From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Psychology

2013

Experiments On The Ethological Organization Of Motor Cortex

Aflalo, Tyson

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp015712m6576

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

EXPERIMENTS ON THE ETHOLOGICAL ORGANIZATION OF MOTOR CORTEX

TYSON N.S. AFLALO

A DISSERTATION PRESENTED TO THE

FACTULTY OF PRINCETON UNIVERSITY

IN THE CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF PSYCHOLOGY

ADVISOR: MICHAEL S.A. GRAZIANO

November 2012

© Copyright by Tyson N.S. Aflalo, 2012. All rights reserved.

Abstract

This thesis explored several aspects of the cortical organization and representation of movement in monkey motor cortex. The experiments were based on the finding that long-duration stimulation applied to the monkey motor cortex evokes coordinated movements that can resemble actions from the monkey's behavioral repertoire (Graziano, Taylor, and Moore 2002). The thesis presents three experiments.

First, we used electrical stimulation of the monkey motor cortex to more carefully quantify the properties of stimulation-evoked movements. We found that arm movements evoked by stimulation share many properties with natural voluntary movements including bell-shaped velocity profiles, proper speed-distance scaling, inter-joint coordination to stabilize the position of the hand in space, and some ability to compensate for the effects of a weight attached to the hand. Movements evoked by stimulation were topographically organized on the cortical surface by the set of body parts recruited, the final position of the hand in space, and the type of behavioral action evoked.

Second, we used single neuron recording to study the tuning of single neurons in primary motor cortex during naturalistic arm movements. We found that the behavior of single neurons over the entire naturalistic dataset was best explained by tuning to a preferred posture of the arm. On average, the preferred posture of single units and the posture evoked by electrical stimulation at the same cortical site were in rough agreement. The neurons were also tuned to other aspects of movement. One possibility is that control signals represented by motor cortex neurons are multiplexed and tuned to

the varied features that characterize statistical biases in the monkey's natural motor behavior.

Third, we used computational modeling methods to study the principles of topographic organization in the motor cortex. We found that the overall topographic organization of motor cortex was largely explained by a model that mapped a high-dimensional representation of the behavioral repertoire onto the two dimensional cortical surface. We were thus able to provide a parsimonious account for many otherwise idiosyncratic features of cortical organization.

Acknowledgements

To T. Mole, J. McCain, S. Mixalot, and others…

Table of Contents

Abstract iii
Acknowledgments v

Chapter 1 : Introduction 1

Introduction Overview 1

Probing motor cortex with electrical stimulation 1

Single neurons and primary motor cortex 6

Topographic organization of motor cortex 8
Somatotopic organization of the primate motor cortex 9
Parcellation of the frontal motor areas into multiple motor areas 10

Complex movements evoked by long-duration stimulation 14
Probing function with long-duration stimulation 15
Cortical output shaped by sensory feedback 18
Ethological actions evoked by stimulation 20

Overview of dissertation studies 25

Chapter 2 : Arm movements evoked by electrical stimulation 28

Introduction 28

Materials and Methods 29
Surgery 29
Neuronal recording and stimulation 30
Location of stimulation sites 33
Measurement of limb position 34
Experimental Sessions 37

Results 41
Stimulation-evoked hand trajectories: basic description 41
Convergence of the hand toward a location 46
Effect of a weighted bracelet on the final stimulation-evoked height of the hand 49
Movement of individual joints 56
Interaction between joints stabilizes the hand for some stimulation sites 58
Movement categories arranged in separate cortical zones 60
Topography of hand location across the cortical surface 62

Discussion 65
Are the stimulation-evoked movements open-loop or closed-loop? 66
Motor cortex topography may be determined by multiple competing factors 67

Chapter 3 : Relationship between unconstrained arm movements and single neuron firing in the macaque motor cortex 68

Introduction 68

Methods 70
Surgery 70
Neuronal recording 70
Measurement of limb position 73
Description of movements in the data set 74
Preliminary analysis to specify somatotopic location 80
EMG recording 80
Regression models 81

RESULTS 86
Speed tuning 86
Direction tuning 90
End-point tuning 104
End-posture tuning 117
Comparison of end-posture tuning and end-point tuning 130
Neurons are more end-posture than start-posture tuned 132
Combined tuning: speed, direction, end point, and end posture 138
Simulated neurons 139
Muscle activity 141
Comparison Between Neurons and Stimulation 144

DISCUSSION 148
Interpretation of R^2 148
Local direction tuning and global posture tuning 150
Partial Match Between Neurons and Stimulation 152

Chapter 4 Possible origins of the complex topographic organization of motor cortex: reduction of a multidimensional space onto a 2-dimensional array .. 153

Introduction 152

Methods 157
Somatotopy 158
Hand Position 160
Ethological Category 161
Input Vectors 168
Learning Rules 169

Results **171**
Somatotopic Organization 171
Clustering of Ethological Function 172
Final Hand Positions 173
Alternate Maps 175
Lesions and use-dependant changes 179
Discussion **183**

Chapter 5 : Concluding Comments **189**

Chapter 6 : References **199**

Chapter 1 : Introduction

Introduction Overview

What aspects of movement are represented in the primate motor cortex and how are these movement representations arranged on the cortical surface? Does motor cortex directly control individual muscles, does it control more abstract features of movement such as direction of hand movement in space, does it specify higher-level motor objectives, or does it control movement in some other manner? Scientifically, these questions have been addressed in two primary ways: by electrically stimulating the motor cortex and observing the resulting movements or by correlating the activity of individual neurons with various aspects of motor behavior. The first two sections of this introduction provide a basic overview of how these two methodologies have advanced our understanding of the cortical representation of movement. The third section looks at various accounts of how these motor representations are topographically arranged on the cortical surface. The fourth section reviews work performed in our lab using long-duration stimulation to evoke complex behavioral actions from the monkey's motor cortex. The fifth section contains an overview of the studies contained in this thesis.

Probing motor cortex with electrical stimulation

Electrical stimulation as a tool to probe cortical function began when Fritz and Hitzig

(1870) applied galvanic stimulation, comprised of brief electrical pulses from a battery, over the surface of the dog's brain. Expanding on this initial discovery, Ferrier applied faradic stimulation, a form of alternating current generated by an induction coil, to a variety of animal species including macaque monkeys. The use of galvanic stimulation by Fritz and Hitzig limited the duration of stimulation to brief pulses effectively inducing muscular twitches. The use of faradic stimulation allowed Ferrier to stimulate for longer durations at higher amplitude without damaging the stimulated region of cortex. Such high-intensity long-duration stimulation evoked complex coordinated movements that could resemble movements in the animal's natural repertoire or in Ferrier's (1874) words "Many of the movements (evoked by stimulation) such as those of the hands, the legs, the facial muscles and the mouth have the aspect of purpose volition and are of the same nature as those which the animal makes in its ordinary intelligent actions."

The exploration of complex naturalistic movements was largely abandoned after Ferrier. Instead, subsequent researchers applied stimulation for shorter duration or at lower amplitude to evoke simple more isolated movements about one or more joints. The modus operandi of the time was cortical localization. By application of "minimal stimulation" researchers were able to identify the "Primary Movement … which is represented above all others at one particular spot in the cortex" (Beevor and Horsley 1887) and thus generate more fine-grained descriptions of movement representations on the cortical surface. Following the work of Campbell (1905), the motor regions of the frontal lobe were partitioned into a premotor subdivision that was thought to be responsible for movement planning and a primary motor subdivision thought to be

responsible for movement execution (See section "The topographic organization of motor cortex"). Subsequent to this partitioning, the use of electrical stimulation to probe cortical output was largely limited to primary motor cortex and thus primary motor cortex will be our focus for the remainder of this section.

The drive for ever finer accounts of cortical localization was largely influenced by the pioneering work of Leyton and Sherrington (1917). Sherrington, whose "main purpose being 'localisation' of the primary movement" used stimulation "not far above threshold value" to map movement representations in motor cortex. They found that "the individual movements, elicited by somewhat minutely localised stimulations, are, broadly speaking, fractional, in the sense that each, though co-ordinately executed, forms, so to say, but a unitary part of some more complex act, that would, to attain its purpose, involve combination of that unitary movement with others to make up a useful whole." In other words, the movements evoked by threshold stimulation form a basis-set of "elementary movements," each localized to a spot on cortex, that can be combined in a variety of ways to produce a diverse set of complex ethologically meaningful actions. Different actions are the result of different spatial and temporal combinations of the "large, though exhaustable, number" of elementary movements.

Following Sherrington, researchers sought to use ever more refined measurement and stimulation techniques to more precisely determine how isolated an affect stimulation can have on the motor periphery (e.g. Chang, Ruch, and Ward 1947, Landgren and Phillips 1962). The logic was that if stimulation could evoke activity in only a single muscle, than the basic element

of the motor repertoire, the "elementary movement", is the activation of a single muscle. Following the inventions of electromyograms, a method to record electrical activity from muscles, and intracortical microstimulation, a method to inject current through fine microwires inserted into cortex (previous investigations applied stimulation to the cortical surface leading to more diffuse spread of current), researchers could inject single or brief trains of microsecond long pulses of current into the projecting layers of cortex and carefully measure minute changes in muscle activity. Employing such a technique, Asanuma and Rosen (1972) concluded that, at least for the representation of the distal musculature of the hand, "there exists a finely grained mosaic organization within the depth of the motor cortex. Each efferent zone is confined to a small cortical volume and projects to a specific motoneuron pool." This represented a particularly extreme view and subsequent authors employing threshold stimulation report substantial overlap of the arm and hand representations on the cortical surface (e.g. Donoghue, Leibovic, and Sanes 1992, Kwan et al. 1978).

Fetz and Cheney (1980) devised a clever method to look at the descending influence of single cortical neurons on the motor periphery without the use of stimulation. They simultaneously recorded the spiking activity of single neurons in primary motor cortex and the activity of muscles in the hand and wrist. By averaging muscle activity time-locked to each neuronal spike (a technique called spike-triggered averaging), they were able to determine the relatively direct influence of motor cortex neurons on the recorded set of muscles. The majority of individual neurons affected multiple muscles. This included coactivation of antagonist muscle

pairs, activation of an agonist and suppression of its antagonist, or more idiosyncratic combinations of facilitation and suppression (Cheney, Fetz, and Palmer 1985, Fetz and Cheney 1980). McKiernan et al. (1998) recorded from a larger set of muscles that crossed the hand, wrist, elbow, and shoulder, and found that single neurons can affect muscle activity of both the arm and hand and are thus able to more holistically organize arm and hand movements. The notion that primary motor cortex controlled individual muscles seemed increasingly untenable.

Any lingering debates about distinct cortical representations of muscles have been settled by the recent studies of Rathelot and Strick (2006, 2009). Using a retrograde-virus that was transported from the injection site in muscles to motor cortex, they were able to study the topographic organization of cortical neurons that project directly to motor neuron pools. The cortical representation of the distal musculature of the hand was completely overlapping on the cortical surface. This provides clear evidence that manipulating stimulation parameters to produce the most minute affect possible on the motor periphery (e.g. by limiting current amplitude and stimulation duration) will provide an inaccurate assessment of movement representations on the cortical surface. This topic will be reintroduced in the section "Probing function with long-duration stimulation."

This section focused on how stimulation has been used to understand the influence of a localized patch of cortex on the motor periphery. The historical progression from Ferrier to the work of Asanuma can be viewed as a progression from the idea that a small cortical patch initiates complex actions, to the idea that a cortical patch produces individually meaningless

movements, to the idea that a cortical patch directly activates individual muscles. While the idea that cortex directly controls muscles has persisted, further experimentation and innovations in methods have shown that localized patches of cortex and even single neurons make divergent projections to many muscles and are thus in a position to more holistically control movements of the body.

Single neurons and primary motor cortex

So far this thesis has reviewed the results of electrical stimulation applied to motor cortex. An altogether different approach to understanding the cortical motor system is to correlate the activity of single neurons with various aspects of motor behavior. Much of the work in primary motor cortex has centered on the question of whether single neurons encode high-level aspects of motor behavior (e.g. "movements") or low-level motor output variables such as muscle activation or force. Initial single neuron recordings suggested that primary motor cortex primarily encoded muscle activation or joint force (Evarts 1968). Subsequent work by Georgopoulos and colleagues suggested that motor cortex encodes movement at a relatively high-level specifying a desired direction of hand motion in space (Georgopoulos et al. 1982, Georgopoulos, Schwartz, and Kettner 1986). One complication is that the preferred direction of a neuron is not a stable quantity. For example, the preferred direction is known to change with the starting position of the hand (Caminiti et al. 1991), configuration of the arm (Scott and Kalaska 1995), and even hand speed (Churchland and Shenoy 2007). One possibility is that

what on the surface appears to be tuning to the direction of the hand in space is simply the consequence of the highly correlated nature of motor variables. For example, muscles activations also appear tuned to direction of the hand in space (Kurtzer et al. 2006). There has been a recent resurgence in the idea that primary motor cortex encodes low-level motor parameters (Ajemian et al. 2008, Holdefer and Miller 2002, Morrow, Pohlmeyer, and Miller 2009, Todorov 2000). However, when tasks are employed that allow for a careful dissociation of more high-level movement-like versus low-level muscle-like representations, single units in primary motor cortex show tuning to both (Kakei, Hoffman, and Strick 1999, Umiltà et al. 2008). For example, Kakei et al. (1999) tested the tuning of single M1 neurons during movements of the wrist as they systematically varied the orientation of the forearm. In this way, a movement of the wrist upwards in space would involve flexion of the wrist when the arm was in a supinated posture or extension of the wrist when the forearm was in a pronated posture. They found that roughly half of the responsive neurons fired most when the wrist moved upwards in space independent of whether a flexion or extension motion was necessary to achieve this movement direction.

It does not seem that any single tuning function can describe the varied behavior of neurons in motor cortex. Cortical motor representations likely reflect idiosyncratic aspects of movements that have proved especially useful in the animal's previous motor experience (e.g. Huber et al. 2012, Kleim, Barbay, and Nudo 1998, Martin, Engber, and Meng 2005, Nudo et al. 1996). For instance, Nudo et al. (1996) observed that "movement combinations that were used

more frequently after training were selectively magnified in their cortical representations." Electrical stimulation applied to the cat motor cortex will ordinarily evoke multi-joint movements. However, if the motor behavior of a cat is restricted in early infancy to be devoid of multi-joint movements, stimulation will evoke more isolated movements about a joint (Martin, Engber, and Meng 2005). These results suggest that cortical motor representations come to reflect statistical biases in the animal's natural behavior. Along these lines, it is interesting to note that the majority of natural primate motor behavior is spent performing small movements about a relatively fixed posture (Graziano et al. 2004). Do single neuron properties during naturalistic movements reflect this statistical bias in primate behavior? Previous studies use extensively trained monkeys working in restrictive behavioral paradigms. Do the observed neural tuning functions of motor cortex reflect this extensive training? Would the same tuning functions still be observed in untrained monkeys performing unconstrained naturalistic movements?

Topographic organization of motor cortex

Thus far we have looked at various accounts of how motor cortex influences the motor periphery. In this section we turn our focus to how motor representations are topographically arranged on the cortical surface.

Somatotopic organization of the primate motor cortex

Beevor and Horsley (1887) were the first authors to perform a careful examination of the cortical representation of movements evoked by threshold stimulation and in so doing were the first authors to generate a clearly delineated map of the body arranged across cortex – a somatotopic map. This somatotopy is characterized by a vertical stacking of body parts with representations of the face in a ventral position progressing to representations of the hand, then the upper arm, and finally of the legs in a dorsal position (Figure 1-1A.) This aspect of topographic organization is the most consistently reported organizational feature of motor cortex (for an example in the human brain see Figure 1-1B.) The basic dorsal-ventral organization in the primate was expanded by Woolsey et al. (1952) to include a posterior to anterior progression of distal to axial musculature (Figure 1-1D). They further describe a complete re-representation of the body located medially in the hemisphere, anterior to the precentral leg representation, that they equated with Penfield's "supplementary motor area" (Penfield and Welch 1951). At least three additional complete somatotopic representations of the body have been described within the medial wall of the hemisphere in and around the cingulate sulcus (Dum and Strick 1991, He, Dum, and Strick 1995, Matelli, Luppino, and Rizzolatti 1991). Importantly, the schematic depictions represented in figures fail to capture the extensive overlap of motor representations. For example, Woolsey states: "It must be emphasized, however, that (the motor map) is an inadequate representation of the localization pattern since (it) cannot indicate the successive overlap which is so characteristic a feature of cortical representation."

In addition to overlap, several systematic deviations from an orderly representation of the body on the cortical surface have been found. The orderly progression of proximal to distal muscles as one moves dorsal to ventral within arm and hand representations of primary motor cortex is better described by a core-surround organization with the distal musculature of the hand and wrist represented centrally surrounded by the proximal musculature of the elbow and shoulder (He, Dum, and Strick 1993, Kwan et al. 1978, Meier et al. 2008, Park, Belhaj-Saif, and Cheney 2004, Park et al. 2001) (see Figure 4-1C). Classic motor maps also represent a segment of the body at a single cortical location while in actuality there are multiple re-representations of a body segment. For example, there are at least 3 separate hand areas on the lateral motor strip that project to the spinal cord, two in Brodmann's area 6 and one in Brodmann's area 4 (e.g. Dum and Strick 2002b, Dum and Strick 2005b) (see Figure 4-1F).

Parcellation of the frontal motor areas into multiple motor areas

Campbell (1905), based on cytoarchitecture, was the first to partition the frontal motor areas into what has come to be known as primary motor cortex (PMcx) and premotor cortex (PreMcx) (see Figure 1-1C). Brodmann (1909) largely agreed with Campbell's subdivisions identifying Area 4 with PMcx and Area 6 with PreMcx. Vogt and Vogt (1919) further subdivided PreMcx into rostral and caudal subfields based on subtle differences in cytoarchitecture and myeloarchitecture (Figure 1-1E), a pattern largely confirmed by later

investigators (e.g. Von Bonin and Bailey 1947).

Early stimulation studies supported the parcellation of motor cortex into separate fields as suggested by anatomical differences. In PMcx, electrical stimulation thresholds were lower and the evoked movements were simpler involving more isolated movements about a joint whereas evoked movements in PreMcx were more complex often involving a greater number of body parts (Bucy 1933, Foerster 1936, Vogt and Vogt 1919). Penfield and Brody (1937), based on electrical stimulation in human surgical patients, challenged this characterization suggesting that what appeared to be "complex movements" were in actuality the result of stimulation-induced seizures. They concluded that PMcx and the caudal PreMcx were part of a single motor field while rostral PreMcx had no obvious motor function. Woolsey (1952) came to a similar conclusion based on stimulation experiments in the monkey. While movements in the rostral section of the unified motor map did involve more body segments, this was seen as the natural consequence of the more integrative nature of movements involving the axial musculature.

The current accepted parcellation of PreMcx was proposed by Matelli et al. (Matelli, Luppino, and Rizzolatti 1985) based on cytochrome oxidase histochemistry and confirmed by a host of other techniques (Geyer et al. 1998, Geyer et al. 2000, Luppino et al. 1991, Matelli, Luppino, and Rizzolatti 1991, Wu, Bichot, and Kaas 2000) (Figure 1-1F). This includes 7 distinct motor fields on the lateral convexity of the hemisphere that they labeled F1-F7 (with the F signifying the physical location of the motor fields in (F)rontal cortex.) In the nomenclature that dominates science produced in the United States and in parenthesis the shorthand notation

along with Matelli et al.'s "F" notation, these regions are (see Figure 1-1F) the primary motor cortex (PMcx / F1), supplementary motor area (SMA / F3), pre-SMA (F6), dorsal-caudal subdivision of premotor cortex (PMdr / F2), dorsal-rostral subdivision of premotor cortex (PMdc / F7), ventral-caudal subdivision of premotor cortex (PMvc / F4), and the ventral-rostral subdivision of premotor cortex (PMvr / F5.) In addition, there are at least three motor fields within the medial wall of the hemisphere in and around the cingulate sulcus with each field containing a complete somatotopic representation of the body (Dum and Strick 1991, He, Dum, and Strick 1995, Matelli, Luppino, and Rizzolatti 1991).

Each motor subregion receives distinct inputs from posterior parietal cortex, prefrontal cortex, and subcortical areas (e.g. Dum and Strick 2005a, Hoover and Strick 1993, Rizzolatti 1998) indicating that these regions are involved in different aspects of motor behavior. However, determining the exact functional role of these areas, and how they work together to produce voluntary motor behavior, has remained a challenge.

One common proposed organizational scheme is that PMcx and the various regions of PreMcx exist in a hierarchical arrangement with PreMcx responsible for planning various aspects of motor behavior and PMcx responsible for executing these plans (e.g. Campbell 1905, Fulton 1935, Wise 1985). However, anatomical tracer studies between PreMcx and the spinal cord have clearly established that many portions of PreMcx make dense direct descending projections to the spinal cord (Dum and Strick 1991, He, Dum, and Strick 1993, 1995) and are thus able to directly influence motor behavior independent of PMcx. This opens the possibility that the

different regions that project to the spinal cord are not hierarchically arranged but instead emphasize different aspects of motor behavior. Direct projections to the spinal cord are not found in PMdr and the pre-SMA. These anterior fields cannot directly affect motor output and thus are presumably at a higher hierarchical level.

Another scheme is to attempt to assign some specific computational mechanism or control function to each premotor field as part of a single unified scheme for the control of the body. For instance, Jeannerod (1981) proposed an influential model in which reach-to-grasp actions are broken into an arm transport component and a grip formation component. This proposed dichotomy is theorized to exist at the level of the organization of the frontal motor areas with PMd responsible for planning aspects of hand transport (e.g. Battaglia-Mayer et al. 2001, Caminiti, Ferraina, and Mayer 1998, Johnson et al. 1996, Marconi, Genovesio, and Battaglia-Mayer 2001, Wise and Boussaoud 1997) and PMv responsible for planning aspects of grasp (e.g. Fogassi et al. 2001, Jeannerod et al. 1995, Rizzolatti 1998, Sakata et al. 1997). However, the anatomy and functional properties do not support a strict segregation. For example, PMd makes direct projections to the hand representations of the spinal cord and PMcx (Dum and Strick 2005b, He, Dum, and Strick 1993), stimulation in PMd can evoke movements of the hand (Godschalk et al. 1995, Raos et al. 2003, Rizzolatti et al. 1988), and single neurons in PMd are tuned to aspects of grasp (Boudreau et al. 2001, Hendrix, Mason, and Ebner 2009, Raos et al. 2004, Stark, Asher, and Abeles 2007).

A consistent complication is that each premotor region seems to have a constellation of

properties that defy attempts to assign simple computational functions. For instance, the SMA has been implicated in temporally sequencing motor actions (Kermadi et al. 1997, Lee and Quessy 2003, Tanji and Kurata 1979, 1982, Tanji and Shima 1994), internally initiated versus sensory-cued actions (Halsband, Matsuzaka, and Tanji 1994, Kermadi et al. 1997, Okano and Tanji 1987, Romo and Schultz 1987, Thaler et al. 1995), bimanual coordination (Brinkman 1984, Serrien et al. 2002, Tanji, Okano, and Sato 1988), and arbitrary sensory-motor mappings (Mushiake, Inase, and Tanji 1991, Tanji and Kurata 1985, Tanji, Okano, and Sato 1988) but is also active during "simple" motor actions (Picard and Strick 2003, Tanji and Kurata 1979, 1982, Thaler et al. 1995). It does not seem that any single "computation" can account for the behavior of each premotor field.

Complex movements evoked by long-duration stimulation

Electrical stimulation, when applied on the time-scale of natural motor behaviors, can evoke complex coordinated movements from the monkeys behavioral repertoire (Graziano, Taylor, and Moore 2002). The movements evoked by long-duration stimulation are notable on several accounts. First, stimulation in the arm and hand region of motor cortex causes the arm to move from any initial starting state toward a common invariant final posture. These movements are never restricted to individual joints much less individual muscles. Further, they involve complex temporal trajectories to bring the arm to the final posture and thus are not simple prolonged rotations about the recruited joints. Second, many stimulated sites evoke movements

that resemble actions from the monkey's behavioral repertoire e.g. defensive, feeding, etc. Correspondingly, stimulation in the classically defined hand and arm regions can often incorporate additional body segments such as the neck, face, mouth, axial musculature, etc. such that the resulting combination of movements more completely represents a meaningful behavior. For example, stimulation in the "hand-to-mouth" region evokes movements of the hand and arm, but also rotation of the head, an opening of the mouth, and positioning of the tongue as though to receive food. Third, different stimulation-evoked behaviors are localized to regions that extend across classically defined primary motor and premotor cortex challenging traditional views of the organization of the frontal motor areas.

Probing function with long-duration stimulation

Over the past several decades, prior to the electrical stimulation studies performed by Graziano and colleagues, investigators have primarily relied on the use of threshold stimulation, in which current amplitude and stimulation duration are minimized, to understand the output properties of motor cortex (see section: "Probing motor cortex with electrical stimulation.") Threshold stimulation is traditionally defined as the set of stimulation parameters that evoke a motor response 50% of the time. While this might be expected to reveal the strongest or most direct influence of the stimulated region on the motor periphery it may not reveal the true causal influence under more general circumstances. Consider the case where a cortical neuron projects to the motor neuron pools of two muscles. Unless all the various properties that influence the

efficacy of evoking a response (e.g. number of synaptic connections, efficacy of synapses, excitability of the motor neurons, background activity of the motor neuron pool, etc.) conspire to equate the two pathways, it should be possible to evoke a response in one muscle and not the other. For these reasons, threshold-stimulation may underestimate how many muscles or joints a limited portion of cortex influences. Further, previous investigators using intracortical microstimulation to map the output properties of motor cortex used brief (e.g. 20ms) trains of stimulation to evoke brief muscle twitches. The short time-scale stimulation durations do not allow for meaningful temporal evolution of movements and thus will enforce the idea that a localized patch of cortex activates muscles but does not produce more complicated movements.

The essential factor in evoking complex coordinated movements from motor cortex is the duration of stimulation. In the facial representation of motor cortex, Huang et al. (1989) evoked simple muscle twitches in the facial musculature when applying stimulation for short (35 ms) duration. These simple muscle twitches evolved into coordinated rhythmic jaw-movements when the duration of stimulation was extended to 3 seconds. In the arm and hand representation, stimulation for a duration of 500ms, the approximate duration of natural arm movements, was necessary to evoke complex-coordinated movements to a posture (Graziano, Taylor, and Moore 2002). Stimulation applied for shorter durations evoked only a truncated version of the full movement suggesting that muscle-twitches evoked by short-duration stimulation may only represent the leading component of a more complex movement.

A common argument against the use of long-duration stimulation is that it causes

physiological spread of current along trans-synaptic pathways and is thus at odds with the objective of determining the unique contribution of the stimulated point in cortex. The counter-argument is that the contribution of the stimulated region is determined by its influence on the networks to which it projects. Long-duration stimulation, by recruiting these networks, is essential to understanding the function of the stimulated region. This perspective would be challenged if the spread of activity were non-specific, that, similar to a seizure, activation spreads across cortex in an indiscriminate fashion. Several recent studies have shown that the physiological spread of current is restricted to regions that form a well-defined functional circuit. For example, electrical stimulation in V1 activates retinotopically matched subregions of extrastriate visual areas that are directly innervated by V1 (Tolias et al. 2005). Similarly, Stepniewska and colleagues (2011) used a combination of stimulation and optical imaging to show that long-duration stimulation of posterior parietal cortex causes localized activation in functionally matched regions of motor cortex. Further, the duration of stimulation (they tried 60ms to 500ms) did not affect the spatial extent of the activated region.

More generally, long-duration stimulation has proved a useful technique in establishing or corroborating the function of a variety of cortical areas outside the motor domain (e.g. Bruce et al. 1985, Freedman, Stanford, and Sparks 1996, Gottlieb, Bruce, and Macavoy 1993, Romo et al. 1998, Salzman, Britten, and Newsome 1990, Tehovnik and Lee 1993). For instance, stimulation of MT will bias a monkey's decision about the direction of moving dots in a manner congruent with the preferred direction of motion of neurons recorded at the same cortical

location (Salzman, Britten, and Newsome 1990). Likewise, in a flutter discrimination task, mechanical stimulation of a finger can be swapped with microstimulation of the finger's representation in primary sensory cortex (Brodmann's Area 3b) without any decrement in the animal's performance and without additional training suggesting that the percept evoked by electrical stimulation of cortex is similar, or is at least treated as functionally equivalent, to the percept evoked by mechanical stimulation (Romo et al. 1998).

Cortical output shaped by sensory feedback

Many studies that have used electrical stimulation to probe the causal influence of motor cortex on the muscle periphery have made the implicit assumption that the affects of stimulation at a localized point of cortex are stable (e.g. Cheney, Fetz, and Palmer 1985, Donoghue, Leibovic, and Sanes 1992, Park, Belhaj-Saif, and Cheney 2004). This work suggests that the output of primary motor cortex influences the motor periphery in a largely feed-forward manner. On the other hand, complex movements evoked by long-duration stimulation are characterized by convergence of the arm from any initial starting state toward a common invariant final posture (Graziano, Taylor, and Moore 2002). Thus, stimulation will evoke flexion at a joint or extension at a joint as necessary to bring the arm from its initial state to the final posture. The mechanism of this convergence, at least in part, relies on sensory feedback remapping the influence of stimulation on muscle activation (Armstrong and Drew 1985a, 1985b, Gellhorn

1948, Gellhorn 1949, Gellhorn and Hyde 1953, Lemon, Johansson, and Westling 1995, Sanes, Wang, and Donoghue 1992). A particularly clear example is given Graziano et al. (2004). They measured stimulation-evoked biceps and triceps EMG activity while holding the elbow fixed at angles subtending the full working range of the joint. They found that biceps and triceps activity was influenced by the position of the elbow with some sites showing a reversal: when the elbow was flexed the triceps was more strongly activated whereas when the elbow was extended the biceps was more strongly activated. Stimulation recruited muscles so as to drive the arm to an intermediate posture. The authors further showed that sensory modifications of EMG responses could be attained whether long-duration stimulation or stimulus-triggered averaging was used. Stimulus-triggered averaging involves averaging muscle activity time-locked to the injection of single low-amplitude pulses of current. This technique minimizes spatial and temporal summation at the targets of the stimulated site and therefore emphasizes the relatively direct affect of the stimulated region on the motor periphery (Cheney and Fetz 1985). The correspondence between stimulus-triggered averaging and spike-triggered averaging (Cheney and Fetz 1985) suggests that even the direct influence of motor cortical neurons on the muscle periphery are remapped by sensory feedback. Finally, long-duration stimulation often resulted in the simultaneous activation of biceps and triceps indicating that the viscoelastic properties inherent to muscles also contribute to establishing the stimulation-evoked posture (Bizzi et al. 1992, Hill 1970, Polit and Bizzi 1979, Richardson et al. 2005, Scott, Brown, and Loeb 1996).

Control of the arm in part relies on cortical control of spinal networks that coordinate the

action of many muscles (Berniker et al. 2009, Bizzi et al. 2008, Giszter, Mussa-Ivaldi, and Bizzi 1993, Krishnamoorthy, Scholz, and Latash 2007, Ting and Macpherson 2005). Numerous computational models have demonstrated how spinal interneuron systems under guidance of sensory feedback can perform a sophisticated array of motor behaviors (Bullock and Contreras-Vidal 1993, Cisek, Grossberg, and Bullock 1998, Feldman 1986, Gribble et al. 1998, Loeb, Brown, and Cheng 1999, Raphael, Tsianos, and Loeb 2010). It should not be surprising that stimulation can evoke rather sophisticated motor behaviors if the affects of stimulation are mediated in part by such spinal systems (or similarly arranged systems within cortex and subcortical motor centers).

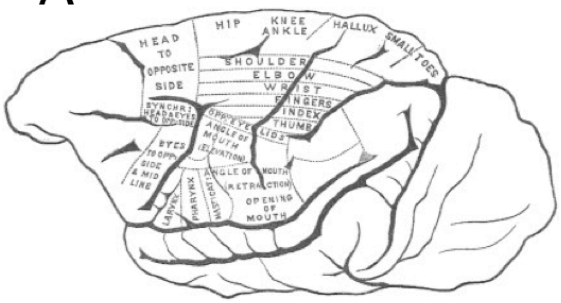
Ethological actions evoked by stimulation

Complex movements evoked by long-duration stimulation can resemble actions from the monkey's behavioral repertoire. Graziano et al. (2002) report on three types of evoked behaviors: hand-to-mouth movements, central-space manipulation movements, and defensive movements. To what degree does the behavior suggested by the evoked-movement correspond to other properties of the stimulated region? This question has been most carefully addressed in the cortical area from which stimulation evokes defensive movements. This defensive region corresponds roughly to the dorsal portion of area F4 in the ventral premotor cortex. Neurons in F4 have bimodal neurons that respond to tactile and visual stimuli (Fogassi et al. 1996, Gentilucci et al. 1988, Graziano and Gandhi 2000, Rizzolatti et al. 1981) and occasionally

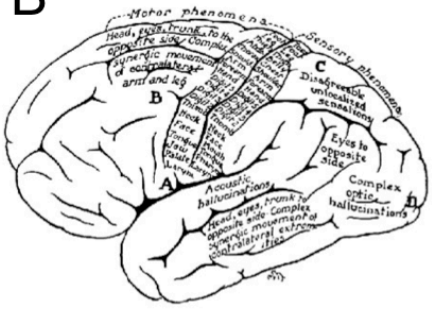
auditory stimuli (Graziano, Reiss, and Gross 1999). Tactile receptive fields are generally located on the arm, face, or torso. Visual receptive fields are often localized to the tactile receptive fields extending out 5-30 cm from the surface of the skin and are especially responsive to visual stimuli approaching the tactile receptive field. The visual receptive fields are often anchored to the tactile receptive field irrespective of the position of the arm or direction of gaze (Fogassi et al. 1996, Graziano, Hu, and Gross 1997). These "body-part centered" neurons were theorized to play a general role in the sensory guidance of movement (Graziano, Hu, and Gross 1997). More recent evidence from the Graziano lab suggests that these sensory properties may be especially important for monitoring a region around the body for purposes of defense. When long-duration stimulation was applied to a region containing sensory receptive fields anchored to the arm, the evoked movement involved a rapid arm withdrawal ending with the arm in a protected position behind the back (Graziano, Taylor, and Moore 2002). When stimulation was applied to a region containing sensory receptive fields localized to the face, the evoked movement involved a facial grimace, an eye-squint, retraction of the ear, a turning of the head away from the receptive field, and movement of the arm up to the side of the face with the palm turned outward (Graziano, Taylor, and Moore 2002). These movements match many properties of naturally occurring defensive movements (Cooke and Graziano 2003). Further, bimodal neurons in this region are active during defensive reactions to an air puff on the face and the magnitude of neural response correlates with the magnitude of the defensive reaction (Cooke and Graziano 2004a). Finally, chemical inactivation of this region led to diminished defensive responses to air puffs while

leaving other motor behavior unaffected. Chemical excitation of this region led to enhanced as well as spontaneous defensive responses while leaving other motor behaviors unaffected (Cooke and Graziano 2004b). These results provide convergent evidence that the region from which stimulation evokes a defensive movement is especially important in organizing natural defensive reactions. In another study, ablation of the forelimb reaching region of the rat, as identified by long-duration stimulation, led to a selective loss of forelimb reaching movements (Ramanathan, Conner, and Tuszynski 2006). Behavioral recovery following ablation was correlated with the reemergence of a cortical region from which stimulation evoked forelimb-reaching movements. These results provide strong evidence that motor cortex is organized into a series of modules that emphasize different aspects of the monkey's natural motor repertoire and that these regions can be identified with the use of electrical stimulation.

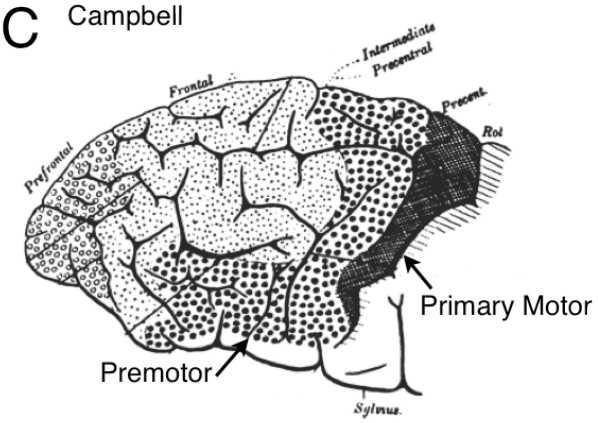
A Beevor and Horsley



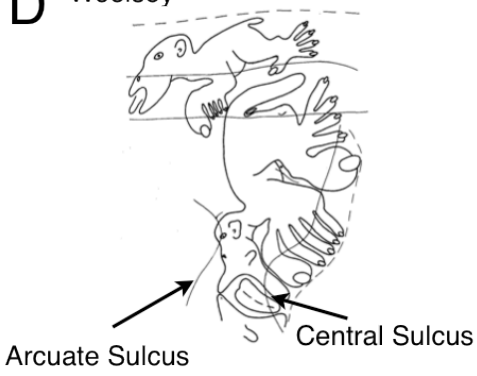
B Foerster



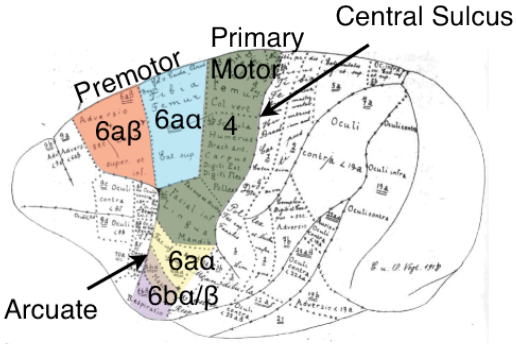
C Campbell



D Woolsey



E Vogt and Vogt



F Matteli et al.

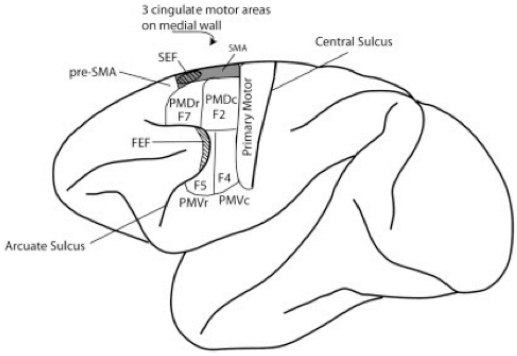


Figure 1-1: Schematic descriptions of the topographic organization of motor cortex. A. Somatotopic organization of the monkey motor areas characterized by a vertical stacking of body parts along the central sulcus (Beevor 1890) B. A similar somatotopic organization in the human brain (Foerster 1936) C. Division of motor cortex into an anterior premotor (Campbell's "intermediate precentral") and posterior primary motor cortex (Campbell's "precentral") (Campbell 1905). D. Woolsey rejects the division of premotor and primary motor cortex instead proposing a single laterally located unified map of the body ("M1") and a second map along the midline ("M2"/ SMA) (Woolsey et al. 1952). E. Division of premotor cortex into anterior and posterior subdivisions (Vogt and Vogt 1919). F. Partitioning of premotor cortex into multiple premotor areas (Luppino et al. 1991, Matelli, Luppino, and Rizzolatti 1985)

Overview of dissertation studies

This thesis describes three studies based on the finding that half-second stimulation can evoke complex movements in motor cortex.

Study 1: Arm movements evoked by electrical stimulation

Long-duration stimulation evokes complex coordinated movements to a posture. We analyzed these movements in greater detail using a high-resolution 3D tracking system. The affects of stimulation are modulated by sensory feedback. We looked at whether the feedback systems engaged by stimulation could show compensation at the level of movement kinematics under two conditions. We tested whether stimulation-evoked movements could compensate for a weight added to the hand, and thus achieve the same final hand position as movements evoked without the weight, or whether the hand would ultimately be dragged down by the added weight. We also tested whether stimulation tended to bring the hand to a similar final location despite trial-by-trial variability in joint angles. Movements evoked by stimulation can resemble movements from the monkey's behavioral repertoire. We stimulated extensively throughout the frontal motor areas to add to the emerging picture of behavioral localization.

Study 2: Relationship between unconstrained arm movements and single neuron firing in the macaque motor cortex

Long-duration stimulation drives the arm to a final invariant posture. The affect of stimulation on muscle activity is remapped by sensory feedback such that a cortical location can excite flexor or extensor muscles as is necessary to move the arm to the final posture. One possible interpretation of these results is that cortical neurons, in part, specify a final desired posture of the arm. We recorded the activity of single neurons as the monkey made naturalistic untrained movements across his workspace. We directly tested whether the behavior of single neurons during these naturalistic movements was explained by tuning to a preferred posture. The tuning model was than compared with the stimulation-evoked posture at the same cortical site.

Many tuning models have been proposed to explain the behavior of neurons during movement. One limitation in previous studies is a tendency to use monkeys that have been extensively trained on highly constrained movement sets. It is possible that the observed neural tuning emerged as a consequence of this training and may not be apparent during natural motor behavior. We tested whether previous tuning models could account for neural behavior during naturalistic untrained movements.

Previous studies have used behavioral paradigms and analysis methods that could emphasize one tuning model at the expense of others. We asked how much of the total neural variance each tuning model could explain during more unstructured, naturalistic motor behavior.

Study 3: Possible origins of the complex topographic organization of motor cortex: reduction of a multidimensional space onto a 2-dimensional array

Long-duration stimulation revealed three forms of topographic organization in the frontal motor areas. This included a rough, highly-overlapping and fractured somatotopic organization, a rough map of the location in space to which the hand moved on stimulation, and a partitioning into zones that emphasized different complex, ethologically relevant movements. We used computational modeling to try to explain the observed patterns of organization. We defined a high-dimensional representation of the monkey's behavioral repertoire that included the set of body parts engaged in the behavior, a final position of the hand, and a behavioral category. We than rendered this high-dimensional description onto a model of a two-dimensional cortical surface subject to the constraint that similar movements in the high-dimensional movement space be close together on the cortical surface thus maximizing functional continuity.

Chapter 2 : Arm movements evoked by electrical stimulation

Introduction

Electrical microstimulation in the motor cortex of monkeys can evoke complex-coordinated movements from the monkey's behavioral repertoire (Graziano, Taylor, and Moore 2002). We sought to more carefully quantify the properties of evoked movements using a 3D tracking system to measure hand position and joint angles. This allowed us to compare various properties of natural voluntary movements with those evoked by stimulation. For example, in normal voluntary movement, the hand follows a smooth trajectory with a bell shaped velocity profile (e.g. Bizzi and Mussa-Ivaldi 1989, Flash and Hogan 1985). If external forces are imposed during arm movement, the movement can show compensation achieving a similar final hand position (Gribble and Ostry 2000, Lackner and Dizio 1994, Shadmehr and Mussa-Ivaldi 1994). When positioning the hand, joint angles are not held fixed, but instead are allowed to covary so long as the position of the hand remains stable (Scholz and Schöner 1999, Todorov and Jordan 2002). Do stimulation evoked movements show similar properties? We further examined the cortical localisation of ethological movement representations over a larger swath of the lateral motor strip than previously explored. The mapping data in the present experiment adds to the emerging picture of the clustering of different types of stimulation-evoked movements.

Materials and Methods

All husbandry, surgical, and behavioral procedures were approved by the Princeton University Institutional Animal Care and Use Committee and the attendant veterinarian and were in accordance with N.I.H. and U.S.D.A. guidelines. We studied two adult male *Macaca fascicularis* (4.5 and 7.0 kg).

Surgery

For each monkey, an initial surgical operation was performed under isoflurane anesthesia and strict aseptic conditions, during which an acrylic skullcap was fixed to the skull with bone screws. A steel bolt for holding the head, and a 2.5 cm diameter steel chamber for neuronal recording and electrical stimulation, were also imbedded in the acrylic. The recording chamber was positioned for a vertical (dorsal-ventral) approach to the precentral gyrus. Each animal recovered from the surgery within a week, but was given two additional weeks to allow the skull to grow tightly around the skull screws. In a subsequent procedure, also under deep anesthesia and aseptic conditions, the recording chamber was opened and a hole approximately 10 mm in diameter was drilled through the layer of acrylic and the bone, exposing the dura.

