

Cortical Control of Arm Movements: A Dynamical Systems Perspective

Krishna V. Shenoy,^{1,2} Maneesh Sahani,^{1,3}
and Mark M. Churchland⁴

¹Departments of Electrical Engineering, ²Bioengineering, and Neurobiology, Bio-X and Neurosciences Programs, Stanford Institute for Neuro-Innovation and Translational Neuroscience, Stanford University, Stanford, California 94305; email: shenoy@stanford.edu

³Gatsby Computational Neuroscience Unit, University College London, London WC1N 3AR, United Kingdom; email: maneesh@gatsby.ucl.ac.uk

⁴Department of Neuroscience, Grossman Center for the Statistics of Mind, David Mahoney Center for Brain and Behavior Research, Kavli Institute for Brain Science, Columbia University Medical Center, New York, NY 10032; email: mc3502@columbia.edu

Annu. Rev. Neurosci. 2013. 36:337–59

First published online as a Review in Advance on
May 29, 2013

The *Annual Review of Neuroscience* is online at
neuro.annualreviews.org

This article's doi:
10.1146/annurev-neuro-062111-150509

Copyright © 2013 by Annual Reviews.
All rights reserved

Keywords

premotor cortex, primary motor cortex, neural control of movement, dimensionality reduction

Abstract

Our ability to move is central to everyday life. Investigating the neural control of movement in general, and the cortical control of volitional arm movements in particular, has been a major research focus in recent decades. Studies have involved primarily either attempts to account for single-neuron responses in terms of tuning for movement parameters or attempts to decode movement parameters from populations of tuned neurons. Even though this focus on encoding and decoding has led to many seminal advances, it has not produced an agreed-upon conceptual framework. Interest in understanding the underlying neural dynamics has recently increased, leading to questions such as how does the current population response determine the future population response, and to what purpose? We review how a dynamical systems perspective may help us understand why neural activity evolves the way it does, how neural activity relates to movement parameters, and how a unified conceptual framework may result.

Contents

INTRODUCTION	338
A DYNAMICAL SYSTEMS	
PERSPECTIVE OF MOTOR	
CONTROL	340
An Alternative to the	
Representational View	340
The Population Dynamical State ...	342
CORTICAL ACTIVITY DURING	
MOVEMENT PREPARATION ...	344
Preparatory Activity as a	
Subthreshold Representation	345
Preparatory Activity as the Initial	
Dynamical State	345
The Dynamics of Preparation	346
THE TRANSITION	
TO MOVEMENT	348
CORTICAL ACTIVITY DURING	
MOVEMENT	349
DISCUSSION	352
FUTURE ISSUES	353

INTRODUCTION

It is difficult to appreciate just how central movement is to everyday life until the ability to move is lost owing to neurological injury or disease. Moving is how we interact and communicate with the world. We move our legs and feet to walk, we move our arms and hands to manipulate the objects around us, and we move our tongues and vocal cords to speak. Movement is therefore also central to self-image and psychological well-being. Decades of research have explored the neural basis of movement preparation, generation, and control. In particular, a substantial body of knowledge about the cortical control of arm movements in rhesus macaques has grown from Evarts' pioneering research (e.g., Evarts 1964, 1968; Georgopoulos et al. 1982, 1986; Kalaska 2009; Tanji & Evarts 1976; Weinrich & Wise 1982; Wise 1985). This knowledge recently helped investigators to design cortically controlled neural prosthetic systems aimed at

restoring motor function to paralyzed patients (for recent reviews, see, e.g., Green & Kalaska 2011, Hatsopoulos & Donoghue 2009).

Extensive as these discoveries have been, and encouraging as these medical applications are, our understanding of the neural control of movement remains incomplete. Indeed, there is remarkably little agreement regarding even the basic response properties of the motor cortex, including PMd and M1 (e.g., Churchland et al. 2010a; Churchland & Shenoy 2007b; Fetz 1992; Graziano 2009, 2011a; Hatsopoulos 2005; Mussa-Ivaldi 1988; Reimer & Hatsopoulos 2009; Scott 2000, 2008; Scott & Kalaska 1995; Todorov 2000 and associated articles). This lack of agreement contrasts starkly with, say, the primary visual cortex, where basic response properties have been largely agreed upon for decades. To understand the motor cortex is thus a major challenge, as well as an essential step toward designing more capable, accurate, and robust neural prostheses (e.g., Gilja et al. 2011, 2012; Shenoy et al. 2011).

