Title
Artificial Sensory Feedback for Neural Prostheses

Author
Dadarlat, Maria C.

Publication Date
2014

Peer reviewed|Thesis/dissertation

Artificial Sensory Feedback for Neural Prostheses

by

Maria C. Dadarlat

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

AND

Artificial Sensory Feedback for Neural Prostheses

To Iunia, Marius, and Voichita Dadarlat

De la o gogolie zburatoare.

Acknowledgments

I have many people to thank for contributing to the work described in this dissertation, but foremost I am indebted to my thesis advisor, Dr. Philip Sabes (Flip), as without him this project would have been a weaker, more chaotic version of what exists today. I would also like to graciously thank my thesis committee—Dr. Michael Brainard, Dr. Jose Carmena, and Dr. Edward Chang—who were supportive and enthusiastic about this project from the beginning to the end. Next, I must acknowledge and give my heartfelt thanks to our funding source, the Defense Advanced Research Projects Agency, who generously sponsored this work. Without their support, we never could have completed such an ambitious project.

In the years I spent here, I greatly benefitted from working alongside my fellow lab-mates, who served as excellent scientific advisors as well as friends. In particular, I would like to thank Joseph G. Makin, for the time and energy he contributed to thinking through the implications and methodology of my analysis. I am also indebted to Joseph O'Doherty, who advised on and helped with many of the practical issues that were crucially important to this thesis, including (importantly) aesthetic considerations in the design of figures. Without him I wouldn't have known how to test a stimulator or if a monkey would ever, *ever* learn to detect ICMS. Next, Matthew Fellows taught me everything I know about training monkeys, and spent many afternoons patiently explaining important topics in systems neuroscience. To him and to Joe Makin I owe the justification of this thesis project.

A boundless thanks is in order to our animal experts, Blakely Andrews and Karen MacLeod, who made working with our animal model infinitely more bearable, and who always made sure my i's were dotted and my t's were crossed. I'd also like to thank my collaborators on this project, Dmitri and Fievel, who were crucial to the experimental outcomes and were generally sweet and good-natured—a real pleasure to work with.

Lastly, I would like to acknowledge my family, my most earnest cheerleaders: Iunia, Voichi, and Marius. Thank you for your sweetness, your interest, and your faith.

Abstract

Artificial Sensory Feedback for Neural Prostheses

by

Maria C. Dadarlat

Doctor of Philosophy in Bioengineering

University of California, San Francisco

Professor Philip Sabes, Chair

When controlling a motor prosthetic device or computer cursor using a brain-machine interface (BMI), users guide their movements relying solely on visual feedback. In contrast, our natural ability to plan and execute movements relies on feedback from multiple sensory signals, and in particular on proprioception—the sense of the bodys position in space. Proprioception guides movement without the need for continuous visual monitoring and is integrated with vision, when both are available, to improve behavioral performance. Such multisensory integration appears to be learned from multisensory experience, a process that theoretically can be driven solely by the shared statistical structure of the inputs. To achieve naturalistic control over prosthetic devices, BMIs will likewise need a proprioceptive feedback signal, one that can ultimately be integrated with natural vision. Here, we demonstrate a novel, learning-based approach to artificial sensory feedback. We show that monkeys can learn to use, in a naturalistic way, a continuous, multi-dimensional, multi-channel intracorti-

cal microstimulation (ICMS) signal that encodes task-relevant feedback. After training with correlated ICMS and visual feedback, monkeys could perform a goal-directed reaching task without vision, using only the ICMS signal. Additionally, the animals integrated the natural and artificial sensory inputs, combining both into a minimum-variance sensory estimate of hand position relative to the target. The ICMS signal further resembled natural sensation in its plasticity. For example, when a misalignment was imposed between vision and ICMS, the ICMS estimate adapted back towards the rewarded visual cue. Furthermore, the monkey's valuation of the accuracy of the ICMS estimate was updated in response to trial conditions. When the ICMS is randomly, but not consistently, perturbed, the monkey's estimate using artificial feedback grows less precise. Finally, we discussed theoretical ways to improve the ICMS signal reliability by changing the parameters of the signal. Together these results demonstrate that a learning-based approach can be used to provide a rich artificial sensory signal, suggesting a new strategy for restoring proprioception to patients with BMIs as well as a powerful new tool for studying the adaptive mechanisms of sensory integration.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Brain-Machine Interfaces

The field of motor neural prosthetics aims to restore sensorimotor function to patients who have suffered traumatic spinal chord injury. Fully repairing these injuries means restoring both the ascending (sensory) and descending (motor) pathways of information flow. This can be accomplished either by trying to repair the severed connections (Angeli et al., 2014; van den Brand et al., 2012) or by artificially re-establishing the link between the brain and the periphery (Carmena et al., 2003; Ethier et al., 2012; Fraser et al., 2009; Gilja et al., 2012; Hatsopoulos et al., 2005; Ifft et al., 2013; Kim et al., 2011; Nishimura et al., 2013; O'Doherty et al., 2009; Orsborn et al., 2014; ODoherty et al., 2011; Shanechi et al., 2014; Suminski et al., 2010).

Systems that perform the latter function are called Brain-Machine Interfaces (BMIs) and consist of a network of devices that record neural activity in the brain, decode the user's motor intentions, and then translate those intentions into control over a wheelchair, robot arm, or even the patient's own limbs (Carmena et al., 2003; Ethier et al., 2012; Fraser et al., 2009; Gilja et al., 2012; Hatsopoulos et al., 2005; Huang et al., 2009; Ifft et al., 2013; Kim et al., 2011; Nishimura et al., 2013; O'Doherty et al., 2009; Orsborn et al., 2014; ODoherty et al., 2011; Shanechi et al., 2014; Suminski et al., 2010; Zimmermann and Jackson, 2014).

Although some portion of BMI research is conducted with human subjects (Donoghue et al., 2007; Hochberg et al., 2006; Kim et al., 2008, 2011; Wolpaw and McFarland, 2004), ethical and practical considerations dictate that much of the preliminary work is performed in non-human primates. These experiments typically involve monkeys moving a cursor around a computer screen by modulating neural activity in primary motor cortex (M1), premotor cortex (PMd or PMv), or parietal cortex (Carmena et al., 2003; Fraser et al., 2009; Gilja et al., 2012; Hatsopoulos et al., 2005; Hauschild et al., 2012; Ifft et al., 2013; Nishimura et al., 2013; O'Doherty et al., 2009; Orsborn et al., 2014; ODoherty et al., 2011; Shanechi et al., 2014; Suminski et al., 2010). Movements made with cursors under BMI control are getting faster and more precise—thanks in part to innovations in decoding algorithms (Bishop et al., 2014; Dangi et al., 2013; Gilja et al., 2012; Ifft et al., 2013; Kao et al., 2013; Orsborn et al., 2014; Shanechi et al., 2014, 2013; Yeom et al., 2014), but somehow they have yet to achieve the fluidity of natural motor control.

Researchers have hypothesized that this lack is due in part to an absence of somatosensory

feedback to the monkeys during BMI control. There's ample evidence that an artificial proprioceptive signal (the sense of the device's position in space) would be useful: monkeys (Suminski et al., 2010) and humans (Gomez-Rodriguez et al., 2011) can improve BMI control if they receive natural proprioceptive feedback in the form of passive movements of the limb. We expect that the benefit of proprioceptive feedback to BMI control will persist even with further refinements in the software and hardware of Brain-Machine interfaces; after all, just as deafferented patients with nominally intact motor systems exhibit motor impairments, a motor neural prosthesis without somatic sensation will remain functionally impaired. *Thus one of the next important steps for neural prosthetic devices is to provide artificial sensory feedback, restoring the flow of information from the somatosensory periphery to the cortex.*

1.2 The importance of somatosensation

Somatosensory inputs such as touch and proprioception are essential for fine control of movements and the dexterous manipulation of objects (Johansson and Flanagan, 2009; Sainburg et al., 1993). Proprioception serves to generate a representation of the body's posture and movements (Kalaska, 1994), information that is crucial for future movement planning and error correction. This data originates as a distributed network of sensors embedded in the muscles, skin, and joints (Jones, 1994) which must be translated in the brain into a lower-dimensional description of limb state and movement direction. By the time information reaches primary somatosensory cortex (S1), cells already seem to encode proprioception in more abstract and concise forms, such as static limb position (Prud'homme and Kalaska,

1994) and dynamic movement direction (Pei et al., 2010; Prud'homme and Kalaska, 1994). This information is further transmitted to parietal cortex, including area 5 (Kalaska, 1996), and then to dorsal premotor cortex (PMd) and motor cortex, which are involved in planning and executing reaching movements (Crammond and Kalaska, 1996; Scott and Kalaska, 1997; Scott et al., 1997).

Somatosensation is essential to normal sensorimotor function: it relieves the need for constant visual attention to manipulated objects and permits movements to be made outside of the visual field (into pockets, etc.) but also enables tasks that are difficult to accomplish even with intact vision. For example, the seemingly simple task of lifting an object or striking a match is more challenging and takes longer to accomplish when attempted without cutaneous sensation from the fingers (Johansson and Westling, 1984; Westling and Johansson, 1984). Furthermore, proprioception enables spatially precise visually-guided movement trajectories that require multi-joint coordination (Sainburg et al., 1993). Lastly, somatosensation is more precise than vision in depth perception (van Beers et al., 2002) and has a shorter feedback latency (Cluff et al., 2014). These attributes of natural proprioception make artificial proprioception a desirable component of neural prosthetic devices.

Multisensory integration of vision and proprioception

Proprioception plays a particularly important role in precise movement planning (Sober and Sabes, 2005, 2003). Humans (Ernst and Banks, 2002; van Beers et al., 1999) and animals (Gu et al., 2008) naturally integrate multiple streams of sensory information (e.g., vision

and proprioception) about relevant parameters (e.g., position, velocity, etc. of the hand) into a unified estimate. As a concrete example, consider the problem of reaching to grasp an object. To accomplish the reach, the brain must combine sensory estimates of the hand and the object into the appropriate motor command. Evidence from behavioral experiments indicates that, at least to a first approximation, this is a statistically optimal process, in that the individual modalities are combined in order to minimize the variance of the integrated sensory estimate (Alais and Burr, 2004; Ernst and Banks, 2002; Fetsch et al., 2012; Gu et al., 2008; van Beers et al., 1999). This process can be expressed mathematically as

$$x_{int} = \frac{\sigma_1^{-2} x_1 + \sigma_2^{-2} x_2}{\sigma_1^{-2} + \sigma_2^{-2}} \tag{1.1}$$

where x_1, x_2, and x_{int} are the mean estimates from the individual modalities and the integrated estimate, respectively. The trial-to-trial variances of these estimates, σ_1^2 and σ_2^2, can also be viewed as measures of the sensory uncertainty associated with a given estimate. Note that the integrated estimate places more weight on the modality that has the smaller variance (least uncertainty). The variance of the integrated estimate, σ_{int}^2, given by Equation 1.2, is guaranteed to be smaller than the variances of the individual estimates, as long as they are finite.

$$\sigma_{int}^{-2} = (\sigma_1^{-2} + \sigma_2^{-2})^{-1} \tag{1.2}$$

This decrease in variance is one reason artificial somatosensation is expected to improve the performance of neural prostheses.

Learning to integrate multisensory signals.

We would like to make sure that any artificial sensory feedback signal provided to a BMI user is similarly integrated optimally with visual feedback, but the mechanisms underlying multisensory integration seems to occur both behaviorally and at the neural level (Fetsch et al., 2012; Gu et al., 2008), the underlying neural mechanisms are unknown.

Fortunately, the process of multisensory integration seems to be learned and remains plastic in adulthood (Burge et al., 2010; Ernst, 2007; Simani et al., 2007; Wallace et al., 2006; Wallace and Stein, 1997), so we can adapt models that plausibly explain this process to the act of learning novel inputs. Our lab has recently proposed one such model of adaptive multisensory integration (Makin et al., 2013). The model stems from the idea that multisensory integration can be viewed as one example of a more general, unsupervised learning problem: latent variable density estimation (LVDE). The goal of LVDE is to extract low dimensional representations of incoming data while retaining as much of the original statistical structure as possible. This model is implemented with a simple neural network that learns LVDE via a biologically plausible Hebbian-like learning rule (Hinton et al., 2006; model illustrated in Figure 1.1). It consists of two populations of input neurons—visual neurons encoding hand position in extrinsic coordinates, and proprioceptive neurons encoding hand position in terms of joint angles—and an output population of multisensory neurons that receive projections from the two input populations. Starting from a state of random connectivity, the links from input to output are learned through exposure to data in which there is a strong correlation between the two input populations arising from the fact that they both represent the same

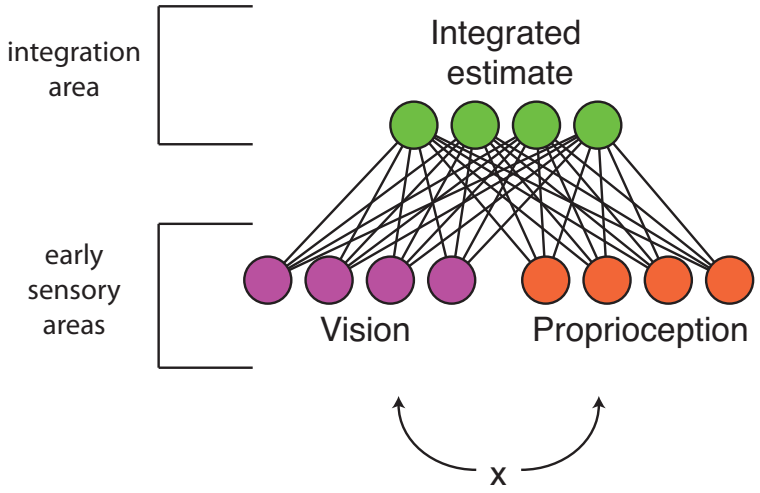

Figure 1.1: **Learning to integrate multisensory signals.** A schematic of a neural network model learns, de novo and in an unsupervised fashion, to integrate visual and proprioceptive feedback of a variable of interest, x, such as the position or velocity of the limb. Learning is driven by correlations between the input populations, which in turn reflect their common encoding of x. This figure is adapted from Makin et al., 2013.

underlying variable(s), x, for example the state of limb. After learning, the network is able to perform minimum-variance cue combination, as well as a range of other movement-related multisensory computations (Makin et al., 2013).

A learning-based approach to artificial feedback

A key insight from the model of Figure 1.1 is that the statistical properties of the input signals, namely the correlation between the activities of the two sensory neural populations, are sufficient to drive the network to learn integrated representations of hand position (Makin et al., 2013), without the need for supervisory signals. We do not know if this model accurately captures the mechanisms implemented in the brain, but its biologically plausible form makes it an exciting candidate and, moreover, it makes several testable predictions of

practical importance for BCI.

One testable prediction that is particularly relevant for artificial somatosensation is that *temporal correlations between two input signals will drive learning and integration.* This suggests a powerful learning-based approach for delivering novel sensory signals to the brain: if one provides a novel stream of information, for example by using intracortical mircostimulation (ICMS), the brain should learn to interpret and integrate that signal as long as it correlates over time with a known sensory signal.

1.3 Intracortical microstimulation

Mimicking natural sensations

Despite the known plasticity of the brain and predictions made by models such as the one described above, the dominant approach to encoding artificial sensation for neural prosthetics has been biomimetic (Fagg et al., 2007)—trying to replicate the patterns of neural activity observed during natural sensory processing. Intuitively, a biomimetic code would be able to instantly provide a rich and easily interpreted artificial sensory feedback signal. Preliminary studies using electrical ICMS show promise, at least for artificial tactile feedback (Berg et al., 2013; Tabot et al., 2013a); however, it not clear how well this approach will extend to encoding more abstract sensations (e.g. proprioception).

First, it is not possible to simply invert the biophysics of neural activity, i.e. to precisely recreate recorded patterns of activity, using presently-available tools. One common approach

is to use electrical stimulation to alter ongoing neural activity. Unlike natural sensory inputs, each pulse of electrical current activates a sphere of neural activity around the stimulation site (Stoney et al., 1968; Tehovnik, 1996; Tehovnik et al., 2006) or, according to another study, a sparse, distributed population of cells whose processes lie proximal to the stimulation site (Histed et al., 2009). In either scenario, it is difficult to target individual neurons without undesired and/or unpredictable collateral activity (Butovas and Schwarz, 2003). This problem will be particularly acute in brain areas that lack a fine-scale topographic map, i.e. where nearby neurons have different response properties, as is the case in the proprioceptive regions of primary somatosensory cortex (Kaas et al., 1979; Weber et al., 2012). Even if very small currents could be used to activate single neurons (Houweling and Brecht, 2008; Voigt et al., 2008), the total number of (directly) activated neurons would be limited by the size of the stimulating array. In the near term, then, it seems infeasible to create targeted spatiotemporal patterns that mimic the precision and complexity of natural sensory activity.

A second obstacle faced by biomimetic approaches is our piecemeal understanding of how those natural patterns of activity encode the sense of touch (Johansson and Flanagan, 2009), proprioception (Prud'homme and Kalaska, 1994), and complex interaction between the two (Rincon-Gonzalez et al., 2012; Warren et al., 2011). Even in cases where elements of this coding are known—e.g. topographic maps for the digits—spinal cord injury with sensory loss raises the additional challenge that the code cannot be mapped using natural stimuli. While the problem could be inverted—stimulating individual sites and asking for reports of

the subsequent percept—this approach would be slow and imprecise.

Yet a third challenge for the biomimetic approach is the cortical remapping that occurs after the loss of sensory afferents. Somatosensory perturbations elicit plastic changes in the adult brain in a matter of months, so that cortical representations of deafferented surfaces become occupied by expanded representations of the surrounding areas (Merzenich et al., 1983a,b). For example, following complete loss of sensory input from the hand, the cortical hand representation ultimately represents somatosensory input from the face (Pons et al., 1991). Further complicating the matter, cortical stimulation itself alters the topography of the sensory cortex (Recanzone et al., 1992), shifting the receptive fields of neighboring cells towards that at the site of stimulation. Thus, neural plasticity must be considering in designing systems for artificial somatosensory feedback: in the absence of a stable cortical map, the target for biomimetic stimulation may itself be non-stationary.

Neural activity evoked by ICMS

Information is transmitted within and across neurons in the brain by chemical and electrical processes. To provide artificial sensory feedback, we need a way to interface with one of these processes and alter ongoing neural activity to encode useful information. One well-studied approach to alter neural activity is Intracortical Microstimulation (ICMS) which consists of delivering small, spatially-restricted electrical currents directly into the cortex. The basic unit of electrical stimulation is a biphasic pulse of electrical current (Figure 1.2b). In anesthetized rats, an ICMS pulse evokes a short (1-2 ms) burst of neural firing followed

by a longer (100 ms) inhibitory period (Butovas and Schwarz, 2003), which mediated in part by electrically coupled inhibitory networks (Butovas et al., 2006). At slightly higher stimulation intensities, the inhibitory period is followed by a rebound response—about 50 ms of excitation. If multiple stimulation pulses are delivered, the inhibitory/excitatory responses interact in a complicated way, up until around 40 Hz where each stimulus pulse evokes one or more spikes against a background of inhibition (Figure 2d in Butovas and Schwarz, 2003). From these results, one can speculate that at higher stimulation frequencies (40 Hz and above), ICMS can evoke a well-defined pattern of spiking in the nearby electrodes. The results are not deterministic, but rather occur with some probability that is a function of stimulation strength and distance from the stimulating electrode.

Cortical response to ICMS is further mediated by the state of the animal. Neural responses to ICMS evoked in awake, resting rats closely resemble those seen during anesthesia; however, during active behavior such as whisking, the duration of the inhibitory period is reduced (Venkatraman and Carmena, 2009). This difference in neural activation may be caused by interactions with ongoing neural activity. Indeed, Venkatraman and Carmena, 2009 show that a computational model of the thalamo-cortical loop accurately captures the transition between activity patterns observed in the two separate states.

Conscious perception of electrical stimulation

The following studies on ICMS detection and discrimination can help us optimize the parameters of stimulation used for artificial sensory feedback.

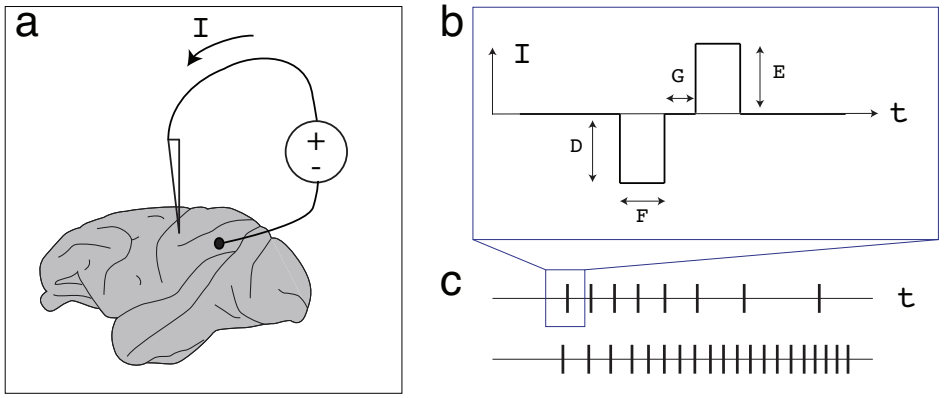

Figure 1.2: **Intracortical microstimulation. a)** A voltage source drives electrical current through a penetrating microelectrode (left). **b)**A typical ICMS pulse consists of a symmetric biphasic pair of current pulses. Each pulse of the pair is of fixed amplitude (D, E) and duration (F), and the two are separated by an interval of zero current amplitude (inter-pulse interval; G). **c)** A stimulus consists of one or more biphasic pulses delivered at variable frequencies (top vs. bottom).

ICMS detection

Humans and animals can "feel" electrical current delivered to certain parts of the surface and depths of the brain; immediately, if delivered to somatosensory (Penfield and Boldrey, 1937) or other primary sensory cortices (vision; Dobelle et al., 1974), but also, with training, to "higher" areas of cortex. This ability to detect an ICMS input is somewhat contingent on the area that is being stimulated. In visual areas, the amount of current required to evoke a sensation rises slightly from lower (V1) to higher (V3) visual areas (Murphey 2007). In somatosensory cortex, excitation of primary areas such as S1 are readily detectible, but training is required to detect stimulation of higher areas in parietal cortex (Doty, 1969).

Within a cortical area, there are several factors that affect ICMS detection. First, stimulation of the surface tends to have a higher threshold (requires more current input for detection) than intracortical stimulation (Doty, 1969). In fact, this threshold decreases from

the uppermost to the deepest layers of cortex (at least in monkey primary visual cortex and rat primary auditory cortex; Koivuniemi et al., 2011; Tehovnik and Slocum, 2009). At a fixed stimulation site, increasing current amplitude increases the probability that the ICMS signal is detected (explicitly shown in Berg et al., 2013; Fitzsimmons et al., 2007; O'Doherty et al., 2009; Tabot et al., 2013b; Zaaimi et al., 2013). Similarly, increasing stimulation frequency—the number of pulses delivered per second—decreases current threshold for detection (Butovas and Schwarz, 2007; Koivuniemi and Otto, 2012), although in rat auditory cortex this improvement saturates at around 80 Hz. With fixed frequency and amplitude, detecting ICMS becomes easier with increasing numbers of stimulation pulses and stimulating electrodes (Butovas and Schwarz, 2007; Tabot et al., 2013b; Zaaimi et al., 2013). Finally, the shape and time-course of the electrical current also influence detectability. In general, symmetric, cathode-leading biphasic pulses are easiest to detect (Koivuniemi and Otto, 2012), except in upper layers of cortex anode-leading pulses are more readily detectible (Tehovnik and Slocum, 2009).

ICMS discrimination

The spatiotemporal pattern of ICMS inputs can be manipulated to elicit dissimilar sensations (Figure 1.2). For example, ICMS inputs with different temporal and spatial patterns are discriminable from one another (Fitzsimmons et al., 2007; Otto et al., 2005; ODoherty et al., 2011; Talwar et al., 2002; but see Ghose and Maunsell, 2012). Specifically, stimulation duration (Fitzsimmons et al., 2007), amplitude (Johnson et al., 2013), and frequency

(Fitzsimmons et al., 2007; Johnson et al., 2013; London et al., 2008; ODoherty et al., 2011; Romo et al., 2000) are all parameters that are discriminable in an ICMS signal. Remarkably, animals can even integrate information about the environment or the state of their own bodies with ICMS inputs. Rats can distinguish if ICMS is delivered at different points along a whisk cycle (Venkatraman and Carmena, 2011), and the location in space of the maximum stimulation frequency of a continuous input (Thomson et al., 2013).

Human studies give us some insight into how this type of discrimination is possible. Stimulation of different locations in the brain evoke sensations at slightly different parts of the body (Heming et al., 2010; Penfield and Boldrey, 1937), though the location of the evoked sensation doesn't always match the receptive field of nearby cells (Heming et al., 2010). Interestingly, humans can distinguish between stimulation with different parameters even if the evoked sensation is not naturalistic. For instance, for a subject stimulated with ECOG electrodes over hand somatosensory cortex, the evoked sensation felt like "wind running down the hand," and this sensation felt "stronger" at a higher stimulation frequencies (100 Hz vs. 75 Hz). Increasing the current amplitude of the stimulus elicited a similar change in the evoked sensation—increasing strength with increasing current (Fridman et al., 2010; Johnson et al., 2013). Thus increasing either amplitude or stimulation frequency seems to have no effect on the quality of the sensation that is evoked, and changes only the perceived strength.

1.4 Candidate neural structures for artificial feedback

We next address the question of which brain area to target for artificial feedback. Consider the neural pathway involved in goal-directed reaching: sensory information first ascends from the periphery via the thalamus to modality-specific primary sensory areas. These, in turn, project to a number of multisensory cortical areas in the parietal lobe. Motor planning appears to occur across this parietal circuit and in the interconnected pre-motor and primary motor cortex in the frontal lobe. As stimulation in higher processing areas is likely to disrupt, rather than enhance, ongoing neural processing, the ideal target for somatosensory stimulation would seem to be one that lies upstream of the multisensory areas involved in movement planning. Furthermore, as most BMI users suffer from paralysis, we limit ourselves to the parts of the somatosensory system above the spinal cord: the thalamus and primary somatosensory cortex.

As the earlier stage of sensory processing, the thalamus is an attractive target for stimulation. The percepts evoked by thalamic stimulation correspond to relatively small and localized portions of the body surface (Heming et al., 2010), implying that each electrode on a stimulating array could be used as a distinct input channel. Additionally, the neural activity evoked by in cortex by stimulation in the thalamus may more closely resemble natural neural activity than that evoked by stimulation of cortex directly (Brockmeier et al., 2012; Choi et al., 2012)—facilitating biomimetic stimulation. Unfortunately, however, the thalamus is difficult to reach using currently available electrode arrays, so may not be the best target for stimulation.

Primary somatosensory cortex (S1) is still an early state of sensory processing, as neurons in S1 respond primarily to touch and proprioceptive information. In primates, S1 is composed of four cortical areas: 3a (mostly proprioceptive), 3b (mostly cutaneous), 1 (mixed responses, but more cutaneous), and 2 (mixed responses, but more proprioceptive; Krubitzer and Kaas, 1990). All of these subdivisions are good candidates for providing artificial somatic sensations, but areas 1 and 2 offer an important practical advantage—easier targeting with electrode arrays because of their superficial location on bank of the post-central gyrus.

S1 also projects directly to areas that are known to perform multisensory integration (A5, PMd) as well as to primary motor cortex, a connection that may play an important role in the ability of somatosensory feedback to elicit rapid corrective responses during movement execution (Omrani et al., 2012). For these and other reasons, most cortical approaches to providing artificial somatic sensation have targeted S1 (Berg et al., 2013; London et al., 2008; O'Doherty et al., 2009; ODoherty et al., 2011; Tabot et al., 2013a). For the work in this thesis, I have taken the same approach, using ICMS to deliver artificial proprioceptive information directly to primary somatosensory cortex.

1.5 Artificial proprioception in a non-BMI task.

ICMS has long been see as a potential tool for providing artificial feedback in neural prosthetics (London et al., 2011; Marzullo et al., 2010; O'Doherty et al., 2009; ODoherty et al., 2011; Santhanam et al., 2006; Venkatraman and Carmena, 2011), but so far efforts to validate its use have been limited, in part because electrical stimulation interferes with the neural

Parameter	Constraint
Current Amplitude (μA)	> detection threshold; can remain fixed
Stimulation Frequency (Hz)	> 100 Hz
Location	multiple electrodes, spatially distinct

Table 1.1: Constraints on stimulation parameters

recordings required to perform a BMI task. Fortunately, there are ongoing hardware- and software-based efforts to resolve these issues (Brown et al., 2008; Chu et al., 2013; Marzullo et al., 2010; Rolston et al., 2009; Venkatraman and Carmena, 2009; Zanos et al., 2011), and future iterations of neural recording and stimulating devices will free researchers of such constraints. In the meantime, we can still test the viability of an ICMS-based sensory feedback signal outside of the context of BMIs by carefully designing a behavioral task where artificial proprioception does not provide a redundant input (as it would for an intact animal who is not performing a BMI task). To circumvent this issue, we designed a behavioral task that required a monkey to reach to an invisible target (described in detail in Chapter 2). For the monkey to successfully complete this task, he needed to know the position of his hand relative to the target—an artificial proprioceptive signal we obligingly provided in the form of spatiotemporally patterned ICMS in S1. We constrained the parameters of the ICMS signal to maximize signal detectability and discriminability, as suggested by the discussion in the previous section (constraints summarized in Table 1.1). Animals are then trained to use the novel ICMS signal by extensive training on *correlated visual and ICMS signals*, which we hypothesized to be sufficient to drive learning and integration of the new input.

1.6 Dissertation summary

In the remainder of this dissertation we experimentally test the suitability of ICMS to encode a learning-based artificial sensory feedback signal. To do so, we first verify that ICMS can be used to encode a two-dimensional signal (hand position) and that this estimate can be naturalistically integrated with visual feedback (Chapter 2). We next examine how estimates of hand position are decoded from the ICMS signal and how the precision of the ICMS signal depends on the number of stimulating electrodes (Chapter 3). Lastly, we compare the plasticity of the ICMS signal to natural sensory inputs by asking the monkey to adapt to an imposed misalignment (Chapter 4). We find that a learning-based ICMS-encoded artificial feedback signal bears many similarities to natural sensory input—it is both integrated with vision and will adapt to an imposed perturbation—validating its use as an artificial feedback signal for neural prostheses and, further, suggesting that it could be used as model system in which to study sensorimotor plasticity and neural encoding.

Resulting publications

At the time this manuscript was completed, none of the chapters had yet been published, but all of the chapters (excluding Chapter 5) contain material that has been submitted for publication or is in preparation.

- Parts of Chapter 1 overlap with material that will be published as part of a book chapter (Dadarlat et al., 2014a).

- Chapter 2 is under review as a journal article (Dadarlat et al., 2014b) and the work it describes was awarded the 2013 Annual BCI Research Award.