Neuronal recording and stimulation

During the daily recording sessions, the monkey sat in a Lexan primate chair with the head restrained by the head bolt. A hydraulic microdrive (Narishige) was mounted to the top of the recording chamber. A steel guide tube (an 18-gauge syringe needle) was lowered through the hole in the skull and into the granulation tissue that lay over the dura. Then the varnish-coated tungsten microelectrode (Frederick Haer Inc., impedance 0.5-2 MOhm) was advanced from the guide tube through the dura and into the brain. Typically an electrode would start with an impedance of approximately 2 MOhm. After repeated use on multiple days, as the insulation began to wear off near the tip, the impedance would begin to drop. When clear single neurons were no longer easily isolable on the electrode, a new one was used for the next penetration.

The introduction of the electrode into the cortex was confirmed by monitoring neuronal activity on an oscilloscope and over a loud speaker. The location of the top of the cortex was estimated as the depth at which spiking neuronal activity was first found. The electrode was typically advanced into the cortex 1 to 2 mm beyond that depth, and then electrical stimulation was tested. On some penetrations, multiple depths were tested, separated by 0.5 mm. When the electrode was advanced beyond the cortex and into the white matter, as indicated by the drop in neuronal cellular spiking, electrical stimulation was no longer tested and the electrode was retracted. A systematic test of different layers of cortex was not attempted in the present

experiment. In previous experiments (Graziano, Taylor, and Moore 2002) our lab found that the movement evoked by 500-ms stimulation is similar as the electrode is advanced perpendicularly through cortex, and that the movement thresholds are typically lower in the deeper layers of cortex.

At each cortical site studied, stimulation was applied by an S88 stimulator and two SIU6 stimulus isolation units (Grass, West Warwick, RI). Stimulation consisted of a train of biphasic pulses presented at 200 Hz. In some cases, 100, 150, and 250 Hz were also tested, causing similar effects (see Results). Each stimulation pulse had a negative followed by a positive phase; each phase was 0.2 ms in duration. In some cases, phases of 0.4 ms were also tested. For most movements, a 500 ms stimulation train was used.

Current was measured via the voltage drop across a 1Kohm resistor in series with the return lead of the stimulus isolation units. For each site we varied the current until an evoked movement was observed. The threshold, the current at which any visible movement was evoked 50% of the time, was determined by two observers. These threshold measurements were thus approximate. The average threshold measured in this fashion was 19 MicroA with a standard deviation (SD) of 14 MicroA. Thresholds were generally lowest (sometimes as low as 5 MicroA) in the anterior bank of the central sulcus. For quantification of the evoked movement, the current was usually set between 25 and 100 MicroA. The current level was adjusted until a clear, consistent, multijoint movement of the arm was obtained, and then the quantitative testing was begun.

Asanuma and Arnold (1975) found that extended trains of cathodal pulses killed the cortical tissue around the electrode tip. With balanced biphasic pulses however, it is possible to stimulate for long durations (seconds) and high currents (milliamps) without measurable damage (e.g. Tehovnik 1996). Our stimulation parameters are within the range for cortical stimulation studies in oculomotor, visual, and somatosensory systems (e.g. Bruce et al. 1985, Freedman, Stanford, and Sparks 1996, Gottlieb, Bruce, and Macavoy 1993, Romo et al. 1998, Salzman, Britten, and Newsome 1990, Tehovnik and Lee 1993). To check whether the current damaged the brain, for each site studied, after stimulating for many trials we switched the amplifier to neuronal-recording mode and confirmed that normal neuronal spiking activity could still be obtained. Second, as Asanuma and Arnold (1975) pointed out, electrical damage to the brain causes the effect of the stimulation to disappear after several trials. We found that the stimulation had a consistent effect over hundreds of trials with no sign of degradation. Finally, in past experiments using the same stimulation techniques (Cooke et al. 2003, Graziano, Taylor, and Moore 2002) on histology our lab found no visible damage to the cortex associated with the stimulation sites. These lines of evidence suggest that our trains of biphasic pulses do not cause extensive cell death, though more subtle damage cannot be ruled out.

Location of stimulation sites

The monkeys were not sacrificed at the termination of this experiment, thus the locations of the stimulation sites were reconstructed through non-histological means. The central and arcuate sulci were located first by shining a bright light on the dura during the initial craniotomy surgery. Both sulci were clearly visible through the dura. The microdrive was then mounted to the recording chamber and the locations of the visualized sulci were measured with the tip of the guide tube. In this way, the locations of the sulci were obtained in microdrive coordinates. The sulcus locations shown in Figure 2-10 are based on this procedure.

After the craniotomy surgery, during the daily experiments, the measured location of the central sulcus was confirmed to within 0.5 mm by recording and stimulating to either side of the sulcus. Just posterior to the sulcus, in primary somatosensory cortex, we observed the expected small tactile receptive fields on the contralateral hand and also the expected lack of effect of electrical stimulation. Just anterior to the sulcus we obtained the expected low stimulation thresholds in primary motor cortex. In monkey 2, the location of the arcuate sulcus was confirmed by stimulating just anterior to it and obtaining no skeletomotor movements, but instead obtaining stimulation-evoked saccadic eye movements presumably in the frontal eye fields. The location of both the central and arcuate sulci were further verified by using the pattern of cellular activity and silence obtained on long electrode penetrations to reconstruct the arrangement of cortex and white matter.

Within the precentral gyrus, the studied area encompassed the arm and hand representation and was bracketed laterally by the face representation and medially by a leg represention. At the most lateral sites shown in Figure 2-10, stimulation sometimes evoked movements of the arm, hand, and mouth. Sites that were even more lateral (not shown) appeared to be in the orofacial region of motor cortex; no arm or hand movements were evoked from these more lateral sites, only facial or oral movements. At the most medial sites shown, stimulation evoked movements of the arm, hand, and leg.

The medial-lateral, anterior-posterior, and dorsal-ventral location of every tested site was supplied to a mapping program to construct a 3-D model of the cortical area studied. The model was then flattened, collapsing sites onto a 2-D surface, and unfolding the anterior bank of the central sulcus. The plots in Figure 2-10 show these 2-D reconstructions, including the gyral surface between the arcuate and central sulci, and the unfolded anterior bank of the central sulcus.

Measurement of limb position

The positions of points on the limb were measured by means of an Optotrak 3020 system (Northern Digital). This system tracks the three-dimensional position of infrared light emitting diodes (LEDs). Each LED could be separately tracked to a spatial resolution of 0.1 mm. The position was measured every 14.3 ms. To create a marker that could be detected by the Optotrak

cameras from any angle, we glued five individual LEDs together to produce an omni-directional marker ball. A marker ball was taped to the monkey's forefinger on the dorsal surface where it would not interfere with grasping; on the thumb, again on the dorsal surface where it would not interfere with grasping; on the back of the hand, between the knuckles of the third and fourth digits; on the lateral aspect of the elbow; and on the shoulder. In addition, 14 individual markers were taped in a double ring around the monkey's wrist, with seven markers per ring and a 1-cm spacing between the rings. A marker was also taped to the monkey's lower jaw, to measure the opening and closing of the mouth and to measure the relative position of the hand and face. A marker was taped to the side of the primate chair to calibrate the position of the monkey with respect to the chair. For the LEDs attached to the arm and hand, the wires were taped in a bundle to the underside of the arm and draped behind the monkey. The monkey's chair was open at the front and side, allowing for almost total range of movement of the arm. The monkey's other arm, ipsilateral to the stimulating electrode, was not studied with Optotrak markers. To ensure that this hand would not reach for the fruit rewards during trials, or tear off the markers taped to the measured hand, this untested hand was fixed to the side of the chair in an arm holder.

The marker balls on the index finger and thumb were used to measure grip aperture. The marker ball on the back of the hand was used to measure hand position.

The double ring of 14 markers around the wrist was subject to a rigid body computation in order to calculate the center point and spatial orientation of the wrist. In this computation, for each time point, a three-dimensional rigid model of the double ring of markers was fitted to the

measured positions of the currently visible markers, using a least-squares method of optimal fit. The orientation and position of the model could then be used to estimate the orientation and center of the wrist. The center of the wrist was taken to be the mean position of the 14 points in the model.

Once the orientation and position of the wrist was calculated, the elbow position in space could then be calculated by assuming that the elbow was a certain distance from the center of the wrist in a direction specified by the orientation of the wrist. This calculation required knowing the distance between the elbow and the wrist markers; this reference distance was measured each day after the ring of markers was put on the wrist, and was typically about 12 cm. The calculated spatial position of the elbow, derived from the wrist markers, was consistently 1-2 cm medial to the marker that was fixed to the lateral aspect of the elbow, as expected. The calculated position was assumed to be closer to the actual point of rotation of the elbow, internal to the arm. For this reason, in reconstructing the posture of the arm, this calculated elbow position was used rather than the location of the marker on the surface of the elbow.

The position of the shoulder in space was calculated by analyzing the position of the elbow over time. Over many time points, the elbow described a portion of a sphere, the origin of which was located at the shoulder joint. For each 3 min block of data, a shoulder position was calculated by fitting a sphere to the data using a least-squares best fit algorithm and using the center of the sphere as the shoulder location. Because the shoulder is capable of small translational movements in addition to rotations, this method of estimating shoulder joint

location is approximate but was sufficient for the purposes of this study. When the shoulder position was calculated multiple times over different time segments, it varied within less than 3 cm. Just as for the elbow, the calculated position of the shoulder joint was assumed to be more accurate than the measured position of the marker ball fixed to the shoulder. The calculated position was at the actual point of rotation, as estimated by the best fit algorithm, whereas the marker ball was fixed to the lateral surface of the shoulder. Thus in reconstructing the posture of the arm, the calculated shoulder position was used.

Three shoulder angles were computed: the elevation; the azimuth; and the "twist" or internal/external rotation of the shoulder joint. We also calculated the flexion of the elbow; the pronation of the forearm; the extension of the wrist; the adduction of the wrist; and the grip aperture. In total, 8 degrees of freedom were calculated for the arm. This model of the arm was verified by applying forward kinematics to estimate the position of the hand. This calculated position of the hand matched the actual, measured position of the hand to within an accuracy of 1.5 cm.

Experimental Sessions

In order to study the effect of different starting postures, stimulation was applied while the monkey performed a simple untrained reaching task. A small piece of fruit was held with forceps by the experimenter at one of many possible locations around the monkey, and the

monkey reached for the fruit. On approximately two thirds of the reaches, stimulation was applied by the experimenter as the hand approached the target location to within about one centimeter but before the monkey had grasped the fruit. The purpose was to stimulate at a moment when the hand was relatively still and at a variety of locations in the workspace, and to do so in a manner that did not entrain the monkey to particular postures. Stimulation was also applied during the monkey's spontaneous reaches that placed the hand in different locations in the workspace, and while the monkey was sitting quietly with the arm stationary. In this way stimulation could be tested during a range of initial postures of the limb. All stimulation trials, whether they took place during a reach to a fruit reward, during spontaneous reaches, or when the monkey was resting, were combined in the data since the resultant stimulation-evoked movement of the arm to a final posture was similar for all of these cases. Data were collected continuously during a 3-min block in which an average of 25 stimulation trials were tested. Two to three blocks were tested for each stimulation site.

On some blocks of trials, a 90-g lead bracelet was wrapped tightly around the contralateral wrist just proximal to the ring of Optotrak markers. A 90 g weight represents about 2% of body weight or about 25% of arm weight (arm weight approximately 300 g). In pilot tests using heavier weights such as 130 g, the monkey was unwilling to lift its hand and therefore we were not able to test different initial hand locations. In contrast, with the 90 g weight, the monkey did reach for fruit rewards and thus we could test a range of initial locations.

For each stimulation site, we first tested the stimulation-evoked movement without a weight on the wrist in a 3-min block of stimulation trials (typically 25 trials). Then the weight was added and a second block of trials was run. This alternation of unweighted and weighted blocks was continued, typically for 4 to 6 blocks.

To help test whether the system showed any evidence of compensation for the weight, we estimated the expected effect that the weight should have if there was no compensation. To perform this calculation, we modeled the physics of the arm. The model involved two hinged segments, an upper arm and a forearm. The shoulder had three degrees of freedom of rotation, and the elbow had one degree of freedom. The model incorporated the gravitational force on these segments, a spring-like muscle force acting around each degree of joint rotation, and a damping term for each degree of joint rotation. Note that this type of model is deterministic; given the lengths and weight distributions of the segments, the force of gravity, and the spring constant, equilibrium angle, and damping force for each joint, Newtonian mechanics fully specifies the equations of motion. We used the Denavit-Hartenberg representation, a standard method for simplifying the equations of a multi-link arm (Denavit 1955), and the recursive Newton-Euler method, a standard method to transform between kinematic variables and dynamic variables in a multi-link arm (Luh and Walker 1980, Walker and Orin 1982). The lengths of the arm segments were taken directly from the optotrak measurements of the arm. Each segment was approximated as a cylinder of uniform width and density. The diameter of each segment was taken as the approximate mean diameter of the monkey's arm (6 cm). The total arm weight

was set to 300 g, which matched our estimate from volumetric measures of the arm. Each degree of freedom of joint rotation was modeled as a spring system and given a spring constant, an equilibrium angle, and a damping term. These parameters were different for different stimulation sites. They were determined in the following way. For each stimulation site tested, we analyzed the data collected from the arm on actual trials when no weight was present. We used the joint angles, angular speed, and angular acceleration, and applied a least-squares method of optimal fit to estimate the spring constants, equilibrium angles, and damping terms that best characterized the movement for those trials. We incorporated those dynamic terms into the model. Using the model, we were able to obtain simulated trajectories that closely matched the real trajectories, as expected.

We then ran the model with an added 90 g weight on the wrist and obtained the calculated trajectory and final position of the hand. This calculation showed, as expected, that the weight should have two principal effects: it should reduce the peak acceleration of movement due to the increased inertial mass; and it should cause the final height of the hand to be lower, due to the greater downward gravitational force on the hand. These effects predicted by the model were always pronounced, due to the fact that that weight was concentrated at the wrist. We then compared this calculated effect of the weight to the actual data collected with the weight fixed to the wrist. This comparison is detailed in the results section.

Results

Stimulation-evoked hand trajectories: basic description

Hand trajectories were measured during electrical stimulation of 91 sites in motor cortex. Figure 2-1 shows examples of hand trajectories evoked by stimulation of 18 typical sites. In each case, stimulation evoked a movement of the arm toward a final posture, and thus a movement of the hand toward a final region of space. As can be seen across the 18 examples, stimulation of different sites in cortex drove the hand toward different spatial regions. As illustrated in Figure 2-1T, when the stimulator was disconnected from the elctrode while otherwise holding all other experimental conditions constant, triggering stimulation had no affect. The pattern obtained during sham stimulation was therefore unlike the pattern obtained during actual brain stimulation at that site (shown in Figure 2-1S), and unlike the pattern found at any stimulation site.

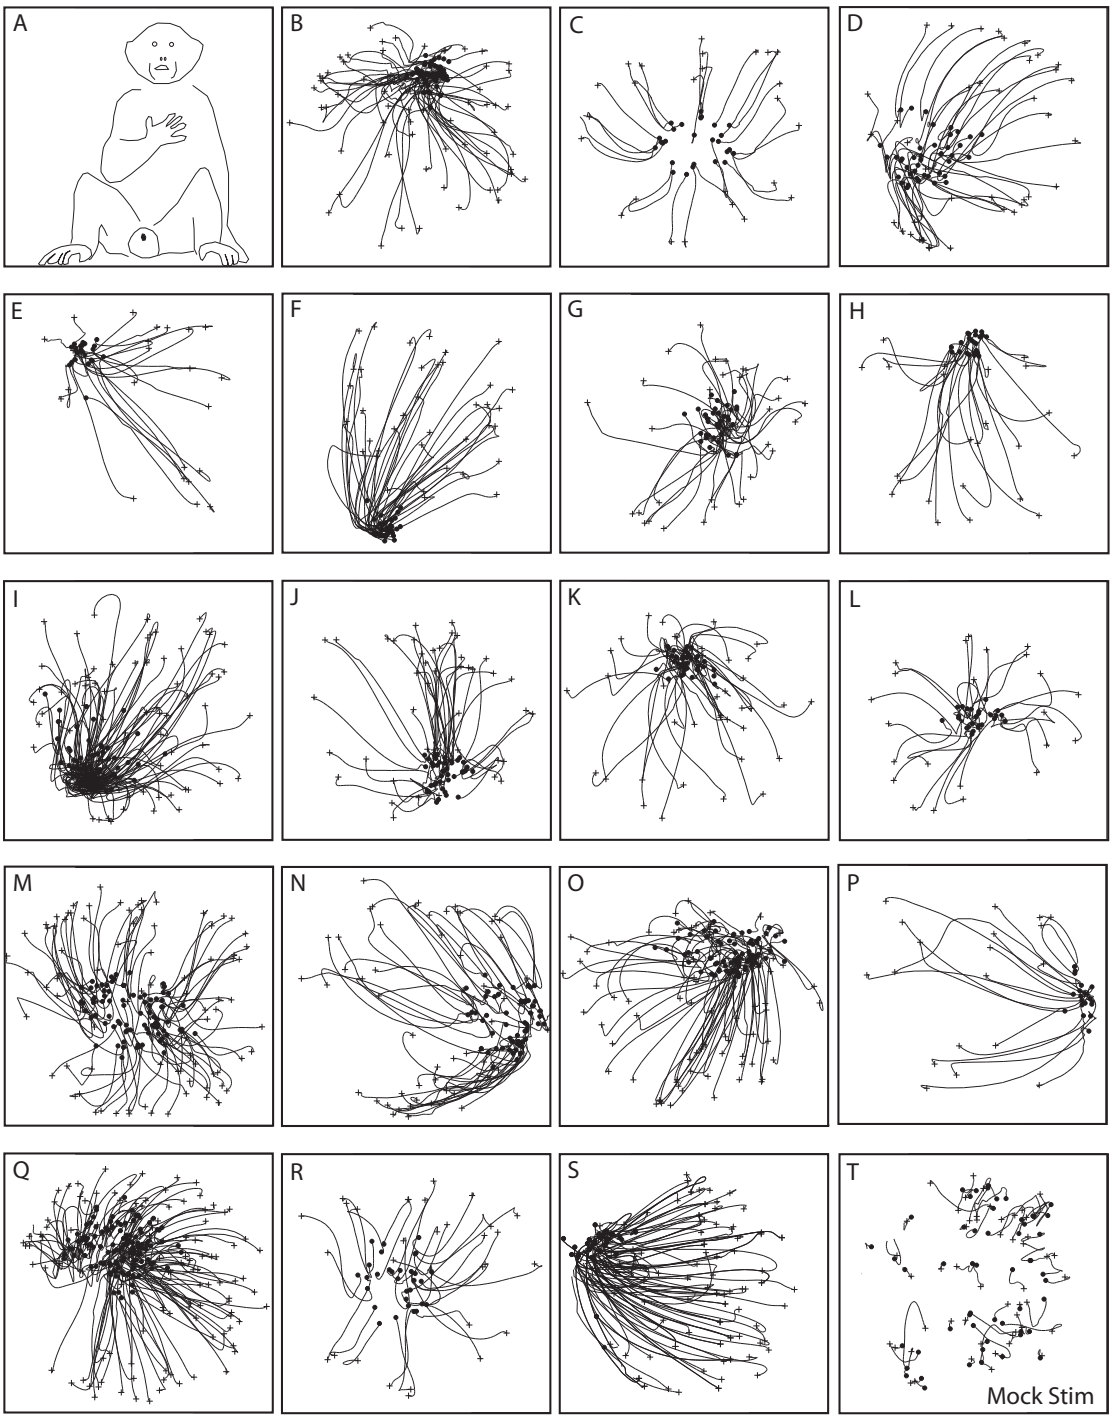

A
B
C
D
E
F
G
H
I
J
K
L
M
N
O
P
Q
R
S
T
Mock Stim

Figure 2-1 : Examples of hand movements evoked by microstimulation in motor cortex. A. The monkey drawing indicates the approximate size, location, and perspective of the monkey within the square frame. The height of the frame represents 50 cm. B-S. Stimulation-evoked hand movements from 18 typical stimulation sites. Each thin black line shows the path of the hand during a stimulation train. The + indicates the start of the movement. The black dot indicates the end of the movement. In a small number of trials, the tracking markers were transiently blocked from the view of the camera due to the specific posture of the limb. In these cases, the trace is interrupted. T. Result of mock stimulation in which the wires to the electrode were disconnected but all other aspects of the testing were the same. The traces in panel S show the result for this same cortical site when the wires were connected.

Figure 2-2A1 shows speed profiles for a typical stimulation site. To calculate latency, we fit a straight line to the pre-stimulation baseline data using a least-squares fit, then fit a straight line to the rising phase of the speed, and determined the time at which the two lines intersected. Calculated in this fashion, the latency was 82.4 ms. (The average latency among the 91 sites was 80.2 ms, with a range of 49.9 to 103.3 and a standard deviation of 10.9).

The speed profiles were usually bell-shaped. This bell-shaped property can be seen more clearly in Figure 2-2A2. Here, the same trials are aligned on peak speed. This example site shows a statistically significant fit to a Gaussian curve (Regression Analysis, $F = 405$, $P < 0.0001$). This fit to a Gaussian was statistically significant for all 91 sites (significance level of 0.05, Bonferroni-adjusted for 91 tests).

Figure 2-2A3 shows the peak speed as a function of the distance that the hand traveled. The relationship is roughly linear, in which greater peak speeds occurred during longer movements. For this example site, the data show a significant linear trend (Regression Analysis,

F = 235, P < 0.0001). A significant linear trend was obtained for 90 of the 91 sites tested (significance level of 0.05, Bonferroni-adjusted for 91 tests).

Figure 2-2B and C show similar results for two more example sites.

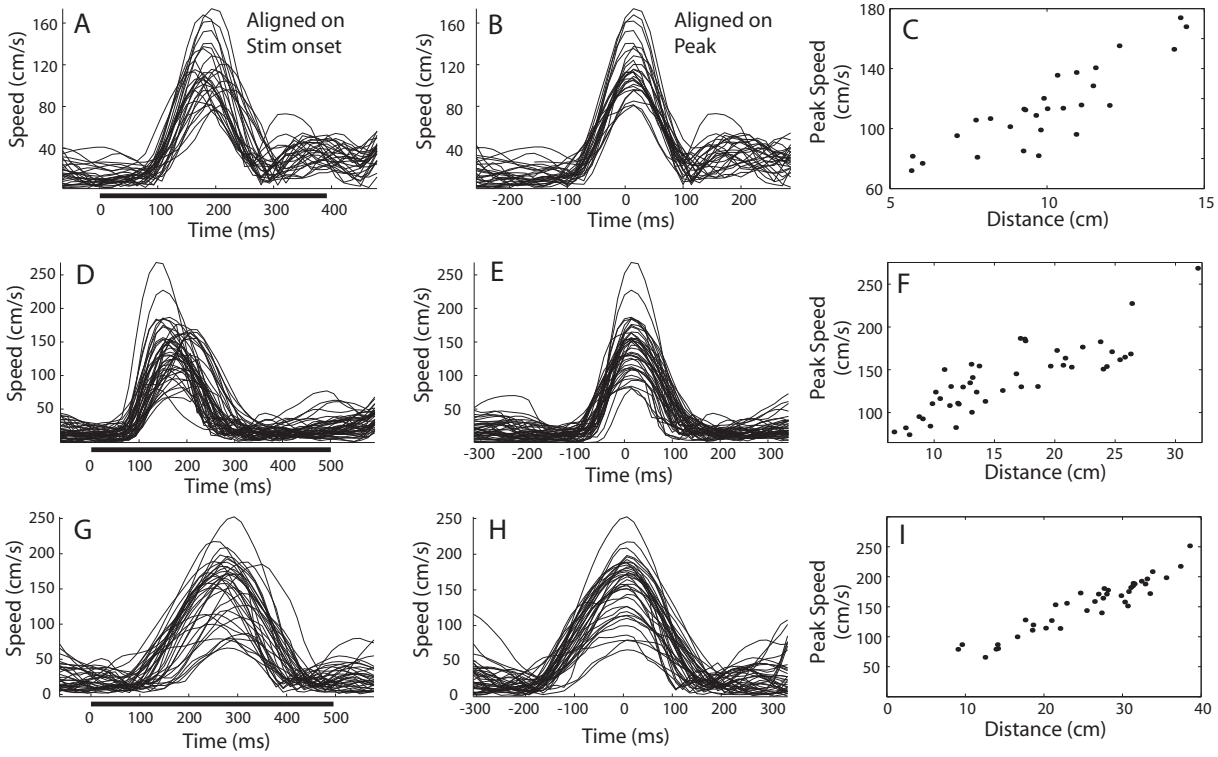


Figure 2-2 : Speed profiles for three typical stimulation sites. A1. Hand speed as a function of time during stimulation, for one stimulation site. Each trace shows the result for one stimulation trial. Speed measured in 14.3-ms increments. Thick black bar at bottom shows the time of the stimulation. A2. Same data as in A1 but the traces are aligned on the time of peak speed. A3. Peak speed during the stimulation-evoked movement as a function of the distance that the hand traveled. B and C. Similar plots for a two more stimulation sites.

We typically tested sites using 200 Hz stimulation, a frequency borrowed from the oculomotor literature (e.g. Robinson and Fuchs 1969). For 6 cortical sites, we examined the effect of different frequencies. Figure 2-3 shows the results for one site tested with stimulation at 100 Hz, 150 Hz, 200 Hz, and 250 Hz. The movement of the hand was similar across these different stimulation frequencies. In each case, the hand converged from a range of initial positions toward a similar final region of space. As shown in Figure 3E, the peak speed of movement varied somewhat with stimulation frequency. The lowest speeds were obtained with 100-Hz stimulation; the highest speeds were obtained with 200-Hz and 250-Hz stimulation. This effect of stimulation frequency on hand speed was significant ($F = 15.53$, $P < 0.0001$). These results suggest that similar movements can be evoked with a broad range of stimulation frequencies, though higher frequencies tended to evoke somewhat higher speeds.

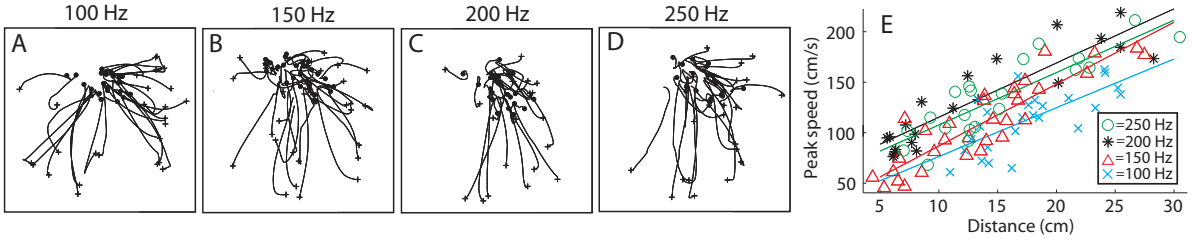


Figure 2-3: Effect of different stimulation frequencies on evoked hand movement for one stimulation site. A-D. Movement of the hand evoked by stimulation at 100, 150, 200, and 250 Hz. See caption to Figure 1 for details of movement traces. E. Peak speed of movement vs. distance that the hand moved for the four different stimulation frequencies.

Convergence of the hand toward a location

The most consistent feature of the stimulation-evoked movements was the convergence of the arm toward a posture, and thus the convergence of the hand toward a region of space. This convergence can be seen in the 18 example sites depicted in Figure 2-1. Figure 2-4A shows this convergence in greater detail for a typical site. The convergence of the hand is shown in each of the three Cartesian dimensions. To investigate the degree of convergence further, we calculated the "mean-distance-to-center" and analyzed how this metric evolved over time during the stimulation. The mean-distance-to-center is a measure of dispersion around the mean. For each time increment we first calculated the mean hand position across all trials. Then, for each trial, we computed the radial distance from the hand to that mean position. Finally, we averaged these distances across trials to arrive at the mean-distance-to-center.

Figure 2-4A4 shows that for this example site, the mean-distance-to-center began at a high level, then dropped during stimulation. This drop in mean-distance-to-center began about 80 ms after the onset of stimulation. It is important to note that this 80 ms is not simply the latency for the hand to move. Rather, it is the latency with which the hand began to converge toward a tighter cluster of locations. To determine if the degree of convergence was statistically significant, we compared the mean-distance-to-center at the start of stimulation and at the end of stimulation (paired t test). The mean-distance-to-center was significantly reduced at the end of

stimulation ($t=14.58$; $P < 0.0001$), indicating that the stimulation caused a significant spatial convergence.

Figure 2-4B and C show similar results for two more example sites. These examples are typical of the population. Of the 91 sites tested, all 91 showed a statistically significant degree of convergence of the hand (significance level of 0.05, Bonferroni-adjusted for 91 tests).

Figure 2-4D shows the mean result for all 91 sites tested. Just as for the individual examples, for the group data the mean-distance-to-center began to drop after stimulation and the latency of this drop was approximately 80 ms.

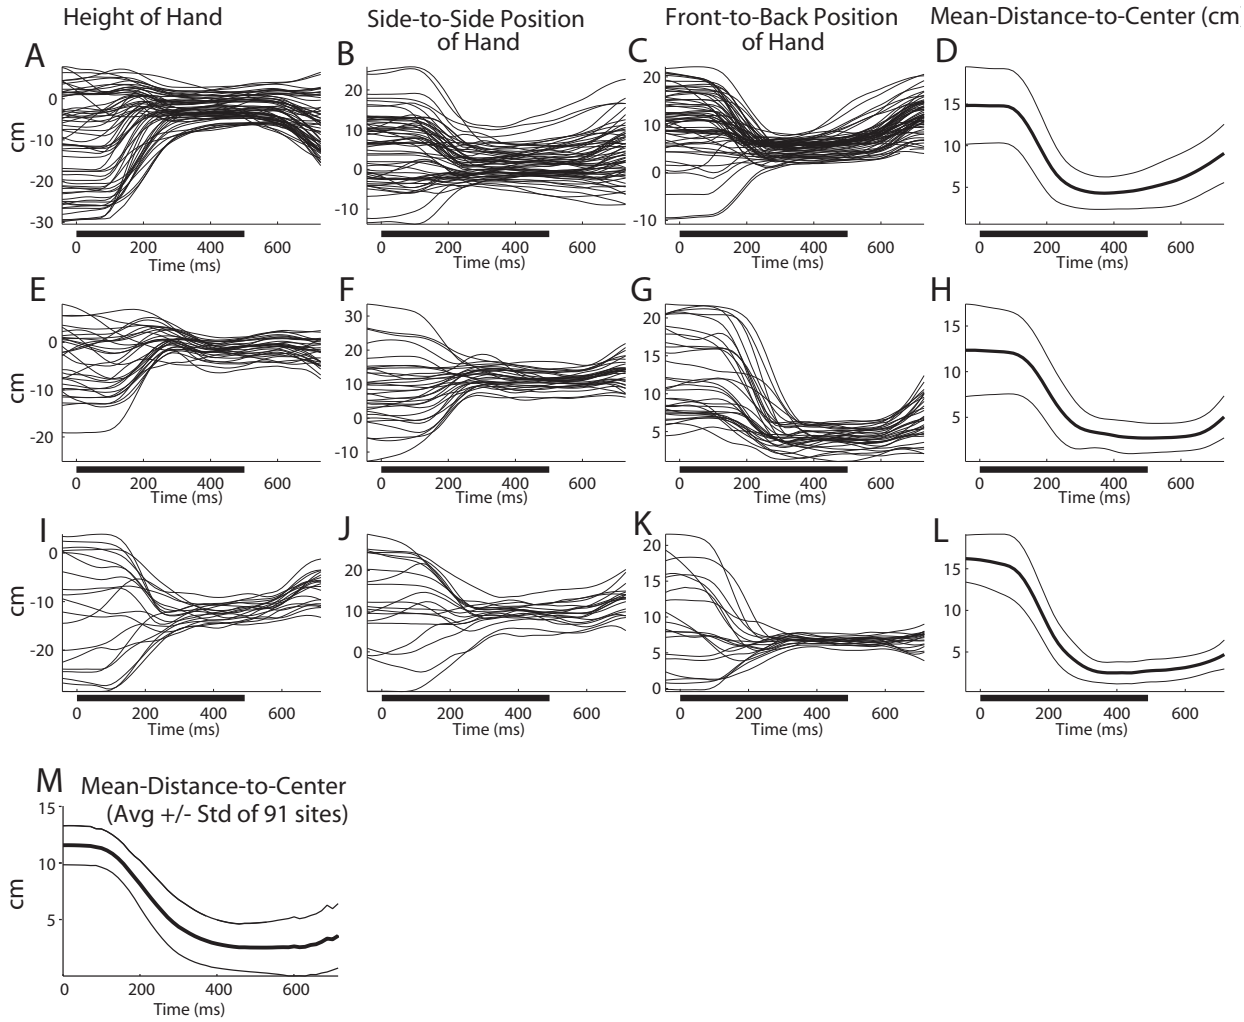


Figure 2-4: Stimulation-evoked convergence of the hand. A1. Height of the hand as a function of time, during stimulation of a typical cortical site. Each line shows data from one stimulation trial. The thick black bar at bottom indicates the time of the stimulation train. A height of 0 indicates the height of the mouth. A2. Data from same stimulation site showing the side-to-side position of the hand. Mouth at 0. A3. Data from same stimulation site showing the front-to-back position of the hand (position along the monkey's line of site). Mouth at 0. Negative numbers indicate positions behind the level of the mouth; positive numbers indicate positions in front of the level of the mouth. A4. Mean-distance-to-center (a measure of spatial dispersion; see explanation in text) +/- standard deviation for this same example site. B and C. Similar plots for two more stimulation sites. D. Group data for 91 sites. Mean and standard error of the mean-distance-to-center.

Effect of a weighted bracelet on the final stimulation-evoked height of the hand

As demonstrated above, for almost all cortical sites, stimulation caused the hand to converge from any initial position toward a final region of space. How is this final position affected by a weight fixed to the arm? We tested sites by fixing a 90 g lead bracelet to the wrist (approximately 25% of arm weight).

Was the peak acceleration significantly reduced by the weight? Each stimulation site was analyzed in the following manner. For each trial, we found the peak acceleration of the hand during the stimulation-evoked movement. We then averaged across trials to obtain the mean of the peak accelerations. We then compared this mean peak acceleration between the weighted and unweighted trials. Figure 2-5 shows group data for the 50 stimulation sites tested. In general, the peak acceleration was lower when the weight was present (Binomial test, $P < 0.0001$). This finding demonstrates that the weight was sufficiently large to measurably reduce the acceleration of the hand during stimulation.

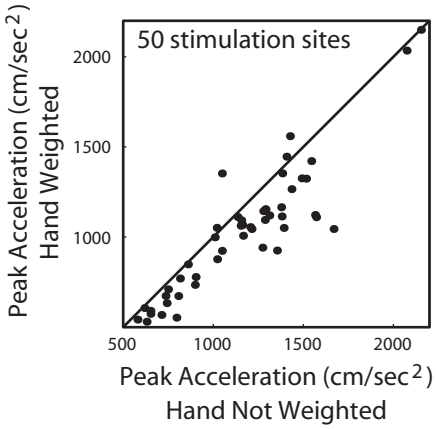


Figure 2-5: Effect of a weighted bracelet on the acceleration of the hand. For each stimulation-evoked movement, the peak acceleration was found. The peak accelerations for different trials were then averaged to find the mean peak acceleration. The mean for the weighted condition (with 90 g lead bracelet) was then plotted against the mean for the unweighted condition. Each dot shows the mean from one stimulation site. Almost all data points are below the y=x line, indicating that the acceleration was smaller when the weight was present. Thus the weight was heavy enough to cause a measurable effect on the movement of the hand.

Was the final position of the hand significantly lowered by the weight? Figure 2-6A shows data from one example site. Figure 2-6A1 shows trials in which no weight was fixed to the arm and Figure 2-6A2 shows interleaved blocks of trials in which the hand was weighted with the 90 g lead bracelet. In both conditions, the hand stabilized at a similar final height by the end of stimulation. Thus the weight did not appear to have a pronounced effect on the final, stimulation-evoked height of the hand, for this particular stimulation site.

We further analyzed the results by examining how the weight affected the position of the hand at different time points throughout the stimulation trial. For each time point, we calculated

the mean difference in height between the weighted condition and the unweighted condition. This time course is plotted in Figure 2-6A3. At the start of stimulation, the mean height of the hand was approximately the same for both weighted and un-weighted conditions (a difference of approximately 0). As the stimulation progressed, the weight began to affect the hand as compared to the unweighted condition. This difference reached a maximum about 300 ms after the start of stimulation, at which time the hand was on average 1.7 cm lower in the weighted condition than in the unweighted condition. This effect of the weight then began to diminish. By the end of stimulation, the height of the hand was once again similar for the weighted and unweighted condition. Thus the weight did affect the movement of the hand during the trial, but by the end of the stimulation train, the weighted hand had reached the same height as the unweighted hand.

We tested whether the effect of the weight on the final height of the hand was statistically significant, in the following manner. We first plotted the final height of the hand (at time 500 ms, the end of stimulation) against the initial height of the hand (at time 0 ms, the start of stimulation). This plot is shown in Figure 2-6A4. We then used an Analysis of Covariance, in which the initial height of the hand was the covariant variable, the final height of the hand was the dependant variable, and the presence or absence of the weight was the independent variable. In this example, the final height of the hand was not significantly different for the weighted and unweighted conditions ($F=0.001$, $P=0.971$). Thus for this particular stimulation site, the hand moved to a similar final height regardless of whether the weight was present or absent.

Figure 2-6B shows similar data from another example site for which the weight did not significantly affect the final height of the hand. Figure 2-6C shows data from a stimulation site for which, unlike the previous two examples, the final position of the hand was significantly reduced by the weight.

A total of 50 stimulation sites were tested with and without a weight fixed to the wrist. For each stimulation site, we calculated the mean distance that the hand was dragged down by the weight at the end of the stimulation train. We calculated whether this effect of the weight was significant using an Analysis of Covariance as outlined in the examples above. We plotted this data on a frequency histogram, shown in Figure 2-6D. Of the 50 sites, 13 (26%) showed a significant drop in hand height caused by the weight and 37 (74%) did not (significance level of 0.05, Bonferroni-adjusted for 50 tests). No sites showed a significant increase in hand height caused by the weight. We also found that no sites showed a significant horizontal shift of hand position due to the weight. Since the weight adds only a vertically downward component of force, a systematic horizontal shift in hand position was not expected. There was no clear topography of sites that showed an effect of the weight and sites that did not; both types of sites were intermingled in cortex. There was also no clear distinction between the types of evoked movements that compensated for the weight and those that did not. There was a trend for final hand positions that were far from the body, with the arm outstretched, to show a smaller effect of the weight (greater compensation). This correlation between hand distance from body and effect

of the weight did not reach significance (mean difference between weighted and unweighted final height vs mean final distance of the hand from the chest, r=0.19, F=1.82, P=0.184).

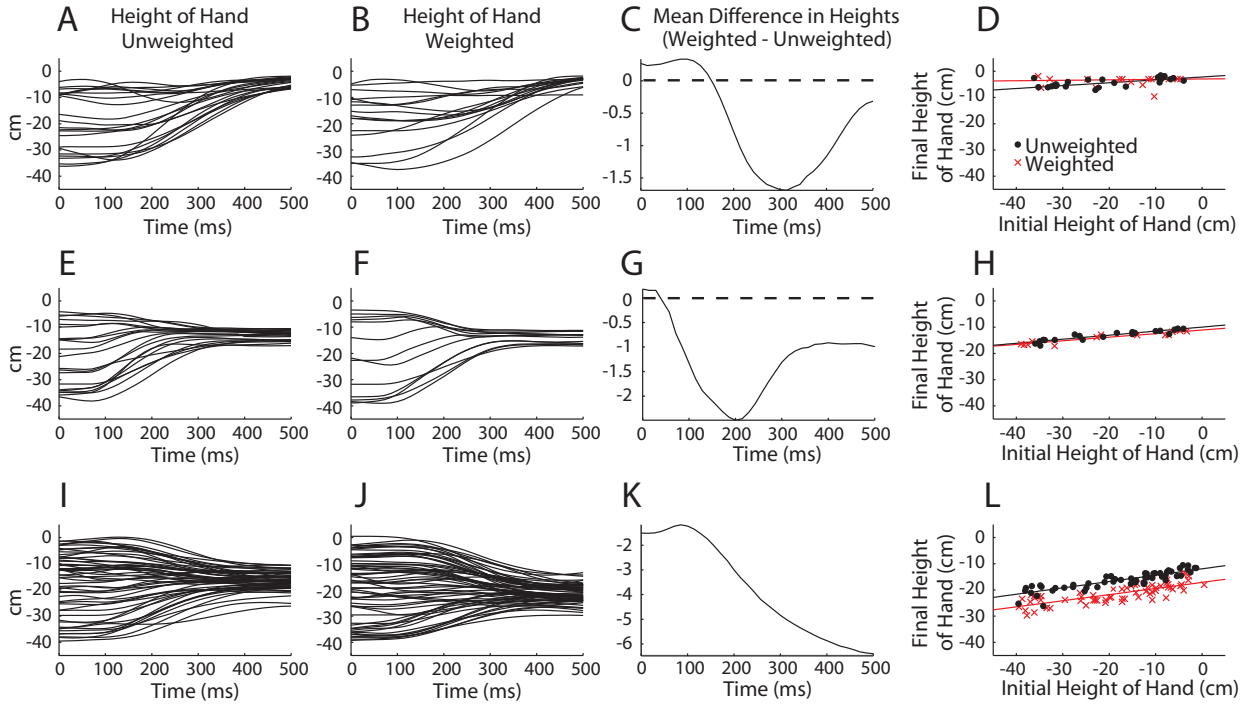


Figure 2-6: Effect of a weighted bracelet on the height of the hand. A1. The height of the hand as a function of time during stimulation for a typical site. Stimulation began at time 0 and ended at time 500. Each trace shows data for one stimulation trial. No weight was fixed to the wrist on these trials. A2. Data from the same cortical site when the weight was fixed to the wrist. A3. Difference in the mean height when the weight was present and when the weight was absent. A difference of 0 indicates no change in height caused by the weight. A negative difference indicates that the hand was lower when the weight was present. A4. The final height of the hand (at time 500) as a function of the initial height of the hand (at time 0). The final height was lower for the weighted condition (red crosses) than for the unweighted condition (black dots). B. Similar plots for a second stimulation site. C. Similar plots for a third stimulation site.

Was the final position of the weighted hand higher than the predicted position? One potential concern with the above test is that the weight may be too small to have a reliable effect on the height of the hand, whether or not compensation is present. We used a mathematical model to approximate the expected effect of a 90 g weight on the movement of the arm assuming no compensation for the weight. We modeled the physics of a two-link arm with three degrees of freedom of shoulder rotation and one degree of freedom of elbow rotation. The model incorporated spring-like muscle forces, damping forces, gravitational forces, and the inertia tensor of the two-link arm. The model is approximate; it provides a rough estimate of the expected effect of the weight, which can then be compared to the actual effect of the weight.

Figure 2-7A-D shows the results of this analysis for one example site. Figure 2-7A shows the measured, actual trajectories of the hand when no weight was present. We used the data from these trials including joint angle, angular speed, and angular acceleration, and applied a least-squares method of optimal fit to estimate the spring constants and damping terms that best characterized the movements during these trials. We incorporated those dynamic terms into the model. Figure 2-7B shows the simulated trajectories obtained from the model, using the same initial positions and speeds as in the actual trials. The model closely matched the actual data, indicating that the fit algorithm was successful. Figure 2-7C shows the measured, actual trajectories of the hand when the weight was added to the hand. Again, the hand converged on approximately the same final height. In this case, the weight appeared to affect the hand initially, pulling it to a slightly lower position, but in the second half of the trial the hand rose up

to a similar final height as without the weight. The mean final height of the hand was 0.4 cm higher in this weighted condition (Figure 2-7C) than in the non-weighted condition (Figure 2-7A). We then ran the model with an added 90 g weight on the wrist and obtained the calculated trajectories of the hand, shown in Figure 2-7D. In the model, the hand converged to a position that was on average 5.8 cm lower than in the actual data. This site is typical in that the model predicted a pronounced effect of the weight.

