applications, afferent interfaces would only be useful when combined with an efferent interface, thereby providing both restored somatosensation and movement (O'Doherty et

al. 2011b; Flesher et al. 2021). However, stimulation in S1 produces large artifacts in recordings from motor cortex (M1). With causal filters, neural signals can be recorded from M1 in a human ~ 0.7 ms after offset of stimulation applied in S1. At low stimulation frequencies, losing the ability to record from M1 for short periods after each pulse will not have much of an impact on decoding performance. When intended cursor velocity was decoded from M1, artificially dropping a random 20% of M1 signals caused only a 10% decrease in decoder performance (Fig. 8 in (Young et al. 2018)). While acausal, time-reversed filtering may allow for slightly earlier recordings, the increased amount of data would likely have a negligible impact on decoding performance.

However, as stimulation protocols become more complicated, with stimulation at high rates and on many electrodes (Bensmaia and Miller 2014; Sombeck and Miller 2019), the percentage of time in which signals can be recorded from M1 will decrease, further decreasing decoder performance. Stimulation at 333 Hz, either on a single electrode or across electrodes, would result in 50% loss of signal, assuming a total blanking duration of 1.5 ms per pulse (Weiss et al. 2018). With some non-trivial amplifier modifications to increase somewhat, the gain during the stimulus artifact, the RRA could potentially enable neural recordings even during the stimulus pulse, albeit at a significantly reduced gain. Although we did not explore them here, there are numerous approaches that could be used to extract neural signal from the artifact if the recorded signal is not saturated: adaptive filtering (Mendrela et al. 2016; Nag et al. 2015), template subtraction (Montgomery Jr, Gale, and Huang 2005; Hashimoto, Elder, and Vitek 2002), independent component analysis (Hyvärinen and Oja 2000; Lemm et al. 2006), linear regression reference (Young et al. 2018), and deep neural networks (Tamada et al. 2020; Zhang and Yu 2018). Of particular note is ERAASR, a technique which uses principal component analysis to exploit the similar structure of the shock artifact sequentially across electrodes, pulses, and then trials (O'Shea

and Shenoy 2017). With these approaches, it may be possible to recover neural signal throughout multi-channel stimulation, thereby enabling full band-width recordings in M1 while providing somatosensory feedback via ICMS in S1. Such technology will likely be necessary to accurately decode motor intent as ICMS feedback becomes more complicated.

Supplementary Materials

Technical description of the rapid-recovery amplifier

We aimed to develop an amplifier which can record neural activity ~ 1 ms after stimulation offset. To do so, the stimulus pulses should not overdrive the amplifier nor saturate the filters, particularly high-pass filters, which by design have a long recovery time constant. The amplifier schematics are shown in Fig. 2.8. The amplifier implements a single-ended configuration with three stages; the first stage uses the OPA140DGK (Texas Instruments, Dallas, TX) operational amplifier (op-amp) and the second and third stages use OPA2277U-EP (Texas Instruments). The first (input) stage and the second stage include high-pass filters and gain compression; the third stage is a linear gain stage.

All stages are supplied with ± 15 V to prevent output saturation and input current shunting. If the op-amp input were to exceed the power supply rails, the input signal would be clamped to ± 15 V due to the electrostatic discharge protection of the chip, resulting in an undesirable transient drop in the input impedance. As a consequence, the input stage might shunt a part of the stimulation current, reducing the current left for stimulation if the stimulator and amplifier share the same electrode. If stimulation uses a separate electrode, the drop in impedance might inject a significant current through the recording electrode.

We used an ac-coupling input capacitance of 1 nF to balance several design objectives. A large capacitance value is advantageous as it reduces the impedance driving the input stage, extends the low-frequency corner, and reduces the residual voltage on the capacitor after a large input artifact, which could saturate the input stage. On the other hand, a small capacitance value is required to limit the charge injected into the brain in case of an op-amp failure, which may short the op-amp input to its ± 15 V supply. The ac-coupling capacitance is biased to ground with a 500 M Ω resistor, which dominates the input impedance. As the input impedance forms a voltage divider together with the electrode and the source's equivalent impedance, the input impedance should be significantly larger than the total impedance driving the amplifier input. The high ac-coupling capacitance and high impedance leads to a long time constant of the first stage (500 ms), which is eliminated by faster (1 ms) ac-coupling dynamics in the second stage. To allow a high input impedance without a large dc voltage offset caused by the op-amp input bias current, the feedback network of the first stage is ac-coupled to ground through a 570 nF capacitor, reducing the dc gain to unity. This capacitance is larger than the ac-coupling input capacitance, but does not affect the maximum injected charge in case of a failure of an integrated circuit.

Saturation of subsequent stages during a stimulus pulse is avoided by nonlinear gain, which compresses large dynamic swings. We implemented a feedback loop in the first stage which reduces the gain to unity when the output exceeds the input by approximately 1.3 V, corresponding the total forward voltage drop of the two series diodes in parallel with the feedback resistor and capacitor. The input to the second amplification stage is ac-coupled with a fast time constant (1 ms) and is clamped by a pair of Schottky diodes limiting the gain for large signals (Mueller et al. 2014). The rest of the second stage and the third stage are linear and provide additional gain for the neural signals.

We used a custom-made breakout board to pass signals from the stimulated channel to the RRA, bypassing the Stim Headstage and Front-end amplifier (Figs. 2.9, 2.10, 2.11). Signals recorded on non-stimulated channels were passed to the Stim Headstage and Front-end amplifier.

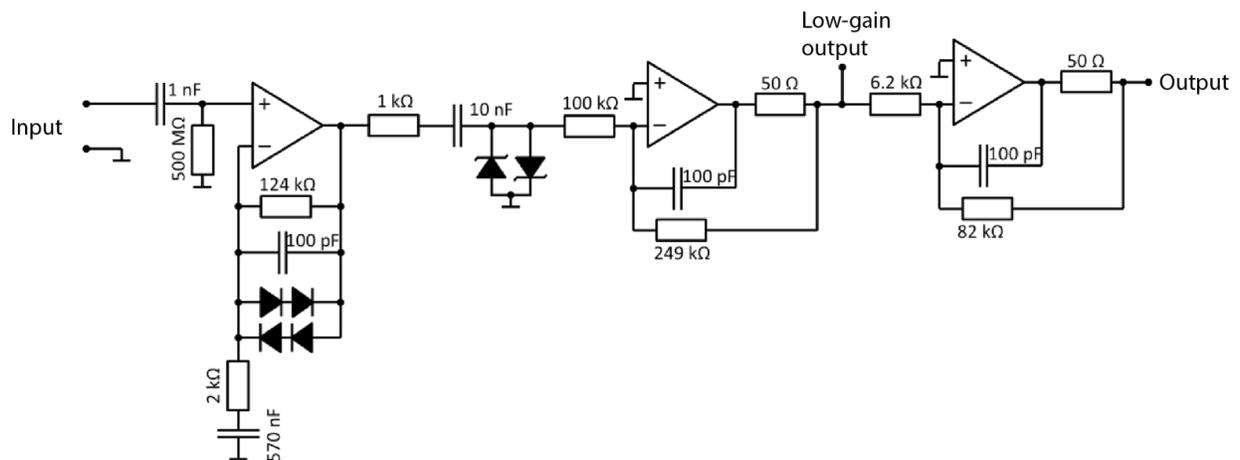


Fig. 2.8. Circuit diagram of the rapid-recovery amplifier. We used the output node (far-right) to record neural signal. The additional low-gain output was included for debugging. The 1 nF and the 10 nF are both 1% metalized polypropylene MKP capacitors from Vishay. The 100 pF are from Panasonic's ech-u1h series. The diodes are SD103A (Vishay) and 1N4148 (Onsemi). The 500 MOhm resistor is from the CRCW series from Vishay.

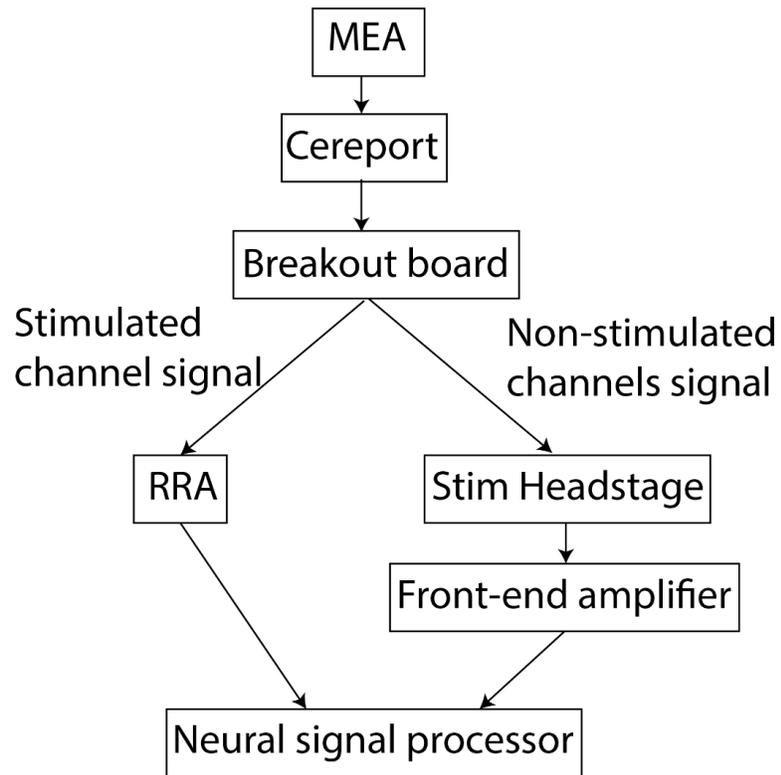


Fig. 2.9. Connection diagram. Neural signal is recorded from a multi-electrode array (MEA). A custom-made breakout board passes only signal from the stimulated channel to the RRA and passes signals from all non-stimulated channels to the Stim Headstage and Front-end amplifier. Signals from all channels are passed to the neural signal processor.

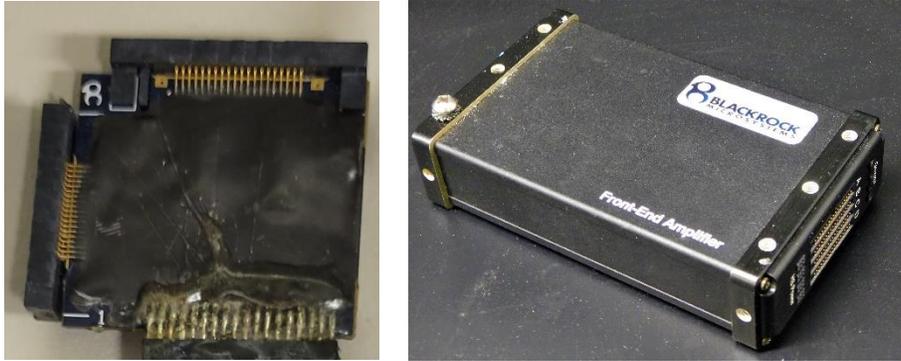


Fig. 2.10. Pictures of Stim Headstage (left) and Front-end amplifier (right).

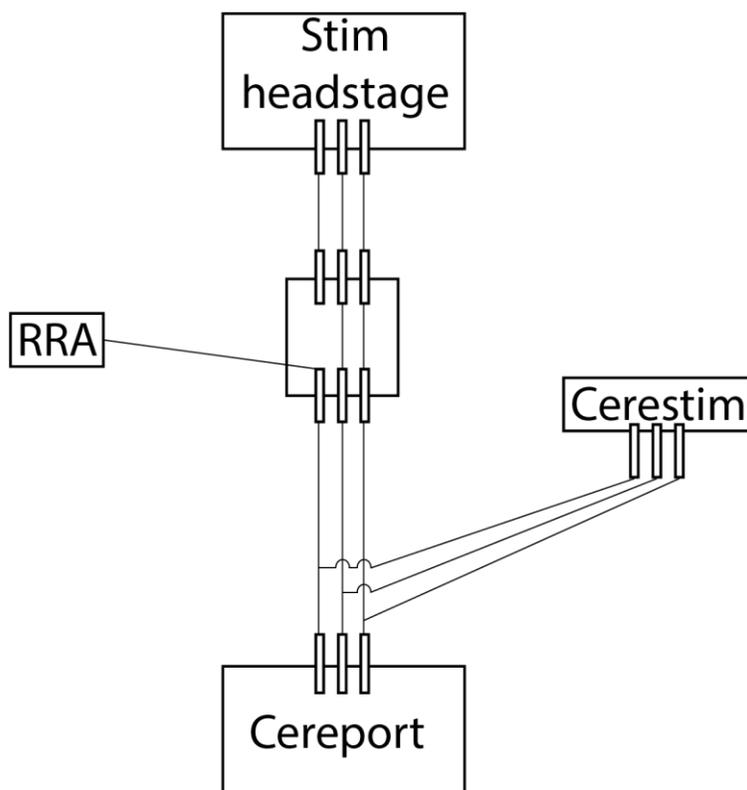


Fig. 2.11. Diagram of breakout board. Neural signals recorded on the stimulated channel are passed to the RRA from the Cereport while all other signals are passed to the Stim Headstage. The channel passed to the RRA was manually selected. All channels are connected to the Cerestim for stimulation (only one channel was stimulated at a time). Only three of the 96 channels are depicted for clarity.

Chapter 3 - Short reaction times in response to multi-electrode intracortical microstimulation may provide a basis for rapid movement-related feedback

Joseph T. Sombeck, Lee E. Miller

Foreword

The following chapter has been adapted from a manuscript published in the *Journal of Neural Engineering* in December 2019. The purpose of this project was to investigate the slow response to intracortical microstimulation (ICMS) previously observed (Godlove, Whaite, and Batista 2014). In these experiments, ICMS elicited response times slower than even vision, which is too slow to control ongoing movements. These results were surprising, as ICMS is applied directly to the brain, and problematic because ICMS needs to provide rapid feedback to replace natural proprioception. In this project, single-electrode stimulation also typically evoked long reaction times. By distributing large amounts of charge over many electrodes, ICMS could evoke reaction times as short or shorter than a proprioceptive cue. These results suggest that future proprioceptive interfaces will need to stimulate on many electrodes simultaneously to replicate the fast feedback provided by natural proprioception.

Abstract

Objective. Tetraplegic patients using brain-machine interfaces (BMIs) can make visually guided reaches with robotic arms. However, restoring proprioceptive feedback to these patients will be critical, as evidenced by the movement deficiencies in patients with proprioceptive loss. Proprioception is critical in large part because it provides faster feedback than vision. Intracortical microstimulation (ICMS) is a promising approach, but the ICMS-evoked reaction time (RT) is typically slower than that to natural proprioceptive and often even visual cues, implying that ICMS feedback may not be fast enough to guide movement. *Approach.* For most sensory modalities, RT decreases with increased stimulus intensity. Thus, it may be that stimulation intensities beyond what has previously been used will result in faster RTs. To test this, we compared the RT to ICMS applied through multi-electrode arrays in area 2 of somatosensory cortex to that of mechanical and visual cues. *Main results.* We found that the RT to single-electrode ICMS decreased with increased current, frequency, and train length. For 100 μA , 330 Hz stimulation, the highest single-electrode intensity we tested routinely, most electrodes resulted in RTs slower than the mechanical cue but slightly faster than the visual cue. While increasing the current beyond 100 μA resulted in faster RTs, sustained stimulation at this level may damage tissue. Alternatively, by stimulating through multiple electrodes (mICMS), a large amount of current can be injected while keeping that through each electrode at a safe level. We found that stimulation with at least 480 μA equally distributed over 16 electrodes could produce RTs as much as 20 ms faster than the mechanical cue, roughly the conduction delay to cortex from the periphery. *Significance.* These results suggest that mICMS may provide a means to supply rapid, movement-related feedback. Future neuroprosthetics may need spatiotemporally patterned mICMS to convey useful somatosensory information.

Introduction

Efferent brain-machine interfaces (BMIs), which decode motor intent from recorded brain activity, can allow a tetraplegic patient to move a robotic arm (Collinger et al. 2013; Hochberg et al. 2012) or even their own arm, using functional electrical stimulation (FES) to cause their paralyzed muscles to contract (Ajiboye et al. 2017; Bouton et al. 2016; Ethier et al. 2012). These BMIs typically rely solely on visual feedback to guide movement, despite the considerable movement deficits suffered by patients without somatosensation (Rothwell et al. 1982; Ghez, Gordon, and Ghilardi 1995). Without cutaneous sensations, subjects exert forces larger than necessary, often crushing delicate objects (Monzée, Lamarre, and Smith 2003), and dexterous manipulation of small objects becomes almost impossible (Johansson and Flanagan 2009). Patients who have lost proprioception are for the most part wheelchair bound and make large reaching errors due to an inability to plan and rapidly update ongoing reaches (Ghez et al. 1990; Sainburg et al. 1995). The relatively slow speed of visual feedback is one of the reasons that it is an inadequate replacement for somatosensation. Additionally, somatosensation is important for prosthesis embodiment (Antfolk et al. 2013). Thus, restoring somatosensation is a critical, yet unmet component of BMI development.

Intracortical microstimulation (ICMS) has the potential to restore somatosensation, having been shown to elicit conscious perceptions in rats (Devecioğlu and Güçlü 2017; Öztürk et al. 2019; Fridman et al. 2010), monkeys (Tabot et al. 2013; London et al. 2008; Romo et al. 2000; O’Doherty et al. 2011b) and humans (Salas et al. 2018; Flesher et al. 2016). Stimulation in tactile cortical areas provide sensations of flutter at a frequency that matches the stimulation frequency (Romo et al. 2000). Consequently, ICMS has been used to provide artificial texture feedback, enabling monkeys to learn to select rewarded virtual objects based on their “feel” (O’Doherty et

al. 2011b). Additionally, the virtual location of the sensation elicited in tactile areas corresponds to the receptive field of neurons in that area (Tabot et al. 2013), enabling a spinal cord injured patient to identify which of multiple robotic fingers were touched (Flesher et al. 2016).

Replicating the sensations corresponding to the more distributed and complex receptive fields of proprioceptive neurons has not been as successful. In an experiment that relied on the ability to learn the meaning of an abstract stimulus, monkeys were able to reach to unseen targets using ICMS feedback about the error vector between the changing hand position and target position (Dadarlat, O'Doherty, and Sabes 2015). This interface, though, required months of training, in contrast to the more rapidly learned, biomimetically-inspired, mapping in tactile areas (Flesher et al. 2016). The long training time required was probably due to the complex learning problem associated with mapping an abstract stimulus to limb state. In an effort to provide a more natural proprioceptive sensation, thereby reducing training time and possibly providing more informative feedback, our group stimulated on small sets electrodes, selected because of their mutually similar responses recorded during arm movements (Tomlinson and Miller 2016). This biomimetic approach predictably biased one monkey's perception of the direction of a coincident mechanical perturbation without any learning, suggesting that it had indeed, evoked a sensation like that of the natural perturbation. However, the effect could not be replicated in subsequent monkeys, for reasons that remain unclear.

In addition to evoking meaningful sensations, afferent interfaces also need to provide fast feedback, like that of somatosensation. Patients without proprioception make their largest errors during rapid movements, in part, because the slow speed of visual feedback limits correction of these movements (Ghez et al. 1990; Sainburg et al. 1995). Slow feedback also limits successful embodiment of a prosthesis (Shimada, 2009). It is reasonable to assume that ICMS could provide

very rapid feedback, as it bypasses the conduction latency from the periphery. However, the response time to single-electrode stimulation in tactile areas is typically slower than that to either tactile or visual cues (Godlove, Whaite, and Batista 2014). Critically, if ICMS is no faster than natural vision, it is unlikely to replace it for guiding rapid reaches or enabling embodiment.

In this paper, we used a reaction time (RT) paradigm as a rapid, sensitive mean to compare the latency of ICMS applied through multi-electrode arrays implanted in area 2 of somatosensory cortex to that of perturbations applied to the hand and to visual cues. Consistent with earlier studies, we found that single-electrode ICMS (sICMS) typically resulted in RTs that were slower than limb perturbations and slightly faster than the visual inputs. On the other hand, multi-electrode ICMS (mICMS) elicited RTs even faster than limb perturbations. We investigated the effect of number of electrodes, total current, and distance between electrodes on the RT to mICMS. The use of many electrodes simultaneously may also allow more the complex spatial patterns of cortical activity typical of natural proprioceptive inputs to be elicited. Our results show that mICMS may be a suitable approach for providing fast feedback in future afferent interfaces.

Methods

Monkeys

All procedures in this study were performed in accordance with the guide for the care and use of laboratory animals and were approved by the institutional animal care and use committee of Northwestern University under protocol #IS00000367. The experiments were performed using two male rhesus macaques (Monkey H: 12.9 kg, Monkey D: 9.8 kg).

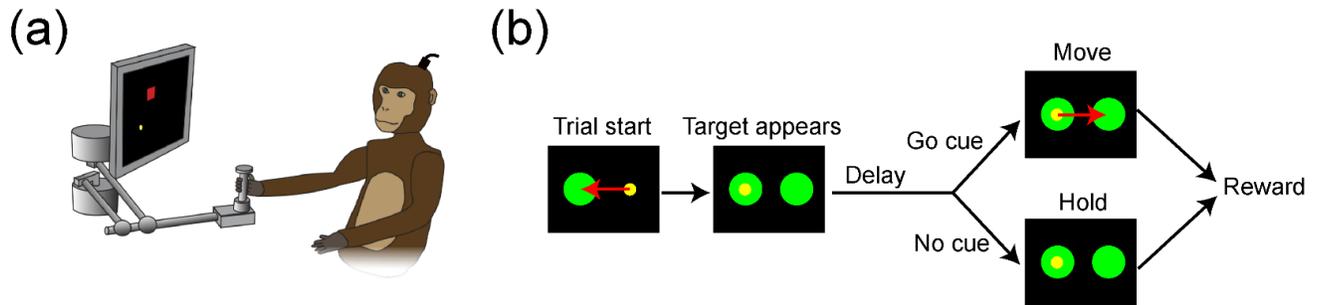


Fig. 3.1. Reaction time task description. (a) Monkeys move the handle of a robotic manipulandum to perform a reaction time task. (b) A trial starts when the monkey moves the cursor into the center target, causing the goal target to appear. After a random delay, the monkey receives a reward for making a reach if a go cue is presented (top), or for holding in the center target if a go cue is not presented (bottom).

Reaction time task

Monkeys held the handle of a two-link planar robotic manipulandum which controlled a cursor on an LCD screen in front of them (Fig. 3.1(a)). In the RT task, monkeys reached to a target in response to a go cue consisting of either a mechanical perturbation of the hand, a change in the color of the targets on the screen, or ICMS in area 2. Each trial began when the monkey moved the cursor into a target at the center of the workspace (Fig. 3.1(b)). At this point, the goal target appeared 9 cm from the center target, so that the monkey could plan where to reach. The cursor then disappeared, to avoid providing visual feedback about the mechanical cue. Monkey H made reaches to the right and monkey D made reaches forward, as these were the directions in which they each moved most rapidly. On 85% of trials for monkey H, a go cue was provided at a random time between 500 and 1500 ms after the goal target appeared. For monkey D, a go cue was provided on 75% of trials between 500 and 2000 ms after the goal target appeared. If the monkeys reached the goal target within a short time window (750 ms for monkey H and 800 ms for monkey D) after go cue onset, they received a water reward. In the rest of the trials, no cue was presented, and the monkeys were rewarded for not moving from the center target. These trials were included to reduce the rate of false starts.

The mechanical cue was a force applied to the hand for 120 ms, including 20 ms rise and fall times. We used a short duration so that the perturbation did not affect the monkey's subsequent reach. There was a white noise audio mask throughout each experiment to prevent monkeys from hearing the motors during mechanical cues. The direction of the mechanical perturbation was perpendicular to the reach direction to simplify determination of movement onset. Stimulation in area 2 likely elicits both proprioceptive and tactile sensations, as area 2 is known to integrate information from both muscle and cutaneous receptors (Weber et al. 2011; Hyvärinen and Poranen

1978; Pons et al. 1985). As such, comparing the RT evoked by a mechanical cue to that of ICMS in area 2 is quite appropriate. The ICMS duration was typically 120 ms for monkey H and 200 ms for monkey D, though we tested the effect of duration in some experiments. For the visual cue, both the center and goal targets changed from red to white. This new color persisted throughout the entire trial. Each cue type (mechanical, visual, and ICMS) was presented alone in a block so that the monkeys knew which cue to attend to. Blocks were presented in random order within a session. In an ICMS block, a single electrode or set of electrodes was stimulated and the stimulation amplitude, frequency, or train length could change, when applicable. There were multiple ICMS blocks within a session.