- Chapters 3 and 4 are part of a manuscript that is in preparation (Dadarlat and Sabes, 2015).

Chapter 2

Encoding artificial proprioception

2.1 Introduction

Humans plan and execute movements under the guidance of both vision and proprioception (Sober and Sabes, 2005, 2003). In particular, maximally precise movements are achieved by combining estimates of limb or target position from multiple sensory modalities, weighing each by its relative reliability (Ernst and Banks, 2002; McGuire and Sabes, 2009; Morgan et al., 2008; van Beers et al., 1999). Furthermore, in the absence of proprioception, even simple multi-joint movements become uncoordinated (Sainburg et al., 1995, 1993). Therefore, we should not expect current brain-machine interfaces (BMIs), which rely on visual feedback alone, to achieve the fluidity and precision of natural movement. It follows that a critical next step for neural prosthetics is the development of artificial proprioception. As a demonstration of the potential value of somatosensory feedback, it has been shown that including

natural kinesthetic feedback improves BMI control in intact monkeys to near-natural levels (Suminski et al., 2010). The ideal artificial proprioceptive signal would be able to fill the same roles that proprioception plays in natural motor control: providing sufficient information to allow competent performance in the absence of other sensory inputs, and permitting multisensory integration with vision to reduce movement variability when both signals are available. Here we present a proof-of-concept study showing that both of these goals can be achieved using multichannel intracortical microstimulation (ICMS). Most efforts to develop artificial sensory signals have taken a biomimetic approach: trying to recreate the patterns of neural activity that underlie natural somatosensation (Choi et al., 2012; Daly et al., 2012; Fagg et al., 2007; Tabot et al., 2013b; Weber et al., 2012). We propose here a complementary approach, which focuses not on reproducing natural patterns of activity, but instead on taking advantage of the natural mechanisms of sensorimotor learning and plasticity. In particular, the process of multisensory integration, where multiple sensory signal are combined to improve the precision of sensory estimates, is learned from cross-modal experience during development (citealpHeld1963,Xu2012) and relies on a continuous process of adaptive recalibration even in adult humans and monkeys (Burge et al., 2008; Simani et al., 2007; Zaidel et al., 2011). Recent theoretical work from our lab suggests that multisensory integration can be learned with experience through a simple Hebbian-like learning rule (Makin et al., 2013). In this model, successful integration of two sensory signals depends not so much on choosing the right patterns of neural activity to encode spatial information, but rather on the presence of spatiotemporal correlations between input signals, which allow downstream

neurons to learn the common underlying cause, e.g. hand position. Following these theoretical principles, we hypothesized that spatiotemporal correlations between a visual signal and novel artificial signal in a behavioral context would be sufficient for a monkey to learn to interpret and integrate the new modality. We tested this hypothesis by delivering real-time, artificial sensory feedback to monkeys via non-biomimetic patterns of ICMS across multiple electrodes in primary somatosensory cortex (S1). The monkeys ultimately learned to extract the task-relevant information from this signal and to integrate this information with natural sensory feedback.

2.2 Methods

Subjects and Implants

Subjects and Implants All animal procedures were performed in accordance with the National Research Councils Guide for the Care and Use of Laboratory Animals and were approved by the UCSF Institutional Animal Care and Use Committee. Two adult male rhesus macaque monkeys (Macaca mulatta) participated in this study. Each was chronically implanted with a 96-channel silicon microelectrode array coated with Iridium Oxide (Blackrock Microsystems, Salt Lake City, UT) over their left primary somatosensory cortices (Brodmann Areas 1, 2; S1). The cells on monkey Fs array had receptive fields spanning the shoulder, back, side of the head, ear and occiput (bottom panels, Figure 2.1) whereas for monkey D most receptive fields spanned the arm and shoulder (top panels, Figure 2.1).

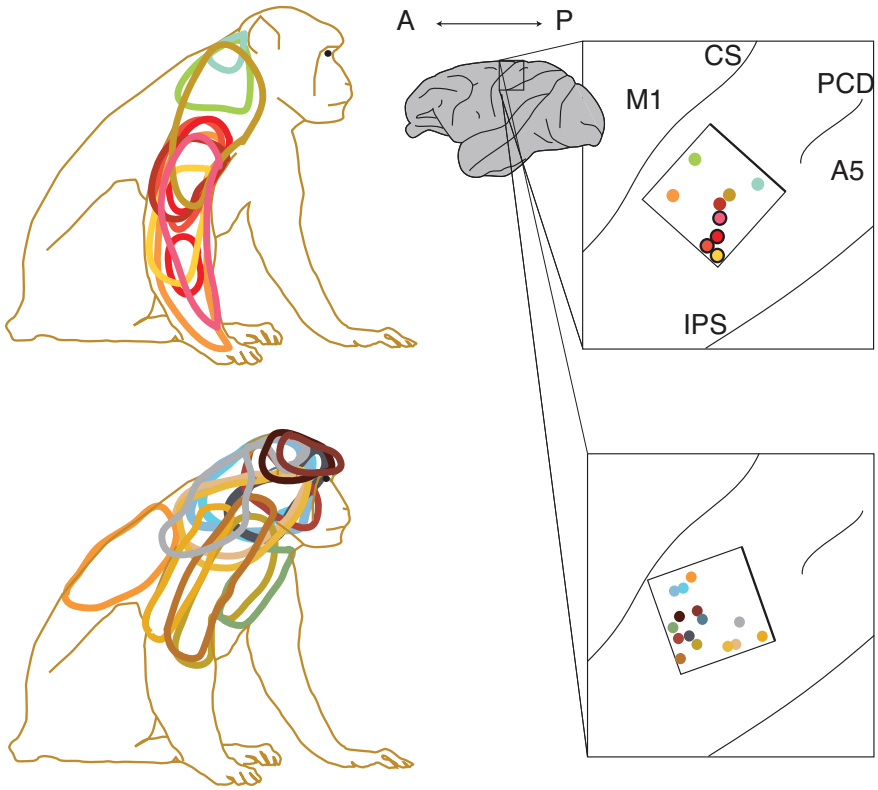

Figure 2.1: **Physiological properties of stimulated somatosensory cortex.** Location of electrode arrays within S1 (right) and example neuronal receptive fields (left) for Monkey D (top) and Monkey F (bottom). Colored circles (right) indicate array locations corresponding to the matching colored receptive fields (left). Neurons responded mainly to light touch; circles with dark borders correspond to cells that responded to limb movements (active and passive).

Behavioral Task

The animals were trained to perform reaches in the horizontal plane to an unseen target in a two-dimensional virtual reality environment, where a mirror and an opaque barrier prevented direct vision of the arm (Figure 2.2). The mirror reflected visual input from a projector, so that the visual cues appeared in the horizontal plane of the reaching hand. Fingertip position was monitored with an electromagnetic position sensor (Polhemus Liberty, Colchester, VT) at 240 Hz.

Each trial consisted of four epochs (Figure 2.3a). i) The monkeys moved the middle

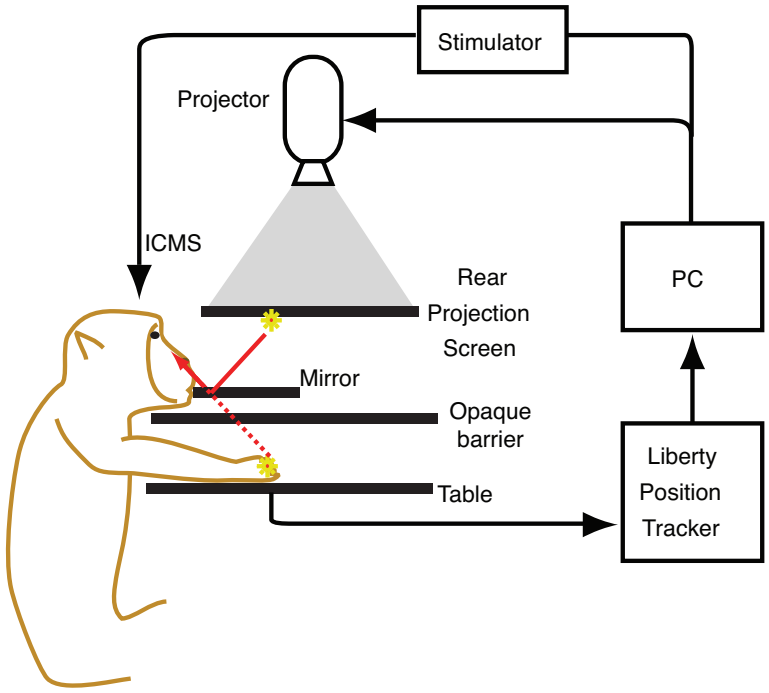

Figure 2.2: **Virtual reality environment.** Animals sit in a virtual reality environment without direct view of the arm. A mirror reflects images on rear-project screen and is adjusted so that visual cues appear in the horizontal plane of the reaching hand. Hand position was tracked electromagnetically (Polhemus Liberty, Colchester, VT), and feedback about the position of the hand, relative to the target, was delivered via a random-dot visual flow field (inset) or via patterned ICMS.

fingertip of their right hand to a fixed start position, located in the center of the screen and indicated by a circular visual target (10 mm radius). ii) After a brief delay (0.25 and 0.5 s for Monkeys D and F, respectively), the target cue was initiated, indicating the movement vector between the monkeys current finger position and the center of the unseen reach target (12 mm radius). Targets were selected uniformly from an annulus centered on the start target with an inner radius of 40 mm and an outer radius of 80 mm. The movement vector cue was provided in the form of a dot-field (VIS), multichannel ICMS (ICMS), or both (VIS+ICMS). The monkeys were required to hold their position during this instructed-delay interval (0.2-

0.7 s and 1-1.5 s, monkeys D and F, respectively). iii) After a go cue (750 Hz tone, 0.5 s) the monkeys made a reach under the guidance of continuously updating VIS, ICMS or VIS+ICMS feedback. iv) After acquiring the target and holding for 400 ms (monkey D) or 500 ms (monkey F), the monkeys received a liquid reward. Trials were terminated without reward if the monkeys moved too early during any of the delay intervals or if they failed to reach the target before a timeout (10 s). Different task parameters were selected for each animal to minimize the number of failed trials (e.g., aborted hold at start or target), and therefore reflect the animals idiosyncratic behavioral tendencies.

Visual feedback For vision, the movement vector was encoded using a random moving-dot flow-field (dot-field) consisting of approximately 600 dots over the visual display (roughly 53 cm x 33 cm, in the reaching plane). Each dot was initialized to a random location on screen, and had a lifetime of 4 seconds (phases randomized), after which it reappeared at a new random location. Each dot in the dot-field moved at the same angle as the movement vector and at a speed proportional to the length of the movement vector, but could not exceed a maximum of 50 cm/s for Monkey D and 40 cm/s for Monkey F. A percentage of the dots moved coherently together in the direction of the continuously updating movement vector. The remaining dots moved in random directions, selected independently and uniformly from the circle. The percentage of dots moving coherently—the dot field coherence—was parametrically varied in order to manipulate the precision of the visual feedback.

ICMS Intracortical microstimulation consisted of biphasic, charge-balanced pulse trains delivered asynchronously to each of eight electrodes in the array. The pulse trains were

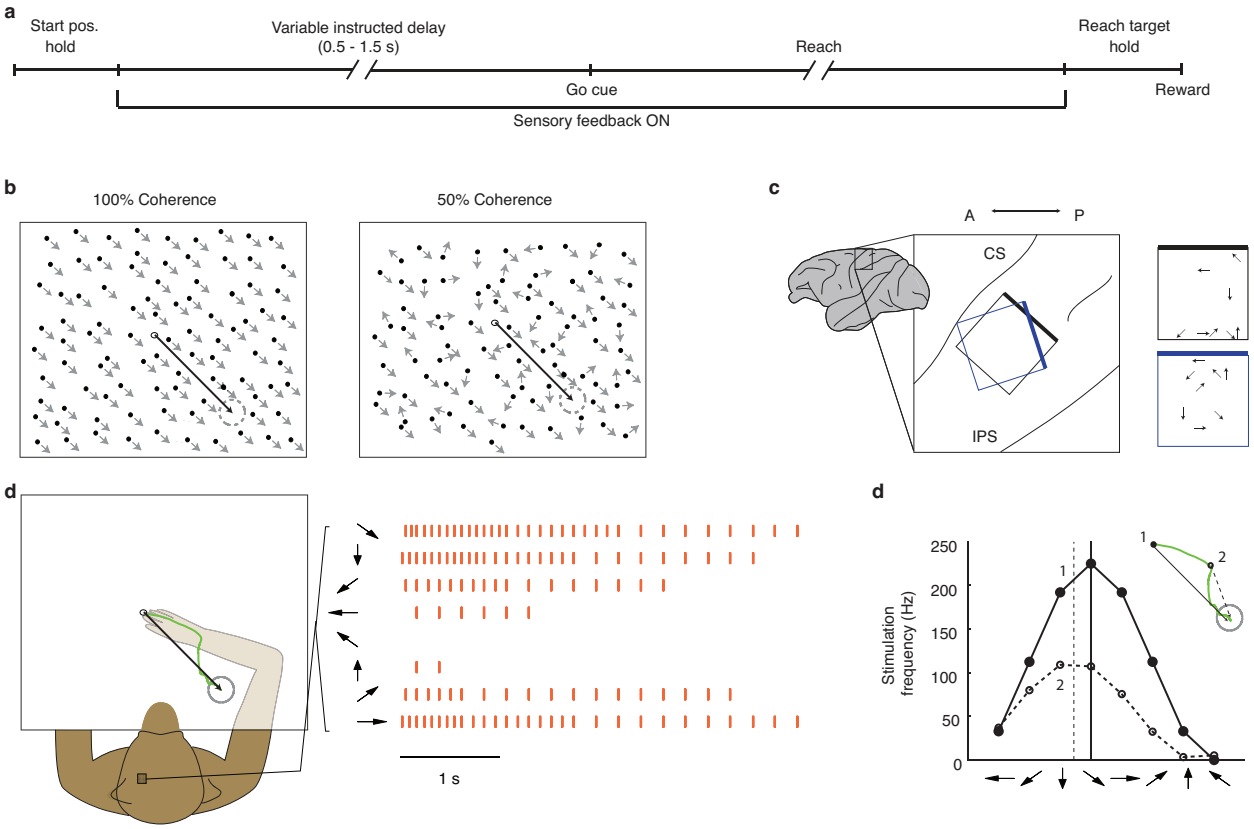

Figure 2.3: **Behavioral task and sensory feedback.** a) Timeline of a behavioral trial. b) Visual feedback of the instantaneous movement vector (black arrow) takes the form of a random moving-dot flow-field (dot-field). c) Implantation site of stimulating electrode arrays for monkeys D (black) and F (blue). CS-central sulcus; IPS-inferior parietal sulcus. Right: the assigned PD of each stimulating electrode is overlaid on its location within the array. d) An example ICMS trial showing the movement vector at the beginning of the reach (black arrow) and the monkeys subsequent movement path (green). At right: ICMS patterns delivered during the trial; each row represents the time-varying stimulation pattern of the electrode with the preferred direction (PD) indicated at left (black arrow). Orange tick marks denote biphasic stimulation pulses, which are shown subsampled for clarity. e) Inset: the instantaneous movement vectors encoded at two time-points during the reach are shown as solid and dashed black arrows. Below, the pattern of stimulation encoding each movement vector is shown across electrodes; arrowheads indicate the PD of each electrode.

cathode-leading and symmetric, with 200 μs/phase and a 250 μs phase separation. The pulse amplitudes varied across electrodes, depending on perceptual threshold (see below) and ranged between 34-60 μA for Monkey D, and 30-80 μA for Monkey F.

ICMS Detection. A preliminary two-alternative forced choice task was used to determine the threshold pulse amplitudes at which the animals could detect ICMS on a given electrode. The monkeys first moved to a fixed start position near the midline (as in main task above) and maintained that position for 0.5 s. Next, there was a 0.5 s instructed delay period during which two reach targets were displayed, to the right and left of midline. The presence of an ICMS pulse train (100 Hz, 0.5 s) cued the animal to reach left; its absence cued the rightward reach. Animals were initially trained on this task using multi-electrode stimulation, and the task was then used to identify electrodes on which ICMS was detectable. Eight such electrodes were identified for Monkey F. For Monkey D, seven such electrodes were identified; the final electrode could not be detected when stimulated alone with amplitudes of less than 60 μA.

Movement Vector Encoding For ICMS, the movement vector was encoded in the spatial and temporal patterns of stimulation across the array (Figure 2.3). Movement vector direction was encoded by the relative stimulation pulse rates across the electrodes. First, each of the eight electrodes was arbitrarily assigned one of eight preferred directions (PD), equally spaced around the circle. Then, the stimulation pulse rate f_i of electrode i was calculated as a function of angle between the direction movement vector, θ, and the electrodes assigned

PD,

$$f_i = \frac{f_0}{2}(1 + cos(PD_i - \theta)) \tag{2.1}$$

The frequency scaling factor, f_0, linearly encoded the movement vector distance, d, within the range of 100-300 Hz:

$$f_0 = \frac{d}{d_{max}} f_{max} \tag{2.2}$$

where d_max was 10 cm for Monkey F and 13.6 cm for Monkey D. The values of d and θ were continuously updated during the reach to provide online feedback (Figure 2.3). where d_{max} was 11.5 cm for Monkey F and 13.6 cm for Monkey D. The values of d and θ were continuously updated during the reach to provide online feedback (Figure 2.3d,e).

Experimental Design

Behavioral sessions were divided into training blocks and testing blocks. The details of those blocks changed during the course of training, as described below.

Behavioral Training. Monkeys first learned to perform the behavioral task using vision alone, with 100% coherence. Next, we began training with only VIS+ICMS trials. A dot-field coherence of 100% was used initially, and that value was slowly decreased across sessions to encourage the animals to use to the ICMS signal. Coherence values were lowered to 25% for Monkey D and 20% for Monkey F, depending on the animals performance level in the VIS conditions. This training regime was employed for approximately 20,000 training trials with Monkey D and 40,000 training trials with Monkey F, at which point the animals showed clear evidence of sensory integration of the VIS and ICMS signals—improved performance on VIS+ICMS trials compared to VIS trials, as evaluated on testing blocks (see below). We

then altered the training regime to include 33% ICMS-only trials and 67% VIS+ICMS trials once the animals were able to perform ICMS-only trials in the testing blocks (see below).

Behavioral Testing. In between blocks of training, approximately every 500-1000 training trials, the animals performed a testing block to quantify performance across all feedback conditions. By the end of the experiment, a total of 11 feedback conditions were used: VIS, VIS+ICMS, and ICMS, with dot-field coherences of [0, 15, 25, 50, 100]% for monkey D and [0, 10, 15, 25, 50]% for monkey F, with the difference between animals reflecting individual performance levels in the VIS condition. At the beginning of the experiment, animals had not been exposed to lower visual coherences and could not perform the task at those coherences. Lower visual coherences for testing were introduced gradually during the course of the experiment as performance improved (see Figures 2.4-2.5). For all testing blocks, all conditions were randomized across trials.

As noted above, testing sessions revealed evidence of sensory integration of the VIS and ICMS signals prior to the animals performing the ICMS-only task. Yet at this stage, the animals still failed to initiate movement on ICMS-only trials, suggesting that they did not generalize the task instructions to trials with no visual cues. We therefore temporarily modified the ICMS-only testing protocol to help them generalize. First, we paired ICMS with a visible target circle during ICMS-only testing trials. The brightness of the target circle was gradually decreased, until it was finally removed entirely, and the monkeys were reaching with ICMS alone (thin vertical green line in Figures 2.4-2.5). Next, the radius of the reach target for ICMS-only trials was temporarily increased from 12 mm to 36 mm—for both

testing and training trials—to avoid discouraging the animal from performing these trials. The radius of the target was then decreased across training sessions until the monkey was completing reaches to a standard 12 mm radius target (thick vertical green lines in Figures 2.4-2.5).

Data Analysis

Behavioral Performance Measures. We quantified the animals ability to use the various sensory cues with the following performance measures (details of each are given below): i) percent correct trials, ii) number of movement sub-segments, iii) normalized movement time, iv) normalized path length, v) mean and variance of initial angle of the movement. The first four metrics assessed animals use of the sensory cues throughout the trial, including movement planning during the instructed delay period and online movement control. The statistics of the initial angle assess only movement planning. The performance summaries in the summary results figures were computed from the last seven testing sessions available for each monkey.

i) Percent correct trials: This is the number of trials in which a monkey acquired a target and received a reward, compared to the total number trials in which a reach movement was initiated. We exclude in this analysis trials where the monkey made errors such as starting a reach before the go cue, not initiating a trial, and swiping through the reach target.

ii) Number of movement sub-segments: This metric quantifies the number of discrete sub-movements in a trial. Starting with the model assumption that sub-movements have bell-shaped velocity profiles (Novak et al., 2002) we identified sub-movements by threshold

crossings of the tangential velocity plot of a trajectory, with a threshold of 20% of the maximum velocity on a given trial.

iii) Normalized movement time: Since maximum movement velocity was largely independent of movement distance (data not shown), targets that were farther away took longer to reach. Therefore, we normalized the movement time by the distance from the starting point to the target.

iv) Normalized path length: Similarly, we normalized the integrated path length by the distance from the starting point to the target.

v) Mean and variance of the initial angle. For each monkey and feedback condition, we first computed a smoothed estimate of the mean initial angle as a function of target angle (robust locally weighted scatterplot smoothing, using the MATLAB smooth function, with a window of 40 data points). The initial angle variance was computed about this mean. Standard errors for the mean variance were estimated via bootstrapping (Efron and Tibshirani, 1993).

Quantifying use of direction and distance from the ICMS signal. A typical reach consists of a long initial movement segment followed by one or more shorter, corrective submovements. The distance and direction of this initial reach can be taken to reflect the monkeys estimate of target distance and direction, as decoded from the sensory information available during the instructed delay period.

i) Direction estimation: We assessed the monkeys ability to estimate target direction from ICMS by regressing initial movement angle against target angle for ICMS-only trials. The

initial movement angle was measured using the first movement sub-segment, as described above. This assay ignores target-dependent biases in initial direction (see above), and is therefore a conservative estimate of the animals ability to decide target direction.

ii) Distance estimation: We assessed the monkeys ability to estimate target distance from ICMS by regressing initial movement distance against target distance for ICMS-only trials. On a subset of trials, however, the initial movement deviated from the norm: animals sometimes made a small initial reach that was followed by several larger corrections. On these trials the distance of the initial reach segment was uncorrelated with movement distance. Therefore, for this analysis we exclude trials for which the first movement segment was not longest segment. This occurred in 30.2% of the trials for Monkey D and 12.2% of the trials for Monkey F.

Model prediction for initial angle variance. Under the model of minimum variance sensory integration (8), we can predict the sensory variability in the bimodal condition from the variability for each unimodal condition. We focus on variability in the animals estimate of the target angle based on the sensory cues during the instructed delay period. Unimodal variances were computed from the variability in initial movement direction, σ^2_{VIS}, and σ^2_{ICMS}, calculated as described for performance measure (v) above. Next, for each coherence level, we predicted the bimodal variance under two limiting assumptions. First, we used the raw initial angle variance directly, which implicitly assumes that all movement variability derives from sensory variability:

$$\sigma^2_{VIS+ICMS} = \frac{\sigma^2_{VIS}\sigma^2_{ICMS}}{\sigma^2_{VIS} + \sigma^2_{ICMS}} \tag{2.3}$$

Second, we assumed that the movement variability was the largest possible value that would still be consistent with the data—this is the smallest initial angle variance across conditions, which we denote with σ^2_{MIN}. Under this model, the bimodal initial angle variance is:

$$\sigma^2_{VIS+ICMS} = \frac{(\sigma^2_{VIS} - \sigma^2_{MIN})(\sigma^2_{ICMS} - \sigma^2_{MIN})}{\sigma^2_{VIS} + \sigma^2_{ICMS} - 2\sigma^2_{MIN}} + \sigma^2_{MIN} \tag{2.4}$$

Model predictions for mean initial angle. The plots of mean initial angle for monkey D show a clear dependence on the feedback type. If we suppose that these differences reflect biases in the sensory estimates of target direction, then the minimum variance model can be used to predict, for each coherence value, the mean initial angle in the VIS+ICMS trials from those measured in the unimodal trials:

$$\bar{\theta}_{VIS+ICMS} = \frac{\sigma^{-2}_{VIS}\bar{\theta}_{VIS} + \sigma^{-2}_{ICMS}\bar{\theta}_{ICMS}}{\sigma^{-2}_{VIS} + \sigma^{-2}_{ICMS}} \tag{2.5}$$

Equation 2.5 can be summarized by the predicted visual cue weighting for each coherence,

$$w_{VIS} = \frac{\sigma^{-2}_{VIS}}{\sigma^{-2}_{VIS} + \sigma^{-2}_{ICMS}} \tag{2.6}$$

which depends only on the unimodal initial angle variances.

We compared these model predictions to empirical values of the visual cue weighting,estimated from the mean initial angles. First, we divided the workspace into octants, and for each octant and feedback condition, we computed the mean difference between the initial angle andthe target angle, we which denote here as $\bar{\delta}_{i,x}$ for octant i and condition x. For each octant and coherence level, we then estimated the visual cue weighting as

$$w_{i,VIS} = \frac{\delta_{i,VIS+ICMS} - \delta_{i,ICMS}}{\delta_{i,VIS} - \delta_{i,ICMS}} \tag{2.7}$$

The standard error $s_{i,VIS}$ of each $w_{i,VIS}$ was estimated from the standard errors of the component means, $\bar{\delta}_{i,x}$, by propagation of errors. Finally, for each coherence level we computed the mean visual cue weighting across octants, with each octant weighted by its standard error:

$$w_{VIS} = \frac{\sum_{i=1}^{8} s_{i,VIS}^{-2} w_{i,VIS}}{\sum_{i=1}^{8} s_{i,VIS}^{-2}} \tag{2.8}$$

2.3 Results

Behavioral task and feedback signals Two rhesus macaques were trained to make instructed-delay center-out reaches to invisible targets (Figure 2.3) in a virtual reality environment (Figure 2.1) guided by feedback that represented the vector (distance and direction) from the middle fingertip to the reach target (black arrows, Figure 2.3b,d). This movement vector was not explicitly shown; instead, it was encoded by one of three feedback types: a visual signal (VIS), a signal delivered through patterned multi-channel ICMS pulse trains (ICMS), or a combination of these two signals (VIS+ICMS). This task was chosen to best test whether the ICMS signal can provide position information that can both be integrated with vision and can replace it. By using natural movement, we obtain the most direct and precise estimates of how well the ICMS signal encodes sensory information about the limb (e.g., not confounded by additional performance noise due to imperfect BMI control). However, natural movement leaves natural proprioception intact, which would make an ICMS signal encoding absolute limb position redundant. By encoding the relative positions of the limb and target, the VIS and ICMS signals provide a feedback variable that is both required to

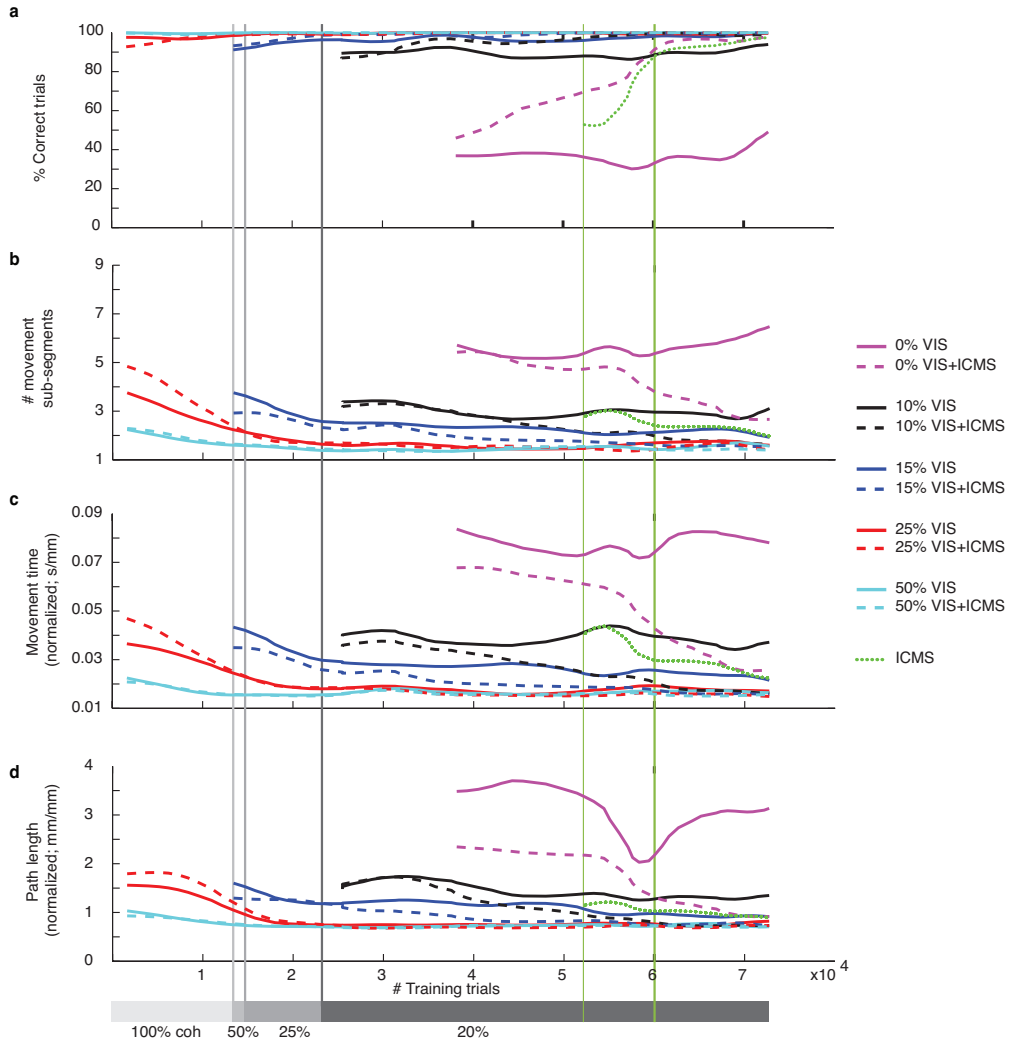

Figure 2.4: **Evolution of performance over training (Monkey D).** Behavioral performance measures are shown as a function of the cumulative number of VIS+ICMS trials performed (training and testing). The data, collected during testing sessions, were smoothed for clarity (Gaussian window with standard deviation of 2.8 training sessions, translating to approximately 2,500 training trials for Monkey D and 2,800 training trials for Monkey F). The visual coherence on training trials was decreased across training sessions (indicated by gray bars at the bottom of the figure and vertical gray lines at the transitions). The left, thin green line denotes the onset of ICMS-only trials, where target sizes were temporarily larger than in the other trial conditions; the right, thick green line denotes the beginning of ICMS-trials with targets of standard size. a) percent correct trials; b) number of movement segments measured online error corrections; c) movement time for the trial is normalized by the initial distance to the reach target; d) path length, normalized as in c.