We compared the data from Figure 2-7C (measured data with weight added) and Figure 2-7D (simulated data with weight added) using an analysis of covariance. The initial height of the hand was the covariant variable and the final height of the hand was the dependant variable. The final height of the hand was significantly higher in Figure 2-7C than in Figure 2-7D ($F = 73.4$, $P < 0.0001$). Thus there is evidence of an active compensation for the weight, in that the electrical stimulation lifted the hand to a significantly greater height than was estimated if no compensation were to occur.

A total of 50 stimulation sites were tested in this fashion and 47 (94%) showed a significant difference (significance level of 0.05, Bonferroni-adjusted for 50 tests). For these sites there was evidence of active compensation for the weight. For 3 sites (6%), however, the weighted hand was not lifted significantly higher than estimated for the case of no compensation.

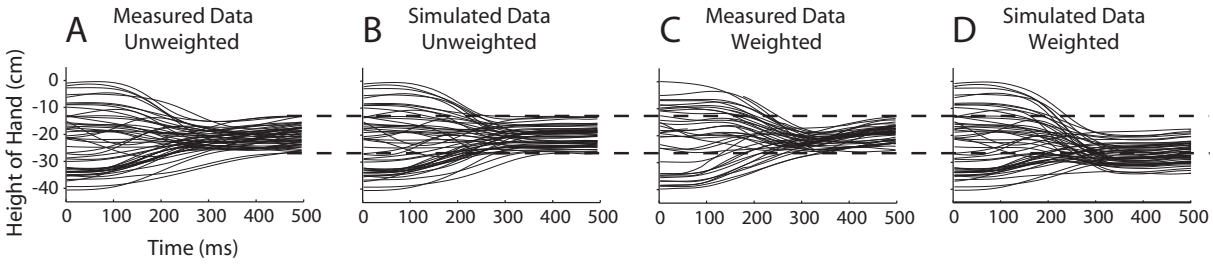


Figure 2-7: Estimating the expected effect of a weighted bracelet on the height of the hand. A. Data from an example site. Y axis shows height of hand relative to mouth height, X axis shows time during stimulation trial. Each trace shows data from one trial. Dotted lines show the range (min and max) of final heights. B. We modeled the physics of the arm and used the data from the trials shown in A to find a best fit for the spring constants and damping forces in the model. The graph shows the calculated trajectories for the model. C. Data from the same example site as in A. On these trials, a 90 g weight was fixed to the wrist. D. The calculated trajectories based on the model, when a 90 g weight was added to the wrist in the model.

Movement of individual joints

In order to probe the movement of the limb to a posture in greater detail, we measured the angles of joints in the arm and the grip aperture. Figure 2-8A shows the results for one stimulation site. Eight graphs are shown, corresponding to the eight degrees of freedom that were measured. During stimulation, the joints rotated from a disparate range of initial angles to a more restricted range of angles. For some joints, such as the elbow, the rotation was to an intermediate angle. When the elbow joint was initially flexed, stimulation caused it to extend toward the final angle; and when it was initially extended, stimulation caused it to flex toward the final angle. For other joints, such as the shoulder azimuth, the rotation was in one direction. For the wrist adduction angle, relatively little range of movement was obtained during

stimulation. This wrist angle generally showed little change compared to other joint angles, even during the monkey's spontaneous movements. Figure 2-8B shows joint angles from another example site.

For each site, we calculated the mean stimulation-evoked angle for each joint. Figure 2-8C shows a frequency histogram of all sites. A range of angles was represented, including extreme angles and intermediate angles.

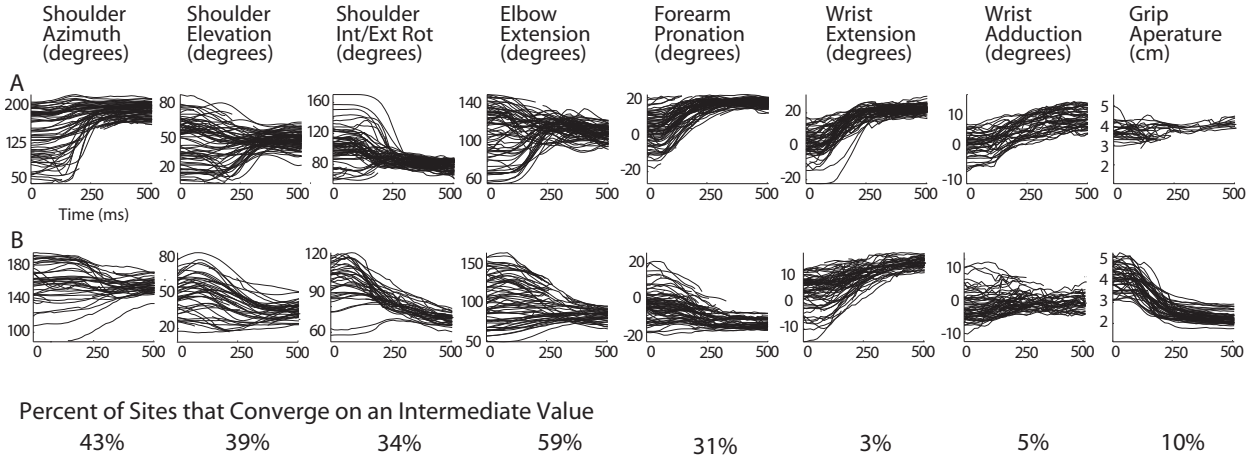


Figure 2-8: Tracking eight degrees of freedom in the arm during stimulation. A. Results for one stimulation site. Each trace shows data for one trial. X axis shows time during stimulation; 0 = start of stimulation, 500 = end of stimulation. Y axis shows joint angle (degrees) or grip aperture (cm). B. Data from a second example site. C. Each graph is a frequency histogram showing the distribution of mean final angles across the 61 stimulation sites tested.

Interaction between joints stabilizes the hand for some stimulation sites

One hypothesis is that stimulation independently drives each joint to a specific angle, and as a result of this aggregate of joint angles, the hand moves to a location in space. Another hypothesis is that the joints move in a coordinated manner, each joint adjusting to compensate for slight deviations in the other joints, in order to bring the hand more specifically toward a desired location. We asked whether the trial-by-trial variability in each joint might compensate for the variability in the other joints in a manner that would help to stabilize the hand in space.

Figure 2-9A shows the results for one example site. For this analysis, we concentrated on four arm angles: shoulder elevation, shoulder azimuth, shoulder internal/external rotation, and elbow flexion. These arm angles, together with the lengths of the arm segments, define the position of the wrist in space. For each of the 37 trials, using the angles reached at the end of stimulation and applying forward kinematics, we calculated the final position of the wrist. These final wrist positions are plotted as black dots in Figure 2-9A.

We then randomly "shuffled" the trials in the following manner. A shuffled trial might contain the shoulder elevation angle from trial 1, the elbow flexion angle from trial 16, etc. The rule for a shuffled trial was that none of the four joint angles that composed the shuffled trial had been collected on the same actual trial. In this fashion, 37 randomly shuffled trials were constructed. We used forward kinematics on these shuffled trials to calculate final wrist positions. If the joints seek their final angles independently, then shuffling the trials in this manner should have little effect on the result. However, if the joints normally interact within a

trial, such that a slight deviation in one joint is compensated by slight deviations in the other joints, then shuffling the trials will remove this inter-joint compensation and result in a wider distribution of final wrist positions. As shown in Figure 2-9A, the shuffled trials (open circles) did show a wider distribution of final wrist positions than the actual trials (filled circles). As a measure of the spread of final positions, we used the mean-distance-to-center. For this example site, the mean-distance-to-center for the shuffled trials was 6.03 cm and the mean-distance-to-center for actual trials was 3.33 cm.

Figure 2-9A shows only one possible re-shuffling of the 37 trials. We tried 40,000 possible re-shuffles, and for each one calculated a mean-distance-to-center. Figure 2-9B shows the distribution. The arrow shows the position of the actual data on this distribution. Of the many ways to re-shuffle the data, all of them resulted in a greater spread of final hand positions than the actual data ($P < 0.0001$). These results show that within a stimulation trial, the joint angles were not independent but instead interacted. Deviations in some joints must have been matched by compensatory deviations in other joints, in a manner that helped to stabilize the hand at a particular location in space.

A similar analysis was performed for each stimulation site. Of 61 sites tested in this manner, 54% had a significant effect similar to the effect shown in Figure 2-9, in which interactions between the joints stabilized the stimulation-evoked hand location (significance level of 0.05, Bonferroni-adjusted for 61 tests). For 46% of the sites, no significant effect of joint interaction was obtained. Thus only approximately half of the sites showed this stabilization of

the hand in space. There was no clear topography of sites that showed significant joint interactions and sites that did not; both types of sites were intermingled in cortex.

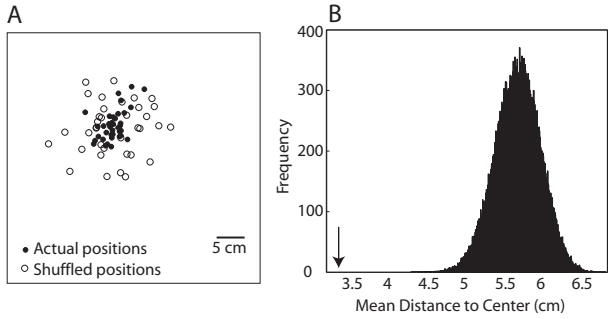


Figure 2-9: Interactions between joints stabilized the hand position. A. Data from one example site. Black dots show final hand positions for 37 stimulation trials, calculated from the measured joint angles. Open circles show final hand positions calculated from 37 "shuffled" trials, using the same data but with the joint angles randomly shuffled across trials. B. 40,000 different shuffles of the 37 trials were tested. For each shuffle, the 37 final hand positions were found through forward kinematics and the mean-distance-to-center was calculated (a measure of the spatial spread of hand positions). These mean-distance-to-centers were then plotted on the frequency histogram shown. The mean-distance-to-center for the actual data is indicated by the arrow.

Movement categories arranged in separate cortical zones

Five categories of movements and their locations on the cortex are shown in Figure 2-10A.

Hand-to-mouth. At some stimulation sites we evoked a characteristic hand-to-mouth movement, defined by the following properties: the grip aperture closed during stimulation as

measured by the optotrak, the hand moved to a location within 1 cm of the mouth as measured by the optotrak, and the mouth opened.

Defensive-like. Another cortical zone could be delineated on the basis of the following specialized properties: the neurons responded to tactile stimuli on the face and arms and visual stimuli near or approaching the tactile receptive fields; and stimulation evoked defensive-like movements including a blink, squint, lifting of the upper lip in a facial grimace, shrugging of the shoulder, and sometimes a blocking movement of the arm. These defensive-like movements were noted here qualitatively and have been extensively studied in a more quantitative fashion in previous experiments (Cooke and Graziano 2004a, 2004b).

Central space + finger movement. A specific type of movement involving both the arm and the hand satisfied the following criteria: the hand moved into a restricted region of space, within 10 cm of a central point in front of the chest (12 cm vertically below the chin); and the grip aperture changed, as measured by the optotrak. Qualitatively, these finger movements resembled an apparent precision grip (thumb against forefinger), a power grip (fist), or a splaying of the fingers accompanied by a turning of the palm toward the face.

Reach. For some cortical sites, the hand moved to a location distant from the body (greater than 10 cm from the central point just in front of the chest) and the grip aperture opened. Subjectively these movements resembled reaching to grasp.

Other outward arm movements. At many sites, stimulation drove the hand to a distal location (greater than 10 cm from the central point just in front of the chest), but without evoking

any opening of the grip aperture. These movements were not clustered in a single zone but instead were scattered, surrounding the reaching sites and the central space/manipulation sites.

Climbing/leaping. In a large medial and anterior region, stimulation evoked climbing-like movements that included all four limbs and the torso. For example, during stimulation of a climbing-like site the monkey might raise both arms placing the hands into upper lateral space, push down on the contralateral leg as though to push off the ground, and raise the ipsilateral leg.

No evoked movement. In monkey 1, in an anterior region of cortex, stimulation did not evoke movements even when the current was raised to 300 microamps.

Topography of hand location across the cortical surface

Figure 2-10B shows one type of topography involving hand position. The stimulation sites were categorized according to the proximity of the hand to a central point just in front of the chest, 12 cm vertically below the chin. Sites for which the hand moved into the region of space within 10 cm of this central point were clustered in a posterior region that extended partly onto the bank of the central sulcus. As described above, for almost all of these sites the fingers also shaped in a specific manner. Sites for which the hand moved outside of this central region were scattered over the remainder of the cortical area.

Figure 2-10C shows another dimension of topography with respect to hand position. The sites were categorized according to their vertical position into three spatial regions separated by

16 cm: high (crosses), middle (open circles), and low (closed circles). The most lateral part of the arm representation was dominated by hand locations in upper space. For electrode penetrations that were more medial along the central sulcus, the hand locations were usually in mid-level space, including hand locations in front of the chest. In the most medial sites along the central sulcus, stimulation sometimes drove the hand into lower space. When the electrode entered a medial and anterior region where both the leg and the arm were represented, the organization was less clear, with hand positions in upper, middle, and lower space intermixed.

We tested the statistical significance of the progression of hand heights along the central sulcus as follows. For each site not within the medial, leg-and-arm representation, we calculated the location of the site as projected onto a line that ran along the central sulcus. We then used an ANOVA to compare the distributions along this line. The separation was statistically significant between upper and middle hand positions ($P < 0.0001$) and also between middle and lower hand locations ($P < 0.0001$). It is important to point out, however, that although the mapping of hand height along the central sulcus is statistically significant, it is also noisy and contains overlap.

We examined other spatial parameters but found no other consistent topography. Figure 2-10D shows the distribution of hand locations in a lateral dimension (contralateral, central, and ipsilateral to the electrode). In Monkey 1, some clustering but no overall topography is apparent. In Monkey 2, some possible topography can be seen in which lateral hand positions were found more anteriorly and central hand positions were found more posteriorly.

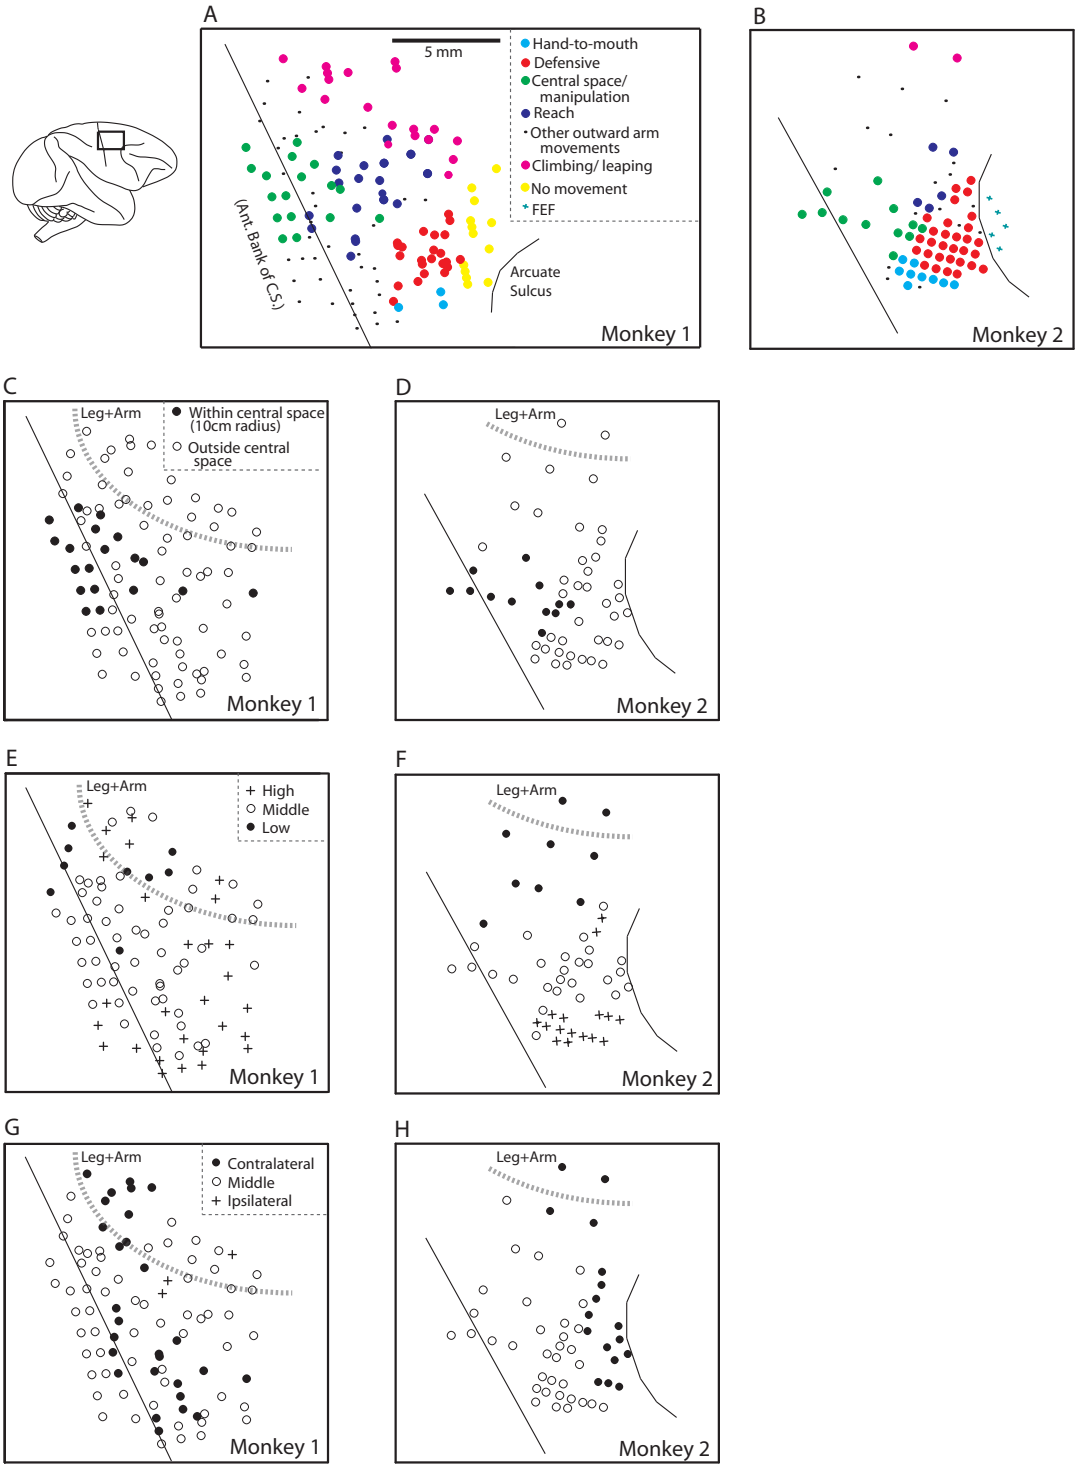

A
5 mm
Hand-to-mouth
Defensive
Central space/ manipulation
Reach
Other outward arm movements
Climbing/ leaping
No movement
FEF
(Ant. Bank of CS.)
Arcuate Sulcus
Monkey 1
B
Monkey 2
C
Leg+Arm
Within central space (10cm radius)
Outside central space
Monkey 1
D
Leg+Arm
Monkey 2
E
Leg+Arm
High
Middle
Low
Monkey 1
F
Leg+Arm
Monkey 2
G
Leg+Arm
Contralateral
Middle
Ipsilateral
Monkey 1
H
Leg+Arm
Monkey 2

Figure 2-10: Topographic arrangement of stimulation effects. Rectangle on schematic brain shows approximate location of studied cortex. See Methods for details on locating the central and arcuate sulcus. A1. Clustering of different categories of movement in Monkey 1. Diagonal line in map indicates lip of central sulcus, and area to left of line indicates unfolded cortex in anterior bank of sulcus. Curved line indicates location of arcuate sulcus as visualized during surgery. Sites are color-coded according to type of complex movement evoked. See text for categorization criteria. A2. Similar map for Monkey 2. Fewer sites were tested in Monkey 2 B1. Sites in Monkey 1 were categorized according to whether the hand moved to a location within a 10 cm radius of a central point just in front of the chest (central point 12 cm vertically below the chin) or moved outside of that central region of space. The dotted line indicates the approximate border between the arm representation (below the line) and a combined leg and arm representation (above the line). Fewer sites are shown here than in A because some of the sites in A, such as many of the sites related to defensive movements, did not involve clear movements of the contralateral arm. B2. Similar map for Monkey 2. C1. Sites in Monkey 1 were categorized according to whether the hand moved to upper space, middle space, or lower space. C2. Similar map for Monkey 2. D1. Sites in Monkey 1 were categorized according to whether the hand moved into the space contralateral to the electrode, into middle space, or into ipsilateral space but showed no clear organization. D2. Similar map for Monkey 2.

Discussion

The present experiments examined arm movements evoked by electrical stimulation in motor cortex using half-second stimulation trains. The movements had a consistent latency of about 80 ms. They generally followed a bell-shaped speed profile in which the peak speed co-varied with movement distance, similar to the profiles of natural movements (e.g. Bizzi and Mussa-Ivaldi 1989, Flash and Hogan 1985). Stimulation with different frequencies had a similar effect, although higher frequencies were usually associated with slightly faster movements. Stimulation moved the arm from any initial position toward a final position. This convergence began with a latency of about 80 ms, suggesting that as soon as the arm began to move, it was already directed in a coherent fashion toward the final position.

Are the stimulation-evoked movements open-loop or closed-loop?

Stimulation drives the arm into a final invariant posture. Stimulation recruits flexor muscles or extensor muscles as appropriate to move the arm from its initial configuration to the final posture (Graziano, Patel, and Taylor 2004) suggesting that cortical output is remapped by sensory input. Here we asked whether such feed-back dependent remapping of cortical stimulation could lead to appropriate compensation for a lead weight attached to the wrist. Because the weight was concentrated at the wrist it was predicted to have a large effect on the movement of the arm. The weight did significantly reduce the acceleration of the movement. Yet for some cortical sites (74%), stimulation ultimately drove the hand to a similar final height as without the weight.

We also found that for some sites (54%), during stimulation, compensatory interactions between the joints helped to stabilize the position of hand and to bring it into a restricted range of positions. This result also suggests some possible refinement of the movement on the basis of feedback, in which the joint angles are adjusted to bring the hand to the correct location. It is worth noting, however, that the complex musculature of the arm makes it difficult to fully interpret this result. Is this effect the result of high-level coordination, or the result of some unanticipated, low-level muscle and torque interactions? One aspect of the result may be relevant to this question. We found the effect to be highly significant for some stimulation sites, and found it to be absent for other stimulation sites, even though similar postures were evoked in both cases. This differing effect for different cortical sites suggests that the stabilization of the

hand by means of joint interactions may have been a consequence of neuronal factors caused by stimulation of particular cortical sites, and not a necessary consequence of the musculature of the arm. However, other interpretations of this result may be possible.

Motor cortex topography may be determined by multiple competing factors

The topographic organization of motor cortex has been described in numerous previous studies (e.g. Donoghue, Leibovic, and Sanes 1992, Dum and Strick 2002b, Foerster 1936, Fritsch and Hitzig 1960, Kwan et al. 1978, Penfield and Boldrey 1937, Strick and Preston 1978, Woolsey et al. 1952). Most descriptions include a rough somatotopic organization with overlap between the representations of different body parts, some fractures in the representations, and some re-representations. These descriptions suggest that there may be other variables influencing the organization of motor cortex in addition to somatotopy. The current study, along with previous work in the lab (Graziano, Taylor, and Moore 2002), has shown the presence of at least three types of organization: a rough somatotopic map of the body, cortical zones representing movements from the monkey's behavioral repetoire, and, noisy maps of hand position in space. In chapter 4 we present a computational model that provides a parsimonious account of the observed organization of motor cortex.

Chapter 3 : Relationship between unconstrained arm movements and single neuron firing in the macaque motor cortex

Introduction

Neurons in the motor cortex of the monkey brain are active during arm movements. This activity is correlated with a range of movement parameters including force, muscle activity, the direction, speed, and position of the hand, joint rotation, and other aspects of movement (Ajemian et al. 2008, Ashe and Georgopoulos 1994, Caminiti, Johnson, and Urbano 1990, Cheney and Fetz 1985, Evarts 1968, Fu, Suarez, and Ebner 1993, Georgopoulos et al. 1992, Georgopoulos, Caminiti, and Kalaska 1984, Georgopoulos et al. 1982, Georgopoulos, Schwartz, and Kettner 1986, Hocherman and Wise 1991, Holdefer and Miller 2002, Kakei, Hoffman, and Strick 1999, 2001, Moran and Schwartz 1999, Reina, Moran, and Schwartz 2001). The search for the "correct" motor parameter is likely an ill-posed experimental question since many different movement parameters are correlated with motor cortex neurons. One possibility is that the motor system uses an "optimal control" method, potentially controlling any parameter, whether spatial, joint-based, or muscle-based, if it is useful for the performance of the current task (Todorov and Jordan 2002).

One limitation in previous studies is a tendency to use highly constrained movement sets and animals that are extensively trained. We recorded from neurons in the motor cortex of untrained monkeys. The arm was free to move spontaneously and naturalistically while the movements were measured. One goal of this study was to compare the results of a naturalistic, untrained movement set to the previous results obtained with structured, trained movement sets. Can the previous findings be replicated in the present movement set, or will removing the behavioral control also remove the previously obtained neuronal tuning curves?

A second goal of this study was to assess the proportion of total neuronal variance attributable to different movement parameters. In most previous experiments, the constrained movement sets were designed to focus on one or a small number of movement parameters. Once all other sources of variance have been removed, minimized, or averaged out of the data, then the particular parameter under study may account for most of the remaining neuronal variance, with an R^2 value that may be as high as 0.9. Such studies, though legitimately addressing a variety of questions, do not address a certain fundamental question. If most sources of neuronal variance are left in the data by using a more natural range of movement, then what proportion of the total neuronal variance can be explained by any single parameter? Are neurons primarily tuned to one parameter, with the majority of their total variance explained by that parameter, or instead are neurons partially tuned to a diversity of parameters, with each parameter capturing a small portion of the total variance?

Long-duration stimulation evokes coordinated movement to a posture. A third goal was to test whether the behavior of single neurons during naturalistic movement could be exlained by tuning to a preferred posture and whether the neuron's preferred posture was matched to the posture evoked by electrical stimulation at the same cortical site.

Methods

All procedures were approved by the Princeton University Institutional Animal Care and Use Committee and the attendant veterinarian and were in accordance with N.I.H. and U.S.D.A. guidelines. We studied the motor cortex in the left hemispheres of two adult male *Macaca fascicularis*.

Surgery

Surgical procedures are described in chapter 2 "Surgery."

Neuronal recording

During the daily recording sessions the monkey sat in a Lexan primate chair with the head restrained by the head bolt. A hydraulic microdrive (Narishige) was mounted to the top of the recording chamber. A steel guide tube (a 22-gauge syringe needle) was lowered through the hole in the skull and into the granulation tissue that lay over the dura. Then the varnish-coated tungsten microelectrode (Frederick Haer Inc., impedance 0.5-2 MOhm) was advanced from the

guide tube through the dura and into the brain. Neural signals were amplified (A-M Systems amplifier, model 1800) filtered (300-5000 Hertz) and recorded at 25,000 hertz. An offline spike-sorting algorithm was used to assign spikes to individual neurons. Typically 1-3 neurons could be reliably isolated on the electrode at one time.

Single neuron activity was sampled at various depths within the motor cortex ranging from the first depth at which any neurons could be isolated to the depth at which the neurons disappeared and the white matter was presumed to begin. A systematic test of different layers of cortex was not attempted in the present experiment. We saw no clear distinction in the tuning properties reported here for shallow or deeper recordings within motor cortex and therefore included all neurons into one analysis population. Neurons were not pre-selected in any way based on response properties. Instead, all neurons that were encountered by the electrode and that could be held long enough for collection of a full data set were included.

Location of recording sites

Neither monkey was sacrificed at the termination of this experiment. For Monkey 1, after the experiment the implant was removed and the brain was scanned in a 3-T Siemens Allegra Magnetic Resonance Imaging (MRI) head scanner using a 16 cm birdcage coil (NMSC-023, Nova Medical, Wakefield, MA). A high-resolution, 0.5 x 0.5 x 0.5 mm anatomical scan of the whole brain was taken (MPRAGE sequence; field of view 128 x 128 mm; matrix 256 x 256; TR = 2,500 ms; TE = 4.4 ms; flip angle 8 degrees). Figure 3-1A shows a surface reconstruction

of the cortex with the studied area in brightened shading. Figure 3-1B shows a sagittal section through the motor cortex with the studied area again in brightened shading.

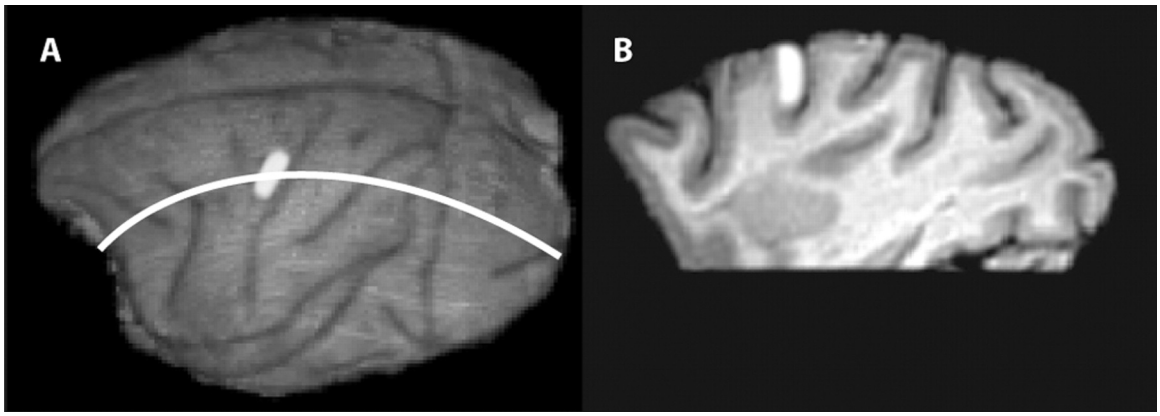


Figure 3-1: Images of the brain of monkey 1 showing the studied area of cortex. A. Magnetic Resonance Imaging (MRI) reconstruction of the cortical surface showing sulcal pattern. Studied region is in lightened shading. The white line indicates the section shown in B. B. Parasagital section (0.5 mm thick) through the studied area of cortex. Studied region is in lightened shading.

To find the motor cortex prior to recording, the central and arcuate sulci were located by shining a bright light on the dura during the initial craniotomy surgery. Both sulci were clearly visible through the dura. The microdrive was then mounted to the recording chamber and the locations of the visualized sulci were measured with the tip of the guide tube. In this way, the locations of the sulci were obtained in microdrive coordinates.

After the craniotomy surgery, during the daily experiments, the measured location of the central sulcus was confirmed to within 0.5 mm by recording and stimulating to either side of the sulcus. Just posterior to the sulcus, in primary somatosensory cortex, we observed the expected

small tactile receptive fields on the contralateral limb and also the expected rarity of effect of intracortical electrical micostimulation even with currents as high as 200 microA (negative leading biphasic pulses, 0.2 ms phase width, 200 Hz, 100-500 ms train durations, 5 – 200 microA current). Just anterior to the sulcus we obtained neuronal responses during hand and arm movement and low microstimulation thresholds, typically less than 20 microA, sometimes as low as 5 microA, as expected from primary motor cortex. The location of the arcuate sulcus was confirmed by microstimulating just anterior to it and obtaining no skeletomotor movements, but instead obtaining stimulation-evoked saccadic eye movements presumably in the frontal eye fields.

The sites tested were located in the arm representation in motor cortex, and were within the anterior bank of the central sulcus or on the cortical surface within 2 mm of the central sulcus. They therefore lay within the boundaries of traditional primary motor cortex. Further information on the location of sites within the functional map was provided by the analyses of the joint angles to which the neurons were responsive, as described below.

Measurement of limb position

The procedure for measuring limb position is described in chapter 2 “Measurement of limb position.“

Description of movements in the data set

During testing of a neuron the monkey was allowed to move its contralateral arm freely to touch and explore parts of the primate chair, to reach for small pieces of fruit held out on the end of forceps, to bring food to its mouth, to retrieve food from its mouth, to hold and examine food in central space, and to rotate and explore food items. Occasionally the monkey also scratched at its skin, scratched rhythmically at a portion of the monkey chair, or attempted to scratch the experimenter with a fast semi-ballistic arm movement. The movement of the arm was recorded through all of these behaviors. Different types of behaviors were not separated in the analysis, partly because one type of behavior tended to grade into another type and the distinction between behaviors could only be made subjectively; and partly because the purpose of the study was to include all arm movements in as large and naturalistic a range as possible given the constraints of the primate chair.

For each neuron, the position of the hand in 3-D space was measured every 14.3 ms during a continuous time interval ranging from 10 to 30 min. The hand position data was smoothed using a cubic spline with a smoothing coefficient of 0.05. The instantaneous hand speed was then calculated using a three-point running average to obtain a smoothed speed profile. Separate movements were extracted from the data set on the basis of a speed analysis. Minima in the speed were identified and the intervals between minima were flagged as potential separate movements. To enter the final data set, the movement had to be at least 150 ms in duration and the peak speed had to be at least 20 cm per sec. These parameters seemed to

successfully divide the data into discrete segments that matched our subjective impression of separate hand movements. The average speed profile is further described in the Results (*Speed Tuning I*).

The various analyses in this study focus on the neuronal activity associated with periods of arm movement as defined by the velocity analysis described above. We originally attempted to analyze both the periods of movement and the interleaved periods of non-movement, but found that neurons were frequently inactive during periods of hand stasis between movements. This lack of activity may be related to the spontaneous nature of the movements. There may have been little planning between movements. Sometimes the arm was relaxed or braced against a part of the chair between movements, and this state of relaxed stasis could not be distinguished from more active stasis in the recorded data. Often the animal performed a mouth movement or a foot movement between arm movements, perhaps drawing attention and motor planning to a different body part. These reasons may have contributed to a reduction in and a variability of neuronal activity between movements. The analyses reported here are therefore focused on the periods of arm movement.

Figure 3-2A shows a typical movement set collected during the testing of a neuron. This set shows 514 separated movement segments that densely sampled the workspace of the hand. Vertically, the movements ranged from 29.7 cm below the mouth to 7.8 cm above the mouth. Horizontally the movements ranged from 19.4 cm on the contralateral side (same side as the studied arm, opposite side to the electrode) to 16.6 cm on the ipsilateral side. In depth (direction

along the monkey's forward line of sight) the movements ranged from 7.5 cm behind the level of the mouth (such as when the monkey reached to its flank or to its ear) to 20.2 cm in front of the mouth (normal for a fully extended reach). The range of movement starting points was not significantly different from the range of ending points.

The average length of a movement was 8.0 cm (standard deviation = 4.9). The average duration was 351 ms (SD 133 ms). The average hand speed was 23.74 cm/sec (SD 9.9). The average peak speed among all movements was 40.7 cm/sec (SD 20.4). The average time between movements was 4.06 sec. This pause between movements ranged from less than 1 sec (40% of pauses, representing brief periods of stasis during ongoing arm movement) to longer than 10 sec (15% of pauses, representing periods when the monkey had stopped making arm actions and was stationary or engaged in movement of other body parts).

For each movement we calculated a standard curvature metric as follows. The straight-line distance between the start and end of the movement was found. The total path-length of the movement was found. The ratio of these two quantities provided a curvature metric in which 1.0 corresponds to no curvature and smaller numbers correspond to increasingly curved movements. The average curvature was 0.91 (SD 0.09), indicating that the movements tended to be straight. In Figure 3-2A, some movements appear to be highly curved. This appearance is a result of collapsing a 3-dimensional movement into a 2-dimensional depiction in which the long axis of the movement is not fully shown.

For each movement we calculated a direction by connecting the start point to the end point and obtaining the azimuth and elevation angles. The distribution of these directions was then examined. The directions appeared to be relatively evenly distributed. The sphere of all possible directions was divided into 20 equal sectors, and the movement directions were distributed over these 20 sectors with all sectors represented. Thus the unstructured movement set was well behaved, followed expected characteristics of normal movement, and included a diverse range in terms of workspace, length, speed, and direction.

Inevitably the movement set included some non-uniformities in the distribution of hand positions and joint angles. For example, hand positions near the face were over-represented. Hand positions in lower space were less well represented. Because the movement sets were spontaneous, we were not able to enforce an even distribution of all types of movement. As a result, there is some concern whether the data set is diverse enough, and the movement parameters independent enough, to allow for meaningful regression analyses. One approach to this issue is addressed in a section of the Results (*Simulated neurons*). We simulated noisy neurons that were tuned in a variety of ways, calculated the expected response of these neurons during the actual recorded movement sets, and then applied the regression models to the results. We found that the movement sets were sufficiently diverse to allow the regression analyses to uncover the correct tuning curves.

Non-uniformities in the movement set present a particular problem for separating the effect of direction tuning from the effect of hand position tuning. For example if a neuron is

tuned to upward motion of the hand, it may appear to prefer an upper hand location simply because that location in space can only be reached by upward motion. This possible confound was addressed partly through the use of the simulated neurons (simulated direction-tuned neurons did not evidence a high degree of tuning to hand position) and partly by means of the analysis of neuronal data in the results section, showing that the preferred direction does not easily account for the hand position tuning (see *End-point tuning*).

Non-uniformities in the movement set may also result in correlations among the different joints of the arm. Such inter-joint correlations would be problematical particularly in the analysis for neuronal tuning to final posture. A neuron may actually be related to one joint, but due to inter-joint correlations, appear to be related to other joints as well. It was therefore important that the movement set contain enough independence among joints that the contributions of each joint could be separated through regression analysis. Figure 3-2B shows the correlations among the joint angles reached at the end of the movements. The joints were not tightly coupled. Some showed correlations with each other, as expected given the known covariance of joints during movement; but the scatter was sufficient to allow a regression analysis to separate the different contributions. This lack of any strong coupling among joints was probably caused by the great diversity of movements in the set, and is presumably the reason for the success of the simulated neuron analysis in which simulated tuning functions could be successfully recovered using our regression analyses. This issue is further addressed in the Results.

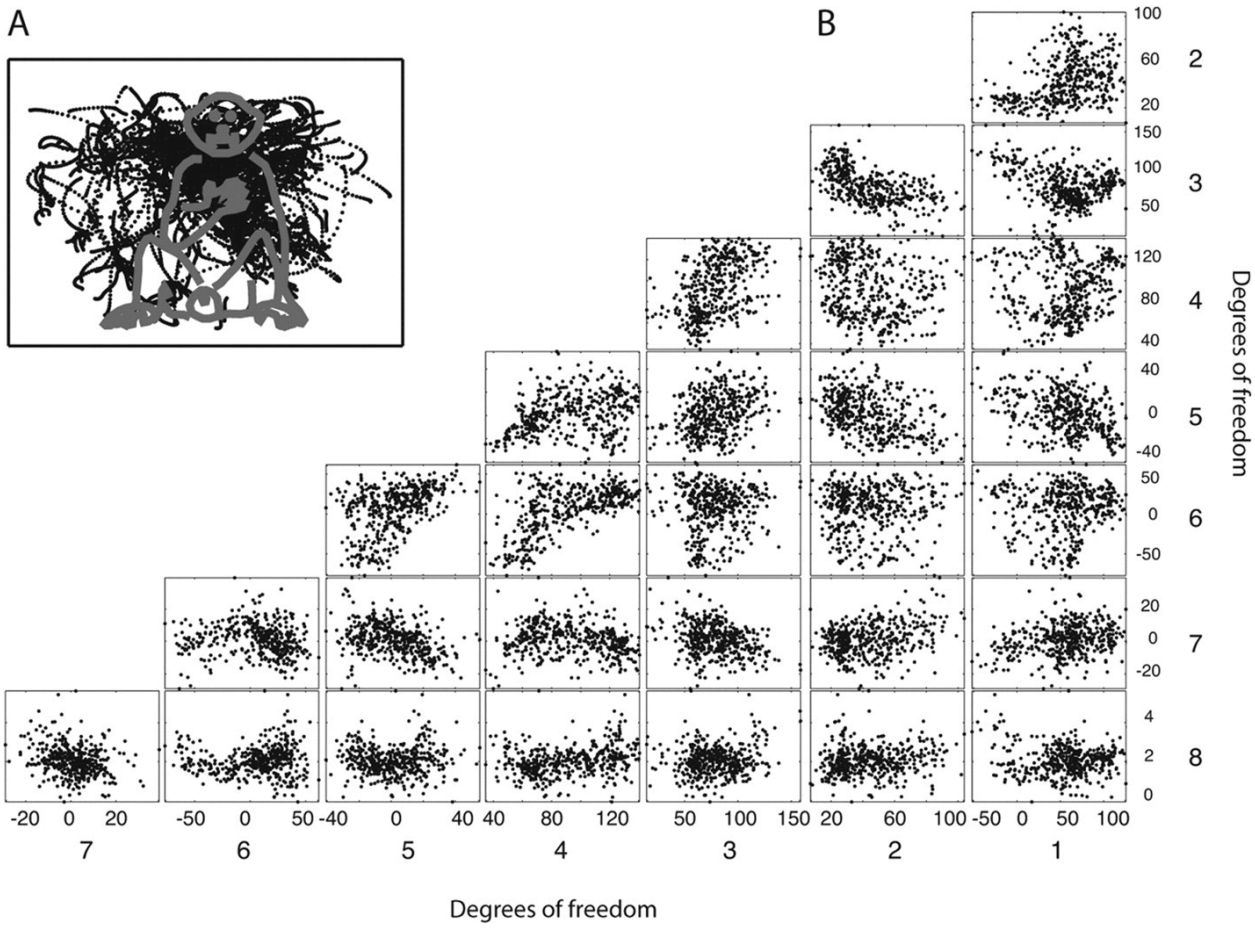


Figure 3-2 : Range of movements in the naturalistic movement set. A. Front view of 514 hand movements made during 10 min of testing one neuron. Each trail of dots = 1 movement measured at 14.3 ms intervals. Frame is 45 cm tall. B. Correlations among the final positions of the eight degrees of freedom among the recorded movements. Degrees of freedom are: 1 = shoulder azimuth; 2 = shoulder elevation; 3 = shoulder internal/external rotation; 4 = elbow extension; 5 = forearm pronation; 6 = wrist extension; 7 = wrist adduction; 8 = grip aperture. Units on X and Y axis of each plot are degrees of joint angle, except for degree of freedom 8, which is expressed in cm.

Preliminary analysis to specify somatotopic location

To further confirm the somatotopic portion of motor cortex that was studied, we performed a preliminary analysis on each neuron. Using a step-wise regression, we obtained the degree of correlation between the neuronal activity and the velocity of each of the 8 measured joints. If we were recording primarily in a distal representation, we would expect to find significant regressions with distal joints including hand aperture, wrist flexion, wrist adduction, and forearm pronation. If we were recording primarily in a proximal representation, we would expect to find significant regressions with promixal joints including elbow flexion and the three degrees of shoulder rotation. Given the known overlap in motor cortex somatotopy, we expected to find neurons related to both proximal and distal joints. The results indicated that 89% of the neurons were significantly related to the proximal joints and 67% were significantly related to the distal joints. These results indicate that the studied neurons were in the forelimb representation in a region that emphasized the proximal joints over the distal joints but represented both.

EMG recording

Fine insulated stainless steel wires were threaded into a 22-gauge syringe needle and inserted into the long biceps. The wires had an exposed tip of 1-2 mm. Three wires spaced about 5 mm apart were inserted to provide input to a differential amplifier and its ground (single

neuron amplifier model 1800, A-M Systems, Sequim, WA). The amplifier filters were set with a low cutoff at 300 Hz and a high cutoff at 5000 Hz. Though some studies use a lower level for the low cutoff (e.g. Cheney et al. 1985), we found that a 300 Hz cutoff was helpful to remove low frequency noise. The EMG signal was digitized at a rate of 20,000 Hz, rectified, and integrated in 14.3 ms time bins.