To train monkeys to respond to the mechanical cue, we paired it with a previously learned audio cue on a large proportion of trials and presented it alone on the remainder. The proportion of trials with only the mechanical cue increased as the monkeys learned. Once they reacted to the mechanical cue alone, we progressively reduced the allotted movement time, causing them to react quicker and make more rapid reaches to receive a reward. The movement time was decreased until the monkeys could no longer successfully complete the task, and then increased by 200 ms. The monkeys then learned to respond to the visual and ICMS cues when paired with the mechanical cue. After about 1 week of training, monkeys began responding to ICMS alone.

Stimulation and data collection

After becoming proficient at the RT task, each monkey was implanted with a 1-mm long, 96-electrode, sputtered iridium-oxide Utah multi-electrode array (Blackrock Microsystems, Salt Lake City, UT) in the proximal arm area of somatosensory cortical area 2. In surgery, we found the arm representation by recording from the cortical surface while manipulating the arm and hand.

For more details on surgical techniques, see (Weber et al. 2011). After the implant surgery, we performed sensory mappings to confirm that recorded neurons had receptive fields corresponding to the proximal arm. On ICMS trials, electrodes were stimulated with pulse trains consisting of cathodal-first, biphasic, 200 μ s pulses, using a Cerestim R96 (Blackrock Microsystems, Salt Lake City, UT). There was 53 μ s between phases in each pulse. In experiments where we stimulated more than 16 electrodes, the electrodes were stimulated in two equal-size groups, separated by a 100 μ s lag, because the stimulator was limited to simultaneous stimulation of 16 electrodes. To control for any potential effect due to asynchronous stimulation in applicable experiments, all sets of electrodes were separated into two groups, regardless of how many electrodes were within each group.

A Cerebus system (Blackrock Microsystems, Salt Lake City, UT) was used to collect handle kinematics and cue onset times. Handle kinematics were recorded at 100 Hz using encoders on the manipulandum joints. Stimulation onset was determined through the sync line from the Cerestim R96 (Blackrock Microsystems, Salt Lake City, UT), and visual cue onset was determined by a photodiode placed near the screen, both sampled at 30kHz. Mechanical cue onset was defined as the time of the command signal to the servo motors.

Experiments with monkey H began 30 months after array implantation, and 1 month after monkey D was implanted. Experiments were performed once a day for 6 months with monkey H (total of 34 sessions), and for 1.5 months (18 sessions) with monkey D. There were 15.2 ± 6.8 (mean \pm standard deviation) successful trials per condition for monkey H and $18.9.0 \pm 6.8$ for monkey D across all conditions and sessions.

Data analysis

RT was defined as the time from cue onset to movement onset. To find movement onset, we first found the time of peak acceleration in the direction of the goal target after cue presentation. We then found movement onset by backtracking until the acceleration dropped below 35 cm/s^2 , roughly two standard deviations above the mean acceleration measured during the hold period. The same threshold was used for all cue modalities across all sessions for both monkeys. RTs above or below two standard deviations from the mean were removed as outliers. There were a small number of outliers for each condition, mostly due to false starts, resulting in extremely fast RTs, and extremely slow RTs, likely due to the monkeys not paying attention during some trials.

Statistical analysis

We used two-sided Welch's t-tests to compare RTs to the mechanical and visual cues within sessions, and to compare sICMS at $100 \mu\text{A}$ for each electrode to the mechanical cues within the same session. We also used Welch's t-tests to compare the mean RT across all electrodes when stimulated at $100 \mu\text{A}$ to the mean RT in response to the mechanical cue, the mean RT to sICMS at $200 \mu\text{A}$ and the mechanical cue, and the RT to mICMS and the mechanical cue. We used paired t-tests to compare the RT to sICMS and mICMS and to adjacent and non-adjacent electrode groups. The effect of amplitude, frequency, or train length on the resulting sICMS RT and the effect of total current and number of electrodes on the resulting mICMS RT were compared with an analysis of variance (ANOVA), with data aggregated across monkeys. We also used an ANOVA to test the effect of electrode position on the RT to sICMS at $100 \mu\text{A}$ for each monkey individually, generating two models. We used an F-test of the significance of the models and each individual parameter, when relevant.

Results

Reaction time in response to natural proprioceptive and visual cues

Monkeys performed an RT task in which they reached from a center target to an outer target when cued with either a mechanical or visual go cue (Fig. 3.1). We tested perturbation forces ranging from 0.1 to 1 N in monkey H, which caused a displacement of the monkey's hand from roughly 0.1 to 0.5 cm forward, and from 0.5 to 4.5 N in monkey D, which moved the monkey's hand to the right from 0.75 to 10 cm. Fig. 3.2 shows the mean RTs in response to mechanical cues of different magnitudes for two representative sessions for each monkey, as well as the mean RT to the visual cue (black dashed line). As anticipated from many earlier studies across sensory modalities (Pins and Bonnet 1996), RT decreased with increasing perturbation magnitude, reaching an asymptotic value at about 160 ms for a 1 N pulse in monkey H and 180 ms for a 3 N pulse in monkey D. We used only these forces for subsequent experiments in the respective monkeys. Across all sessions, the mean RT in response to the mechanical cue was 162 ± 14 ms for monkey H and 186 ± 23 ms for monkey D. For the visual cue, the mean RT for monkey H across all sessions was 250 ± 48 ms, and for monkey D, 380 ± 39 ms. In every session for both monkeys, the mean RT to the mechanical cue was significantly faster than that to the visual cue ($p < 0.01$, Welch's t-test).

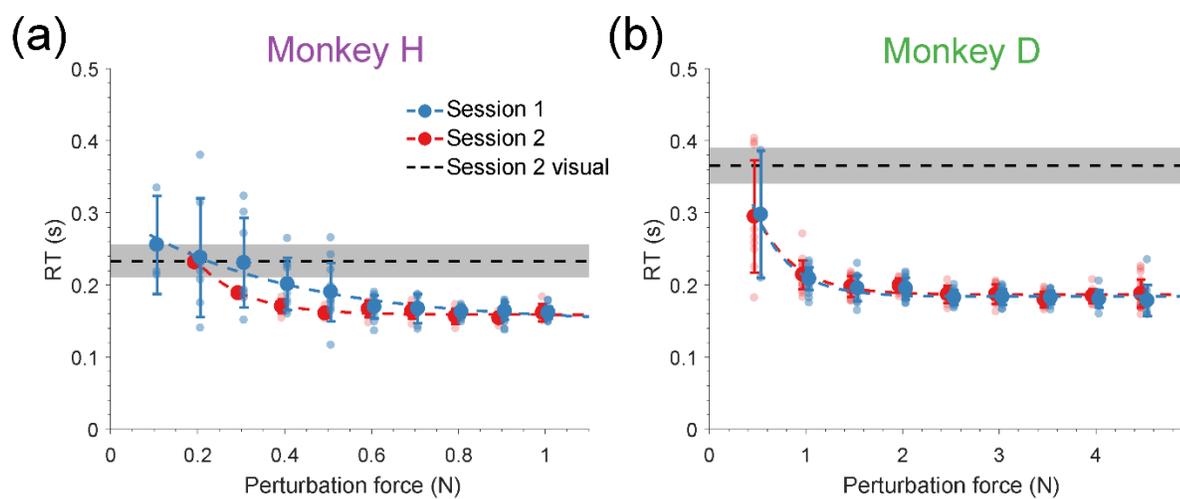


Fig. 3.2. Dependence of reaction time on mechanical cue force magnitude. The reaction time (RT) to the mechanical cue with various forces for two sessions (red and blue) in (a) monkey H and (b) monkey D. Large circles represent the mean RT for each force magnitude. Small circles show the RT for single trials. Colored dashed lines show exponential fits to the data. The RT to the visual cue during session 2 is shown as a horizontal black dashed line. All error bars show standard deviation.

Reaction time in response to single-electrode ICMS

We tested sICMS with a wide range of current amplitudes, frequencies, and train lengths. The range of amplitudes (20-100 μ A) and frequencies (50-500 Hz) included parameters commonly used in ICMS experiments. Initially, we used a maximum of 100 μ A, as this was the largest current tested in earlier ICMS safety studies (Chen et al. 2014; Rajan et al. 2015). We also tested a range of train lengths (75-300 ms) with a maximum longer than the mean RT to ICMS found previously (Godlove, Whaite, and Batista 2014). Thus, we expected the RT to settle for a train length within the range tested. In one series of experiments, we kept two of these parameters constant and measured the RT while varying the third. The RT to stimulation on two representative channels in monkey D for varied train lengths is shown in Fig. 3.3(a). Figs. 3.3(b) and 3.3(c) show the effect on RT of varied frequency and amplitude respectively. Results in monkey H were similar to these. We used an ANOVA to determine the effect of each parameter on RT. Aggregated across both monkeys, we found that increasing each parameter resulted in significantly faster RTs (amplitude: $F(1,30) = 67$; frequency: $F(1,33) = 46$; train length: $F(1,27) = 17.8$; $p \ll 0.01$ for all).

The RT stopped decreasing for train lengths above about 120 ms for monkey H and 200 ms for monkey D (with a frequency of 200 Hz). The RT stopped decreasing at a frequency of about 330Hz (with 250 ms train length for monkey H and 200 ms for monkey D) for both monkeys. However, the RT continued to decrease for amplitudes up to 100 μ A. Therefore, we used this frequency and these train lengths in later experiments and varied amplitude to change stimulation intensity, unless otherwise noted.

Monkey D

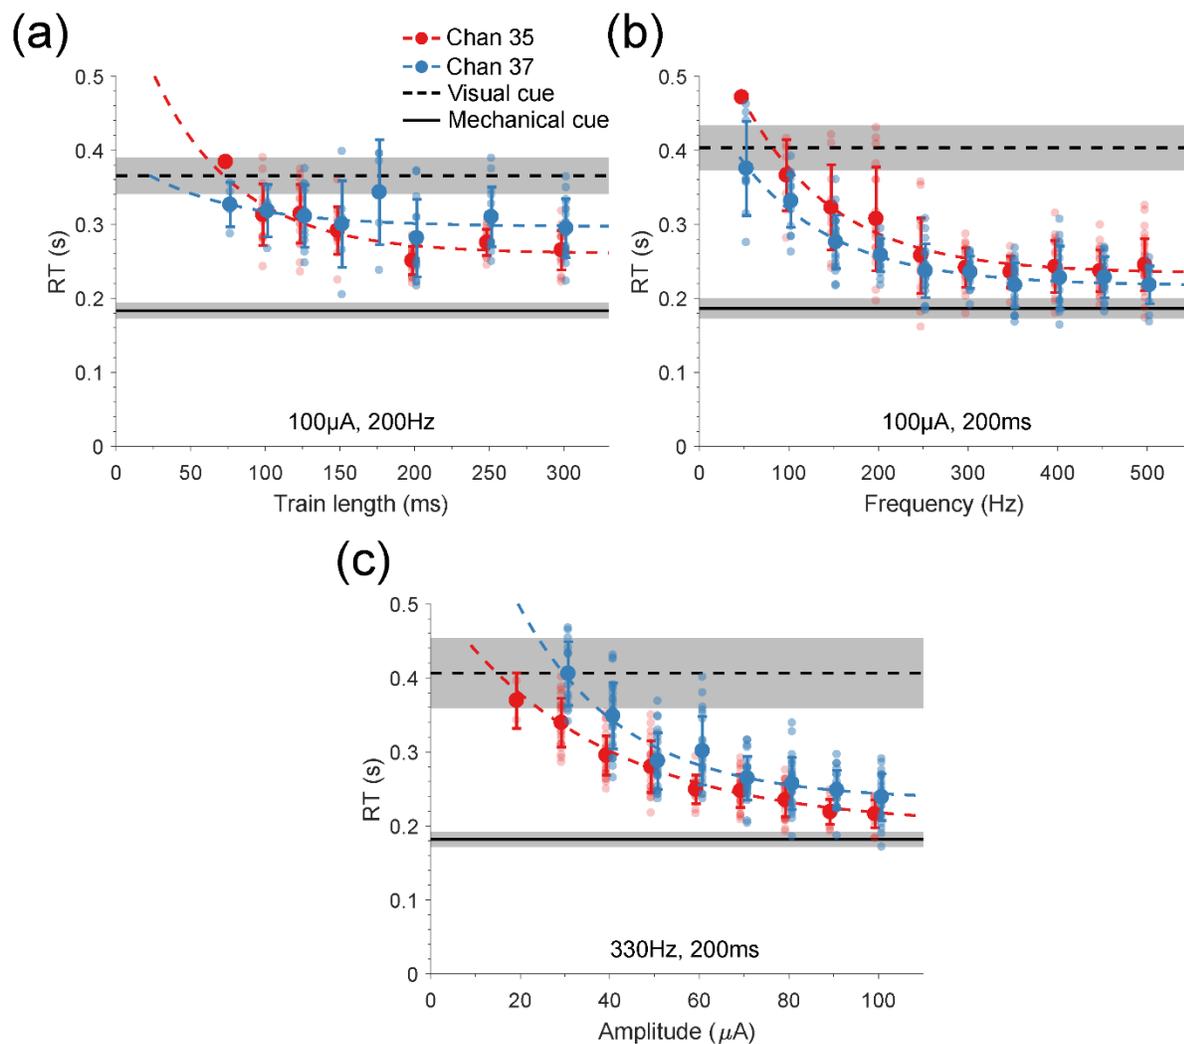


Fig. 3.3. Dependence of reaction time to single-electrode stimulation on stimulation parameters. The reaction time (RT) to single-electrode stimulation while varying (a) train length, (b) frequency, and (c) amplitude and keeping the other two parameters constant are shown for two example channels in monkey D. The fixed parameters are shown in each panel. Black horizontal solid line shows the RT to the mechanical cue during the corresponding session. The mechanical cue was a 1 N, 120 ms pulse for monkey H and 3 N, 120 ms for monkey D. Black horizontal dashed line shows the RT to the visual cue. The same conventions are used as in Fig. 3.2.

We measured the RT for each electrode at 100 μ A across numerous sessions. The mean RT for each electrode in the order they were tested is shown in Fig. 3.4(a) and Fig. 3.4(c) for each monkey. Within each session (divided by vertical dashed lines), we also measured the mean RT in response to the mechanical cue (black solid line) and the visual cue (black dashed line). We compared the single-electrode RT for each electrode to the mechanical cue RT, as this was the fastest natural stimulus. Stimulation on 93 of the 192 individual electrodes resulted in significantly slower RTs than the mechanical cue while 9 electrodes evoked RTs significantly faster than the mechanical cue ($p < 0.05$, two-sided Welch's t-test, Bonferroni correction). 79 electrodes resulted in RTs that were not significantly different than the mechanical cue. The monkeys did not react to stimulation on nine electrodes. Two electrodes caused monkey D to vocalize and were not tested further. Mean RTs for all electrodes and mean RTs to the mechanical cue and visual cue pooled across sessions are summarized in Fig. 3.4(b) and Fig. 3.4(d) for each monkey. The mean single-electrode RT across all electrodes was 199 ± 39 ms for monkey H and 225 ± 53 ms for monkey D. For both monkeys, the RT to single-electrode stimulation across electrodes was significantly slower than that to the mechanical cue in the corresponding session ($p \ll 0.001$, Welch's t-test).

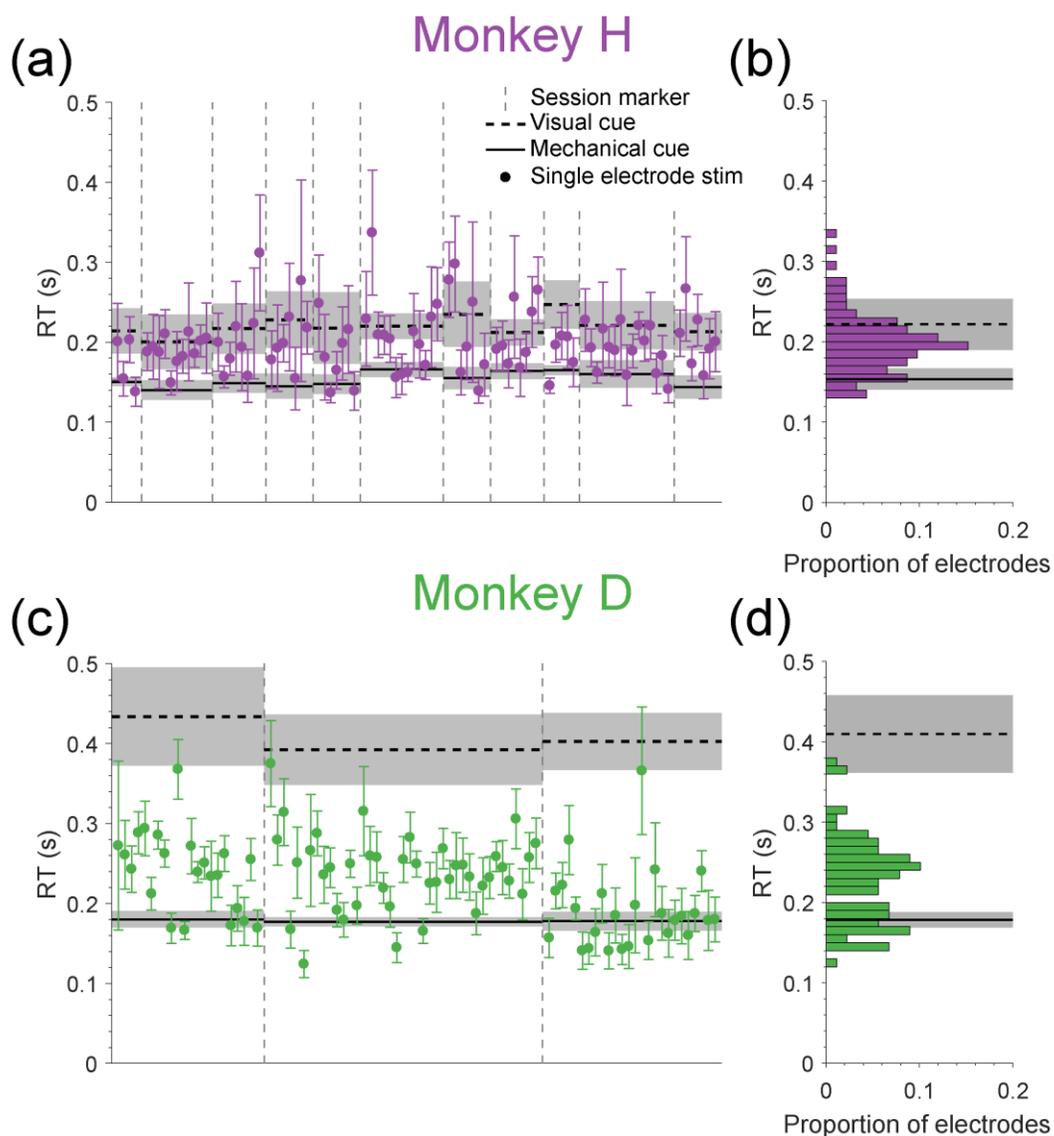


Fig. 3.4. Reaction time to single-electrode stimulation for many electrodes. The mean reaction time (RT) to single-electrode stimulation (circles) for electrodes are shown in the order they were tested across multiple sessions for (a) monkey H and (c) monkey D. The mean RT to the mechanical cue (black solid line) and to the visual cue (black dashed line) are shown for each session. The parameters of the mechanical cue were the same as those of Fig. 3.2. Vertical grey dashed lines denote different sessions. Electrodes were stimulated with $100 \mu\text{A}$ at 330Hz for 120ms for monkey H and for 200ms for monkey D. Error bars show standard deviation. RT data across electrodes is summarized for (b) monkey H and (d) monkey D. The mean RT to the mechanical cue (black horizontal line) and the visual cue (black dashed line) across sessions and the pooled standard deviation are shown for comparison.

The location of the electrode on the array may change the depth of the electrode in cortex, or the apparent location of the sensation on the arm, and thus affect the RT. The position of each array in cortex is shown in Fig. 3.5(a), with labels indicating the central sulcus (CS) and intraparietal sulcus (IPS). The color of each electrode in Fig. 3.5(b) corresponds to the RT measured when that electrode was stimulated at 100 μ A. Black X's indicate electrodes that the monkeys did not respond to, red X's indicate electrodes that caused the animal to vocalize, and white boxes indicate electrodes that were not connected (by design, four on each array). We fit a linear model for each monkey to predict the RT from the position along medial-lateral and anterior-posterior axes of the array, approximately parallel and perpendicular to the IPS respectively. This model was statistically significant in both monkeys, ($p \ll 0.001$, F-test), suggesting that the resulting RT depended on the position of the electrode in cortex. The resulting vector of maximal decrease in RT is shown by the arrows in Fig. 3.5(a). For monkey H, the RT decreased for more lateral electrodes. The RT decreased for posterior electrodes for monkey D.

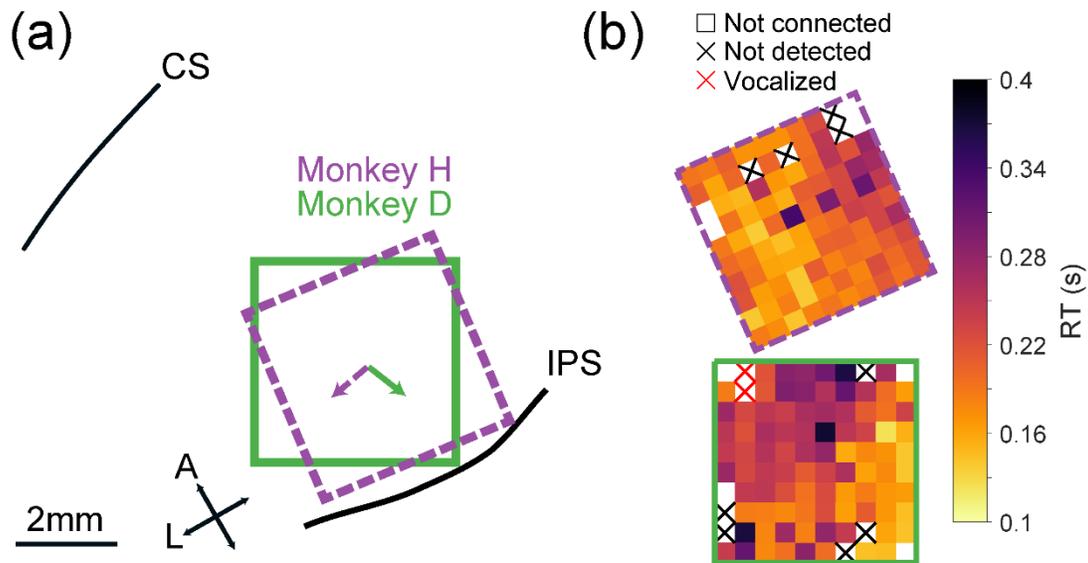


Fig. 3.5. Dependence of reaction time to single-electrode stimulation on electrode position. (a) Schematic of the arrays in cortex for the two monkeys. The purple dashed square shows the array for monkey H, and the green solid square shows the array for monkey D. CS is central sulcus, IPS is intraparietal sulcus, A is anterior, and L is lateral. (b) The reaction time (RT) during single-electrode stimulation for each electrode is shown for monkey H (top) and monkey D (bottom). The color of each electrode denotes the RT to stimulation through that electrode. White squares represent electrodes that were not connected. The monkeys did not respond to stimulation through electrodes denoted with black X's. Red X's denote electrodes which resulted in vocalization.

Reaction time in response to high-amplitude sICMS

A possible explanation for the slow RT in response to ICMS is that the evoked sensation is rather weak, like that of a small force-pulse perturbation. As in our results for the mechanical cue and sICMS, reaction time is typically a saturating function of stimulus intensity for all parameters (Fig. 3.2, Fig. 3.3) (Pins and Bonnet 1996). While large currents may lead to a faster RT, there is a limit to the current that can be delivered safely through any given electrode. To test the effect of even larger currents on RT, we used currents up to 200 μA on single electrodes, near the maximum of our stimulator. The RT for currents from 40 to 200 μA is shown for two example channels in monkey D (Fig. 3.6(a)). RT decreased with increased current as it did for all channels, approaching that of the mechanical cue (black line). Fig. 3.6(b) shows the RT to sICMS at 200 μA across all electrodes compared to that of the mechanical cue in the corresponding session. Across electrodes, the mean RT to sICMS at 200 μA was about 4 ms slower than that to the mechanical cue for monkey H, and about 29 ms slower for monkey D. The RT to 200 μA stimulation was not significantly different than the RT to the mechanical cue in either monkey ($p = 0.38$ for monkey H, $p = 0.11$ for monkey D, Welch's t-test), though we performed only a small number of experiments to limit any damage done to the tissue surrounding the electrodes.