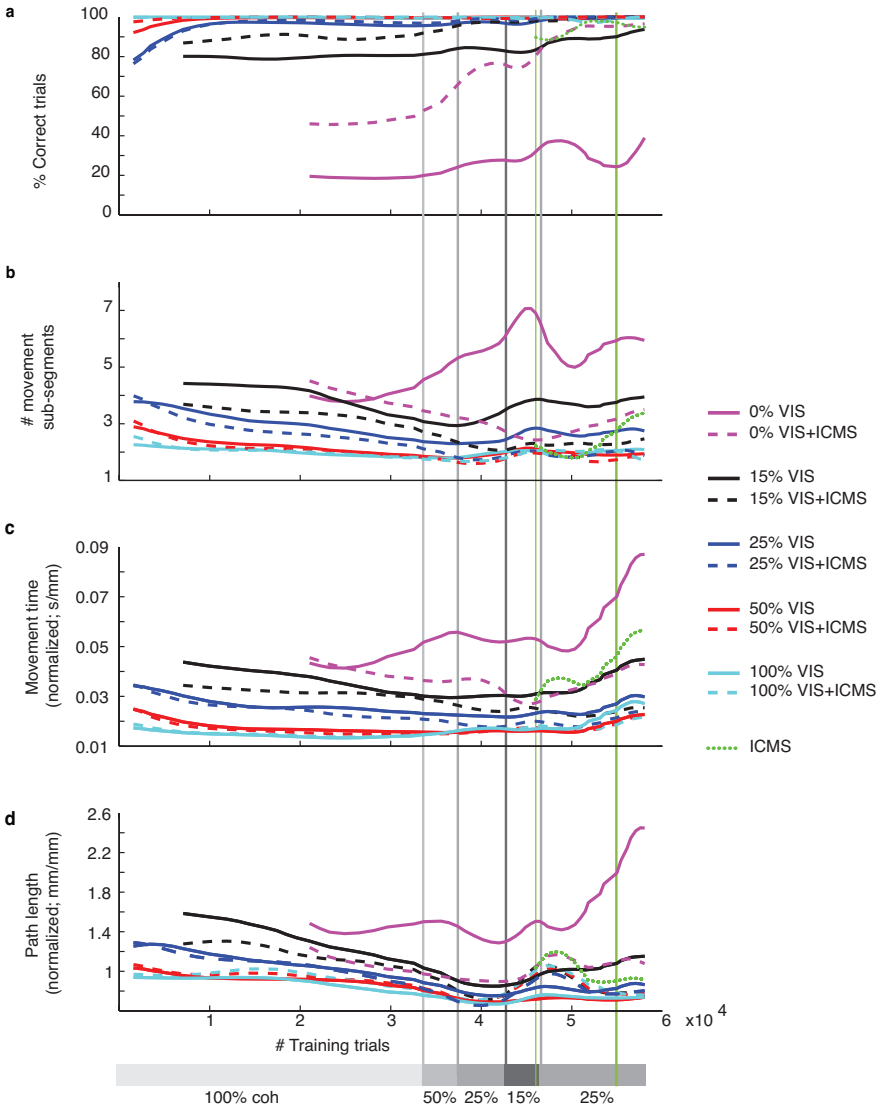

Figure 2.5: **Evolution of performance over training (Monkey F).** Same conventions as Figure 2.4.

complete the task and that is not available from natural sensory signals.

Performance with visual feedback. Monkeys first learned to perform the task with VIS feedback alone. We quantified performance with three behavioral metrics designed to capture how well the animals made use of a sensory signal during reach planning and execution. Performance on VIS-only trials increased monotonically with increasing dot-field coherence

for all behavioral metrics (purple lines, Figure 2.6), demonstrating that differences in performance reflect the precision of the sensory cue.

Learning ICMS After the monkeys could perform the reaching task using visual feedback, we tested our hypothesis that spatiotemporal correlations between vision and ICMS could drive integration of the new sensory modality. We did so by exposing the monkeys to paired, correlated VIS+ICMS feedback signals both during the instructed delay period (as static information) and throughout the reach (dynamically updated feedback). The visual signal was first set to 100% coherence, but was gradually reduced across training blocks to increase the relative value of the ICMS signal (ultimately settling at 20% for Monkey F and 25% for Monkey D). Under this training regime, the animals learned to integrate the two sensory signals, i.e. the addition of ICMS improved performance (see below). Animals needed more explicit instruction to learn to initiate movement on ICMS-only trials (see Methods). Once that was accomplished, the training regime changed to include ICMS-only trials (33%), a pragmatic choice intended to speed learning. We periodically assessed learning (approximately every 500-1000 training trials) by including testing blocks: trials of VIS-only, VIS+ICMS, and ICMS-only trials where the visual dot-field coherence could take a range of values ([0,10,15,25,50] for Monkey F; [0,15,25,50,100] for Monkey D).

Substitution of vision by ICMS We analyzed the data from testing blocks to determine how well the animals could interpret the ICMS signal, using it in place of vision to perform accurate reaches. Once the monkeys began making reaches on ICMS-only trials, they became increasingly proficient across training sessions (Figure 2.4-2.5), and performance on ICMS

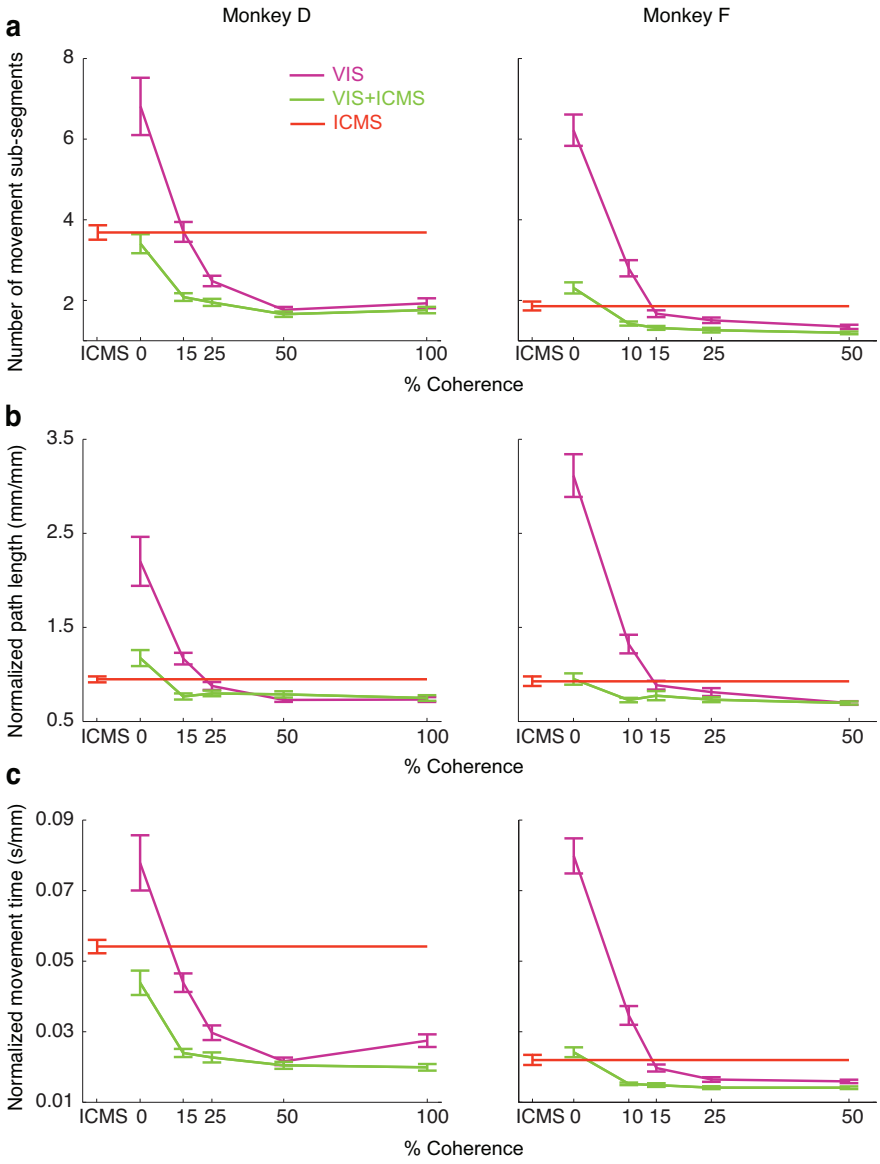

Figure 2.6: **Comparison of task performance across sensory feedback conditions.** Behavioral performance measures are averaged across the last seven testing sessions for each monkey, shown for each sensory feedback type and as a function of visual coherence (for VIS and VIS+ICMS trials). Error bars denote bootstrapped standard error of the mean. The ICMS data points, which are independent of visual coherences, are extended across the plot to aid visual comparison. a) number of movement sub-segments; b) movement path length, normalized by the initial distance to the reach target; c) movement time, normalized as in b.

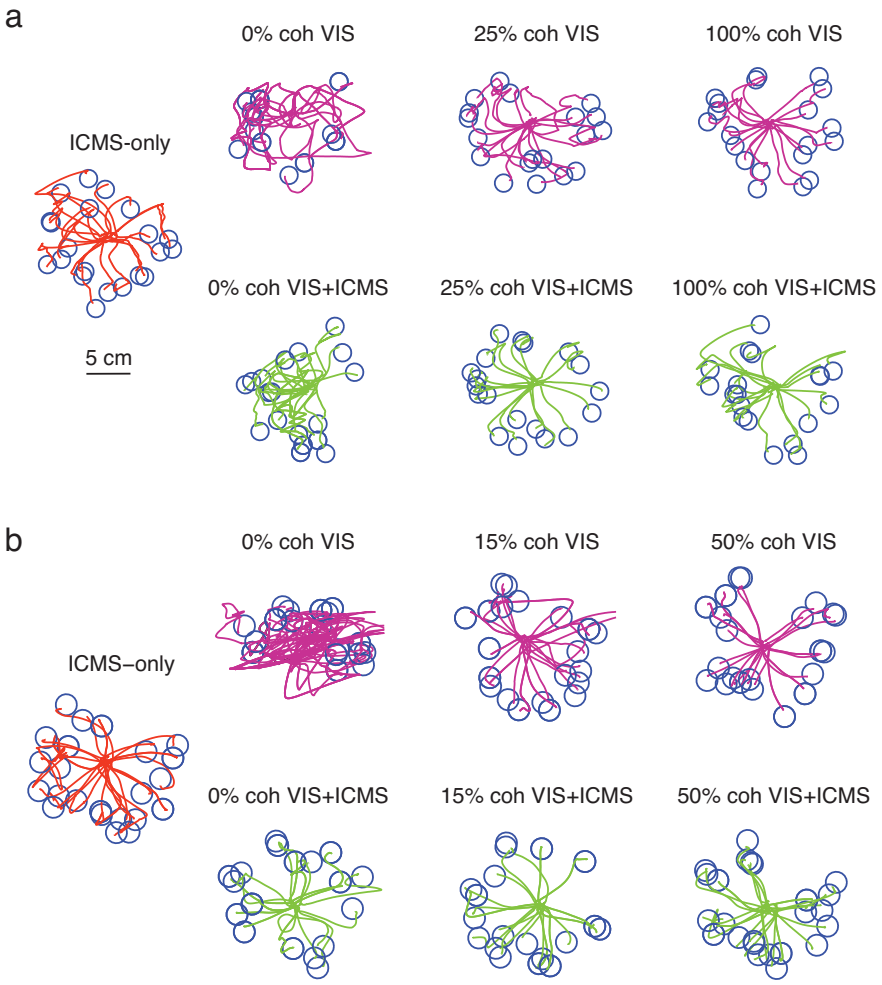

Figure 2.7: **Sample movement paths.** Sample movement paths from randomly selected successful trials for Monkeys D (a) and F (b) for seven feedback conditions. Each reach begins at the fixed central starting point and ends within the unseen reach target (here depicted in blue).

trials was ultimately comparable to performance with low-to-mid visual coherences (15-25% for Monkey D, 15% for Monkey F; orange symbols, Figure 2.6). A more qualitative impression of the performance comparisons can be obtained from sample movement paths for various feedback conditions (Figure 2.7).

We next quantified how well the monkeys made use of the distance and direction information encoded in the ICMS signal. To do this, we analyzed the distance and direction of the

initial movement segment of the reach, which reflects the animals estimates of the required movement vectors derived from sensory feedback during the instructed delay period (Figure 2.8). Monkey F was highly adept at estimating target angle (ICMS, regression $R^2 = 0.948$; Figure 2.8), performing as well with ICMS as with the highest visual coherences (50% VIS, $R^2 = 0.945$; Figure 2.9). Although demonstrating some idiosyncratic biases (see below), monkey D was also able to derive good estimates of target direction from ICMS (ICMS, $R^2 = 0.900$; Figure 2.8a; 100% VIS, $R^2 = 0.957$; Figure 2.9a). Both monkeys were somewhat worse at estimating distance than direction from ICMS.

Due to differences in performance across the workspace, we analyze distance estimation separately for the two half-planes: the more proximal workspace, with target angles [$-\pi$, 0], and the more distal workspace, with target angles [0, π]. For monkey F, distance estimates were equally good across the workspace (gray symbols: [$-\pi$, 0], $R^2 = 0.432$; orange symbols: [0, π], $R^2 = 0.473$; Figure 2.8d) and largely fall within one target radius of correct distance (but not necessarily in the target on each trial, due to directional error), although these values are lower than those the animal achieved with high-coherence VIS feedback ([$-\pi$, 0], $R^2 = 0.716$; [0, π], $R^2 = 0.751$; Figure 2.9d).

Monkey D could accurately estimate distance in the distal half of the workspace (orange symbols: [0, π], $R^2 = 0.494$; Figure 2.8c), but was less able to do so in the proximal half (gray symbols: [$-\pi$, 0], $R^2 = 0.108$; Figure 2.8c). Still, these values are comparable to those the animal achieved with the highest-coherence VIS feedback ([0, π], $R^2 = 0.365$; [$-\pi$, 0], $R^2 = 0.176$; Figure 2.9c). In summary, the task performance observed in the ICMS-only

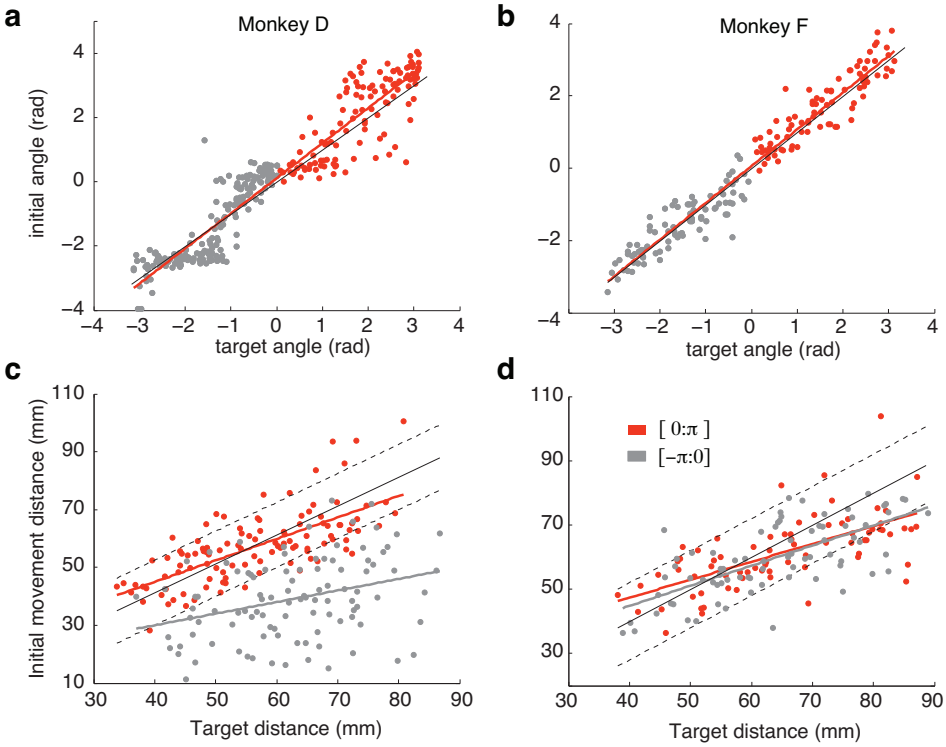

Figure 2.8: **Monkeys estimate both target distance and direction from ICMS.** The panels show correlations, for both direction and distance, between animals initial movements and the target location. Black solid lines are unity and the thick colored lines represent linear fits between the movement and target variable. Fits were performed separately for the distal (orange points, target angle [0:π]) and proximal (gray points, target angle [-π:0]) halves of the workspace. (a,b) Initial movement angle versus target angle for monkey D and F, respectively. (c,d) Initial movement distance versus target distance for monkey D and F, respectively. Region within the dashed black lines falls within the diameter of the target.

condition is driven by the animals ability to decode both distance and direction information from the ICMS signal.

Augmentation of vision by ICMS In addition to serving as the training condition, VIS+ICMS trials provide a test of the animals ability to improve performance by combining information from the two sensory cues. This ability emerged during training, with performance on VIS+ICMS trials becoming progressively better than for VIS trials, even before the animals

could complete reaches with ICMS alone (Figure2.8), supporting the idea that multisensory integration drives ICMS learning (Makin et al., 2013). Moreover, the hallmarks of multisensory integration were observed in the asymptotic performance on VIS+ICMS trials, after learning was complete (green data, Figure 2.6). At intermediate dot-field coherences (10-25%), where performance between the two unimodal cues was similar, VIS+ICMS reaches were significantly better, i.e., faster and straighter. In contrast, at high (50-100%) and low (0%) dot-field coherences, behavior in the bimodal condition approximated that observed with the more reliable of the unimodal cues.

Minimum-variance integration of vision and ICMS We next ask whether the visual and ICMS cues were integrated in an optimal, i.e. minimum variance manner, as is the case for natural visual and somatosensory signals (Ernst and Banks, 2002; van Beers et al., 1999). The answer comes from an analysis of the statistics of the initial reach directions, which give the most direct readout of the animals estimate of target direction following the instructed delay period. The minimum variance model makes specific predictions about both the variance and bias of this estimate, and we consider each in turn. We first consider how the variance of the initial reach angle depends on feedback condition (Figure 2.10a). For VIS trials, the initial angle variance increased dramatically with decreases in coherence, as expected if the variance reflects the residual uncertainty about cue direction after the instructed delay. Variability in the ICMS trial was comparable to VIS trials at 15-25% visual coherence, consistent with the other movement metrics above. From these unimodal variances, we can determine what the initial angle variance should be for the VIS+ICMS condition,

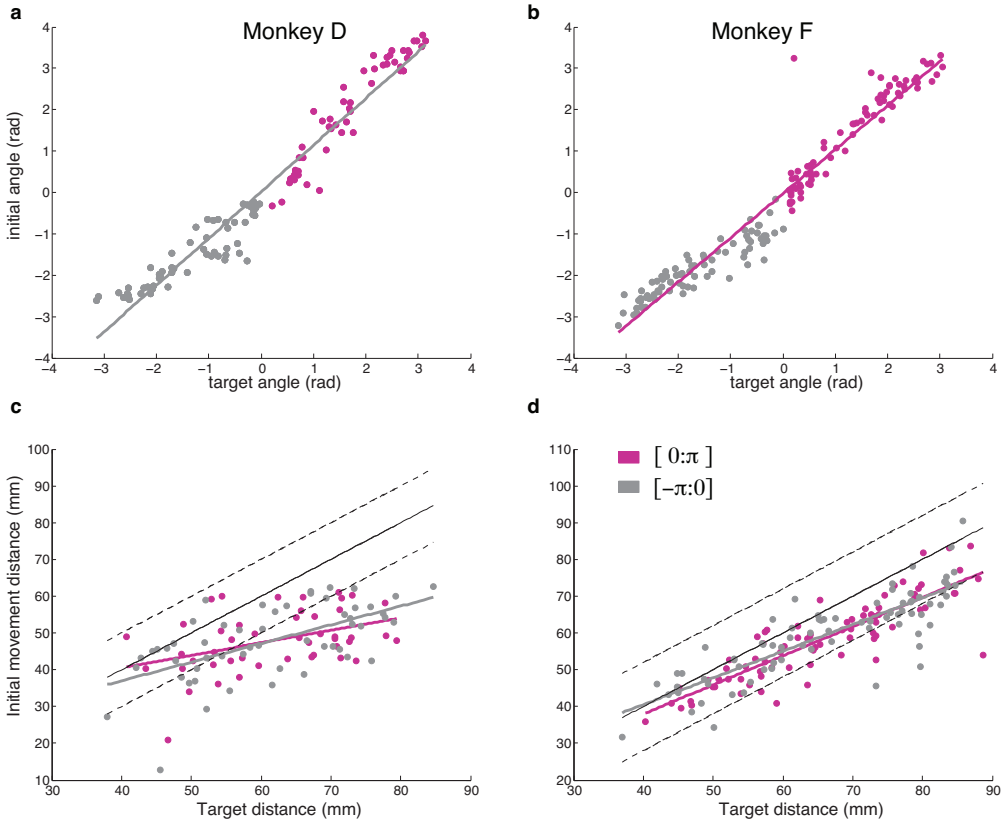

Figure 2.9: **Monkeys estimate both target distance and direction from high coherence VIS cue (50% for Monkey F, 100% coherence for Monkey D.** The panels show correlations, for both direction and distance, between animals initial movements and the target location. Black solid lines are unity and the thick colored lines represent linear fits between the movement and target variable. Fits were performed separately for the distal (orange points, target angle [0:π]) and proximal (gray points, target angle [-π:0]) halves of the workspace. (a,b) Initial movement angle versus target angle for monkey D and F, respectively. (c,d) Initial movement distance versus target distance for monkey D and F, respectively. Region within the dashed black lines falls within the diameter of the target.

under the model of minimum variance integration4. We computed this prediction under two limiting conditions: assuming that the initial angle variance arises only from variability in the sensory estimates of target direction (black dashed line, Figure 2.10a), or assuming that the measurements also include the maximal consistent level of downstream (e.g., motor) variability (see Methods; black dotted line,Figure 2.10a). The empirical variances observed in the VIS+ICMS condition follow the predicted trend closely, and lie between the two limiting predictions in the region where the animal received most of the multisensory training (20-25% coherence; see Methods). This comparison suggests that after training, the animals optimally combined the ICMS signal with vision. We next test this conclusion further by analyzing the pattern of mean initial angles.

The animals exhibited idiosyncratic patterns of mean initial reach angle as a function of target angle. For Monkey D, these patterns were clearly distinct between the VIS and ICMS trials (Figure 2.10b; also see Figure 2.11 for Monkey F, where the patterns are less well defined). Since the required movements are the same across cue conditions, these patterns likely arise from biased estimation of the target direction. Therefore, they offer another opportunity to test whether the VIS and ICMS signals are combined optimally. Minimum variance integration predicts that in the VIS+ICMS condition, as the visual coherence increases from 0% to 100%, the animals should transition from relying primarily on the ICMS cue to primarily on the visual cue. This trend can be seen qualitatively in the pattern of mean initial angles for Monkey D (Figure 2.10b): at 15% coherence, the VIS+ICMS mean was close to that observed with ICMS alone; at 100% it was close to that observed with VIS.

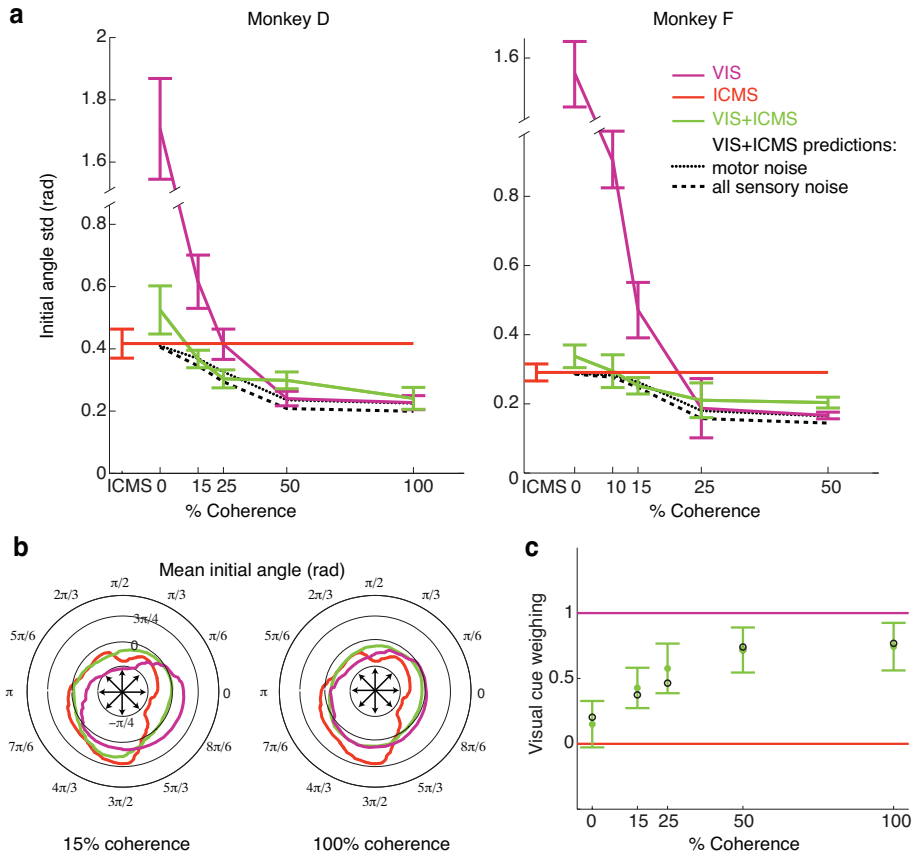

Figure 2.10: **Integration of vision and ICMS minimizes reach variance.** (a) Standard deviation of initial angle relative to target angle as a function of visual coherence for different feedback conditions for each monkey. Standard deviation was calculated after subtracting a smoothed estimate of mean initial angle (panel b); results were qualitative unchanged with only the target angle subtracted (i.e., angle computed with respect to straight-line reach; Figure 2.11). Error bars denote bootstrapped estimates of standard error. Dashed black lines indicate model predictions with no motor noise; dotted black lines indicate model predictions with maximal motor noise. (b) Mean initial angle, with respect to a straight-line reach. Smoothed values are shown on a polar plot as a function of target direction. (c) Visual cue weighting (see Online Methods) for combined VIS+ICMS conditions were closer to zero (ICMS) for low coherence trials and closer to one (VIS) for high coherence trials. Green filled circles: visual cue weighting estimated from data (Equation 2.8); black unfilled circles: minimum variance model prediction (Equation 2.7); error-bars: bootstrapped estimates of standard error. Data is from Monkey D.

The relative weighting of the two modalities can be estimated quantitatively by modeling the VIS+ICMS mean as an affine combination of the unimodal biases (Equation 2.5). As expected, the weighting of the visual cue smoothly transitioned from zero to unity as the visual coherence increased (green symbols, Figure 2.10c). Under the model of minimum variance integration, each sensory cue should be weighted inversely proportional to its variance (Methods, Figure 2.10b). Using the unimodal variance data from Figure 2.10a, we obtained quantitative predictions for the cue weighting in the VIS+ICMS trials (open black circles, Figure 2.10c; Equations 2.7-2.8); these are in good agreement with the empirical data.

2.4 Discussion

We have shown that multi-channel patterned ICMS of primary somatosensory cortex can be used to provide monkeys with continuous information about hand position that enables goal-directed reaching. In particular, the monkeys were able to use ICMS to estimate the distance and direction between their current hand location and the reach target. Furthermore, when both visual and ICMS feedback was available, the monkeys combined these signals to achieve increased levels of task performance, and they did so ator neartheoretical optimal levels, as is observed for natural sensory signals (Ernst and Banks, 2002; van Beers et al., 1999).

What does the ICMS signal convey?

An important finding of this study is that animals can learn to use ICMS as a temporally continuous feedback signal. However, it is possible that the animals only used the ICMS signal to estimate the target location during the instructed delay period, with subsequent

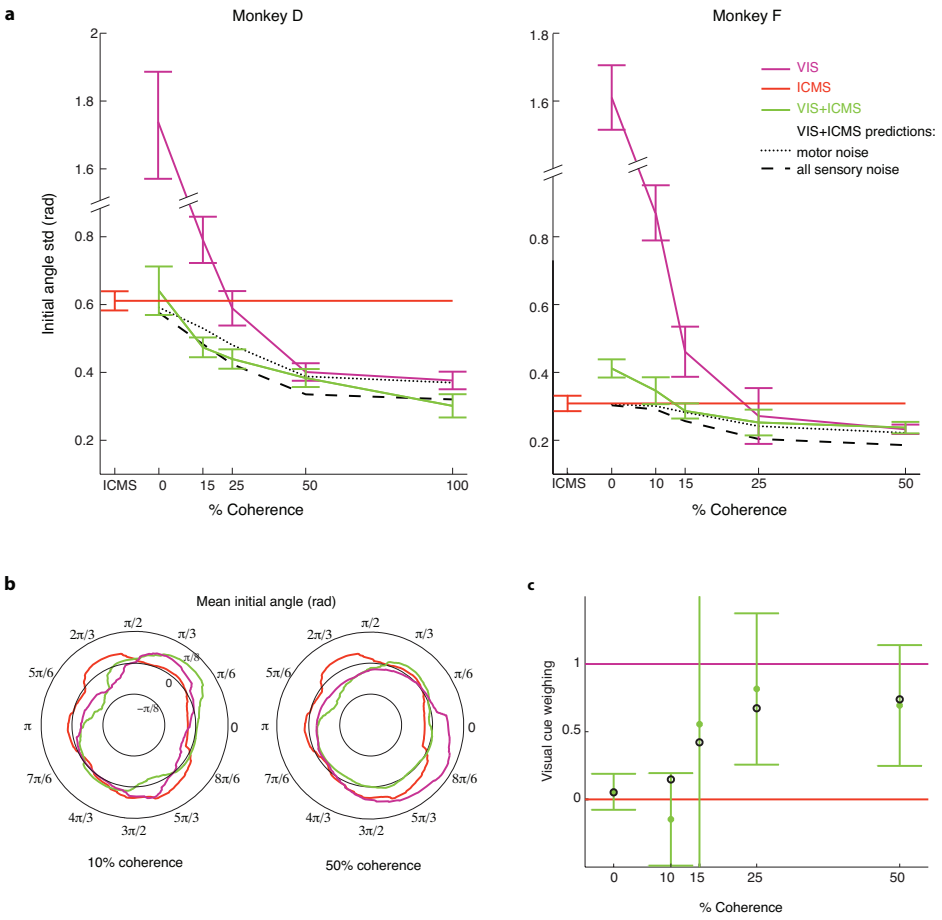

Figure 2.11: **Additional analyses of initial angle.** (a) Standard deviation of initial angle (relative to target angle) for the different trial types and visual coherences. Plots follow the same conventions as Figure 2.10a,b. This figure demonstrates that the qualitative results of the main text are not an artifact of the subtracting the smoothed estimates of mean initial angle used there. Error bars denote bootstrapped standard error. (b) Smoothed mean initial angle for Monkey F, relative to the target angle. Monkey F did not exhibit the marked differences in mean initial angle across feedback types that were observed for Monkey D. (c) Visual cue weighting for Monkey F in the VIS+ICMS trials, as a function of dot-field coherence. The results are consistent with the minimum variance model, however the analysis has poor statistical power due to the similarity in mean initial angle across feedback types. Green filled circles: visual cue weighting estimated from data; black unfilled circles: minimum variance model prediction; error-bars: bootstrapped estimates of standard error.

corrective sub-movements guided either by the remembered target location or simply by a random search. An analysis of corrective sub-movements, Figure 2.12, shows that this is not the case. Sub-movement direction correlates well with the direction of the online movement vector (target current hand position) at the end of the previous sub-movement. In fact, the precision of the corrective movements in the ICMS-only condition is comparable to that seen with high visual coherences, and is considerably lower than that observed with 0% VIS, where no directional information is available, or that which would be expected by chance. These results suggest that the VIS and ICMS cues are being used as online feedback signals. Furthermore, if the animals had simply memorized the location of the target during the instructed delay period, we would expect a decline in precision across sequential corrective movements. Instead, the correlation between cued and executed sub-movements was largely consistent across corrective sub-movement number, with no clear increase in error variance for later sub-movements. These results strongly indicate that the animals are using the online feedback to execute corrective sub-movements.