Regression models

Each neuron was tested for a range of possible tuning functions. Each model of a tuning function was compared to the behavior of the neuron by means of a regression analysis. Some models, such as tuning to speed, involved one free parameter. Other models, such as tuning to posture, involved several free parameters. To avoid inflating the R^2 value with the addition of more parameters, we used the standard Adjusted R^2 metric that takes into account the number of regressors (Cohen et al. 2003). Also following standard practice, we performed a regression on a neuron's data only if the number of data points was twenty or more times the number of regressors, thus minimizing the risk of over-fitting the data. Finally, as a validity check, we performed an analysis on simulated neurons (see Results, *Simulated neurons*). For example, we tested a simulated neuron that was direction tuned. When tested on an actual, recorded movement set and with our regression models, the neuron yielded a high R^2 value for direction

tuning. It yielded a low R^2 value for postural tuning, despite the larger number of regressors in the postural analysis. Likewise, a simulated neuron with no tuning signal yielded a low R^2 value on all regression models regardless of the number of regressors in the models.

For all regression models described below, we followed a standard procedure in which we offset the movement data from the neuronal spike data by a latency. In this procedure, the kinematic data was taken from the movement period defined by the velocity analysis and the neuronal spike data was taken from an equivalent time period offset by a latency. We tested latencies between –286 and 286 ms at 14.3 ms increments (the temporal resolution of the movement data) and used the result of the regression analysis to choose the latency that optimized the R^2 between neuronal activity and movement. The reason for this approach is that in general a latency will exist between the activity of a cortical neuron and the movement of the arm, and different neurons may have different latencies.

In most analyses, the neuronal firing rate was averaged over the duration of each movement, and therefore no additional smoothing of the neuronal data was required. However, for the analyses of hand speed, the instantaneous neuronal firing rate was required. For these analyses we smoothed the firing rate data using a 10 Hz upper cutoff. The reason for using this frequency cutoff was to ensure that our results were directly comparable to previous studies that used a similar technique (Moran and Schwartz 1999, Reina, Moran, and Schwartz 2001).

The four main regression models used in this paper are described below and other models are described in the relevant sub-sections of the Results.

Direction tuning

Each neuron was tested for direction tuning in the following manner. For each movement, we calculated a mean firing rate of the neuron (spikes per sec during the movement, offset by a latency as described above). Each hand movement was assigned a direction in Cartesian space based on the vector connecting the beginning and end point of the movement. Firing rate was modeled as a function of the angular deviation ($\Delta\theta$) between this movement vector and a preferred direction.

$$Firing\ Rate = \mathrm{A}\cos(\Delta\theta) + \mathrm{B}$$

A regression analysis was used to find the optimal preferred direction and coefficients, following the method of Georgopoulos et al. (1986). The regression analysis provided an R^2 value indicating how much of the variance in neuronal activity could be attributed to the direction-tuning model.

Direction tuning that rotates with the start position of the hand

Each neuron was also tested with a second direction-tuning model. In this model, the preferred direction vector (PD) was obtained using a cosine tuning model of neuronal firing rate,

as above. However, PD was not assumed to be fixed in all regions of the workspace. Instead it was assumed to vary depending on the start position of the arm. The start vector (ST) was defined as the vector pointing from the shoulder to the position of the hand at the start of the movement. The direction of ST was described by two parameters, the start azimuth θ and the start elevation ϕ. PD was assumed to vary linearly with ST. In this model the azimuth of PD was linearly related to the azimuth of ST with a slope defined by the parameter AZ, and the elevation of PD was linearly related to the elevation of ST with a slope defined by the parameter EL. Thus if both AZ and EL are zero, then PD is indeed independent of ST and always points in the same direction regardless of starting position. If AZ and EL are both 1, then PD rotates exactly in tandem with ST, remaining at a fixed spatial relationship with respect to the arm. In this model, the equation for PD as a function of AZ, θ, EL, and ϕ, is:

$$PD = \begin{pmatrix} \cos(AZ\theta)\cos(EL\phi) & \sin(AZ\theta) & \cos(AZ\theta)\sin(EL\phi) \\ \sin(AZ\theta)\cos(EL\phi) & -\cos(AZ\theta) & \sin(AZ\theta)\sin(EL\phi) \\ \sin(EL\phi) & 0 & \cos(EL\phi) \end{pmatrix} \begin{pmatrix} PD_x^o \\ PD_y^o \\ PD_z^o \end{pmatrix}$$

where PD^0 indicates the preferred direction when the hand is located at θ=0 and ϕ=0.

End-point tuning

For this model of neuronal tuning all data concerning the direction or trajectory of the movement was discarded and only the end point of the movement was considered. Firing rate was modeled as a Gaussian function of these end points in Cartesian space. In the following equation, x1, x2, and x3 refer to the three Cartesian coordinates of the end point of the movement; P1, P2, and P3 refer to the coordinates of the peak of the Gaussian; the standard deviations of the Gaussian around that peak are indicated by σ1, σ2, and σ3; the height of the Gaussian is given by A; and the floor of the Gaussian is B. A non-linear regression technique (Bates and Watts 2007) was used to fit this equation to the data for each neuron.

$$Firing\ Rate = \mathrm{A}e^{\frac{(x_1-P_1)^2}{2{\sigma_1}^2}+\frac{(x_2-P_2)^2}{2{\sigma_2}^2}+\frac{(x_3-P_3)^2}{2{\sigma_3}^2}} + \mathrm{B}$$

End-posture tuning

This model followed the same general equation as the end-point model except that it involved the 8 dimensions of arm posture space (x1 through x8) rather than the 3 dimensions of Cartesian space. Firing rate was modeled as a Gaussian function that had a peak at a specific, preferred posture. Again a non-linear regression technique was used to fit the model to the data for each neuron.

$$Firing\ Rate = \mathrm{A}e^{\Sigma_{i=1:8}\frac{(x_i-P_i)^2}{2{\sigma_i}^2}} + \mathrm{B}$$

RESULTS

Speed tuning

Speed tuning I

We first tested whether previous standard methods of obtaining speed tuning from motor cortex neurons would also reveal speed tuning in the present, unconstrained movement set. In the first test of speed tuning we used a procedure based partly on that of Moran and Schwartz (1999). In this method, for each neuron, many movements are averaged together to produce a mean speed profile and a mean firing-rate profile. The two profiles are then compared to each other using a regression analysis to determine if the average hand speed tends to track the average firing rate throughout the movement.

Figure 3-3 shows the results for one example neuron. The data on the changing position of the hand through time was segmented into separate movements (see methods for details). Because of the unconstrained nature of the animal's behavior, different movements were of different durations. To average across movements we first normalized the length of movements by dividing each movement into 20 time bins. For each time bin, the instantaneous firing rate and the hand speed were calculated. These numbers were then averaged across the 321 movements tested for this neuron. The thin line in Figure 3-3A shows the average hand speed rising and falling during the movement in a smooth, bell-shaped velocity profile typical of

normal movement. The thick line shows the average neuronal activity, also rising and falling during movement.

We performed a regression analysis to compare the average hand speed to the average firing rate. Following a standard method, we used a fixed time window of kinematic data that corresponded to the hand movement, and used a time window of equal duration but adjustable start time for the neuronal data. The temporal offset between the neuronal and kinematic data was then optimized to yield the maximum R^2 value. A similar offset optimization was used for all analyses (see Methods). For this neuron, the regression analysis returned the largest R^2 value at a temporal offset of -71 ms with neuronal activity preceding movement. The R^2 value for this example neuron was 0.87 and the regression was highly significant ($P = 3.5 \times 10^{-9}$).

Figure 3-3B shows a histogram of R^2 values for all 64 cells tested in this manner. The mean R^2 value was 0.71, and 93% of the neurons had a statistically significant regression against speed (0.05 significance level Bonferroni corrected for the number of neurons). These findings roughly match the findings of Moran and Schwartz (1999) who also reported a high degree of speed tuning. The results indicate that the same type of neuronal tuning to speed found in previous experiments can be obtained in a freely moving condition.

The above method of testing for speed tuning has a limitation. The method involves averaging across many movements to obtain relatively smooth, well-behaved curves from which the between-movement variance has been removed. If the averaging is not performed, and

therefore the full range of neuronal variance is left in the analysis, what percentage of the total variance will be attributable to speed tuning? This question is addressed in the next section.

Speed tuning II

To examine speed tuning in the raw data without using averaging to remove variance, we used a technique based partly on Ashe and Georgopoulos (1994). Figure 3-3C shows the result for one example neuron. For this neuron, 339 movements were studied. These movements, placed end-to-end with periods of non-movement removed, spanned 129 seconds. During this span of time, did the instantaneous firing rate of the neuron correlate with the instantaneous hand speed? We treated each 14.3 ms time bin within each movement as a distinct data point in the regression. For this neuron, 9018 time bins were used. The regression between neuronal activity and hand speed (offset by an optimized latency; see Methods) was highly significant ($P = 1.5 \times 10^{-17}$). Yet the R^2 value was 0.03. Of the variance in neuronal activity during hand movement, only three percent could be attributed to speed tuning.

Figure 3-3D shows the R^2 values for all 64 cells tested. The mean R^2 was 0.01 and 98% of the neurons had a statistically significant regression (0.05 significance level Bonferroni corrected for the number of neurons). These results show that almost all neurons are tuned to hand speed and that this tuning accounts for a small proportion of the total neuronal variance.

Many different analytic methods can be used to test for speed tuning. We tested two additional regression analyses with similar results. We compared the peak speed of each

movement (hand speed during the 14.3 ms time bin around the peak speed) to the firing rate associated with that peak speed (during the 14.3 ms time bin offset from the peak speed by a latency optimized for each neuron). This regression returned a mean R^2 value of 0.03. We also compared the mean speed during each movement to the mean firing rate associated with the movement (during a time window of the same duration as the movement and offset by a latency optimized for each neuron). This regression returned a mean R^2 value of 0.02. All of these methods converged on a similar result: The neuronal firing rate varies from movement to movement, and 1-3% of this variance is attributable to speed tuning. Only when the between-movement variance is removed by averaging (as in Speed Tuning I) will the remaining variance be well correlated with hand speed. These results do not show that neurons have no real speed tuning. Rather, they are genuinely speed tuned, accounting for a small proportion of the total neuronal variance. Previous studies report that motor cortex neurons are tuned both to speed and direction. Can a large proportion of the remaining variance be attributed to direction tuning? The following sections examine tuning to direction and also to velocity.

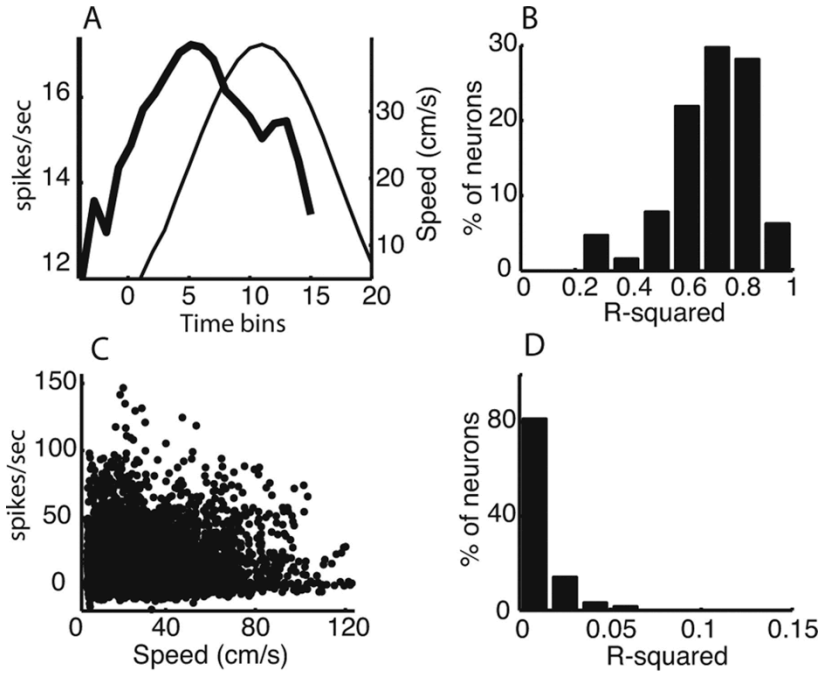


Figure 3-3 : Speed tuning of motor cortex neurons. A. Firing rate and hand speed averaged over 321 movements for one example neuron. When neuronal data was shifted forward by 71 ms (the optimal time lag for this neuron), the two curves matched closely with an R^2 of 0.87. B. Frequency histogram of R^2 values for all 64 neurons tested as in A. C. Firing rate vs hand speed for 9018 time bins during hand movement studied for one example neuron. The two variables showed a small but highly significant correlation ($R^2 = 0.03$, $P = 1.5 \times 10^{-17}$). D. Frequency histogram of R^2 values for all 64 neurons tested as in C.

Direction tuning

Direction tuning I: Local

Figure 3-4A shows 26 movements selected from the 320 spontaneous movements performed by the monkey during the testing of one example neuron. These 26 movements were selected from the full set of movements on the basis of start location (within a central ball of space with radius 5 cm) and length (between 6 and 15 cm). These movements therefore roughly

approximated the center-out movement set commonly used to test direction tuning in motor cortex neurons (e.g. Georgopoulos, Schwartz, and Kettner 1986). For each of these movements we defined the direction to be the vector pointing from the start location to the end location of the movement. Also for each movement a mean neuronal firing rate was calculated. The mean firing rate was based on a time interval equal to the movement interval but offset from the movement interval by an optimized temporal lag (see Methods). We performed a regression analysis to determine the preferred direction of the neuron. We used a standard cosine tuning model (Georgopoulos, Schwartz, and Kettner 1986) in which the firing rate of the neuron is proportional to the cosine of $\Delta\theta$, the angle between the preferred direction and the actual direction of movement.

Figure 3-4B shows the result of this analysis for an example neuron. On average, the firing rate was high during movements near the preferred direction (thus with low $\Delta\theta$) and low during movements far from the preferred direction (with high $\Delta\theta$). The regression fitting to the cosine model returned an R^2 value of 0.43 and was highly significant ($P = 0.0001$).

Figure 3-4C shows the R^2 values for all cells tested. The mean R^2 value was 0.42, and 68% of the cells showed a significant fit to the cosine model of directional tuning (0.05 significance level Bonferroni corrected for the number of neurons). The present results are therefore similar to previous results using a center-out task, in that most neurons showed a significant fit to a cosine tuning function.

Just as in the case for speed tuning, the R^2 values obtained for the present analysis showing strong direction tuning must be interpreted cautiously. The movement set is limited such that the direction of movement represents a main source of variance and thus results in a relatively high R^2. Other sources of variance, including differing starting positions, starting postures, and movement distances, though present in this selected data set, are minimized. If the entire range of movements were used, and thus all sources of variance were admitted to the data, how much of the total variance would be attributable to direction tuning? This question is addressed in the next section.

Direction tuning II: Global

Figure 3-4D shows the full set of 320 movements performed by the monkey during the testing of one example neuron. We performed a regression analysis to determine how well the neuronal firing during the full range of movements fit a cosine model of direction tuning. Figure 3-4E shows the result for the same example neuron shown in Figure 3-4B. The regression fit to the cosine model was highly significant ($P = 0.00008$). The R^2 value for the global movement set, however, was 0.05. Thus direction tuning accounted for only about 5% of the total variance in this neuron's behavior. Figure 3-4F shows the R^2 values for all 64 cells tested. The mean R^2 value was 0.08, and 63% of the cells showed a significant fit to the cosine model of directional tuning (0.05 significance level Bonferroni corrected for the number of neurons).

These results show that most neurons in motor cortex are indeed direction tuned, but by itself direction tuning accounts for only about 8% of the global neuronal variance. Other parameters must account for the remaining variance. Numerous studies have shown that hand position and arm posture have a large effect on the activity of motor cortex neurons (Caminiti, Johnson, and Urbano 1990, Kettner, Schwartz, and Georgopoulos 1988, Scott and Kalaska 1997, Scott and Kalaska 1995, Sergio and Kalaska 1997). These effects of posture may be even greater in the present data set because a greater range of postures was included.

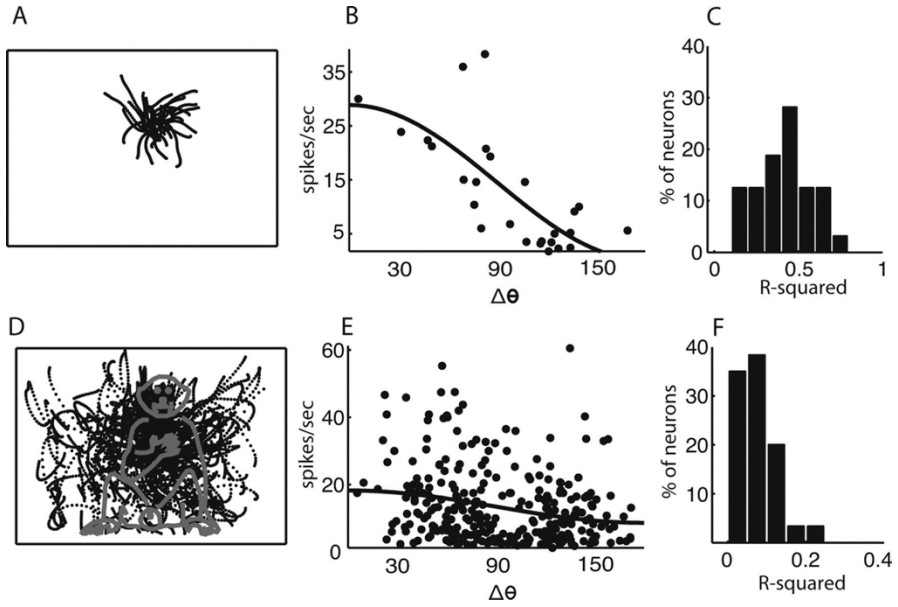


Figure 3-4 : Direction tuning of motor cortex neurons. A. Front view of 26 selected hand movements made during 10 min of testing one neuron. Each trail of dots = 1 movement measured at 14.3 ms intervals. Frame is 45 cm tall. Each movement shown originated within a 5-cm radius sphere of central space and was between 6 and 15 cm in length. B. Tuning of an example neuron to direction, based on selected movement set. X axis shows angular difference between the direction of each movement and the preferred direction; Y axis shows mean firing rate during each movement; for cosine tuning to direction, R^2 = 0.43, P = 0.0001. C. Frequency histogram of R^2 values for all neurons tested as in B. D. Front view of full set of 320 hand movements made during testing of one neuron. E. Direction tuning of an example neuron (same neuron as in B), based on full movement set. R^2 = 0.05, P = 0.00008. Note that a new preferred direction was obtained by regression, and therefore the data points shown in B do not plot to the same location on the x axis as in E. F. Frequency histogram of R^2 values for all neurons tested as in E.

Direction tuning III: rotation of preferred direction with changes in starting hand position

The previous section examined whether each neuron was tuned to a single preferred direction. One possibility, however, is that a single preferred direction might not capture the full extent of direction tuning for a typical neuron. Perhaps direction tuning would capture a greater

share of the total neuronal variance if we used a more sophisticated tuning model in which the preferred direction rotates systematically as the starting position of the hand changes.

Caminiti et al. (1990) tested direction tuning in motor cortex neurons when the hand was placed at three different starting positions separated horizontally. Averaged across neurons, the preferred direction rotated in the same direction that the starting position of the arm rotated. (See also Wu and Hastopoulos 2006 for a similar finding of shoulder-centered directional tuning.) We asked whether a similar result could be confirmed in the present data set, and whether this model of a rotating preferred direction could capture more of the total neuronal variance.

We analyzed each neuron as follows (see Methods for equations). First, for each movement, we defined the start vector to be the straight-line vector from the monkey's shoulder to the position of the hand at the start of the movement. The direction of this start vector was defined by a specific start azimuth angle and start elevation angle. We then fit the neuronal data to a regression model that included a cosine tuning to a preferred direction, but in which the preferred direction rotated in the azimuth in a manner linearly related to the start azimuth angle of the arm and also rotated in elevation in a manner linearly related to the start elevation angle of the arm. The ratio between the rotation of the preferred vector of the neuron and the start vector of the arm was parameterized by two numbers, AZ for the azimuth ratio and EL for the elevation ratio. For example, if the regression settled on an AZ = 0, this result would indicate that the preferred direction for the neuron remained constant in different regions of space and did not rotate in the azimuth regardless of the start vector. If the regression returned an AZ > 0, this

result would indicate that the preferred direction of the neuron tended to rotate in the same direction as the starting vector of the arm. An AZ of 1 would indicate that the preferred vector rotated exactly in tandem with the start vector of the arm. An AZ < 0 would indicate that the preferred vector rotated in the opposite direction as the start vector of the arm. Given the noisy nature of neuronal signals, a range of AZ values was expected, and a null result would emerge as a distribution of AZ values centered around 0. Similar relationships would pertain to the elevation of the preferred vector and the elevation of the start vector, characterized by the parameter EL.

Figure 3-5A shows the results of this regression for all neurons. The black bars show the results for the azimuth. The AZ values peaked near 1. The mean AZ value was 0.88, and the distribution was significantly greater than 0 ($t = 3.87$; $P = 0.00039$). This result replicates the finding of Caminiti et al. (1990) in three ways. First, the preferred direction rotated when the start position of the hand was changed. Second, the extent and direction of rotation of the preferred direction varied considerably from neuron to neuron, as shown by the wide distribution of AZ values. Third, on average across neurons, the preferred direction rotated in the same direction and by roughly the same extent (approximately a ratio of 1) as the start vector. On this final point our results are even more clear than the results of Caminiti et al., possibly because Caminiti et al. used only three start positions of the hand whereas we used a continuous range of start positions over the entire workspace, improving our sampling.

The experiment of Caminiti et al. examined changes in start position in the azimuth only. Here we were able to examine changes in both azimuth and elevation, and the open bars in Figure 3-5A show the result for elevation. The mean EL value was 0.54, and the distribution was significantly greater than 0 ($t = 2.28$; $P = 0.029$). This result indicates that the preferred direction tended to rotate in elevation in the same direction that the start vector rotated. However, the distribution was not as clearly shifted toward 1 for elevation as it was for azimuth.

These results show that the preferred direction that best fits a neuron, as extracted by regression analysis, tends to be one that changes depending on the start position of the hand, rotating roughly in tandem with the arm. This trend is more clear and highly significant for rotations in the azimuth, and is less clear though still significant for elevation. Figure 3-5B shows the distribution of R^2 values obtained with this model of a rotating preferred direction. The mean R^2 value is 0.13. This mean R^2 value is significantly greater than the mean R^2 of 0.08 obtained with the previous, simpler model of a fixed preferred direction and shown in Figure 3-4F (comparison of the two R^2 distributions, $t = 3.5$; $P = 0.00072$).

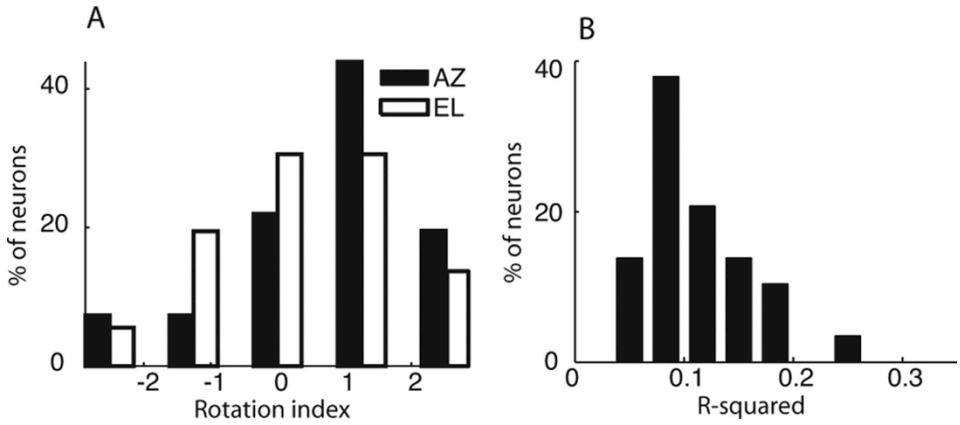


Figure 3-5 : Rotation of preferred direction with starting hand position. A. A preferred direction model was tested in which the preferred direction of a cell was not fixed but instead could rotate as the start position of the hand rotated. Frequency histogram shows AZ and EL values for all cells tested. The AZ rotation index indicates the ratio between the starting azimuth angle of the arm and the azimuth angle of the neuron's preferred direction. The peak in AZ near 1 indicates that the preferred direction of most neurons tended to rotate in the same direction and by a similar amount as the starting angle of the arm. Similarly the EL rotation index indicates the ratio between the starting elevation angle of the arm and the elevation angle of the neuron's preferred direction. B. Frequency histogram of R^2 values for all cells tested with this model of a rotating preferred direction.

Direction tuning IV: changes of preferred direction with changes in initial posture

Scott and Kalaska (1997, 1995) showed that the preferred direction of a neuron tends to change when the initial posture of the arm is changed. In their experiment the start posture of the arm was varied by altering the height of the elbow, essentially increasing the elevation angle of the shoulder and the internal rotation angle of the shoulder in tandem, to keep the hand in roughly the same spatial location. Here we examine whether a similar result applies to the present naturalistic movement set.

For each neuron we performed the following analysis. We analyzed the posture of the arm at the start of each movement. We focused on two aspects of arm posture: the elevation angle of the shoulder and the internal rotation angle of the shoulder. We constructed a single parameter, the "raised elbow" parameter, in which both of these shoulder angles varied together. A high value of "raised elbow" indicates that the shoulder elevation was high and the shoulder was internally rotated, essentially bringing the arm to a "chicken wing" posture with the elbow in upper space. Likewise a low value of "raised elbow" indicates that the shoulder elevation was low and the shoulder was externally rotated, essentially bringing the elbow to a low position. All movements were divided into two equal groups along a median split: those with a high "raised elbow" value at the start of the movement, and those with a low "raised elbow" value at the start of the movement. For each group of movements, we separately calculated a preferred direction using a regression analysis as above (*Direction tuning I*). We then compared the two preferred directions by calculating the angular difference between them.

Figure 3-6A shows the result for one example neuron. The preferred direction was different for the trials with a high value of "raised elbow" than for the trials with a low value of "raised elbow," with a change in preferred direction of 62 degrees (thin vertical line in Figure 3-6A). Is this change in preferred direction the result of the change in arm posture, or is it merely the result of noise in measuring an unreliable preferred direction? To address this question we took the same total movement set and divided it in half randomly, calculating a preferred direction for each half and then finding the difference between the two preferred

directions. We performed this random division of the data 200 times. The results are shown in the histogram in Figure 3-6A. When the movements were randomly assigned to two groups, the difference in preferred direction tended to be small, with a mean difference of 24 degrees. When the movements were assigned to two groups on the basis of arm posture, the difference obtained was far from the randomized mean, with a Z score of 2.81. For this particular cell, the start posture of the arm had a significant effect on the preferred direction, beyond that expected by chance (P = 0.0051). This example neuron was not typical; it showed an unusually clear effect.

Figure 3-6B shows the distribution of Z scores obtained for all neurons. If the posture of the arm had no significant effect on direction tuning, then the Z scores should be equally distributed about 0. The distribution is broad and extensively overlaps 0, but is significantly greater than 0 (mean = 0.51; t = 2.89; P = 0.0057). Thus the effect of starting posture on the preferred direction is variable from neuron to neuron, but across the population of neurons is statistically significant. These results replicate the finding of Scott and Kalaska (1997, 1995), indicating that the same effect obtained in a constrained movement set can also be obtained in a free movement paradigm. The results of this section and the previous section (*Direction tuning III: rotation of preferred direction with changes in starting hand position*) together confirm, in the free-moving paradigm, one of the basic properties of motor cortex neurons that had been described previously: direction tuning is not consistent across all parts of space or all configurations of the arm. Instead, it is most consistent locally, and changes with global changes

of hand position or arm posture (Caminiti, Johnson, and Urbano 1990, Kettner, Schwartz, and Georgopoulos 1988, Scott and Kalaska 1997, Scott and Kalaska 1995, Sergio and Kalaska 1997).

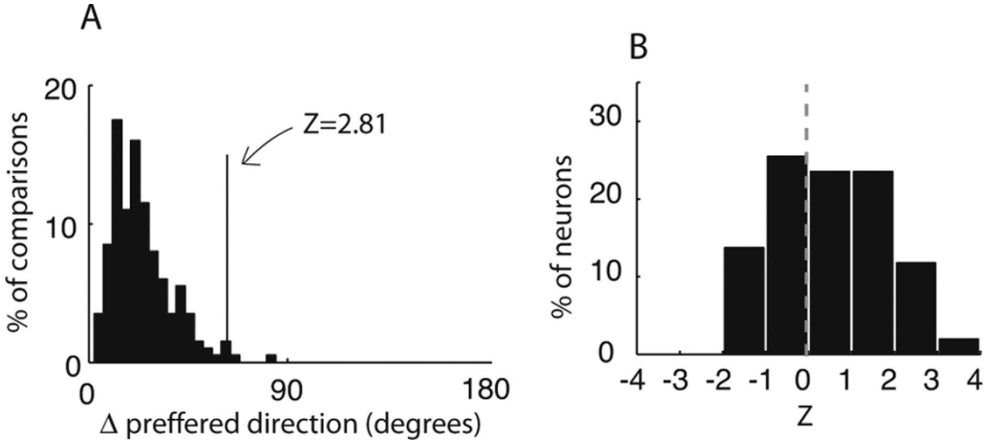


Figure 3-6 : Change in preferred direction with starting arm posture. A. Results for one neuron. All movements were divided on a median split according to a "raised elbow" parameter (see *Direction tuning IV: changes of preferred direction with changes in initial posture*). The preferred direction was calculated separately for the two sets of movements and the difference in preferred direction ($\Delta\theta$) is shown as the vertical black line. To assess the reliability of this result the movements were also randomly divided into two groups and a $\Delta\theta$ was calculated. This random division was performed 200 times and the results are shown as a frequency histogram. On this randomized distribution, the $\Delta\theta$ of the non-random, posture-based split was significantly above the mean ($z = 2.81$, $P = 0.0051$). B. Z scores for all cells tested as in A. This distribution of Z scores was significantly greater than 0 (mean = 0.51; $t = 2.89$; $P = 0.0057$).

Direction tuning V: Velocity tuning

The above sections explore speed tuning and direction tuning separately. It has also been shown that motor cortex neurons are tuned to velocity, or to the combination of speed and direction (Moran and Schwartz 1999, Paninski et al. 2004, Reina, Moran, and Schwartz 2001).

Since speed accounted for almost none of the overall variance, we expected a velocity model to account for about the same amount of variance as direction tuning alone. We tested two velocity models.

First, we performed the same regression as in *Direction Tuning II: Global* but the direction vectors in the regression equation were replaced by velocity vectors (the velocity vector for each movement was defined as the mean hand velocity during the movement, whose components were the mean x, y, and z hand velocity). The regression model returned a preferred velocity vector for each neuron with a cosine tuning to that preferred velocity. This model of velocity tuning resulted in a mean R^2 value of 0.09, not significantly different from the value obtained for regression against direction alone (0.08).

Second, we performed the same regression as in *Direction tuning III: rotation of preferred direction with changes in starting hand position* but again with the direction vectors in the regression equation replaced by hand velocity vectors. This regression model returned a preferred velocity vector that could rotate depending on the start position of the hand. This model of velocity tuning resulted in a mean R^2 value of 0.11, also not significantly different from the value obtained for regression against direction alone (0.13).

Direction tuning VI: modulation by a hand position signal

The activity of motor cortex neurons is modulated not only by the direction of a reach but also by the position of the hand (Caminiti, Johnson, and Urbano 1990, Kettner, Schwartz, and

Georgopoulos 1988, Scott and Kalaska 1995, Sergio and Kalaska 1997, Sergio and Kalaska 2003). This hand position signal is often included in the direction-tuning regression model as a linear modulation of the baseline firing rate of the neuron or as a linear modulation of the gain of the direction tuning response.

We examined a model of this type that took into account both direction and hand position. We began with the regression equation from *Direction tuning III*, in which the preferred direction was allowed to rotate depending on the starting position of the hand. This model, as described above, resulted in a mean R^2 value of 0.13. It was the best of the direction tuning models in the sense that it captured the greatest percentage of the variance. We then modified the regression in two ways. The A term in the model (see Methods) was replaced by a term that depended linearly on the x, y, and z starting position of the hand. In this manner the gain of the direction tuning response could be modulated by hand position. Likewise, the B term in the model was replaced by a term that depended linearly on the x, y, and z starting position of the hand, thereby allowing the baseline of the direction tuning response to be also modulated by hand position. This model therefore examined the amount of neuronal variance that could be attributed to a combination of direction tuning and a linear effect of starting hand position. Because the model took into account more than direction tuning, it was expected to capture more of the neuronal variance than did direction tuning alone.

The mean R^2 value for this combined regression was 0.16. One interpretation is that direction tuning has improved by taking linear hand position trends into account. This

interpretation however is not correct. The model is not a direction-tuning model; it is a mixed model. The 16% of the variance is not attributable to direction tuning alone. Instead, the results indicate that direction tuning captures about 13% of the variance; and that adding a linear dependence of firing rate on starting hand position can add about an additional 3% of the variance.

End-point tuning

End-point tuning I: Gaussian tuning

In a standard test of direction tuning, the hand begins in a central location and reaches to a variety of peripheral targets. The movements therefore differ both in direction and in final hand position. Many studies have examined whether the tuning obtained in these tasks is truly a tuning to direction or instead a tuning to the final position of the hand (Caminiti, Johnson, and Urbano 1990, Fu, Suarez, and Ebner 1993, Kettner, Schwartz, and Georgopoulos 1988). The studies test a variety of starting and ending positions, such that both hand position and direction can be independently varied. These studies find that direction tuning is genuine and persists independently of hand position, though there is also some effect of final hand position. These experiments however were inadvertently designed to favor the effect of direction tuning and to diminish the effect of end-point tuning, in the following manner. Most of these studies sampled the full range of possible directions, thereby maximizing the chance of

detecting direction tuning, even if that tuning is broad. The studies, however, typically explored only a small range of positions near central space. For example, a central cube of space that is 10 cm in width comprises only about 5% of the volume of the workspace used by the monkey in the present experiment. In one of the most extensive previous tests of the workspace (Caminiti, Johnson, and Urbano 1990), three cubes of space, each 10 cm across, were tested, totaling about 15% of the total workspace. If a neuron is broadly tuned to hand end point, such that the tuning curve can be seen only over the full workspace and varies little over the limited range tested, then end-point tuning would be minimized in these previous experiments. Here we examine whether neurons are significantly end-point tuned over the global movement set. It is important to note that we do not treat direction tuning and end-point tuning as rival hypotheses. Neurons may in principle be significantly tuned to direction in a manner that emerges over local regions of space, and in an independent fashion significantly tuned to end point in a manner that emerges over a more global expanse of the workspace.

In order to test for end-point tuning, we modeled the firing rate of the neuron as a Gaussian function of end point in which the Gaussian was peaked at a preferred spatial location of the hand (see Methods). Movements that ended near that preferred end point should be associated with high firing rates and movements that ended far from that preferred end point should be associated with low firing rates. For each neuron, we obtained an R^2 value indicating how well this end-point model fit the neuronal data. Figure 3-7A shows the distribution of R^2

values for the population of neurons. On average, the end-point model accounted for 22% of the total neuronal variance.

Figure 3-7B-D show the distribution of preferred hand positions for the tested neurons, as obtained with the end-point tuning model. These preferred hand positions were dispersed through the space around the animal. For some neurons, the regression analysis found a preferred hand position at the edge of the normal range of the workspace. In such a case, the neuron essentially had a monotonic tuning to hand position, preferring one extreme side of the workspace over the other. Most neurons, however, had a preferred hand position inside the workspace and therefore had a non-linear response function that fell off to either side of the preferred location. The tuning functions were generally broad: the mean width of the Gaussian tuning curve at half-height was 18 cm, about half the range of the workspace. Because of the broadness of these tuning curves, the end-point tuning is unlikely to be a major source of signal in experiments that test a small part of the workspace in front of the monkey. If only a small volume of the workspace is tested, most neurons will have a preferred hand position outside that tested region, and the tuning to end point will either be missed (if the tested part of space happens to overlap a relatively flat part of the tuning curve) or will be measured as a graded, monotonic preference for one side of the tested space over the other (if the tested part of space overlaps with the rising slope of the tuning curve).

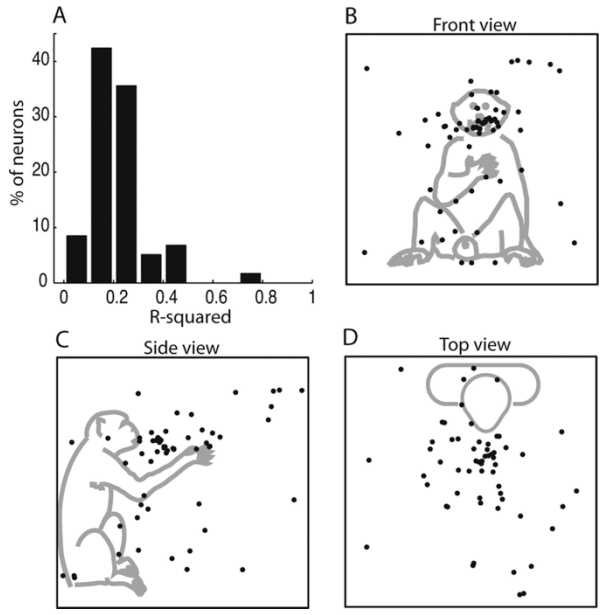


Figure 3-7 : Tuning of neurons to hand end point. A. Each neuron was tested with a tuning model in which the neuron fired most during movements that terminated with the hand at or near a specific location in space, and fired progressively less during movements for which the hand terminated progressively farther from the preferred location. The graph shows a frequency histogram of R^2 values for all cells tested with this end-point model. B-D. Preferred hand positions as determined by the end-point tuning model, displayed from 3 perspectives. The schematic monkey drawing indicates approximate scale and orientation.

Figure 3-8A1 shows data from an example neuron. First the end-point tuning was found using the regression analysis described above. This analysis obtained a Gaussian surface in 3-D space for which the peak of the Gaussian corresponded to the neuron's preferred end point. In this graph, the x axis represents the distance in cm between the end point of each movement and the preferred end point of the neuron. The y axis represents the firing rate of the neuron during the movement. On average, movements that terminated near the preferred posture (distance of 0) had higher firing rates, and movements that terminated progressively farther from the preferred

posture had progressively lower firing rates. This trend is noisy for two reasons. First, only about 24% of this neuron's variance was attributable to the Gaussian tuning to the preferred posture. Second, the standard deviation of the Gaussian tuning was not equal in all spatial directions. This cell was more sharply tuned in the x dimension that in the y and z dimension. In terms of the Gaussian tuning, a cm in the x direction is not equivalent to a cm in the y or z direction. To better show the Gaussian fit to the data, Figure 3-8B1 shows a replotting of the same graph. In this plot, the x axis represents the distance between the end point of each movement and the preferred end point of the neuron, but this distance now is expressed in standard deviations of the Gaussian tuning function. Again, on average, movements that terminated near the preferred posture (within a standard deviation of the peak of the Gaussian) had higher firing rates, and movements that terminated progressively farther from the preferred posture had progressively lower firing rates. This graph shows in a more direct way the relationship between firing rate and the proximity of the movement end point to the preferred end point.

Figure 3-8C1 shows rasters of neuronal activity for the same example neuron. Each line in the raster display shows spike data during the analysis window for a single movement. Since different movements were of different durations in the naturalistic movement set, these lines in the raster display are of different lengths. The first raster display shows data from the 10% of movements that terminated nearest to the preferred position; the second raster display shows data from the 10% of movements that terminated farthest from the preferred position. These rasters

show that the firing rate was variable from movement to movement, but that the neuron tended to fire more during movements that terminated near the preferred hand position (mean of 49.8 spikes per second, SD = 20.6) and tended to fire less during movements that terminated far from the preferred hand position (mean of 13.9 spikes per second, SD = 12.4) with a significant different between these two firing rates ($t = 11.3$, $P = 4X10^{-20}$).

Is the tuning to final hand position merely an artifact of direction tuning? The region of final hand positions that the neuron appears to prefer may have been approached predominantly from the preferred direction. For example, a neuron that prefers upward movements might appear to prefer an upper location in space. This explanation, however, is unlikely for several reasons. First, it predicts that the preferred hand locations should cluster systematically around the edges of the workspace, but as shown in Figure 3-7, the neurons preferred a range of final positions including many within the workspace. Second, as described in a later section of the results (*Simulated neurons*), we simulated a direction-tuned neuron and tested it on the actual movements measured from the monkey. The simulated neuron performed well on a direction-tuning model and showed little tuning to final hand position. A neuronal signal that is known to be direction tuned, therefore, does not produce an artifactual hand position tuning.

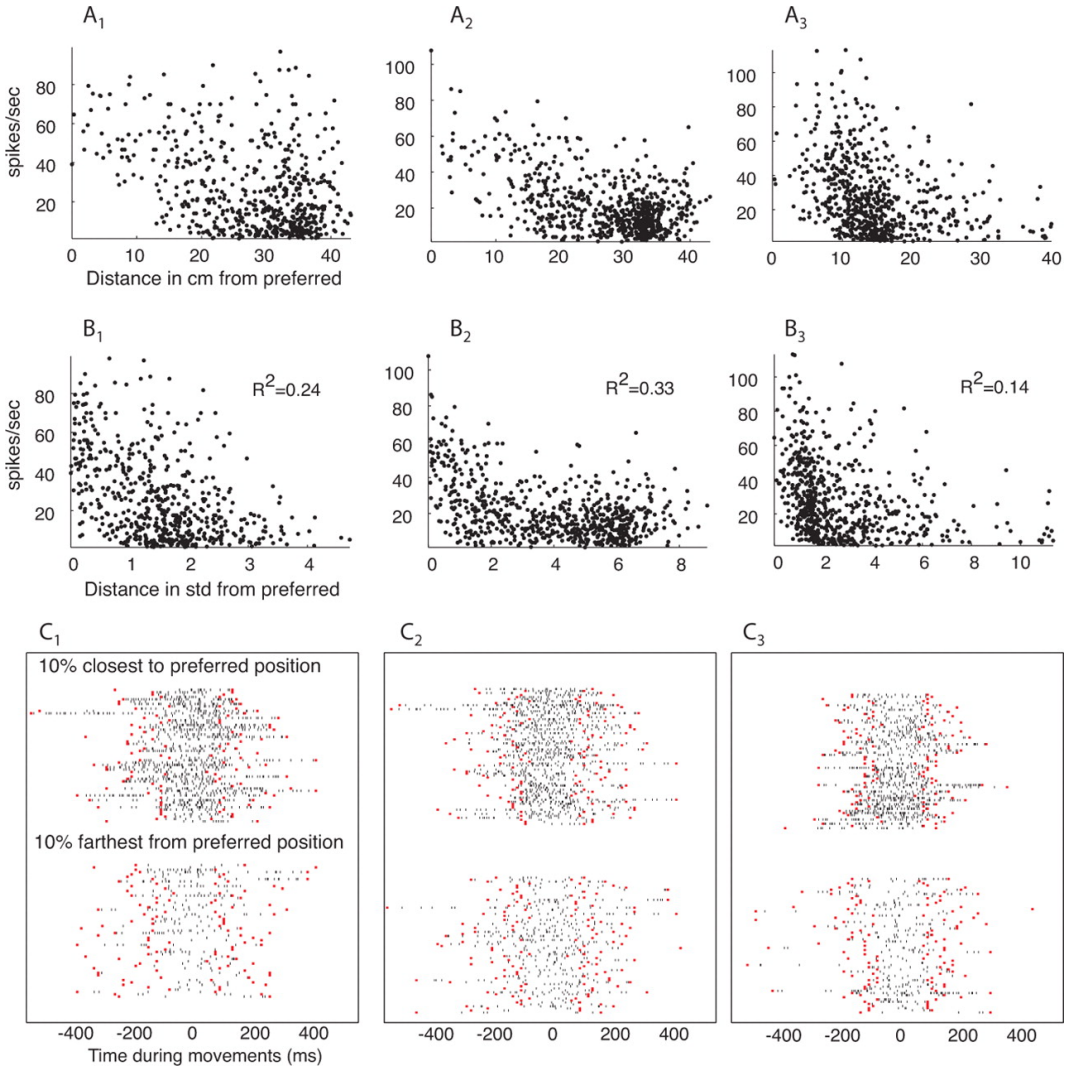


Figure 3-8 : End-point tuning of three example neurons. A1-3. Mean firing rate of example neurons 1-3 during each movement as a function of the distance (cm) between the end of the movement and the preferred end point determined by regression analysis. The neurons fired more during movements that terminated closer to the preferred end point. B1-3. Same data as in A but with the x axis re-plotted in units of the standard deviations of the Gaussian tuning curve, displaying the tuning to end point more clearly. C1-3. Rasters showing high neuronal activity of example neurons 1-3 during the 10% of movements that terminated nearest to the preferred end point, and low neuronal activity during the 10% of movements that terminated farthest from the preferred end point. Red tic marks indicate start and end of movement.