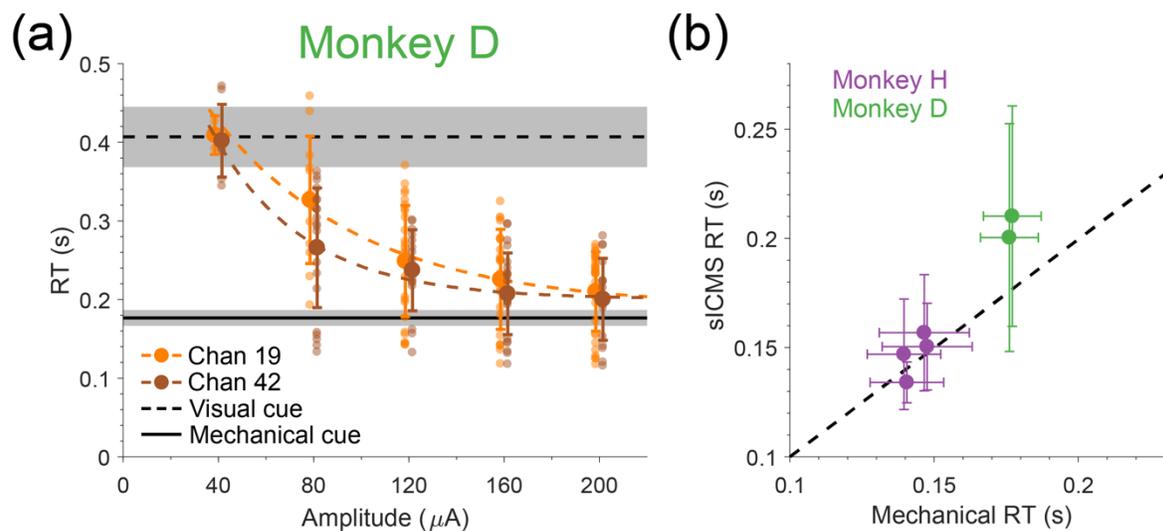


Fig. 3.6. Reaction time to single-electrode stimulation at large currents. (a) The reaction times (RTs) to stimulation amplitudes up to 200 μA are shown for two example channels in monkey D using the same figure conventions and mechanical cue parameters as in Fig. 3.2. (b) The mean RT to stimulation at 200 μA for 4 channels in monkey H (purple) and 2 channels in monkey D (green) compared to the mean mechanical cue RT during corresponding sessions are shown. RTs to the mechanical cue have been jittered slightly to avoid overlap. Black dashed line represents unity. All error bars show standard deviation.

Reaction time in response to multi-electrode stimulation

Because the RT continued to decrease for currents up to 200 μA during sICMS for at least some of the channels, it may be that increasing the stimulation amplitude further would reliably elicit RTs as fast as or faster than the mechanical cue. However, since the current was already at the limit of what is considered safe for the electrode-tissue interface (Chen et al. 2014), we did not want to increase it further. Instead, we stimulated on multiple electrodes simultaneously in order to inject a large amount of total current while keeping the current through each electrode small. We randomly selected sets of 16 electrodes in each monkey and measured the RT to this mICMS. We used a train length of 120 ms for both monkeys, matching the duration of the mechanical cue. RTs for representative sets of electrodes at various total currents (160-800 μA) are shown in Fig. 3.7(a)-(b). Insets show the locations of the electrodes on the array, oriented as in Fig. 3.5. For every set of electrodes in monkey H, RT decreased with increasing total current. For monkey D, the RT decreased up to a total current of 480 μA . Unexpectedly, the RT began to increase for currents above 480 μA for 3 out of 4 electrode sets tested.

The mean RT for each set of electrodes compared to the mean RT to the mechanical cue in the corresponding session is shown in Fig. 3.7(c) for each monkey. For each electrode set, the fastest mean RT across the tested total currents is shown. For monkey H, stimulating with 16 electrodes resulted in significantly faster RTs than the mechanical cue ($p = 0.0037$, Welch's t-test), with a mean difference of 27 ms. There was no significant difference between the RT to 16-electrode stimulation and the mechanical cue in monkey D ($p = 0.965$, Welch's t-test).

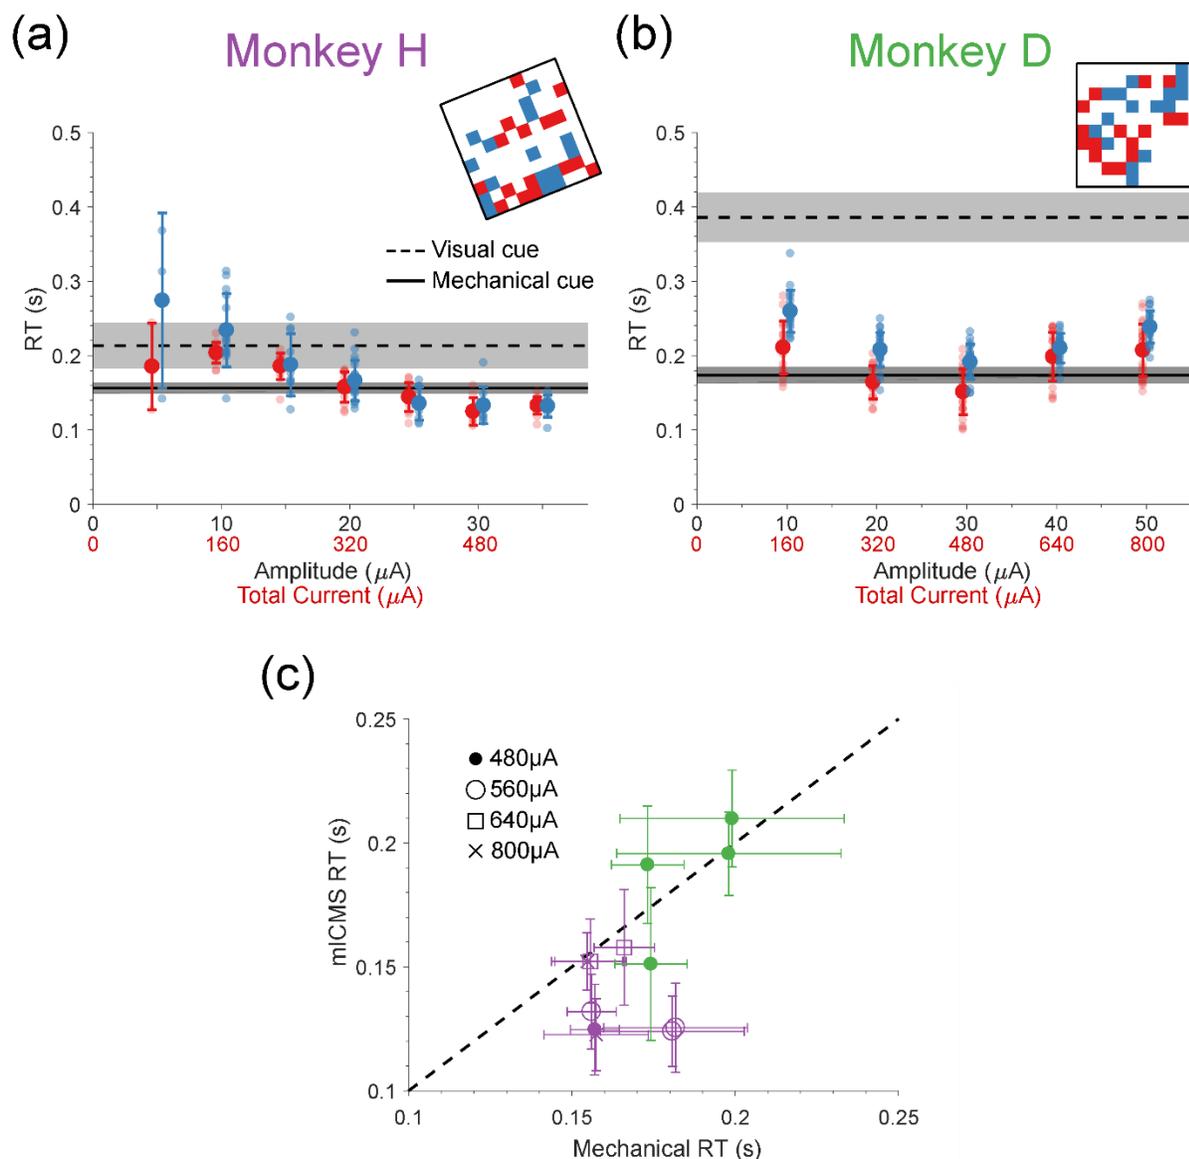


Fig. 3.7. Reaction time to 16-electrode stimulation. The reaction time (RT) to simultaneous stimulation of two example sets of 16 electrodes at various total currents are shown for (a) monkey H and (b) monkey D. The same figure conventions and mechanical cue parameters as Fig. 3.2 are used, except without exponential fits. Insets denote the position of the electrodes on the array for each set, oriented as in Fig. 3.5c. (c) The mean RT to stimulation of 8 sets of electrodes in monkey H (purple) and 4 sets in monkey D (green) compared to the mean mechanical cue RT during corresponding sessions are shown. The amplitude which resulted in the fastest mean RT was used. RTs to the mechanical cue have been jittered slightly to avoid overlap. Black dashed line represents unity. All error bars show standard deviation.

Interestingly, neither monkey responded to stimulation with 80 μA distributed over 16 electrodes, even though 80 μA was typically detected during single-electrode stimulation. This implies that total current does not fully predict the resulting RT. We investigated the effect of total current and number of stimulation electrodes more thoroughly (Fig. 3.8). On each trial, we chose a random number of electrodes (4, 6, 8, 12, and 24) and total current (240, 360, and 480 μA). Then, we chose a random set of electrodes and distributed the total current equally across those electrodes. Because our stimulator was limited to 16 simultaneous channels, we stimulated the electrodes in two groups, separated by a 100 μs lag, even when stimulating on less than 16 channels. Again, we used a train length of 120 ms to match the duration of the limb perturbation. Resulting RTs to all combinations of electrodes and currents are shown in Fig. 3.8 for four sessions for monkey H and two sessions for monkey D. The RTs to the mechanical cue (black solid line) and the visual cue (black dashed line) pooled across sessions are also shown for each monkey. We used an ANOVA, combining data across monkeys, total currents, and number of electrodes, to determine the effect of these parameters on the resulting RT. RT consistently decreased when we stimulated with more current for the same number of electrodes ($F(1,26) = 17.5$, $p = 2.9\text{E-}4$). Interestingly, the RT increased as the number of electrodes increased ($F(1,26) = 19.7$, $p = 1.5\text{E-}4$), an effect that was more pronounced at smaller total currents.

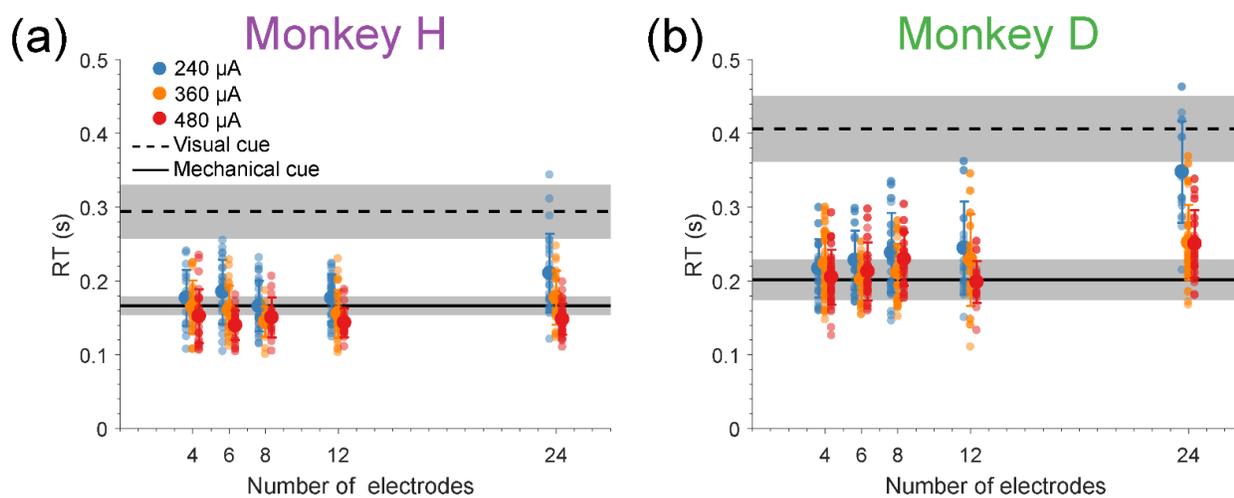


Fig. 3.8. Dependence of multi-electrode reaction time on total current and number of electrodes. The mean reaction time (RT) to multi-electrode stimulation with different total currents and number of electrodes is shown for (a) monkey H and (b) monkey D. The same figure conventions and mechanical cue parameters as Fig. 3.2 are used.

Comparison of mICMS to sICMS

We wanted to determine whether the apparent advantage mICMS offers might be eliminated when compared to the single most sensitive electrode in the group. To test this, we measured the RT to stimulation on groups of 2 or 3 electrodes at 100 μ A per electrode. We compared the resulting mICMS RT to the fastest 100 μ A sICMS RT within each group (Fig. 3.9(a)). sICMS RT was measured in sessions 1-2 weeks prior to the corresponding mICMS RT in monkey H and less than a week prior in monkey D. We used a train length of 120 ms for monkey H and 200 ms for monkey D to match the train lengths used when testing sICMS. Across both monkeys, we found that mICMS resulted in significantly faster RTs than sICMS ($p \ll 0.001$, paired t-test).

Effect of inter-electrode distance on mICMS reaction time

Previous experiments showed that the distance between electrodes does not affect how they interact during mICMS (Kim, Callier, Tabot, Tenore, et al. 2015; Zaaimi et al. 2013). However, this may not be true when stimulating with the large currents in this study, which will activate a larger population of neurons surrounding each electrode (Stoney, Thompson, and Asanuma 1968). To test this, we paired the groups of electrodes from Fig. 3.9(a), such that one group contained only adjacent electrodes, while the other was composed of non-adjacent electrodes. These groups were tested in adjacent blocks to decrease any intra-session effect. Electrodes within paired groups were matched to have approximately equal RTs during sICMS, and we measured the RT when stimulating all electrodes within a group at 100 μ A per electrode and 120 ms (monkey H) and 200 ms (monkey D). Fig. 3.9(b) shows the RT to stimulation of adjacent groups of electrodes and the corresponding non-adjacent groups. Across multiple sessions, both monkeys, and 30 pairs

of electrode groups, the mean RT to stimulation on adjacent electrodes was 7 ms slower than the mean RT to stimulation on non-adjacent electrodes, a difference which was not significant ($p = 0.071$, paired t-test).

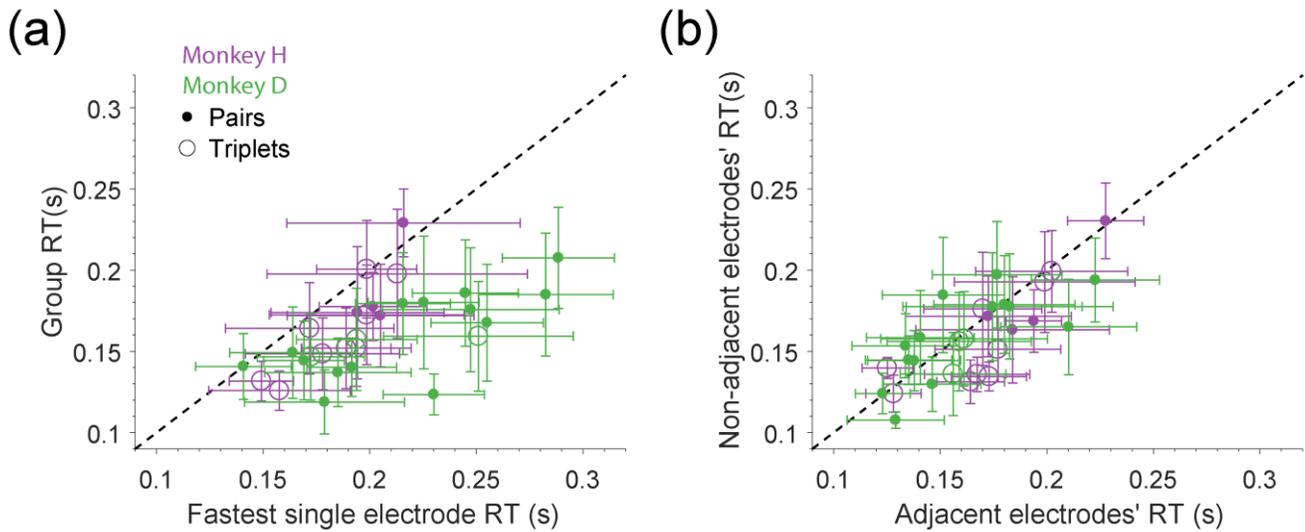


Fig. 3.9. Effect of distance between electrodes on multi-electrode reaction time. (a) The mean reaction time (RT) to simultaneous stimulation on pairs (filled circles) or triplets (large open circles) of electrodes with $100 \mu\text{A}$ per electrode is compared to the fastest RT during single-electrode stimulation at $100 \mu\text{A}$ for electrodes within the group of electrodes. (b) The mean RT to simultaneous stimulation on groups of adjacent electrodes compared to non-adjacent electrodes. Groups of electrodes were paired such that the RT to single-electrode stimulation for electrodes within each pair was approximately the same.

Discussion

Summary of results

In a series of experiments in two monkeys, we investigated the RT in response to both sICMS and mICMS applied through multi-electrode arrays implanted in area 2 of somatosensory cortex. We found that the RT to sICMS typically decreased with increased stimulation amplitude, frequency, and train length. Even at large stimulation parameters, the RT for most individual electrodes was slower than to mechanical cues. Increasing the stimulation amplitude to 200 μ A resulted in RTs only slightly slower than that to mechanical cues, though currents this large may cause damage to tissue surrounding the electrodes. However, mICMS elicited RTs as fast as or faster than mechanical cues with safe levels of current through each electrode. Together, these results suggest that it may be possible to use mICMS to provide fast, artificial feedback, and thereby restore proprioception.

Reaction time to single-electrode stimulation

For most modalities, RT decreases with increased stimulus intensity, settling at some minimum latency (Pins and Bonnet 1996) (Fig. 3.2). We found that increasing the stimulation amplitude, frequency, and train length of ICMS all resulted in faster RTs (Fig. 3.3), consistent with their effect on detection thresholds (Kim, Callier, Tabot, Gaunt, et al. 2015; Butovas and Schwarz 2007). Nonetheless, the RT to sICMS remained slower than that to limb perturbations (Fig. 3.4), implying that the evoked sensation for many electrodes was still weaker than natural stimuli. We wanted to determine a rough estimate of the magnitude of sensation caused by sICMS. To do so, we assumed that the magnitude of sensation caused by sICMS was roughly equal to the force of the mechanical cue which resulted in the same RT, after adding 20 ms to account for the conduction

delay between the periphery and cortex. The mean RT to sICMS at 100 μ A corresponded to a 0.2 N mechanical force for monkey H and 0.7 N for monkey D, moving the hand about 0.1 and 1.5 cm respectively. This rough estimate implies that sICMS does not cause a large sensation. This direct comparison is difficult for a variety of reasons, including the fact that the mechanical cues move the whole arm while sICMS may elicit a sensation only about a small part of the arm (Salas et al. 2018).

While sICMS with large currents might provide fast feedback (Fig. 3.6), this approach may damage tissue and cause neural loss surrounding the stimulated electrodes (McCreery, Pikov, and Troyk 2010). McCreery et al. found that stimulation at even 20 μ A applied eight hours a day for 30 days caused loss of neurons around the stimulated electrodes. The amount of damage depended on the current, as 10 μ A applied for the same duration did not cause significant neural loss. In contrast, a recent study found that ICMS applied daily for four hours, five days per week, for six months, caused only a small amount of neural loss, even at 100 μ A (Rajan et al. 2015). Instead, most of the tissue damage was due to implanting, residence, and explanting the array. Results from the 100 μ A condition in this study should be interpreted cautiously, as only a few tissue samples were analyzed. Overall, while sICMS with large currents might feasibly provide fast feedback, the safety of this approach is a concern which warrants further study.

Furthermore, large currents may reduce the effectiveness of feedback, well before causing damage. Stimulation applied to area MT, an area involved in processing visual motion, has been used to predictably bias a monkey's report of the direction of a noisy motion signal (Murasugi, Salzman, and Newsome 1993). With increasing stimulation amplitude, the bias became larger, up to currents of about 80 μ A. At that point, the monkeys' ability to identify the correct direction of motion dropped to chance. Since increasing the stimulation current leads to direct activation of

neurons farther away from the stimulated electrode (Stoney et al. 1968), this result is most likely explained by the increased activation of neurons with differing receptive field properties.

Effect of electrode location on sICMS reaction time

In both monkeys, the location of the electrode affected the RT (Fig. 3.5). One possible explanation for the change in RT across the array might be the depth of the electrode tips, the result of the array not conforming perfectly to the curved cortical surface. We would expect that any depth-related changes would be primarily in the anterior-posterior direction, with increasing distance from the greater convexity of the IPS. In monkey D, we observed that electrodes nearer the IPS had faster RTs, electrodes which would likely be shallower compared to more distant electrodes. This result stands in contrast with studies where stimulation in deeper cortical layers resulted in lower detection thresholds than in shallower layers (Tehovnik and Slocum 2009; Koivuniemi and Otto 2011), though one study observed the opposite effect (DeYoe, Lewine, and Doty 2005). Computational models of cortical stimulation predict that neurons in layer 5 have a lower activation threshold than those in layer 2/3, implying that stimulation in deeper layers would activate more neurons than stimulation in shallower layers (Aberra, Peterchev, and Grill 2018). The depth of the electrode tip would probably not explain the medial-lateral gradient in monkey H. Because our arrays are in the proximal arm area, it may be that more lateral electrodes elicited sensations closer to the hand than did medial electrodes. Perhaps monkey H responded faster to cues near his hand than on his arm. It is likely that a combination of factors, including electrode depth and sensation location, affect RT. More experiments will be required to determine those factors.

Reaction time to multi-electrode stimulation

By stimulating on multiple electrodes simultaneously, we were able to evoke RTs in monkey H about 27ms faster, on average, than the mechanical cue, roughly the conduction delay between the periphery and cortex (Fig. 3.7). In monkey D, there was no difference between the mechanical cue RT and mICMS at 480 μ A. Unexpectedly RT increased for total currents beyond 480 μ A. We have no clear explanation for this observation, but speculate that it might be due to eliciting something like a startle response, or the result of mICMS actually delaying the monkey's reach planning, in a manner similar to stimulation in pre-motor cortex (Churchland and Shenoy 2007). This effect may also be due to increased activation of inhibitory circuits at higher currents, as interneurons likely have higher activation thresholds than do pyramidal neurons (Overstreet, Klein, and Helms Tillery 2013). We used random sets of electrodes during mICMS. It may be that choosing electrodes based on the sensation they elicit would have resulted in faster RTs in monkey D. Regardless, the mean mICMS RT in monkey D was considerably faster than sICMS. Even at 200 μ A, sICMS RT was about 30 ms slower than in response to the mechanical cue.

While mICMS can elicit fast RTs with relatively low current through each electrode, there could still be the concern that tissue damage may be caused by the summed current at a return electrode. In our study, the current was returned through a large titanium pedestal placed on the monkey's skull, and the current density through the pedestal was very low. mICMS has also induced effects such as discomfort (suggested by vocalization of the animal), muscle twitches, and seizures in previous studies (Parker et al. 2011; Chen et al. 2014). In these studies, transient effects occurred when a large amount of charge was injected simultaneously, or when electrodes were damaged before implantation. We did not induce any such transient effects in our experiments during mICMS, possibly because the current we injected was smaller than in the reported cases,

where at least 1600 μA total was required to induce such effects. Two individual electrodes did cause monkey D to vocalize. Since these two electrodes were positioned next to each other, it may be that this area of the array was damaged or happened to be in a sensitive region of cortex. These explanations seem unlikely as one of the electrodes was actively recording from a neuron, the impedances of these electrodes were similar to other electrodes on the array ($\sim 50\text{k}\Omega$), and electrodes surrounding these two did not elicit any transient effects.

Effect of number of stimulation electrodes and total current

We found that mICMS did decrease the RT compared to sICMS (Fig. 3.9(a)). This is consistent with previous studies measuring the detection threshold due to stimulation in area 3b/1 (Kim, Callier, Tabot, Tenore, et al. 2015) and area 2 (Zaaimi et al. 2013). However, at a constant total current, we found that increasing the number of electrodes resulted in slower RTs (Fig. 3.8). This effect may be due to the current on some electrodes falling below an activation threshold required to contribute to the overall sensation. This seems to conflict with our earlier observation that sub-detection threshold currents on multiple electrodes sum supralinearly (Zaaimi et al. 2013). Instead, our result may be more similar to what Kim et al. found: each electrode had an independent effect on sensitivity, even for sub-threshold currents (Kim, Callier, Tabot, Tenore, et al. 2015). However, it is difficult to compare the current study to that of either Kim et al. or the earlier sub-threshold detection study of Zaaimi et al., as our current study did not measure detection thresholds during single-electrode stimulation. It is apparent that the current-per-electrode needs to be carefully considered when designing stimulation patterns, as current that is too high may damage tissue while too low a current may not provide robust sensation.