Another key result of the paper is that with only eight electrodes, we are able to deliver continuous spatial information with a reliability comparable to that achieved with the visual cue. Initial angle estimation with the ICMS signal has the same variance as that observed with 15-20% visual coherence and is only about three times greater than that observed with the highest coherence for both animals. Furthermore, when the signals are used online, ICMS performance is even closer to that achieved with visual feedback (Figure 2.12). The better performance (relative to vision) of ICMS during corrective movements could be due to either

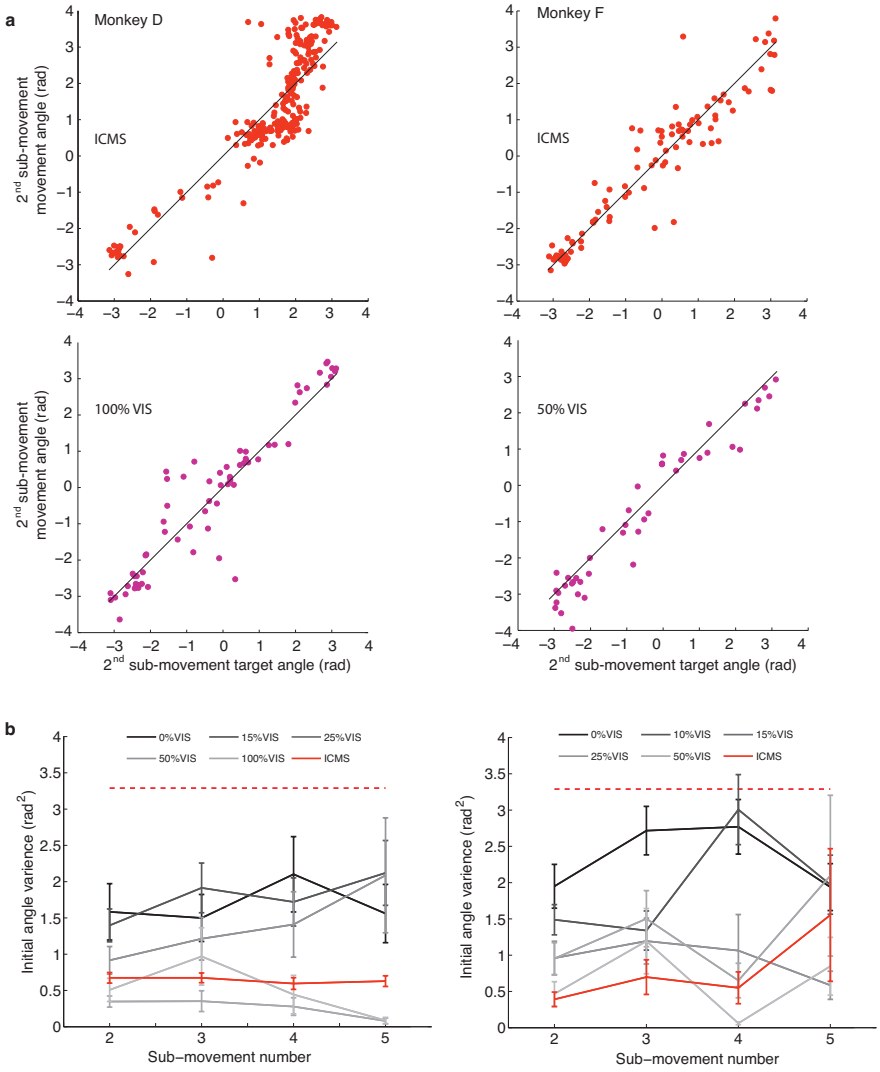

Figure 2.12: **Directed error correction.** a) Angle of the second sub-movement as a function of instantaneous movement vector angle in trials that required error correction. Top: ICMS-only trials; Bottom: VIS-only trials at high visual coherence (50% for Monkey F; 100% for Monkey D). Black line: unity. b) Error variance (rad2) in sub-movement angle estimation.

the inherent delays in visual feedback, the relative importance of somatosensory feedback for online movement controls, the shorter integration times available for online corrections, or a greater contribution of motor noise.

Learning to integrate a novel sensory signal The ability to use the ICMS signal was necessarily learned during the training process, since, by design, the patterns of ICMS do not mimic naturally occurring signals in the brain. This learning could be driven by several possible mechanisms of learning. In the experiment, the visual and ICMS signals changed in a correlated fashion. Previous modeling work from our lab showed that in a network with unsupervised, Hebbian-like learning, such correlations are sufficient to learn optimal integration (Makin et al., 2013). Other learning mechanisms may also have contributed, including error-corrective or reinforcement learning (Cheng and Sabes, 2007; Izawa and Shadmehr, 2011) of a sensory-to-motor mapping from ICMS to the appropriate movement. Our experimental design cannot definitively distinguish between these possibilities, but the emergence of multisensory integration before animals could perform with ICMS alone (Figures 2.4-2.5) suggests that unsupervised, multisensory learning played a large, if not dominant, role.

While there is evidence of multisensory integration at all visual coherences (Figure 2.6), optimal performance in the VIS+ICMS condition was only observed at mid-level visual coherences; performance at the lowest and highest coherence levels only approached optimality (Figure 2.10). A likely explanation is that monkeys learned to integrate vision and ICMS optimally when vision matched or was close to the training coherence (20% for monkey F; 25% for monkey D). In fact, a similar effect was observed in the network model of unsuper-

vised sensory integration (Makin et al., 2013): small but apparent departures from optimal integration were seen when the unimodal variances deviated too far from the regime in which the network was trained. Those observations are qualitatively consistent with the present results.

ICMS feedback as a tool for studying multisensory circuits

Animals ability to integrate ICMS signals and use them to control movement offers a novel and potentially powerful tool for studying information processing in sensorimotor circuits. For example, the posterior parietal cortex uses sensory feedback for a variety of multisensory computations, including estimation of the position of the limb and the location of targets (Batista et al., 1999; Battaglia-Mayer et al., 2000; Bremner and Andersen, 2012; Graziano et al., 2000; Kalaska et al., 1990). Computational models have been developed to demonstrate how neural circuits could perform these operations (Deneve et al., 2001; Ma et al., 2006; Makin et al., 2013), but testing these models has proven difficult. The challenge stems, in large part, from the fact that brain areas within the PPC exhibit complex, heterogeneous and partially redundant spatial representations (Chang and Snyder, 2010; Marzocchi et al., 2008; McGuire and Sabes, 2011) and interact in a complex network (Andersen and Buneo, 2002; Battaglia-Mayer et al., 2003; Kalaska, 1996; Sabes, 2011; Wise et al., 1997), often with overlapping function (Levy et al., 2007; Yttri et al., 2013). It may not be possible to discover how information is processed within this complex circuit only by manipulating the distal sensory inputs. ICMS feedback, on the other hand, affords the experimenter proximal control of the afferent signal. This has two pertinent advantages. First, the anatomical

origin of the signal can be controlled. In this particular study, we stimulated areas of somatosensory cortex that we know project directly to multisensory areas as such area 5 and VIP (Lewis and Van Essen, 2000; Pearson and Powell, 1978); however, stimulation could be performed in other brain areas that participate more strongly in other cortical circuits, e.g., in area 7 to study spatial representations in the circuits for saccadic eye movements. Second, because the ICMS signal bypasses peripheral receptors and subcortical processing, it gives the experimenter finer control over the statistics of the signal. Signal statistics play a big role in current models of multisensory neuronal processing (Ma et al., 2006; Makin et al., 2013), but the signal manipulations that would be most diagnostic for these models would be challenging or impossible to achieve with natural stimulation. With ICMS feedback, the manipulations become tractable, for example, changing the correlational structure between neurons in a given area or between neurons representing two sensory signals, and gaining full control over the timing and context of subjects exposure to the signal.

Furthermore, we think the learning observed in the present study taps into the same mechanisms of plasticity that drives other forms of multisensory learning, such as intersensory calibrationthe ability of two sensory modalities to come back into alignment following a perturbation (Burge et al., 2008; Redding and Wallace, 2002; Sabes, 2011; Zaidel et al., 2011). ICMS feedback offers an ideal tool for studying these mechanisms, e.g., for testing models of learning (Makin et al., 2013) in real neural circuits, for the reasons described above local access and control of statisticsand because we can observe changes in the behavioral and electrophysiological state of the animal from the very first exposure to the novel signal.

Neuroprosthetic applications

In this study, we show that ICMS can be used to deliver a task-relevant feedback signal that guides online, multi-dimensional movement control. We chose to encode the position of the hand relative to the reach target, as opposed to an absolute proprioceptive signal, because this allowed us to study artificial sensation in a simple, natural task, without having to suppress natural proprioception. Because the reach target was never visible, estimation of the relative hand position was required in order to perform the task. To use this approach to provide proprioceptive feedback from a prosthetic device, the ICMS signal would instead encode the state of the device with respect to the body, for example joint or endpoint position or velocity. Because these variables are also available via visual feedback, the same learning mechanisms should apply.

We expect that ICMS feedback could play the same role for BMI control that proprioceptive feedback does for normal movement control. During natural movement, vision and proprioception make comparable contributions to limb state estimation (Sober and Sabes, 2005, 2003; van Beers et al., 1996, 1999). While proprioceptive loss does not have a substantial effect on the simplest reaches when visual feedback is available, it does impair performance of movements that require inter-joint coordination, which include most activities of daily life (Sainburg et al., 1995). We have shown that the eight-channel ICMS signal used here can provide online feedback with reliability that is also comparable to vision. We expect that increasing the number of electrodes and optimizing the encoding scheme will further improve the quality and information capacity of this signal.

Our learning-based approach can be contrasted with a biomimetic approachthe attempt to reproduce natural patterns of sensory-evoked neural activity. In practice, a truly biomimetic stimulation scheme is not attainable, due to a range of technical and scientific considerations, such as the lack of access to the full neuronal population, inability to translate electrical stimulation into naturalistic neural activity patterns, an incomplete understanding of neural encoding mechanisms, and difficulty in characterizing those patterns in patients with sensory loss. Our approach circumvents these issues by taking advantage of the inherent plasticity of the brain. On the other hand, completely disregarding prior knowledge about natural neural signals, such as somatotopic organization, will impair initial performance and may limit the rate or extent of learning. In fact, learning-based and biomimetic approaches (Choi et al., 2012; Daly et al., 2012; Fagg et al., 2007; Tabot et al., 2013b; Weber et al., 2012) are highly complementary, and systems optimized for clinical applications would likely benefit from taking a combined approach.

Chapter 3

Optimizing the ICMS signal

3.1 Introduction

In the previous chapter we saw that intracortical microstimulation (ICMS) could be used to encode an abstract, two-dimensional variable in real time. Monkeys were asked to perform a reaching task to an unseen target guided by a vector describing the distance and direction between their current hand position and the center of the reach target. The ICMS signal encoded the amplitude of this "movement vector" in the mean stimulation frequency across all of the stimulating electrodes, where larger vector amplitudes resulted in higher mean frequencies and smaller resulted in lower. Movement vector direction, θ, was encoded in the relative firing rate across electrodes by assigning each electrode a "preferred direction" (PD) chosen from the set of cardinal and ordinal directions. Then, during a behavioral trial, the stimulation frequency of each electrode was set to be proportional to the difference between

it's PD and the direction of the movement vector.

Monkeys ultimately learned to perform the reaching task using this artificial signal; in fact, ICMS was ultimately as precise as visual feedback at low-to-middle levels of reliability. For our second monkey (Monkey F), this translated into a target angle estimation with near-zero bias and an error variance of 0.09 rad^2. But despite being functionally informative, this level of performance is more than twice as variable (or half as precise) as performance with the highest level of visual coherence. How can that be? Is this an inherent limitation on the precision of an artificial sensory signal, or does it reflect some factor of our implementation, such as the range of stimulation frequencies or the number of electrodes that we used to encode the signal?

The process of translating an external vector into movement takes many steps, and noise can corrupt the signal at each step. For simplicity, we will concentrate on encoding of a single variable, movement vector direction, θ. Consider the process: at each point during a reach, θ is calculated and then translated into stimulation frequencies across eight electrodes (see Equations 3.1 and 3.2). Electrical stimulation evokes neural activity, but not deterministically, but with some noise added (see Butovas and Schwarz, 2003). Some portion of this neural population projects onto downstream areas in cortex, which are responsible for reading out (decoding) the input neural activity and estimating θ. Finally, the new estimate of θ must be incorporated into a motor plan, and the movement must be accordingly adjusted.

The issues evoked by this process can be formalized into two, related questions whose answers will help us optimize encoding of an artificial signal:

- First, how is the ICMS signal being decoded?
- Given this neural decoding algorithm, what factors of ICMS encoding limit the signal's precision?

To answer the first question, we randomly perturbed the ICMS signal while the monkey was performing the behavioral task by silencing (setting to zero amplitude) a subset of stimulating electrodes. We then compared the patterns of errors he made under perturbation to predictions made several possible decoding algorithms. Although the ICMS signal encoding contains redundancies (eight parameters with only two unknown), we expected that a minimum-variance solution would use information from all eight electrodes (*à* la multisensory integration seen in Ernst and Banks, 2002; van Beers et al., 1996) and so for simplicity restricted our analysis to the following three decoding algorithms: the analytical solution (ANA), a center-of-mass decode (COM), and a winner-take-all approach (WTA). Please keep in mind that these strategies needn't be implemented by the monkey at the cognitive level– they can be thought of as describing the computations performed by downstream neurons on the incoming, multi-channel signal.

1. Analytic: The monkey computes the (correct) analytic answer to decode θ from the set of stimulation frequencies, SF (where SF is an 1xN vector, where N is the number of stimulating electrodes). Then target angle is estimated as:

$$\hat{\theta} = \tan^{-1}\left(2(SF - \frac{1}{2})\begin{bmatrix} \cos(PD_1) & ... & \cos(PD_N) \\ \sin(PD_1) & ... & \sin(PD_N) \end{bmatrix}^{+}\right)$$

where $^+$ marks the pseudoinverse (see Eqns. 3.1-3.2 for a description of ICMS encoding).

2. The monkey uses a center-of-mass strategy (COM), weighing each electrode by its stimulating frequency. Then, target angle is estimated as:

$$\hat{\theta} = \frac{\sum_{i=1}^{8} PD_i f_i}{\sum_{i=1}^{8} f_i}$$

where PD_i is the direction assigned to be encoded by the electrode i and f_i is the stimulation frequency of that electrode.

3. The monkey uses a winner-take-all strategy (WTA), moving in the preferred direction of the electrode with the highest stimulation frequency. Then target angle is estimated as:

$$\hat{\theta} = PD(\max[f])$$

where f is the set of stimulation frequencies across all eight electrodes and PD is the set of directions encoded by its corresponding electrode.

To answer the second question, we modeled the effect of several of the parameters used to encode the ICMS signal (number of stimulating electrodes, stimulation frequency, number of neurons activated by stimulation) on behavioral performance and compared the experimental and theoretical results. To explicitly compare the effect of number of stimulating electrodes on ICMS signal precision (as measured by behavior), we trained the monkey on a degraded ICMS signal and examined his level of improvement.

3.2 Methods

In this section, I have tried to include enough information to make this chapter parse-able without external reference, but I have omitted many details. Please refer to the methods section in Chapter 2 for a more detailed description.

Behavioral task

All animal procedures were performed in accordance with the National Research Councils Guide for the Care and Use of Laboratory Animals and were approved by the UCSF Institutional Animal Care and Use Committee. One adult male rhesus macaque monkey was implanted with a 96-channel silicon microelectrode array coated with Iridium Oxide (Blackrock Microsystems) over left primary somatosensory cortex (Brodmann Areas 1, 2; S1). The animal was trained to perform reaches in the horizontal plane to an unseen target in a two-dimensional virtual reality environment, where a mirror and an opaque barrier prevented direct vision of the arm. The mirror reflected visual input from a projector, so that the visual cues appeared in the horizontal plane of the reaching hand.

Each trial consisted of four epochs: i) The monkey moved to a start position, located in the center of the workspace (circle of 1 cm radius; Figure 3.1a). ii) The monkey was required to maintain its position for an instructed delay period (1-1.5 s) during which the monkey received either visual (VIS), ICMS, or bi-modal (VIS+ICMS) feedback encoding a movement vector (the distance and direction between his fingertip and the center of the reach target). Reach target positions were selected uniformly from an annulus centered on the start target

with an inner radius of 40 mm and an outer radius of 80 mm (Figure 3.1a). iii) After a go cue the monkey was free to reach under the guidance of continuously updating VIS, ICMS or VIS+ICMS feedback. iv) After acquiring the target the monkey received a liquid reward.

Visual feedback The movement vector was encoded with a random moving-dot flow-field, where each dot in the field moved at a speed proportional to the length of the movement vector. A percentage of the dots moved coherently in the direction of the continuously updating movement vector, while the remaining dots moved in random directions. In the experiments described below the dot field coherence was always set at twenty percent–the value at which vision and ICMS seemed to be equally reliable (see Chapter 2).

ICMS Intracortical microstimulation consisted of biphasic, charge-balanced pulse trains delivered asynchronously to each of eight electrodes in the array. The pulse amplitudes varied across electrodes, depending on perceptual threshold (established using a two-target forced choice task) and ranged between 30-80 μA. Eight electrodes, each of which could be individually detected, were chosen to encode the ICMS signal.

The movement vector was encoded via the spatial and temporal patterns of stimulation across the eight stimulating electrodes. Movement vector direction was encoded by the relative stimulation pulse rates across the electrodes. First, each of the electrodes was arbitrarily assigned one of eight preferred directions (depicted as back arrows in Chapter 3, e.g. Figure 3.1b), equally spaced around the circle. Then, the stimulation pulse rate f_i of electrode i was calculated as a function of angle between the direction movement vector, θ, and the

electrodes assigned PD,

$$f_i = \frac{f_0}{2}(1 + cos(PD_i - \theta)) \tag{3.1}$$

The frequency scaling factor, f_0, linearly encoded the movement vector distance, d, within the range of 100-300 Hz:

$$f_0 = \frac{d}{d_{max}} f_{max} \tag{3.2}$$

where d_{max} was 10 cm. The values of d and θ were continuously updated during the reach to provide online feedback.

Experimental Design

The monkey had previously learned to use the eight-electrode ICMS signal (see Chapter 2). Furthermore, he had shown signs of integrating vision and ICMS to gain improved performance on trials using combined feedback (vs. either uni-modal feedback signals), particularly when the visual dot-field coherence was fixed at twenty percent.

Instantaneous degradation experiments.

To assess how the monkey used the ICMS signal, we looked at the causal relationship between a degraded ICMS signal and task performance. We degraded the IMCS signal by "silencing" (setting to zero μA current amplitude) a subset of electrodes. The signal was only degraded during ICMS-only trials. To avoid adaptation or learning in response to the degraded signal, we degraded the signal *instantaneously*—randomly interspersing trials with degraded and intact conditions. The signal could be degraded by silencing two, four, six, or zero (a control)

stimulating electrodes. There were two separate sets of degradation conditions (Figure 3.1b), each of which silenced different subsets of electrodes. Only one set of conditions was ever used per behavioral session.

Behavioral sessions.

Behavioral sessions consisted of randomly interspersed VIS+ICMS and ICMS-only trials. Degradation sessions consisted of one-half VIS+ICMS trials and one-half ICMS-only trials, where the ICMS signal condition was drawn uniformly from a set of four possibilities: 8 electrodes, 6 electrodes, 4 electrodes, or 2 electrodes. There were two sets of degradation conditions, which silenced different subsets of electrodes (see Figure 3.1b). Training sessions consisted of VIS+ICMS trials with fully intact ICMS. These sessions were inserted before each of the last two testing sessions to battle what what appeared to be a gradual decline in ICMS-only performance across behavioral sessions (see Figure 3.1c). The experiment took place over eight behavioral sessions in the following sequence: Set 1, Set 1, Set 2, Set 2, Training, Set 1, Training, Set 1.

ICMS recovery

We tested if the animal could be re-trained to agilely use a degraded ICMS signal by pairing the degraded signal with vision. Training sessions consisted of 66.7% VIS+ICMS trials and 33.3% ICMS-only. There were three total training sessions. On the first two sessions, the experimenter had mistakenly degraded the signal to five (rather than four, as had previously been planned) electrodes. On the third session, the degraded signal consisted of the cardinal

directions, as planned (see Figure 3.7). Although further training sessions were scheduled, conditions outside of the experimenter's control forced an early termination of the study.

Data analysis

Data metrics

Regressions All regressions shown were conducted on ICMS-ony conditions (excluding any VIS+ICMS data points). Linear regressions were done using the "regress function" in MATLAB. Expontial regression were simply linear regression using log(y) as the dependent variable.

Fraction correct trials: Fraction correct trials is computed as the sum of completed trials divided by the number of trials the monkey attempts to complete. This calculations excludes errors due to factors such as not initiating a trial, starting to move before the go tone, and swiping through the reach target.

Number of movement sub-segments: This metric quantifies the number of discrete sub-movements in a trial. Starting with the model assumption that sub-movements have bell-shaped velocity profiles (51) we identified sub-movements by threshold crossings of the tangential velocity plot of a trajectory, with a threshold of 20% of the maximum velocity on a given trial.

Initial movement angle: The initial movement angle is the angle of the monkey's initial movement (the direction of the endpoint of the first sub-movement).

Initial movement angle error: Initial angle error is computed as the difference between

the angle of the monkey's initial movement angle and the angle of the reach target.

$$err = (initial\ mvmt.\ angle) - (target\ angle)$$

Behavioral performance as a function of trial parameters: We looked at how task performance was affected by five parameters: reach target radial distance, target x-position, target y-position, target angle, and the duration of the instructed-delay period. For all but the first parameter, we divided the parameter space into 10 equal bins, and calculated the average initial angle error and fraction correct trials for each bin. For radial target distance we needed to first square the values to ensure a uniform distribution of points. The calculations for two-dimensional plots were similar, but divided the trials into a (10x10) matrix of possible parameter combinations.

Modeling

Modeling target angle estimation To differentiate between the three possible decoding strategies (Analytic solution, Center-of-Mass, and Winner-Take-All), we modeled target angle estimation under the strategies during both intact and degraded input, using the same trial parameters (target locations) we used in the experiments. We simulated "silencing" electrodes by removing those preferred directions from the following calculations.

In the following equations, $\hat{\theta}$ is movement angle, PD_i is the direction assigned to be encoded by the electrode i, and f_i is the stimulation frequency of that electrode.

We estimated movement angle for COM as :

$$\hat{\theta} = \frac{\sum_{i=1}^{8} PD_i f_i}{\sum_{i=1}^{8} f_i}$$

To properly compute a circular center-of-mass, we rotated the preferred directions of the electrodes (originally in the range $[-\pi, \pi]$) to be centered on the target angle. For example, for a target angle of 3 radians, the electrodes whose PD were less than $(\pi - 3)$rad were shifted up by 2π.

For WTA, movement angles were estimated as:

$$\hat{\theta} = PD(max[f])$$

Finally, for ANA, movement angle is estimated as:

$$\hat{\theta} = \tan^{-1}\left(2\left(\begin{bmatrix} f_1 \\ \\ f_N \end{bmatrix} - \frac{1}{2}\right)\begin{bmatrix} \cos(PD_1) & ... & \cos(PD_N) \\ \\ \sin(PD_1) & ... & \sin(PD_N) \end{bmatrix}^{+}\right)$$

where $^{+}$ marks the pseudoinverse.

Later, when we look at the effect of ICMS parameters on the precision of the signal, f_i, the stimulation frequency on each electrode, is replaced by g_i, the set of neural responses evoked by each electrode. We find g_i by first choosing the number of neurons, Q, that are activated by a single electrode (this could range anywhere from 5-500 neurons—see Histed et al., 2009 and Tehovnik, 1996), and then assuming that each neuron will fire with a rate drawn from a poisson distribution with a mean equal to the stimulation frequency ,f_i, of the nearby electrode (a common noise model used to estimate neural response to natural stimuli). Thus we get a Qx1 vector of neural responses on each electrode, g_i, where $g_{i,j} = Pois(f_i)$ is the response of neuron j on electrode i. Then, our equations for each decoding scheme are slightly modified from before to:

- COM

$$\hat{\theta} = \frac{\sum_{i=1}^{8} PD_i \sum_{j=1}^{Q} g_{i,j}}{\sum_{i=1}^{8} \sum_{j=1}^{Q} g_{i,j}}$$

- WTA

$$\hat{\theta} = PD(\max \left[\sum_{j=1}^{Q} g_{1,j}, ..., \sum_{j=1}^{Q} g_{N,j} \right])$$

- ANA

$$\hat{\theta} = \tan^{-1} \left(2(\begin{bmatrix} \bar{g_1} \\ \\ \bar{g_N} \end{bmatrix} - \frac{1}{2}) \begin{bmatrix} \cos(PD_1) & ... & \cos(PD_N) \\ \\ \sin(PD_1) & ... & \sin(PD_N) \end{bmatrix}^{+} \right)$$

To get a robust estimate of performance, we generated 10000 artificial trials for these simulations, drawing the target angle in each trial uniformly from [-π,π]. Electrodes were "silenced" for degradation trials by setting the responses of neurons on the silenced electrodes to zero.

3.3 Results

Decoding the movement vector from ICMS

Performance metrics

The monkey made the quickest, most direct reaches to targets under the VIS+ICMS condition. With ICMS alone—even with all eight electrodes—he performed somewhat worse than with VIS+ICMS in all of the performance metrics we employ. In ICMS-only trials, the monkey's behavioral performance improved with increasing number of stimulating electrodes (except for total movement distance). For instance, he completed a larger fraction of trials (Figure 3.1c) with more electrodes: at his worst (2 electrodes) he completed less than 10%

of trials, while at his best (with eight electrodes) he completed nearly 95% of trials. The reaches he made with more electrodes were also faster (Figure 3.1e) and required fewer online corrections (Figure 3.1f) than those with fewer electrodes.

Once interesting and unexpected result was the sensitivity of behavioral performance to the identity of the electrodes we silenced (set identity; Figure 3.1b) and to extended testing on degraded stimuli. First, we see that the monkey tended to perform better with Set 1 electrodes than with Set 2 electrodes (Figure 3.1g). We can conclude that thought we thought all electrodes were equal, some electrodes were actually more equal than others. Furthermore, when degradation sessions were presented in sequence, behavioral performance across sessions grew continuously worse in all conditions; however, we found that performance could be rescued by inserting a training session (all VIS+ICMS with intact ICMS; small bold T at bottom of Figure 3.1g) between degradation sessions. Including a training session largely recovered behavioral performance (see 4, 6, and 8 electrodes), even elevating it above the original levels (see 4 and 6 electrodes). This sensitivity of the ICMS percept to degradation was very surprising, as more than half (62.5 %) of all trials on these days featured intact ICMS. It seems that confidence in the ICMS signal (and thus signal reliability) is not only determined by the trial-to-trial reliability of the signal but is also learned over time.

Error in target angle estimation

To determine how ICMS signal is decoded by the monkey, we examined plots of initial movement angle as a function of true target angle during intact and degraded ICMS conditions

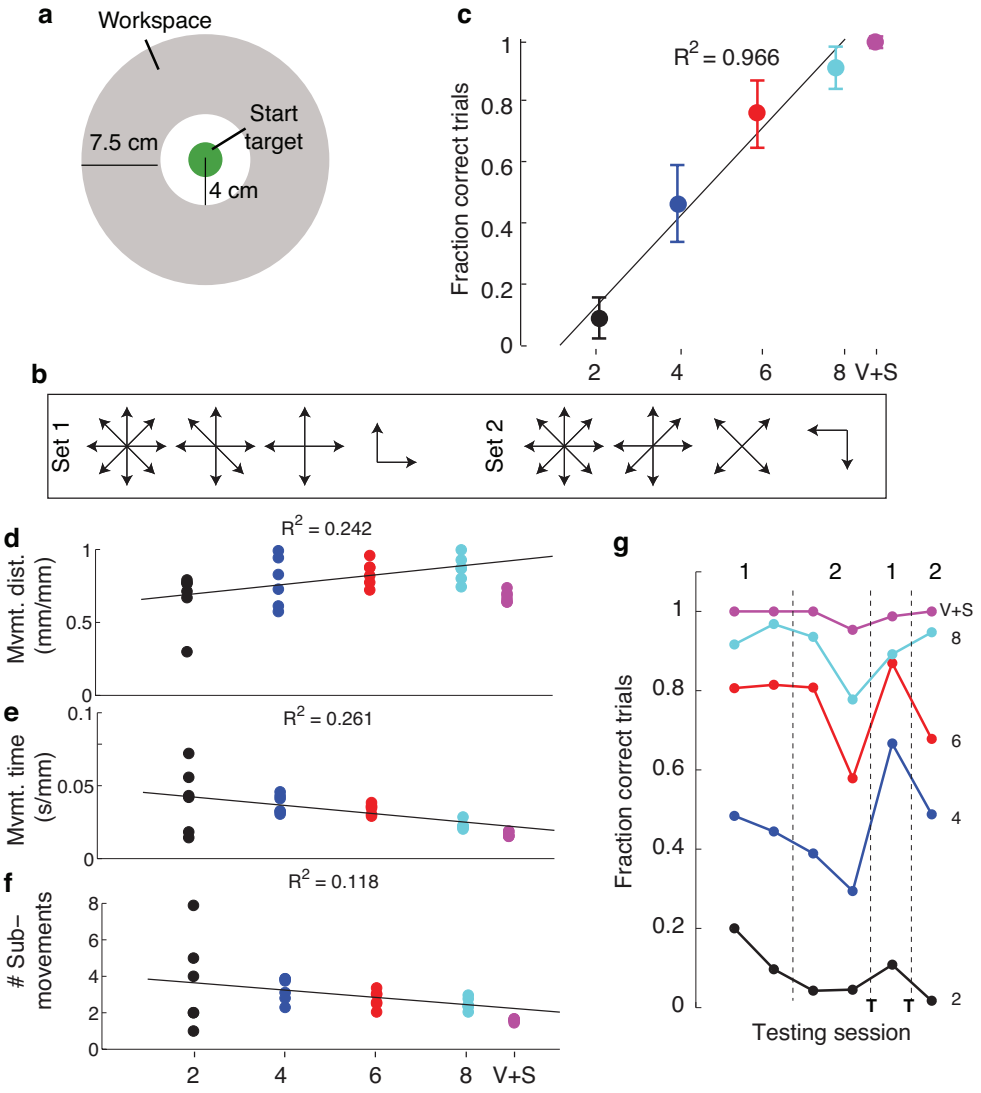

Figure 3.1: **Instantaneous degradation of the ICMS signal.** R^2 values and fit lines report a linear regression on the ICMS-only conditions (no VIS+ICMS). a) Behavioral task layout. Reach target centers are drawn uniformly from an annulus (gray area) surrounding the start target location (green circle). b) The two sets of degradation conditions. In this schematic, each arrow indicates that the electrode with the matching preferred direction had non-zero current amplitude in that condition. Missing arrows indicate silenced electrodes. c) Fraction correct within a degradation condition, averaged across sessions and degradation sets. d) Normalized movement distance (path length divided by target distance). e) Normalized movement time (time divided by target distance). f) Number of movement sub-sections. g) Percent correct trials across training sessions. At the top the set of degradation conditions used in each session (set 1 or set 2) is labeled. Dashed gray lines mark switches in conditions. Small bold T's denote training sessions interspersed between testing sessions. At right each condition is labeled: V+S is paired, 20% coherence; numbers indicate number of stimulating electrodes (8, 6, 4, or 2).