Finally, Figure 3-9 shows a third analysis to address this question. Figure 3-9A1 contains data from one example neuron. The dots show the final hand positions for the 10% of movements that terminated closest to the preferred hand position. Figure 3-9B1 shows the directions of the movements that terminated in those locations. This movement set included a range of directions, and did not approach the end points from predominantly one direction. We used this limited movement set to obtain a preferred direction (using a standard cosine tuning model as in *Direction tuning I*). The direction-tuning R^2 value was 0.18. (The relatively low R^2 value, below the mean of 0.42 obtained for local direction tuning in *Direction Tuning I*, probably reflects the fact that these movements, though more local than a full movement set, are none-the-less more dispersed than in our test of local direction tuning). The preferred direction among these movements, plotted as a thick line in Figure 3-9A1, was aimed down and to the left, whereas the preferred hand location was up and to the right. For this neuron, the preferred direction was pointed the wrong way to account for the preferred hand location. Figure 3-9 also shows a second example cell. In general the preferred direction had no obvious relationship to the preferred position. Thus although neurons are to some degree direction tuned (as outlined in the previous sections), and are to some degree tuned to the end point of the movement, the preference for a particular final position of the hand apparently cannot be explained as a consequence of a preference for a particular direction of motion. This question is further addressed in the next section.

End-point tuning II: directional convergence

Previous experiments suggest that neurons rarely show a preference for movements directed toward a specific location in space (Caminiti, Johnson, and Urbano 1990, Sergio and Kalaska 2003). The preferred direction of a neuron does not change systematically across space in such a way that it aims toward a point of convergence. Instead, when the preferred direction is assessed in different regions of space, it tends to change somewhat idiosyncratically. On average across neurons, it tends to rotate in the same direction as the rotation of the shoulder. Similar results were confirmed in the present experiment (*Direction tuning III: rotation of preferred direction with changes in starting hand position*). How can this result, a failure of the preferred direction to converge to a point, be reconciled with the result of the previous section, that neurons fire most during movements that end near a preferred location in space? On first sight the two results appear to be contradictory.

There are, however, at least two ways to define end point tuning. One is a proximity model; the other is a directional model. In the proximity model, tested in the previous section, a neuron will fire best during movements that terminate near a preferred position, regardless of the direction of the movement. For example, a movement that starts at the preferred position and ends one cm away might result in a high firing rate, even though the movement is directed away from the preferred position. Likewise, a movement that starts 20 cm from the preferred position and ends 15 cm from it might result in a low firing rate, though the movement is directed toward the preferred position. This model, in which trajectory information is discarded and only the

final position enters the equation, accounted for 22% of the neuronal variance on average, as described above.

In a directional model of end point tuning, in contrast, the neuron fires most during movements directed toward the preferred position, and fires least during movements directed away from the preferred position, regardless of whether the movement terminates near or far from the preferred position. For example, a movement that starts at the preferred position and ends one cm away might result in a low firing rate, even though the movement ends near the preferred position. Likewise, a movement that starts 20 cm from the preferred position and ends 15 cm from it might result in a high firing rate, even though the movement ends far from the preferred position. This model, in which the final position of the hand at the end of each movement is discarded and only the direction enters the equation, has generally been found to account for little of the neuronal behavior (Caminiti, Johnson, and Urbano 1990, Sergio and Kalaska 2003).

To determine whether our data set was consistent with previous findings, we also tested a directional model of end-point tuning. In this model, for each movement, two vectors were defined. The first vector connected the start point of the movement to the end point of the movement. The second vector connected the start point of the movement to a preferred end point of the neuron. The angular deviation between these vectors was defined as $\Delta\theta$. In the regression equation, the firing rate was modeled as a cosine function of $\Delta\theta$ (Firing rate = A cos $[\Delta\theta]$ + B). For each neuron, the regression obtained the preferred end point that best fit this

model. For a neuron that follows this model, the firing rate should be highest when the movement is directed toward the preferred end point and lowest when the movement is directed away from the preferred end point. This model resulted in a mean R^2 value across neurons of 0.07, less than the mean R^2 value of 0.22 obtained for the Gaussian tuning to end point described in the previous section.

It is instructive to compare the three general classes of direction model tested in the present experiment. First, a standard direction model, in which the preferred direction is the same everywhere in space (*Direction Tuning II: Global*), accounted for 8% of the variance on average. Second, a direction model in which the preferred direction could rotate in different regions of space (*Direction tuning III: rotation of preferred direction with changes in starting hand position*) accounted for 13% of the variance. This rotating model tended to select a preferred direction that rotated roughly with the shoulder angle. This rotation of the preferred direction results in a diverging pattern, in which preferred directions in adjacent parts of space point slightly away from each other. It is not surprising therefore that the third model (of the present section), in which the preferred directions in different regions of space were constrained to converge toward a common location, performed somewhat less well than either of the previous ones, accounting for 7% of the variance.

These results together with the results of the previous sections suggest that neurons are to some extent tuned to end point (mean $R^2 = 0.22$). This end-point tuning, however, should not be construed as a set of preferred directions aimed at a common location in space. The neurons are

not well tuned to a converging pattern of preferred directions. Instead the neurons are tuned to any movements that terminate near a preferred location. The spatial tuning curve is broad, the preferred region is large, the natural hand movements are often relatively short (mean length of 8 cm), and therefore a variety of movements aimed in all directions can terminate within the preferred region yielding a high firing rate. In addition, superimposed on this spatial tuning for end point, the neurons are tuned to some extent to direction (mean R^2 of approximately 0.08 to 0.13 depending on the model used), where the preferred direction has little obvious relationship to the preferred hand location.

End-point tuning III: Start-point or end-point tuning?

As described above, the firing rate of the neurons tended to be larger for movements that terminated near a preferred hand location, and this correlation accounted for about 22% of the neuronal variance. However, because the monkey's spontaneous movements were relatively short (mean of 8 cm), the movements tended to start and end in a similar region of space. Start point and end point, therefore, were correlated across the global movement set. Were the neurons tuned specifically to the *end* point of the movement, or was neuronal activity equally determined by the proximity of the *start* point to the preferred location?

To approach this issue we performed the following analysis. For each neuron we selected the 10% of movements that terminated nearest to the preferred position. Figure 3-9B1 shows this set of movements for one example neuron. These movements contain some variance

in ending and starting position. Because this local region of space around the preferred location had a diameter that approximated the average length of a movement, therefore within this movement set no significant correlation existed between start point and end point. We then asked, within this set of movements near the preferred location, whether the distance between the *start* point of a movement and the preferred point was significantly related to firing rate (using a linear regression between firing rate and distance from start point to preferred point). This was not typically the case. Across neurons, the mean R^2 value for this linear regression was 0.02, and 7% of the neurons showed a significant relationship. In contrast, we asked whether the distance between the *end* point of a movement and the preferred point was significantly correlated with firing rate. A relationship was obtained in which the firing tended to be higher for movements that terminated closer to the preferred point. The mean R^2 value for this linear regression was 0.20, and 65% of the neurons showed a significant relationship. These results suggest that the proximity of the end points to a preferred location significantly affected the firing rate of neurons, and that the proximity of the start points to the preferred location was at least a much weaker influence. This question of whether the neurons preferentially encode the end state of the movement or whether the firing rate equally reflects the initial state is further addressed below in two sections (*Neurons are more end-posture than start-posture tuned: I. Analysis of temporal offsets;* and *Neurons are more end-posture than start-posture tuned: II. Analysis of movements near the preferred posture*).

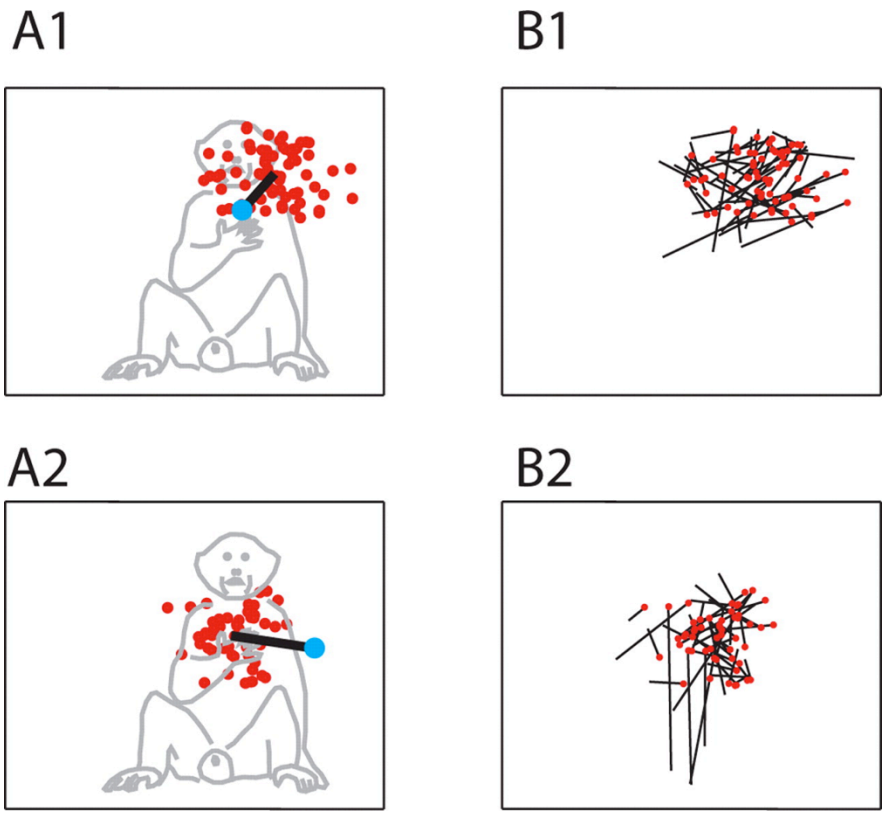


Figure 3-9 : Direction tuning does not explain end point tuning. A1-2, data from two example neurons. The dots show the final hand position for the 10% of movements that terminated closest to the calculated preferred position. The line shows the preferred direction of movement as calculated using the same local set of movements. B1-2 data from the same neurons shown in A. Each line depicts a movement (straight line connecting the start and end point of the movement) for the same movements whose end positions are shown in A.

End-posture tuning

Electrical stimulation of motor cortex on a behaviorally relevant time scale can evoke movements that terminate in specific postures (Graziano 2006, Graziano, Aflalo, and Cooke 2005, Graziano, Patel, and Taylor 2004, Graziano, Taylor, and Moore 2002). For example,

stimulation of one location in cortex might bring the hand to the mouth with the elbow in lower space, the forearm supinated such that the palm faces the mouth, the wrist slightly flexed, the fingers in a grip posture, and the mouth open. The effect of stimulation is therefore not merely to bring the hand to a specific end point, but to bring the many joints of the arm to a specific final configuration or posture. Observational studies of the normal spontaneous behavior of monkeys indicate that a high proportion of time is spent using the arm as a postural device (Graziano, Patel, and Taylor 2004, Graziano, Taylor, and Moore 2002). Psychophysical studies in humans suggest some degree of encoding of the final posture of a movement (e.g. Desmurget and Prablanc 1997, Feldman 1986, Rosenbaum et al. 1995). We therefore asked to what degree neurons in motor cortex are tuned to the posture reached by the arm at the end of each movement.

Eight degrees of freedom of the arm were monitored, including grip aperture and seven joint angles. These degrees of freedom define an 8-D posture space. We modeled neuronal firing rate as a Gaussian surface in 8-D space whose peak corresponds to the preferred end posture. Movements that terminate at a posture near the peak of the Gaussian should be associated with a high neuronal firing rate, and movements that terminate at a posture far from the peak of the Gaussian should be associated with a low neuronal firing rate. For each neuron, we obtained an R^2 value indicating how well this model fit the neuronal data. The distribution of R^2 values across the population of neurons is shown in Figure 3-10 (red bars). The mean R^2 value was 0.36. For comparison, Figure 3-10 also shows the distribution of R^2 values for speed

tuning, for the complex direction tuning in which the preferred direction could rotate with the starting position of the hand, and for end-point tuning in 3-dimensional Cartesian space. Of the models of neuronal tuning tested here, the one that performed best involved so called intrinsic coordinates (in this case joint angles) rather than extrinsic coordinates (direction or Cartesian hand location), in agreement with previous suggestions that motor cortex emphasizes intrinsic coordinates (Scott and Kalaska 1997). These tuning curves to end posture are examined in greater detail in the following six sections.

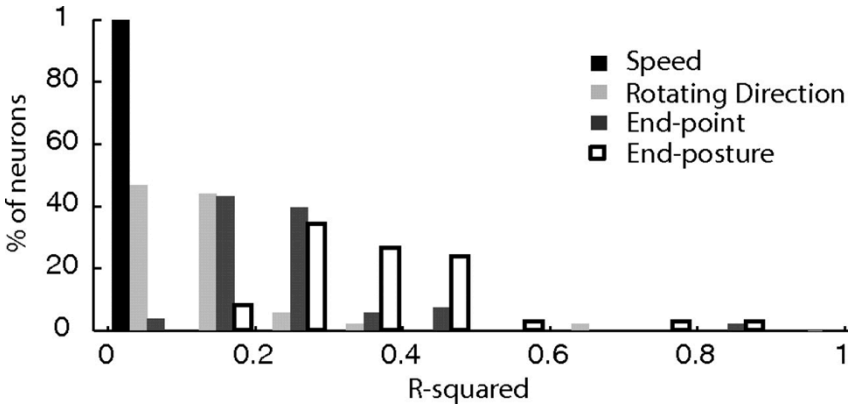


Figure 3-10 : Comparison of four regression models. Each neuron was tested with a tuning model in which the neuron fired most during movements that terminated with the 8 joints of the arm at or near a specific postural configuration. This end-posture tuning is shown as a frequency histogram of R^2 values. Also shown is a frequency histogram of speed tuning (from Figure 3-3D), direction tuning in which the preferred direction was allowed to rotate depending on starting hand position (from Figure 3-5B), and tuning to the 3-dimensional end point of the movement (from Figure 3-7A).

Examples of end-posture tuning

Figure 3-11A1 shows data from an example neuron. First the end-posture tuning was found using the regression analysis described above. This analysis obtained a Gaussian surface in 8-D posture space for which the peak of the Gaussian corresponded to the neuron's preferred posture. In this graph, the x axis represents the distance in posture space between the final posture of each movement and the preferred posture of the neuron. It is important to note that this distance is not the distance of the hand from a preferred location in space, but rather the distance of the arm from a preferred configuration in posture space. Thus the units cannot be in Cartesian cm. Instead, we have expressed the units in standard deviations of the Gaussian fitting function. In this fashion, all 8 dimensions of posture space can be expressed in the same units, and a one-dimensional graph can be presented. The advantage of the one-dimensional graph is that it shows the Gaussian fit to the data. On average, movements that terminated near the preferred posture (within a standard deviation of the peak of the Gaussian) had higher firing rates, and movements that terminated progressively farther from the preferred posture had progressively lower firing rates. About 33% of this neuron's variance was attributable to the Gaussian tuning to the preferred posture ($R^2 = 0.33$, $P < 0.0001$).

Figure 3-11B1 shows rasters of neuronal activity for the same example neuron. Each line in the raster display shows spike data during the analysis window for a single movement. Since different movements were of different durations in the naturalistic movement set, these lines in the raster display are of different lengths. The first raster display shows data from the 10% of

movements that terminated nearest to the preferred posture; the second raster display shows data from the 10% of movements that terminated farthest from the preferred posture. These rasters show that the firing rate was variable from movement to movement, but that the neuron clearly fired more during movements that terminated near the preferred posture. Although only 33% of the neuronal variance was attributable to end-posture tuning, the effect is clearly visible in the rasters.

Figure 3-11C1-E1 show more data from the same example neuron. Each stick figure shows the configuration of the arm at the termination of a movement; the figure shows the 10% of movements that terminated closest to the preferred posture; and the three panels show a front, side, and top view of the postures. This display does not capture the full 8-dimensional posture space. Several degrees of freedom, especially the internal/external rotation of the shoulder and the pronation/supination of the forearm, are not easily visible. However the graph does give some sense of the range of postures preferred for this neuron. The neuron is more narrowly tuned to some joints (such as the shoulder elevation) than to others (such as the shoulder azimuth). As a result, this range of preferred postures does not correspond neatly to a small cluster of hand positions in space. Instead, a large and irregular region of hand space is represented by these postures. This finding that postural tuning is more sharply peaked for some degrees of freedom than for others is discussed in greater detail in the section *Posture subspaces*. For this neuron, the preferred posture involved a raised elevation of the shoulder lifting the hand mainly into upper space; a range of azimuth angles that placed the hand in a band of frontal

space; a partially extended elbow that placed the hand away from the body; an extended wrist angle; and a grip aperture (not shown) that was on average 2 cm. The posture is similar to a naturalistic reach in order to grasp or manipulate a small object in upper frontal space.

Figure 3-11 also shows the end-posture tuning for three other example neurons.

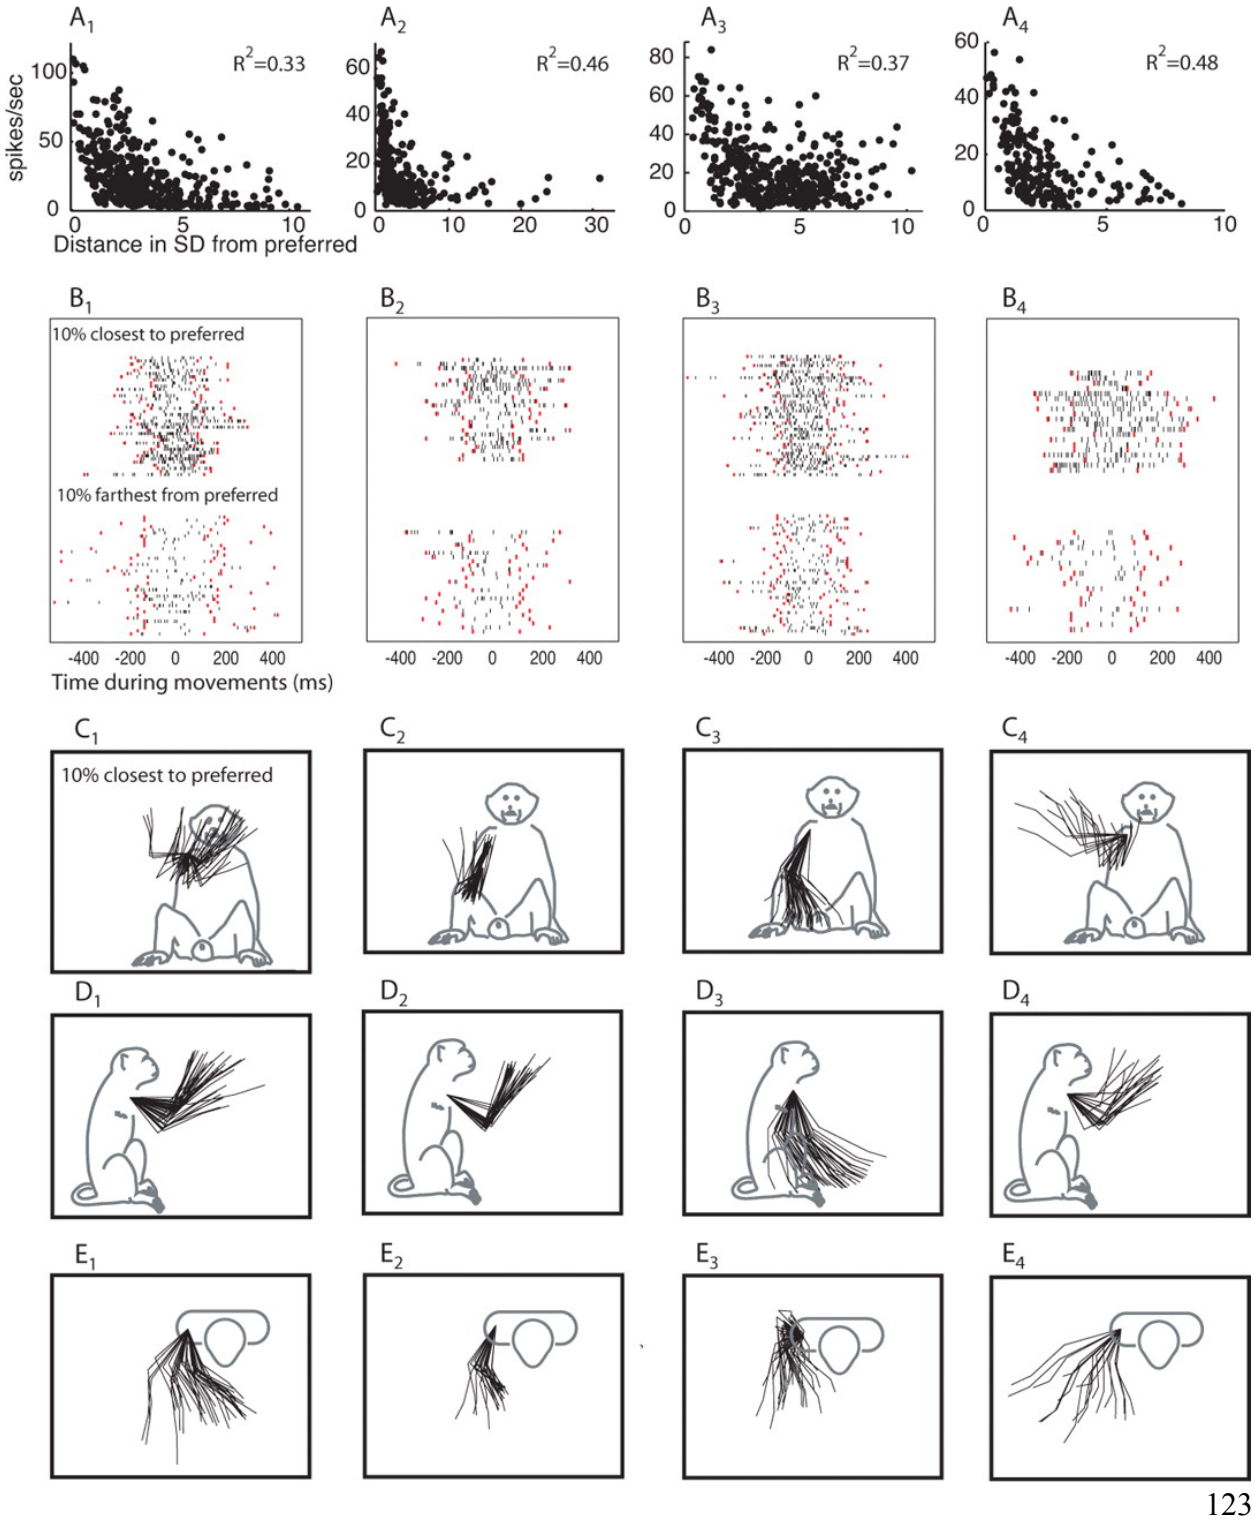

A1
A2
A3
A4
R2=0.33
R2=0.46
R2=0.37
R2=0.48
spikes/sec
Distance in SD from preferred
B1
B2
B3
B4
10% closest to preferred
10% farthest from preferred
Time during movements (ms)
C1
C2
C3
C4
D1
D2
D3
D4
E1
E2
E3
E4

Figure 3-11 : End-posture tuning of four example neurons. A1-4. Data from example neurons 1-4. For each neuron, the preferred multi-joint posture was determined by regression analysis. The final posture of each movement was compared to the preferred posture. The distance between them was calculated in 8-dimensional posture space. Distance was measured in units of standard deviations of the Gaussian tuning function, in order to express all 8 dimensions in posture space in equivalent units. This distance is plotted on the X axis and firing rate during the movement is plotted on the Y axis. B1-4. Rasters showing high neuronal activity during the 10% of movements that terminated nearest to the preferred end posture, and low neuronal activity during the 10% of movements that terminated farthest from the preferred end posture. Red tic marks indicate start and end of movement. C1-4. Stick-figure drawings showing front view of the end postures for the 10% of movements that terminated nearest to the preferred end posture. Three joints are shown: shoulder, elbow, and wrist. The schematic monkey drawing indicates approximate scale and orientation. D1-4: same as C but side view. E1-4: same as C but top view.

Range of preferred joint angles

Some experiments have examined linear models of tuning to joint rotation (e.g. Reina, Moran, and Schwartz 2001). In a linear model, a neuron will prefer rotation of a joint to one or another extreme. The present model of end-posture tuning is non-linear; it assumes that neurons fire best during rotation from any direction toward a preferred joint angle. However, the preferred angle could in principle lie at or beyond the edge of the range of movement. In this case, the non-linear model used here would approximate the linear models used elsewhere. An important question, therefore, is whether the preferred joint angles as determined by the regression model tend to fall at the extreme of the range of motion or within the range of motion. Figure 3-12 shows the result. To examine those neurons that had clearly preferred postures, we ranked neurons by the sharpness of the Gaussian tuning function and arbitrarily chose the 50% of neurons that were most sharply peaked, thus the neurons for which the end posture was most

clearly specified. These neurons are represented in Figure 3-12. Each frequency histogram shows data for one joint. For the most proximal joints, especially the three shoulder joints, neurons were tuned to a range of different preferred angles including those in the middle of the range. For the more distal joints, neurons were more likely to be tuned to an extreme. For example, neurons tended to prefer a closed grip aperture. Note that the same neuron might be tuned to an intermediate shoulder angle and a closed grip. It was not the case that some neurons were tuned to proximal joints and others to distal joints; rather, neurons tended to be tuned to a combination of many joints, an issue discussed in greater detail below in *Posture subspaces*.

The plot in Figure 3-12 provides some indication of the postures most commonly preferred by neurons. Many neurons preferred a shoulder azimuth around 90 degrees, orienting the arm into the front portion of space rather than to the side (0 degrees); a shoulder elevation of about 80 degrees and an elbow extension of about 130 degrees, which places the arm in a raised and extended posture; the wrist extended; and the grip closed. These most common joint angles correspond to a typical arm posture of the monkey, as if it had reached out to grip an object in central space. A range of other postures was also represented, as is evident from the range of joint angles in Figure 3-12. Since electrical stimulation in different regions of the precentral gyrus tended to evoke different common postures of the arm (Graziano, Taylor, and Moore 2002), it is possible that neuronal recordings from other cortical regions in the precentral gyrus might have revealed a different distribution of preferred postures.

To compare directly to a linear model of neuronal tuning in joint space, we also tested neurons using a simple 8-dimensional linear regression comparing mean firing rate during a movement with mean joint angular speed. Such a regression model assumes that a neuron prefers movement of a joint in one direction only, toward one extreme angle or the other. The mean R^2 for this linear model was 0.14, significantly less than the mean R^2 of 0.36 for the Gaussian end-posture model ($t = 7.0$; $P = 3.2 \times 10^{-9}$). This result confirms that neurons are tuned to movement through joint space, but that the linear model misses a common feature of the tuning, a tuning to joint rotations that terminate at an intermediate posture.

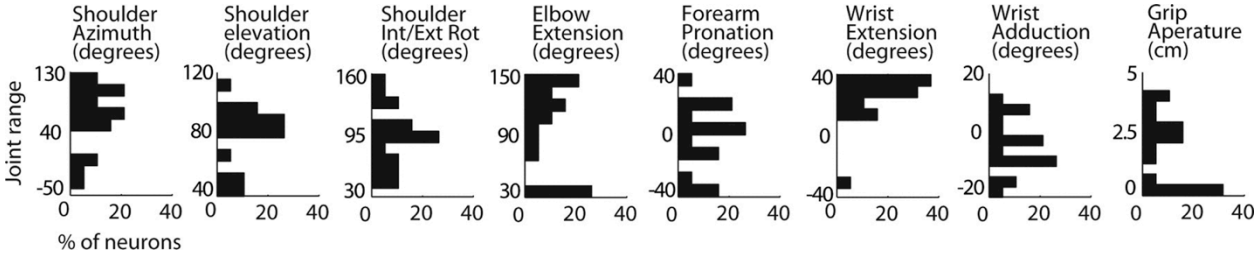


Figure 3-12 : Distribution of joint angles preferred by neurons. For each joint, a frequency histogram shows the proportion of neurons tuned to a particular final joint angle. Only the 50% most sharply tuned neurons were plotted for each graph. Most neurons were tuned to more than one joint and are therefore represented on more than one histogram.

Posture subspaces

As described above, for each neuron a regression model was used to obtain the best fitting Gaussian tuning function in 8-D posture space. On examining the resultant Gaussian tuning functions, we found that neurons were not equally well tuned to all eight dimensions of

posture space. We examined the sharpness of tuning to each of the eight degrees of freedom. Figure 3-13A shows the result for one example neuron. For each joint we calculated the width of the tuning curve at half height, in degrees of joint angle. We then divided this width by the range of movement measured during the animal's spontaneous behavior, resulting in a tuning width index that could be compared across joints. A tuning width index of 1 indicates that the width of the tuning curve was about the same as the full range of movement. A tuning width index <1 indicates a narrower tuning curve; a tuning width index >1 indicates a tuning curve so broad that it essentially was not relevant over the normal range of movement of the joint. If the index was 1 or less, we classified the neuron as well tuned to that joint. On this criterion, this example neuron was well tuned to four joints of the eight (four joints show a tuning width index <1). This neuron was tuned to shoulder azimuth, internal/external shoulder rotation, wrist extension, and grip aperture. It was therefore tuned to a mixture of proximal and distal joints. It was essentially un-tuned to the other joints. Figure 3-13B-D shows several more example neurons tuned to different combinations of joints. Figure 3-13E shows the mean across neurons. The joints are arranged in this graph in order from most sharply tuned to least sharply tuned. On average neurons were well tuned to four joints (tuning width index <1) and relatively poorly tuned to the remaining joints (tuning width index >1). These results show that it is not quite correct to view the neurons as tuned to a preferred final posture in an 8-dimensional posture space. Rather, the neurons are tuned within posture subspaces, each neuron primarily tuned to a subset of joints and less well tuned to the remaining joints.

One caveat of the above analysis is that the joints may be correlated with each other, hence a neuron that is genuinely tuned to one joint may appear to be tuned to other, correlated joints. To address this caveat we performed an alternative, conservative analysis that removed the covariance among joints. For each neuron, we first tested the end-posture model with a regression analysis as described above, and arrived at an R^2 value indicating how much of the total variance could be explained by that model. We then performed an n-1 regression to determine how much of that explained variance could be attributed solely to each joint, once the variance explained by the remaining seven joints had been removed. In this manner the correlation among joints was explicitly taken into account and removed from the analysis. We then asked, for each joint, whether it made a statistically significant unique contribution to the R^2. This method is conservative, since it eliminates signal that overlaps among joints and considers only the component of signal that is unique to each joint. Yet using this regression method we obtained a similar result. Each neuron on average was significantly tuned to four joints of the eight.

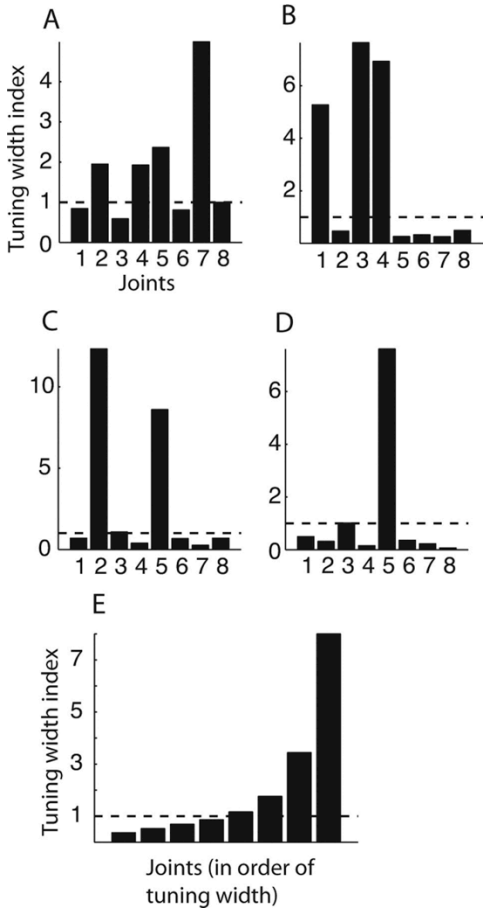


Figure 3-13 : Tuning of neurons to subsets of joints. A. Data from one example neuron. The end-posture regression analysis returned a set of preferred joint angles to which the cell was tuned in a Gaussian manner. For each joint a tuning width index was calculated. Tuning width index = width of Gaussian tuning curve at half height in degrees of joint angle / natural range of joint motion measured during movement. Tuning index of 1 indicates that the width of the tuning curve at half height approximated the normal range of joint motion. Lower tuning index corresponds to sharper tuning. The graph shows that the neuron was sharply tuned to four joints and poorly tuned to the remaining four joints. The joints are (in order) shoulder azimuth; shoulder elevation; shoulder external/internal rotation; elbow extension; forearm pronation; wrist extension; wrist adduction; grip aperture. B. An example neuron sharply tuned to 5 of the 8 joints. C. An example neuron sharply tuned to 6 of the 8 joints. D. An example neuron sharply tuned to 7 of the 8 joints. E. For each neuron the joints were arranged from most sharply tuned to least sharply tuned; the results were then averaged across neurons. On average neurons were well tuned to four of the eight joints.

Comparison of end-posture tuning and end-point tuning

The previous sections examined two types of end-state tuning: tuning to end point and tuning to end posture. Do the hand positions associated with the preferred postures tend to match the hand positions associated with the preferred end points? Such a match would provide a validity check, showing that two different analysis methods arrived at a similar result.

This question is complicated by two caveats. First, as described above, the neurons are not typically narrowly tuned to all 8 angles that make up the posture of the arm. Each neuron tends to be narrowly tuned to an average of 4 joints of the 8. A set of postures near the preferred posture therefore corresponds to a narrow range of angles for some joints and a wide range of angles for other joints. This in turn can correspond to a large, irregular set of hand positions. (See Figure 3-11 for examples.)

Second, a hand position associated with a preferred posture can also be reached by postures that are not preferred by the neuron. With 8 degrees of freedom in the arm, any hand position can be achieved through a large range of postures.

For these two reasons, there is no clean mapping between hand position and arm posture. Therefore a strict correspondence between the hand positions preferred by the neurons and the arm postures preferred by the neurons should not be expected. With these two caveats in mind, however, it may still be possible to find some degree of correspondence between the cloud of hand positions associated with the preferred posture and the cloud of hand positions associated with the preferred end point.

For each neuron we obtained the 10% of movements that terminated closest to the preferred posture of the neuron, as determined by the 8-dimensional regression against end posture (e.g. Figure 3-11C1). We obtained the mean x, y, and z location of the hand at the end of these 10% of movements. For the same neuron, we obtained the 10% of movements that terminated closest to the preferred hand position as determined by the 3-dimensional regression against end point. We obtained the mean x, y, and z location of the hand at the end of these 10% of movements. For each neuron we therefore obtained a mean hand position associated with the preferred posture (posture-defined position) and a mean hand position associated with the preferred end point (end-point-defined position). We then asked whether the posture-defined position and the end-point-defined position were correlated across neurons. In the vertical dimension, the posture-defined position and the end-point-defined hand position were highly significantly correlated ($R = 0.62$, $P = 0.00003$). In the medial-lateral dimension the correlation was also highly significant ($R = 0.58$, $P = 0.0001$). Finally, in the front-to-back dimension the correlation was also highly significant ($R = 0.43$, $P = 0.007$). These results indicate that the preferred hand positions obtained through the end-point model tended to roughly match the hand positions associated with the preferred postures obtained through the end-posture model. One possibility is that the tuning to end point, accounting for a mean of 22% of the neuronal variance, was partly a noisy reflection of a more robust tuning to arm posture, accounting for 36% of the neuronal variance.

Neurons are more end-posture than start-posture tuned

Neurons are more end-posture than start-posture tuned: I. Analysis of temporal offsets

Numerous studies have reported that neuronal activity in motor cortex is modulated by posture (Caminiti, Johnson, and Urbano 1990, Kettner, Schwartz, and Georgopoulos 1988, Paninski et al. 2004, Scott and Kalaska 1995, Sergio and Kalaska 1997, Sergio and Kalaska 2003). The present study supports the same finding, showing in addition that over the larger workspace the postural tuning curves can take the form of a broad, non-linear tuning to a preferred posture. One question, however, is whether this postural tuning reflects the active control of movement, such as movement planning and execution, or whether it reflects feedback signals that inform the motor cortex about the state of the limb. One way to address this question is to ask whether the neurons are tuned specifically to the final posture of a movement, as if in preparation of that end state, or whether the neurons are equally tuned to the starting posture of the movement. We therefore performed the regression analysis described above (*End-posture tuning*) but regressing against the start posture of each movement rather than the end posture to determine if neurons show a similar Gaussian tuning, and a similar degree of fit, to a start-posture model.

The mean R^2 for the start-posture tuning model was 0.33, nearly as high as the R^2 for end-posture tuning (0.36). At first glance this high value would seem to suggest that the neurons are modulated by current or recent posture rather than preferentially encoding the anticipated posture. An analysis of the temporal offsets, however, revealed a different reason for this high

R^2 value. In this section we first describe the temporal offsets for several regression models and then return to the question of start-posture coding.

Figure 3-14 shows the optimized temporal offsets obtained for five of the above regression analyses. In this graph, negative values represent offsets in which the window for neuronal activity was shifted earlier than the movement window. For example, Figure 3-14A shows the offsets obtained for direction tuning (following the method in *Direction Tuning II: Global*). The mean temporal offset for direction tuning was -36 ms (SD = 117 ms) and was not significantly different from 0, indicating that on average the regression against direction was optimized when using neuronal activity that was approximately coincident with the movement. Figure 3-14 also shows the result for the more complex direction tuning model in which the preferred direction could rotate depending on starting hand position (Figure 3-14B); speed tuning as in *Speed Tuning I* (Figure 3-14C); and tuning to end posture (Figure 3-14D). In all of these cases, the offsets were distributed around and not significantly different from 0.

Previous studies reported offsets in which neuronal activity significantly preceded movement (Moran and Schwartz 1999, Paninski et al. 2004, Reina, Moran, and Schwartz 2001). In our data, the temporal offsets for direction tuning trended in that direction but were not significantly different from 0. One possible reason is that in a trained movement task in which the animal responds to task cues, the movement involves a relatively consistent reaction time and planning period, whereas in the present study in the spontaneous movement set the animal may be guiding the movements more on the fly and with less preplanning, resulting in shorter and

more variable temporal offsets. As a result, the lag between neuronal activity and movement no longer reaches significance.

Figure 3-14 also shows that the temporal offsets for speed and direction models were more tightly distributed than for the arm posture model. This pattern resembles previous findings (Paninski et al. 2004). One interpretation is that the control of direction and speed can be exercised only in association with the movement itself since these parameters are not defined when the arm is stationary, whereas the active control of posture must be exercised at all times. Therefore one might expect the neuronal control of posture to extend over a greater range of times than the neuronal control of direction, resulting in a spread-out distribution of temporal offsets. A second speculation (proposed by Paninski et al. 2004) is that the modulation of neuronal activity by posture reflects both feedforward and feedback signals and therefore has a greater range of temporal offsets. As described next, however, other aspects of the data suggest that the postural tuning reflected relatively more of a feedforward signal than a feedback signal.

Figure 3-14E shows the temporal offsets for start-posture tuning. These offsets were shifted with respect to the others, and were highly significantly below 0 ($t = 4.94$, $P = 0.00004$). A comparison between Figure 3-14D and Figure 3-14E shows the difference between end-point tuning, for which the temporal lags were distributed broadly around the movement period, and start-point tuning, for which the temporal lags mainly preceded the movement. To further explore this effect we extended the range of possible offsets downward from –286 ms to -1000

ms and re-ran the start-posture analysis. Figure 3-14F shows the resulting distribution of temporal offsets. The mean is - 413 ms.

These results help to answer the question raised at the beginning of this section. Are the neurons tuned to the anticipated, upcoming posture, as though contributing to the control of the movement, or are they equally tuned to the prior, starting posture, as though reflecting a feedback signal? The activity during a movement is indeed tuned to the final posture of the movement. But it is not really tuned to the start posture. Rather, to obtain a high R^2 value for the start posture, it is necessary to select neuronal activity from well before the start of the movement. In effect, neuronal activity during the *previous* movement is correlated with the end posture *of that previous movement*. This result suggests that the neuronal tuning to posture reflects relatively more an anticipation of the future state of the limb rather than a feedback signal about a recent state of the limb. Both influences probably contribute, but the planned future posture appears to be the larger influence according to this analysis. The next section provides a second, independent approach to the same question and arrives at the same answer.

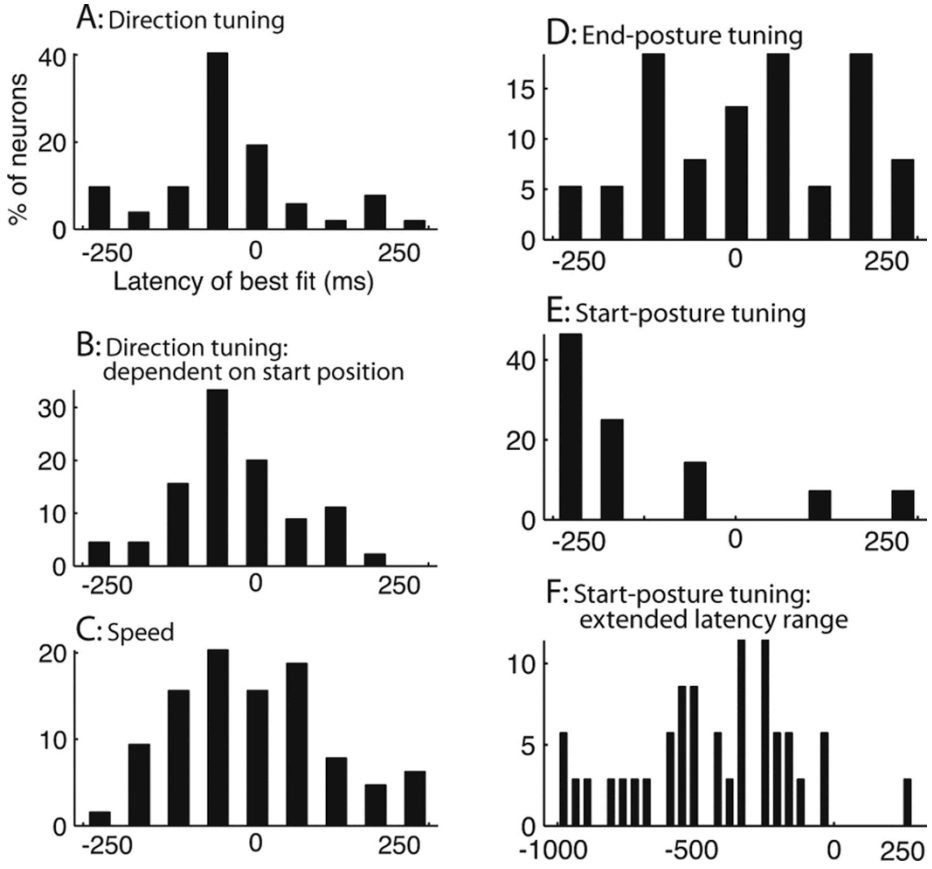


Figure 3-14 : Temporal offsets for six regression models. For each regression model, the kinematic data was selected from a fixed time period during the movement and the neuronal data was selected from a time period equal in duration but offset by a temporal lag. For each neuron, the temporal offset that optimized the regression was selected. These optimal temporal offsets are shown in the frequency histogram in which the x axis represents temporal offset. Negative values indicate neuronal activity leading movement. A. Temporal offsets for direction tuning following the method in *Direction Tuning II: Global.* B. Temporal offsets for direction tuning following the method in *Direction tuning III: rotation of preferred direction with changes in starting hand position*. C. Temporal Offsets for speed tuning following the method in Speed Tuning I. D. Temporal offsets for end-posture tuning. E. Temporal offsets for start-posture tuning. F. Temporal offsets for start-posture tuning in which the range of possible offsets was extended downward to –1000.