Effect of distance between electrodes

At the largest currents in our study, adjacent electrodes may have activated overlapping populations of neurons. This might be expected to cause weaker sensations than from non-adjacent electrodes. However, previous studies have suggested that this is not the case (Zaaimi et al. 2013; Kim, Callier, Tabot, Tenore, et al. 2015). We found that groups of electrodes that were adjacent elicited slightly, though not significantly, slower RTs than non-adjacent groups of electrodes (Fig. 3.9(b)), consistent with the earlier studies. With at least 400 μm between electrodes, as with a Utah array, the distance between electrodes does not seem to be an important consideration when designing stimulation patterns.

Implications for neuroprosthetics

Current BMIs rely solely on visual feedback to correct movements, which is too slow to update rapid reaches (Ghez et al. 1990; Sainburg et al. 1995). Restoring fast somatosensory feedback to users should improve prosthetic control (Shanechi et al. 2017) and may enable users to develop a stronger sense of embodiment of the prosthesis (Shimada, 2009). Even though sICMS and stimulation through mini-electrocorticography arrays can elicit sensations spanning quite a range of different qualities (Flesher et al. 2016; Lee et al. 2018; Salas et al. 2018), these approaches may not be able to provide fast somatosensory feedback (Godlove, Whaite, and Batista 2014; Caldwell et al. 2019). We show that mICMS can be used to trigger movement at very short latencies, making it potentially suitable for providing rapid somatosensory feedback.

In addition to the more robust sensations it appears to provide, mICMS seems well suited to recreating the spatially complex patterns of cortical activity that are characteristic of the somatosensory response to limb movement (Soso and Fetz 1980; Tomlinson and Miller 2016;

Prud'homme and Kalaska 1994). One such proprioceptive interface provided target-proximity feedback, using eight electrodes with arbitrarily assigned “preferred direction” (error) vectors (Dadarlat, O'Doherty, and Sabes 2015). Two monkeys learned to use the system, but required months of training and still made movements that were much slower than typical. A different, more biomimetic approach from our group used stimulation on small sets of electrodes with similar directional tuning properties in an effort to elicit naturalistic sensations of limb movement. This approach caused a predictable bias of a monkey's perception of a coincident mechanical perturbation without any training (Tomlinson and Miller 2016). However, it failed to do so in three other monkeys. Although there is justifiable concern that synchronous stimulation of many neurons with mICMS may evoke artificial sensations (Tan et al. 2014), it is hard to imagine, given existing methods, an alternative means to activate the cortical circuits needed to mimic the spatially complex patterns of neural activity evoked by limb movements. Unlike the paradigms that have been used in an effort to mimic tactile stimulation with single electrodes (Romo et al. 2000; O'Doherty et al. 2012; Tabot et al. 2013), multiple electrodes will likely be required to provide useful proprioceptive feedback.

Most applications of ICMS for touch have used single electrodes to deliver simple, punctate sensations (Tabot et al. 2013; O'Doherty et al. 2011b). However, any realistic object manipulation or haptic exploration will result in many contacts across the hand and fingers, possibly even the forearm. To provide robust cutaneous sensations about the whole hand, stimulation could be applied through multiple sets of electrodes, where each set elicits a localized sensation. This approach was tested recently in a spinal cord injured patient, where force applied to the fingers of a prosthetic hand was mapped to stimulation of sets of electrodes that evoked sensations in the corresponding finger of the patient (Flesher et al. 2016). With this interface, the

patient could correctly identify which robotic fingers were touched, even when they were touched in pairs, although the latter was less accurate. This approach was extended to provide haptic feedback while a patient controlled movement of the prosthesis (Flesher et al. 2019). With haptic feedback provided by mICMS and visual feedback, the patient was able to grasp objects faster than with visual feedback alone. Whether for touch or proprioception, future neuroprosthetics will most likely need spatially and temporally patterned mICMS to provide natural, robust somatosensory sensation. Such results will likely be necessary to improve motor control.

Chapter 4 - Modeling proprioceptive feedback evoked by intracortical microstimulation using a topographical variational autoencoder

Joseph T. Sombeck, Kyle Blum, Nathan Schimpf, Max Grogan, A. Aldo Faisal, Lee E. Miller

Foreword

We previously tried to measure the naturalistic effects of ICMS in area 2 and found conflicting results across monkeys. In one monkey, stimulation through 4 electrodes predictably biased the monkey's perception of the direction of a perturbation applied to his hand. In three other monkeys, this approach biased the monkeys' perceptions, but not in a predictable manner. The differences across monkeys could be due to different topographical organizations in area 2 and electrode placement. The purpose of this project was to simulate the proprioceptive effect of ICMS in a model of proprioceptive cortex to explore the effect of topography and electrode placement on the evoked sensation. By using a topographical variational autoencoder that included lateral effects to impose topography, I generated an 80x80 grid of artificial neurons in which neurons with similar preferred directions formed clusters. Artificial stimulation applied simultaneously at the center of many of these clusters resulted in effects of similar magnitude to natural limb movements and in directions predicted by the mean PD of the stimulated sites. This result implies that stimulation will need to be applied through many carefully chosen electrodes to evoke naturalistic sensations of limb movement. We are preparing this manuscript for submission.

Abstract

Tetraplegic patients can make visually guided reaches using efferent brain-computer interfaces (BCIs). However, their reaches are slow and imprecise, like the deficits of people who can generate movement but have lost somatosensation. While intracortical microstimulation (ICMS) has been used to provide tactile feedback to both monkeys and human BCI users, restoring proprioceptive feedback with ICMS has proven more difficult. The difference in success between touch and proprioception may be because early tactile cortical areas have a distinct somatotopy while proprioceptive cortical areas do not. Without a well-defined somatotopy, ICMS may activate neurons with very different encoding properties, resulting in unnatural sensations or perceived hand movements in unpredictable directions. To explore these effects, we trained a topographical variational autoencoder to regenerate arm joint velocities. The network included lateral excitation and inhibition, which caused clusters of neurons with similar PDs to form in the map. We simulated the effect of ICMS on movement-related map activity, which we interpreted by reading out the difference between decoded hand velocity with and without ICMS. Increasing ICMS amplitude increased both the magnitude of this difference-vector and the directional error between it and the PD of neurons at the stimulus site. At 100 μA , single-location stimulation evoked an effect much smaller than that of a limb movement and with a mean directional error of $\sim 85^\circ$ across tested sites. To reduce the error substantially, we needed to apply low-current stimulation within 100 μm of the center of a cluster of neurons. Stimulating many cluster centers simultaneously resulted in a mean directional error as low as 10° and a magnitude comparable to limb movements. The results from this study show that future proprioceptive interfaces will need to stimulate on numerous, carefully chosen electrodes, each with sufficiently small currents to activate neurons with homogeneous PDs.

Introduction

Current state-of-the-art brain-computer interfaces (BCIs) have progressed to a point where persons with paralysis due to spinal cord injury can reach toward and grasp objects using highly anthropomorphic robotic arms (Collinger et al. 2013; Wodlinger et al. 2014; Hochberg et al. 2012; Flesher et al. 2021). In addition to restoring movement per se, BCIs will need to restore somatosensation, without which, normally-controlled movements are impossible (Ghez, Gordon, and Ghilardi 1995; Sainburg, Poizner, and Ghez 1993). While good progress has been made using intracortical microstimulation (ICMS) to restore a modicum of contact location and force feedback, enabling users to grasp and move objects faster than with visual feedback alone (Flesher et al. 2021), much less progress has been made for proprioception, the sense of position and movement of the body. Limited proprioceptive sensations have been evoked by ICMS, as reported by a human participant, though these sensations are typically small and correspond to only a small portion of the whole arm (Salas et al. 2018). Proprioceptive-like interfaces using an arbitrary stimulus require many months for monkeys to learn (Dadarlat, O'Doherty, and Sabes 2015), and more biomimetic approaches have not been reliable (Tomlinson and Miller 2016). Without this advance, reaches made by with BCIs will presumably be slow and imprecise, like those made by people who have lost proprioception (Sainburg et al. 1995; Ghez, Gordon, and Ghilardi 1995).

The large difference in efficacy between ICMS for touch and proprioception may be due to the different topographical organizations of the corresponding cortical areas. In early tactile cortical areas, such as area 3b and 1, neurons with similar receptive field properties tend to be located close together (Mountcastle 1957; Callier, Suresh, and Bensmaia 2019; Chen et al. 2001). Stimulation in these areas typically results in sensations localized at the receptive fields recorded on the stimulated electrodes (Tabot et al. 2013; Berg et al. 2013). We frequently encounter simple

punctate tactile stimuli in the real world, as well as the lab, that activate a very circumscribed cortical area. In contrast, limb movements, even simple ones, evoke a complex spatial pattern of activity across proprioceptive cortex. This difference is due, in part, to the nature of the primary proprioceptive receptors, muscle spindles and Golgi tendon organs embedded in muscles that span as many as two and even three joints (Proske and Gandevia 2012; Houk, Rymer, and Crago 1981), and are then mapped onto the two-dimensional cortical surface. This results in a less well-defined somatotopy in proprioceptive cortical areas than in tactile areas (Iwamura, Iriki, and Tanaka 1994; Pons et al. 1985). The complex organization in proprioceptive areas may be a key reason why evoking naturalistic sensations of proprioception has been difficult.

Receptive fields in the tactile and visual systems have been studied for decades by monitoring the effects of well-localized (in both time and space) stimuli (Hubel and Wiesel 1962; Mountcastle 1957; Ringach 2004; Gilbert and Wiesel 1992; Sripathi et al. 2006; DiCarlo, Johnson, and Hsiao 1998). This approach is largely unavailable to the study of proprioception, which must rely on active or passive limb movements that stretch multiple muscles simultaneously. In this study, we made use of a self-organizing map (“topoVAE”) that was recently developed to study the organization of proprioceptive cortex (Blum, Grogan, et al. 2021). The topoVAE generates a cortical layer which embeds arm kinematic information while recreating three properties of neurons in actual S1. Just as actual proprioceptive neurons, the firing of artificial neurons in the cortical layer is sinusoidally tuned to movement direction (Prud'homme and Kalaska 1994; London and Miller 2012), artificial neurons with similar preferred directions (PDs) tend to be clustered together (Weber et al. 2011), and the overall PD distribution tends to be bimodal, likely as a result of limb biomechanics (Versteeg, Chowdhury, and Miller 2021).

The topoVAE allowed us to simulate the proprioceptive sensations evoked by ICMS on top of those due to on-going movement. We used a separate biophysical model to estimate the evoked firing rates of neurons based on their distance from the stimulus site and the stimulus intensity (Kumaravelu et al. 2022). We inferred the effect of stimulation by measuring the difference between the predicted hand velocity from the cortical map with and without stimulation. The magnitude of this difference-vector and the angle between it and the PD measured at the stimulus site increased with increasing stimulus amplitude. Stimulating multiple locations simultaneously with small amplitudes reduced this directional error, but it remained large. To reduce the error substantially, we needed to apply stimulation within 100 μm of the center of a cluster of neurons with homogeneous PDs. The results from this study show that future proprioceptive interfaces will need to stimulate on multiple, carefully chosen electrodes, each with sufficiently small currents to activate neurons with homogeneous PDs.

Methods

Topographical variational autoencoder

We trained a topographical variational autoencoder (topoVAE) to reconstruct arm joint velocities while capturing the stochasticity inherent to cortical neurons (for more details, see (Blum, Grogan, et al. 2021)). The encoder of the topoVAE contained two fully connected layers of size 20 and 40 neurons, using tanh activation functions, and a final readout layer using ReLU activation functions to prevent firing rates below 0 (Fig. 4.1). Based on the density of cortical neurons, we estimated that an 80x80 grid of artificial output neurons would span the 4x4 mm area of a Utah microelectrode array (Blackrock Neurotech, Salt Lake City, UT; see (Blum, Grogan, et al. 2021) for more details). To model lateral connectivity between neurons in the cortical layer, we

included a Mexican-hat function in the cost function (Amari 1977). The lateral effect between two neurons (γ) was:

$$\gamma = \left(1 - \frac{d^2}{2\sigma^2}\right) e^{-\frac{d^2}{2\sigma^2}},$$

where d is the Euclidean distance between the neurons. With $\sigma = 100 \mu\text{m}$, the lateral interaction transitions from excitation to inhibition, reaching a maximum at $200 \mu\text{m}$ vanishing beyond about $400 \mu\text{m}$ (Fig. 4.2). We modelled the firing of neurons in the cortical map as Poisson processes to capture the inherent stochasticity of cortical neurons. To prevent sparse and implausibly high firing rates, we included a KL divergence term in the loss function that penalized firing rates far from the expected distribution (Poisson, with mean rate of 20 Hz). We decoded joint angular velocities from the simulated firing rates using a linear readout. To analyze the effect of stimulation with respect to hand movement we converted joint angular velocities to Cartesian hand velocities.

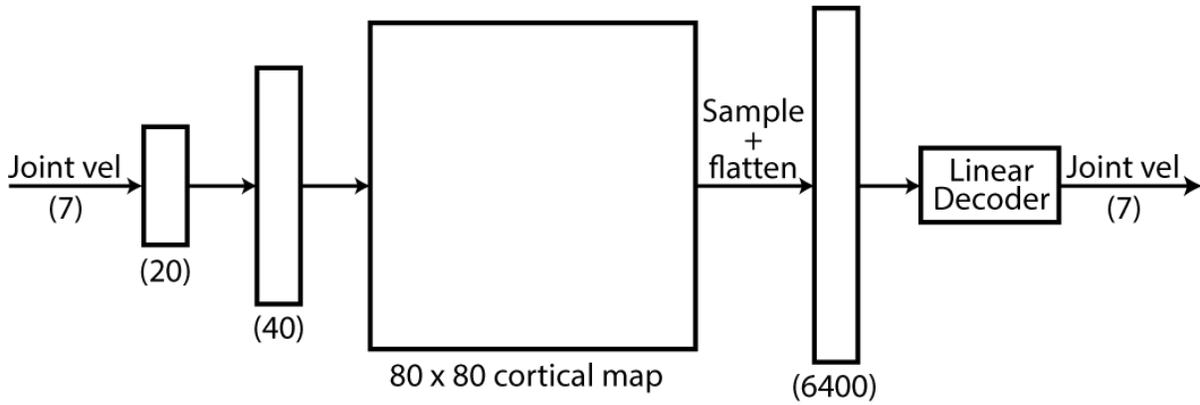


Fig. 4.1. Architecture of the topoVAE model. The topoVAE encoded joint velocities via two fully connected feed-forward layers of size 20 and 40 neurons using the tanh activation function. The final readout layer of the encoder consisted of an 80x80 grid of neurons. To prevent firing rates below 0 in this cortical map, cortical neurons used a ReLU activation function. We sampled firing rates using a Poisson model and then linearly decoded joint velocities from these firing rates.

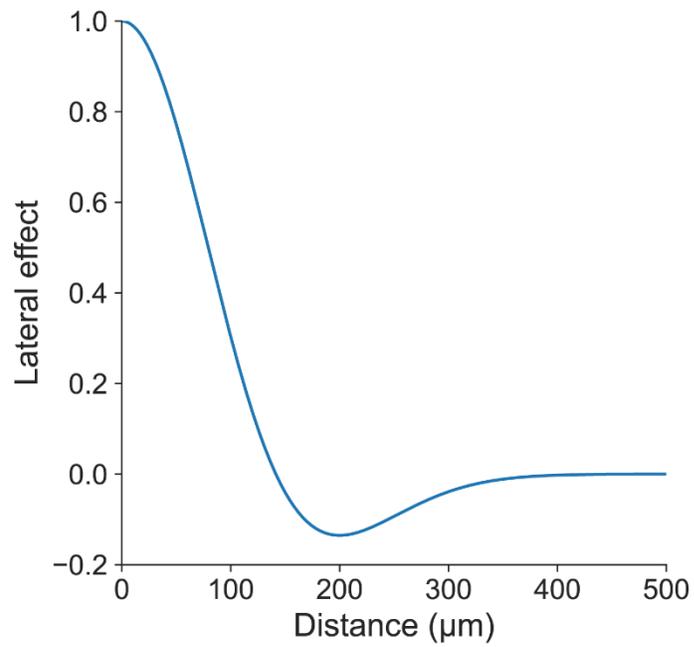


Fig. 4.2. Lateral effect between neurons in the cortical map. We modeled cortical connectivity by including a Mexican-hat function in the cost function. The lateral effect between two neurons consisted of local excitation and intermediate inhibition.

Arm kinematics dataset

We collected data from a single 12.9 kg male rhesus macaque to use as training data for the topoVAE. We performed all procedures in this study in accordance with the Guide for the Care and Use of Laboratory Animals. The institutional animal care and use committee of Northwestern University approved all procedures in this study under protocol #IS00000367.

The monkey held a two-link planar manipulandum and made reaching movements to targets presented on a screen. The monkey needed to reach to four successive randomly located targets to receive a liquid reward. We used a custom motion tracking system to track 10 markers painted on the monkey's right arm (see (Chowdhury, Glaser, and Miller 2020) for more details). We used Opensim (Delp et al. 2007) and a 3D musculoskeletal model of the macaque arm with seven degrees-of-freedom to compute the angular velocities of the shoulder, elbow, and wrist joints, and Cartesian velocities of the hand (Chan and Moran 2006). We binned velocities at 50 ms, then convolved these data with a non-causal Gaussian of width 100 ms. Because the rectangular workspace resulted in more right/left than forward/backward movements, we subsampled the joint angular velocity data to create a more nearly uniform distribution of movement directions. For each trial of simulated ICMS, we randomly selected 250 ms of consecutive data from the full joint angular velocity dataset to generate cortical firing rates.

Training hyperparameters

We z-scored the angular joint velocities and trained the topoVAE with 80% of the subsampled angular velocity dataset. We tested reconstruction accuracy on the remaining 20%. The topoVAE model was implemented in Python (van Rossum 1995) using PyTorch (Paszke et al. 2019). All models were trained for 3,000 epochs using the Adam optimizer with a learning rate

of 10^{-6} and a batch size of 1,024. The cost function weighted the lateral connection term as 20x that of the reconstruction term and the KL divergence as 0.0005x of the reconstruction term, values chosen to put each loss term on roughly the same scale. During training, we used dropout regularization (97%) to prevent overfitting. To compare the firing rates evoked by stimulation to those of the cortical neurons, we needed to convert the instantaneous output of cortical neurons to firing rates in 50 ms bins. We scaled the output of neurons in the cortical map by 10 as this resulted in firing rates that matched that of actual cortical neurons.

We evaluated goodness-of-fit of the reconstructed joint velocities by computing the variance accounted for (VAF), which assesses how well our reconstruction matched the original signal instead of simply how well correlated the signals were (Fagg et al. 2009). After training, the VAF of the reconstructed joint velocities ranged from 0.77 to 0.97 across joints (mean = 0.90) and the VAF of the reconstructed hand velocities was 0.83 for both the forward/backward and leftward/rightward components.

Stimulation experiments and statistical analysis

Because stimulation activates many neurons around the site of stimulation, we characterized the movement-related firing rate at each location with a multi-unit PD computed as a weighted sum of the firing rates of nearby neurons. The weight decreased linearly from 1 to 0 between 0 and 160 μm . The weight for neurons beyond 160 μm was zero. We modeled this weighted firing rate as a function hand velocity using a Poisson generalized linear model:

$$f \sim \text{Poisson}(\lambda), \lambda = e^{X\beta},$$

where f is the firing rate, X is a matrix of Cartesian hand velocities, and β is a matrix of model parameters. We defined the multi-unit PD as the inverse tangent of the ratio of the y- and x-velocity

parameters in β at each location in the map (Chowdhury, Glaser, and Miller 2020; Versteeg et al. 2021).

In all experiments, we artificially stimulated with 200 Hz trains lasting 200 ms (40 pulses). Increasing stimulation frequency increased the magnitude of the effect, but not its direction. When we stimulated multiple locations simultaneously, we first chose a single location, then randomly chose other locations having PDs within ± 23 degrees of the first location. Unless otherwise specified, we repeated each stimulus condition and location 50 times to compute mean effects. We used two-way ANOVAs to analyze our results across conditions. Due to the large number of sets of locations we tested for each experiment, statistical tests returned near zero p-values even when the effect size was small.

Neurons with similar PDs in the cortical layer of the topoVAE formed clusters. Cluster size will presumably interact with stimulus intensity to determine effect magnitude and directional error. While we could have computed new cortical maps with different lateral effect scales (σ in the Mexican-hat function) to test this effect, this would have resulted in different topographical organization due to the stochastic learning process which could affect the efficacy of stimulation. Instead, we simply scaled the dimensions of the original cortical map from 4000 μm to alter the effective size of clusters. When we did so, we adjusted the magnitude of the inferred effect to account for changes in the density of neurons.

Results

Artificial stimulation pipeline

To simulate the effect of stimulation within proprioceptive cortex, we generated an 80x80 cortical map with a topographical variational autoencoder (topoVAE) that included both lateral

inhibition and excitation and was trained to regenerate joint angular velocities during reaching movements (Blum, Grogan, et al. 2021). The map recapitulated a number of the features present in recordings from proprioceptive regions of cortex (Blum, Grogan, et al. 2021). After training the network, we artificially stimulated it using an activation function derived from a biophysical model of cortical stimulation, shown in Fig. 4.3c (Kumaravelu et al. 2022). The ICMS-induced activity was on top of activity due to arm movements that had previously been recorded as a monkey performed a random-target task while holding a robotic manipulandum. For each stimulus trial, we computed firing rates with and without ICMS (Fig. 4.3d). Artificial stimulation at 200 Hz and 20 μ A caused a large increase in the firing rates of neurons within ~ 300 μ m of the site of stimulation. We linearly decoded hand velocity from the firing rates and computed the difference with and without stimulation. We analyzed the magnitude and direction of this difference-vector to quantify the effect of stimulation.

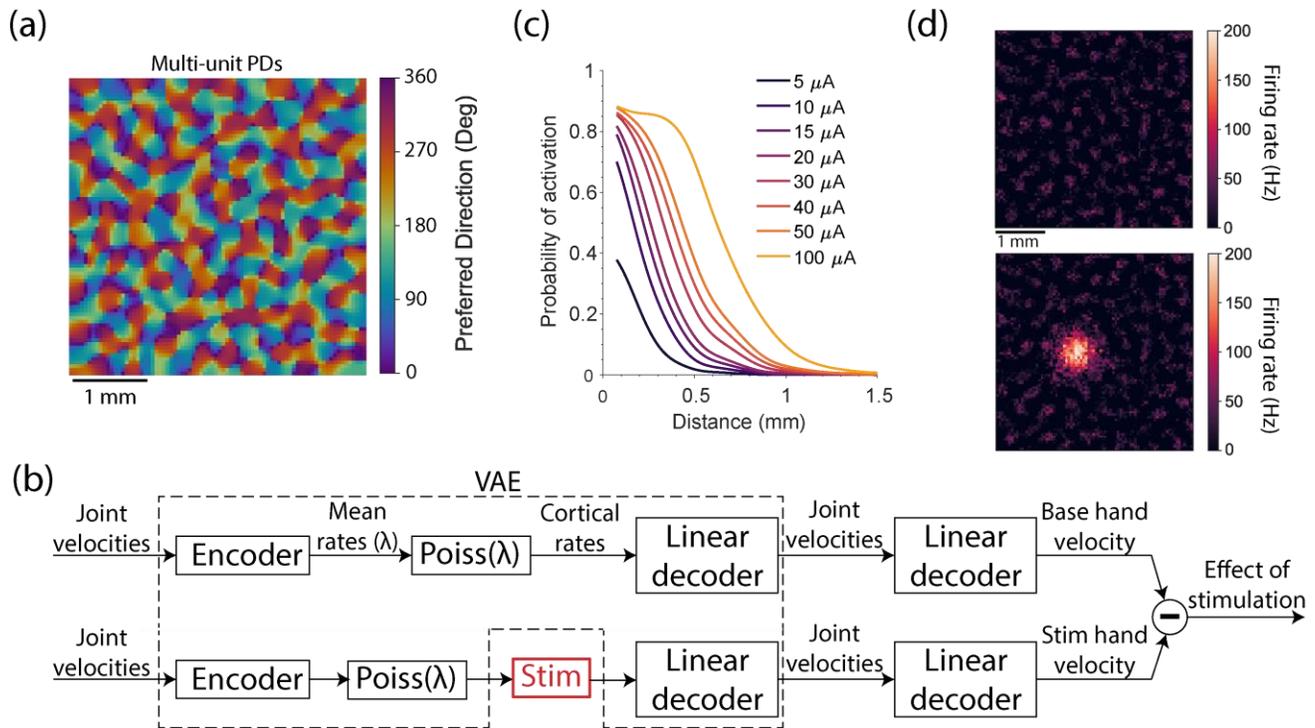


Fig. 4.3. Simulating the effect of ICMS in proprioceptive cortex using a topographic variational autoencoder (topoVAE). (a) Map of the multi-unit PDs of the neurons in the cortical layer of the topoVAE. The color of each cell represents the multi-unit PD calculated at the corresponding location. (b) Pipeline to generate decoded hand velocities during movement. The topoVAE generates two sets of cortical firing rates from the joint velocities, one with, and one without ICMS, and hand velocity is linearly decoded from both. The effect of stimulation was taken as the difference between the velocity with and without stimulation. (c) Activation function used to modulate the firing rates of neurons as a function of ICMS amplitude. (d) The firing rates of artificial neurons averaged over a 50 ms period without (top) and with (bottom) stimulation at 200 Hz and 20 μA . The color of each cell represents the firing rate of the corresponding neuron.