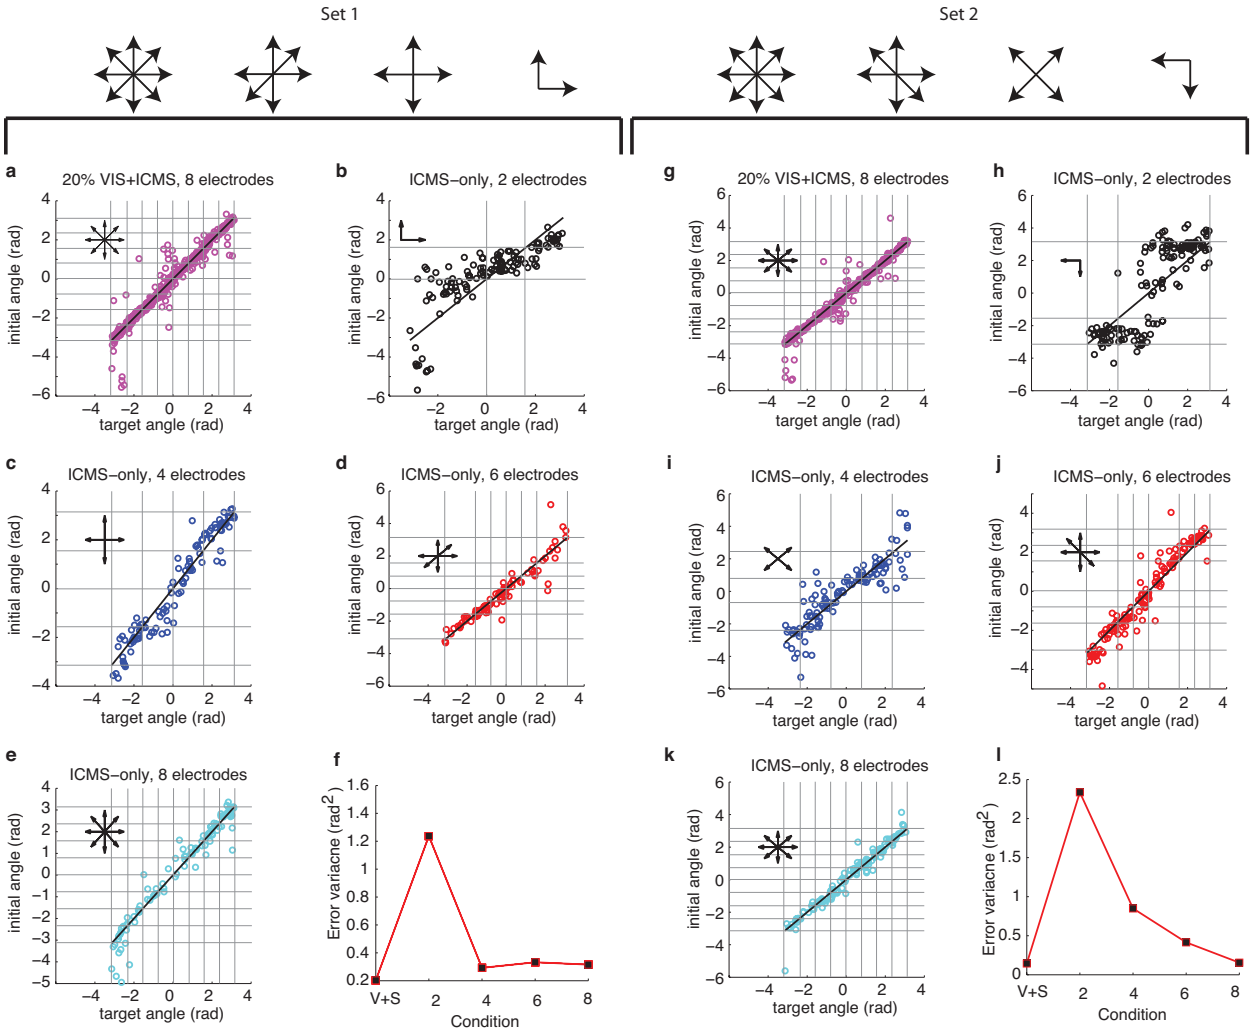

Figure 3.2: **Initial angle estimation as a function of reach target angle** Behavioral task performance across degradation conditions. Data points are from all attempted trials, not only those that the monkey completed correctly. The left column plots data from behavioral session using degradation set 1, and the right plots that from set 1. In each of the initial angle plots the electrodes encoding the ICMS signal are plotted at upper left (black arrows, where each arrow indicates the PD of one of the non-zero amplitude stimulating electrodes). At bottom middle and far right is plotted the correlation coefficient for each condition. Note each subfigure has different axis labels. Thin grey lines mark the PDs of active electrodes.

(Figure 3.2). To increase our sample size, we plotted initial movement angles from all attempted trials, even ones where the monkey did not complete the reach to the target. As expected, the variance of errors in estimating target angle was higher when fewer electrodes were encoding the ICMS signal, at least for the degradation conditions in Set 2 (Figure 3.2l). In contrast, for Set 1, error variance seems relatively fixed from four to eight electrodes (Figure 3.2f). Furthermore, initial angle estimates seemed biased towards the PDs of the remaining electrodes, as shown clearly in Figure 3.2b,h. In this condition, initial angle estimates were somewhat accurate in the range between the two electrodes' preferred directions ($[0, \pi/2]$), but they were biased away from the true answer, towards these preferred directions outside of that range.

We tried to explain this data by estimating the pattern of errors that would be seen during degraded ICMS trials under three decoding strategies, ANA, COM, and WTA (Figure 3.3), translating stimulating frequencies directly into estimated movement angles (see Methods). Although we made estimates for all of the ICMS degradation conditions, here we only plotted estimates of movement angles from the two-electrode condition from Sets 1 (Figure 3.3a) and 2 (Figure 3.3b), as these were the conditions with the most distinct error patterns. Both the ANA and COM models captured the pattern of errors seen in Figure 3.2 remarkably well, particularly in the two-electrode condition for Set 1 (Figure 3.3a). In contrast, for Set 2, all of the models captured the two-electrode condition equally well (Figure 3.3b). The patterns of errors for four and six active electrodes (not shown) were more subtle and difficult to compare directly, but they are better approximate (globally) by either ANA or COM estimates than

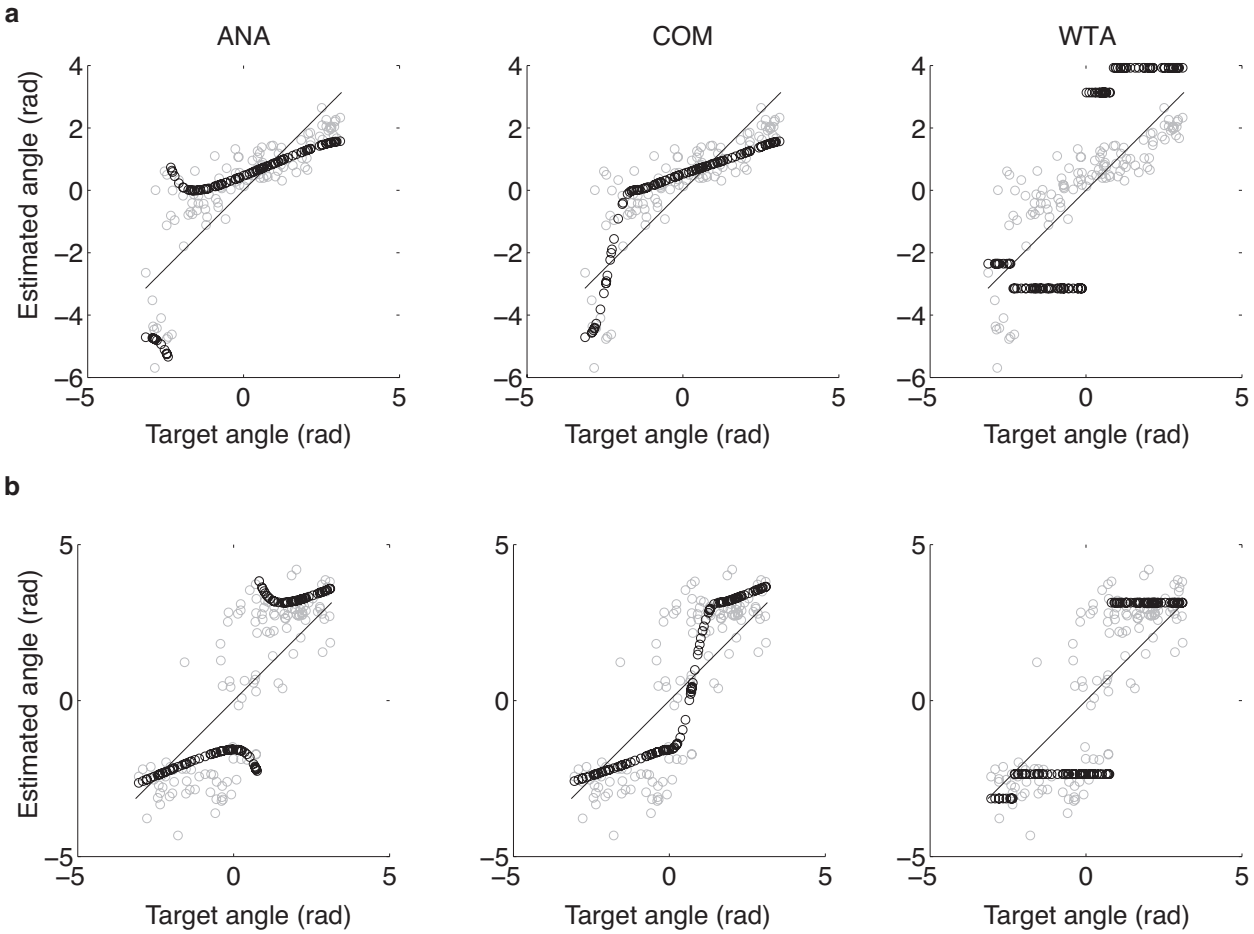

Figure 3.3: **Modeling Initial angle estimation** Target angles match the data in Figure 3.2b,h. Grey circles: true experimental data. Black circles: simulated movement angles for each trial under ANA (left), COM (center), and WTA (right). Black lines denote unity (the encoded target angles). a) 2-electrode condition for Set 1. b) 2-electrode condition for Set 2.

by WTA (see 3.4). We conclude that it is most likely that the monkey used one of the more sophisticated decoding algorithm, ANA or COM, to estimate the movement vector angle encoded in the ICMS signal; it is difficult to differentiate between the two at this point since their accuracy was so similar.

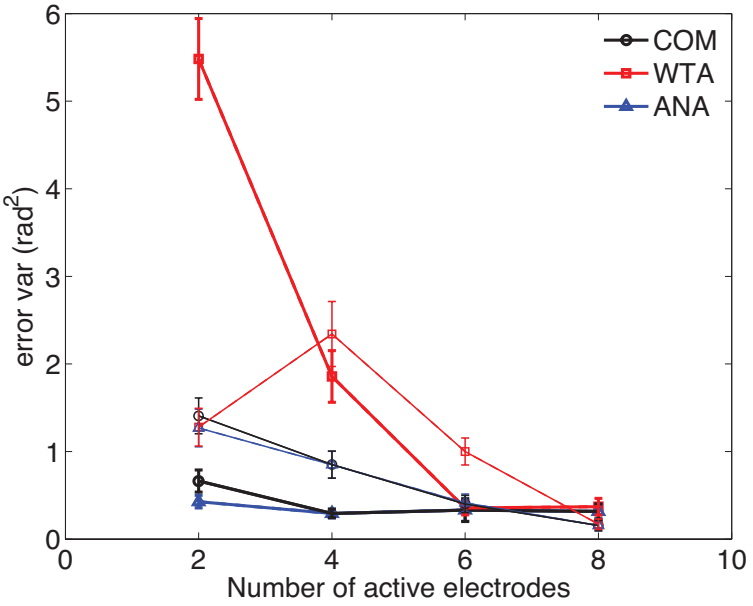

Figure 3.4: **Error variance of model predictions** Variance of errors between experimental (plotted in Figure 3.2) and the model (Figure 3.3) estimates of movement angle. Thick lines represent results for data from Set 1; thin lines from Set 2. Error vars denote standard error of the mean estimate for error variance, computed by a bootstrap method.

Limits on ICMS precision

The models described above assume a noiseless decode of movement vector angle from the set of stimulation frequencies, but, of course, this is not the case. The ICMS signal incurs noise in translating from stimulation frequency to firing rate of activated neurons, in the transmission of that information from the stimulated cortex to downstream areas responsible for decoding, and in translating the decoded estimate into actual movement. Other sources of noise could include interference between electrodes that activate overlapping populations of neurons and neural encoding of other parameters, such as natural proprioception (Pei et al., 2010; Prud'homme and Kalaska, 1994). Unfortunately, there is little we can do to try to minimize some of these sources of noise, but we *can* try to optimize the parameters of the

ICMS signal in order to maximize signal precision. In particular, we can model the expected effect of the following parameters on ICMS reliability: the number of stimulating electrodes (N), the stimulation frequency, and the number of neurons that are activated per electrode (Q). It is obvious how to manipulate the former parameters, and the latter can be controlled indirectly, by adjusting the stimulating current amplitude, which increases the probability of spiking in nearby neurons (Butovas and Schwarz, 2003). We incorporated these parameters into the decoding models as described in the Methods section, and plotted the results in Figure 3.5.

The predictions made by the models are somewhat complicated, and no single parameter acts independently, but there are several general remarks that we can make. First, increasing the number of stimulating electrodes (x-axis on all plots) will improve the precision of the ICMS estimate, but with diminishing returns as performance asymptotes. In fact, the *rate* at which performance asymptotes is correlated highly with stimulation frequency (compare rates for different stimulation frequencies in Figure 3.5), while the *level* seems to be set by the number of activated neurons (compare Figure 3.5c and d). Furthermore, surprisingly, if noise is added to the model, than COM makes better, more precise estimates than ANA in all cases but for very few electrodes with high firing rates (see Figure **??**b,d).

To test the predictions of these models for the effect of stimulation frequency on behavioral performance, we plotted the data from the instantaneous degradation experiments as a function of the workspace (Figure 3.6a-d) and then projected that data onto the parameters of the movement vector: distance and angle (Figure 3.6f-j). As our encoding scheme mapped

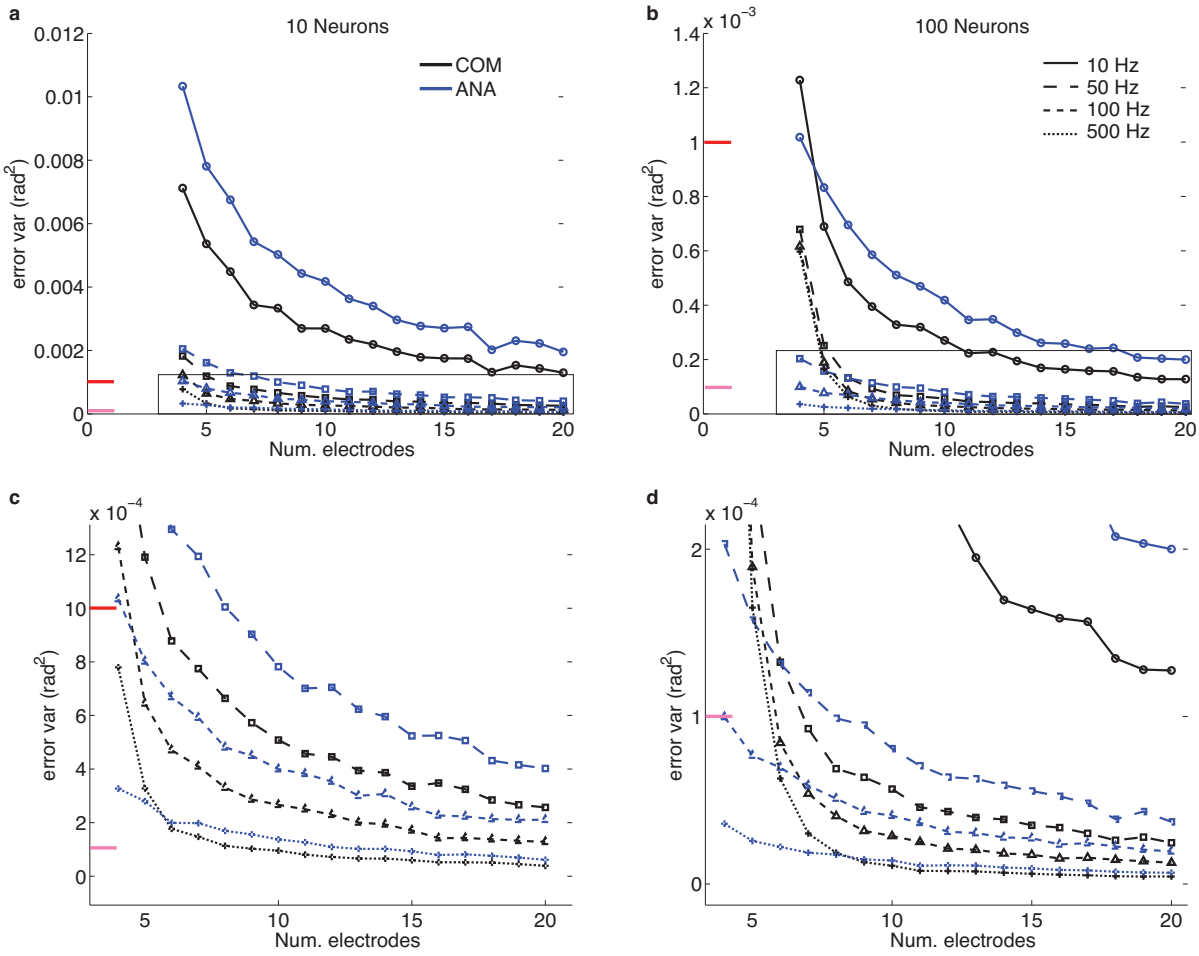

Figure 3.5: **Expected effect of ICMS parameters on signal precision.** The expected accuracy of movement vector angle estimation using COM (black) and ANA (blue) models as a function of ICMS signal encoding parameters. a-b) Interaction between number of neurons activated (a: 10, b: 100), number of stimulating electrodes, and stimulation frequency. c-d) Magnified version of top panels (areas marked by thin rectangle on top). Red and pink lines along vertical axis mark the same values across all panels.

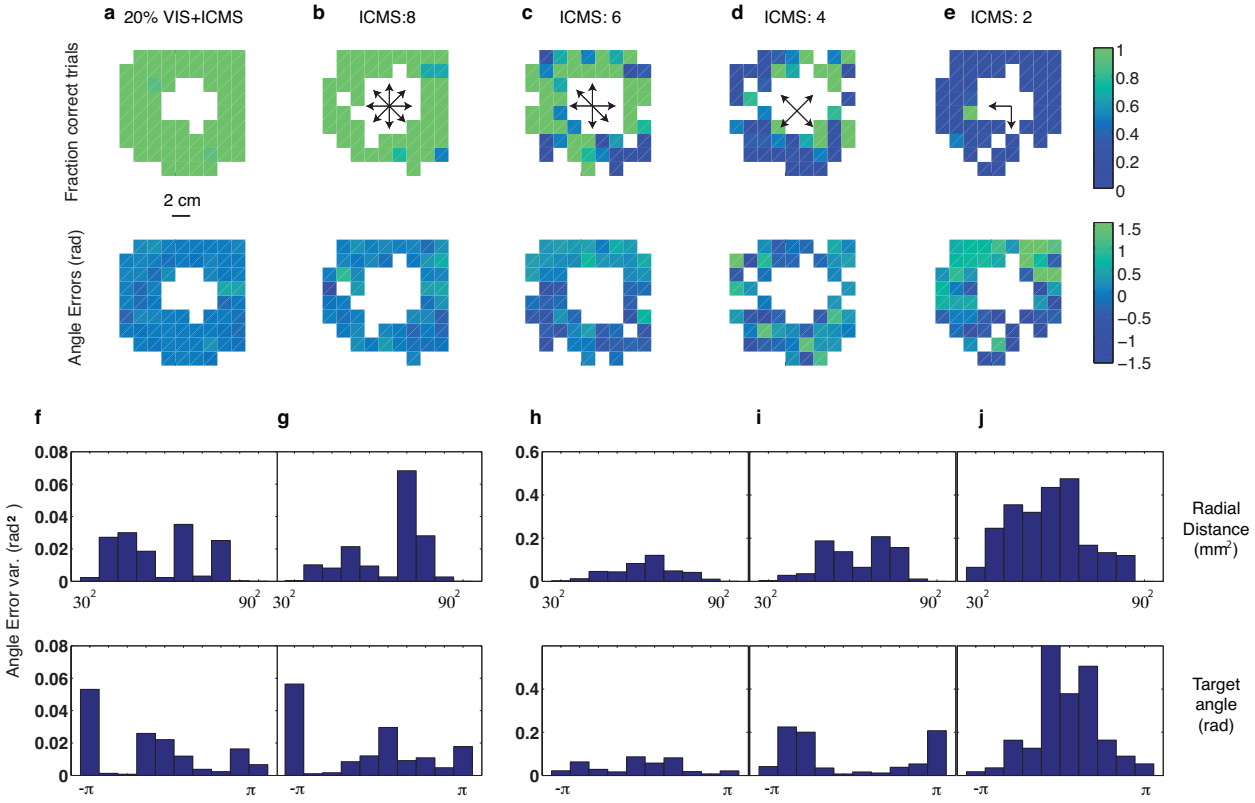

Figure 3.6: **Behavioral performance across the workspace.** Fraction correct trials (top) and average angle errors (bottom) are plotted as a function of [x,y] position in the workspace (Figure 3.1a). a-e) This data shows averages across the three session using Set 2 degradation conditions (see black arrow in center of each top plot). f-j) Projection of error variance from data in (a-e) onto individual parameters: radial distance and target angle.

distance onto mean stimulation frequency, this was an direct method of testing the effect of frequency on signal precision. In this experiment, the hand-to-target distance range of 30 mm to 90 mm mapped onto a stimulation frequency range of 80 to 240 Hz. We can see that in this range, there was no systematic effect of frequency on error variance when ICMS is intact (Figure 3.6g, top panel). At the very least, error variance did not decrease with increasing stimulation frequency, as was predicted by the neural decoding models. Similarly, there is no systematic error in initial angle estimation as a function of target direction.

In contrast, degrading the ICMS signal by silencing electrodes profoundly effected the monkey's ability to perform the behavioral trials. Scanning the plots from left to right (a-e) visually exemplifies the degradation of performance and accumulation of bias that was incurred with fewer electrodes. By the time only two electrodes were left, the monkey was able to complete just a few trials: those in which the reach target happened to lie close to the PD of one of the active electrodes—the direction in which the monkey tends to head in most trials with this condition (Figure 3.6e).

What can explain the discrepancies between the model predictions and our experimental results? To start with, the error variance predicted by these models (ranging from $2^{-}5$ to 0.01 rad^2 in Figure 3.5c) is orders of magnitude smaller than that in the experiments described in Chapter 2 (compare to an experimental ICMS error variance of 0.09 rad^2), even after removing the maximum variance possibly attributable to motor noise (then 0.07 rad^2) and assuming a paltry 10 neurons activated per stimulating electrode. One possibility is that, within the range of stimulation frequencies used during this experiment to encode target angle (80-240 Hz), there is little noise reduction to be gained, and we would need a larger number of trials to be able to observe the small differences in performance at the extremes of this range. Of course, another possibility is that the empirical error variance reflects the sum of higher-magnitude, non-motor noise sources that were in included in the model. In the latter case, we would expect little gain from increasing stimulation frequency past 80 Hz. These results agreed with a rat model of ICMS detection, where increasing stimulation frequency up to but not beyond 80 Hz increased the probability of detection (Koivuniemi

and Otto, 2012).

We next experimentally tested the effect of electrode number on ICMS reliability. The instantaneous degradation experiments can't address this issue as any drop in performance could be due to the novelty of the degraded signal and not reflect any inherent limits to the information that can be encoded with a smaller number of electrodes. Therefore, we trained a monkey on a degraded ICMS signal(dICMS; four active electrodes whose preferred directions matched the cardinal directions) paired with a twenty-percent coherence visual dot-field. Each training session consisted of 66% VIS+dICMS trials and 33% dICMS-only trials. We then compared the behavioral performance before training (instantaneous degradation) and on the third training day (Figure 3.7; Tables 3.1-3.2).

With four active electrodes (encoding the cardinal directions), the monkey only completed 53.4% of trials during the instantaneous degradation experiments. Surprisingly, after only three days of training with a multisensory signal (20% visual coherence dot-field paired with the four electrodes the encoding cardinal directions), the monkey's performance recovers to 74.8% correct trials (Figure 3.7 and Table 3.1). Furthermore, he is now completing trials across the entire workspace without obvious systematic bias (compare panels c and d in Figure 3.7). The error variance for dICMS was even lower than ICMS using eight active electrodes (0.204 rad^2 after training vs. 0.136 rad^2 during instantaneous degradation trials), though this effect may reflect an ongoing re-stabilization of the IMCS percept as it was still higher than the error variance found for intact ICMS in Chapter 2 (0.09 rad^2). Pairing vision with the degraded signal not only improves understanding of the signal itself, but "recovers"

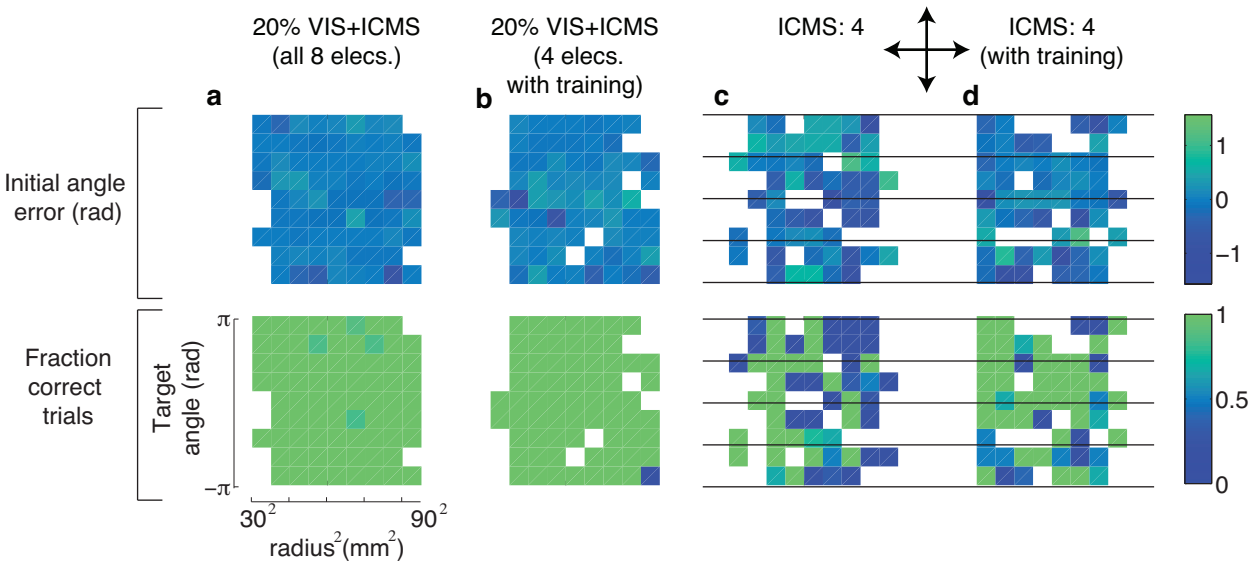

Figure 3.7: **Error bias by target radius and angle: before and after training.** Plots of initial angle errors and fraction correct trials as a function of reach target radius and angle. Top: error in initial movement angle, calculated as (initial movement angle - reach target angle). Bottom: fraction completed trials (of total attempted). Parameters (error and fraction correct trials) are calculated as a function of squared radius to maintain uniform sampling of points across space. Left half: comparison of performance on 20% coherence vision paired with (far left) all eight electrodes and (middle left) a subset of four electrodes, after training on the latter. Right half: comparison of performance on ICMS-only trials when only four electrodes are encoding the movement vector either instantaneously (middle right) or after training on paired vision and degraded ICMS (far right). Black lines overlaid on plots indicate the PDs of the active electrodes.

the multisensory signal. By the third recovery day, trials performed with 20% vision paired with four electrodes were completed at 99.1%, only half a percent lower than 20% vision paired with all eight electrodes (99.6%). Furthermore, the error variance in initial angle estimation fell to 0.16 rad^2, as compared to 0.2 rad^2 of ICMS with eight electrodes. These results are summarized in Table 3.1.