Neurons are more end-posture than start-posture tuned: II. Analysis of movements near the preferred posture

With relatively short movements (mean length of 8 cm) the end posture of a movement might often be similar to the start posture. Movements that end in the preferred region may therefore also tend to start in the preferred region and remain in the preferred region throughout the movement. Though the firing rate is typically high during these movements near the preferred posture, is it because the start posture was near the preferred posture, or because the end posture was near the preferred posture?

To approach this issue we performed the following analysis. For each neuron we selected the 10% of movements that terminated nearest to the preferred posture. These movements were typically associated with a high firing rate, since they terminated near the peak of the postural tuning curve (for examples see Figure 3-11B1-4). The purpose of the analysis was to examine this set of movements clustered around the preferred posture, and to determine whether the firing rate was more related to the proximity between the start posture and the preferred posture, or to the proximity between the end posture and the preferred posture.

We performed two linear regressions. First, we asked whether the distance between the start posture of a movement and the preferred posture was significantly related to firing rate. This was not the case. Across neurons, the mean R^2 value for this linear regression was 0.02, and none of the neurons showed a significant relationship. The firing rate during movement tended to be high because the movements were near the preferred posture, but the firing rate tended to

be uniformly high regardless of the proximity of the start posture of the movement to the preferred posture.

In contrast, using the same subset of movements, we asked whether the distance between the end posture of a movement and the preferred posture was significantly correlated with firing rate. A relationship was obtained in which the firing tended to be higher for movements that terminated closer to the preferred posture. The mean R^2 value for this linear regression was 0.4, and 71% of the neurons showed a significant relationship.

Thus among the subset of movements near the preferred posture, the firing rate was dependant on the end posture and independent of the start posture. This result, like the result of the previous section, strongly suggests that the postural tuning obtained here reflects primarily an anticipation of the planned state of the limb rather than a feedback signal about the recent state of the limb.

Combined tuning: speed, direction, end point, and end posture

We combined speed tuning, direction tuning, end-point tuning, and end-posture tuning in a hierarchical regression to determine how much of the total variance can be explained by these variables. We first regressed against end posture; then regressed the residual against end point; then regressed the second residual against the direction-tuning model in which the preferred direction could rotate systematically with changes in the initial hand position; and finally we

regressed the third residual against the mean speed of the hand during each movement. Because of the possible partial overlap of the different tuning models, the combined R^2 value is expected to be less that the sum of the separate R^2 values given in the previous sections. Figure 3-15 shows the R^2 values obtained by this sequential regression. On average, 44% of the variance was explained by the combination of parameters. The remainder of the variance may be attributable to the many other parameters critical to movement but not tested here; to the changing motivational or attentional state of the animal; and to noise in the measurement of arm movement and neuronal activity.

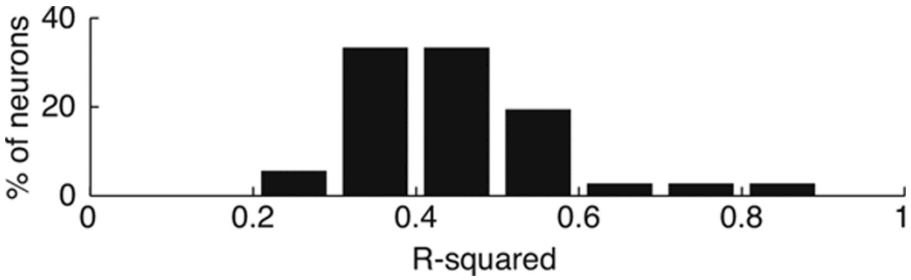


Figure 3-15 : Frequency histogram of R^2 values for hierarchical regression. This regression model included tuning to a preferred final posture; tuning to a preferred final hand position; tuning to a preferred direction that can rotate with changes in the starting hand position; and a linear dependence on mean speed.

Simulated neurons

To test the validity of the above regression models, we modeled artificial neurons. One artificial neuron was direction tuned. The data was generated using the cosine direction-tuning

model with Gaussian noise added to create a neuron that was noisily tuned to a preferred direction. We simulated the firing rate that this artificial neuron should have during a set of actual, recorded hand movements. These firing rates were then subjected to the regression analyses described above. Figure 3-16A shows the result. When tested with a direction-tuned model, the artificial direction tuned neuron showed a mid-range R^2 value consistent with its noisy tuning. When tested with the other regression models, it showed near zero R^2 values. Thus when testing a neuron that is truly direction tuned, the analyses extracted a directional tuning signal and not a spurious hand position tuning or posture tuning signal.

In a similar manner we generated an artificial neuron that was tuned to final hand location in space, and an artificial neuron that was tuned to final posture of the arm. As shown in Figure 3-17B and C, the R^2 value for each type of neuron was highest for the matching type of regression model and near zero for the non-matching regression models. Finally we generated an artificial neuron that was not tuned. Its firing rate was randomized. Figure 3-17D shows that all regression models returned a low R^2 value for this artificial neuron, with no tendency for the models that had more regressors to extract a higher R^2 value.

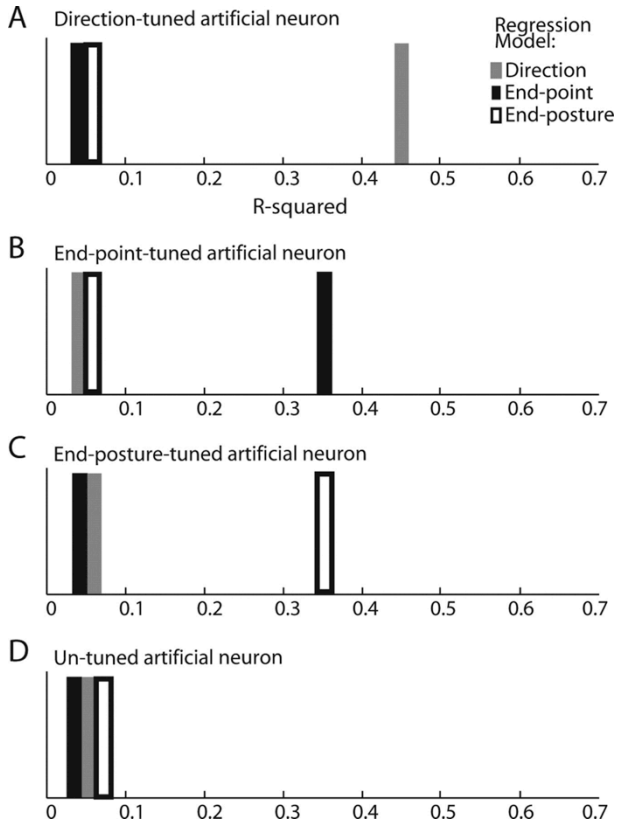


Figure 3-16 : Analysis of simulated noisy neurons. A. Simulated direction-tuned neuron tested on direction, end-point, and end-posture regression models. B. Simulated end-point tuned neuron tested on the same three models. C. Simulated end-posture tuned neuron tested on the same three models. D. Simulated neuron with no tuning on the same three models.

Muscle activity

One possibility is that a neuron may not be tuned to posture per se but instead may be correlated to muscle activity which in turn may be correlated with posture. To address this question we repeated our analysis on EMGs recorded from the biceps to see if a muscle activity

model could account for our results. We predicted that the EMG signal from the biceps would show good tuning to final posture, but that the tuning would differ from the typical tuning of our motor cortex neurons in two ways. First, the neurons often were related to movement of more than one joint whereas the biceps EMG should be primarily related to the elbow. Second, the neurons often preferred movement toward an intermediate joint angle whereas the biceps EMG was expected to prefer a flexed elbow angle. In these ways the muscle signal was expected to be simpler than the neuronal signals.

As expected, the biceps EMG was largest during flexion of the elbow joint. Figure 3-17A shows EMG traces during rotations of the elbow. Flexions of the elbow were associated with higher EMG activity than extensions of the elbow. Just as for the neuronal analysis, we performed a linear regression of the biceps EMG against joint angular velocity with eight degrees of joint motion. The total R^2 value for this regression was 0.41. Figure 3-17B shows the individual contributions of each joint when the covariance with the remaining seven joints was regressed out. As expected, the elbow joint shows the dominant effect. The three degrees of shoulder rotation show a small effect, probably because the movement of the shoulder is to a small degree mechanically coupled to the elbow. The four distal degrees of freedom show essentially no effect. The overall pattern shows the activity of the biceps is related primarily to rotation of the elbow joint.

We then analyzed the biceps EMG using the model of a tuning to a preferred final posture. Because the biceps EMG is largest during elbow movements toward flexion, the signal

was expected to be relatively well tuned to final posture, in which the preferred final posture was a flexed elbow. The R^2 value for this tuning to final posture was 0.49. The EMG signal was indeed tuned to a flexed elbow posture. The preferred elbow angle obtained by the regression fit was 2 degrees, outside the normal mechanical range of the elbow (~30-180 degrees). That is, the biceps EMG was on average larger for movements that terminated in more flexed postures, resulting in an extrapolated preferred posture more flexed than the arm could actually go. In contrast, for most neurons, the regression analysis obtained a preferred final posture within the normal mechanical range of the elbow. Only one neuron had a preferred elbow angle outside of that range. The preferred final elbow angle for the biceps EMG was so extreme, that it was highly significantly smaller than that obtained for the neurons ($P = 10^{-8}$).

These results show, as expected, that the biceps EMG activity was related primarily to the elbow and was tuned primarily to flexion, not to movement toward a posture in the normal range of the arm. In these ways the pattern of results for the muscle was simpler and much clearer than the pattern described above for the neurons in motor cortex. It cannot be ruled out that some of the neuronal signals studied here are correlated primarily with muscle activity. Indeed, considerable direct control of muscle activity seems likely. However, to the extent that each neuron was related to many joints, and to the extent that neurons often preferred movement of joints to intermediate angles rather than to extremes, the pattern of results was too complex to be entirely explained by a simple linear relationship to muscle activity.

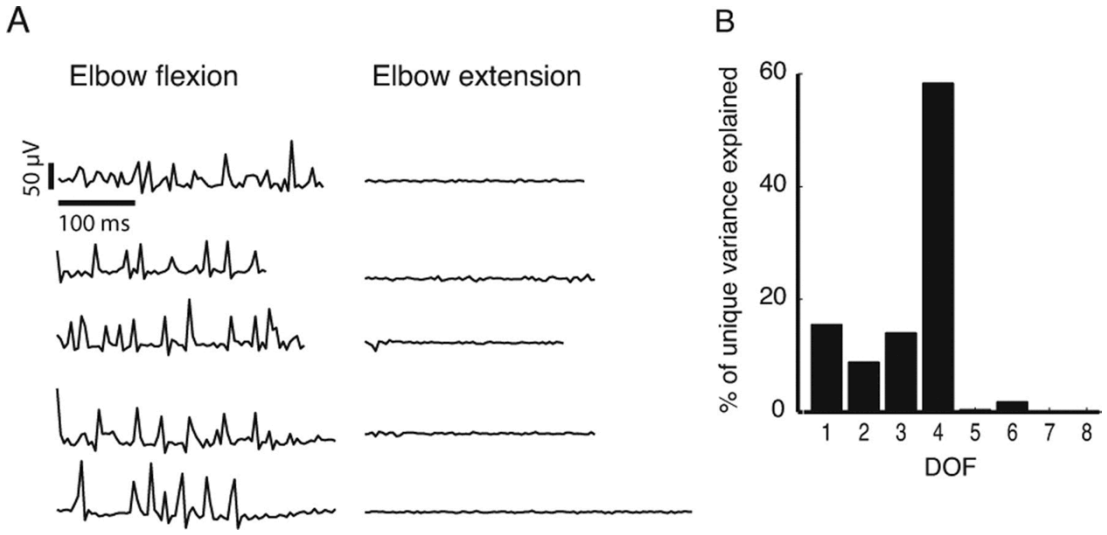


Figure 3-17 : EMG results for biceps during naturalistic movements. A. Right: EMG traces during movements involving elbow flexion. Left: EMG traces during movements involving elbow extension. B. An 8 dimensional regression against joint velocity was performed and returned an $R^2 = 0.41$. The unique contribution for each joint, when the contributions of the other 7 joints were regressed out, was calculated. The figure shows the percent of contribution to the R^2 for each joint. As expected, the elbow joint showed the largest effect. Joints on X axis are, in order: shoulder azimuth; shoulder elevation; shoulder external/internal rotation; elbow extension; forearm pronation; wrist extension; wrist adduction; grip aperture.

Comparison Between Neurons and Stimulation

Figure 3-18A shows data from an example neuron in a format similar to Figure 3-11A1-4. Approximately 40% of this neuron's variance was attributable to the Gaussian tuning to the preferred posture ($R^2 = 0.38$; $P < 0.0001$). After recording from this neuron, before withdrawing the electrode, we electrically stimulated the same site in cortex with a 500-ms train of biphasic pulses (0.2 ms per phase, negative phase leading, 200 Hz, 50 µA), and the arm moved to a final

posture. Such stimulation-evoked postures were detailed in Chapter 2. In the present example, data were collected from 30 stimulation trials, and the mean final posture was obtained. The red vertical line in Figure 3-18A shows the location of this stimulation-evoked posture relative to the neuron's preferred posture. The stimulation-evoked posture was near the peak of the neuron's Gaussian tuning function, 0.4 SD away from the neuron's preferred posture.

The postures evoked at all other cortical sites tested are also shown on the same graph as vertical black lines. They are all farther from the neuron's preferred posture (>1 SD away from the preferred posture). Thus, the neuron's preferred posture matched best to the posture obtained on stimulation of the same cortical site and matched less well to postures obtained on stimulation of all other sites. This result shows a specific correspondence between the properties of the neuron and the results of stimulating the same site in cortex.

To parameterize this result, we ranked the concordant stimulation site (stimulation site at the same cortical location as the neuron) in comparison to the discordant stimulation sites (stimulation sites at different cortical locations from the neuron) in terms of how closely they matched the preferred posture of the neuron. For the example neuron in Figure 3-18A, the concordant stimulation site ranked first, in the 90th-to-100th percentile range. For each neuron, we obtained a percentile in this fashion. If there was no overall match between neuron properties and stimulation results, then the percentiles should be randomly distributed between 0 and 100. If there was a match between neuron properties and stimulation results, then the percentiles should be skewed to the high end. As shown in the population data in Figure 3-18B, the percentiles are

not randomly distributed. They are skewed to the high end, with the distribution significantly above the 50th percentile (binomial test, $P < 0.0001$). Indeed, half of the neurons are above the 80th percentile.

These results indicate that there is a significant degree of match between the postures preferred by neurons as determined by the regression analysis and the postures evoked by electrical stimulation of the same sites in cortex. However, the match is not absolute, with neurons showing a range of other percentiles.

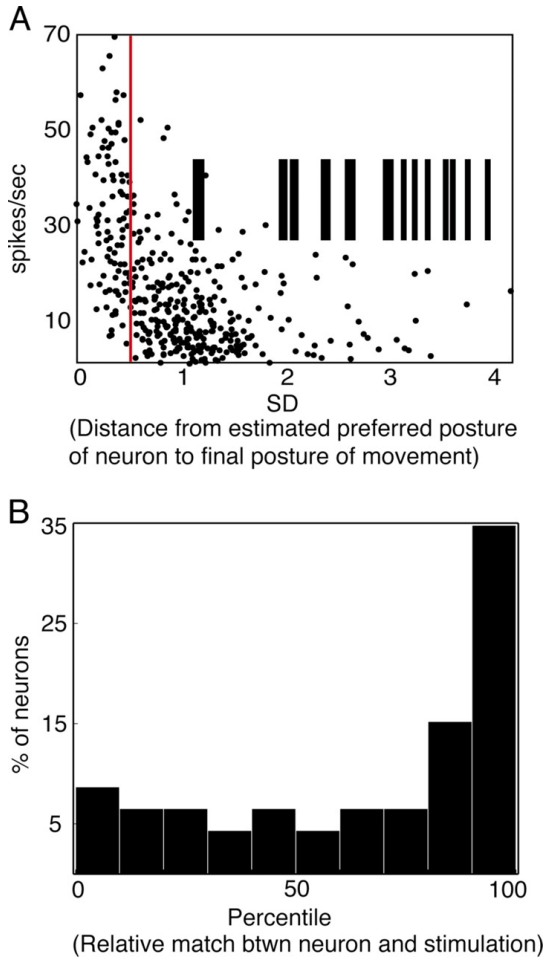


Figure 3-18 : Comparison of neuronal tuning and stimulation-evoked postures. (A) Data from an example neuron. The preferred final posture was determined by regression analysis. The final posture of each movement was compared to the preferred posture. The distance between them was calculated in 8D posture space. Distance was measured in units of standard deviations of the Gaussian tuning function to express all eight dimensions in posture space in equivalent units. This distance is plotted on the x axis, and firing rate during the movement is plotted on the y axis. The data fit the Gaussian tuning function with an R^2 value of 0.38 ($P < 0.0001$). Electrical stimulation of the same cortical site evoked a final posture. The mean stimulation-evoked posture is plotted (red vertical line) and was near the preferred posture of the neuron. The mean stimulation-evoked postures from other cortical sites (black vertical bars) were farther from the preferred posture of this neuron. (B) For each neuron a percentile was computed indicating how closely the neuron's preferred posture matched the posture evoked by stimulation of the same cortical site as compared to stimulation of other cortical sites. This distribution is not uniform; it is significantly greater than 50% (binomial test, $P < 0.0001$).

DISCUSSION

The present results suggest that the standard tuning properties of motor cortex neurons described previously in trained tasks are also present during semi-naturalistic, unstructured arm movements. However, the amount of neuronal variance explained by standard directional tuning may be less than previously suspected while the variance explained by postural tuning may be more. On average speed tuning accounted for about 1% of the total variance; simple direction tuning for 8%; a more complex version of direction tuning for 13%; tuning to the final position of the hand for 22%; and tuning to the final posture of the arm for 36%.

Interpretation of R^2

All R^2 values in the present study are below 50%. Do these seemingly low values indicate a weak or invalid result? A high R^2 value for a particular source of variance will be obtained only if all other sources of variance are minimized. In the case of speed tuning it is possible, by averaging out the variance from other sources, to obtain an R^2 value that surpasses 90% (*Speed Tuning I*). Yet when all sources of variance are admitted, although speed tuning remains significant, it explains only about 1% of the variance. In this case the R^2 value can

change for some neurons from 90% to 1% depending on the manner of the test, while still reflecting the same, valid speed-tuning signal.

Tuning to a final multijoint posture of the arm accounted for the largest share of the neuronal variance: 36%. An incorrect interpretation might suggest that this model cannot be valid since it explained far less than 100% of the variance and even less than 50%. Yet the fact that the postural tuning captures a third of the variance, in a grab bag with all other sources of variance present, suggests that it is a major contributor to the signal.

Presumably the use of the arm for complex, normal behavior requires the control of numerous variables including hand speed, joint angle, muscle force, final posture, path curvature, grip size, the relative timing of arm and grip movement, and others. These variables are relevant to different extents in different movements. They are all presumably reflected in motor cortex neuronal activity. The variance in neuronal activity should be related in some proportion to each of these control variables, each one capturing a small part of the total variance. The present findings are consistent with the view that neurons are tuned to many control variables important in the animal's repertoire, and that posture is an especially prominent control variable capturing about a third of the total variance.

One major caveat is that many of the movement variables are likely to be correlated with each other. If a neuron contributes to the control of variable A, but variable A is correlated with variable B, then the neuron may appear on analysis to be tuned to B. Therefore our results showing that neurons are tuned to many parameters must be taken cautiously. It cannot be

excluded that some of the tuning to different movement variables reported here is a result of correlations among the variables. This difficulty of correlations among movement variables has been noted before and has led to some alternative interpretations to direction tuning (Scott and Kalaska 1997, Scott and Kalaska 1995). Arguably the correlative nature of single-neuron experiments has been the main obstacle to interpreting single-neuron results in motor cortex.

Local direction tuning and global posture tuning

In the now classical paradigm of Georgopoulos and colleagues (Georgopoulos et al. 1982, Georgopoulos, Schwartz, and Kettner 1986), monkeys perform a center-out task, reaching from a central location to surrounding locations. Most neurons are broadly direction tuned, preferring one direction of reach and firing less during neighboring directions. Furthermore, when different starting and ending positions are tested, direction tuning still predominates (e.g. Caminiti, Johnson, and Urbano 1990, Fu, Suarez, and Ebner 1993, Kettner, Schwartz, and Georgopoulos 1988). Yet the present experiment suggests that in a more global range of movements, the neuronal variance is only about 8% attributable to this type of direction tuning. Other types of tuning, such as end-point and end-posture tuning, explain more of the variance.

These seemingly disparate results are not, however, in contradiction. In a center-out task and its variants, the direction of the hand is typically explored through a full range of angles. Yet the posture of the arm is tested over a narrow range. In most experiments, the hand is limited to

a central region of space comprising less than 5% of the volume tested here. Since the tuning curves for posture and hand position are broad, they will drive little of the neuronal variance if tested over this small range. As a result, when testing local regions of space with a full range of hand directions, direction tuning will tend to predominate.

Our results replicate previous findings that the neurons are locally but not globally direction tuned. The directional tuning changes when the hand is placed in a new starting position or when the arm is given a new starting posture (Caminiti et al. 1990; Scott and Kalaska 1995, 1997; Sergio and Kalaska 2003). The present findings further suggest that the behavior of the neurons over the global movement set may be better described as a broad tuning to a preferred final posture.

One suggests that neurons compensate for a force load on the hand in a manner that emphasizes either postural coding or directional coding (Kurtzer, Herter, and Scott 2005). Other recent work suggests that in the human, postural coding may be emphasized in the right hemisphere in motor cortex and directional coding may be emphasized in the left hemisphere in motor cortex, though both types of processing are of course represented in both hemispheres (Bagesteiro and Sainburg 2002). Exactly how these results relate to the present study is not yet clear. Most normal actions require both a postural component and a directional component, as well as control of other aspects of movement. Yet it may be that some actions emphasize postural control whereas others emphasize directional control.

Partial Match Between Neurons and Stimulation

Our results show a significant degree of match between the postures preferred by neurons and postures evoked by electrical stimulation of the same cortical site. Yet the match is not absolute. The mismatch may have a variety of sources. First, many neurons in motor cortex may simply not be posture-tuned. Indeed, some of the neurons showed low R^2 values for posture tuning, <0.2. Second, even the neurons that are posture-tuned are noisy, and thus the estimate of their preferred posture may not be precise. Third, stimulation directly activates a ball of neighboring neurons that might have a diversity of properties; thus, stimulation may cause an averaging or summing of these properties. These sources of error are likely to have diluted the match between neuronal properties and stimulation effects.

It has long been known that stimulation is helpful in discovering basic motor or sensory properties encoded in cortex, especially because the technique is causal rather than correlational. But because of the limitations of the technique, especially the direct stimulation of nearby neurons with diverse properties, the strongest evidence is always provided by a convergence of different techniques on a similar answer. In the present case a convergence of stimulation effects and single-neuron properties strongly supports the hypothesis that motor cortex neurons control movement partly by specifying a final posture.

Chapter 4 Possible origins of the complex topographic organization of motor cortex: reduction of a multidimensional space onto a 2-dimensional array

Introduction

Many different variants of functional topography have been proposed for the lateral motor strip in the primate precentral gyrus. Some of these variants are shown in Figure 4-1. These proposed organizations include versions of a somatotopic map (A-C) (Beevor and Horsley 1887, Foerster 1936, Kwan et al. 1978, Park et al. 2001, Woolsey et al. 1952), and divisions of the precentral gyrus into primary motor cortex and several sectors of caudal and rostral premotor cortex (D-F) (Campbell 1905, Dum and Strick 2002a, 2005b, Matelli, Luppino, and Rizzolatti 1985, Preuss, Stepniewska, and Kaas 1997). Two additional types of organization were suggested by recent experiments using electrical stimulation. We mapped the motor cortex in monkeys using long-train electrical stimulation to evoke complex, multi-joint movements (Cooke and Graziano 2003, Graziano, Aflalo, and Cooke 2005, Graziano, Patel, and Taylor 2004, Graziano, Taylor, and Moore 2002). As expected, we obtained a rough, overlapping somatotopic organization consistent with previous studies, with the feet in a medial location and

the mouth in a lateral location. In addition, within the arm and hand representation, we obtained a rough map of the location in space to which the hand moved on stimulation (Figure 4-1G, H, see also Figure 2-10 for data from second monkey.) This hand location map encompassed the traditional primary motor cortex and the caudal parts of premotor cortex. We also obtained a partitioning of motor and premotor cortex into zones that emphasized different complex, ethologically relevant movements (Figure 4-1; see chapter 2). These additional topographies were unexpected but may help to explain the noise and overlapping within the somatotopic progression. We hypothesized that the simultaneous presence of several noisy and intermeshed types of topographic organization might be the result of a competition among conflicting mapping requisites.

(Kohonen 1989, 2001) suggested that cortical maps may self-organize to optimize nearest neighbor relationships. In this view, neurons that process similar information are located near each other, minimizing wiring length. Self-organization can lead to fractured or seemingly disordered maps in certain cases, due to the fact that the cortex is two-dimensional yet the relevant parameter space may be of higher dimensionality. For example, the principle of self-organizing maps was used to explain the complex organization of primary visual cortex (Durbin and Mitchison 1990), in which line orientation and ocular dominance interact to produce a pattern of cortical pinwheels (Obermayer and Blasdel 1993). Likewise, some of the apparent disorder in the tonotopic map in primary auditory cortex has been attributed to a self-

organization that is under the influence of many dimensions reduced onto a two-dimensional sheet (Schreiner 1995).

In the present study we modeled motor cortex topography using a Kohonen network. We gave the map an initial somatotopic organization. We then allowed the network to re-organize under the influence of two additional mapping requisites: hand location in space, and ethologically relevant action categories that combine more than one body part. We hypothesized that the constraint of optimizing nearest-neighbor relationships in this multi-dimensional space would force the map to express many of the specific, complex features found in the monkey motor cortex.

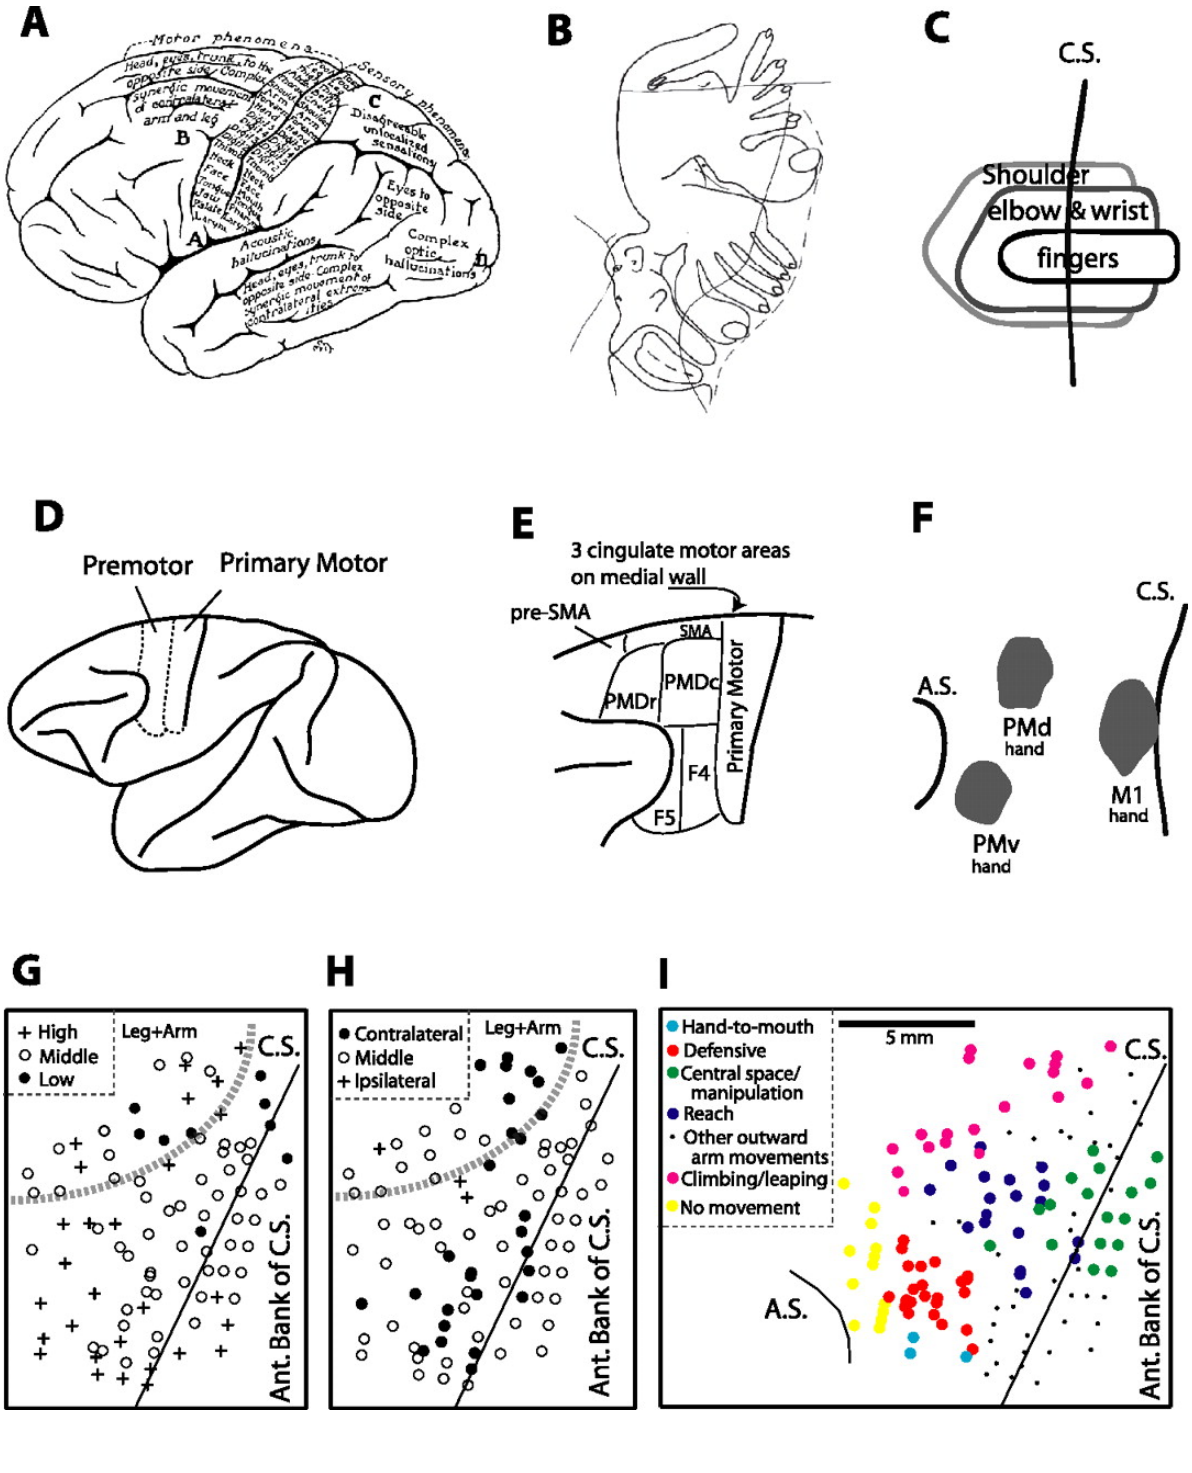

A
B
C
C.S.
Shoulder
elbow & wrist
fingers
D
Premotor
Primary Motor
E
3 cingulate motor areas
on medial wall
pre-SMA
SMA
PMDr
PMDc
Primary Motor
F4
F5
F
C.S.
A.S.
PMd
hand
M1
hand
PMv
hand
G
+ High
○ Middle
● Low
Leg+Arm
C.S.
Ant. Bank of C.S.
H
● Contralateral
○ Middle
+ Ipsilateral
Leg+Arm
C.S.
Ant. Bank of C.S.
I
Hand-to-mouth
Defensive
Central space/
manipulation
Reach
Other outward
arm movements
Climbing/leaping
No movement
5 mm
C.S.
A.S.
Ant. Bank of C.S.

Figure 4-1 : Eight different possible organizations to the lateral motor cortex in the monkey brain. A. A sequence of body parts arranged vertically with the feet in a medial location and the mouth in a lateral location(Foerster 1936, Fritsch and Hitzig 1960, Fulton 1938). B. A subsequent refinement of the map including a horizontal component to the organization, in which the axial musculature is represented anteriorly and the distal musculature posteriorly (Woolsey et al. 1952). C. Another proposed organization in which a posterior hand representation is partly surrounded by an arm representation (Kwan et al. 1978, Meier et al. 2008, Park et al. 2001). D. Division of the precentral gyrus into two strips: primary motor cortex that tended to represent individual body parts, and the more complex premotor cortex (Campbell 1905). E. Partitioning of premotor cortex into a dorsal and ventral sector, each sector subdivided into a caudal and a rostral part (Matelli, Luppino, and Rizzolatti 1985, Preuss, Stepniewska, and Kaas 1997). F. Three regions in the lateral motor strip that project directly to the hand portion of the spinal cord: the primary motor hand area, the ventral premotor hand area, and the dorsal premotor hand area (Dum and Strick 2002a, 2005b). G. Rough map of the height of the hand as evoked by electrical stimulation. Data from one monkey (Graziano, Aflalo, and Cooke 2005). Upper hand locations are in a ventral location, lower hand locations are in a dorsal location. Within the arm+leg region of the somatotopy, the mapping of the hand height breaks down. H. Data from the same monkey as in G, but now showing the lateral position of hand evoked by electrical stimulation. I. Cortical patches that, when stimulated, result in different ethologically relevant movements. The details differ among monkeys but the overall pattern is consistent. See Figure 2-10 for data from a second monkey.

Methods

The Kohenen network is a standard tool for solving the problem of dimensionality reduction, or the problem of representing a multidimensional space on a lower dimensional space such that neighbor relationships are optimized. For the present purpose, the self-organizing Kohonen network is not meant to model the specific neuronal interactions or learning algorithms of the brain. Rather the network is merely an analytic tool that optimizes topographic continuity. No assumptions are made here about the learning algorithms that might be implemented in the brain to produce smoothly varying topographic maps.

The final product of a traditional self-organizing map algorithm is an array of nodes that provides a structured representation of the input data. Each node represents a subset of the input

data, and nodes that are topographically near each other in the array tend to represent similar input data. This topographic feature is the natural result of the learning algorithm: when an input is presented to the network, the node that best matches that input, and the topographically surrounding nodes, are updated to more closely represent that input. This local smoothing effect leads to a global structuring of the data.

In our study, the array of nodes was a 36x20 two-dimensional grid that formed the artificial motor cortex. Each node represented a movement. This movement representation was expressed in the form of an 18 dimensional vector, called a codebook vector. Each vector specified the body parts that were being moved (ten dimensions, one for each body part), the position that the hand reached in Cartesian space (three dimensions), and the ethological category to which the movement belonged (five dimensions corresponding to five categories). These aspects of the codebook vector are explained in the sections below.

Somatotopy

Figure 4-2A shows Woolsey's representation of the somatotopic organization of motor cortex in the monkey (Woolsey et al. 1952). This somatotopy is one of the most complete in the literature because it covers most of the precentral gyrus, and shows a horizontal (trunk to hand) as well as vertical (head to toe) organization. This somatotopy is of course only approximate, and leaves out details of overlapping and repeated representations commonly reported in the

monkey. We took this simple, global somatotopy as a starting point for the self-organizing map. The rationale for beginning with a simple somatotopy was the developmental studies of Martin et al. (2005) showing that motor cortex may indeed begin with a pre-wired somatotopy that is then altered and blurred through learning. As described in the results, during learning, this initial simple somatotopy developed into a more complex one with overlapping and repeated representations.

The overlay in Figure 4-2A depicts ten discrete somatotopic subregions corresponding to the tongue, lips, jaw, upper face, neck, torso, arm, hand, leg, and toe, approximately matching the underlying map. A schematized, blocked version of the Woolsey map was used for simplicity because the exact borders of the initial somatotopic components turned out to have little effect on the final state of the network. These somatotopic regions covered a rectangular area with a missing wedge in the posterior lateral corner. Therefore our node array had this same shape of a rectangular area with a missing wedge in the posterior lateral corner. To operationalize the somatotopy in our node array, a 10-element codebook vector was created with each element corresponding to a particular body part. This vector took the form of [tongue, lips, jaw, upper face, neck, torso, arm, hand, toe, leg]. As an initial state, before any learning, each node in the array was assigned a codebook vector in which body parts took a value of either 0, indicating that the node did not represent that particular body part, or 1, indicating that the node did represent the particular body part. In this way, the array was initialized to match the Woolsey map. For example, nodes in the hand area were initially characterized as

[0,0,0,0,0,0,0,0,1,0,0]. During learning, these initial values could change in a continuous fashion. In this way, the initial, cleanly segregated somatotopy could be altered and nodes could come to partially represent any combination of body parts.

Hand Position

Hand position was represented by adding three dimensions (x, y, and z position) to the 10-dimensional somatotopic description. All hand positions were encoded in cm relative to the mouth, with positive x as the vertical height in centimeters above the mouth, positive y as the horizontal distance in the contralateral direction away from the mouth, and positive z as the distance forward of the mouth along the line of sight. For example, a node representing hand position in lower, central space might be characterized as [-20, 0, 0]. Hand position was only included for those nodes at which the arm component of the somatotopy was non-zero. If the arm component had a value of 0, then the values for hand position were left empty and did not actively contribute to the formation of the topographic map. Hand positions were initialized in the node array as small random numbers uniformly distributed between -.1 and .1. Thus, initially, the array contained no topographic map of hand position.

Ethological Category

Finally, movements were categorized according to an ethological description of the movement. These ethological categories were based on our previous finding of a set of discrete movement categories evoked from specific sites in motor cortex. Each category of movement included a range of different movements all similar in that they served a common behavioral function. To represent these categories, we used a five dimensional vector, each dimension corresponding to a particular category of behavior. For example, [1,0,0,0,0] indicated a movement that belonged entirely to the first ethological category but not to any of the other categories. Ethological categories were initialized in the array as small random numbers uniformly distributed between -0.1 and 0.1. Thus, initially, the array contained no systematic representation of ethological categories. The five categories are described in the following sub-sections.

Hand-to-mouth movements:

Hand-to-mouth movements are common in the natural behavior of monkeys (Graziano et al. 2004) and were also obtained on electrical stimulation of certain regions of the motor cortex (Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). These movements include a closure of the hand into a grip posture, a movement of the hand toward the mouth, a movement of the mouth, and a rotation of the head to bring the mouth in alignment with the hand. In our self-organizing network, this type of movement was defined as combining the

tongue, lips, jaw, neck, arm, and hand. The distribution of hand positions associated with hand-to-mouth movements was defined as having a mean position of [0, 0, 3] with standard deviations [1, 1, 1]. This distribution approximates the distribution of hand-to-mouth movements obtained on stimulation of the monkey cortex, and is shown by the light blue dots in Figure 4-2B-D. A node in the motor cortex array that represents specifically and only a hand-to-mouth movement might have a codebook vector of [(1,1,1,0,1,0,1,1,0,0)(0.5,0,1)(1,0,0,0,0)] where the first set of parentheses shows the somatotopic components, the second set shows the hand location, and the third set shows the ethological category. Including an explicit representation of ethological category in the codebook array allows movements that are different from each other to none-the-less belong to the same ethological category; as a result, during re-organization, these movements tend to become represented near each other.

Defensive movements:

Three different but related types of movement were included in the ethological category of defensive movement. The first type was a facial defensive movement. This movement was studied in monkeys exposed to an air puff on the face (Cooke et al. 2003), and was also observed on chemical or electrical stimulation of specific parts of motor cortex (Cooke and Graziano 2004a, 2004b, Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). The movement included a blink, a squinting of the musculature around the eyes, a lifting of the upper lip and sometimes opening of the jaw, a folding of the ear against the side of the head, and a retraction of the head. In our self-organizing network, therefore, this type of defensive

movement was defined as combining the lips, jaw, upper face, and neck. No arm movement or hand position was included. A node in the array that represents this kind of defensive movement might have a codebook vector of [(0,1,1,1,1,0,0,0,0,0)(-,-,-)(0,1,0,0,0)], where the three dimensions of hand position in space are left blank because the arm is not involved (arm component = zero).

The second type of defensive movement included a facial and arm movement. This movement was also observed in monkeys exposed to an air puff to the face (Cooke and Graziano 2003), and was observed on chemical or electrical stimulation of specific parts of motor cortex (Cooke and Graziano 2004a, 2004b, Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). The movement included all the components of a facial defensive movement, but also included a shrugging of the shoulder and a lifting of the hand to a region of space near the face as if to block an impending impact. In our self-organizing network, therefore, this type of defensive movement was defined as combining the lips, jaw, upper face, neck, torso, and arm. It was also defined as involving a final hand position in upper space near the side of the face. A node in the array that represents this kind of defensive movement might have a codebook vector of [(0,1,1,1,1,1,1,0,0,0)(2,8,0)(0,1,0,0,0)].

The third type of defensive movement was an arm retraction. This movement was typical of monkeys exposed to air puff on the arm or hand (Cooke and Graziano 2003), and was also observed on chemical or electrical stimulation of specific regions in cortex (Cooke and Graziano 2004a, 2004b, Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). This

movement involved a rapid withdrawal of the hand to a region at the side or back of the torso, accompanied by a twisting motion of the torso. In our self-organizing network, therefore, this type of defensive movement was defined as combining the torso and arm, and involving hand positions in the lower space near the side or back of the torso. The hand positions related to defensive movement, as defined for the purposes of this model, are shown in Figure 4-2B-D (red dots). A node in the array that represents this kind of defensive movement might have a codebook vector of [(0,0,0,0,0,1,1,0,0,0)(-20,8,0)(0,1,0,0,0)].

Central space manipulation:

In monkeys, manipulation of objects such as food or toys typically occurs in a large ball of central space in front of the chest (Graziano et al. 2004). Stimulation of specific regions in the motor cortex of monkeys evoked movements of the fingers and wrist that resembled manipulation, and also evoked accompanying movements of the arm that brought the hand into central space (Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). In our self-organizing network this category of movement was defined as combining the hand and arm to varying proportions. Hand Positions were given a mean position of [-12, 0, 6] with the standard deviations [3, 3, 3] (Figure 4-2B-D, green dots). A node that represents a manipulation movement might have a codebook vector of [(0,0,0,0,0,0,1,1,0,0)(-10,0,10)(0,0,1,0,0)].

Reaching:

Another salient type of movement observed on stimulation of motor cortex was a reaching-like movement in which the hand shaped as if to grasp an object, the arm projected outward, and the shoulder and torso rotated to contribute to the reach (Graziano, Aflalo, and Cooke 2005). In our self-organizing network this category of movement was defined as combining the hand, arm, and torso. The mean hand position was defined as [-7, 0, 16]. The outer boundary of this distribution of hand positions was made to follow a curve that roughly matched the expected boundary of a monkey's reach (outer diameter of curve 30 cm) (see Figure 4-2B-D, dark blue dots). A node in the array that represents a reach-to-grasp movement might have a codebook vector of [(0,0,0,0,0,1,1,1,0,0)(-6,0.2,17)(0,0,0,1,0)].

Climbing:

A final type of complex movement observed on stimulation of motor cortex was a climbing-like movement that included all four limbs and the torso (Graziano, Aflalo, and Cooke 2005). For example, during stimulation of a climbing-like site the monkey might raise both arms placing the hands into upper lateral space, push down on the contralateral leg as though to push off the ground, and raise the ipsilateral leg. In our self-organizing network this category of movement was defined as combining the arm, leg, and torso. Hand position was given a mean of [-5,7,7] with standard deviation of [4, 2, 2] (see Figure 4-2B-D, pink dots with black border). A node in the array that represents a climbing movement might have a codebook vector of [(0,0,0,0,0,1,1,0,1,0)(-4,8,7)(0,0,0,0,1)].