Much of the controversy over motor cortex responses has hinged on the question of whether the cortical activity codes (or represents) muscle action on the one hand or higher-level movement parameters such as effector velocity on the other. **Figure 1** illustrates the dichotomy. Cortical activity passes, via the spinal cord, to the muscles, which contract to move the arm; but the temporal patterns of muscle activity and hand movement differ. Which signal is found in the cortex? Does the firing of cortical cells drive muscle contraction with little intervening translation, so that cortical activity resembles muscle activity; or does it encode the intended movement end point or path, to be transformed by the spinal cord into commands that contract the muscles? Studies correlating neural activity with electromyographic (EMG) muscle activity or with movement kinematics (factors such as velocity and position) have proven equivocal; investigators have seen both patterns (for a recent review, see Kalaska 2009). Just as critically, the activity of most neurons is poorly explained by either pattern (e.g., Churchland & Shenoy 2007b, Graziano

M1: primary motor cortex

PMd: premotor cortex, dorsal aspect

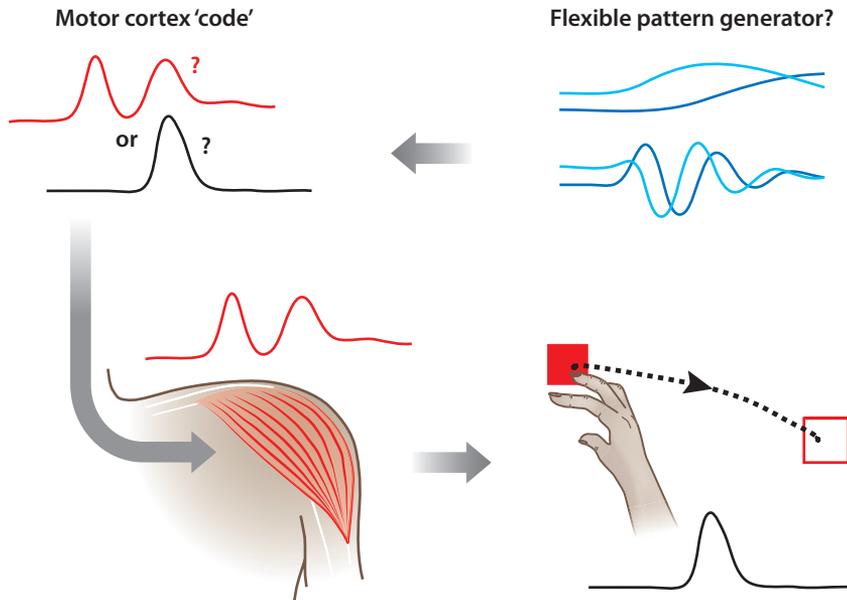


Figure 1

Schematic illustrating the focus of the representational perspective and of the dynamical systems perspective. The traditional perspective has concentrated on the representation or code employed by the motor cortex. For example, does the motor cortex (*upper left panel*) code muscle activity (*red trace*) or reach velocity (*black trace*)? Thus, the traditional perspective attempts to determine the output or controlled parameters of the motor cortex. The dynamical systems perspective focuses less on the output itself and more on how that output is created (*upper right panel*). It attempts to isolate the basic patterns (*blue*) from which the final output might be built. It further attempts to understand the dynamics that produced that set of patterns and the role of preparatory activity in creating the right set of patterns for a particular movement. The red trace indicates the activity of the deltoid versus time during a rightward reach (e.g., Churchland et al. 2012). The black trace is the hand velocity for that same reach; the black trace between the beginning and ending reach targets is the hand path. The light and dark blue traces (*upper right*) illustrate a potential dynamical basis set from which the red trace might be built.

2011b, Scott 2008). Thus, the controversy has continued.

In fact, determining the ‘code’ or ‘representation’ in motor cortex is but half the challenge. Whatever the cortical output, its temporal pattern must be generated by the circuitry of the cortex and reciprocally connected subcortical structures. Where is this flexible pattern generator—which can produce the wide variety of motor commands necessary to drive our large repertoire of movements—to be found? Is it upstream of M1, handing down a ‘motor program’ to be executed there (Miles & Evarts 1979), or is the pattern at least partly generated in M1 itself? Questions such as these suggest a different way to study the motor cortex, shifting the focus

from the meaning of the output to the nature of the dynamical system that creates the required, precisely patterned, command (e.g., Graziano 2011b). A core prediction of this perspective is that activity in the motor system reflects a mix of signals: Some will be outputs to drive the spinal cord and muscles, but many will be internal processes that help to compose the outputs but are themselves only poorly described in terms of the movement. They may, for instance, reflect a much larger basis set of patterns from which the eventual commands are built (see **Figure 1**). Some of these internal signals may well correlate coarsely with movement parameters: For example, one of the blue traces in **Figure 1** resembles hand position, whereas