Effect of stimulation at two example locations

In initial experiments, we repeatedly stimulated a single location within a cluster of neurons with similar PDs (Fig. 4.4a, teal cross). When stimulating at $10 \mu\text{A}$, neurons within $\sim 200 \mu\text{m}$ (teal circle) had at least a 50% probability of being activated by each pulse in the stimulation train. The effects of the activated neurons on inferred hand velocity are denoted on the left of Fig. 4.4b as black vectors, with a length determined by the difference in firing rate with and without stimulation. The resulting distribution in this example was weakly bimodal, with the primary mode in direction 210° . The net effect of stimulation, summed across all neurons, was consistently near 210° across trials (Fig. 4.4b, right), near the 184° preferred direction at the center of the stimulus site.

The pink cross in Fig. 4.4a shows a second stimulus location, this one centered on a cortical location with virtually the same PD as in Fig. 4.4b but surrounded by neurons with more heterogeneous PDs. Similar stimulation at this location evoked effects in individual neurons with a broader range of PDs (Fig. 4.4b, left), resulting in a net velocity vector that was shorter than that of the first location (Fig. 4.4b, right). Across trials, the direction of the net effect varied considerably more than that of the first location and was roughly opposite that of the PD at the stimulated location.

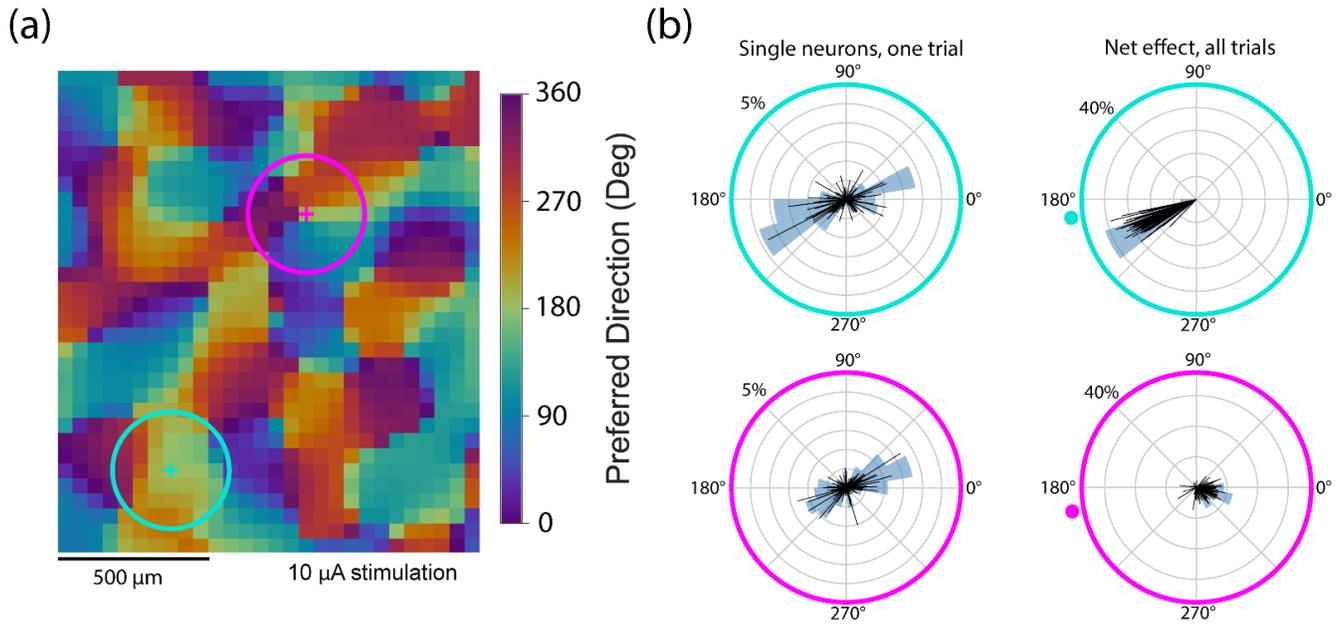


Fig. 4.4. Effect of stimulation at two example locations. (a) Subset of the cortical layer of the topoVAE as in Fig 1a. Teal and pink crosses represent two stimulus locations. Neurons within the colored circles around each cross had at least a 50% chance to be activated by a 10 μ A pulse. (b) The effect of each activated neuron on decoded hand velocity for a single trial when the teal location was stimulated (top, left) and when the pink location was stimulated (bottom, left) is represented by a black vector, and summarized and by the blue histogram. The summed effect of all neurons is represented for each trial by the black vectors and histograms (right). Vectors that reach the outermost circle have a magnitude of 5% (left) or 40% (right) of the median hand speed. Teal and pink dots indicate the PDs of the two stimulus sites.

Effect of increasing amplitude on the magnitude and direction of the effect of stimulation

Increasing ICMS amplitude increases the distance at which neurons will be activated and will, given the topography of the map, increasingly recruit neurons with differing encoding properties (including PDs). We quantified this effect across 250 stimulation locations and a wide range of amplitudes (5-80 μA), computing the mean effect at each location across 50 repetitions. The magnitude of the velocity readout effect increased sublinearly with increased ICMS amplitude (Fig. 4.5a). To analyze the effect across locations with different PDs, we computed the directional error between the PD of the stimulus site and the direction of the velocity readout. This directional error is shown for all trials during 5 μA stimulation in Fig. 4.5b. Stimulation resulted in effects that were slightly biased toward the PD at the stimulus site. The mean directional error is shown across a wide range of currents in Fig. 4.5c. Compared to 0 μA , where errors are due to inherent stochasticity in the cortical firing rates, stimulation at 5 μA significantly lowered the directional error (paired t-test, $p \cong 0$), resulting in effect directions that were nearer to the PD of the stimulated location than a random effect. As amplitude increased from 5 μA , directional error also increased.

While larger currents are necessary to induce larger sensation, the larger directional error they induce limits the utility of that approach. We wondered whether distributing current over multiple electrodes could evoke large effects while maintaining low directional error. To test this, we simultaneously stimulated 250 sets of four locations, with each site in a set having a similar PD (tolerance of $\pm 23^\circ$). As expected, increasing the current at each location resulted in larger effect magnitude (Fig. 4.5a). Using a two-way ANOVA, we determined that the magnitude of the effect increased significantly with both increased total current ($F(1, 2996) = 7450, p \cong 0$) and number of stimulus sites ($F(1, 2996) = 1139, p \cong 0$). Additionally, the effect of increasing

amplitude was larger when the number of simultaneous stimulus sites was increased from 1 to 4 ($F(1,2996) = 193, p \cong 0$). However, quadrupling the number of simultaneous stimulation sites did not quadruple the magnitude of the effect, likely due to overlapping population of activated neurons from the four sites.

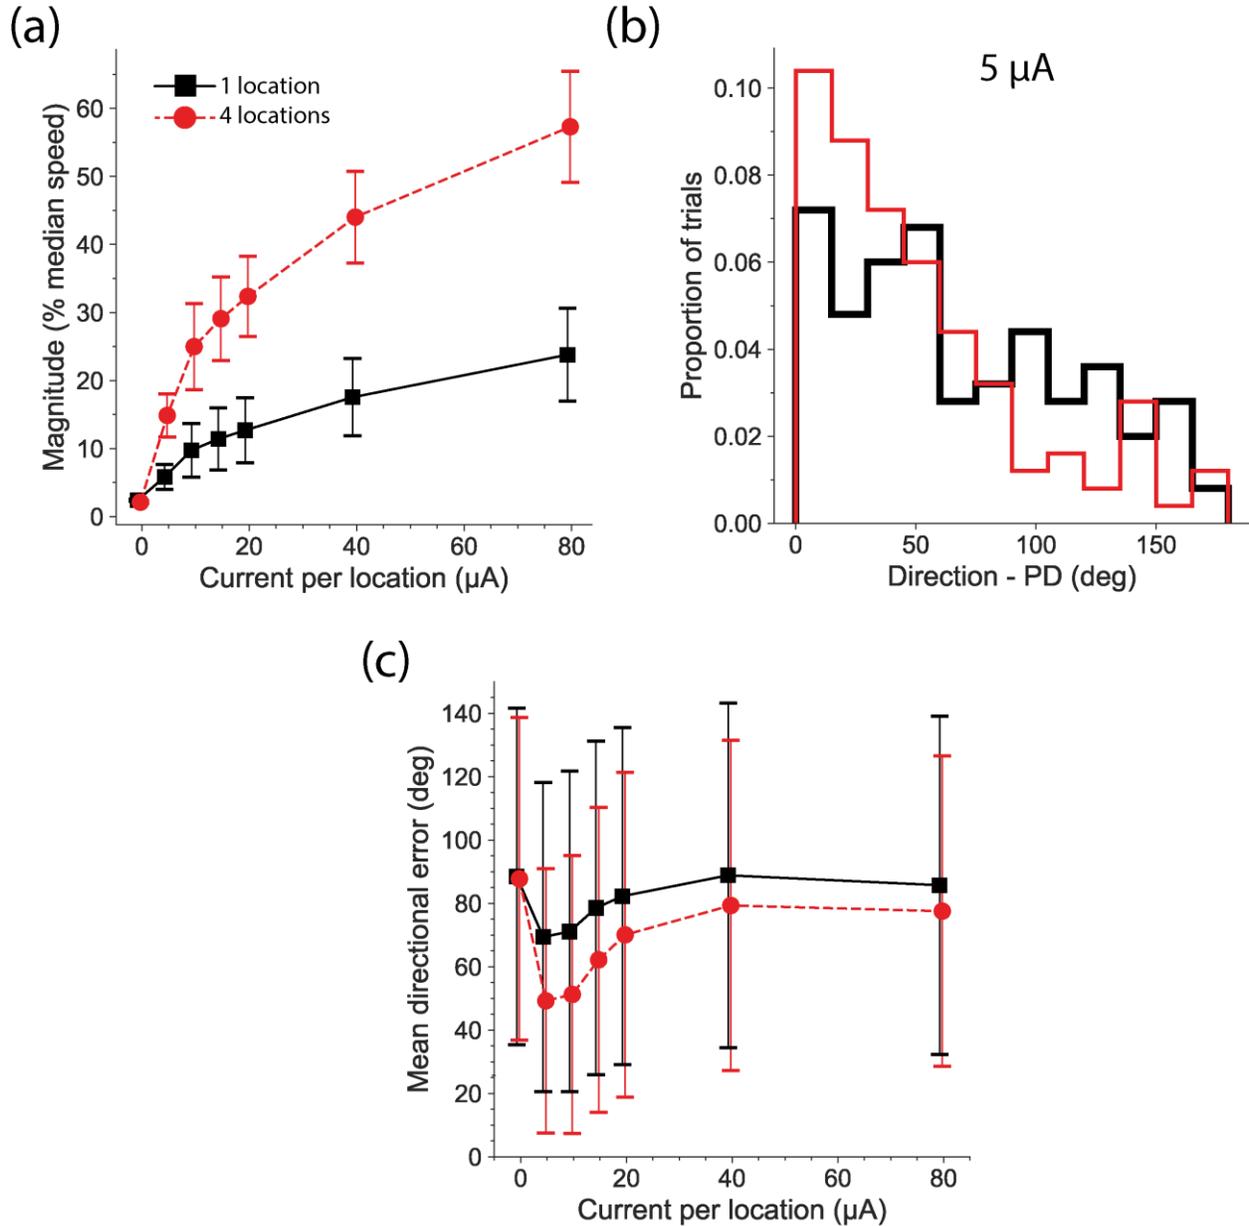


Fig. 4.5. Effect of stimulation across locations. (a) The mean magnitude of the effect of stimulation across 250 stimulus locations for different amplitudes and numbers of sites. Magnitude is normalized by the median speed of the monkey's hand during natural movements. All error bars represent standard deviation. (b) The difference between the direction of the effect at each location and the mean PD of the stimulus site(s) when 1 site (black) or 4 sites (red) are stimulated with 5 μA . (c) The mean difference between the inferred direction and the mean PD of the stimulated location(s) across amplitudes and numbers of simultaneous stimulus locations.

Stimulating at multiple locations also resulted in lower directional errors compared to the same total current at a single location (Figs. 4.5b,c; two-way ANOVA, $F(1, 2997) = 114, p \cong 0$). However, increasing amplitude per location from 5 to 40 μA significantly increased directional error as with single electrodes ($F(1, 2997) = 63, p \cong 0$). Overall, the mean PD of the stimulated locations was closest to the direction of the effect of stimulation when stimulating multiple locations simultaneously each with small currents.

Efficacy of stimulation at cluster centers

As we saw with the two example locations (Fig. 4.4), the PD distribution of the neurons surrounding the stimulus site impacts the magnitude and directional error dramatically. We wondered how the efficacy of stimulation changes as the site of stimulation is moved away from cluster centers. To test this, we stimulated 100 locations with amplitudes spanning 5-20 μA , beginning at cluster centers. We then progressively moved the stimulus site out from the center, causing the magnitude of the effect to decrease (Fig. 4.6a). Using a two-way ANOVA, we determined that magnitude decreased significantly with increased distance from the cluster center ($F(1, 1997) = 12, p = 6.3 * 10^{-4}$).

As expected, increasing distance from the center also significantly increased directional error (Fig. 4.6b; two-way ANOVA, $F(1, 1997) = 97, p \cong 0$). There was little effect 100 μm from the center, likely because the surrounding neurons still had a homogeneous PD distribution. Increasing current also increased directional error, an effect that was independent of the positioning within the cluster. This suggests that electrodes should be placed within $\sim 100 \mu\text{m}$ of the center of clusters for ICMS to evoke effects with the lowest directional error.

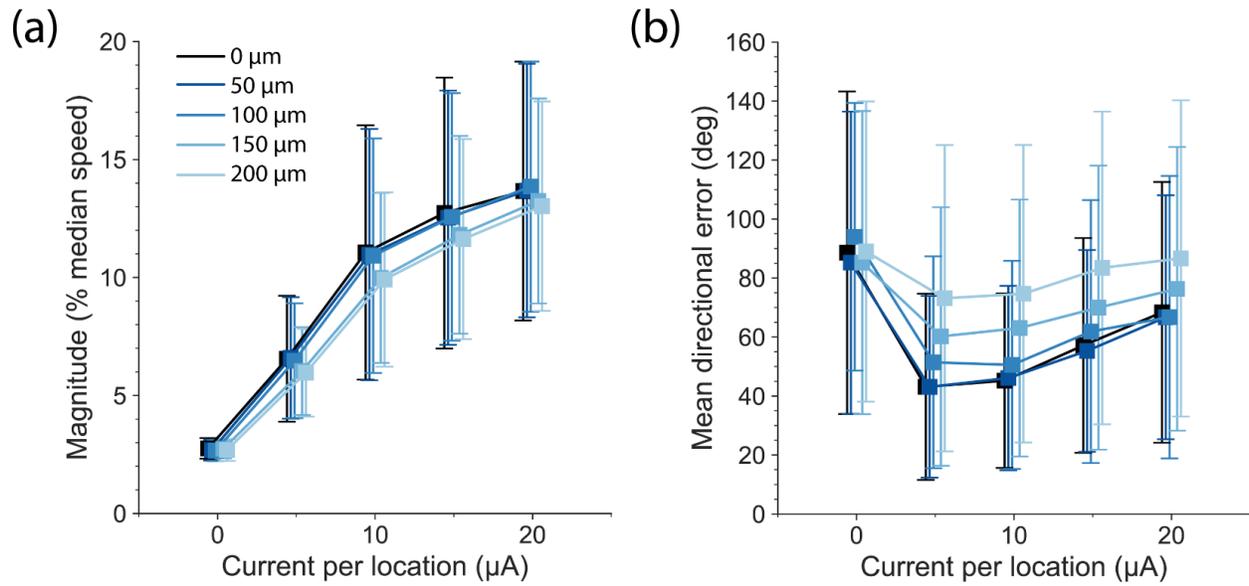


Fig. 4.6. Efficacy of stimulation at progressively greater distances from cluster centers. (a) Magnitude of the effect of stimulation for different distances from the center of clusters with homogeneous PDs. (b) Directional error with increasing distance from the cluster centers.

To determine how effective stimulation could be in ideal conditions, we stimulated 20 sets of up to 24 clusters with similar PDs simultaneously. As with fewer sites, magnitude increased with amplitude (Fig. 4.7a; two-way ANOVA, $F(1,316) = 175$, $p \cong 0$) and number of sites ($F(1,316) = 154$, $p \cong 0$). The amplitude effect increased sublinearly with the number of sites ($F(1,316) = 19$, $p = 2 * 10^{-5}$), likely due to sites activating overlapping populations of neurons.

Stimulating 4 cluster centers resulted in directional errors of $\sim 20^\circ$ at $5 \mu\text{A}$, decreasing to $\sim 5^\circ$ when 24 sites were stimulated (Fig. 4.7b). Directional error increased as current per location increased, though was still less than 40° when stimulating with $15 \mu\text{A}$. Increasing the number of simultaneous stimulus sites decreased directional error when stimulating at $5 \mu\text{A}$ and $10 \mu\text{A}$ (two-way ANOVA, $F(1,315) = 21$ and 12.5 , $p = 6 * 10^{-6}$ and $p = 0.0005$ respectively), but not at larger amplitudes ($p > 0.05$). Overall, stimulating many cluster centers simultaneously each with small currents activates neurons with similar PDs, resulting in large magnitude effects in a direction near the mean PD of the neurons at the stimulus sites.

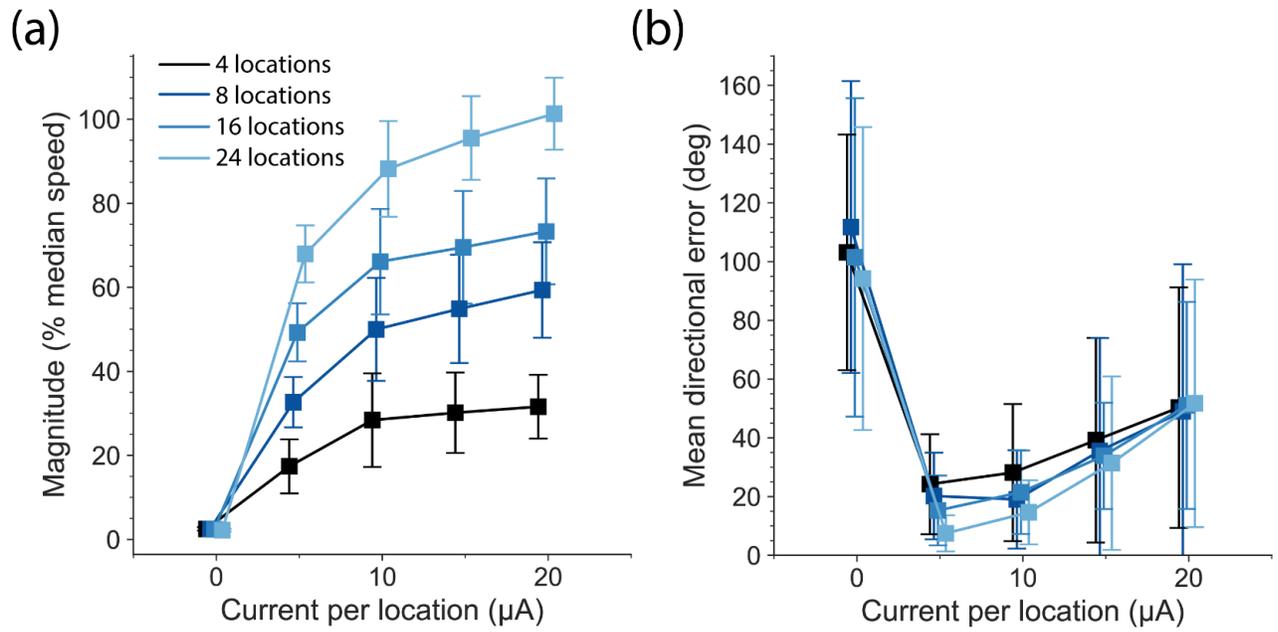


Fig. 4.7. Efficacy of many-location stimulation applied at cluster centers. (a) Magnitude of the effect of stimulation when stimulation was applied at many cluster centers across amplitudes. (b) Directional error with increasing number of simultaneous stimulus sites and amplitude.

Effect of cluster size on the magnitude and direction of the effect of stimulation

We also wondered what effect the size of clusters would have on the efficacy of stimulation and the precision with which electrodes would need to be positioned. We expected stimulation within larger clusters to activate more neurons with homogeneous PDs, resulting in larger magnitude effects and lower directional error. To test this, we altered the dimensions of our cortical map to effectively increase or decrease the size of clusters (Fig. 4.8a). We then tested the effect of stimulation at 100 locations. The mean adjusted magnitude increased with increased cluster size, as stimulation activated a more homogeneous population of neurons (Fig. 4.8b). Using a two-way ANOVA, we determined that increasing the cluster size increases the magnitude of the effect of stimulation across amplitudes ($F(1, 2397) = 23, p \cong 0$). Likewise, directional error decreased substantially with increased map size for currents $< 40 \mu\text{A}$, again due to increased activation of more homogeneous neurons. (Fig. 4.8c; $F(1, 2397) = 94, p \cong 0$).

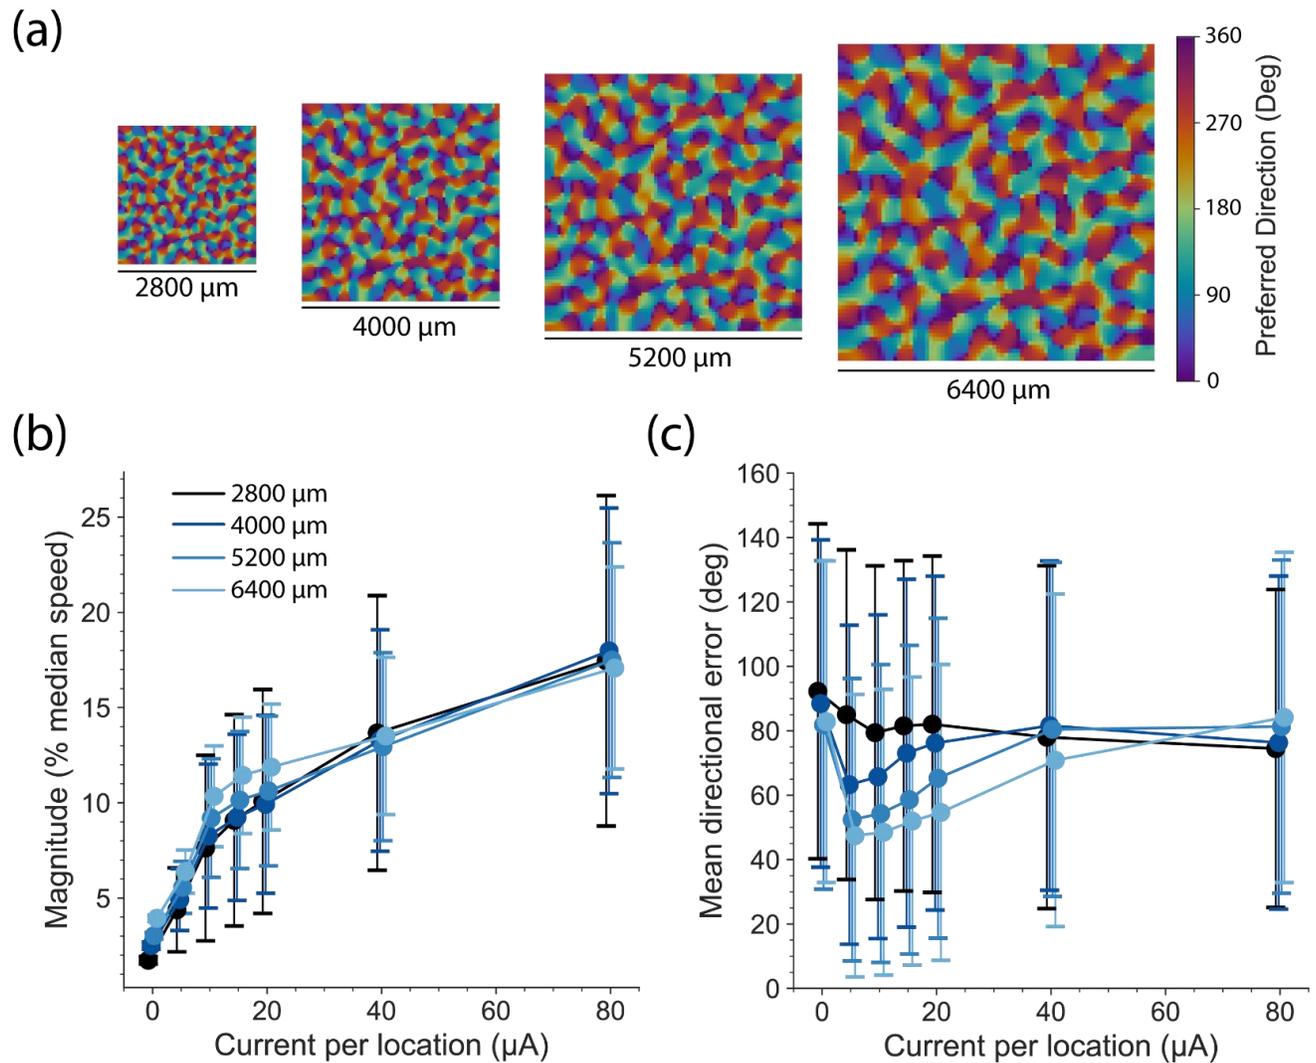


Fig. 4.8. Effect of cluster size. (a) Map representing the PDs of cortical neurons. To avoid the need to recompute an entire cortical model, with PD clusters that would have differed as a result of the stochastic learning process, we simply scaled the map up and down in size from 4000 μm to alter the effective size of clusters. (b) Magnitude of the effect of stimulation for different cluster sizes across a range of amplitudes. (c) Directional error across cluster sizes.