The definitions of movements given above incorporate specific details taken from the monkey data. For example, for purposes of the model, the central-space/manipulation movements are defined in a specific manner and given specific characteristics borrowed from the movements evoked from monkey cortex and the movements observed in the spontaneous behavior of monkeys. One possible concern is that the details of these movements might not be perfectly accurate. Should the central-space/manipulation movements be centered exactly at a mean position of [-12, 0, 6] with standard deviations [3, 3, 3]? Should reaching movements extend to a furthest limit of 30 cm? Would altering these details of the model slightly cause a substantial change in the results? As described in the results section, the model turned out to be highly robust despite changes in these details of the movement definitions. As long as hand-to-mouth movements are roughly near the mouth, and a high percentage of manipulation takes place roughly in front of the chest, and the other movement categories are also roughly in the expected range, then the result is essentially the same.

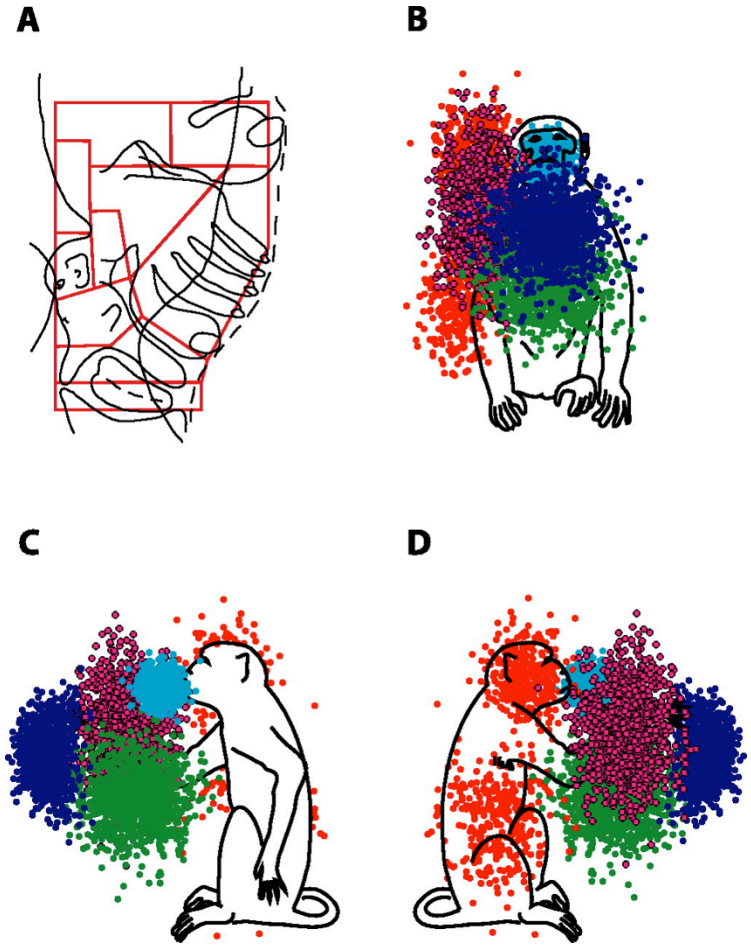


Figure 4-2 : Some movement definitions used in the motor cortex model. A. Woolsey's map of the somatotopic organization in monkey motor cortex (Woolsey et al. 1952), overlaid with blocked regions showing the schematized somatotopy used in the present motor cortex model. B-D. Three views of a schematized monkey showing the hand locations defined for different ethological categories of movement. Light blue = hand-to-mouth; dark blue = reach; red = defense; green = central space/manipulation; pink with black border = climbing.

Input Vectors

As described above, each node in the network was assigned an initial codebook vector. The codebook vector indicated the movement or combination of movements currently represented by that location in cortex. We then defined a set of input vectors. The input vectors were the movement set on which the network was trained. Each input vector represented a specific movement and took the same form as a codebook vector. For example, an input vector of [(1,1,1,0,1,0,1,1,0,0)(0.5,0,1)(1,0,0,0,0)] would represent a hand-to-mouth movement.

The network was trained on 12,800 different input vectors. These input vectors included examples drawn in equal proportion from the ethological categories. To form a more complete, realistic movement set, the input vectors also included: movements of the foot and leg; movements of the mouth that combined the lips, jaw, and tongue (to mimic chewing); and movements of the arm to a sampling of locations in space. Monkeys spend a high proportion of time with the hand placed in three typical positions: the space just in front of the chest (such as for manipulation), the space just in front of the mouth (such as for manipulating food at the mouth), and lower space (such as for bracing the weight of the body during sitting) (Graziano et al. 2004). We constructed the input array to contain an equal number of vectors in which hand location was in central space, in front of the mouth, and in lower space, in order to include the normal range of a monkey's movements.

The exact proportions of these different types of movements turned out not to be critical. As described in the results section, the final map was similar even when the proportion of

movements was altered from equal proportions of different movement types to proportions that more closely approximated the normal repertoire of a monkey.

Learning Rules

For details of the self-organizing learning algorithm and its implementation, see Kohonen (2001). Here we describe only those aspects of the technique for which there was more than one possible way to apply the method, requiring us to choose the approach that suited our particular application.

Before the learning algorithm was applied, each dimension of the input array was normalized such that values varied between -1 and 1. In this manner no dimension of the array could dominate the learning process (e.g. ensuring that the height of the hand had no more influence over map topology than any other dimension in the input vector).

The learning algorithm employed was the batch computation variant of the self-organizing map. In the batch computation algorithm, the entire set of input data is employed during each update step of the nodes of the network. The network is updated over multiple steps (in our case 2000 steps). The alternative approach, an update rule in which the different input vectors are presented separately and in sequence, can produce results that are biased towards the particular inputs that happen to be presented early. The batch computation approach does not suffer from this problem of bias toward early inputs.

A defining feature of the self-organizing map is that as each node of the map updates, it also updates its surrounding nodes. The extent to which nodes influence each other is given by this "neighborhood" function of which several forms can be applied. Following one standard method, we used the Gaussian neighborhood function that takes the form:

$$h_{c,i} = \exp\left(\frac{-d^2(c,i)}{2\sigma^2(t)}\right)$$

where c and i specify two nodes, d(c,i) specifies the physical distance between the two nodes in the array, and σ (t) specifies the time dependent standard deviation of the Gaussian. The time dependence on sigma allows for a large initial neighborhood for the purposes of a global reordering of the network. The neighborhood then shrinks with time to allow for convergence upon a stable solution. In our case, following standard conventions, the initial value of sigma was set to half the network size and fell off linearly for each iteration of the algorithm to a final value of 1.

Often, various weighting schemes are used with self-organizing maps. Certain dimensions are weighted more than others, effectively making them dominant, when it is thought that they should have more effect on network topography. We employed a standard weighting scheme in our network in which the somatotopic components of the codebook vector were updated at a rate 30% slower than the other components, providing for a more stable somatotopic

organization. This weighting avoided situations in which the network, during reorganization, developed unphysiological somatotopies such as a foot representation in a ventral position, or a tongue representation in a dorsal position.

All the methods for a Kohonen self-organizing network can be obtained from Kohonen (2001). In addition, a self-organizing network toolbox with standard programs can be obtained at http://www.cis.hut.fi/projects/somtoolbox/.

Results

Somatotopic Organization

Figure 4-3 shows the results from one learning set. Figure 4-3A shows the initial somatotopic organization in the node array. For consistency with the spatial arrangement found in the monkey brain, in the following description "posterior" refers to nodes on the right of the array, "anterior" to nodes on the left, "dorsal" and "medial" to nodes toward the top, and "ventral" and "lateral" to nodes toward the bottom of the array.

Figure 4-3 B-K show the final somatotopic arrangement after learning. For example, B shows the tongue representation, in which "hot" colors show the nodes that most strongly represent the tongue and "colder" colors show nodes that represent the tongue less strongly. These somatotopic panels show that the somatotopy became less discrete as a result of learning. The map became extensively overlapping.

The map contained more somatotopic overlap in more anterior areas and contained a more discrete somatotopy in the more posterior areas. Along the posterior edge of the map, a relatively discrete progression can be seen. This progression includes a mouth representation at the bottom, then a region that emphasizes the hand but also weakly represents the arm, then a region that emphasizes the arm but also weakly represents the hand, then a region that represents the foot and leg.

The hand representation became divided into three regions (Figure 4-3 H). One hand representation was located in the posterior part of the array, as if corresponding to the primary motor hand area; the second hand representation was located in an anterior region within the dorsal half of the array, as if corresponding to the dorsal premotor hand area; and the third hand representation was located in an anterior region at the ventral edge of the array, as if corresponding to the ventral premotor hand area. (Compare to Figure 4-1 F, based on results from the monkey brain.)

Clustering of Ethological Function

Figure 4-3 L shows the clustered arrangement of ethological movement categories after learning. Central-space/manipulation movements were represented in a posterior region, overlapping the posterior hand representation. Reaching movements were represented in an anterior dorsal region, overlapping the anterior-dorsal hand representation. Hand-to-mouth movements were represented in an anterior ventral location, overlapping the anterior ventral

hand representation. Defensive movements were represented in an anterior region between the hand-to-mouth representation and the reaching representation. Climbing movements were represented dorsal to the reaching movements, in the most dorsal and anterior corner of the map. (Compare to Figure 4-1 I, based on results from the monkey brain.)

The particular pattern shown in Figure 4-3 L, in which the different ethological categories are organized into a specific spatial arrangement, results from a bias provided by the initial somatopic map. For example, the hand-to-mouth movements included the tongue, lips, jaw, neck, hand, and arm; the reaching movements included the arm, hand, and torso; given the initial somatotopic locations for these body parts, the reaching movements inevitably came to be represented in a location dorsal to the hand-to-mouth movements. Similar initial somatotopic biases explain the relative location for the other ethological categories.

Final Hand Positions

Figure 4-3 M-O show the final topographic arrangement of hand position. In these maps, certain array locations are blanked out because the arm was not represented at these sites and therefore no hand location could be defined. The clearest organization was obtained for hand height. Upper hand positions were represented mainly in a ventral anterior area that included the hand-to-mouth zone; mid-height locations were represented in an intermediate location that included the central-space-manipulation zone; and lower hand locations were represented in a

posterior dorsal area, completing the topographic progression. In addition to this map of hand height, there was also a second representation of upper hand locations in the anterior dorsal region of the array, corresponding to the climbing zone. (Compare to Figure 4-1G, based on results from the monkey brain.)

There was also some degree of mapping of the lateral-to-medial location of the hand, but with a reversal. Hand locations near the vertical meridian tended to be represented in posterior parts of the array, especially those corresponding to the central-space-manipulation zone. Lateral hand locations tended to be represented in a more anterior strip of the array. In the reaching zone, in a dorsal anterior region, the map reversed and hand locations again tended to fall near the vertical meridian, representing hand locations distant from the body but near the midline. (Compare to Figure 4-1H, based on results from the monkey brain.)

There was no clear systematic representation of the z location of the hand (distance forward of the body along the line of sight), though in accordance with the tendency of a Kohonen map, the array showed a clumping in which similar z values were represented in neighboring nodes.

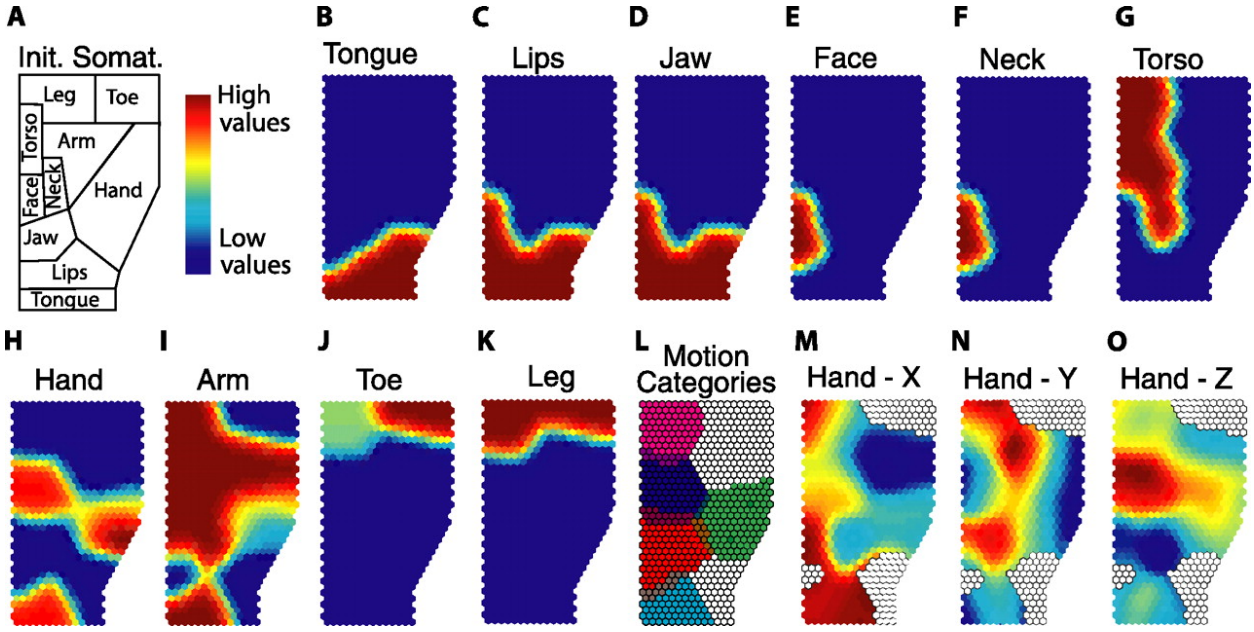


Figure 4-3 : Final state of the self-organizing map model. In this learning set, the different movement types were presented to the network in equal proportion. A. The initial somatotopic arrangement of body parts before reorganization. B-K. Somatotopic arrangement of the ten body parts after reorganization. L. Arrangement of the 5 ethological categories after reorganization. In this panel, light blue = hand-to-mouth; dark blue = reach; red = defense; green = central space/manipulation; pink = climbing. Regions of overlap have intermediate colors. M-O. Maps of hand location after reorganization. X= hand height, warm colors = greater height; Y= lateral location of hand, warm colors equal more lateral locations; Z=distance of hand from body along line of sight; warm colors = more distant locations.

Alternate Maps

One possible concern with the result obtained in Figure 4-3 is that the model was trained on a movement set in which different movement categories were equally represented. The input array contained an equal proportion of all categories of movement. Monkeys do not perform these movements in equal proportions. For example, manipulation of objects in central space is a relatively high-frequency behavior, while reaching outward to grasp makes up a smaller

proportion of the animal's spontaneous behavior (Graziano et al. 2004). Will the final configuration of the map depend on the exact proportions of these different movements in the training set? Figure 4-4 shows the results for a training set in which the different movement categories were represented in unequal proportions based on the observations of monkey's spontaneous behavior in the home cage (see figure caption for details). The results are nearly identical. The sizes of the different representations are slightly altered. For example the region representing central-space/manipulation is slightly larger in Figure 4-4 than in Figure 4-3, reflecting the larger proportion of that category of movement in the learning set. The overall pattern of results, however, is similar.

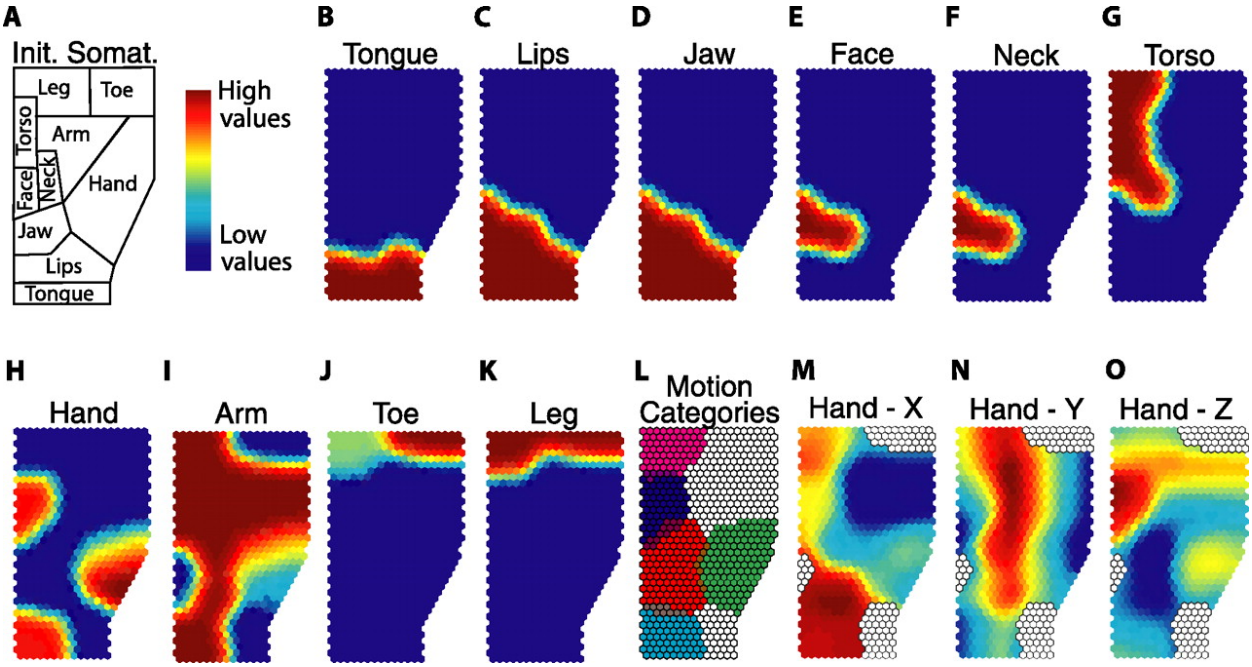


Figure 4-4 : Alternate version of the self-organizing map model. In this learning set, different movement types were presented to the network in unequal proportions. The central-space/manipulation movements were increased in proportion by 50%; the reaching movements were decreased by 20%; the hand-in-lower space movements were increased by 50%; the chewing movements and the leg and foot movements were decreased by 20%. A similar result was obtained here as in Figure 4-3.

To further probe the robustness of the final map, we tried other variants of the learning set and variants of the definitions of the movements. Figure 4-5 shows the arrangement of the ethological movement categories obtained from some of these variants (see figure caption for details). The dimensionality reduction arrived at the same overall pattern. The final map is therefore highly constrained. If topographic continuity is to be optimized, and if the map is to place importance on somatotopic organization, hand location, and ethological category, then the final arrangement that compromises among those different influences is essentially determined.

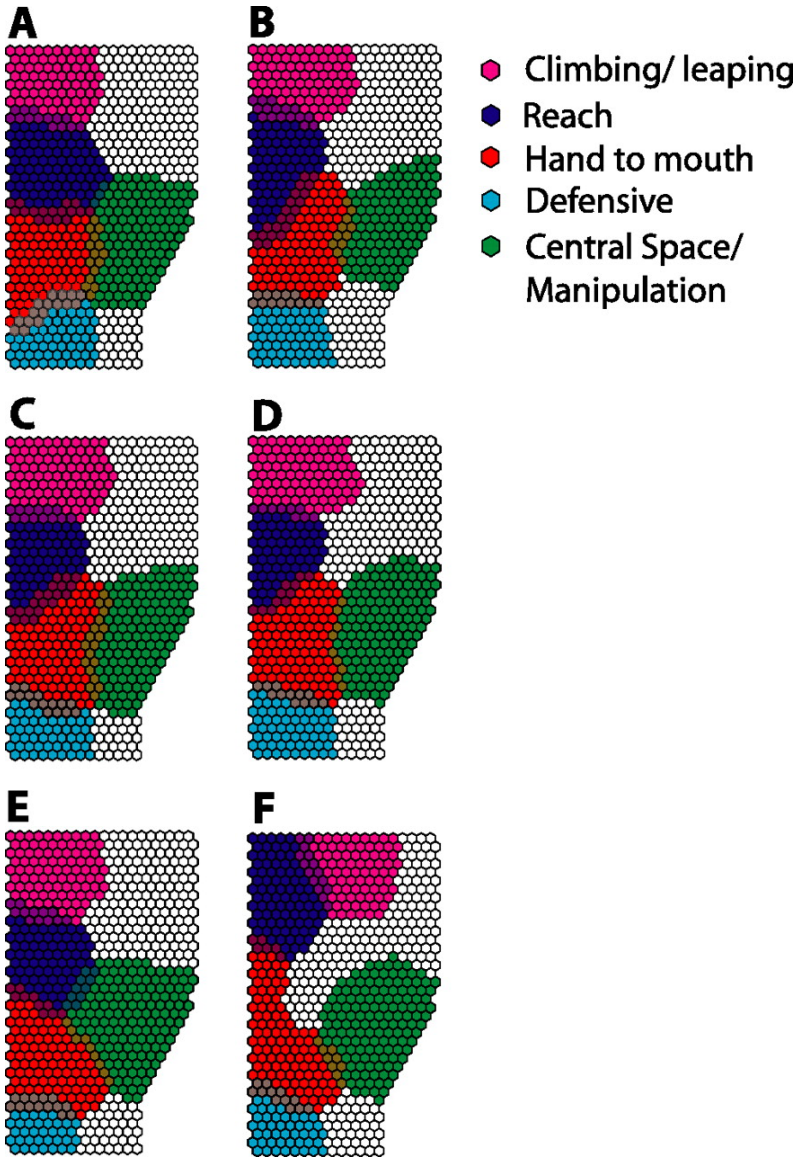


Figure 4-5 : Six variants of the self-organizing map model. Despite the variations in initial state, the results were similar. In each case, the model was the same as shown in Figure 4-3 except as follows: A. Mean hand location for each of the five movement types was shifted in a random direction by 5 cm. B. The standard deviation of hand location for each of the five movement types was doubled. C-F. The proportions of the movement types in the input array were altered randomly by 20%. In one variant (F), the climbing zone moved to a more posterior region and the reaching zone moved to the dorsal anterior corner of the array. This variant was therefore similar but not identical to the others. This result shows that changing the proportions of the movement types can cause small variations in the final results, but the essential topography remains similar.

Lesions and use-dependant changes

Previous studies have shown that the topography in motor cortex can change with experience or with lesions to parts of motor cortex (e.g. Classen et al. 1998, Karni et al. 1995, Kleim, Barbay, and Nudo 1998, Nudo and Milliken 1996, Nudo et al. 1996, Pascual-Leone et al. 1995). Can similar changes be observed in the present model of motor cortex?

Use-dependant changes in hand representation. The results described above suggest that when the model is trained on a movement set with a larger proportion of manipulation movements, the corresponding representation in the map becomes enlarged. Figure 4-6 examines this effect more systematically. We first trained the model as before on the movement set used in Figure 4-3. Once the model settled on its final state, we then allowed the model to continue learning on a second movement set. Figure 4-6A shows the result when the second movement set was the same as the first one. There is no change in the size of the manipulation zone (compare to Figure 4-3 L). Figure 4-6B shows the result when the second movement set included a 50% increase in the proportion of central-space/manipulation movements. The corresponding representation in the map expanded. Figure 4-6C shows the result when the second movement set included a 50% decrease in the proportion of central-space/manipulation movements. The corresponding representation in the map shrank.

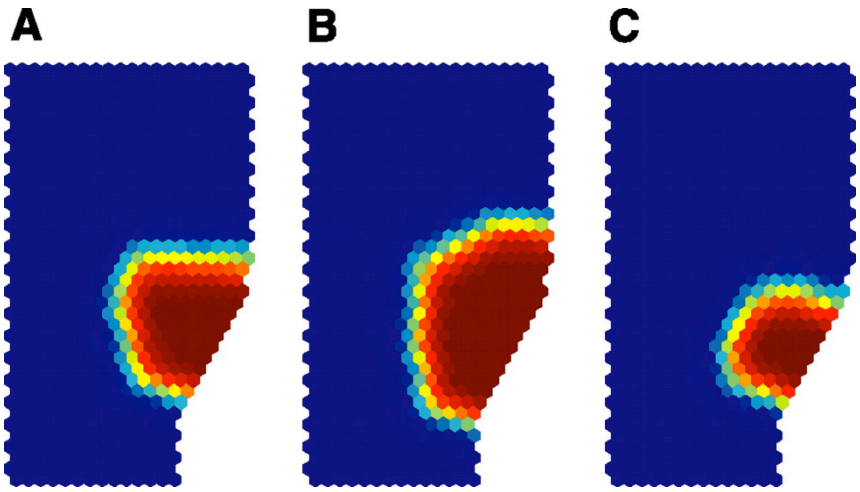


Figure 4-6 : Usage-dependant changes in the map. The model was trained on the same movement set used in Figure 4-3, then trained again on a second movement set. A. Final state of map model when the second movement set was the same as the first. The representation of one ethological movement category, central-space/manipulation, is shown. B. Final state of the map model when the second movement set was the same as the first except that the number of central-space/manipulation movements was increased by 50%. C. Final state of the map model when the second movement set was the same as the first except that the number of central-space/manipulation movements was decreased by 50%.

"Lesion" to the hand representation. We first let the model self-organize as in Figure 4-3. Once the model settled on a final map, we blanked out a set of nodes in the hand representation. These blanked-out nodes were no longer functional parts of the array. Figure 4-7A shows the hand representation and the lesion (black area) before any further learning had occurred. Figure 4-7B shows a similar plot for the arm representation. We then allowed the model to continue learning on the same standard movement set on which it had originally been trained. Figure 4-7C shows the re-organized hand representation. Essentially the residual hand

representation around the edge of the lesion expanded to form a complete hand representation in an adjacent location. This re-organization of the hand representation also caused a re-organization in other parts of the map. This shifting of the rest of the map is illustrated in Figure 4-7D, in which the arm representation receded to allow space for the expansion of the hand representation.

One caveat is important to note here. The Kohonen network seeks an optimal map configuration. When map conditions are changed, such as when the usage of one type of movement is increased, or when a piece of the map is "lesioned," then the Kohonen algorithm seeks on a new optimal map configuration. In this sense, the present model shows use-dependant and lesion-dependant changes. However, the process by which the Kohonen network finds the optimal map configuration is a purely computational one and is presumably different from the mechanisms of plasticity in the brain. The finding that the model settles on a similar final state as obtained in physiological studies of the brain suggests that the plastic mechanisms in the brain, whatever they may be, may also tend to seek continuity in mapping multiple parameters onto a two-dimensional sheet.

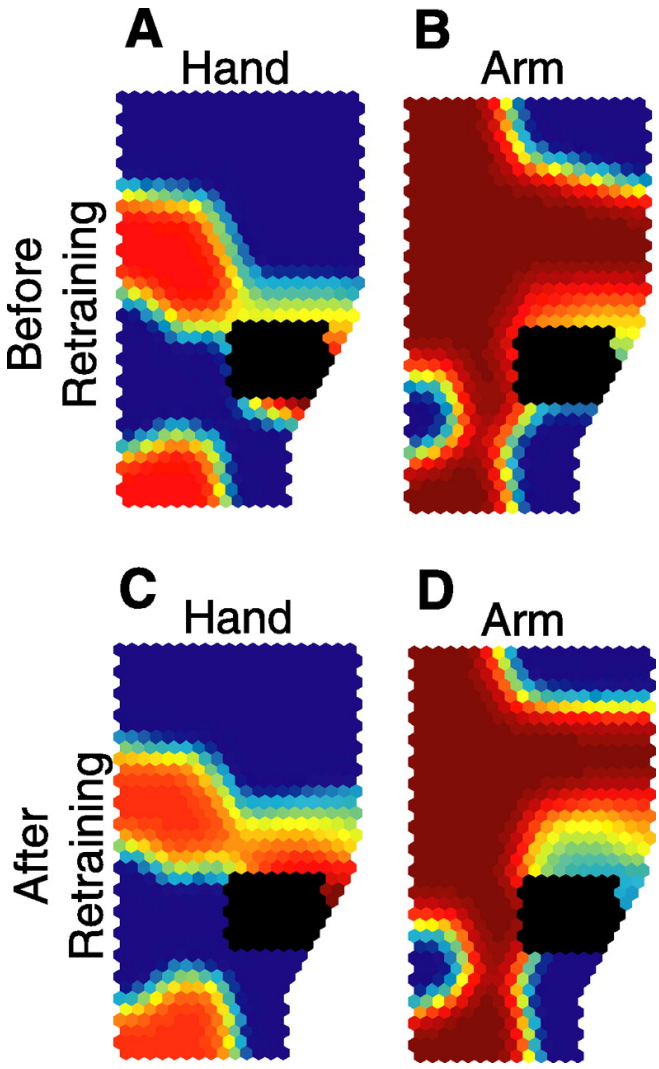


Figure 4-7 : Effect of a "lesion" on the map model. A. Final state of the hand representation using parameters and learning set as in Figure 4-3, with black area showing the site of the lesion to be made. B. Same as A but for arm representation. C. State of hand representation after lesion and further learning on the same movement set. The lesioned hand representation expanded. D. State of arm representation after lesion and further learning on the same movement set. The arm representation shifted to allow room for the expanded hand representation.

Discussion

The motor cortex map generated by our self-organizing network resembled the actual maps obtained from the lateral motor cortex of monkeys in five specific ways, listed below.

Simpler movements and clearer somatotopy in a posterior "primary-like" strip.

Traditionally the posterior portion of motor cortex, "primary motor" cortex, is thought to contain a relatively more discrete somatotopy in which the major body parts are separated, whereas the more anterior strip, considered the "premotor" cortex, is thought to contain a more overlapping somatotopy and to control more complex movements (e.g. Campbell 1905, Fulton 1938, Rizzolatti and Luppino 2001). Both strips project directly to the spinal cord (Bortoff and Strick 1993, Dum and Strick 1996, 2002b, 2005b, He, Dum, and Strick 1993, Wu, Bichot, and Kaas 2000), calling into question whether they form a true hierarchy or simply emphasize different types of movement.

The map generated by our self-organizing network contained similar topographic trends. In the posterior half of the array, there was a somatotopic sequence of four major representations including the tongue, lips and jaw (Figure 4-4B-D); the hand (Figure 4-3 H); the arm (Figure 4-3 I); and the leg and foot (Figure 4-3 J, K). In the anterior regions, the overlap among body segments was more extensive and the somatotopy more fractured. The reason for this trend in

the self-organizing map is clear: Some of the movements in our model were especially complex in that they required coordination among different body parts. For example, reaching involved action of the hand, arm, and torso. Defensive movements involved action of the face, rotation of the head, movement of the torso, and movement of the arm. Climbing was perhaps the most integrated, involving the arm, torso, and leg. These movements involved the axial musculature, because the trunk and neck form the connecting links between different body segments. The initial somatotopy was arranged with the axial musculature in an anterior region and the distal musculature in a posterior region. As a result, during learning, the complex movements that link more than one body segment gravitated to the anterior regions of the map. Thus in our model, in its final state, one can distinguish a posterior strip that is "primary-like" in that it contains a relatively discrete somatotopy, representing body segments in a separate manner; and an anterior strip that is "premotor-like" in that it contains a more integrated, blurred somatotopy and represents movements of greater inter-segment complexity. However, no processing hierarchy is implied by the map. The anterior strip does not control the posterior strip.

Cortical zones for ethologically relevant movement types.

Our previous stimulation experiments indicated that the precentral gyrus contains a mosaic of subregions or hot spots that emphasize different movement categories (Cooke and Graziano 2003, Graziano, Aflalo, and Cooke 2005, Graziano, Patel, and Taylor 2004, Graziano, Taylor, and Moore 2002). Stimulation in each subregion tended to evoke a specific type of movement recognizable from the monkey's behavioral repertoire. These five movement

categories, obtained empirically on stimulation, were used to train our self-organizing map. During learning, the different movement types became clumped into distinct subregions of the array. Remarkably, the topographic arrangement of the subregions in the self-organizing map closely resembled the arrangement observed in the actual monkey brain. This arrangement of ethological zones is constrained by the initial somatotopy and the subsequent attempt of the network to optimize nearest neighbor relationships.

Three hand representations.

Three main hand representations have been described in the lateral precentral gyrus of monkeys; one in a posterior area in primary motor cortex; one in dorsal premotor cortex (probably in the caudal division, or PMDc); and one in ventral premotor cortex. The three hand areas were defined by their projections to the hand motoneuron pools in the spinal cord (Dum and Strick 2002b, 2005b, He, Dum, and Strick 1993). Our self-organizing map also converged on three hand representations (Figure 4-4H). The reason why the network arrived at three hand representations is that the hand participated in three distinct movement categories. One hand representation was on the posterior edge of the array, emphasizing the distal movements used during manipulation of objects. The second hand area was in the dorsal anterior portion of the array, emphasizing the control of reaching to grasp. The third hand area was in the ventral anterior portion of the array, emphasizing interactions between the hand and the mouth.

Noisy maps of hand location.

We previously found that stimulation of different sites in the monkey motor cortex drove the hand to different locations in space, and that the evoked hand locations formed a map across the lateral motor strip, encompassing the primary motor and caudal premotor areas (Graziano, Aflalo, and Cooke 2005, Graziano, Taylor, and Moore 2002). The map contained considerable overlap. The height of the hand was most consistently mapped; the lateral position of the hand was less clearly ordered in the map; and the forward distance of the hand along the line of sight showed no consistent topography. Our self-organizing map converged on a topographic arrangement of hand position that roughly resembled this arrangement seen in monkeys.

Learning-dependant reorganization.

Our self-organizing network began with a discrete, well-defined somatotopy and then became reorganized during training on a set of complex movements. Data from motor cortex in humans, monkeys, cats and rats suggest that a similar process occurs during development and learning, in which an initial somatotopy reorganizes around the types of movement that are common in the animal's repertoire. In the cat, at an early stage in development the motor cortex contains a discrete somatotopic map in which joints have separate representation, but during practice with more complex reaching movements the map develops overlapping representations of functionally related joints (Martin, Engber, and Meng 2005). In humans, rats, and monkeys, practice on a motor task can cause changes in the size and degree of overlap of the cortical representations of the relevant body parts (e.g. Classen et al. 1998, Karni et al. 1995, Kleim,

Barbay, and Nudo 1998, Nudo and Milliken 1996, Nudo et al. 1996, Pascual-Leone et al. 1995). In humans, practice on a task that combines movements of more than one body part can cause the representations of those body parts to move toward each other (Liepert, Terborg, and Weiller 1999). These and other studies suggest that motor cortex topography is plastic and constantly reorganizing, and movement components that normally occur together during behavior can develop representations that are near each other or overlapping in cortex. This topographic reorganization has been suggested to be a physical manifestation of the "motor engram," where memory for complex motor skills is stored (Monfils, Plautz, and Kleim 2005). We suggest that it is precisely these motor engrams that are triggered by electrical stimulation. Complex, behaviorally relevant movements, electrically evoked from specific sites in motor-related cortical areas, have now been observed in monkeys, rats, cats, and prosimians (Ethier, Brizzi, and Darling 2006, Haiss and Schwarz 2005, Stepniewska, Fang, and Kaas 2009).

Differences between the model and the monkey data.

In the monkey, the hand representation is reported to be surrounded in a semicircle by a forearm and upper arm representation (Figure 4-1C) (Kwan et al. 1978, Meier et al. 2008, Park et al. 2001). In our model, the arm representation also tended to wrap partially around the hand representation (Figure 4-4H,I). This wrap-around was less complete in our model than in the descriptions from the monkey. One possible reason is that the model contained only a crude somatotopy including a hand and an arm, but not the finer gradations of wrist, forearm, elbow,

and shoulder. Thus the wrapping of the wrist and forearm representation around the hand was not possible here.

A second difference between the model and the data is that the model resulted in nearly identical maps on different runs, despite changes in the proportions of different movements, yet the maps in the monkey appear to be variable among individuals. Some possible reasons for the greater variation in monkeys are: the exact size and shape of the relevant motor strip may be different in different individuals, leading to different proportions internal to the strip; and the actual motor map may be influenced by many more factors than the three modeled here.

Despite these differences, by using only three main mapping factors in competition with each other the model was able to produce a topography that captures some of the essential features of the monkey motor cortex. We suggest that the organization in motor cortex cannot be understood as a single map of one kind or another, but rather can be partially understood as a result of an initial, underlying somatotopy that is re-organized under the competing influences of other mapping requisites, including a cluster map of ethologically relevant movements and a topographic map of hand position in space. Other variables relevant to the animal's movement repertoire may also influence the topography within motor cortex in ways not tested in the present model.

Chapter 5 : Concluding Comments

In this thesis I described three studies inspired by the finding that long-duration stimulation in the monkey motor cortex evokes complex coordinated movements that can resemble actions from the monkey's behavioral repertoire. In the first study we showed that movements evoked by long-duration stimulation can match features of natural voluntary movements at the level of kinematic trajectories and the ability to stabilize the final position of the hand against externally imposed and internal variability. These movements were topographically organized on the cortical surface by the set of body parts recruited, the final position of the hand in space, and the type of behavioral action evoked. In the second study we found that the behavior of single neurons during unstructured naturalistic movements was best explained by tuning to a preferred posture. In many cases, the preferred posture of a neuron and the posture evoked by electrical stimulation at the same cortical site were in rough agreement. Neurons were also tuned to other aspects of movement (e.g. direction of hand motion in space) as observed in previous studies, however, these alternate tuning models were generally poor at explaining the behavior of neurons over the full naturalistic dataset. Instead, these models performed best when the full dataset was limited to only include movements in a restricted part of the animals' workspace. In the third study we showed that many features of the complex topographic organization of the lateral motor strip emerge as the natural consequence of representing a high-dimensional description of the monkey's behavioral repertoire onto a two-dimensional cortical surface subject

to the constraint of local continuity of representation.

Feedforward versus feedback control. Is motor cortex fundamentally a feed-forward system producing movement by preplanning the exact spatiotemporal profile of muscle activations necessary to produce a desired movement or does the motor system use a feedback control strategy? A persisting view is that the motor system is organized in a heriarchical feed-forward manner with premotor regions instructing motor cortex and motor cortex controlling muscles. Long-duration stimulation drives the arm from any intitial state to a final invariant posture (Graziano, Taylor, and Moore 2002). Muscle activation patterns resulting from stimulation are remapped by sensory feedback in such a way as to drive the arm to this final posture (Graziano, Patel, and Taylor 2004). These results suggest that motor commands are generated from the dynamic interplay between cortical activity and sensory feedback challenging classic descriptions of feed-forward motor control. In chapter 2 we demonstrated that movements evoked by stimulation can show compensation for a weight added to the arm with the hand achieving the same final position as the unweighted condition. The movements evoked by stimulation showed inter-joint coordination such that the final position of the hand was relatively stable despite trial-to-trial variability for each individual joint. These results provide evidence that feedback mechanisms evident in the sensory dependent remapping of muscle activation can manifest in meaningful compensation at the level of behavioral output.

Postural coding by primary motor cortex. As described in Chapter 3, many single neurons recorded in primary motor cortex were tuned to posture but they were also tuned to other motor variables, especially when analyzed in a restricted region of space. Neural representations of movement are labile and, through experience, come to reflect aspects of movements that have proved especially useful in the animal's previous motor experience (e.g. Huber et al. 2012, Kleim, Barbay, and Nudo 1998, Martin, Engber, and Meng 2005, Nudo et al. 1996). Many of the actions performed by monkeys (and humans) require placing the arm, including both proximal and distal joints, in a canonical posture and making smaller adjustments to that posture. For example, manipulating an object between the hand and the teeth involves a canonical arm posture including the shoulder in a low elevation, the elbow flexed, the forearm supinated, the wrist slightly flexed, and the grip closed. Within this canonical posture, smaller adjustments are made to each of these joints to manipulate the object at the mouth. An example from the human motor repertoire is typing: the arm remains in a canonical posture with the upper arm vertically downward, the elbow flexed to approximately 90 degrees, the forearm pronated, and the wrist straight. Superimposed on this relatively stable posture is a set of smaller directional adjustments. Arguably in this type of behavior the proximal joints require more stabilization than the distal joints, but all joints whether proximal or distal participate in the canonical posture and also participate in the finer directional adjustments around that posture needed to perform the action. We speculate that so many actions in the primate repertoire have this architecture of a global posture and local directional adjustments to that posture that neurons in motor cortex may

have come to reflect that structure in their movement tuning.

Much of the literature in motor control asks how the hand is moved from point A to point B. Is it moved via a postural controller (Feldman 1986, Feldman and Levin 2009, Rosenbaum et al. 1995) or a directional controller? We suggest, however, that almost all normal behavior requires a much greater range of objectives than merely a cursor-like movement of the hand from point to point. Control of posture is not merely a means of moving the hand around, but can be a prominent part of the behavioral task. The arm configuration provides the correct hand orientation, stability along the correct axis required for the task, a maximum of comfort, and prevention of different parts of the arm from colliding with the body or with nearby objects. For example, bringing the hand to the mouth is not merely a matter of translating the hand, cursor-like, to the mouth; it involves a characteristic whole posture of the arm. Moreover, in some types of behavior, such as nudging something with the elbow, or flinging up the arm into a protective posture to block a threat, the position of the hand in space is relatively unimportant and the posture of the arm becomes the primary task requirement. For these reasons it does not seem surprising that motor cortex should contain such robust signals concerning the configuration of the arm, or that postural signals and directional signals should be combined as they are in the behavior.

Is postural control the basis of all motor actions? The equilibrium point hypothesis (EPH), an influential and controversial model in motor control, proposes that motor commands result from specifying desired postures and allowing the discrepancy between the current and

desired posture to drive motor output (Feldman 1986, Feldman and Levin 2009). For instance, the EPH proposes that in order to produce a precise force against an object, the motor system plans a "virtual" posture that penetrates the object (Pilon, De Serres, and Feldman 2007). The further into the object the virtual posture is specified, the greater the force. In our view, the finding that posture is a prominent part of the cortical representation of movement is a consequence of its importance in the behavioral repertoire. However, the motor behavior of primates is diverse and different neural representations may be appropriate for different behaviors. For instance, behaviors like walking and running rely on the motor system producing internally driven rhythmic oscillations. Control of grasp requires force and tactile feedback to maintain appropriate grip strength. Force control may involve specialized circuits that directly regulate force based on force and tactile sensory feedback or it may involve an intermediate stage involving a postural controller. Given the complexities involved in teasing out these alternatives (e.g. Ostry and Feldman 2003) our data does not directly adjudicate between these possibilities.

Population coding in primary motor cortex. Neurons were often tuned to a subset of the full number of measured joint angles and the tuning to posture was often broad; the neuron fired most for its preferred posture but also fired at an elevated level for neighboring postures. How might this fragmentary representation be utilized during the control of natural behavior? Perhaps this question can best be understood in relation to the population coding mechanism proposed by

Georgopoulos and colleagues (1986). The tuning of individual neurons to direction of motion is broad (Georgopoulos et al. 1982). Each neuron fires most for its preferred direction and progressively less as the direction of movement deviates from the preferred direction. How can such broadly tuned neurons move the hand in a specific direction? Georgopoulos et al. (1986) represented each neuron as a vector whose direction was determined by the neuron's preferred direction and whose amplitude was determined by the neuron's firing rate. When the vectors of a population of neurons were added together, the resulting vector, the "population vector", correctly predicted the direction of movement of the monkey's hand. They concluded that movements of the arm were specified by the roughly linear combination of a large population of directionally selective neurons. The underlying principal is that each neuron nudges the hand in a specific direction and only through the combined action of many neurons can coherent behavior unfold. This general principal can be extended to other forms of cortical representation. As opposed to nudging the hand in a specific direction, our results suggest that each neuron nudges the system to produce a fragment of behavior that is partly characterized by movement to a posture. This neural representation is presumably learned and reflects statistical biases in the animal's previous motor experience that is largely characterized by performing small movements about a relatively fixed posture. Can the population coding method extend to more idiosyncratic learned movement representations? Ethier et al. (2006) used electrical stimulation to evoke complex multi-joint movements from the cat motor cortex. When they stimulated two cortical sites simultaneously they were also able to evoke complex multi-joint movements. The

movement trajectories and muscle activation profiles when stimulating two cortical sites could be predicted from the linear combination of the results of stimulating each site in isolation. The multi-joint movements evoked by Ethier et al. presumably emerged as a result of the cat's previous motor experience (Martin, Engber, and Meng 2005). These results suggest that fragments of behavior, gained through experience, can be combined in a relatively straightforward way to produce novel movements. In its basic form, this population coding scheme can be likened to that proposed by Leyton and Sherrington (1917) nearly 100 years ago. Each neuron is tuned to a fragment of a more complex behavioral action. These fragments, or what Sherrington called elementary movements, can then be combined to produce meaningful actions. Our results suggest that a component of the elementary movement involves movement toward a desired posture (or postural subset).