	Percent correct trials			
Condition	Instantaneous	# trials	After Training	# trials
20% coh. vision + 8 ICMS electrodes	99.6%	451	–	–
20% coh. vision + 4 ICMS electrodes	–	–	99.1%	213
8 ICMS electrodes	92.4%	92	–	–
4 ICMS electrodes (cardinal)	53.4%	103	74.8%	99
4 ICMS electrodes (ordinal)	39.6%	111	–	–

Table 3.1: Percent of correctly completed trials by task condition. We compare performance levels before (instantaneous) and after (after training) being trained on a degraded signal. The values in conditions for which we have no data are left blank.

	Error variance (rad^2)			
Condition	Instantaneous	# trials	After Training	# trials
20% coh. vision + 8 ICMS electrodes	0.202 (0.203)	450 (451)	–	–
20% coh. vision + 4 ICMS electrodes	–	–	0.159 (0.161)	211 (213)
8 ICMS electrodes	0.250 (0.316)	85 (92)	–	–
4 ICMS electrodes (cardinal)	0.136 (0.292)	55 (103)	0.204 (0.344)	74 (99)
4 ICMS electrodes (ordinal)	0.197 (0.849)	44 (111)	–	–

Table 3.2: Variance of initial angle errors (initial movement angle - target angle) by task condition. The first number in each column describes the variance of errors in correctly complete trials while the second, parenthesized, value refers to the variance of errors for all attempted trials. We compare performance levels before (instantaneous; left) and after (after training; right) being trained on a degraded signal. The values in conditions for which we have no data are left blank.

3.4 Discussion

Encoding an ICMS signal

The question then remains, will having more electrodes really help? We have some indication from this training experiment that the monkey substantially improves at using a four-electrode ICMS signal with training, but we didn't get a chance to determine the behavioral asymptote. Furthermore, the error variances during the instantaneous degradation to the four-electrode condition is strangely low—even lower than with the intact ICMS signal,

making it difficult to understand the data. But remember, that data was collected while the ICMS percept was unstable (Figure 3.1g), and in fact the training on the degraded input did not occur until after the experiments described in Chapter 4, which could have further destabilized the ICMS signal. If we instead compare the error variance at the endpoint of training with four electrodes to that at the end of training with eight active electrodes (at the end of Chapter 2), the number are closer to what one might expect: 0.09 rad^2 with eight electrodes and 0.2 rad^2 with four. If we believe these numbers, ICMS with four electrodes at that stage in training was about half as precise as ICMS with eight electrodes.

Although we don't know how precise the ICMS signal could ultimately become with only four electrodes, these results do tell us something about the decoding model that is being used. Consider that the Center-of-Mass decode precision improved by three to then times when going from four to eight stimulating electrodes (Figure 3.5). In contrast, for the Analytic decode, the precision only doubled in that range, strongly implying that this is the decoding algorithm used by the monkeys. The results of Figure 3.4 confirm this hypothesis, where the Analytic solution was slightly better at predicting the experimental pattern of errors than were either the Center-of-Mass or Winner-Take-All decodes. Strangely, the Analytic solution is not the most desirable outcome, as including noise in the decoding model makes it slightly worse than COM at estimating the true target angle.

Improving the reliability of artificial sensation

These results of the modeling study, however flawed in magnitude, give several suggestions for improving ICMS precision. Assuming a ANA decoding algorithm, it is easier (and more effective) to increase the stimulation frequency or amplitude than to increase the number of stimulating electrodes. Indeed, according to the models, increasing stimulation frequency from 50 to 100 Hz was more effective than increasing the number of stimulating electrodes from 10 to 20, although this improvement seemed to saturate quickly after 100 Hz. Similarly, increasing the number of activated neurons from 10 to 100 yielded more than a fivefold improvement in the asymptotic level of performance, implying that increasing stimulation amplitude would improve signal reliability as long as stimulating electrodes were placed at a distance that would prevent interference.

Stability of the ICMS percept

If the ICMS signal were stable across task conditions, performance levels should remain the same, regardless if the monkey reaches near vs. far, reaches right vs. left, or reaches ballistically vs. slowly. As we saw in Figure 3.6, the error in target angle estimation is impervious to variation in trial parameters for VIS+ICMS and ICMS-only with 8 electrodes, and only once the signal was degraded did systematic biases begin to appear. These results are important, as they say that the monkey was equally able to decode inputs at high (around 240 Hz for trials with the most distance targets) and low (around 100 Hz for trials with the closest targets) stimulation frequencies in decoding. Furthermore, as performance

was consistent across instructed delay periods (1-1.5 s; data not shown), the monkey was able to decode all of the necessary information within the first second of presentation. In the experiments from Chapter 2, we saw a similar independence of trial parameters for both monkeys, indicating that the ICMS percept was, indeed, stable across trial conditions.

On the other hand, the ICMS signal percept was sensitive to training conditions. The fraction of correct trials with intact ICMS is lower than that reported in Chapter 2, likely due to experiments involving instantaneous degradation of the ICMS signal, which actively harmed the ICMS percept (see Figure 3.1g). This sensitivity reflects an active mechanism of plasticity, whereby an estimate of the accuracy of a sensory sensory signal is updated more slowly than the instantaneous estimate of cue reliability (Zaidel et al., 2011). Although detrimental to performance in this particular experiment, plasticity of the ICMS signal and adaptation to ongoing events is a desirable attribute of the signal, as it better approximates natural sensation.

Chapter 4

Plasticity of Artificial Signals

4.1 Introduction

The utility of an artificial sensory signal depends on its ability to function like natural sensory signal. In Chapter 2, that meant that ICMS could serve as an informative signal on its own and that it was integrated with vision when both sensory inputs were available. But integration is only one component of normal multisensory processing: multisensensory systems are also constantly adapting, calibrating their estimates to one another and to the outside world (Simani et al., 2007; Zaidel et al., 2013). Such ongoing plasticity lets us react quickly and flexibly to external perturbations; without it, we wouldn't be able to recover from injury or other unexpected changes in our sensory or motor systems.

Sensory recalibration is thought to be driven by (at least) two forces: error-corrective learning and internal realignment (Cheng and Sabes, 2007; Izawa and Shadmehr, 2011; Shad-

mehr et al., 2010; Simani et al., 2007; Zaidel et al., 2013). Error-corrective learning describes adaptation in response to an external error signal, e.g. how close was the endpoint of your reach to the target you were reaching for? Internal realignment of signals describes adaptation in response to a misalignment of sensory cues, where the two cues are thought to adapt towards each other to reduce any discrepancies (Burge et al., 2010; Zaidel et al., 2011). The latter form occurs even in the absence of any active movements, in response to passive exposure to misaligned sensory cues (Cressman and Henriques, 2010; Zaidel et al., 2011).

The process of multisensory calibration is made apparent in experiments that impose an artificial shift between two sensory cues (Burge et al., 2008; Cressman and Henriques, 2010; Izawa and Shadmehr, 2011; Simani et al., 2007; Zaidel et al., 2013, 2011). A common form of this experiment involves putting prism goggles on human subjects, which shift visual feedback with respect to proprioception. At first, people make biased reaches in the direction of the integrated estimate, often having to make large corrections just to reach their target; however, humans can adapt to this shift on the order of just a few trials. To make accurate reaches under the imposed shift, the estimates of individual sensory cues must be re-calibrated. The dynamics of this process seems to depend both on the identity of the cue the is rewarded and the relative reliability of the two cues (Zaidel et al., 2013). Surprisingly, in many experiments involving adult animals or humans the two cues never achieve internal consistency—that is, they don't adapt all the way (Burge et al., 2010; Zaidel et al., 2011; visual inspection of Fernández-Ruiz and Díaz, 1999; Knudsen and Knudsen, 1990; Tseng et al., 2007).

A useful (naturalistic) artificial sensory feedback signal would react to an external per-

turbation in a similar fashion. If we impose a shift between vision and ICMS, we would expect sensory adaptation by the mechanisms of internal realignment and error-correction, the dynamics of which would depend on the relative reliability of our cues and the experimental paradigm (Zaidel et al., 2013). We can further expect that the signal would not fully adapt, instead saturating at some point.

The design of an experiment to test artificial adaptation should be guided by the following factors:

1. Multisensory integration was most "optimal" around the visual training coherence (twenty percent), so we should use this value in any experiment involving adaptation; at this level, ICMS and vision are are about equally reliable.

2. We should like to to test adaptation of the *ICMS signal*, so the experiment should be designed to maximize that possibility. One way to do so would be by setting vision as the rewarded cue and forcing ICMS to adapt.

Our experiment to examine the plasticity of the ICMS signal would then look as follows. The monkey performs reaches under two sensory modalities of equal reliability: vision and ICMS. After a block of baseline trials where vision and ICMS agree, a shift is imposed (IMCS is shifted relative to vision). The monkey will continue to make reaches using both vision and shifted ICMS, the the signal encoded by the visual cue will be the one that is "correct." To estimate changes in unisensory cues, we will interleave trials with vision-only and ICMS-only. To isolate the rate of error-based learning, we can further test the monkey on a block of ICMS-only trials.

Showing that the ICMS percept successfully adapts under the pressures of multisensory calibration would extend its potential uses from an artificial feedback signal for Brain-Machine Interfaces to a unique tool to study the sensorimotor processing and plasticity in the brain.

4.2 Methods

Monkeys performed the task described in this chapter after the instantaneous degradation study described in Chapter 3.

Behavioral Experiment

Subject and Implant

All animal procedures were performed in accordance with the National Research Councils Guide for the Care and Use of Laboratory Animals and were approved by the UCSF Institutional Animal Care and Use Committee. One rhesus macaque participated in this study. He had previously been chronically implanted with a 96-channel silicon microelectrode array coated with Iridium Oxide (Blackrock Microsystems) over left primary somatosensory cortices (Brodmann Areas 1, 2; S1).

Behavioral Task

The monkey had previously been trained to perform instructed-delay reaches to an unseen target in a two-dimensional virtual reality environment (Figure 4.2a), as described previously.

Each trial consisted of four epochs: i) The monkey moved his right hand (tip of the middle finger) to a fixed start position in the center of the screen (a visible circle of 10 mm radius). ii) An instructed delay period began (1-1.5 s) during which the target cue was presented but the monkey was prohibited from moving. The target cue encoded a movement vector between the monkeys current finger position and the center of the unseen reach target (12 mm radius) and was provided in the form of a visual dot-field (VIS), a multichannel ICMS signal (ICMS), or both (VIS+ICMS). Targets were selected uniformly from an annulus centered on the start target with an inner radius of 40 mm and an outer radius of 80 mm. iii) After a go cue, the monkeys made a reach under the guidance of continuously updating VIS, ICMS or VIS+ICMS feedback. iv) After acquiring the target the monkey received a liquid reward. Trials were terminated without reward if the monkeys moved too early during any of the delay intervals or if they failed to reach the target before a timeout (10 s).

Visual feedback The movement vector was encoded visually using a random moving-dot flow-field (dotfield) consisting of approximately 600 dots over the visual display. A percentage of the dots moved coherently together in the direction of the continuously updating movement vector (the dot-field coherence); the rest moved in random directions. All of the dots moved at the same speed, proportional to the length of the movement vector.

ICMS The monkey involved in this study had previously been trained to interpret the ICMS signal, and we use the same parameters of stimulation and encoding scheme as has been described previously (see Chapter 2 Methods).

Briefly: Intracortical microstimulation consisted of biphasic, charge-balanced pulse trains

delivered asynchronously to each of eight electrodes in the array. The movement vector was encoded in the spatial and temporal patterns of stimulation across the array. Direction was encoded by the relative stimulation pulse rates across the electrodes. First, each of the eight electrodes was arbitrarily assigned one of eight preferred directions (PD), equally spaced around the circle. Then, the stimulation pulse rate f_i of electrode i was calculated as a function of angle between the direction movement vector, θ, and the electrodes assigned PD,

$$f_i = \frac{f_0}{2}(1 + cos(PD_i - \theta)) \tag{4.1}$$

The frequency scaling factor, f_0, linearly encoded the movement vector distance, d, within the range of 100-300 Hz:

$$f_0 = \frac{d}{d_{max}} f_{max} \tag{4.2}$$

where d_max was 11.5 cm for Monkey F. The values of d and θ were continuously updated during the reach to provide online feedback.

Imposing a shift between ICMS and vision

The ICMS signal was "shifted" relative to vision by reassigning the preferred directions (PD) of each of the electrodes: $PD' = PD + \pi/4$. The change in encoding scheme effectively imposes a shift of $-\pi/4$ on ICMS: if under the new encoding scheme the set of electrodes were encoding a movement vector with direction zero, then under the old encoding scheme (which the monkey is using upon application of the shift) the electrodes would be encoding a movement vectors of (zero - $\pi/4$).

Experimental Design

In the experiments described below, vision is the rewarded cue whenever a shift is imposed, meaning that vision encodes the correct, "true" movement vector while ICMS encodes one that is shifted.

Multisensory integration Before conducting adaptation experiments, we confirmed that the monkey performed multisensory integration of vision and ICMS. To test multisensory integration, the monkey spent a behavioral session performing reaches using randomly interleaved VIS-only, VIS+ICMS, and ICMS-only trials (2:1:1 trial ratio). 100 correct trials into the behavioral session, a shift of -$\pi/4$ was imposed in ICMS. The shift remained for a series of four subsequent behavioral sessions.

Multisensory adaptation The previous experiment (testing multisensory integration) provided a baseline measure for individual cue estimates before adaptation. To test for adaptation of the ICMS percept, we exposed the monkey to shifted feedback over the course of 850 VIS+ICMS trials. No unisensory trials were presented during this time period. After training, we again tested unisensory percepts, randomly interleaving these trials with VIS+ICMS trials (2:1:1 VIS+ICMS:VIS:ICMS).

Estimating the time-course of adaptation To estimate the time course of adaptation, we rapidly switched between blocks of not-shifted and shifted trials (100 trials long). Each block consisted of randomly-interleaved VIS+ICMS, VIS-only, and ICMS-only trials ($5/7 : 1/7 : 1/7$). To analyze adaptation, we superimposed trials from each block in order of their appearance in the block (first ICMS-only trial, second ICMS-only trial, third ICMS-only

trial, and etc.), regardless of the true trial number.

Error-corrective learning To isolate error-corrective learning, we trained the animal on shifted ICMS-only trials (approximately 300), where vision (had it been present) would have been the rewarded cue.

Data Analysis

Behavioral Performance Measures. We quantified the monkey's ability to use the various sensory cues by the mean and variance of errors in the estimates of target angle, a parameter this monkey has been shown to estimate accurately (Chapter 2). Error (in rad) is calculated as: err = (initial movement angle - target angle), where initial movement angle is calculated as described below. This metric assess only the monkey's ability to interpret ICMS during the Instructed Delay period, and ignores any online corrections he may make. Furthermore, we only analyzed target estimates on trials the monkey successfully completes, as the trials are relatively easy to complete and failing to do so often reveals a lack of motivation that would only add noise to the data.

Initial movement angle was found by first calculating the number of movement sub-segments that composed a reach. Starting with the model assumption that sub-movements have bell-shaped velocity profiles we identified sub-movements by threshold crossings of the tangential velocity plot of a trajectory, with a threshold of 20% of the maximum velocity on a given trial. The angle between the starting point and the endpoint of the first sub-movement is then taken to be the initial movement angle.

Permutation tests: mean estimate different from zero To test if the single-modality estimates were different at each time point, we performed a two-sided permutation test (N = 10000) comparing each distribution with a normal distribution of the same variance but mean zero.

Time point of adaptation (long study) To determine when adaptation had taken place across a set of 1000 training trials, I found the trial number where the mean of a sliding fifty-sample window became (consistently) significantly different from the mean of the first fifty data points (post-shift baseline). Datasets were compared using a two-sided permutation test (N = 10000).

4.3 Results

Multisensory integration of vision and ICMS

In Chapter 2, two monkeys ultimately learned to performed optimal multisensory integration of vision and ICMS, at least during trials when the visual signal coherence approximated the value they used during behavioral training. We made these conclusions based on target estimation error variances and small biases, but did not test for integration via any experimental manipulations. Here, we confirmed in one monkey that multisensory integration takes place during paired VIS+ICMS trials by imposing a shift (see Methods) of $-\pi/4$ rad. between the movement vectors encoded in vision and ICMS (bottom panel of Figure 4.2b).

Imposing a shift visibly biased the monkey's initial estimate of target angle away from the

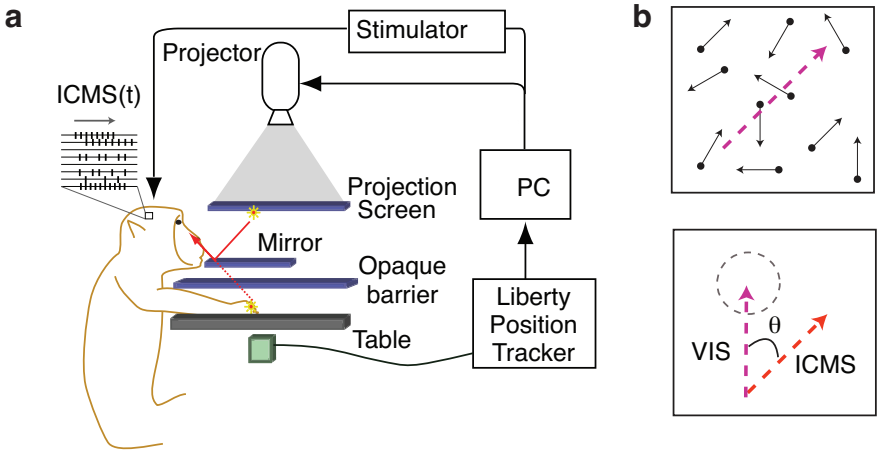

Figure 4.1: **Behavioral setup.** a) Behavioral setup. A monkey performs reaches in a virtual-reality environment guided by time-varying visual and ICMS feedback. b) Top: visual feedback encodes a movement vector (purple arrow) using a random dot-field at 20% coherence (20/100 dots move in the same direction; the rest move in random directions). Bottom: ICMS can be manipulated to encode a different movement vector than vision. In this example, the ICMS movement vector (orange arrow) is rotated counter-clockwise by $\pi/4$ (45 degrees) relative to the visual movement vector (purple arrow). Vision encodes the true location of the target (grey dashed circle).

true value, yielding curved movement paths (Figure 4.2a). To verify the consistency of this effect, we plotted the time course of VIS+ICMS angle estimation errors just before and after a shift was imposed (Figure 4.2b). While before the shift the monkey's errors centered on zero, after the shift the animal had a consistent negative bias, indicating integration of vision and ICMS. A low percentage of VIS+ICMS trials (only 1/4 of total trials) ensured that no adaptation took place during the course of this experiment, and a two-sided permutation test found no significant difference in mean values between the first twenty and last twenty points in the shifted data series ($p = 0.714$). Linear regression on each set of data (pre- and post-shift) also verified that there is no change is mean error over time (black lines, Figure 4.2b).

As the VIS+ICMS target angle errors showed no adaptation over a short timecourse,

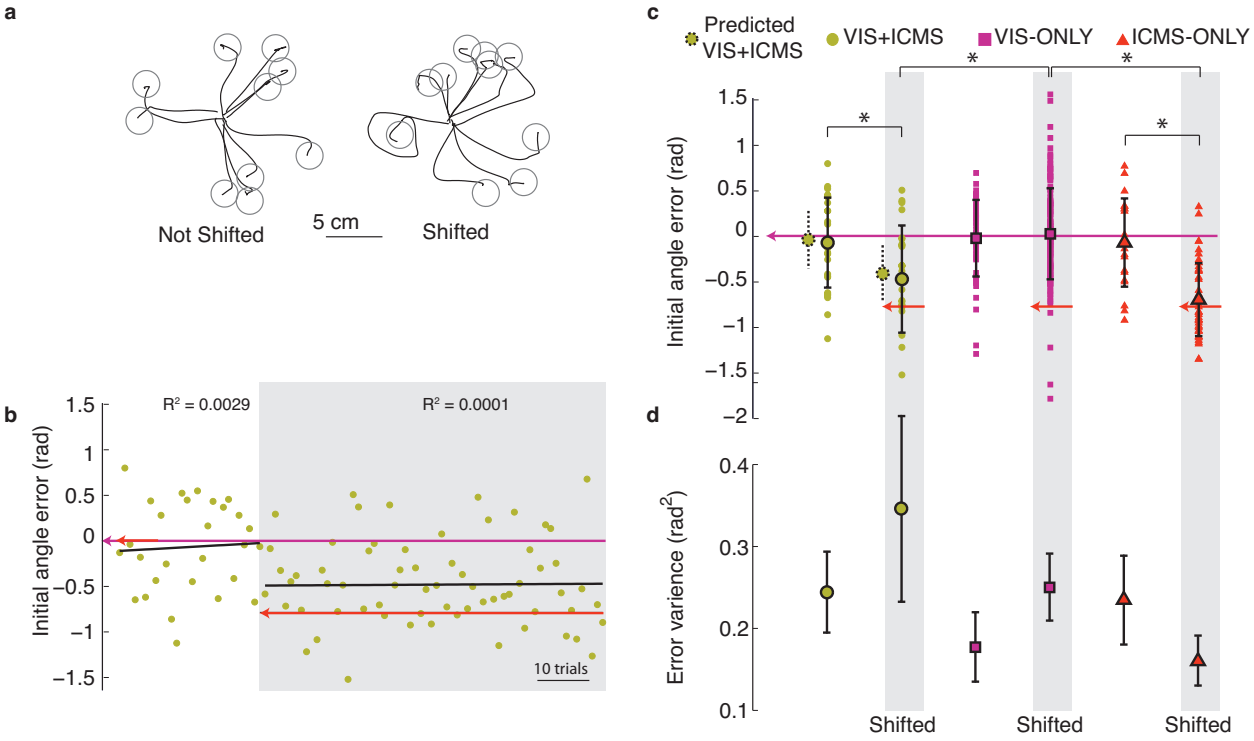

Figure 4.2: **Multisensory integration of vision and shifted ICMS.** a) Sample movement paths for VIS+ICMS trials when vision and ICMS agree (left: Not Shifted) and once a shift is imposed (right: Shifted). Target locations are shown here for clarity but were not present during the trials. b) Time-series plot of error in target angle estimation for VIS+ICMS trials (yellow circles) before a shift was imposed (left; white background) and after (right; gray background). The purple arrow indicates the expected error given the visual dot field (where zero error matches the true target location). The orange arrow indicates the expected error given the ICMS signal. Black lines and R^2 values depict linear regression on data. There was no significant difference in error between the beginning and end of the shift period depicted here. c) Comparison of sensory estimates before (white background) and after (gray background) a shift was imposed. Error bars denote standard deviation. Yellow circles surrounded by dashed lines display optimal integrated estimates (and std.) given unisensory estimates. Overlaid (blue curve) is the variance of errors in each condition (blue line) and estimates for optimal integration (dashed-blue circles). Unisensory trials (VIS-only, purple squares; ICMS-only, orange triangles) were randomly interspersed with the bimodal trials (VIS+ICMS, orange circles) plotted in (d). Data is averaged across each period separately. Arrows are as in (d). * indicates a significant differences as determined by permutation tests, p-values corrected for multiple comparisons. d) Variance in initial angle estimation. Data repeats what is plotted in panel (c).

we could average the time-series data across the two conditions to get an estimate of the accuracy (bias) and precision (inverse variance) of each of the sensory estimates (VIS-only, ICMS-only, and VIS+ICMS) before and after a shift was imposed (Figure 4.2c). Before the shift, the errors hovered around a mean of zero. After the shift, the visual estimate remained at zero but incurred an increase in variance (Figure 4.2d). Both the ICMS and VIS+ICMS estimates were significantly different after the shift ($p \leq 0.01$), moving closer towards the encoded ICMS signal ($-\pi/4$; orange arrow). They also became significantly different from the VIS estimate alone ($p \leq 0.01$).

Optimal multisensory integration (described in Equation 4.3) predicts that the monkey's estimate of target angle during VIS+ICMS trials ($\hat{\theta}_{VIS+ICMS}$) will be a weighted average of his estimates based on each unisensory input ($\hat{\theta}_{VIS}$ and $\hat{\theta}_{ICMS}$).

$$\theta_{VIS+ICMS} = \frac{\sigma_{VIS}^{-2}}{\sigma_{VIS}^{-2} + \sigma_{ICMS}^{-2}}\hat{\theta}_{VIS} + \frac{\sigma_{ICMS}^{-2}}{\sigma_{VIS}^{-2} + \sigma_{ICMS}^{-2}}\hat{\theta}_{ICMS} \tag{4.3}$$

Here, $\hat{\theta}_i$ and σ_i^{-2} represent the mean and variance for unisensory estimate i. The variance of this optimal estimate is calculated as:

$$\sigma_{VIS+ICMS}^{-2} = (\sigma_{VIS}^{-2} + \sigma_{ICMS}^{-2})^{-}1 \tag{4.4}$$

We compared the empirical values and the predicted, optimal, values of $\theta_{VIS+ICMS}$ and $\sigma_{VIS+ICMS}^2$ (Figure 4.2). Instead of explicitly calculating θ, we looked at the pattern of angle estimation errors across the workspace. The mean errors in VIS+ICMS (yellow circles) closely match the optimal estimate (yellow circles with dashed lines) both before and after the shift was imposed; however, the variance of the errors is much larger than would be

expected under optimal integration. Perhaps as a result of the increased visual variance, the monkey ended up weighing the visual signal less than would be expected from optimal integration ($w_{VIS,emp.} = 0.31$ vs. $w_{VIS,opt.} = 0.39$)—meaning that the monkey showed a small preference towards ICMS. Of course, given the variance of the empirical values, this distinction is likely artificial.

Multisensory adaptation

To drive sensory adaptation we trained the monkey on 850 VIS+ICMS trials with shifted ICMS, as shown in the time-series plot of errors in target angle estimation (Figure 4.3a). No unisensory trials were presented during the training period, but we took estimates of unisensory perception just before the shift (BS), just after the shift (AS), and after the extended period of training (T). Before a shift was imposed, all of the sensory estimates had distributions centered on zero (Figure 4.3; compare the three distributions on the far left with solid black line indicating zero error).

The mean visual estimate increases slightly from BS to T (Table 4.1), but the distributions don't seem to be significantly different from each other (p = [0.534, 0.126, 0.149]) or from zero (p = [0.682, 0.476, 0.076]; see Methods). As the data is very noisy, it is not clear if the visual estimate really remained at zero, as suggested by the statistical analysis, or if there was a small positive shift in the estimate.

In contrast, after extended training on VIS+(shifted ICMS), the mean estimates for both VIS+ICMS and ICMS-only have shifted partially towards vision at zero error (Figure 4.3b-c).

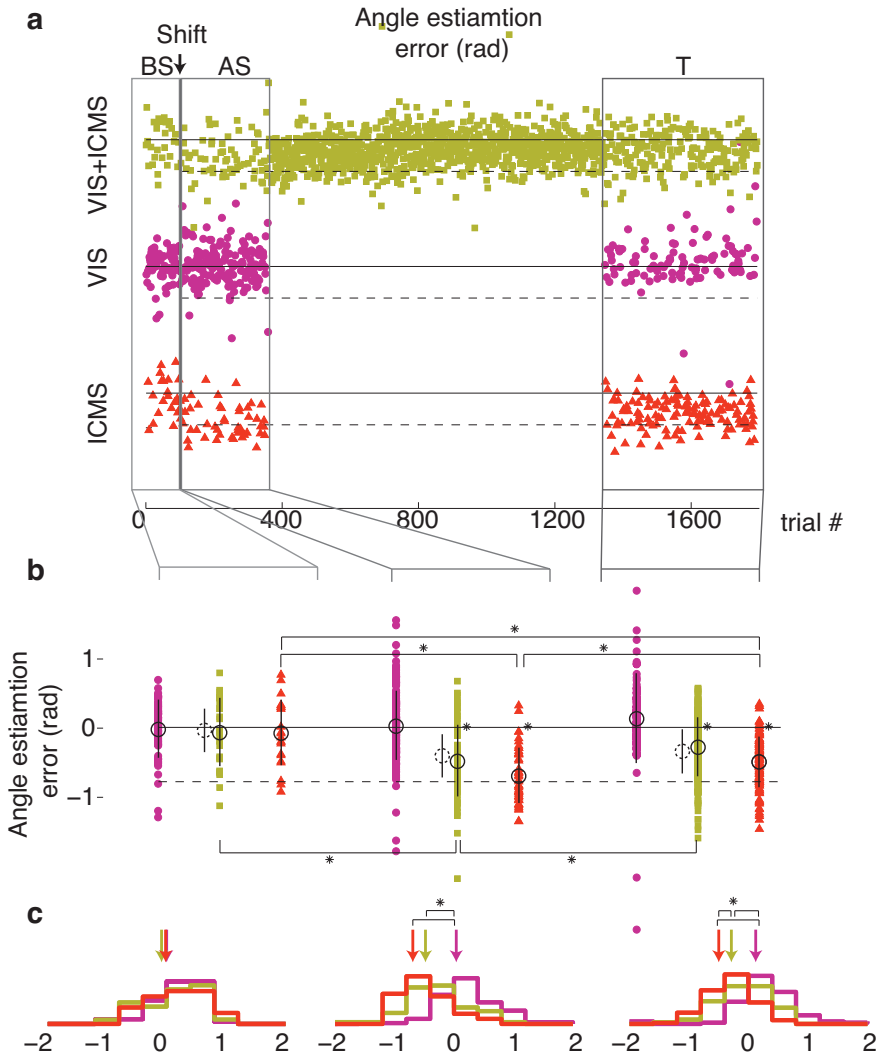

Figure 4.3: **Adaptation of the ICMS sensory percept.** a) Time-series data of angle estimation error for VIS+ICMS (top), VIS (middle), and ICMS (bottom); data is plotted by trial number. ICMS was shifted by $-pi/4$ at the point labeled as "Shift." BS: before shift; AS: after shift; T: after training on shifted data. VIS+ICMS data in boxes BS and AS was previously plotted in Figure 4.2. b) Summary of data shown in (a). Error bars denote standard deviation. VIS: purple squares; VIS+ICMS: yellow circles; ICMS: orange triangles. Dashed circles: minimum-variance estimate for VIS+ICMS. Solid black line at zero: rewarded estimates. Dashed black line at $-\pi/4$: imposed ICMS shift. Black stars near zero line: the mean of that distribution is significantly different from zero. The presence of a horizontal comparison bar denotes a significant difference between the means of the distributions. c) Distribution of means across sensory conditions at each time-point. Arrows denote center-of-mass of histogram. The presence of a horizontal comparison bar denotes a significant difference between the means of the distributions.

The adaptation in ICMS was not complete, extending to only 28.4% of the initial shift. The lack of total adaptation is not likely due to the noise of the sensory system, since the error estimates remained significantly different from zero-mean at time at time T ($p < 0.001$), but rather must reflect some underlying mechanism of adaptation.

Determining the time-course of adaptation proved difficult, as the variance in the VIS+ICMS was considerable and the difference in means before and after adaptation, though significant, was small (0.2 rad difference; Table 4.1). A naive attempt (described in the Methods section) found that a sliding window of fifty data points was significantly different from the starting point at around 500 trials after the shift takes place, approximately 150 trials after training commences. This estimate is rough, but implies that ICMS adapts to a new estimate in a relatively short time period, after which it saturates.