Discussion

We simulated the effect of stimulating within proprioceptive cortex using an autoencoder network (the “topoVAE”) trained to regenerate arm kinematics (Blum, Grogan, et al. 2021). topoVAE maps form small domains of neurons with similar preferred directions, which are themselves organized into pinwheel-like patterns similar to those found in the visual cortex (Malonek, Tootell, and Grinvald 1994; Kang, Shelley, and Sompolinsky 2003), but without any discernible long-range organization. We artificially stimulated the map using a wide range of conditions like those that might be used to implement an afferent BCI. We decoded hand velocity from the simulated cortical output as a proxy of the sense of limb movement that a user might experience. Increasing current amplitude increased the magnitude of the effect of stimulation, but also increased the directional error of the decoded velocity estimate with respect to the PD of the neurons at the immediate stimulus site. Increasing the number of simultaneous stimulus locations decreased this error. The directional error could be decreased substantially by stimulating near the center of clusters of neurons with homogeneous PDs. Here, we discuss the implications of our results for designing proprioceptive interfaces, comparisons between our simulation study and psychophysical experiments, and limitations of our study.

Developing proprioceptive interfaces

Proprioception is vital for motor control: Individuals who have lost proprioception but retain the ability to contract their muscles are typically wheelchair bound and make slow and imprecise reaches (Ghez et al. 1990; Sainburg et al. 1995). With state-of-the-art BCIs, spinal-cord injured patients can make reaches with a robotic limb or even their own arm (Ajiboye et al. 2017; Collinger et al. 2013; Ethier et al. 2012; Bouton et al. 2016). However, because current BMIs do

not restore proprioception, the reaches made by these users will, in most cases, also suffer the limitations faced by deafferented individuals. It will be important to restore proprioceptive feedback to these patients if the full benefit of the motor interfaces is to be realized.

Early tactile cortical areas (Brodmann areas 3b and 1) have a well-defined somatotopy (Mountcastle 1957; Callier, Suresh, and Bensmaia 2019; Chen et al. 2001). ICMS applied to these areas typically evokes sensations at the same location as the receptive fields recorded on the stimulated electrodes (Tobot et al. 2013; Flesher et al. 2016). The sensations evoked by ICMS may be similar to flutter sensations, as monkeys could accurately compare the frequency of ICMS to the frequency of a mechanical vibrator (Romo et al. 1998). Monkeys can also accurately discriminate frequency of two ICMS trains (Romo et al. 2000; Kim, Callier, Tobot, Gaunt, et al. 2015). In an initial bidirectional BCI, monkeys selected virtual objects with two different artificial “textures,” conveyed by different stimulation frequency, by controlling a virtual arm with signals recorded from motor cortex (O’Doherty et al. 2011a). Further experiments have shown that modulating frequency changes the sensation quality and intensity in a highly electrode-dependent manner (Callier et al. 2020), implying that monkeys could discriminate frequency by simply discriminating intensity. These early experiments have culminated in current afferent interfaces which can elicit tactile sensations in humans (Salas et al. 2018; Osborn et al. 2021). When using current bidirectional BCIs, which provide control of a robotic limb and force and contact location feedback, a spinal cord injured patient grasped and move objects faster than with visual feedback alone (Flesher et al. 2021).

Stimulation in proprioceptive areas has not been as effective, possibly because the somatotopy in proprioceptive areas is less well-defined. The poorly defined somatotopy in proprioceptive areas is partly because signals arise largely from receptors embedded in multiple

muscles in the arm, some spanning multiple joints, then need to be mapped onto the two-dimensional cortical surface. Because of this, arm movements, which contract and stretch multiple muscles, evoke a complex pattern of activity across proprioceptive cortex. Nonetheless, proprioceptive-like sensations have been reported by a human participant at currents slightly above those required to elicit tactile sensation from the same electrode array. (Salas et al. 2018). Oddly, these were in response to what was likely stimulation in area 1. The evoked sensations were relatively weak and loosely correlated with stimulus amplitude, but larger amplitudes tended to shift the perceived modality from cutaneous to proprioceptive.

In a series of experiments, monkeys were trained to report the direction of a mechanical perturbation applied to their hand (Tomlinson and Miller 2016). In some cases, the mechanical perturbation was paired with stimulation through four electrodes with similar PDs in an attempt to bias the monkey's report of the perturbation direction. In one monkey, the bias direction could be predicted based on the PD of the stimulated electrodes. However, in three other monkeys it could not be.

In the face of these challenges to take a fully biomimetic approach to a proprioceptive BCI (Bensmaia and Miller 2014), it has nonetheless proven possible to monkeys to learn an arbitrary mapping between ICMS and the error vector between their hand position and the position of invisible targets (Dadarlat, O'Doherty, and Sabes 2015). After a few months of training, monkeys were able to use the ICMS feedback to reach the target, though their movements were slow. Given this success, it seems likely that a mapping which uses stimulation patterns that evoke more naturalistic sensations of proprioception would reduce the training time and improve the quality of feedback.

Our simulation experiments show that low-amplitude stimulation applied near multiple PD cluster centers may be necessary to evoke robust, meaningful sensations of limb movement (Fig. 4.7). Stimulation efficacy depended on the PD distribution of neurons near the stimulated location (Figs. 4.6 and 4.7) and on the size of clusters of neurons with homogeneous PDs (Fig. 4.8). It may be that by chance, the electrodes in the successful monkey described above were placed in optimal locations. Alternatively, that monkey may have had larger PD clusters leading to more homogeneous activation.

The need to stimulate with very low currents on multiple electrodes to achieve adequate accuracy, raises the concern that it may be impossible to achieve adequate perceived intensity. Previous experiments have reached different conclusions about the linearity of perceived effects across more than one electrode. Single-electrode detection thresholds typically range from ~10-40 μA (Kim, Callier, Tabot, Tenore, et al. 2015; Sombeck and Miller 2019) and can even be as low as 5 μA (Zaaimi et al. 2013). In one set of experiments, multi-electrode stimulation was detected more reliably than was predicted by a model which assumed independence across electrodes. Monkeys were able to detect multi-electrode stimulation reliably even when stimulation applied through each electrode alone was subthreshold (Zaaimi et al. 2013). However, instead of this supralinear effect, other experiments found a close match to a model assuming independence (Kim, Callier, Tabot, Tenore, et al. 2015). A third study suggested even sublinear summation, as monkeys could readily detect 80 μA of stimulation through a single electrode, but not if the same total current distributed over 16 electrodes (Sombeck and Miller 2019). Since changing amplitude affects both the magnitude and directional error, future proprioceptive interfaces will need to carefully tune the amplitude of stimulation.

Parallels to psychophysical work in area MT

The study of Tomlinson and Miller described above (Tomlinson and Miller 2016), was based on much earlier experiments that sought to bias monkeys' perception of the direction of moving dots by stimulating in area MT (Murasugi, Salzman, and Newsome 1993; Salzman et al. 1992). Clear parallels exist between our simulation study and both of these psychophysical studies. Area MT is involved in processing the direction of motion of a visual stimulus and lacks sensitivity to color or form (Albright, Desimone, and Gross 1984; Maunsell and Van Essen 1983; Van Essen, Maunsell, and Bixby 1981). Neurons with similar direction preference are organized into columns in area MT, which are themselves organized in pinwheel-like structures (Albright, Desimone, and Gross 1984; Burkhalter, Van Essen, and Maunsell 1981; Malonek, Tootell, and Grinvald 1994), much like the organization of neurons in our generated cortical map (Fig. 4.3a). Stimulating with small currents ($\leq 40 \mu\text{A}$) near the center of a column in area MT induced a perceptual bias with a direction that could be predicted by the PDs of neurons in the stimulated column. As the electrode was progressively moved away from the center, the effect of stimulation nearly disappeared (Murasugi, Salzman, and Newsome 1993). Furthermore, at $80 \mu\text{A}$, monkeys no longer could discern the direction of moving dots, presumably because of the activation of many neurons with a wide distribution of encoding properties. We observed a similar effect, whereby directional error increased with increasing current, and as we moved the stimulus site away from the center of a cluster of neurons (Fig. 4.6).

Further evidence that stimulation amplitude affects the nature of the evoked sensation comes from a human's reports of the phosphenes evoked by stimulating V1. Stimulation applied at amplitudes near the detection threshold evoked phosphenes of various colors, but the color of the phosphenes changed to white, greyish, or yellowish as amplitude increased and additional

neurons were recruited (Schmidt et al. 1996). Directional error in our study reached chance level at $\sim 40 \mu\text{A}$, possibly because clusters in our simulated map are smaller than those in area MT. Indeed, while clusters in our generated map spanned $\sim 250 \mu\text{m}$, those in area MT of owl monkeys are $\sim 400 \mu\text{m}$ wide (Malonek, Tootell, and Grinvald 1994). Future proprioceptive interfaces will need to be designed with carefully choose stimulus locations to evoke effects in directions predicted by the neurons recorded on the stimulated electrodes.

Assumptions and limitations

In this study, we simulated the effects of ICMS using an artificially generated cortical map. We wondered how the topography of the map would interact with stimulus parameters in determining the efficacy of stimulation. While we used joint velocities as input to our cortical maps, neurons in actual S1 receive information from muscle spindles and Golgi tendon organs in addition to joint receptors (Proske and Gandevia 2012; Houk, Rymer, and Crago 1981; Houk and Simon 1967). To determine whether the choice of inputs affected our results, we also generated maps with muscle velocities. The two types of maps were similar, in part because joint velocities completely determine muscle velocities, with an approximately linear relation assuming a constant joint lever arm (Chowdhury, Glaser, and Miller 2020). In area 2, where our previous psychophysical experiments were performed (Tomlinson and Miller 2016; Sombeck and Miller 2019), neurons receive cutaneous inputs as well as those from muscle receptors. As we did not include cutaneous information in our current model, it may be more applicable to area 3a, which receives primarily muscle inputs. Future stimulation experiments should target area 3a, though its position at the bottom of the sulcus in the macaque brain makes access with multi-electrode arrays difficult.

To determine neuronal activation due to stimulation, we used a biophysical model of cortex (Kumaravelu et al. 2022). Although the topoVAE itself included realistic lateral connections during training, these played no direct role during stimulation, with neurons activated only directly. In real cortex, responses evoked within ~350 μm are primarily direct, as they are largely unaffected when synapses were blocked (Histed, Bonin, and Reid 2009). Furthermore, 80% of spikes evoked near the stimulus site of the biophysical model were directly evoked. However, transsynaptically evoked activity can propagate millimeters from the site of stimulation potentially leading to overestimates of the specificity of our modeled effects (Hao, Riehle, and Brochier 2016; Sombeck et al. 2021; Voigt, Yusuf, and Kral 2018; Logothetis et al. 2010). In one example, 67% of spikes recorded from the descending pyramidal tract evoked by stimulation in motor cortex had latencies consistent with transsynaptic activation (Jankowska, Padel, and Tanaka 1975). Given the precision of the evoked tactile perceptions in both monkey and human experiments, this seems not to be a major effect (Flesher et al. 2016; Tabot et al. 2013). Future modeling efforts might include synapses within a single cortical layer or across multiple cortical layers to infer the effect of stimulation more accurately.

Chapter 5 - Discussion

The preceding chapters described my work towards developing stimulation patterns to provide proprioceptive feedback via intracortical microstimulation (ICMS) in somatosensory cortex (S1). One reason providing proprioceptive feedback has proven more difficult than providing tactile feedback is the less well-defined topography in proprioceptive cortical areas. This topography arises from the need to map receptors from many muscles onto the two-dimensional cortical surface. Because of this, even simple limb movements evoke a complex spatial pattern of neural activity across proprioceptive cortex. To mimic this pattern, stimulation likely needs to be provided through many electrodes, each with small amplitudes. In support of this, I measured the spatiotemporal neural response to ICMS in Chapter 2, finding that decreasing stimulation amplitude activated a more spatially restricted population of neurons, neurons that may have similar encoding properties in proprioceptive cortex. To test whether multi-electrode stimulation could provide rapid feedback, I compared the reaction times evoked by stimulation to that evoked by a mechanical perturbation applied to the hand in Chapter 3. Multi-electrode stimulation evoked reaction times as short or shorter than the mechanical cue, implying that this approach can provide feedback rapid enough to replace natural proprioception. Finally, I simulated the sensation evoked by multi-electrode ICMS using a model of proprioceptive cortex in Chapter 4. Stimulation at many sites with small amplitudes produced inferred effects with smaller directional error than did stimulation at single sites.

These results add to previous evidence for the need for multi-electrode stimulation in future afferent interfaces. Single electrode stimulation in early tactile areas evokes sensations typically corresponding to a single finger or a small patch of the hand (Flesher et al. 2016; Tabot et al. 2013;

Fifer et al. 2020), and sensations of limb movement reported by a human participant corresponded only to a small portion of the arm (Salas et al. 2018). To provide feedback about the whole arm and/or hand, these small projected fields will need to be combined by stimulating through many electrodes simultaneously. Additionally, multi-electrode ICMS may also increase the number of distinguishable levels of intensity that could be provided to a user. The just noticeable difference (JND) is the amount stimulus intensity needs to be changed for the difference to be detected. For most natural stimuli, the JND increases with increased intensity. The JND for ICMS was constant across a wide range of amplitudes, measured as $\sim 15\text{-}30\ \mu\text{A}$ across monkeys and humans (Flesher et al. 2016; Kim, Callier, Tabot, Gaunt, et al. 2015). With a current range of $\sim 20\text{-}100\ \mu\text{A}$, corresponding to detection thresholds and the safety maximum, respectively (Rajan et al. 2015; Chen et al. 2014), afferent interfaces would be able to deliver 3-6 distinguishable levels of feedback. Multi-electrode stimulation may increase that number, as this would increase the maximum total current that could be injected from $100\ \mu\text{A}$ through a single electrode to something like the $480\ \mu\text{A}$ used in Chapter 3.

Taken together, these results point to the need for multi-electrode stimulation in future afferent interfaces. The rest of this chapter will discuss the development of biomimetic, multi-electrode stimulation patterns to provide proprioceptive feedback, differences in ICMS-evoked activity compared to activity evoked by natural stimuli, and methods to simultaneously decode neural activity in motor cortex while stimulating S1.

Multi-electrode stimulation to provide proprioceptive feedback

Two previous interfaces used vastly different approaches to provide proprioceptive feedback. In one approach, monkeys were required to learn an arbitrary mapping between stimulation and the error vector between the position of an invisible target and hand position

(Dadarlat, O'Doherty, and Sabes 2015). Stimulation was applied through 8 electrodes, and each electrode was arbitrarily assigned a direction. The frequency of stimulation for each electrode was modulated by projecting the error vector onto each electrode's assigned direction. The maximum frequency scaled with the magnitude of this error vector. After months of training, monkeys learned to use this interface to reach to invisible targets, though it took them twice as long to reach the targets than when provided full visual feedback. Additionally, monkeys would regularly stop during movements to interpret ICMS feedback, resulting in roughly double the number of submovements for each reach than when provided full visual feedback.

It may be stimulation patterns that evoke naturalistic sensations of proprioception would be easier to learn and provide better feedback than arbitrary stimulation patterns. We previously tried to measure the naturalistic effects of ICMS in area 2 of S1 (Tomlinson and Miller 2016). To do so, we trained monkeys to report the direction of a mechanical perturbation by reaching to the one of two opposing targets that was nearer the perturbation direction. During the mechanical perturbation, we applied stimulation at 20 μ A through each of 4 electrodes with similar preferred directions (PDs) in area 2 of S1. In one monkey, stimulation through 6 out of 7 sets of electrodes biased the monkey's perception of the perturbation in the direction of the mean PDs recorded on the stimulated electrodes. In three other monkeys, this stimulation biased the perception of the perturbation but in directions we were unable to predict. It remains unclear why this approach predictably biased one monkey's perception but not the other three. Based on the results in Chapter 4, it may be that the one monkey where stimulation caused predictable biases had larger clusters of neurons with similar PDs than the other monkeys. Alternatively, electrodes may have been placed nearer to cluster centers, as both manipulations substantially reduced directional error between the inferred effect direction and the mean PD at the stimulated sites.

The method of naturalistic stimulation to provide proprioceptive feedback proposed in Chapter 4 builds on this previous attempt (Tomlinson and Miller 2016). Through our recording (Chapter 2) and modeling efforts (Kumaravelu et al. 2022), we observed that stimulation at 20 μA can activate neurons hundreds of microns away from the site of stimulation. This activated population will likely have a broad distribution of encoding properties due to the complex spatial organization of neurons in proprioceptive areas (Pons et al. 1985). In my simulations in Chapter 4, 20 μA stimulation at four sites, each with similar PDs, resulted in a large mean directional error between the inferred sensation and the mean PD of neurons recorded at the stimulus location. This implies that this approach would not reliably evoke sensations of hand movement in the direction of the mean PD of the stimulated electrodes, consistent with the random biases we observed in three out of four monkeys.

To activate a more spatially restricted population of neurons, amplitude should be reduced to 5-10 μA . Since reducing amplitude also reduces the magnitude of the evoked sensation, stimulation will need to be applied through many electrodes simultaneously to increase the total current injected. In Chapter 3, reaction times in response to 10 μA stimulation through 24 electrodes simultaneously (240 μA total) were longer than that to a mechanical perturbation applied to the hand, and only slightly shorter to that of a visual cue for one monkey. Instead of increasing amplitude, frequency could be increased to increase sensation intensity, as this manipulation increases the detectability of stimulus trains (Kim, Callier, Tabot, Gaunt, et al. 2015; Fridman et al. 2010) and shortens reaction times (Chapter 3). In addition to changing intensity, modulating stimulation frequency may change the quality of sensation. During low-frequency stimulation in early tactile cortical areas, human participants reported mostly sensations of pressure, tapping, and “sparkling,” while they reported more warm, buzzing, and tingling

sensations during high-frequency stimulation (Hughes et al. 2021). The effects of changing frequency on sensation quality in proprioceptive cortex warrants further investigation before frequency can be used to change sensation intensity.

Ideally, stimulation electrodes would be placed in the center of clusters of neurons with similar PDs, so that the stimulation would primarily activate neurons with similar encoding properties. Instead of isolating single units on each electrode, multi-unit recordings could be used to better estimate the encoding properties of the population of neurons around each electrode. This approach may be insufficient, though, because the distance at which neural activity is recorded is relatively small compared to the extent of stimulation. To find cluster centers, previous studies moved the stimulation electrode until neurons recorded along 200 μm of electrode travel had similar PDs (Murasugi, Salzman, and Newsome 1993). The stimulation electrode was then placed in the center of this region. Since we cannot move the electrode after implanting in monkeys or humans, we may need to implant electrode arrays with small interelectrode distances to measure the encoding properties of neighboring neurons. This approach may enable better selection of stimulation electrodes. Stimulation through many carefully chosen electrodes each with small amplitudes may closely mimic the complex spatial pattern of activity evoked by naturalistic limb movements, and thus may evoke naturalistic sensations.

Since multi-electrode stimulation will likely deliver a large amount of charge in future afferent interfaces, it is reasonable to question whether this approach is safe. In a study about the safety of ICMS, electrode arrays were implanted into S1. Stimulation was delivered through a maximum of 12 electrodes simultaneously and up to amplitudes of 100 μA for five days a week over six months (Rajan et al. 2015). The arrays were then explanted and histology on the cortical tissue was performed. Most of the damage to tissue was due to implanting and explanting the

arrays, not stimulation (Rajan et al. 2015). From this study, the damage to tissue surrounding stimulation electrodes seems minimal, even during multi-electrode stimulation at large amplitudes.

Framework for designing biomimetic stimulation patterns

Developing stimulation patterns which mimic the spatiotemporal neural response to natural stimuli is difficult in part because of the complex relationship between ICMS and evoked activity. While it is difficult to record evoked activity in vivo, this relationship from ICMS to spikes can be modeled using biophysical models of cortical stimulation (Aberra, Peterchev, and Grill 2018). In order to generate stimulation patterns that evoke a specific pattern of neural activity, we need the opposite mapping, one from desired spike patterns to the necessary stimulus pattern. A genetic algorithm is well suited for searching the stimulus parameter space to find these patterns because the relationship between ICMS and evoked activity is highly nonlinear (Davis 1991). Stimulation patterns produced by a genetic algorithm evoked activity in modeled neurons that was more similar to naturally occurring activity than that evoked by a linear mapping between neural activity and stimulation amplitude (Kumaravelu et al. 2020). This resulted in an encoder that converted desired neural activity into stimulus patterns which best produced that activity.

The accuracy of stimulation patterns generated by this encoder in part depends on how well the biophysical model of cortical stimulation models the relationship between ICMS parameters and neural firing. This model used highly realistic neurons that were adapted from 3D digital reconstructions of neurons in rats (Markram et al. 2015). Stimulation in this model reproduced three observations from in vivo ICMS experiments: the strength-duration relationship, the current-distance relationship, and the preferential activation of axons (Aberra, Peterchev, and Grill 2018). To provide further evidence that the modeled activity was similar to that of actual cortical stimulation, we attempted to compare our physiological recordings of the evoked activity to the

modeled evoked activity. Through this process, we realized that the biophysical model generated directly evoked spikes, those evoked by the stimulus pulse itself, while the spikes we recorded in actual cortex had latencies consistent with transsynaptically evoked activity. Because of this, we could not meaningfully compare our recordings. Since we cannot record directly evoked spikes physiologically, it may be that the only way to test the sensations evoked by this encoder is to use human reports or infer the sensation in monkeys.

Rapid assessment of the evoked direction of limb movement

It would be worthwhile to compare the perceptual effects across many different stimulation patterns to find those that evoke more naturalistic sensations of proprioception. To thoroughly explore the stimulation parameter space, it is important to rapidly assess the perceptual effects of each pattern. The experiments we performed to measure biases caused by ICMS, though, took multiple days for only a few stimulation conditions (Tomlinson and Miller 2016). With this approach, it would be infeasible to test the perceptual effects for many stimulation patterns. Additionally, this task only could infer the direction of the bias in one of two opposing directions, providing a very limited view on the evoked sensation.

To more rapidly assess the direction of the sensation evoked by ICMS, I developed a “ring-reporting” task. In this task, monkeys reported the direction of a perturbation by reaching to an outer ring. Monkeys were initially trained to reach in the direction of a mechanical perturbation. After they became proficient at this, ICMS in area 2 replaced the mechanical perturbation on a small proportion of trials. If ICMS evoked a sensation of limb movement similar to the mechanical perturbation, then the monkey would reach to the outer ring in the direction of the ICMS-evoked sensation.

With the ring-reporting task, it was difficult to determine how to reward the monkey on ICMS trials, as the correct reach direction was not known before stimulating. I first rewarded the monkeys regardless of the direction they reached, either by rewarding on every ICMS trial or on a random proportion of ICMS trials. Because the direction of their reach did not matter, the monkeys quickly learned to reach in the same direction on every ICMS trial to receive a reward. I also tried rewarding on none of the ICMS trials, which caused the monkeys to hold still in the middle of the workspace waiting for a mechanical perturbation.