Dynamical system: a physical system whose future state is a function of its current state, its input, and possibly some noise

another resembles a filtered version of hand velocity. But such coincidental correlations may not generalize across different tasks and need not constitute a representation of the movement parameters that is actively used by the brain (e.g., Churchland & Shenoy 2007b, Fetz 1992, Todorov 2000). Indeed, the dynamical systems perspective predicts that the evolution of neural activity should be best captured not in terms of movement parameter evolution, but in terms of the dynamical rules by which the current state causes the next state.

The dynamical systems perspective is not new to motor neuroscience. Brown, a student of Sherrington, argued in 1914 that internal pattern generators are at least as important as feedforward reflex arcs (Brown 1914, Yuste et al. 2005). Since then, the approach has shaped our understanding of central pattern generators (e.g., Grillner 2006, Kopell & Ermentrout 2002) and the brain stem circuitry that guides eye movements (e.g., Lisberger & Sejnowski 1992, Skavenski & Robinson 1973). In studying the motor cortex, Fetz (1992) argued for the dynamical systems perspective 20 years ago in an article entitled, “Are Movement Parameters Recognizably Coded in the Activity of Single Neurons?” He noted that

over the last three decades this formula [recording single neurons in behaving animals] has generated numerous papers illustrating neurons whose activity appears to code (i.e., to covary with) various movement parameters or representations of higher-order sensorimotor functions. . . . the search for neural correlates of motor parameters may actually distract us from recognizing the operation of radically different neural mechanisms of sensorimotor control. (p. 77)

The same point has been reiterated recently by Cisek 2006b, who summarizes that “the role of the motor system is to produce movement, not to describe it” (p. 2843). The dynamical systems perspective is also reflected in recent attempts to understand motor cortex as it relates to optimal feedback control (e.g., Scott

2004, Todorov & Jordan 2002). Indeed, the dynamical systems perspective may be experiencing a renaissance in neuroscience as a whole (e.g., Briggman et al. 2005, 2006; Broome et al. 2006; Mazor & Laurent 2005; Rabinovich et al. 2008; Stopfer et al. 2003; Yu et al. 2006), largely as the result of the widening adoption of multichannel recording techniques (e.g., Churchland et al. 2007, Harvey et al. 2012, Maynard et al. 1999), machine-learning based algorithms for estimating the population state from those recordings (e.g., Yu et al. 2009), and the computational resources necessary for data analysis and the exploration of plausible models. Just as importantly, there are growing bodies of neural data that are difficult to interpret from a purely representational framework but may be more approachable when dynamical systems concepts are brought to bear (e.g., Ganguli et al. 2008, Machens et al. 2010).

In this review we focus on one such body of literature, that from the field of motor control. We focus less on the role of dynamics in the context of sensory feedback (e.g., Scott 2004) and more on the internal neural dynamics that occur during movement preparation and the subsequent dynamics that translate preparatory activity into movement activity.

A DYNAMICAL SYSTEMS PERSPECTIVE OF MOTOR CONTROL

An Alternative to the Representational View

In principle, the representational and dynamical perspectives are compatible: The first seeks to determine the parameters controlled by cortical output, whereas the latter seeks to determine how that output is generated. However, in practice, adoption of the representational perspective has led to attempts to explain most neural activity in terms of tuning for movement parameters. That is, studies have sought to describe the firing (r) of each neuron (n) in the motor cortex as a function of various parameters ($param_i$) of an upcoming

or concurrent movement:

$$r_n(t) = f_n(\text{param}_1(t), \text{param}_2(t), \dots). \quad 1.$$

Admittedly, the available range of parameters is extensive, so such models may be adjusted to exhibit considerable richness (e.g., Fu et al. 1995, Hatsopoulos & Amit 2012, Pearce & Moran 2012, Reimer & Hatsopoulos 2009, Wang et al. 2010). Possible covariates include the intended target location, the kinematics of the hand or of the joints, the activity of individual muscles or synergistic groups, the activity of proprioceptors, the predicted end point error, and many others. These parameters may also be filtered, allowing for varying time lags, differentiation, or integration of the corresponding time-dependent signals. The common theme, however, is that neuronal activity should be understood in terms of such representational functions.