We next compared the relationship across sensory modalities within each time point (Figure 4.3). As discussed previously, the three sensory estimates were all zero-mean before the shift, and only ICMS and VIS+ICMS changed just after the shift. At AS, the three distributions were significant different from each other (Figure 4.3c), and this relationship persisted after ICMS had adapted (still present at time T). Furthermore, the near-optimality of the mean VIS+ICMS estimate persisted after training (Figure 4.3b, dashed circles just left of VIS+ICMS distribution); however, as at the previous time points, the empirical data had higher variance than would be expected under minimum-variance integration. The empirical visual weight increased to 0.34 after training, up from 0.31 just after the shift, and far above the optimal weight of 0.24.

	Error variance (rad^2)			Error mean (rad)		
Condition	BS	AS	T	BS	AS	T
VIS	0.177	0.250	0.423	-0.019	0.029	0.135
ICMS	0.225	0.160	0.134	-0.073	-0.694	-0.497
VIS+ICMS	0.244	0.268	0.182	-0.067	-0.481	-0.279

Table 4.1: Variance of initial angle errors (initial movement angle - target angle) by task condition. We compare performance levels before the shift (BS), just after the shift (AS), and after training on the shifted signal (T).

Our rough estimate of the time-course of adaptation (by 150 trials) implies that no further adaptation occurred in the remaining 700 training trials—that ICMS adaptation under bimodal (VIS+ICMS) exposure had reached saturation. To test if further adaptation could be driven by error-corrective learning, and to find the time-course and extent of that learning, we trained the animal on shifted ICMS-only trials. To our surprise, the ICMS signal failed to adapt any further under this training paradigm (Figure 4.4). The mean estimate was stable across days (compare T1 to T2, $p = 0.637$) and within a behavioral session (compare T1 to T3 $p = 0.880$; Figure 4.4a-b), and the pattern of errors had similar distributions at all three time points (Figure 4.4c).

To obtain a better estimate of the adaptation time course and dynamics (pictured in Figure 4.3), we repeated the shift experiment ten times, though on a shorter time scale—switching between blocks of 100 non-shifted trials and 100 shifted trials (Figure 4.5a). 5/7 of the trials in each block were VIS+ICMS, 1/7 were ICMS-only, and 1/7 were VIS-only. We superimposed the data from each block, separated by trial type (Figure 4.5b,c), and fit a linear (shown) and exponential (not-shown) regression to the data. Neither gave a particularly good fit (see R^2 for both linear and exponential fits on the plots)—not because

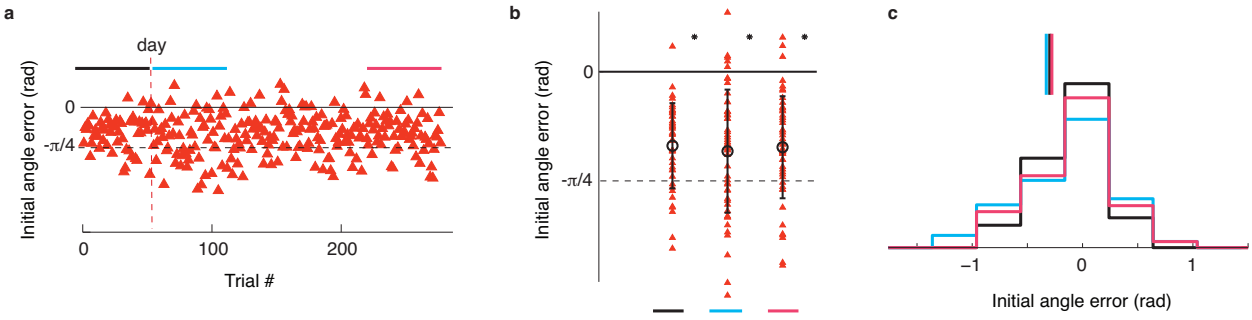

Figure 4.4: **ICMS adaptation under error-corrective learning** a) Time-series plot of angle estimation errors in ICMS-only trials. ICMS is shifted ($-\pi/4$;dashed black line) relative to vision. Data collection began after ICMS has previously adapted. Thick colored bars on top mark out different time points of interest in the data. Dashed red line marks a break separates data from two behavioral sessions. b) Summary of data from selected time points in (a). Error bars denote standard deviation. Black stars denote the mean of the distribution is significantly different from zero ($p =< 0.001$, two-sided permutation test). There are no significant difference between the means of the distributions ($p = [0.637,0.880,0.730]$, two-sided permutation test). c) Distribution of errors from three time points, overlaid.

the data is better fit by some other curve, but because it is static over this time period.

In fact, the mean error across the shifted and not-shifted blocks is remarkably stable (Figure 4.5d), in particular in the VIS+ICMS condition where we have more samples per block. The mean VIS+ICMS estimates in the not-shifted condition have a positive bias, reflecting the adaptation ICMS had previously undergone (Figure 4.3). Indeed, the VIS+ICMS bias equalled that of ICMS-only in the non-shifted condition (0.27 vs. 0.25 rad) and had an integrated estimate of -0.32 rad during the shifted condition. The ICMS bias observed during the non-shifted trials matched the expected value given the level of adaptation seen during shifted trials (36% during non-shifted vs. 38% adaptation during shifted). This level of adaptation was higher than what was observed during the earlier adaptation experiment, where the mean ICMS estimate had adapted to -0.50 rad. from an initial estimate at -0.69 rad., a total of 28% towards the rewarded direction at zero radians. Here, during the fast-

shift experiments, mean ICMS estimates were at -0.42 radians. Perhaps during the course of this experiments leading up to this (data not shown) ICMS had continued to adapt towards vision. The mean visual estimate was the same across shifted and unshifted conditions ($p = 0.430$), and seems to have incurred a positive bias (on average 0.12 rad), though the distribution of errors was not found to be statistically different from a normal distribution centered at zero with the same variance as vision.

4.4 Discussion

How does ICMS compare to natural sensory signals?

These results are noisy, and it is difficult to compare the details of ICMS signal plasticity to that of natural sensory inputs. If we compare the processes at a courser level, we can say that natural and artificial adaptation share some of the same characteristics. For example, the ICMS estimate does adapt in response to a imposed shift, though it adapts more slowly than natural sensation (humans performing a reaching task adapt on the order of tens rather than hundreds of trials; Fernández-Ruiz and Díaz, 1999; Taylor et al., 2014; Tseng et al., 2007). Our estimated adaptation timescale of hundreds of trials matches more closely the adaptation rate of monkeys passively exposed to misaligned input who have not yet fully adapted after 500 trials (Zaidel et al., 2011)[1]. Furthermore, the ICMS estimate adapted to a lesser extent than proprioception in those same monkeys (compare our 28.4% adaptation to a proprioceptive adaptation of 35% to 72% for monkeys) but, oddly, matched the extent

[1] Admittedly, this estimate could be a reflection of adaptation saturation and not the speed of adaptation.

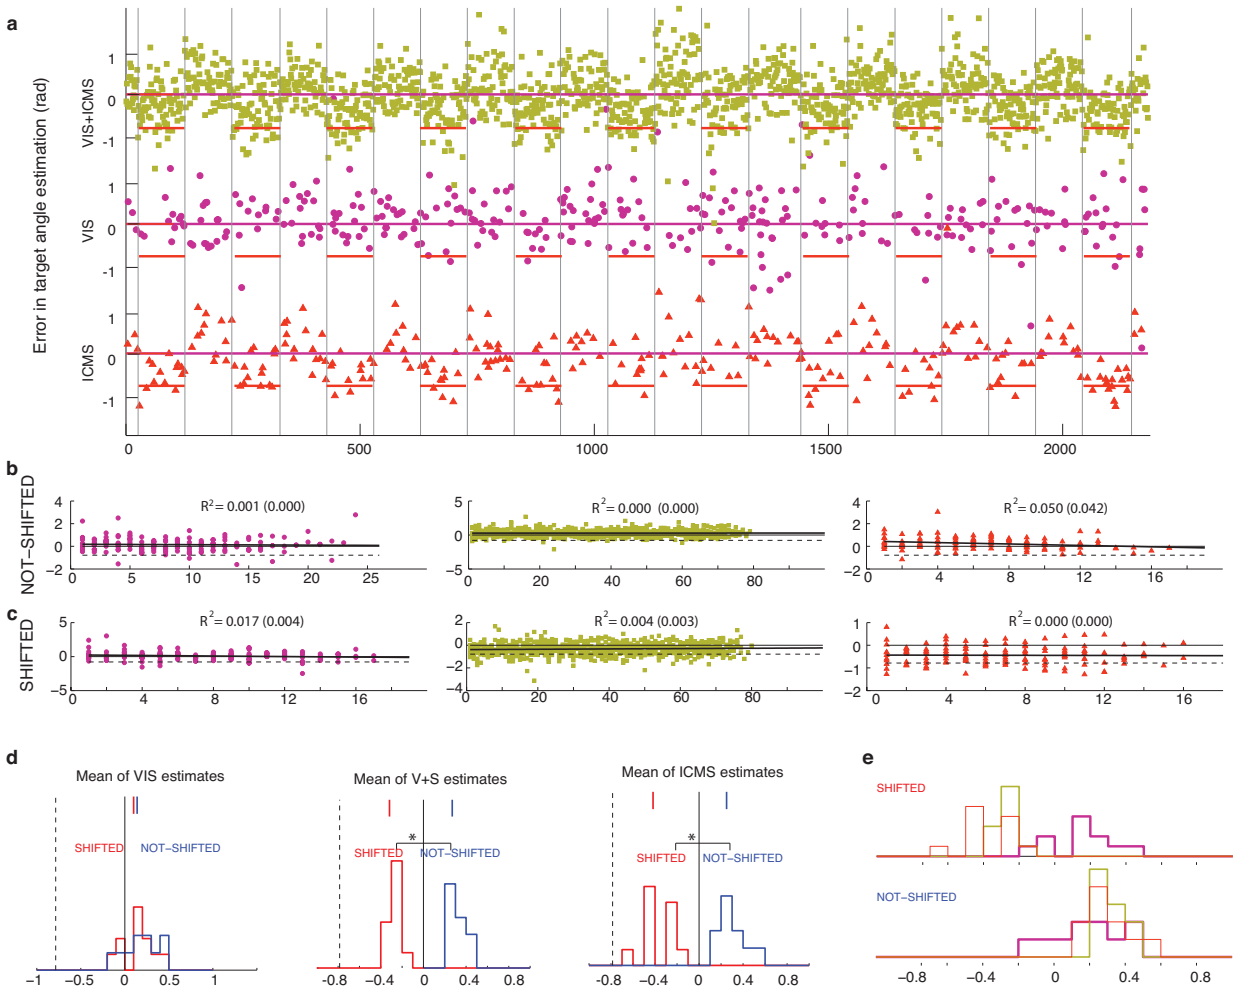

Figure 4.5: **Stability of adapted percepts.** a) Time-series data from VIS (middle), VIS+ICMS (top), and ICMS-only (bottom) trials. Orange lines mark out blocks of shifted trials. b) Superimposed data from all of the non-shifted blocks. Thick lack line: Linear regression; think black line: zero error; dashed thin black line: extent of shift ($-\pi/4$). R^2 values for linear and (exponential) fits to the data. c) Superimposed data from all of the shifted blocks. Thick lack line: Linear regression; think black line: zero error; dashed thin black line: extent of shift ($-\pi/4$). R^2 values for linear and (exponential) fits to the data. d) Distributions of unisensory means from not-shifted vs. shifted blocks. Top: Solid bars denote center-of-mass of histogram; dashed line: extent of shift ($-\pi/4$). $\star$ indicates $p < 0.001$, assessed by a two-sided permutation test. e) Distribution of means across sensory conditions in the top: shifted, and bottom: not-shifted conditions.

in humans performing the same unrewarded task (20% and 27% in humans; both monkey and human values are from Zaidel et al., 2011).

Natural and artificial adaption aren't wholly similar: unlike plasticity in natural sensory systems, we could not drive adaptation in response to error corrective feedback. There are three possible explanations for this result, which I'll list in order of increasing probability. The first possibility is that the result should be taken at face value—the artificial percept only adapts in response to misaligned inputs and does not change in response to an error signal in trials. Though this explanation is unlikely, we don't have any literature to draw upon in asserting that the ICMS signal should, in fact, adapt in response to trial error. The second possibility is that the negative result reflects a flaw in the experimental design: we tested for error-corrective learning after adaptation during VIS+ICMS trials had already saturated. It is plausible that, after saturation, no further adaptation could take place. If we had performed the same experiment just after the shift was first imposed, we may have had a better chance document error-corrective learning. The last possibility is that we saw no error-based adaption because the error signal in this behavioral task is very weak: the target was never explicitly shown and all feedback ceased as soon as the monkey reached the target. In this case, the monkey had to infer error from the length of his reach path and the difference between the direction of his initial movement and the direction of the rewarded endpoint (see Figure 4.2c for sample movement paths). A weak error signal would translate into a slow learning rate, so although error-corrective learning took place during training, no change could be detected over the 300 trials of the experiment.

Rapid switching experiments: a failure to adapt

Why did the ICMS signal fail to adapt on the shorter timescale experiments (100 trials)? According to our (admittedly rough) estimate, adaptation seemed to have occurred by 150 training trials. Then we should have seen at least a small change over the time-course of 100 trials, but only if the experiments were identical. One important difference between the two experiments was the relative ratio of VIS+ICMS to VIS-only and ICMS-only trials. We saw adaptation occur when the monkey was trained exclusively on VIS+ICMS trials. In contrast, only 5/7 of the trials were VIS+ICMS in the short-timescale experiments. Is it possible that having VIS+ICMS only 71% of the time fails to drive adaptation? Intuitively, single-modality trials should only improve the accuracy of the individual estimates. In VIS-only trials, error feedback should work to counteract any bias incurred from VIS+IMCS trials (see section below on Visual yoking). In ICMS-only trials, error feedback should further drive the ICMS percept towards the rewarded, clockwise direction. Instead, including these trials seems to forestall adaption—supporting our previous assertion that the error-feedback signal is weak or even absent in these trials.

Relative weighting of vision and ICMS during shifted input

The results of Chapter two showed that ICMS was equally informative as vision at twenty percent coherence; yet, when a shift between senses was imposed, the monkey gave a weight of 0.3 to vision and 0.7 to ICMS (assuming an affine combination). A lopsided weighing persisted throughout the adaptation experiment, including the experiments in Figure 4.5

involving adaptation on a short timescale, where the estimated visual weight was only $w_{VIS} = 0.21$. These weights are not unexpected, as they reflect the instantaneous variability of each of the sensory cues during the adaptation experiment, where vision proved more variable than ICMS once a shift was imposed (see Table 4.1), but they are inconsistent with our previous evaluation of relative vision and ICMS reliability.

Visual "yoking"

That vision was the less reliable signal in these experiments is inferable from the small positive bias incurred in its mean estimate during adaptation (0.14 rad after extended adaptation and 0.15 rad during the short-term adaptation experiments, though neither were signficant). These results are consistent with the the supervised model of inter-sensory calibration proposed by Zaidel et al., 2013 (Model 2 in Figure 1c), where the process of adaptation relies both on the relative reliability and accuracies of the two cues. In this model, the goal of adaptation is to correct the integrated estimate, not only the inaccurate cue. If the accurate cue is less reliable (e.g. vision in our experiment), the integrated estimate will lie closer to the more reliable cue (ICMS). Then in an attempt to rectify the integrated estimate, vision will be "yoked" (pulled) away from the correct answer, resulting in a transient positive bias. As discussed in the introduction, the process of adaptation is thought to consist of both inter-sensory calibration and error correction, so we would expect to see this result only if the rate of internal recalibration were lower than error-corrective learning and if adaptation saturated before the misalignment could be fully corrected. Of course, this requirement is

inconsistent with the data we have collected regarding error-corrective learning, as the visual mean after adaptation never became statistically different from zero, the positive bias may illusory.

Minimum-variance multisensory integration

Figure 4.2b reveals that although the mean multisensory estimates closely matched the optimal estimates, the variance of those estimates were greater, in all cases exceeding at least one of the single modality variances (see Table 4.1). Although these is no obvious explanation for this result, the adaptation studies were conducted after the instantaneous degradation experiments from Chapter 3, which we saw to de-stabilize his percept of IMCS. Perhaps the transient destabilization resulted in a long-term perturbation of minimum-variance integration.

Conclusions

Although the results presented here may frustrate the reader (and the experimenter) who desired to carefully characterize the process of artificial sensory adaptation, we think they at least qualitatively substantiate an important point: that an artificial sensory signal will adapt in response to a misalignment between itself and a natural sensory signal. That it does so validates its use both as supplemental feedback for neural prosthetics and as a model, however difficult, with which to study sensorimotor plasticity in the brain.

Chapter 5

Conclusions

This thesis work makes several significant contributions to the field of neural prosthetics, but we've left many questions unanswered and even unasked. After a brief summary of the contributions of this thesis (below), we will address some remaining topics of interest.

5.1 Contributions of this thesis

First, we showed that monkeys could learn to use, in a naturalistic way, a continuous, multi-dimensional, multi-channel intracortical microstimulation (ICMS) signal that encoded task-relevant feedback (Chapter 2). After training with correlated ICMS and visual feedback, monkeys could perform a goal-directed reaching task without vision, using only the ICMS signal. Furthermore, the animals learned to integrate the natural and artificial sensory inputs, combining both into a minimum-variance sensory estimate of hand position relative to the target. The ICMS signal described here was not proprioception in the traditional

sense, as hand position was described relative to a moving reference frame (the reach target), but by fixing the reference frame, this same signal could be applied to encode artificial proprioception for prosthetic limbs—in that case encoding the position of the limb relative to the body.

We also showed that the ICMS percept shared some of the mechanisms of plasticity of natural sensory signals observed in adults. First, we found that the brain keeps a long-term estimate of ICMS precision that can be degraded and restored, based on the quality of the ICMS inputs that the monkey receives (Chapter 3). Furthermore, we found that the ICMS signal would adapt in response to a perturbation (Chapter 4). When we shifted the ICMS signal relative to a rewarded, visual cue, the ICMS estimate adapted towards in response. These results highlight the suitability of an ICMS-based approach to studying neural circuits and systems.

Lastly, we found that the variables encoded in the multi-channel ICMS signal seemed to be decoded by downstream neurons using an algorithm that approximated the analytic solution to the problem (Chapter 3), and suggested several simple methods to maximize the precision of an ICMS signal.

5.2 Learning to use an artificial sensory signal

We showed animals can successfully learn to use a two-dimensional artificial ICMS signal in a naturalistic way, being able to use it in absence of additional sensory input, and being near-optimally integrated with vision. These results, summarized in Chapter 2, describe

the a monkey's ability to use a learning-based ICMS signal after an average of 65,000 trials spread over the course of months of behavioral training (n = 2). If an average trial lasted 2.5 s (1.25 s of which was instructed delay), then these monkeys were exposed to 45 continuous hours of paired vision and ICMS before they optimally integrated the two signals; however, monkeys began to show signs of integration (where VIS+ICMS trials outperformed VIS-only trials) maybe half-way through training, say after 22 hours and 30 minutes (see Chapter 2). Is that a lot or a little? It's hard to compare directly, but the timescale in not far what what is seen for the development of multisensory integration in adult superior colliculus (a common model system).

Multisensory integration seems to be a learned phenomena, absent in newborn animals and developed only after exposure to multisensory inputs (Xu et al., 2012). Although this normally happens during a plastic period early in life, the ability to develop integrative responses is maintained into adulthood (Yu et al., 2010). Yu and colleagues found that cells in the superior colliculus of dark-reared adult cats began to develop integrative responses to multisensory inputs after around 3,600 exposures (100 ms each; 6 minutes total). By 37,800 trials (1 hour and 3 minutes), cells already exhibited superadditivity—the cell's response to the multisensory signal was more than the sum of its response to each unisensory signal. Though an order of magnitude less than our timescales, the two are not inconsistent due to the low number of conditions the experimenters presented (only two). In contrast, our exposure signal could take a continuous range of values ($\theta \in [-\pi, \pi]$ rad; $dist. \in [0, 10]$cm), presenting a larger range of values that must be learned, and, as the number of integrative

neurons grow with multisensory exposure, a longer total training period.

These responses were contingent on the multisensory inputs being present simultaneously. If vision and audition were spatially (but not temporally) coincident, neurons would develop responses to each unisensory input but did not integrate the two (Yu et al., 2010). We can infer, then, that training on paired VIS+ICMS trials was necessary for the development of visual-ICMS integration in our monkeys. Had the monkeys learned to use the ICMS in absence of the visual dot-field, they probably wouldn't have been able to integrate, optimally or otherwise, vision and ICMS.

Remarkably, the experiments described above were conducted under anesthesia, meaning that cells in cat superior colliculus developed multisensory receptive fields and integrative capabilities under anesthesia—implying that the spatiotemporal correlations between inputs were both necessary and sufficient to drive learning. These results, that spatiotemporal correlations drive integration, agree with the predictions of a neural-network model that learns to integrate multisensory inputs (Makin et al., 2013)—the same model that inspired our learning-based approach.

5.3 The limits of ICMS-based signal precision

The suitability of learning-based ICMS as artificial sensory feedback for neural prosthetics depends very much on its ability to encode useful information—how many variables can we encode and with what precision? We concluded in Chapter 2 that the ICMS signal was as precise as vision at twenty to twenty-five percent visual coherence, but there are

several reasons to suspect that this value could ultimately be even higher. First of all, multisensory integration was closes to being optimal (minimum-variance) near twenty percent visual coherence—the visual signal with which ICMS was paired during training. We argued in Chapter 2 that, under a different training regime, integration would have looked optimal at all visual coherence levels. Then, if learning is highly influenced by the conditions of training, conceivably the ICMS signal ended up being only as precise as twenty percent visual coherence because it was the signal it was trained on.

We have some evidence that the reliability of the ICMS signal was matched with the reliability of it's partner visual signal: in Monkey D we saw an immediate drop in the number of trials the animal could complete at low coherence VIS+ICMS once we started training ICMS with fifteen percent visual coherence (see Chapter 2, Figure 2.4). His performance was rescued as soon as we returned to training with vision at twenty-five percent coherence. Perhaps the next step in improving the reliability of the ICMS signal will be simply to pair it with a high-precision visual signal.

5.4 How sensitive is the ICMS signal to design choices?

As we've discussed, learning plays an essential role in in interpreting artificial sensory feedback delivered using ICMS. But how important were the particulars (the design choices we made) of the artificial signal to the final outcome? For example, we delivered the ICMS

signal to primary somatosensory cortex (S1), a cortical area known to part of the pathway involved in multisensory integration (Sabes, 2011) that projects both to higher cortical areas as well as to motor cortex and the spinal chord. The cells in S1 encode touch sensation, position of the limbs (proprioception), and even the direction of objects moving across the skin (Krubitzer and Kaas, 1990; Pei et al., 2010)—the building blocks of somatosensation. Had the signal been delivered to a cortical area outside of this pathway, we might not have seen evidence of multisensory integration; had the signal been delivered to a cortical area further along the pathway, we might have succeeded only in disrupting ongoing sensory encoding.

In addition, although our signal was non-biomimetic in the traditional sense (we didn't try to replicate the spatiotemporal patterns of neural activity seen during natural sensory processing), we borrowed many concepts observed in natural sensory encoding. For example, by assigning each electrode a direction to encode (its preferred direction, PD), we end up imposing artificial tuning curves on on nearby neurons. Since stimulation frequency was proportional to the (circular) distance between the electrode's PD and the movement vectors directions, and stimulating at higher frequencies drives neurons to fire at higher frequencies (Butovas and Schwarz, 2003), each electrode effectively imposes a tuning curve on the set of nearby neurons.

Furthermore, since we have electrodes with unique preferred directions, stimulation across the set of electrodes—each of which drives its local population in proportion to its stimulation frequency—effectively established a population encoding of movement direction and distance. Looking across the population of neurons, one would see a gradation in firing rate that

is proportional to the cosine of the (circular) distance between the movement vector and the preferred direction imposed by the nearby stimulating electrode. Thus the population activity, though unfaithful to the native receptive fields of the stimulated electrodes, will look remarkably like a neural population encoding natural movement (see Figure 1, Pei et al., 2010). Although this sort of mimicry may not be necessary for artificial sensory feedback signal to be learnable, it might have helped speed the process.

5.5 Moving beyond a two-dimension signal.

Your (natural) arm provides many sensory feedback signals that would be equally useful for a robotic arm: at minimum we need information about touch (pressure and location) and a sense of the position and movement of the limb in space (of the fingers, the forearm, etc.). These two basis senses might be best encoded using different approaches, as they are essentially of different character.

Encoding touch information using ICMS

Touch is a widely distributed percept that might be best described a matrix of inputs that can be sparse (a single fingertip grazes a table) or full (your forearm rests on the table). To abstract this information into fewer variables would not be easy: we would have to encode, separately for each the body part, the extent and magnitude of the sensation. A biomimetic approach might be the most straightforward here, since the body is already conveniently mapped out by the receptive fields of neurons in somatosensory cortex (Kaas

et al., 1979); however, we can only sample a subset of that space with currently available electrode arrays. It may prove just as easy to take a learning-based approach to touch, assigning a "receptive field" to each electrode on a stimulating array. These receptive fields can vary in size depending on the part of the body they encode, having finer gradation for areas like the fingertips and sampling areas like the shoulder more coarsely.

Encoding proprioception using ICMS

Proprioception is a more tractable somatosensation to encode artificially. Although the sense of the body's position in space begins as a distributed signal at the periphery (e.g. the set of lengths of muscle spindles and joint capsules across the entire arm), it can abstracted into a lower dimensional space without loss of information. For example, hand position can be described by two variables ([x,y] or [r,θ]); joint angles (of the wrist, elbow, shoulder) by three. We can then feasibly use a learning-based approach (much like the one taken in this thesis) to encode all of the the necessary proprioceptive information.

Learning to use a high-dimensional artificial signal

One criticism of the learning-based approach is that although it is possible to learn a two-dimensional artificial signal, humans (and monkeys) would be unable to learn any high-dimensional variables because keeping track of many artificial variables is too high of a cognitive load. This unfortunate view misrepresents the neural operations underlying artificial processing, which likely rely on many of the same mechanisms of plasticity as natural

sensorimotor function. As an example of learning to effortlessly perform a complicated task, consider learning to play a sport (swinging a tennis racket or shooting a free-throw) or even learning to play a video game. As a concrete example, consider Super Mario 64. The first few times you want to execute a long jump with a punch at the end, you must look carefully at the controller and press the control stick plus Z plus the A button, and then press the B button once you've reached your enemy—quite a complicated task. If this state of cognitive attention to motor control persisted forever, we would never play video games—they would be far too tedious; happily, we quickly adapt and are able to transition from thinking "let's get that guy" to executing a jump and punch, without concentrating on the intervening steps.

The classical studies of Bach-y-Rita have shown that a similar learning procedure applies to adaptation to artificial sensory information being input at the periphery. In the first of his experiments, blind subjects were taught to detect visual objects using a sensory substitution system that translated video input from a camera into a matrix of tactile inputs on the subjects back. Users of this system had great success, learning to recognize faces, partially occluded objects, and shadows. Most importantly, they reported the external localization of stimuli, meaning that the objects they sensed seemed to come directly from the camera rather than having to interpreted from the tingling or touch on their backs (Bach-y Rita et al., 1969). We expect a similar outcome for artificial sensory feedback delivered using ICMS: subjects may *begin* by using a cognitive strategy, but ultimately the interpretation will be sub-conscious, allowing the user to focus on the task at hand rather than the array

of sensory information he receives.

5.6 Will learning-based ICMS ever feel "natural?"

Sensations evoked by electrical stimulation of somatosensory cortex

Learning-based approaches are commonly criticized for not being intuitive to use, as the percepts they evoke may not match the information they encode. This may be true, but, generally speaking, stimulation of the human somatosensory system (excluding the periphery) tends to produce "unnatural" percepts, regardless of stimulation patterns (Heming et al., 2010). These percepts may be naturalistic in a certain sense—they are "felt" on the parts of the body that are close or similar to the receptive fields of nearby neurons. Wilder Penfield and Edwin Boldrey (1937) are often cited for their famous paper reporting that electrical stimulation of the cortical surface produces sensations in body parts such as the fingers, hand, face, and etc., depending on the site of stimulation (see also Heming et al., 2010), and this statement is often interpreted to mean (and cited to support the notion) that stimulation of cortex produces naturalistic *sensations*. Instead, Penfield reports that the patient,

> ..when asked what the sensation was like, replied "like going to sleep." When asked it is was numbness or tingling, she replied "Both."

204 of the recorded responses in this patient felt like "tingling or electricity," 131 felt like numbness, and forty-nine felt like a movement was evoked (though none was observed; Penfield and Boldrey, 1937, pg. 433). Other, minor sensations (at relatively low frequency) included pain, the feeling of cold, a sense of blood rushing, thickness of the tongue and swelling of the tongue. A similar result is seen in other parts of the somatosensory circuit: although they were able to sometimes match the location of receptive fields of stimulated neurons, ICMS of human thalamus produced (predominantly) a "tingle" sensation, even when the patterns of stimulation matched those previously recorded in the thalamus (Hem-ing et al., 2010).

On the other hand, ICMS does seem to evoke natural sensations in animal experiments. ICMS has been shown to bias the perception of natural stimuli in the direction of the receptive field of the stimulated cell in both the visual (Salzman et al., 1990, 1992) and somatosensory (Tomilson et al., 2013) systems. Furthermore, ICMS has been successfully substituted for natural sensation in behavioral tasks: animals trained on natural sensation can (quickly) generalize to the new input (Otto et al., 2005; Romo et al., 2000, 1998; Tabot et al., 2013a; Tomilson et al., 2013; Venkatraman and Carmena, 2011).

Why should this difference in evoked sensation exist between ICMS in human and animals? One might think that it lies in the type of electrode that was used: a surface or macro-electrode used in humans might evoked a different percept than a micro-electrode used in animal studies. Indeed, in human thalamus, electrical stimulation with a micro-electrode evokes smaller projected fields—the size and location of the percept evoked by

IMCS—than with a macro-electrode, but the predominant sensation evoked by both was still "tingle" (Heming et al., 2010). Perhaps the reason the sensations evoked in the animal studies described above seemed to feel natural is that the percepts evoked by ICMS shared important characteristics with the natural signals (e.g. in human studies cited above there were similarities between the receptive fields of recorded neurons and the projected fields of ICMS) even if sensations themselves were not wholly the the same. In fact, in the one study which thoughtfully probed the matter, rats who transitioned smoothly between natural whisker deflections and ICMS of barrel cortex were able to differentiate between the two if they were rewarded for doing so (Venkatraman and Carmena, 2011).