Instead of rewarding independently of reach direction, I lastly used my predicted direction for each stimulus pattern as the rewarded direction. While this approach does allow monkeys to learn the mapping between stimulus and reach direction, it would be difficult to do so because of the numerous stimulation patterns tested in a single session. With this reward strategy, monkeys still reached in the same direction on every ICMS trial. This may be because the monkeys had already learned to reach in a single direction on ICMS trials from previous reward strategies. Alternatively, if my predictions were inaccurate, the rewarded direction would not match the direction of the ICMS-evoked. In this scenario, the monkey would be rewarded on random trials where my predictions happened to align with direction of the evoked sensation. When I tested randomly rewarding the monkey, he reached in the same direction on every ICMS trial. The only conclusion from these experiments was that monkeys can discriminate ICMS in area 2 from the mechanical perturbation, implying that these two sensations are different.

Targeting other proprioceptive areas

Part of the reason ICMS in area 2 evokes unnatural or unpredictable sensations may be because it is a mixed-modality area. Area 2 receives both cutaneous and proprioceptive inputs, and neurons in this area respond to cutaneous stimuli as well as joint and muscle manipulations

(Padberg et al. 2019; Pons et al. 1985; Iwamura et al. 1993). Because of this, stimulation in area 2 will likely activate both cutaneous and proprioceptive neurons, resulting in mixed-modality sensations. Additionally, the presence of both cutaneous and proprioceptive information makes the hand portion of area 2 ideally suited for stereognosis, the ability to recognize the form objects without visual input (Yau et al. 2016). We have further speculated that the arm portion of area 2 plays a role in mapping the location of encountered objects with respect to the body (Blum, Versteeg, et al. 2021). This object-based representation may make area 2 a poor target for proprioceptive interfaces that aim to evoke simpler sensations of arm movement.

Instead of area 2, proprioceptive interfaces could target area 3a of S1, as it primarily receives inputs from joint and muscle receptors, with only a small proportion of cutaneous inputs (Iwamura et al. 1983; Krubitzer et al. 2004). Unfortunately, this area is located at the bottom of the central sulcus, making it difficult to implant chronic arrays. When acute electrodes were inserted into area 3a, monkeys readily detected and discriminated trains of different frequencies (London et al. 2008). If chronic electrode arrays could be implanted into this area, something that we are pursuing, then area 3a is an appealing target for eliciting proprioceptive sensations.

Combining biomimetic and learning approaches

Ultimately, it will be difficult to mimic the complex spatial pattern of activity evoked by limb movements with the spatially imprecise activation caused by ICMS (Histed, Bonin, and Reid 2009). Because of this, the sensation evoked by biomimetic approaches will most likely not be fully natural. Users, then, would need to learn a mapping between stimulation and feedback. It can take a few months of training to learn an arbitrary mapping between stimulation and even simple, two-dimensional feedback (Dadarlat, O'Doherty, and Sabes 2015). Even after months of training, monkeys made slower reaches with ICMS feedback than with visual feedback. It may take much

longer to learn to use interfaces which provide feedback about the seven-dimensional kinematics of the whole arm. Instead, stimulation patterns which evoke more naturalistic sensations may be easier to learn, and thus reduce the training time required to use these interfaces. Future experiments could determine whether monkeys learn to use biomimetic stimulation patterns faster than nonbiomimetic patterns. If they do, then this would imply that the biomimetic patterns elicit more naturalistic sensations than nonbiomimetic patterns, perhaps opening the door for a combination approach in humans.

Statistics of ICMS-evoked activity are inherently unnatural

The preceding sections discuss using ICMS to mimic the spatiotemporal pattern of naturally evoked cortical activity. Even with complex patterns of ICMS, the resulting sensation still may be unnatural, as the ICMS-evoked activity is fundamentally different than activity evoked by natural stimuli. As an example, for most natural stimuli, JND increases with increasing intensity. This is because increasing intensity of natural stimuli results in both more neural activity and more noise, reflected in increases in both the mean firing rate of neurons and the variance of that rate respectively (Ekman 1959; Johnson 1980b). In contrast, the JND for ICMS was constant across a wide range of amplitudes (Kim, Callier, Tabot, Gaunt, et al. 2015; Flesher et al. 2016), implying that ICMS amplitude affects the mean and variance of the firing rate of neurons differently than do natural stimuli. The results in Chapter 3 support this hypothesis: the mean rate recorded during trains of stimulation increased with increased amplitude, but the variance across trains was constant. These two relationships result in a constant JND across amplitudes. This example illustrates that ICMS-evoked activity is fundamentally different than naturally evoked activity.

Electrical stimulation activates neurons at a tightly fixed latency after each stimulus pulse, resulting in unnaturally synchronized activity across the activated population of neurons. This synchrony may be why it is difficult for stimulation in peripheral nerves to evoke natural tactile percepts (Tan et al. 2014). Similarly, when electrical stimulation is applied to the muscles, this synchrony results in jerky movements and high muscle fatigue (Hughes, Guo, and DeWeerth 2010). In the case of muscle stimulation, evoked movements can be made more fluid and fatigue can be decreased by stimulating on multiple functionally similar electrodes asynchronously, thereby introducing asynchrony between the fibers activated by different electrodes (Hughes, Guo, and DeWeerth 2010). When stimulating early tactile cortical areas, such asynchronous stimulation did not change the detection threshold compared to synchronous stimulation, implying that the intensity of the evoked sensation was unaffected by slightly altering the timing of each pulse (Kim, Callier, Tabot, Tenore, et al. 2015).

Synchrony across neurons can be further reduced even during single electrode stimulation by replacing each pulse in the stimulus train with a series of high-frequency pulses (~1 kHz), each with increasing amplitude. Because the amplitude of stimulation increases during this burst, neurons with different activation thresholds will be activated at different times during the burst. Neurons will not fire multiple times during each burst because the pulses within each burst occur at such a high frequency. In computer simulations and in explanted nerves, trains of high-frequency bursts evoked less synchronous activity than did conventional trains (Formento et al. 2020). While it is currently not possible to record neural activity during each burst in a monkey, thus preventing direct confirmation, the perceptual effects of these trains could be inferred in monkeys or reported by humans.

Rebound excitation may be problematic for future afferent interfaces

While future afferent interfaces should develop stimulation patterns that provide rich, informative feedback during grasping, at some point the stimulation train will stop, such as when the hand is not grasping an object. The evoked sensation should also stop when the train stops. However, human participants have reported sensations lasting beyond the end of stimulus trains (Hughes, Flesher, and Gaunt 2021). I observed greatly increased firing rates after the end of high frequency trains in Chapter 2, which may be the cause of these persistent sensations. As this rebound excitation primarily occurred at high stimulation frequencies, there may be a maximum stimulus frequency that interfaces should employ.

Decoding motor intent during multi-electrode ICMS

Almost all applications of afferent interfaces would only be useful if paired with an interface which restores motion. For spinal cord injured patients, motor intent can be decoded from signals recorded in motor cortex (M1) (Ethier et al. 2012; Collinger et al. 2013). However, stimulation applied in S1 to provide somatosensory feedback produces large electrical artifacts in M1 recordings, likely reducing the accuracy of motor decoders. To remove the artifact, the signal could be blanked during the stimulus pulse, after which a low-pass filter could be applied (Weiss et al. 2018). With this approach, neural signals in M1 were recorded ~ 0.7 ms after stimulation offset in S1, resulting in ~ 1.5 ms of lost data for each stimulus pulse. At low stimulation frequencies the impact of losing this data will be small. Artificially dropping a random 20% of M1 signals only caused a 10% decrease in performance when decoding intended cursor velocity (figure 8 in (Young et al. 2018)), an amount that corresponds to the amount of data lost when stimulating at ~ 130 Hz.

Functional afferent interfaces would likely provide more stimulation during limb movements than have the early proof of principle interfaces, resulting in more data loss during

these critical periods. Furthermore, as stimulation protocols become more complicated, with higher frequencies and stimulation over many electrodes, the percentage of corrupted M1 data will increase. Stimulation delivered at 333 Hz would obscure ~50% of signal in M1, resulting in a 30% decrease in decoding performance during a simple task (figure 8 in (Young et al. 2018)). Fortunately, if the amplifier does not saturate, there are many techniques that could extract neural data during stimulation: adaptive filtering (Mendrela et al. 2016; Nag et al. 2015), template subtraction (Montgomery Jr, Gale, and Huang 2005; Hashimoto, Elder, and Vitek 2002), independent component analysis (Hyvärinen and Oja 2000; Lemm et al. 2006), linear regression reference (Young et al. 2018), and deep neural networks (Tamada et al. 2020; Zhang and Yu 2018). Neural activity has been recorded during the stimulus pulse by using principal component analysis to exploit the similar structure of the shock artifact sequentially across electrodes, pulses, and then trials (O'Shea and Shenoy 2017). With these techniques, it may be possible to record full-bandwidth signal in M1 while applying complex spatiotemporal patterns of stimulation in S1.

Final Conclusions

Proprioception is vital for controlling movements: Without it, the movements made by spinal cord injured patients with state-of-the art motor interfaces will be slow and imprecise. While tactile interfaces have provided force and contact location feedback to spinal cord injured patients, proprioceptive interfaces have been more difficult to develop. In this work, I developed stimulation techniques for providing proprioceptive feedback in a brain-computer interface. I measured the spatiotemporal neural response to ICMS on all electrodes on a multi-electrode array, including on the stimulated channel (Chapter 2). High-amplitude single electrode stimulation caused widespread activation that cannot recreate the complex pattern of activity observed in proprioceptive cortical areas during limb movements. Instead, stimulation would need to be

provided through multiple electrodes to mimic the natural spatial pattern. I tested the ability for multi-electrode ICMS to provide rapid feedback (Chapter 3) and modeled the inferred sensation using an artificially generated cortical map (Chapter 4). In both studies, I found that multi-electrode stimulation with small amplitudes through many electrodes was more effective than single electrode stimulation. In this final chapter, I discussed this approach and the development of biomimetic stimulation patterns generally. Given the difficulty in evoking naturalistic sensations of proprioception and measuring the perceptual effects of ICMS in proprioceptive areas, it is likely that future afferent interfaces will need to combine biomimetic patterns with learning approaches to provide informative feedback without requiring months of training.

References

- Aberra, Aman S, Angel V Peterchev, and Warren M Grill. 2018. 'Biophysically realistic neuron models for simulation of cortical stimulation', *Journal of Neural Engineering*, 15: 066023.
- Ajiboye, A. Bolu, Francis R. Willett, Daniel R. Young, William D. Membedger, Brian A. Murphy, Jonathan P. Miller, Benjamin L. Walter, Jennifer A. Sweet, Harry A. Hoyer, Michael W. Keith, P. Hunter Peckham, John D. Simeral, John P. Donoghue, Leigh R. Hochberg, and Robert F. Kirsch. 2017. 'Restoration of reaching and grasping movements through brain-controlled muscle stimulation in a person with tetraplegia: a proof-of-concept demonstration', *The Lancet*, 389: 1821-30.
- Albright, Thomas D, Robert Desimone, and Charles G Gross. 1984. 'Columnar organization of directionally selective cells in visual area MT of the macaque', *Journal of Neurophysiology*, 51: 16-31.
- Allison-Walker, Tim, Maureen A Hagan, Nicholas SC Price, and Yan T Wong. 2021. 'Microstimulation-evoked neural responses in visual cortex are depth dependent', *Brain Stimulation*, 14: 741-50.
- Amari, Shun-ichi. 1977. 'Dynamics of pattern formation in lateral-inhibition type neural fields', *Biological cybernetics*, 27: 77-87.
- Antfolk, Christian, Marco D'Alonzo, Birgitta Rosén, Göran Lundborg, Fredrik Sebelius, and Christian Cipriani. 2013. 'Sensory feedback in upper limb prosthetics', *Expert Review of Medical Devices*, 10: 45-54.
- Bensmaia, Sliman J., and Lee E. Miller. 2014. 'Restoring sensorimotor function through intracortical interfaces: progress and looming challenges', *Nature Reviews Neuroscience*, 15: 313-25.
- Berg, J. A., Iii J. F. Dammann, F. V. Tenore, G. A. Tabot, J. L. Boback, L. R. Manfredi, M. L. Peterson, K. D. Katyal, M. S. Johannes, A. Makhlin, R. Wilcox, R. K. Franklin, R. J. Vogelstein, N. G. Hatsopoulos, and S. J. Bensmaia. 2013. 'Behavioral Demonstration of a Somatosensory Neuroprosthesis', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 21: 500-07.
- Blum, Kyle P, Christopher Versteeg, Joseph Sombeck, Raed H Chowdhury, and Lee E Miller. 2021. 'Proprioception: a sense to facilitate action.' in, *Somatosensory Feedback for Neuroprosthetics* (Elsevier).

- Blum, Kyle P., Max Grogan, Yufei Wu, J. Alex Harston, Lee E. Miller, and A. Aldo Faisal. 2021. 'Predicting proprioceptive cortical anatomy and neural coding with topographic autoencoders', *bioRxiv*: 2021.12.10.472161.
- Bouton, Chad E., Ammar Shaikhouni, Nicholas V. Annetta, Marcia A. Bockbrader, David A. Friedenberg, Dylan M. Nielson, Gaurav Sharma, Per B. Sederberg, Bradley C. Glenn, W. Jerry Mysiw, Austin G. Morgan, Milind Deogaonkar, and Ali R. Rezai. 2016. 'Restoring cortical control of functional movement in a human with quadriplegia', *Nature*, 533: 247-50.
- Brown, Edgar A, James D Ross, Richard A Blum, Yoonkey Nam, Bruce C Wheeler, and Stephen P DeWeerth. 2008. 'Stimulus-artifact elimination in a multi-electrode system', *IEEE Transactions on Biomedical Circuits and Systems*, 2: 10-21.
- Burkhalter, A, DC Van Essen, and JHR Maunsell. 1981. "Patterns of 2-deoxyglucose labeling in extrastriate visual cortex of unstimulated and unidirectionally stimulated macaque monkeys." In *Sot. Neurosci. Abstr*, 172-98.
- Butovas, Sergejus, Sheriar G. Hormuzdi, Hannah Monyer, and Cornelius Schwarz. 2006. 'Effects of Electrically Coupled Inhibitory Networks on Local Neuronal Responses to Intracortical Microstimulation', *Journal of Neurophysiology*, 96: 1227-36.
- Butovas, Sergejus, and Cornelius Schwarz. 2003. 'Spatiotemporal Effects of Microstimulation in Rat Neocortex: A Parametric Study Using Multielectrode Recordings', *Journal of Neurophysiology*, 90: 3024-39.
- . 2007. 'Detection psychophysics of intracortical microstimulation in rat primary somatosensory cortex', *European Journal of Neuroscience*, 25: 2161-69.
- Caldwell, David J, Jeneva A Cronin, Jing Wu, Kurt E Weaver, Andrew L Ko, Rajesh PN Rao, and Jeffrey G Ojemann. 2019. 'Direct stimulation of somatosensory cortex results in slower reaction times compared to peripheral touch in humans', *Scientific Reports*, 9: 3292.
- Callier, Thierry, Nathan W Brantly, Attilio Caravelli, and Sliman J Bensmaia. 2020. 'The frequency of cortical microstimulation shapes artificial touch', *Proceedings of the National Academy of Sciences*, 117: 1191-200.
- Callier, Thierry, Aneesha K Suresh, and Sliman J Bensmaia. 2019. 'Neural Coding of Contact Events in Somatosensory Cortex', *Cerebral Cortex*.
- Chan, Sherwin S, and Daniel W Moran. 2006. 'Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces', *Journal of Neural Engineering*, 3: 327.

- Chandrasekaran, Santosh, Stephan Bickel, Jose L Herrero, Joo-won Kim, Noah Markowitz, Elizabeth Espinal, Nikunj A Bhagat, Richard Ramdeo, Junqian Xu, and Matthew F Glasser. 2021. 'Evoking highly focal percepts in the fingertips through targeted stimulation of sulcal regions of the brain for sensory restoration', *Brain Stimulation*, 14: 1184-96.
- Chen, Kevin H., John F. Dammann, Jessica L. Boback, Francesco V. Tenore, Kevin J. Otto, Robert A. Gaunt, and Sliman J. Bensmaia. 2014. 'The effect of chronic intracortical microstimulation on the electrode–tissue interface', *Journal of Neural Engineering*, 11: 026004.
- Chen, Li Min, Robert M Friedman, Benjamin M Ramsden, Robert H LaMotte, and Anna Wang Roe. 2001. 'Fine-scale organization of SI (area 3b) in the squirrel monkey revealed with intrinsic optical imaging', *Journal of Neurophysiology*, 86: 3011-29.
- Chen, Xing, Feng Wang, Eduardo Fernandez, and Pieter R Roelfsema. 2020. 'Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex', *Science*, 370: 1191-96.
- Chowdhury, Raed H, Joshua I Glaser, and Lee E Miller. 2020. 'Area 2 of primary somatosensory cortex encodes kinematics of the whole arm', *Elife*, 9: e48198.
- Churchland, Mark M., and Krishna V. Shenoy. 2007. 'Delay of Movement Caused by Disruption of Cortical Preparatory Activity', *Journal of Neurophysiology*, 97: 348-59.
- Collinger, Jennifer L., Brian Wodlinger, John E. Downey, Wei Wang, Elizabeth C. Tyler-Kabara, Douglas J. Weber, Angus J. C. McMorland, Meel Velliste, Michael L. Boninger, and Andrew B. Schwartz. 2013. 'High-performance neuroprosthetic control by an individual with tetraplegia', *The Lancet*, 381: 557-64.
- Dadarlat, Maria C., Joseph E. O'Doherty, and Philip N. Sabes. 2015. 'A learning-based approach to artificial sensory feedback leads to optimal integration', *Nature Neuroscience*, 18: 138-44.
- Davis, Lawrence. 1991. 'Handbook of genetic algorithms'.
- Delp, Scott L, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. 2007. 'OpenSim: open-source software to create and analyze dynamic simulations of movement', *IEEE Transactions on Biomedical Engineering*, 54: 1940-50.
- DeMichele, GA, and PR Troyk. 2003. "Stimulus-resistant neural recording amplifier." In *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No. 03CH37439)*, 3329-32. IEEE.

- Devecioğlu, İsmail, and Burak Güçlü. 2017. 'Psychophysical correspondence between vibrotactile intensity and intracortical microstimulation for tactile neuroprostheses in rats', *Journal of Neural Engineering*, 14: 016010-10.
- DeYoe, Edgar A, Jeffrey D Lewine, and Robert W Doty. 2005. 'Laminar variation in threshold for detection of electrical excitation of striate cortex by macaques', *Journal of Neurophysiology*, 94: 3443-50.
- DiCarlo, James J, Kenneth O Johnson, and Steven S Hsiao. 1998. 'Structure of receptive fields in area 3b of primary somatosensory cortex in the alert monkey', *Journal of Neuroscience*, 18: 2626-45.
- Ekman, Gösta. 1959. 'Weber's law and related functions', *The Journal of Psychology*, 47: 343-52.
- Epstein, Charles M. 1995. 'A simple artifact-rejection preamplifier for clinical neurophysiology', *American Journal of EEG Technology*, 35: 64-71.
- Ethier, C., E. R. Oby, M. J. Bauman, and L. E. Miller. 2012. 'Restoration of grasp following paralysis through brain-controlled stimulation of muscles', *Nature*, 485: 368-71.
- Fagg, Andrew H, Gregory W Ojakangas, Lee E Miller, and Nicholas G Hatsopoulos. 2009. 'Kinetic trajectory decoding using motor cortical ensembles', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17: 487-96.
- Fifer, Matthew S, David P McMullen, Tessy M Thomas, Luke Osborn, Robert W Nickl, Daniel N Candrea, Eric A Pohlmeier, Margaret C Thompson, Manuel Anaya, and Wouter Schellekens. 2020. 'Intracortical microstimulation elicits human fingertip sensations', *medRxiv*.
- Flesher, S. N., J. L. Collinger, S. T. Foldes, J. M. Weiss, J. E. Downey, E. C. Tyler-Kabara, S. J. Bensmaia, A. B. Schwartz, M. L. Boninger, and R. A. Gaunt. 2016. 'Intracortical microstimulation of human somatosensory cortex', *Science Translational Medicine*, 8: 361ra141-361ra141.
- Flesher, Sharlene N, John E Downey, Jeffrey M Weiss, Christopher L Hughes, Angelica J Herrera, Elizabeth C Tyler-Kabara, Michael L Boninger, Jennifer L Collinger, and Robert A Gaunt. 2019. 'Restored tactile sensation improves neuroprosthetic arm control', *bioRxiv*: 653428.
- . 2021. 'A brain-computer interface that evokes tactile sensations improves robotic arm control', *Science*, 372: 831-36.
- Forget, R, and Y Lamarre. 1987. 'Rapid elbow flexion in the absence of proprioceptive and cutaneous feedback', *Hum Neurobiol*, 6: 27-37.

- Formento, Emanuele, Edoardo D'Anna, Sandra Gribi, Stéphanie P Lacour, and Silvestro Micera. 2020. 'A biomimetic electrical stimulation strategy to induce asynchronous stochastic neural activity', *Journal of Neural Engineering*, 17: 046019.
- Freeman, John A. 1971. 'An electronic stimulus artifact suppressor', *Electroencephalography and clinical neurophysiology*, 31: 170-72.
- Fridman, Gene Y., Hugh T. Blair, Aaron P. Blaisdell, and Jack W. Judy. 2010. 'Perceived intensity of somatosensory cortical electrical stimulation', *Experimental Brain Research*, 203: 499-515.
- Georgopoulos, Ap, Jf Kalaska, R. Caminiti, and Jt Massey. 1982. 'On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex', *The Journal of Neuroscience*, 2: 1527-37.
- Ghez, C., J. Gordon, M. F. Ghilardi, C. N. Christakos, and S. E. Cooper. 1990. 'Roles of proprioceptive input in the programming of arm trajectories', *Cold Spring Harbor Symposia on Quantitative Biology*, 55: 837-47.
- Ghez, Claude, James Gordon, and MARIA FELICE Ghilardi. 1995. 'Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy', *Journal of Neurophysiology*, 73: 361-72.
- Gilbert, Charles D, and Torsten N Wiesel. 1992. 'Receptive field dynamics in adult primary visual cortex', *Nature*, 356: 150-52.
- Godlove, Jason M., Erin O. Whaite, and Aaron P. Batista. 2014. 'Comparing temporal aspects of visual, tactile, and microstimulation feedback for motor control', *Journal of Neural Engineering*, 11: 046025.
- Gustafsson, B., and E. Jankowska. 1976. 'Direct and indirect activation of nerve cells by electrical pulses applied extracellularly', *The Journal of Physiology*, 258: 33-61.
- Hao, Yaoyao, Alexa Riehle, and Thomas G. Brochier. 2016. 'Mapping Horizontal Spread of Activity in Monkey Motor Cortex Using Single Pulse Microstimulation', *Frontiers in Neural Circuits*, 10.
- Hashimoto, Takao, Christopher M Elder, and Jerrold L Vitek. 2002. 'A template subtraction method for stimulus artifact removal in high-frequency deep brain stimulation', *Journal of Neuroscience Methods*, 113: 181-86.

Helms, J, J Müller, F Schön, L Moser, W Arnold, T Janssen, R Ramsden, C Von Ilberg, J Kiefer, and T Pfennigdorf. 1997. 'Evaluation of performance with the COMBI 40 cochlear implant in adults: a multicentric clinical study', *ORL*, 59: 23-35.

Histed, Mark H., Vincent Bonin, and R. Clay Reid. 2009. 'Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation', *Neuron*, 63: 508-22.

Hochberg, Leigh R, Daniel Bacher, Beata Jarosiewicz, Nicolas Y Masse, John D Simeral, Joern Vogel, Sami Haddadin, Jie Liu, Sydney S Cash, and Patrick Van Der Smagt. 2012. 'Reach and grasp by people with tetraplegia using a neurally controlled robotic arm', *Nature*, 485: 372.

Houk, James, and William Simon. 1967. 'Responses of Golgi tendon organs to forces applied to muscle tendon', *Journal of Neurophysiology*, 30: 1466-81.

Houk, JC, WZ Rymer, and PE Crago. 1981. 'Nature of the dynamic response and its relation to the high sensitivity of muscle spindles to small changes in length.' in, *Muscle receptors and movement* (Springer).

Hu, Jia Ming, Mei Zhen Qian, Hisashi Tanigawa, Xue Mei Song, and Anna Wang Roe. 2020. 'Focal electrical stimulation of cortical functional networks', *Cerebral Cortex*, 30: 5532-43.

Hubel, D. H., and T. N. Wiesel. 1962. 'Receptive fields, binocular interaction and functional architecture in the cat's visual cortex', *The Journal of Physiology*, 160: 106-54.