The behavioral organization of the frontal motor areas. In chapter 4 we modeled the topographic organization of motor cortex by projecting a high-dimensional representation of movements in the monkey's behavioral repertoire onto a two-dimensional cortical surface subject to the constraint of local continuity of representation. The cerebral cortex is organized such that similar functions are generally spatially near each other. A possible adaptive advantage of this organizational feature is that it minimizes wiring length between interconnected neurons thus lowering metabolic demands and preserving limited cranial space (Cherniak 1994, Kaas 2000, Klyachko and Stevens 2003, Vanessen 1997, Young 1992). Whatever the underlying

cause, a principle of maximum functional continuity can explain features of cortical organization at the largest and smallest scales. At the largest scale, function is clumped into cortical sectors devoted mainly to vision, audition, movement, and other major information categories. At a slightly smaller scale, within the visual modality, cortical areas may be clustered together by similar function (Kolster, Mandeville, and Arsenault 2009, Wandell, Dumoulin, and Brewer 2007). At an even smaller scale, within a cortical area, the principle of functional continuity can explain topography such as retinotopic or somatotopic maps (Kaas and Catania 2002, Kohonen 1982, Saarinen and Kohonen 1985). When the functional properties of a cortical region are known, computational modeling methods that attempt to represent the high-dimensional functional space onto the two-dimensional cortical surface while maximizing functional continuity are able to recreate many features of the observed topographic organization of the cortical region. For example, Durbin and Michison (1990) were able to recreate the basic pin-wheel topographic organization of V1 by mapping a four-dimensional functional space, characterized by retinotopic location and orientation, onto a 2 dimensional model of cortex. This suggests that if the functional properties of a cortical area are known, rendering the functional description onto the cortical surface should reproduce features of the actual topographic organization of the cortical region. Our model was able to capture many features of the observed organization of the monkey motor cortex and thus helps support the idea that motor cortex is in many ways a map of the behavioral repertoire.

It is generally accepted that the frontal motor areas are divided into functionally and

anatomically distinct regions (e.g. primary motor cortex and the multiple premotor regions), however, the underlying reason for these subdivisions has remained ellusive. Each cortical field is fairly heterogeneous in its representations and the differences between regions are more graded than absolute. These findings do not support either a clear feed-forward hierarchical organization or the idea that each region performs some specific computation as part of a single unified system for the control of the body. One possibility is that the different properties of each region reflect the different types of sensory information, sensory-motor transformations, and motor output variables necessary to perform particular behavioral actions (Graziano 2009, Graziano 2006). For example, when manipulating an object, e.g. tearing the skin off a piece of fruit, the control system must correctly apply forces to the object to e.g. remove the skin while not crushing the fruit. Electrical stimulation in the "hand region" of motor cortex evokes the canonical posture associated with manipulation of objects in addition to a shaping of the hand (Graziano et al. 2004, Graziano, Taylor, and Moore 2002). A large cortical representation of the distal musculature of the hand, strong proprioceptive and cutaneous feedback, direct projections to alpha motor neurons and partial coding of force and muscle activation may reflect the properties especially important for dexterous manipulation. Recent stimulation results suggest that posterior parietal cortex is similarly arranged into subregions emphasizing different complex behavioral actions (Kaas, Gharbawie, and Stepniewska 2011, Stepniewska, Fang, and Kaas 2005, Stepniewska, Fang, and Kaas 2009) suggesting a general principal whereby cortical regions responsible for motor planning and execution are best described as a series of subregions

emphasizing control of different types of behavioral actions.

After 130 years it now appears that the classically described hierarchically ordered, somatotopic organization of motor cortex is better described as a map of the complex, behavioral repertoire. In motor cortex, neural representations of movement are multiplexed, tuned to postural subsets in addition to other aspects of movement. Presumably, these learned fragmentary movement representations tend to organize themselves such that movement representations, which are especially useful for a particular type of behavior, tend to cluster near each other on the cortical surface. Despite advances in understanding the ethological organization of motor cortex, many questions remain. Are the neural pathways responsible for transforming the electrical stimulation of cortex into coordinated movements mediated through transcortical pathways or are they mediated by cortex's influence on subcortical and spinal centers? What factors determine the formation of these regions of behavioral specialization? To what degree are they present at birth and to what degree are they learned through experience? Is the human motor cortex organized into regions emphasizing different types of behavioral actions? How are the fragmentary movement representations learned? How are the movement representations within a behavioral region coordinated to produce meaningful behavior? How do different regions of behavioral specialization coordinate with each other?

Chapter 6 : References

Ajemian R, Green A, Bullock D, Sergio L, Kalaska JF, and Grossberg S (2008) Assessing the Function of Motor Cortex: Single-Neuron Models of How Neural Response Is Modulated by Limb Biomechanics. *Neuron* 58, 414-428.

Armstrong DM and Drew T (1985a) Electromyographic Responses Evoked in Muscles of the Forelimb by Intracortical Stimulation in the Cat. *J Physiol* 367, 309-326.

Armstrong DM and Drew T (1985b) Forelimb Electromyographic Responses to Motor Cortex Stimulation During Locomotion in the Cat. *J Physiol* 367, 327-351.

Asanuma H and Arnold AP (1975) Noxious Effects of Excessive Currents Used for Intracortical Microstimulation. *Brain Res* 96, 103-107.

Ashe J and Georgopoulos AP (1994) Movement Parameters and Neural Activity in Motor Cortex and Area 5. *Cereb Cortex* 4, 590-600.

Bagesteiro LB and Sainburg RL (2002) Handedness: Dominant Arm Advantages in Control of Limb Dynamics. *J Neurophysiol* 88, 2408-2421.

Bates DM and Watts DG. (2007) *Nonlinear Regression Analysis and Its Applications*. Wiley-Interscience.

Battaglia-Mayer A, Ferraina S, Genovesio A, Marconi B, Squatrito S, Molinari M, Lacquaniti F, and Caminiti R (2001) Eye–Hand Coordination During Reaching. Ii. An Analysis of the Relationships between Visuomanual Signals in Parietal Cortex and Parieto-Frontal Association Projections. *Cereb Cortex* 11, 528-544.

Beevor CE (1890) An Experimental Investigation into the Arrangement of the Excitable Fibres of the Internal Capsule of the Bonnet Monkey (Macacus Sinicus). *Philosophical Transactions of the Royal Society of London. B* 181, 49-88.

Beevor CE and Horsley V (1887) A Minute Analysis (Experimental) of the Various Movements Produced by Stimulating in the Monkey Different Regions of the Cortical Centre for the Upper Limb, as Defined by Professor Ferrier. *Philosophical Transactions of the Royal Society B* 178, 153-167.

Berniker M, Jarc A, Bizzi E, and Tresch MC (2009) Simplified and Effective Motor Control Based on Muscle Synergies to Exploit Musculoskeletal Dynamics. *Proc Natl Acad Sci* 106, 7601-7606.

Bizzi E, Cheung VCK, D&Apos;Avella A, Saltiel P, and Tresch M (2008) Combining Modules for Movement. *Brain Research Reviews* 57, 125-133.

Bizzi E, Hogan N, Mussa-Ivaldi FA, and Giszter S (1992) Does the Nervous System Use Equilibrium-Point Control to Guide Single and Multiple Joint Movements? *Behav Brain Sci* 15, 603-613.

Bizzi E and Mussa-Ivaldi FA. (1989). *Handbook of Neuropsychology, Vol. 2*. edited by Boller E and Grafman J.

Bortoff GA and Strick PL (1993) Corticospinal Terminations in Two New-World Primates: Further Evidence That Corticomotoneuronal Connections Provide Part of the Neural Substrate for Manual Dexterity. *J Neurosci* 13, 5105-5118.

Boudreau MJ, Brochier T, Paré M, and Smith AM (2001) Activity in Ventral and Dorsal Premotor Cortex in Response to Predictable Force-Pulse Perturbations in a Precision Grip Task. *J Neurophysiol* 86, 1067-1078.

Brinkman C (1984) Supplementary Motor Area of the Monkey's Cerebral Cortex: Short- and Long-Term Deficits after Unilateral Ablation and the Effects of Subsequent Callosal Section. *J Neurosci* 4, 918-929.

Bruce CJ, Goldberg ME, Bushnell MC, and Stanton GB (1985) Primate Frontal Eye Fields. Ii. Physiological and Anatomical Correlates of Electrically Evoked Eye Movements. *J Neurophysiol* 54, 714-734.

Bucy PC (1933) Electrical Excitability and Cyto-Architecture of the Premotor Cortex in Monkeys. *Archives of Neurology & Psychiatry*.

Bullock D and Contreras-Vidal JL (1993) How Spinal Neural Networks Reduce Discrepancies between Motor Intention and Motor Realization. *Variability and Motor Control*, 183-221.

Caminiti R, Ferraina S, and Mayer AB (1998) Visuomotor Transformations: Early Cortical Mechanisms of Reaching. *Curr Opin Neurobiol* 8, 753-761.

Caminiti R, Johnson P, Galli C, Ferraina S, and Burnod Y (1991) Making Arm Movements within Different Parts of Space: The Premotor and Motor Cortical Representation of a Coordinate System for Reaching to Visual Targets. *J Neurosci* 11, 1182-1197.

Caminiti R, Johnson PB, and Urbano A (1990) Making Arm Movements within Different Parts of Space: Dynamic Aspects in the Primate Motor Cortex. *J Neurosci* 10, 2039-2058.

Campbell A. (1905) *Histological Studies on the Localization of Cerebral Function*. Cambridge, UK: Cambridge University Press.

Chang H, Ruch TC, and Ward AA (1947) Topographical Representation of Muscles in Motor Cortex of Monkeys. *J Neurophysiol* 10, 39-56.

Cheney PD and Fetz EE (1985) Comparable Patterns of Muscle Facilitation Evoked by Individual Corticomotoneuronal (Cm) Cells and by Single Intracortical Microstimuli in Primates: Evidence for Functional Groups of Cm Cells. *J Neurophysiol* 53, 786-804.

Cheney PD, Fetz EE, and Palmer SS (1985) Patterns of Facilitation and Suppression of Antagonist Forelimb Muscles from Motor Cortex Sites in the Awake Monkey. *J Neurophysiol* 53, 805-820.

Cherniak C (1994) Component Placement Optimization in the Brain. *J Neurosci* 14, 2418-2427.

Churchland MM and Shenoy KV (2007) Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. *J Neurophysiol* 97, 4235-4257.

Cisek P, Grossberg S, and Bullock D (1998) A Cortico-Spinal Model of Reaching and Proprioception under Multiple Task Constraints. *J Cogn Neurosci* 10, 425-444.

Classen J, Liepert J, Wise SP, Hallett M, and Cohen LG (1998) Rapid Plasticity of Human Cortical Movement Representation Induced by Practice. *J Neurophysiol* 79, 1117-1123.

Cohen J, Cohen P, West SG, and Aiken LS. (2003) *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*. 3rd ed. Mahwah, NJ: Lawrence Erlbaum Associates.

Cooke DF and Graziano MSA (2003) Defensive Movements Evoked by Air Puff in Monkeys. *J Neurophysiol* 90, 3317-3329.

Cooke DF and Graziano MSA (2004a) Sensorimotor Integration in the Precentral Gyrus: Polysensory Neurons and Defensive Movements. *J Neurophysiol* 91, 1648-1660.

Cooke DF and Graziano MSA (2004b) Super-Flinchers and Nerves of Steel: Defensive Movements Altered by Chemical Manipulation of a Cortical Motor Area. *Neuron* 43, 585-593.

Cooke DF, Taylor CSR, Moore T, and Graziano MSA (2003) Complex Movements Evoked by Microstimulation of the Ventral Intraparietal Area. *Proc Natl Acad Sci* 100, 6163-6168.

Denavit J (1955) A Kinematic Notation for Lower-Pair Mechanisms Based on Matrices. *J App Mech* 22, 215-222.

Desmurget M and Prablanc C (1997) Postural Control of Three-Dimensional Prehension Movements. *J Neurophysiol* 77, 452-464.

Donoghue JP, Leibovic S, and Sanes JN (1992) Organization of the Forelimb Area in Squirrel Monkey Motor Cortex: Representation of Digit, Wrist, and Elbow Muscles. *Exp Brain Res* 89, 1-19.

Dum R and Strick PL. (2005a) "Motor Areas in the Frontal Lobe: The Anatomical Substrate for the Central Control of Movement." In *Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions*, edited by Riehle A and Vaadia E. Boca Raton, FL: CRC Press.

Dum RP and Strick PL (1991) The Origin of Corticospinal Projections from the Premotor Areas in the Frontal Lobe. *J Neurosci* 11, 667-689.

Dum RP and Strick PL (1996) Spinal Cord Terminations of the Medial Wall Motor Areas in Macaque Monkeys. *J Neurosci* 16, 6513-6525.

Dum RP and Strick PL (2002a) Motor Areas in the Frontal Lobe of the Primate. *Physiology & Behavior* 77, 677-682.

Dum RP and Strick PL (2002b) Motor Areas in the Frontal Lobe of the Primate. *Physiol Behav* 77, 677-682.

Dum RP and Strick PL (2005b) Frontal Lobe Inputs to the Digit Representations of the Motor Areas on the Lateral Surface of the Hemisphere. *J Neurosci* 25, 1375-1386.

Durbin R and Mitchison G (1990) A Dimension Reduction Framework for Understanding Cortical Maps. *Nature* 343, 644-647.

Ethier C, Brizzi L, and Darling W (2006) Linear Summation of Cat Motor Cortex Outputs. *J Neurosci* 26, 5574-5581.

Evarts EV (1968) Relation of Pyramidal Tract Activity to Force Exerted During Voluntary Movement. *J Neurophysiol* 31, 14-27.

Feldman AG (1986) Once More on the Equilibrium-Point Hypothesis (Lambda Model) for Motor Control. *J Motor Behav* 18, 17-54.

Feldman AG and Levin MF (2009) The Equilibrium-Point Hypothesis--Past, Present and Future. *Progress In Motor Control: A Multidisciplinary Perspective* 629, 699-726.

Ferrier D. "Experiments on the Brain of Monkeys.--No. I." In *Proceedings of the Royal Society of London* 409-430, 1874.

Fetz EE and Cheney PD (1980) Postspike Facilitation of Forelimb Muscle Activity by Primate Corticomotoneuronal Cells. *J Neurophysiol* 44, 751-772.

Flash T and Hogan N (1985) The Coordination of Arm Movements: An Experimentally Confirmed Mathematical Model. *J Neurosci* 5, 1688-1703.

Foerster O (1936) The Motor Cortex in Man in the Light of Hughlings Jackson's Doctrines. *Brain* 59, 135-159.

Fogassi L, Gallese V, Buccino G, Craighero L, Fadiga L, and Rizzolatti G (2001) Cortical Mechanism for the Visual Guidance of Hand Grasping Movements in the Monkey: A Reversible Inactivation Study. *Brain* 124, 571-586.

Fogassi L, Gallese V, Fadiga L, Luppino G, Matelli M, and Rizzolatti G (1996) Coding of Peripersonal Space in Inferior Premotor Cortex (Area F4). *J Neurophysiol* 76, 141-157.

Freedman EG, Stanford TR, and Sparks DL (1996) Combined Eye-Head Gaze Shifts Produced by Electrical Stimulation of the Superior Colliculus in Rhesus Monkeys. *J Neurophysiol* 76, 927-952.

Fritsch G and Hitzig E (1960) Ueber Die Elektrishe Erregarkeit Des Grosshirns (on the Electrical Excitability of the Cerebrum). *Some Papers on the Cerebral Cortex (Translated by G von Bonin)*, 73-96.

Fu QG, Suarez JI, and Ebner TJ (1993) Neuronal Specification of Direction and Distance During Reaching Movements in the Superior Precentral Premotor Area and Primary Motor Cortex of Monkeys. *J Neurophysiol* 70, 2097-2116.

Fulton J (1935) A Note on the Definition of the "Motor" and "Premotor" Areas. *Brain* 58, 311-316.

Fulton J. (1938) *Physiology of the Nervous System*. New York: Oxford University Press.

Gellhorn E (1948) The Influence of Alterations in Posture of the Limbs on Cortically Induced Movements. *Brain* 71, 26-33.

Gellhorn E (1949) Proprioception and the Motor Cortex. *Brain* 72, 35-62.

Gellhorn E and Hyde J (1953) Influence of Proprioception on Map of Cortical Responses. *J Physiol* 122, 371-385.

Gentilucci MM, Fogassi LL, Luppino GG, Matelli MM, Camarda RR, and Rizzolatti GG (1988) Functional Organization of Inferior Area 6 in the Macaque Monkey. I. Somatotopy and the Control of Proximal Movements. *Exp Brain Res* 71, 475-490.

Georgopoulos AP, Ashe J, Smyrnis N, and Taira M (1992) The Motor Cortex and the Coding of Force. *Science* 256, 1692-1695.

Georgopoulos AP, Caminiti R, and Kalaska JF (1984) Static Spatial Effects in Motor Cortex and Area 5: Quantitative Relations in a Two-Dimensional Space. *Exp Brain Res* 54, 446-454.

Georgopoulos AP, Kalaska JF, Caminiti R, and Massey JT (1982) On the Relations between the Direction of Two-Dimensional Arm Movements and Cell Discharge in Primate Motor Cortex. *J Neurosci* 2, 1527-1537.

Georgopoulos AP, Schwartz AB, and Kettner RE (1986) Neuronal Population Coding of Movement Direction. *Science* 233, 1416-1419.

Geyer S, Matelli M, Luppino G, Schleicher A, Jansen Y, Palomero-Gallagher N, and Zilles K (1998) Receptor Autoradiographic Mapping of the Mesial Motor and Premotor Cortex of the Macaque Monkey. *J Comp Neurol* 397, 231-250.

Geyer S, Matelli M, Luppino G, and Zilles K (2000) Functional Neuroanatomy of the Primate Isocortical Motor System. *Anat Embryol (Berl)* 202, 443-474.

Giszter SF, Mussa-Ivaldi FA, and Bizzi E (1993) Convergent Force Fields Organized in the Frog&Apos;S Spinal Cord. *J Neurosci* 13, 467-491.

Godschalk MM, Mitz ARA, Van Duin BB, and Van Der Burg HH (1995) Somatotopy of Monkey Premotor Cortex Examined with Microstimulation. *Neurosci Res* 23, 269-279.

Gottlieb JP, Bruce CJ, and Macavoy MG (1993) Smooth Eye Movements Elicited by Microstimulation in the Primate Frontal Eye Field. *J Neurophysiol* 69, 786-799.

Graziano MS. (2009) *The Intelligent Movement Machine: An Ethological Perspective on the Primate Motor System*. New York: Oxford University Press.

Graziano MS and Gandhi S (2000) Location of the Polysensory Zone in the Precentral Gyrus of Anesthetized Monkeys. *Exp Brain Res* 135, 259-266.

Graziano MS, Reiss LA, and Gross CG (1999) A Neuronal Representation of the Location of Nearby Sounds. *Nature* 397, 428-430.

Graziano MSA (2006) The Organization of Behavioral Repertoire in Motor Cortex. *Annu Rev Neurosci* 29, 105-134.

Graziano MSA, Aflalo TN, and Cooke DF (2005) Arm Movements Evoked by Electrical Stimulation in the Motor Cortex of Monkeys. *J Neurophysiol* 94, 4209-4223.

Graziano MSA, Cooke DF, Taylor CSR, and Moore T (2004) Distribution of Hand Location in Monkeys During Spontaneous Behavior. *Exp Brain Res* 155, 30-36.

Graziano MSA, Hu XT, and Gross CG (1997) Visuospatial Properties of Ventral Premotor Cortex. *J Neurophysiol* 77, 2268-2292.

Graziano MSA, Patel K, and Taylor C (2004) Mapping from Motor Cortex to Biceps and Triceps Altered by Elbow Angle. *J Neurophysiol* 92, 395-407.

Graziano MSA, Taylor CSR, and Moore T (2002) Complex Movements Evoked by Microstimulation of Precentral Cortex. *Neuron* 34, 841-851.

Gribble PL and Ostry DJ (2000) Compensation for Loads During Arm Movements Using Equilibrium-Point Control. *Exp Brain Res* 135, 474-482.

Gribble PL, Ostry DJ, Sanguineti V, and Laboissiere R (1998) Are Complex Control Signals Required for Human Arm Movement? *J Neurophysiol* 79, 1409-1424.

Haiss F and Schwarz C (2005) Spatial Segregation of Different Modes of Movement Control in the Whisker Representation of Rat Primary Motor Cortex. *J Neurosci* 25, 1579-1587.

Halsband UU, Matsuzaka YY, and Tanji JJ (1994) Neuronal Activity in the Primate Supplementary, Pre-Supplementary and Premotor Cortex During Externally and Internally Instructed Sequential Movements. *Neurosci Res* 20, 149-155.

He SQ, Dum RP, and Strick PL (1993) Topographic Organization of Corticospinal Projections from the Frontal Lobe: Motor Areas on the Lateral Surface of the Hemisphere. *J Neurosci* 13, 952-980.

He SQ, Dum RP, and Strick PL (1995) Topographic Organization of Corticospinal Projections from the Frontal Lobe: Motor Areas on the Medial Surface of the Hemisphere. *J Neurosci* 15, 3284-3306.

Hendrix CM, Mason CR, and Ebner TJ (2009) Signaling of Grasp Dimension and Grasp Force in Dorsal Premotor Cortex and Primary Motor Cortex Neurons During Reach to Grasp in the Monkey. *J Neurophysiol* 102, 132-145.

Hill A. (1970) *First and Last Experiments in Muscle Mechanics*. Cambridge, UK: Cambridge University Press.

Hocherman S and Wise SP (1991) Effects of Hand Movement Path on Motor Cortical Activity in Awake, Behaving Rhesus Monkeys. *Exp Brain Res* 83, 285-302.

Holdefer RN and Miller LE (2002) Primary Motor Cortical Neurons Encode Functional Muscle Synergies. *Exp Brain Res* 146, 233-243.

Hoover JE and Strick PL (1993) Multiple Output Channels in the Basal Ganglia. *Science* 259, 819-821.

Huang CS, Hiraba H, Muray GM, and Sessle BJ (1989) Topographical Distribution and Functional-Properties of Cortically Induced Rhythmical Jaw Movements in the Monkey (Macaca-Fascicularis). *J Neurophysiol* 61, 635-650.

Huber D, Gutnisky DA, Peron S, Oconnor DH, Wiegert JS, Tian L, Oertner TG, Looger LL, and Svoboda K (2012) Multiple Dynamic Representations in the Motor Cortex During Sensorimotor Learning. *Nature* 484, 473-478.

Jeannerod M, Arbib MA, Rizzolatti G, and Sakata H (1995) Grasping Objects: The Cortical Mechanisms of Visuomotor Transformation. *Trends Neurosci* 18, 314-320.

Johnson PB, Ferraina S, Bianchi L, and Caminiti R (1996) Cortical Networks for Visual Reaching: Physiological and Anatomical Organization of Frontal and Parietal Lobe Arm Regions. *Cereb Cortex* 6, 102-119.

Kaas J (2000) Why Is Brain Size So Important: Design Problems and Solutions as Neocortex Gets Bigger or Smaller. *Brain and Mind* 1, 7-23.

Kaas J and Catania K (2002) How Do Features of Sensory Representations Develop? *Bioessays* 24, 334-343.

Kaas JH, Gharbawie OA, and Stepniewska I (2011) The Organization and Evolution of Dorsal Stream Multisensory Motor Pathways in Primates. *Frontiers in neuroanatomy* 5, 34.

Kakei S, Hoffman D, and Strick PL (1999) Muscle and Movement Representations in the Primary Motor Cortex. *Science* 285, 2136-2139.

Kakei S, Hoffman D, and Strick PL (2001) Direction of Action Is Represented in the Ventral Premotor Cortex. *Nat Neurosci* 4, 1020-1025.

Karni A, Meyer G, Jezzard P, Adams MM, Turner R, and Ungerleider LG (1995) Functional Mri Evidence for Adult Motor Cortex Plasticity During Motor Skill Learning. *Nature* 377, 155-158.

Kermadi II, Liu YY, Tempini AA, and Rouiller EME (1997) Effects of Reversible Inactivation of the Supplementary Motor Area (Sma) on Unimanual Grasp and Bimanual Pull and Grasp Performance in Monkeys. *Somatosensory & motor research* 14, 268-280.

Kettner RE, Schwartz AB, and Georgopoulos AP (1988) Primate Motor Cortex and Free Arm Movements to Visual Targets in Three-Dimensional Space. Iii. Positional Gradients and Population Coding of Movement Direction from Various Movement Origins. *J Neurosci* 8, 2938-2947.

Kleim JA, Barbay S, and Nudo RJ (1998) Functional Reorganization of the Rat Motor Cortex Following Motor Skill Learning. *J Neurophysiol* 80, 3321-3325.

Klyachko VA and Stevens CF (2003) Connectivity Optimization and the Positioning of Cortical Areas. *PNAS* 100, 7937-7941.

Kohonen T (1982) Self-Organized Formation of Topologically Correct Feature Maps. *Biol Cybern* 43, 59-69.

Kohonen T. (1989) *Self-Organization and Associative Memory*. Berlin, Germany: Springer.

Kohonen T. (2001) *Self-Organizing Maps*. Berlin, Germany: Springer.

Kolster H, Mandeville J, and Arsenault J (2009) Visual Field Map Clusters in Macaque Extrastriate Visual Cortex. *J Neurosci* 2009, 7031-7039.

Krishnamoorthy V, Scholz JP, and Latash ML (2007) The Use of Flexible Arm Muscle Synergies to Perform an Isometric Stabilization Task. *Clin Neurophysiol* 118, 525-537.

Kurtzer I, Herter TM, and Scott SH (2005) Random Change in Cortical Load Representation Suggests Distinct Control of Posture and Movement. *Nat Neurosci* 8, 498-504.

Kurtzer I, Pruszynski J, Herter T, and Scott S (2006) Primate Upper Limb Muscles Exhibit Activity Patterns That Differ from Their Anatomical Action During a Postural Task. *J Neurophysiol* 95, 493-504.

Kwan HC, Mackay WA, Murphy JT, and Wong YC (1978) Spatial Organization of Precentral Cortex in Awake Primates. Ii. Motor Outputs. *J Neurophysiol* 41, 1120-1131.

Lackner JR and Dizio P (1994) Rapid Adaptation to Coriolis Force Perturbations of Arm Trajectory. *J Neurophysiol* 72, 299-313.

Landgren S and Phillips C (1962) Minimal Synaptic Actions of Pyrimidal Impulses on Some Alpha Motoneurones of the Baboon&Apos;S Hand and Forearm. *J Physiol*.

Lee D and Quessy S (2003) Activity in the Supplementary Motor Area Related to Learning and Performance During a Sequential Visuomotor Task. *J Neurophysiol* 89, 1039-1056.

Lemon RN, Johansson R, and Westling G (1995) Corticospinal Control During Reach, Grasp, and Precision Lift in Man. *J Neurosci* 15, 6145-6156.

Leyton A and Sherrington CS (1917) Observations on the Excitable Cortex of the Chimpanzee, Orang-Utan and Gorilla. *Q J Exp Physiol* 11, 135-222.

Liepert J, Terborg C, and Weiller C (1999) Motor Plasticity Induced by Synchronized Thumb and Foot Movements. *Exp Brain Res* 125, 435-439.

Loeb GE, Brown IE, and Cheng EJ (1999) A Hierarchical Foundation for Models of Sensorimotor Control. *Exp Brain Res* 126, 1-18.

Luh J and Walker M (1980) On-Line Computational Scheme for Mechanical Manipulators. *Journal of Dynamic Systems Measurement and Control* 102, 69-76.

Luppino GG, Matelli MM, Camarda RMR, Gallese VV, and Rizzolatti GG (1991) Multiple Representations of Body Movements in Mesial Area 6 and the Adjacent Cingulate Cortex: An Intracortical Microstimulation Study in the Macaque Monkey. *J Comp Neurol* 311, 463-482.

Marconi B, Genovesio A, and Battaglia-Mayer A (2001) Eye–Hand Coordination During Reaching. I. Anatomical Relationships between Parietal and Frontal Cortex. *Cortex; a journal devoted to the study of the nervous system and behavior*.

Martin JH, Engber D, and Meng Z (2005) Effect of Forelimb Use on Postnatal Development of the Forelimb Motor Representation in Primary Motor Cortex of the Cat. *J Neurophysiol* 93, 2822-2831.

Matelli M, Luppino G, and Rizzolatti G (1985) Patterns of Cytochrome Oxidase Activity in the Frontal Agranular Cortex of the Macaque Monkey. *Behav Brain Res* 18, 125-136.

Matelli M, Luppino G, and Rizzolatti G (1991) Architecture of Superior and Mesial Area 6 and the Adjacent Cingulate Cortex in the Macaque Monkey. *J Comp Neurol* 311, 445-462.

Mckiernan BJ, Marcario J, Karrer J, and Cheney PD (1998) Corticomotoneuronal Postspike Effects in Shoulder, Elbow, Wrist, Digit, and Intrinsic Hand Muscles During a Reach and Prehension Task. *J Neurophysiol* 80, 1961-1980.

Meier JD, Aflalo TN, Kastner S, and Graziano MSA (2008) Complex Organization of Human Primary Motor Cortex: A High-Resolution Fmri Study. *J Neurophysiol* 100, 1800-1812.

Monfils M-H, Plautz EJ, and Kleim JA (2005) In Search of the Motor Engram: Motor Map Plasticity as a Mechanism for Encoding Motor Experience. *The Neuroscientist* 11, 471-483.

Moran DW and Schwartz AB (1999) Motor Cortical Representation of Speed and Direction During Reaching. *J Neurophysiol* 82, 2676-2692.

Morrow MM, Pohlmeyer EA, and Miller LE. (2009) "Control of Muscle Synergies by Cortical Ensembles." In *Progress in Motor Control*. 179-199.

Mushiake HH, Inase MM, and Tanji JJ (1991) Neuronal Activity in the Primate Premotor, Supplementary, and Precentral Motor Cortex During Visually Guided and Internally Determined Sequential Movements. *J Neurophysiol* 66, 705-718.

Nudo RJ and Milliken GW (1996) Reorganization of Movement Representations in Primary Motor Cortex Following Focal Ischemic Infarcts in Adult Squirrel Monkeys. *J Neurophysiol* 75, 2144-2149.

Nudo RJ, Milliken GW, Jenkins WM, and Merzenich MM (1996) Use-Dependent Alterations of Movement Representations in Primary Motor Cortex of Adult Squirrel Monkeys. *J Neurosci* 16, 785-807.

Obermayer K and Blasdel GG (1993) Geometry of Orientation and Ocular Dominance Columns in Monkey Striate Cortex. *J Neurosci* 13, 4114-4129.

Okano KK and Tanji JJ (1987) Neuronal Activities in the Primate Motor Fields of the Agranular Frontal Cortex Preceding Visually Triggered and Self-Paced Movement. *Exp Brain Res* 66, 155-166.

Ostry DJ and Feldman AG (2003) A Critical Evaluation of the Force Control Hypothesis in Motor Control. *Exp Brain Res* 153, 275-288.

Paninski L, Fellows MR, Hatsopoulos NG, and Donoghue JP (2004) Spatiotemporal Tuning of Motor Cortical Neurons for Hand Position and Velocity. *J Neurophysiol* 91, 515-532.

Park M, Belhaj-Saif A, and Cheney PD (2004) Properties of Primary Motor Cortex Output to Forelimb Muscles in Rhesus Macaques. *J Neurophysiol* 92, 2968-2984.

Park MC, Belhaj-Saif A, Gordon M, and Cheney PD (2001) Consistent Features in the Forelimb Representation of Primary Motor Cortex in Rhesus Macaques. *J Neurosci* 21, 2784-2792.

Pascual-Leone A, Nguyet D, Cohen LG, Brasil-Neto JP, Cammarota A, and Hallett M (1995) Modulation of Muscle Responses Evoked by Transcranial Magnetic Stimulation During the Acquisition of New Fine Motor Skills. *J Neurophysiol* 74, 1037-1045.

Penfield W and Boldrey E (1937) Somatic Motor and Sensory Representation in the Cerebral Cortex of Man as Studied by Electrical Stimulation. *Brain* 60, 389-443.

Penfield W and Welch K (1951) The Supplementary Motor Area of the Cerebral Cortex; a Clinical and Experimental Study. *A.M.A. archives of neurology and psychiatry* 66, 289-317.

Picard N and Strick PL (2003) Activation of the Supplementary Motor Area (Sma) During Performance of Visually Guided Movements. *Cereb Cortex* 13, 977-986.

Pilon J-F, De Serres SJ, and Feldman AG (2007) Threshold Position Control of Arm Movement with Anticipatory Increase in Grip Force. *Exp Brain Res* 181, 49-67.

Polit A and Bizzi E (1979) Characteristics of Motor Programs Underlying Arm Movements in Monkeys. *J Neurophysiol* 42, 183-194.

Preuss T, Stepniewska I, and Kaas J (1997) Movement Representation in the Dorsal and Ventral Premotor Areas of Owl Monkeys: A Microstimulation Study. *J Comp Neurol* 371, 649-675.

Ramanathan D, Conner JM, and Tuszynski MH (2006) A Form of Motor Cortical Plasticity That Correlates with Recovery of Function after Brain Injury. *Proc Natl Acad Sci U S A* 103, 11370-11375.

Raos V, Franchi G, Gallese V, and Fogassi L (2003) Somatotopic Organization of the Lateral Part of Area F2 (Dorsal Premotor Cortex) of the Macaque Monkey. *J Neurophysiol* 89, 1503-1518.

Raos V, Umiltá M-A, Gallese V, and Fogassi L (2004) Functional Properties of Grasping-Related Neurons in the Dorsal Premotor Area F2 of the Macaque Monkey. *J Neurophysiol* 92, 1990-2002.

Raphael G, Tsianos GA, and Loeb GE (2010) Spinal-Like Regulator Facilitates Control of a Two-Degree-of-Freedom Wrist. *J Neurosci* 30, 9431-9444.

Rathelot J-A and Strick PL (2006) Muscle Representation in the Macaque Motor Cortex: An Anatomical Perspective. *Proc Natl Acad Sci U S A* 103, 8257-8262.

Rathelot J-A and Strick PL (2009) Subdivisions of Primary Motor Cortex Based on Cortico-Motoneuronal Cells. *Proc Natl Acad Sci U S A* 106, 918-923.

Reina GA, Moran DW, and Schwartz AB (2001) On the Relationship between Joint Angular Velocity and Motor Cortical Discharge During Reaching. *J Neurophysiol* 85, 2576-2589.

Richardson AG, Slotine J, Bizzi E, and Tresch M (2005) Intrinsic Musculoskeletal Properties Stabilize Wiping Movements in the Spinalized Frog. *J Neurosci* 25, 3181-3191.

Rizzolatti G (1998) The Organization of the Cortical Motor System: New Concepts. *Electroencephalogr Clin Neurophysiol* 106, 283-296.

Rizzolatti G and Luppino G (2001) The Cortical Motor System. *Neuron* 31, 889-901.

Rizzolatti G, Scandolara C, Matelli M, and Gentilucci M (1981) Afferent Properties of Periarcuate Neurons in Macaque Monkeys. Ii. Visual Responses. *Behav Brain Res* 2, 147-163.

Rizzolatti GG, Camarda RR, Fogassi LL, Gentilucci MM, Luppino GG, and Matelli MM (1988) Functional Organization of Inferior Area 6 in the Macaque Monkey. Ii. Area F5 and the Control of Distal Movements. *Exp Brain Res* 71, 491-507.

Robinson DA and Fuchs AF (1969) Eye Movements Evoked by Stimulation of Frontal Eye Fields. *J Neurophysiol* 32, 637-648.

Romo R, Hernández A, Zainos A, and Salinas E (1998) Somatosensory Discrimination Based on Cortical Microstimulation. *Nature* 392, 387-390.

Romo R and Schultz W (1987) Neuronal Activity Preceding Self-Initiated or Externally Timed Arm Movements in Area 6 of Monkey Cortex. *Exp Brain Res* 67, 656-662.

Rosenbaum DA, Loukopoulos LD, Meulenbroek RGJ, Vaughan J, and Engelbrecht SE (1995) Planning Reaches by Evaluating Stored Postures. *Psychol Rev* 102, 28-67.

Saarinen J and Kohonen T (1985) Self-Organized Formation of Colour Maps in a Model Cortex. *Perception* 14, 711-719.

Sakata H, Taira M, Kusunoki M, Murata A, and Tanaka Y (1997) The Tins Lecture the Parietal Association Cortex in Depth Perception and Visual Control of Hand Action. *Trends Neurosci* 20, 350-357.

Salzman CD, Britten KH, and Newsome WT (1990) Cortical Microstimulation Influences Perceptual Judgements of Motion Direction. *Nature* 346, 174-177.

Sanes JN, Wang J, and Donoghue JP (1992) Immediate and Delayed Changes of Rat Motor Cortical Output Representation with New Forelimb Configurations. *Cereb Cortex* 2, 141-152.

Scholz JP and Schöner G (1999) The Uncontrolled Manifold Concept: Identifying Control Variables for a Functional Task. *Exp Brain Res* 126, 289-306.

Schreiner CE (1995) Order and Disorder in Auditory Cortical Maps. *Curr Opin Neurobiol* 5, 489-496.

Scott S, Brown IE, and Loeb GE (1996) Mechanics of Feline Soleus .1. Effect of Fascicle Length and Velocity on Force Output. *J Muscle Res Cell Motil* 17, 207-219.

Scott S and Kalaska JF (1997) Reaching Movements with Similar Hand Paths but Different Arm Orientations .1. Activity of Individual Cells in Motor Cortex. *J Neurophysiol* 77, 826-852.

Scott SH and Kalaska JF (1995) Changes in Motor Cortex Activity During Reaching Movements with Similar Hand Paths but Different Arm Postures. *J Neurophysiol* 73, 2563-2567.

Sergio L and Kalaska JF (1997) Systematic Changes in Directional Tuning of Motor Cortex Cell Activity with Hand Location in the Workspace During Generation of Static Isometric Forces in Constant Spatial Directions. *J Neurophysiol* 78, 1170-1174.

Sergio LE and Kalaska JF (2003) Systematic Changes in Motor Cortex Cell Activity with Arm Posture During Directional Isometric Force Generation. *J Neurophysiol* 89, 212-228.

Serrien DJ, Strens LHA, Oliviero A, and Brown P (2002) Repetitive Transcranial Magnetic Stimulation of the Supplementary Motor Area (Sma) Degrades Bimanual Movement Control in Humans. *Neurosci Lett* 328, 89-92.

Shadmehr R and Mussa-Ivaldi FA (1994) Adaptive Representation of Dynamics During Learning of a Motor Task. *J Neurosci* 14, 3208-3224.

Stark E, Asher I, and Abeles M (2007) Encoding of Reach and Grasp by Single Neurons in Premotor Cortex Is Independent of Recording Site. *J Neurophysiol* 97, 3351-3364.

Stepniewska I, Fang P-C, and Kaas JH (2005) Microstimulation Reveals Specialized Subregions for Different Complex Movements in Posterior Parietal Cortex of Prosimian Galagos. *Proc Natl Acad Sci U S A* 102, 4878-4883.

Stepniewska I, Fang P-CY, and Kaas JH (2009) Organization of the Posterior Parietal Cortex in Galagos: I. Functional Zones Identified by Microstimulation. *J Comp Neurol* 517, 765-782.

Stepniewska I, Friedman RM, Gharbawie OA, Cerkevich CM, Roe AW, and Kaas JH (2011) Optical Imaging in Galagos Reveals Parietal-Frontal Circuits Underlying Motor Behavior. *Proc Natl Acad Sci U S A* 108, E725-732.

Strick PL and Preston JB (1978) Multiple Representation in Primate Motor Cortex. *Brain Res* 154, 366-370.

Tanji J and Kurata K (1979) Neuronal Activity in the Cortical Supplementary Motor Area Related with Distal and Proximal Forelimb Movements. *Neurosci Lett* 12, 201-206.

Tanji J and Kurata K (1982) Comparison of Movement-Related Activity in Two Cortical Motor Areas of Primates. *J Neurophysiol* 48, 633-653.

Tanji J and Shima K (1994) Role for Supplementary Motor Area Cells in Planning Several Movements Ahead. *Nature* 371, 413 - 416.

Tanji JJ and Kurata KK (1985) Contrasting Neuronal Activity in Supplementary and Precentral Motor Cortex of Monkeys. I. Responses to Instructions Determining Motor Responses to Forthcoming Signals of Different Modalities. *J Neurophysiol* 53, 129-141.

Tanji JJ, Okano KK, and Sato KCK (1988) Neuronal Activity in Cortical Motor Areas Related to Ipsilateral, Contralateral, and Bilateral Digit Movements of the Monkey. *J Neurophysiol* 60, 325-343.

Tehovnik E (1996) Electrical Stimulation of Neural Tissue to Evoke Behavioral Responses. *J Neurosci Methods* 65, 1-17.

Tehovnik E and Lee K (1993) The Dorsomedial Frontal-Cortex of the Rhesus-Monkey - Topographic Representation of Saccades Evoked by Electrical-Stimulation. *Exp Brain Res* 96, 430-442.

Thaler D, Chen YC, Nixon PD, Stern CE, and Passingham RE (1995) The Functions of the

Medial Premotor Cortex. I. Simple Learned Movements. *Exp Brain Res* 102, 445-460.

Ting LH and Macpherson JM (2005) A Limited Set of Muscle Synergies for Force Control During a Postural Task. *J Neurophysiol* 93, 609-613.

Todorov E (2000) Direct Cortical Control of Muscle Activation in Voluntary Arm Movements: A Model. *Nat Neurosci* 3, 391-398.

Todorov E and Jordan MI (2002) Optimal Feedback Control as a Theory of Motor Coordination. *Nat Neurosci* 5, 1226-1235.

Umiltà MA, Escola L, Intskirveli I, Grammont F, Rochat M, Caruana F, Jezzini A, Gallese V, and Rizzolatti G (2008) When Pliers Become Fingers in the Monkey Motor System. *PNAS* 105, 2209-2213.

Vanessen D (1997) A Tension-Based Theory of Morphogenesis and Compact Wiring in the Central Nervous System. *Nature* 385, 313-318.

Vogt C and Vogt O (1919) Allgemeinere Ergebnisse Unserer Hirnforschung. *J. Psychol. Neurol.*, 279–439.

Von Bonin G and Bailey P. (1947) *The Neocortex of Macaca Mulatta.* Urbana, Illinois: University of Illinois Press.

Walker M and Orin DE (1982) Efficient Dynamic Computer-Simulation of Robotic Mechanisms. *Journal of Dynamic Systems Measurement and Control* 104, 205-211.

Wandell BA, Dumoulin SO, and Brewer AA (2007) Visual Field Maps in Human Cortex. *Neuron* 56, 366-383.

Wise SP (1985) The Primate Premotor Cortex - Past, Present, and Preparatory. *Annu Rev Neurosci* 8, 1-19.

Wise SP and Boussaoud D (1997) Premotor and Parietal Cortex: Corticocortical Connectivity and Combinatorial Computations. *Annu Rev Neurosci* 29, 25-42.

Woolsey CN, Settlage PH, Meyer DR, Sencer W, Pinto Hamuy T, and Travis AM (1952) Patterns of Localization in Precentral and "Supplementary" Motor Areas and Their Relation to the Concept of a Premotor Area. *Association for Research in Nervous and Mental Disease* 30, 238-264.

Wu C, Bichot N, and Kaas J (2000) Converging Evidence from Microstimulation, Architecture, and Connections for Multiple Motor Areas in the Frontal and Cingulate Cortex of Prosimian Primates. *J Comp Neurol* 423, 140-177.

Young MP (1992) Objective Analysis of the Topological Organization of the Primate Cortical Visual-System. *Nature* 358, 152-155.