Differences in neural activity patterns evoked by IMCS and natural input

> The current cannot impose upon neurons the spatiotemporally coded and integrated output they normally achieve; it can only drive them in bizarre and nonsensical synchrony. Thus any subtle, highly integrated neural effects resulting from the stimulation must ensue only because the neural systems downstream ... receiving the nonsense signal are able to transform it into an effective neural code.
>
> – Robert Doty, 1969

The differences between sensations evoked by natural sensory input and ICMS likely mirror differences between the evoked spatiotemporal patterns of activity (see Butovas and

Schwarz, 2003 for an example of ICMS-evoked neural activity). As Robert Doty stated, it then falls to the neurons downstream of the ICMS input to try to interpret the foreign patters of activity; perhaps, failing to do so, we then interpret the inputs as tingle. Of course, with practice animals improve at detecting these patterns—current thresholds for ICMS detection fall over time. But these improvements comes at the expense of detecting low-threshold visual stimuli at the same receptive field (Ni and Maunsell, 2010), as if the downstream neurons specialize in detecting certain patterns, however unnatural to begin with, that reflect the type of inputs they regularly receive. This thesis is supported by evidence that, before training, more "naturalistic" patterns of ICMS-evoked neural activity have lower thresholds for detection (Brugger et al., 2012). In this study, current-source density analysis was used to determine that activity patterns evoked by natural whisker stimulation most closely resembled those evoked by ICMS of layer IV neurons—the layer with the lowest current thresholds for detection.

Delivering naturalistic sensory information to the brain

The logical conclusion to this argument might seem to be that the only way to encode natural, intuitive information via electrical stimulation would be to precisely mimic natural patterns of activity. In fact, what I am arguing is that the term "naturalistic" merely reflects an endpoint: it's a descriptive term that compares the input a neuron is currently receiving to the distribution of inputs the neuron regularly receives. A brain can be trained to detect completely novel patterns of input (ala Ni and Maunsell, 2010), and we spend much of our

lives, especially during development, doing just that.

Training is even required to use the most successful neural prosthetic device—the cochlear implant. Patients who are implanted do not instantaneously regain audition; instead, users must practice with the device, relearning how to hear (e.g. Fryauf-Bertschy et al., 1992). We expect a similar training procedure will be required to interpret artificial somatosensory input: initially extraneous ICMS inputs will become meaningful over time. Importantly, these inputs won't simply be interpretable with attention: humans won't simply learn that the set of stimulation frequencies [5, 10, 54, 100] is encoding a hand position of [0,4], they will come to *feel*, instantaneously, that their hand is at [0,4].

Bibliography

Alais, D. and Burr, D. (2004). Ventriloquist Effect Results from Near-Optimal Bimodal Integration. *Current Biology*, 14:257–262.

Andersen, R. A. and Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual review of neuroscience*, 25:189–220.

Angeli, C. A., Edgerton, V. R., Gerasimenko, Y. P., and Harkema, S. J. (2014). Altering spinal cord excitability enables voluntary movements after chronic complete paralysis in humans. *Brain*, 137:1394–1409.

Bach-y Rita, P., Collins, C., and Saunders, F. (1969). Vision Substitution by Tactile Image Projection. *Nature*, 221.

Batista, A. P., Buneo, C. A., Snyder, L. H., and Andersen, R. A. (1999). Reach plans in eye-centered coordinates. *Science (New York, N.Y.)*, 285:257–260.

Battaglia-Mayer, A., Caminiti, R., Lacquaniti, F., and Zago, M. (2003). Multiple levels of representation of reaching in the parieto-frontal network. *Cerebral cortex (New York, N.Y. : 1991)*, 13(10):1009–22.

Battaglia-Mayer, A., Ferraina, S., Mitsuda, T., Marconi, B., Genovesio, A., Onorati, P., Lacquaniti, F., and Caminiti, R. (2000). Early coding of reaching in the parietooccipital cortex. *Journal of neurophysiology*, 83:2374–2391.

Berg, J. A., Dammann, J. F., Tenore, F. V., Tabot, G. A., Boback, J. L., Manfredi, L. R., Peterson, M. L., Katyal, K. D., Johannes, M. S., Makhlin, A., Wilcox, R., Franklin, R. K., Vogelstein, R. J., Hatsopoulos, N. G., and Bensmaia, S. J. (2013). Behavioral demonstration of a somatosensory neuroprosthesis. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 21:500–507.

Bishop, W., Chestek, C. C., Gilja, V., Nuyujukian, P., Foster, J. D., Ryu, S. I., Shenoy, K. V., and Yu, B. M. (2014). Self-recalibrating classifiers for intracortical brain-computer interfaces. *Journal of neural engineering*, 11:026001.

Bremner, L. R. and Andersen, R. A. (2012). Coding of the Reach Vector in Parietal Area 5d. *Neuron*, 75:342–351.

Brockmeier, A. J., Choi, J. S., Emigh, M. S., Li, L., Francis, J. T., and Principe, J. C. (2012). Subspace matching thalamic microstimulation to tactile evoked potentials in rat somatosensory cortex. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 2957–2960.

Brown, E. A., Ross, J. D., Blum, R. A., Nam, Y., Wheeler, B. C., and DeWeerth, S. P. (2008). Stimulus-artifact elimination in a multi-electrode system. *IEEE Transactions on Biomedical Circuits and Systems*, 2:10–21.

Brugger, D., Pidpruzhnykova, G., and Schwarz, C. (2012). Increasing the spatial resolution of cortical microstimulation. In *Annual meeting, Society for Neuroscience*, page 792.14.

Burge, J., Ernst, M., and Banks, M. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision*, 8:1–19.

Burge, J., Girshick, A. R., and Banks, M. S. (2010). Visual-haptic adaptation is determined by relative reliability. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30:7714–7721.

Butovas, S., Hormuzdi, S. G., Monyer, H., and Schwarz, C. (2006). Effects of electrically coupled inhibitory networks on local neuronal responses to intracortical microstimulation. *Journal of neurophysiology*, 96(3):1227–36.

Butovas, S. and Schwarz, C. (2003). Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings. *Journal of neurophysiology*, 90(5):3024–39.

Butovas, S. and Schwarz, C. (2007). Detection psychophysics of intracortical microstimulation in rat primary somatosensory cortex. *The European journal of neuroscience*, 25(7):2161–9.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., and Nicolelis, M. A. L. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1.

Chang, S. W. C. and Snyder, L. H. (2010). Idiosyncratic and systematic aspects of spatial representations in the macaque parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 107:7951–7956.

Cheng, S. and Sabes, P. (2007). Calibration of visually-guided reaching is driven by error corrective learning and internal dynamics. *Journal of neurophysiology*, 97:3057–3069.

Choi, J. S., Distasio, M. M., Brockmeier, A. J., and Francis, J. T. (2012). An electric field model for prediction of somatosensory (S1) cortical field potentials induced by ventral posterior lateral (VPL) thalamic microstimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20:161–169.

Chu, P., Muller, R., Koralek, A., Carmena, J. M., Rabaey, J. M., and Gambini, S. (2013). Equalization for intracortical microstimulation artifact reduction. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 245–248.

Cluff, T., Crevecoeur, F., and Scott, S. H. (2014). A perspective on multisensory integration and rapid perturbation responses. *Vision research*.

Crammond, D. J. and Kalaska, J. F. (1996). Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*, 108:45–61.

Cressman, E. K. and Henriques, D. Y. P. (2010). Reach adaptation and proprioceptive recalibration following exposure to misaligned sensory input. *Journal of neurophysiology*, 103(4):1888–95.

Dadarlat, M. C., O'Doherty, J. E., and Sabes, P. N. (2014a). A Learning-Based Approach to Artificial Sensory Feedback. In *Brain-Computer Interface Research: A State of the Art Summary 3*, chapter 4. Springer.

Dadarlat, M. C., O'Doherty, J. E., and Sabes, P. N. (2014b). A learning-based approach to artificial sensory feedback leads to optimal integration. *Nature Neuroscience*, page under review.

Dadarlat, M. C. and Sabes, P. N. (2015). Encoding and Decoding Artificial Sensation. In *IEEE EMBS Neural Engineering Conference*, page in preparation.

Daly, J., Liu, J., Aghagolzadeh, M., and Oweiss, K. (2012). Optimal space-time precoding of artificial sensory feedback through mutichannel microstimulation in bi-directional brain-machine interfaces. *Journal of neural engineering*, 9(6):065004.

Dangi, S., Orsborn, A. L., Moorman, H. G., and Carmena, J. M. (2013). Design and Analysis of Closed-Loop Decoder Adaptation Algorithms for Brain-Machine Interfaces. *Neural Computation*, 25:1693–1731.

Deneve, S., Latham, P. E., and Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature neuroscience*, 4:826–831.

Dobelle, W. H., Mladejovsky, M. G., and Girvin, J. P. (1974). Artifical vision for the blind: electrical stimulation of visual cortex offers hope for a functional prosthesis. *Science (New York, N.Y.)*, 183:440–444.

Donoghue, J. P., Nurmikko, A., Black, M., and Hochberg, L. R. (2007). Assistive technology and robotic control using motor cortex ensemble-based neural interface systems in humans with tetraplegia. *The Journal of physiology*, 579:603–611.

Doty, R. (1969). Electrical stimulation of the brain in behavioral context. *Annual review of psychology*.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap.* Chapman & Hall, Boca Raton.

Ernst, M. O. (2007). Learning to integrate arbitrary signals from vision and touch. *Journal of vision*, 7:7.1–14.

Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415:429–433.

Ethier, C., Oby, E. R., Bauman, M. J., and Miller, L. E. (2012). Restoration of grasp following paralysis through brain-controlled stimulation of muscles.

Fagg, A. H., Hatsopoulos, N. G., de Lafuente, V., Moxon, K. A., Nemati, S., Rebesco, J. M., Romo, R., Solla, S. A., Reimer, J., Tkach, D., Pohlmeyer, E. A., and Miller, L. E.

(2007). Biomimetic brain machine interfaces for the control of movement. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27:11842–11846.

Fernández-Ruiz, J. and Díaz, R. (1999). Prism adaptation and aftereffect: specifying the properties of a procedural memory system. *Learning & memory (Cold Spring Harbor, N.Y.)*, 6(1):47–53.

Fetsch, C. R., Pouget, A., DeAngelis, G. C., and Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience*, 15(1):146–154.

Fitzsimmons, N. a., Drake, W., Hanson, T. L., Lebedev, M. a., and Nicolelis, M. a. L. (2007). Primate reaching cued by multichannel spatiotemporal cortical microstimulation. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27(21):5593–602.

Fraser, G. W., Chase, S. M., Whitford, A., and Schwartz, A. B. (2009). Control of a brain-computer interface without spike sorting. *Journal of neural engineering*, 6:055004.

Fridman, G. Y., Blair, H. T., Blaisdell, A. P., and Judy, J. W. (2010). Perceived intensity of somatosensory cortical electrical stimulation. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 203(3):499–515.

Fryauf-Bertschy, H., Tyler, R. S., Kelsay, D. M., and Gantz, B. J. (1992). Performance over time of congenitally deaf and postlingually deafened children using a multichannel cochlear implant. *Journal of speech and hearing research*, 35:913–920.

Ghose, K. and Maunsell, J. H. R. (2012). A strong constraint to the joint processing of pairs of cortical signals. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(45):15922–33.

Gilja, V., Nuyujukian, P., Chestek, C. a., Cunningham, J. P., Yu, B. M., Fan, J. M., Churchland, M. M., Kaufman, M. T., Kao, J. C., Ryu, S. I., and Shenoy, K. V. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature neuroscience*, 15:1752–7.

Gomez-Rodriguez, M., Peters, J., Hill, J., Schölkopf, B., Gharabaghi, a., and Grosse-Wentrup, M. (2011). Closing the sensorimotor loop: haptic feedback facilitates decoding of motor imagery. *Journal of neural engineering*, 8(3):036005.

Graziano, M. S., Cooke, D. F., and Taylor, C. S. (2000). Coding the location of the arm by sight. *Science (New York, N.Y.)*, 290:1782–1786.

Gu, Y., Angelaki, D. E., and Deangelis, G. C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nature neuroscience*, 11:1201–1210.

Hatsopoulos, N. H. N., Mukand, J. M. J., Polykoff, G. P. G., Friehs, G. F. G., and Donoghue, J. D. J. (2005). Cortically controlled brain-machine interface. *2005 IEEE Engineering in Medicine and Biology 27th Annual Conference*.

Hauschild, M., Mulliken, G. H., Fineman, I., Loeb, G. E., and Andersen, R. A. (2012). Cognitive signals for brain-machine interfaces in posterior parietal cortex include continuous 3D trajectory commands.

Heming, E., Sanden, A., and Kiss, Z. H. T. (2010). Designing a somatosensory neural prosthesis: percepts evoked by different patterns of thalamic stimulation. *Journal of neural engineering*, 7(6):064001.

Hinton, G. E., Osindero, S., and Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural computation*, 18:1527–1554.

Histed, M. H., Bonin, V., and Reid, R. C. (2009). Direct activation of sparse, distributed populations of cortical neurons by electrical microstimulation. *Neuron*, 63(4):508–22.

Hochberg, L. R., Serruya, M. D., Friehs, G. M., Mukand, J. A., Saleh, M., Caplan, A. H., Branner, A., Chen, D., Penn, R. D., and Donoghue, J. P. (2006). Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442:164–171.

Houweling, A. R. and Brecht, M. (2008). Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*, 451(7174):65–8.

Huang, D., Lin, P., Fei, D.-Y., Chen, X., and Bai, O. (2009). Decoding human motor activity from EEG single trials for a discrete two-dimensional cursor control. *Journal of neural engineering*, 6:046005.

Ifft, P. J., Shokur, S., Li, Z., Lebedev, M. a., and Nicolelis, M. a. L. (2013). A brain-machine interface enables bimanual arm movements in monkeys. *Science translational medicine*, 5:210ra154.

Izawa, J. and Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS computational biology*, 7(3):e1002012.

Johansson, R. S. and Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature reviews. Neuroscience*, 10:345–359.

Johansson, R. S. and Westling, G. (1984). Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*, 56:550–564.

Johnson, L. a., Wander, J. D., Sarma, D., Su, D. K., Fetz, E. E., and Ojemann, J. G. (2013). Direct electrical stimulation of the somatosensory cortex in humans using electrocorticography electrodes: a qualitative and quantitative report. *Journal of neural engineering*, 10(3):036021.

Jones, L. a. (1994). Peripheral mechanisms of touch and proprioception. *Canadian journal of physiology and pharmacology*, 72(5):484–7.

Kaas, J. H., Nelson, R. J., Sur, M., Lin, C. S., and Merzenich, M. M. (1979). Multiple representations of the body within the primary somatosensory cortex of primates. *Science (New York, N.Y.)*, 204:521–523.

Kalaska, J. F. (1994). Central neural mechanisms of touch and proprioception. *Canadian journal of physiology and pharmacology*, 72:542–545.

Kalaska, J. F. (1996). Parietal cortex area 5 and visuomotor behavior. *Canadian journal of physiology and pharmacology*, 74:483–498.

Kalaska, J. F., Cohen, D. A. D., Prud'homme, M., and Hyde, M. L. (1990). Parietal area 5 neuronal activity encodes movement kinematics, not movement dynamics. *Experimental Brain Research*, 80:351–364.

Kao, J. C., Nuyujukian, P., Stavisky, S., Ryu, S. I., Ganguli, S., and Shenoy, K. V. (2013). Investigating the role of firing-rate normalization and dimensionality reduction in brain-machine interface robustness. In *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 293–298.

Kim, S.-P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., and Black, M. J. (2008). Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of neural engineering*, 5:455–476.

Kim, S. P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., Friehs, G. M., and Black, M. J. (2011). Point-and-click cursor control with an intracortical neural interface system by humans with tetraplegia. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19:193–203.

Knudsen, E. I. and Knudsen, P. F. (1990). Sensitive and critical periods for visual calibration of sound localization by barn owls. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 10:222–232.

Koivuniemi, A., Wilks, S. J., Woolley, A. J., and Otto, K. J. (2011). *Multimodal, longitudinal assessment of intracortical microstimulation.*, volume 194. Elsevier B.V., 1 edition.

Koivuniemi, A. S. and Otto, K. J. (2012). The depth, waveform and pulse rate for electrical microstimulation of the auditory cortex. *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, 2012:2489–92.

Krubitzer, L. A. and Kaas, J. H. (1990). The Organization Marmosets and Connections of Somatosensory Cortex in. *Journal of Neuroscience*, 10(3):952–974.

Levy, I., Schluppeck, D., Heeger, D. J., and Glimcher, P. W. (2007). Specificity of human cortical areas for reaches and saccades. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27:4687–4696.

Lewis, J. W. and Van Essen, D. C. (2000). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *The Journal of comparative neurology*, 428(1):112–37.

London, B. M., Jordan, L. R., Jackson, C. R., and Miller, L. E. (2008). Electrical stimulation of the proprioceptive cortex (area 3a) used to instruct a behaving monkey. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 16(1):32–6.

London, B. M., Torres, R. R., Slutzky, M. W., and Miller, L. E. (2011). Designing stimulation patterns for an afferent BMI: Representation of kinetics in somatosensory cortex.

Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference, 2011:7521–4.

Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9:1432–1438.

Makin, J. G., Fellows, M. R., and Sabes, P. N. (2013). Learning Multisensory Integration and Coordinate Transformation via Density Estimation. *PLoS Computational Biology*, 9.

Marzocchi, N., Breveglieri, R., Galletti, C., and Fattori, P. (2008). Reaching activity in parietal area V6A of macaque: Eye influence on arm activity or retinocentric coding of reaching movements? *European Journal of Neuroscience*, 27:775–789.

Marzullo, T. C., Lehmkuhle, M. J., Gage, G. J., and Kipke, D. R. (2010). Development of closed-loop neural interface technology in a rat model: combining motor cortex operant conditioning with visual cortex microstimulation. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 18(2):117–26.

McGuire, L. M. M. and Sabes, P. N. (2009). Sensory transformations and the use of multiple reference frames for reach planning. *Nature neuroscience*, 12(8):1056–61.

McGuire, L. M. M. and Sabes, P. N. (2011). Heterogeneous Representations in the Superior Parietal Lobule Are Common across Reaches to Visual and Proprioceptive Targets. *Journal of Neuroscience*, 31(18):6661–6673.

Merzenich, M. M., Kaas, J. H., Wall, J., Nelson, R. J., Sur, M., and Felleman, D. (1983a). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. *Neuroscience*, 8:33–55.

Merzenich, M. M., Kaas, J. H., Wall, J. T., Sur, M., Nelson, R. J., and Felleman, D. J. (1983b). Progression of change following median nerve section in the cortical representation of the hand in areas 3b and 1 in adult owl and squirrel monkeys. *Neuroscience*, 10:639–665.

Morgan, M. L., Deangelis, G. C., and Angelaki, D. E. (2008). Multisensory integration in macaque visual cortex depends on cue reliability. *Neuron*, 59(4):662–73.

Ni, A. and Maunsell, J. (2010). Microstimulation Reveals Limits in Detecting Different Signals from a Local Cortical Region. *Current Biology*, 20(9):824–828.

Nishimura, Y., Perlmutter, S. I., and Fetz, E. E. (2013). Restoration of upper limb movement via artificial corticospinal and musculospinal connections in a monkey with spinal cord injury. *Frontiers in neural circuits*, 7:57.

Novak, K. E., Miller, L. E., and Houk, J. C. (2002). The use of overlapping submovements in the control of rapid hand movements. *Experimental brain research*, 144(3):351–64.

O'Doherty, J. E., Lebedev, M. A., Hanson, T. L., Fitzsimmons, N. A., and Nicolelis, M. A. L. (2009). A brain-machine interface instructed by direct intracortical microstimulation. *Frontiers in integrative neuroscience*, 3:20.

Omrani, M., Diedrichsen, J., and Scott, S. H. (2012). RAPID FEEDBACK CORRECTIONS DURING A BIMANUAL POSTURAL TASK.

Orsborn, A. L., Moorman, H. G., Overduin, S. A., Shanechi, M. M., Dimitrov, D. F., and Carmena, J. M. (2014). Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, 82:1380–1393.

Otto, K. J., Rousche, P. J., and Kipke, D. R. (2005). Microstimulation in auditory cortex provides a substrate for detailed behaviors. *Hearing research*, 210(1-2):112–7.

ODoherty, J. E., Lebedev, M. a., Ifft, P. J., Zhuang, K. Z., Shokur, S., Bleuler, H., and Nicolelis, M. a. L. (2011). Active tactile exploration using a brainmachinebrain interface. *Nature*, 479(7372):228–231.

Pearson, R. C. and Powell, T. P. (1978). The cortico-cortical connections to area 5 of the parietal lobe from the primary somatic sensory cortex of the monkey. *Proceedings of the Royal Society of London. Series B, Containing papers of a Biological character. Royal Society (Great Britain)*, 200(1138):103–8.

Pei, Y. C., Hsiao, S. S., Craig, J. C., and Bensmaia, S. J. (2010). Shape invariant coding of motion direction in somatosensory cortex. *PLoS Biology*, 8.

Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain: A journal of neurology*.

Pons, T. P., Garraghty, P. E., Ommaya, A. K., Kaas, J. H., Taub, E., and Mishkin, M. (1991). Massive cortical reorganization after sensory deafferentation in adult macaques. *Science (New York, N.Y.)*, 252:1857–1860.

Prud'homme, M. J. and Kalaska, J. F. (1994). Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of neurophysiology*, 72(5):2280–301.

Recanzone, G., Merzenich, M., and Dinse, H. (1992). Expansion of the cortical representation of a specific skin field in primary somatosensory cortex by intracortical microstimulation. *Cerebral Cortex*, 2(3):181–196.

Redding, G. M. and Wallace, B. (2002). Strategic calibration and spatial alignment: a model from prism adaptation. *Journal of motor behavior*, 34:126–138.

Rincon-Gonzalez, L., Naufel, S. N., Santos, V. J., and Helms Tillery, S. (2012). Interactions between tactile and proprioceptive representations in haptics. *Journal of motor behavior*, 44:391–401.

Rolston, J. D., Gross, R. E., and Potter, S. M. (2009). A low-cost multielectrode system for data acquisition enabling real-time closed-loop processing with rapid recovery from stimulation artifacts. *Frontiers in neuroengineering*, 2:12.

Romo, R., Hernández, a., Zainos, a., Brody, C. D., and Lemus, L. (2000). Sensing without touching: psychophysical performance based on cortical microstimulation. *Neuron*, 26(1):273–8.

Romo, R., Hernández, A., Zainos, A., and Salinas, E. (1998). Somatosensory discrimination based on cortical microstimulation. *Nature*, 392(6674):387–390.

Sabes, P. N. (2011). Sensory integration for reaching : models of optimality in the context of behavior and the underlying neural circuits. *Progress in Brain Research*, 191.

Sainburg, R. L., Ghilardi, M. F., Poizner, H., and Ghez, C. (1995). Control of limb dynamics in normal subjects and patients without proprioception. *Journal of neurophysiology*, 73(2):820–35.

Sainburg, R. L., Poizner, H., and Ghez, C. (1993). Loss of proprioception produces deficits in interjoint coordination.

Salzman, C., Britten, K., and Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 364:174–177.

Salzman, C. D., Murasugi, C. M., Britten, K. H., and Newsome, W. T. (1992). Microstimulation in Visual Area MT : Effects Discrimination Performance on Direction. *Journal of Neuroscience*, 12(June):2331–2355.

Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., and Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442:195–198.

Scott, S. and Kalaska, J. (1997). Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *Journal of neurophysiology*, 77(2):826–52.

Scott, S. H., Sergio, L. E., and Kalaska, J. F. (1997). Reaching movements with similar hand paths but different arm orientations. II. Activity of individual cells in dorsal premotor cortex and parietal area 5. *Journal of neurophysiology*, 78:2413–2426.

Shadmehr, R., Smith, M. A., and Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience*, 33:89–108.

Shanechi, M. M., Hu, R. C., and Williams, Z. M. (2014). A cortical-spinal prosthesis for targeted limb movement in paralysed primate avatars. *Nature communications*, 5:3237.

Shanechi, M. M., Williams, Z. M., Wornell, G. W., Hu, R. C., Powers, M., and Brown, E. N. (2013). A Real-Time Brain-Machine Interface Combining Motor Target and Trajectory Intent Using an Optimal Feedback Control Design. *PLoS ONE*, 8.

Simani, M., McGuire, L., and Sabes, P. (2007). Visual-shift adaptation is composed of separable sensory and task-dependent effects. *Journal of neurophysiology*, pages 2827–2841.

Sober, S. and Sabes, P. (2005). Flexible strategies for sensory integration during motor planning. *Nature neuroscience*, 8(4):490–497.

Sober, S. J. and Sabes, P. N. (2003). Multisensory integration during motor planning. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(18):6982–92.

Stoney, S. D., Thompson, W. D., and Asanuma, H. (1968). Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current. *Journal of neurophysiology*, 31:659–669.

Suminski, A. J., Tkach, D. C., Fagg, A. H., and Hatsopoulos, N. G. (2010). Incorporating feedback from multiple sensory modalities enhances brain-machine interface control. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30:16777–16787.

Tabot, G. a., Dammann, J. F., Berg, J. a., Tenore, F. V., Boback, J. L., Vogelstein, R. J., and Bensmaia, S. J. (2013a). 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(45):18279–18284.

Tabot, G. a., Dammann, J. F., Berg, J. a., Tenore, F. V., Boback, J. L., Vogelstein, R. J., and Bensmaia, S. J. (2013b). 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.

Talwar, S. K., Xu, S., Hawley, E. S., Weiss, S. a., Moxon, K. a., and Chapin, J. K. (2002). Rat navigation guided by remote control. *Nature*, 417(6884):37–8.

Taylor, J. a., Krakauer, J. W., and Ivry, R. B. (2014). Explicit and Implicit Contributions to Learning in a Sensorimotor Adaptation Task. *Journal of Neuroscience*, 34(8):3023–3032.

Tehovnik, E. J. (1996). Electrical stimulation of neural tissue to evoke behavioral responses. *Journal of neuroscience methods*, 65(1):1–17.

Tehovnik, E. J. and Slocum, W. M. (2009). Depth-dependent detection of microampere currents delivered to monkey V1. *The European journal of neuroscience*, 29(7):1477–89.

Tehovnik, E. J., Tolias, A. S., Sultan, F., Slocum, W. M., and Logothetis, N. K. (2006). Direct and indirect activation of cortical neurons by electrical microstimulation. *Journal of neurophysiology*, 96:512–521.

Thomson, E. E., Carra, R., and a.L. Nicolelis, M. (2013). Perceiving invisible light through a somatosensory cortical prosthesis. *Nature Communications*, 4:1482.

Tomilson, T., Ruiz-Torres, R., and Miller, L. (2013). Multi-electrode stimulation in somatosensory area 2 induces a natural sensation of limb movement. In *Annual meeting, Society for Neuroscience*, page 835.03.

Tseng, Y.-W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R., and Bastian, A. J. (2007). Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *Journal of neurophysiology*, 98:54–62.

van Beers, R. J., Sittig, a. C., and Denier van der Gon, J. J. (1996). How humans combine simultaneous proprioceptive and visual position information. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 111(2):253–61.

van Beers, R. J., Sittig, A. C., and Gon, J. J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of neurophysiology*, 81:1355–1364.

van Beers, R. J., Wolpert, D. M., and Haggard, P. (2002). When feeling is more important than seeing in sensorimotor adaptation. *Current biology : CB*, 12(10):834–7.

van den Brand, R., Heutschi, J., Barraud, Q., DiGiovanna, J., Bartholdi, K., Huerlimann, M., Friedli, L., Vollenweider, I., Moraud, E. M., Duis, S., Dominici, N., Micera, S., Musienko, P., and Courtine, G. (2012). Restoring voluntary control of locomotion after paralyzing spinal cord injury. *Science (New York, N.Y.)*, 336:1182–5.

Venkatraman, S. and Carmena, J. (2009). Behavioral modulation of stimulus-evoked oscillations in barrel cortex of alert rats. *Frontiers in integrative ...*, 3(June):1–10.

Venkatraman, S. and Carmena, J. M. (2011). Active sensing of target location encoded by cortical microstimulation. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 19(3):317–24.

Voigt, B. C., Brecht, M., and Houweling, A. R. (2008). Behavioral detectability of single-cell stimulation in the ventral posterior medial nucleus of the thalamus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28:12362–12367.

Wallace, M. T., Carriere, B. N., Perrault, T. J., Vaughan, J. W., and Stein, B. E. (2006).

The development of cortical multisensory integration. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 26:11844–11849.

Wallace, M. T. and Stein, B. E. (1997). Development of multisensory neurons and multisensory integration in cat superior colliculus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 17:2429–2444.

Warren, J. P., Santello, M., and Tillery, S. I. (2011). Effects of fusion between tactile and proprioceptive inputs on tactile perception. *PLoS ONE*, 6.

Weber, D. J., Friesen, R., and Miller, L. E. (2012). Interfacing the somatosensory system to restore touch and proprioception: essential considerations. *Journal of motor behavior*, 44:403–18.

Westling, G. and Johansson, R. S. (1984). Factors influencing the force control during precision grip. *Experimental brain research*, 53:277–84.

Wise, S. P., Boussaoud, D., Johnson, P. B., and Caminiti, R. (1997). Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annual review of neuroscience*, 20:25–42.

Wolpaw, J. R. and McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 101:17849–17854.

Xu, J., Yu, L., Rowland, B. a., Stanford, T. R., and Stein, B. E. (2012). Incorporating cross-modal statistics in the development and maintenance of multisensory integration. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(7):2287–98.

Yeom, H. G., Kim, J. S., and Chung, C. K. (2014). High-Accuracy Brain-Machine Interfaces Using Feedback Information. *PloS one*, 9(7):e103539.

Yttri, E. a., Liu, Y., and Snyder, L. H. (2013). Lesions of cortical area LIP affect reach onset only when the reach is accompanied by a saccade, revealing an active eye-hand coordination circuit. *Proceedings of the National Academy of Sciences of the United States of America*, 110:2371–6.

Yu, L., Rowland, B. a., and Stein, B. E. (2010). Initiating the development of multisensory integration by manipulating sensory experience. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(14):4904–13.

Zaaimi, B., Ruiz-Torres, R., Solla, S. a., and Miller, L. E. (2013). Multi-electrode stimulation in somatosensory cortex increases probability of detection. *Journal of neural engineering*, 10(5):056013.

Zaidel, A., Ma, W. J., and Angelaki, D. E. (2013). Supervised calibration relies on the multisensory percept. *Neuron*, 80(6):1544–57.

Zaidel, A., Turner, A. H., and Angelaki, D. E. (2011). Multisensory calibration is independent

of cue reliability. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(39):13949–62.

Zanos, S., Richardson, A. G., Shupe, L., Miles, F. P., and Fetz, E. E. (2011). The neurochip-2: An autonomous head-fixed computer for recording and stimulating in freely behaving monkeys. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19:427–435.

Zimmermann, J. and Jackson, A. (2014). Closed-loop control of spinal cord stimulation to restore hand function after paralysis. *Frontiers in Neuroscience*, 8:87.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

Author Signature  Date 09/16/14