Hughes, Aaron C, Liang Guo, and Stephen P DeWeerth. 2010. "Interleaved multichannel epimysial stimulation for eliciting smooth contraction of muscle with reduced fatigue." In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, 6226-29. IEEE.

Hughes, Christopher L, Sharlene N Flesher, and Robert A Gaunt. 2021. 'Effects of stimulus pulse rate on somatosensory adaptation in the human cortex', *bioRxiv*.

Hughes, Christopher L, Sharlene N Flesher, Jeffrey M Weiss, Michael Boninger, Jennifer L Collinger, and Robert A Gaunt. 2021. 'Perception of microstimulation frequency in human somatosensory cortex', *Elife*, 10: e65128.

Hyvärinen, Aapo, and Erkki Oja. 2000. 'Independent component analysis: algorithms and applications', *Neural networks*, 13: 411-30.

Hyvärinen, J, and A Poranen. 1978. 'Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey', *The Journal of Physiology*, 283: 539-56.

- Iwamura, Y, M Tanaka, M Sakamoto, and O Hikosaka. 1983. 'Functional subdivisions representing different finger regions in area 3 of the first somatosensory cortex of the conscious monkey', *Exp Brain Res*, 51: 315-26.
- Iwamura, Yoshiaki, Atsushi Iriki, and Michio Tanaka. 1994. 'Bilateral hand representation in the postcentral somatosensory cortex', *Nature*, 369: 554-56.
- Iwamura, Yoshiaki, Michio Tanaka, Masahiro Sakamoto, and Okihide Hikosaka. 1993. 'Rostrocaudal gradients in the neuronal receptive field complexity in the finger region of the alert monkey's postcentral gyrus', *Experimental Brain Research*, 92: 360-68.
- Jankowska, E., Y. Padel, and R. Tanaka. 1975. 'The mode of activation of pyramidal tract cells by intracortical stimuli', *The Journal of Physiology*, 249: 617-36.
- Johansson, Roland S, and J Randall Flanagan. 2009. 'Coding and use of tactile signals from the fingertips in object manipulation tasks', *Nature Reviews Neuroscience*, 10: 345.
- Johnson, KO. 1980a. 'Sensory discrimination: decision process', *Journal of Neurophysiology*, 43: 1771-92.
- . 1980b. 'Sensory discrimination: neural processes preceding discrimination decision', *Journal of Neurophysiology*, 43: 1793-815.
- Jung, Hyunjun, Jintae Kim, and Yoonkey Nam. 2018. 'Recovery of early neural spikes from stimulation electrodes using a DC-coupled low gain high resolution data acquisition system', *Journal of Neuroscience Methods*, 304: 118-25.
- Kang, Kukjin, Michael Shelley, and Haim Sompolinsky. 2003. 'Mexican hats and pinwheels in visual cortex', *Proceedings of the National Academy of Sciences*, 100: 2848-53.
- Kim, Sungshin, Thierry Callier, Gregg A. Tabot, Robert A. Gaunt, Francesco V. Tenore, and Sliman J. Bensmaia. 2015. 'Behavioral assessment of sensitivity to intracortical microstimulation of primate somatosensory cortex', *Proceedings of the National Academy of Sciences*, 112: 15202-07.
- Kim, Sungshin, Thierry Callier, Gregg A. Tabot, Francesco V. Tenore, and Sliman J. Bensmaia. 2015. 'Sensitivity to microstimulation of somatosensory cortex distributed over multiple electrodes', 9.
- Koivuniemi, A. S., and K. J. Otto. 2011. 'Asymmetric Versus Symmetric Pulses for Cortical Microstimulation', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19: 468-76.

Kral, Andrej, Michael F Dorman, and Blake S Wilson. 2019. 'Neuronal development of hearing and language: cochlear implants and critical periods', *Annual review of neuroscience*, 42: 47-65.

Kral, Andrej, and Gerard M O'Donoghue. 2010. 'Profound deafness in childhood', *New England Journal of Medicine*, 363: 1438-50.

Krubitzer, Leah, Kelly J Huffman, Elizabeth Disbrow, and Gregg Recanzone. 2004. 'Organization of area 3a in macaque monkeys: contributions to the cortical phenotype', *Journal of Comparative Neurology*, 471: 97-111.

Kumaravelu, Karthik, Joseph Sombeck, Lee E Miller, Sliman J Bensmaia, and Warren M Grill. 2021. 'Stoney vs. Histed: Quantifying the Spatial Effects of Intracortical Microstimulation', *bioRxiv*.

———. 2022. 'Stoney vs. Histed: Quantifying the spatial effects of intracortical microstimulation', *Brain Stimulation*, 15: 141-51.

Kumaravelu, Karthik, Tucker Tomlinson, Thierry Callier, Joseph Sombeck, Sliman J Bensmaia, Lee E Miller, and Warren M Grill. 2020. 'A comprehensive model-based framework for optimal design of biomimetic patterns of electrical stimulation for prosthetic sensation', *Journal of Neural Engineering*, 17: 046045.

Lee, Brian, Daniel Kramer, Michelle Armenta Salas, Spencer Kellis, David Brown, Tatyana Dobrova, Christian Klaes, Christi Heck, Charles Liu, and Richard A Andersen. 2018. 'Engineering artificial somatosensation through cortical stimulation in humans', *Frontiers in Systems Neuroscience*, 12: 24.

Lemm, Steven, Gabriel Curio, Yevhen Hlushchuk, and K-R Muller. 2006. 'Enhancing the signal-to-noise ratio of ICA-based extracted ERPs', *IEEE Transactions on Biomedical Engineering*, 53: 601-07.

Logothetis, Nikos K., Mark Augath, Yusuke Murayama, Alexander Rauch, Fahad Sultan, Jozien Goense, Axel Oeltermann, and Hellmut Merkle. 2010. 'The effects of electrical microstimulation on cortical signal propagation', *Nature Neuroscience*, 13: 1283-91.

London, B. M., L. R. Jordan, C. R. Jackson, and L. E. Miller. 2008. 'Electrical Stimulation of the Proprioceptive Cortex (Area 3a) Used to Instruct a Behaving Monkey', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 16: 32-36.

London, Brian M., and Lee E. Miller. 2012. 'Responses of somatosensory area 2 neurons to actively and passively generated limb movements', *Journal of Neurophysiology*, 109: 1505-13.

- Malonek, D, RBH Tootell, and Amiram Grinvald. 1994. 'Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT', *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 258: 109-19.
- Margalit, Shany Nivinsky, and Hamutal Slovin. 2018. 'Spatio-temporal characteristics of population responses evoked by microstimulation in the barrel cortex', *Scientific Reports*, 8.
- Markram, Henry, Eilif Muller, Srikanth Ramaswamy, Michael W Reimann, Marwan Abdellah, Carlos Aguado Sanchez, Anastasia Ailamaki, Lidia Alonso-Nanclares, Nicolas Antille, and Selim Arsever. 2015. 'Reconstruction and simulation of neocortical microcircuitry', *Cell*, 163: 456-92.
- Maunsell, John H, and David C Van Essen. 1983. 'Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation', *Journal of Neurophysiology*, 49: 1127-47.
- Maynard, Edwin M, Craig T Nordhausen, and Richard A Normann. 1997. 'The Utah intracortical electrode array: a recording structure for potential brain-computer interfaces', *Electroencephalography and clinical neurophysiology*, 102: 228-39.
- McCreery, Douglas, Victor Pikov, and Philip R Troyk. 2010. 'Neuronal loss due to prolonged controlled-current stimulation with chronically implanted microelectrodes in the cat cerebral cortex', *Journal of Neural Engineering*, 7: 036005.
- McElvain, Lauren E, Martha W Bagnall, Alexandra Sakatos, and Sascha Du Lac. 2010. 'Bidirectional plasticity gated by hyperpolarization controls the gain of postsynaptic firing responses at central vestibular nerve synapses', *Neuron*, 68: 763-75.
- Mendrela, Adam E, Jihyun Cho, Jeffrey A Fredenburg, Vivek Nagaraj, Theoden I Netoff, Michael P Flynn, and Euisik Yoon. 2016. 'A bidirectional neural interface circuit with active stimulation artifact cancellation and cross-channel common-mode noise suppression', *IEEE Journal of Solid-State Circuits*, 51: 955-65.
- Michelson, Nicholas J., James R. Eles, Alberto L. Vazquez, Kip A. Ludwig, and Takashi D. Y. Kozai. 2018. 'Calcium activation of cortical neurons by continuous electrical stimulation: Frequency dependence, temporal fidelity, and activation density', *Journal of Neuroscience Research*.
- Molineux, Michael L, John E McRory, Bruce E McKay, Jawed Hamid, W Hamish Mehaffey, Renata Rehak, Terrance P Snutch, Gerald W Zamponi, and Ray W Turner. 2006. 'Specific T-type calcium channel isoforms are associated with distinct burst phenotypes in deep cerebellar nuclear neurons', *Proceedings of the National Academy of Sciences*, 103: 5555-60.

Montgomery Jr, Erwin B, John T Gale, and He Huang. 2005. 'Methods for isolating extracellular action potentials and removing stimulus artifacts from microelectrode recordings of neurons requiring minimal operator intervention', *Journal of Neuroscience Methods*, 144: 107-25.

Monzée, Joël, Yves Lamarre, and Allan M Smith. 2003. 'The effects of digital anesthesia on force control using a precision grip', *Journal of Neurophysiology*, 89: 672-83.

Mountcastle, Vernon B. 1957. 'Modality and topographic properties of single neurons of cat's somatic sensory cortex', *Journal of Neurophysiology*, 20: 408-34.

Mueller, Jerel K, Erinn M Grigsby, Vincent Prevosto, Frank W Petraglia, Hrishikesh Rao, Zhi-De Deng, Angel V Peterchev, Marc A Sommer, Tobias Egner, and Michael L Platt. 2014. 'Simultaneous transcranial magnetic stimulation and single-neuron recording in alert non-human primates', *Nature Neuroscience*, 17: 1130-36.

Murasugi, Cm, Cd Salzman, and Wt Newsome. 1993. 'Microstimulation in visual area MT: effects of varying pulse amplitude and frequency', *The Journal of Neuroscience*, 13: 1719-29.

Nag, Sudip, Sujit Kumar Sikdar, Nitish Vyomesh Thakor, Valipe Ramgopal Rao, and Dinesh Sharma. 2015. 'Sensing of stimulus artifact suppressed signals from electrode interfaces', *IEEE Sensors Journal*, 15: 3734-42.

Nowak, L. G., and J. Bullier. 1998a. 'Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. I. Evidence from chronaxie measurements', *Experimental Brain Research*, 118: 477-88.

———. 1998b. 'Axons, but not cell bodies, are activated by electrical stimulation in cortical gray matter. II. Evidence from selective inactivation of cell bodies and axon initial segments', *Experimental Brain Research*, 118: 489-500.

O'Doherty, J. E., M. A. Lebedev, Li Zheng, and M. A. L. Nicolelis. 2012. 'Virtual Active Touch Using Randomly Patterned Intracortical Microstimulation', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20: 85-93.

O'Shea, Daniel J., and Krishna V. Shenoy. 2017. 'ERAASR: An algorithm for removing electrical stimulation artifacts from multielectrode array recordings'.

O'Doherty, Joseph E, Mikhail A Lebedev, Peter J Ifft, Katie Z Zhuang, Solaiman Shokur, Hannes Bleuler, and Miguel AL Nicolelis. 2011a. 'Active tactile exploration using a brain-machine-brain interface', *Nature*, 479: 228-31.

O'Doherty, Joseph E., Mikhail A. Lebedev, Peter J. Ifft, Katie Z. Zhuang, Solaiman Shokur, Hannes Bleuler, and Miguel A. L. Nicolelis. 2011b. 'Active tactile exploration using a brain-machine-brain interface', *Nature*, 479: 228-31.

Osborn, Luke E, Breanne P Christie, David P McMullen, Robert W Nickl, Margaret C Thompson, Ambarish S Pawar, Tessy M Thomas, Manuel Alejandro Anaya, Nathan E Crone, and Brock A Wester. 2021. "Intracortical microstimulation of somatosensory cortex enables object identification through perceived sensations." In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 6259-62. IEEE.

Otto, Kevin J., Patrick J. Rousche, and Daryl R. Kipke. 2005. 'Cortical microstimulation in auditory cortex of rat elicits best-frequency dependent behaviors', *Journal of Neural Engineering*, 2: 42-51.

Overstreet, C. K., J. D. Klein, and S. I. Helms Tillery. 2013. 'Computational modeling of direct neuronal recruitment during intracortical microstimulation in somatosensory cortex', *Journal of Neural Engineering*, 10: 066016.

Öztürk, Sevgi, Ismail Devecioğlu, Mohammad Beygi, Ahmet Atasoy, Şenol Mutlu, Mehmed Özkan, and Burak Güçlü. 2019. 'Real-Time Performance of a Tactile Neuroprosthesis on Awake Behaving Rats', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 27: 1053-62.

Padberg, Jeffrey, Dylan F Cooke, Christina M Cerkevich, Jon H Kaas, and Leah Krubitzer. 2019. 'Cortical connections of area 2 and posterior parietal area 5 in macaque monkeys', *Journal of Comparative Neurology*, 527: 718-37.

Parker, Rebecca A., Tyler S. Davis, Paul A. House, Richard A. Normann, and Bradley Greger. 2011. 'The functional consequences of chronic, physiologically effective intracortical microstimulation.' in *Progress in Brain Research* (Elsevier).

Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, and Luca Antiga. 2019. 'Pytorch: An imperative style, high-performance deep learning library', *Advances in neural information processing systems*, 32.

Pins, Delphine, and Claude Bonnet. 1996. 'On the relation between stimulus intensity and processing time: Piéron's law and choice reaction time', *Perception & Psychophysics*, 58: 390-400.

Pons, T. P., P. E. Garraghty, C. G. Cusick, and J. H. Kaas. 1985. 'The somatotopic organization of area 2 in macaque monkeys', *The Journal of Comparative Neurology*, 241: 445-66.

- Proske, Uwe, and Simon C. Gandevia. 2012. 'The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force', *Physiological Reviews*, 92: 1651-97.
- Prud'homme, M. J., and J. F. Kalaska. 1994. 'Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements', *Journal of Neurophysiology*, 72: 2280-301.
- Rajan, Alexander T., Jessica L. Boback, John F. Dammann, Francesco V. Tenore, Brock A. Wester, Kevin J. Otto, Robert A. Gaunt, and Sliman J. Bensmaia. 2015. 'The effects of chronic intracortical microstimulation on neural tissue and fine motor behavior', *Journal of Neural Engineering*, 12: 066018.
- Ringach, Dario L. 2004. 'Mapping receptive fields in primary visual cortex', *The Journal of Physiology*, 558: 717-28.
- Roe, Anna W, Leonardo Chelazzi, Charles E Connor, Bevil R Conway, Ichiro Fujita, Jack L Gallant, Haidong Lu, and Wim Vanduffel. 2012. 'Toward a unified theory of visual area V4', *Neuron*, 74: 12-29.
- Rolston, JD, RE Gross, and SM Potter. 2009. 'A low-cost multielectrode system for data acquisition and real-time processing with rapid recovery from stimulation artifacts', *Frontiers in Neuroengineering*, 2: 1-17.
- Romo, Ranulfo, Adrián Hernández, Anótonio Zainos, and Emilio Salinas. 1998. 'Somatosensory discrimination based on cortical microstimulation', *Nature*, 392: 387-90.
- Romo, Ranulfo, Adrián Hernández, Antonio Zainos, Carlos D. Brody, and Luis Lemus. 2000. 'Sensing without Touching: Psychophysical Performance Based on Cortical Microstimulation', *Neuron*, 26: 273-78.
- Rothwell, JC, MM Traub, BL Day, JA Obeso, PK Thomas, and CD Marsden. 1982. 'Manual motor performance in a deafferented man', *Brain*, 105: 515-42.
- Sainburg, R. L., M. F. Ghilardi, H. Poizner, and C. Ghez. 1995. 'Control of limb dynamics in normal subjects and patients without proprioception', *Journal of Neurophysiology*, 73: 820-35.
- Sainburg, R. L., H. Poizner, and C. Ghez. 1993. 'Loss of proprioception produces deficits in interjoint coordination', *Journal of Neurophysiology*, 70: 2136-47.
- Salas, Michelle Armenta, Luke Bashford, Spencer Kellis, Matiar Jafari, HyeongChan Jo, Daniel Kramer, Kathleen Shanfield, Kelsie Pejasa, Brian Lee, Charles Y. Liu, and Richard A. Andersen.

2018. 'Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation', *Elife*, 7.
- Salzman, C Daniel, Chieko M Murasugi, Kenneth H Britten, and William T Newsome. 1992. 'Microstimulation in visual area MT: effects on direction discrimination performance', *Journal of Neuroscience*, 12: 2331-55.
- Schmidt, EM, MJ Bak, FT Hambrecht, CV Kufta, DK O'rourke, and P Vallabhanath. 1996. 'Feasibility of a visual prosthesis for the blind based on intracortical micro stimulation of the visual cortex', *Brain*, 119: 507-22.
- Serruya, Mijail D, Nicholas G Hatsopoulos, Liam Paninski, Matthew R Fellows, and John P Donoghue. 2002. 'Instant neural control of a movement signal', *Nature*, 416: 141-42.
- Shanечи, Maryam M, Amy L Orsborn, Helene G Moorman, Suraj Gowda, Siddharth Dangi, and Jose M Carmena. 2017. 'Rapid control and feedback rates enhance neuroprosthetic control', *Nature Communications*, 8: 13825.
- Softky, William R, and Christof Koch. 1993. 'The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs', *Journal of Neuroscience*, 13: 334-50.
- Sombeck, Joseph, Juliet Heye, Karthik Kumaravelu, Stefan M Goetz, Angel V Peterchev, Warren M Grill, Sliman J Bensmaia, and Lee E Miller. 2021. 'Characterizing the short-latency evoked response to intracortical microstimulation across a multi-electrode array', *bioRxiv*.
- Sombeck, Joseph T, and Lee E Miller. 2019. 'Short reaction times in response to multi-electrode intracortical microstimulation may provide a basis for rapid movement-related feedback', *Journal of Neural Engineering*, 17: 016013.
- Soso, MJ, and EE Fetz. 1980. 'Responses of identified cells in postcentral cortex of awake monkeys during comparable active and passive joint movements', *Journal of Neurophysiology*, 43: 1090-110.
- Sripati, Arun P, Takashi Yoshioka, Peter Denchev, Steven S Hsiao, and Kenneth O Johnson. 2006. 'Spatiotemporal receptive fields of peripheral afferents and cortical area 3b and 1 neurons in the primate somatosensory system', *Journal of Neuroscience*, 26: 2101-14.
- Stoney, S. D., W. D. Thompson, and H. Asanuma. 1968. 'Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current', *Journal of Neurophysiology*, 31: 659-69.

- Swadlow, Harvey A. 1990. 'Efferent neurons and suspected interneurons in S-1 forelimb representation of the awake rabbit: receptive fields and axonal properties', *Journal of Neurophysiology*, 63: 1477-98.
- Tabot, Gregg A., John F. Dammann, Joshua A. Berg, Francesco V. Tenore, Jessica L. Boback, R. Jacob Vogelstein, and Sliman J. Bensmaia. 2013. 'Restoring the sense of touch with a prosthetic hand through a brain interface', *Proceedings of the National Academy of Sciences of the United States of America*, 110: 18279-84.
- Tamada, Daiki, Marie-Luise Kromrey, Shintaro Ichikawa, Hiroshi Onishi, and Utaroh Motosugi. 2020. 'Motion artifact reduction using a convolutional neural network for dynamic contrast enhanced MR imaging of the liver', *Magnetic Resonance in Medical Sciences*, 19: 64.
- Tan, Daniel W, Matthew A Schiefer, Michael W Keith, James Robert Anderson, Joyce Tyler, and Dustin J Tyler. 2014. 'A neural interface provides long-term stable natural touch perception', *Science Translational Medicine*, 6: 257ra138-257ra138.
- Tanaka, Yuta, Tomohiro Nomoto, Takuya Shiki, Yuya Sakata, Yoshinori Shimada, Yuki Hayashida, and Tetsuya Yagi. 2019. 'Focal activation of neuronal circuits induced by microstimulation in the visual cortex', *Journal of Neural Engineering*, 16: 036007.
- Taylor, Dawn M, Stephen I Helms Tillery, and Andrew B Schwartz. 2002. 'Direct cortical control of 3D neuroprosthetic devices', *Science*, 296: 1829-32.
- Tehovnik, E. J., A. S. Tolias, F. Sultan, W. M. Slocum, and N. K. Logothetis. 2006. 'Direct and Indirect Activation of Cortical Neurons by Electrical Microstimulation', *Journal of Neurophysiology*, 96: 512-21.
- Tehovnik, Edward J, and Warren M Slocum. 2009. 'Depth-dependent detection of microampere currents delivered to monkey V1', *European Journal of Neuroscience*, 29: 1477-89.
- Tolhurst, David J, J Anthony Movshon, and Andrew F Dean. 1983. 'The statistical reliability of signals in single neurons in cat and monkey visual cortex', *Vision research*, 23: 775-85.
- Tomlinson, Tucker, and Lee E. Miller. 2016. 'Toward a Proprioceptive Neural Interface that Mimics Natural Cortical Activity.' in Jozsef Laczko and Mark L. Latash (eds.), *Progress in Motor Control* (Springer International Publishing: Cham).
- Van Essen, DC, JHR Maunsell, and JL Bixby. 1981. 'The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization', *Journal of Comparative Neurology*, 199: 293-326.
- van Rossum, Guido. 1995. 'Python reference manual', *Department of Computer Science [CS]*.

- Venkatraman, Subramaniam, Ken Elkabany, John D Long, Yimin Yao, and Jose M Carmena. 2008. 'A system for neural recording and closed-loop intracortical microstimulation in awake rodents', *IEEE Transactions on Biomedical Engineering*, 56: 15-22.
- Versteeg, Christopher, Raed H Chowdhury, and Lee E Miller. 2021. 'Cuneate nucleus: The somatosensory gateway to the brain', *Current Opinion in Physiology*, 20: 206-15.
- Versteeg, Christopher, Joshua M. Rosenow, Sliman J. Bensmaia, and Lee E. Miller. 2021. 'Encoding of limb state by single neurons in the cuneate nucleus of awake monkeys', *Journal of Neurophysiology*, 126: 693-706.
- Voigt, Mathias Benjamin, Prasandhya Astagiri Yusuf, and Andrej Kral. 2018. 'Intracortical microstimulation modulates cortical induced responses', *The Journal of Neuroscience*: 0928-18.
- Weber, Douglas J., Brian M. London, James A. Hokanson, Christopher A. Ayers, Robert A. Gaunt, Ricardo R. Torres, Boubker Zaaïmi, and Lee E. Miller. 2011. 'Limb-State Information Encoded by Peripheral and Central Somatosensory Neurons: Implications for an Afferent Interface', *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19: 501-13.
- Weiss, Jeffrey M, Sharlene N Flesher, Robert Franklin, Jennifer L Collinger, and Robert A Gaunt. 2018. 'Artifact-free recordings in human bidirectional brain–computer interfaces', *Journal of Neural Engineering*, 16: 016002.
- Wilson, Blake S, and Michael F Dorman. 2008. 'Interfacing sensors with the nervous system: lessons from the development and success of the cochlear implant', *IEEE Sensors Journal*, 8: 131-47.
- Wodlinger, B, JE Downey, EC Tyler-Kabara, AB Schwartz, ML Boninger, and JL Collinger. 2014. 'Ten-dimensional anthropomorphic arm control in a human brain– machine interface: difficulties, solutions, and limitations', *Journal of Neural Engineering*, 12: 016011.
- Yau, Jeffrey M, Sung Soo Kim, Pramodsingh H Thakur, and Sliman J Bensmaia. 2016. 'Feeling form: the neural basis of haptic shape perception', *Journal of Neurophysiology*, 115: 631-42.
- Young, D, F Willett, WD Memberg, B Murphy, B Walter, J Sweet, J Miller, LR Hochberg, RF Kirsch, and AB Ajiboye. 2018. 'Signal processing methods for reducing artifacts in microelectrode brain recordings caused by functional electrical stimulation', *Journal of Neural Engineering*, 15: 026014.
- Zaaïmi, Boubker, Ricardo Ruiz-Torres, Sara A. Solla, and Lee E. Miller. 2013. 'Multi-electrode stimulation in somatosensory cortex increases probability of detection', *Journal of Neural Engineering*, 10: 056013.

Zhang, Yanbo, and Hengyong Yu. 2018. 'Convolutional neural network based metal artifact reduction in x-ray computed tomography', *IEEE transactions on medical imaging*, 37: 1370-81.

Zhou, Andy, Benjamin C. Johnson, and Rikky Muller. 2018. 'Toward true closed-loop neuromodulation: artifact-free recording during stimulation', *Current Opinion in Neurobiology*, 50: 119-27.