By contrast, the dynamical systems perspective stresses the view that the nervous system is a machine that must generate a pattern of activity appropriate to drive the desired movement. That is, the cortical activity [a time-varying vector $\mathbf{r}(t)$], when mapped to muscle activity [a time-varying vector $\mathbf{m}(t)$] by downstream circuitry,

$$\mathbf{m}(t) = G[\mathbf{r}(t)], \quad 2.$$

must produce forces that move the body in a way that achieves the organism's goals. The mapping $G[\]$ captures the action of all the circuits that lie between the cortex and the muscles, which may themselves implement sophisticated controllers. The dimension of $\mathbf{m}(t)$, set approximately by the number of independent muscle groups or synergies, is much lower than that of $\mathbf{r}(t)$, the number of different neurons in the motor cortex. Thus it is unlikely that $G[\]$ will be invertible. That is, knowledge of the final output alone (e.g., desired muscle activity or kinematics) may be insufficient to determine fully the pattern of neural activity that generated the output. This view thus accords with the observation that the apparent tuning of many neurons changes idiosyncratically with time (Churchland & Shenoy 2007b), with arm

starting location (Caminiti et al. 1991), with posture (Takei et al. 1999, Scott & Kalaska 1995), and with movement speed (Churchland & Shenoy 2007b). More broadly, it may help to understand why, despite many well-designed experiments, the issue of representation in the motor cortex has remained unresolved (e.g., Reimer & Hatsopoulos 2009, Scott 2008). In this view, a confusion of representation is not unexpected: the functions (f_n) of Equation 1 may not exist for any proposed set of movement parameters (Churchland et al. 2010a).

By moving the activity $\mathbf{r}(t)$ to the right-hand side of the equation, the dynamical systems perspective brings into focus the system that must generate that firing pattern (Graziano 2011b). Mathematically, population activity evolves with a derivative $\dot{\mathbf{r}}$, scaled by time constant τ , that is determined by the local circuitry of the motor cortex acting on its current activity through a function $b(\)$ and by inputs that arrive from other areas, $\mathbf{u}(t)$:

$$\tau \dot{\mathbf{r}}(t) = b(\mathbf{r}(t)) + \mathbf{u}(t). \quad 3.$$

With the appropriate input, this dynamical system causes the population activity to trace a path in time that maps through G to generate the correct movement. As this occurs, the neurons in the population may exhibit a variety of response patterns. Some patterns will directly influence the output of G , but others will reflect the act of pattern generation itself. A central aim of the dynamical systems approach is to understand these responses and thus to understand how the dynamics of a neural population produce the temporal patterns needed to drive movement.

Thus, the representational and dynamical perspectives often suggest very different forms of experiment and analysis. If seeking a representation, one asks which parameters are represented by neural activity, in which reference frames, and how these representations are transformed from one reference frame to another (e.g., Andersen & Buneo 2002, Batista et al. 2007, McGuire & Sabes 2011, Mullette-Gillman et al. 2009, Pesaran et al. 2006). Equation 1 suggests that the pattern of neural

State-space

trajectory: evolution of network activity in a space where each axis captures one neuron's response, or a factor shared among neurons

activity in time, and across different movements, should resemble that of the encoded parameters. Thus by designing experiments in which movement variables (e.g., muscle activity) vary systematically, one searches for neural firing patterns that vary in the same way. Conversely, a failure to find neural activity that covaries systematically with muscle activity would be taken to falsify the hypothesis that the cortex is concerned with control of muscles (e.g., Hatsopoulos 2005) or at the very least to suggest a 'messiness' of representation.

From the dynamical standpoint, the possibility that some cortical activity patterns are only indirectly related to the movement and reflect instead internal states of the dynamical process dictates a different approach. At least three practical possibilities present themselves. First, one might seek evidence of this internal state-space via direct visualization (Yu et al. 2009) or by testing the prediction that the population response is more complex than expected given the final output (Churchland & Shenoy 2007b). Second, one might attempt to trace the causality of the dynamical system. One might ask how the population's premovement state is determined (e.g., Churchland et al. 2006c), how this state influences the subsequent neural activity (e.g., Churchland et al. 2010a, 2012), and how variability in this state influences both neural activity and the movement (e.g., Afshar et al. 2011, Churchland et al. 2006a). Finally, and perhaps most challenging, one might seek to characterize the function $b()$ of Equation 3 by mapping out attractor states (whether fixed-points, limit cycles, or more complex), probing the system's transient behavior (e.g., Buesing et al. 2012, Macke et al. 2011, Sussillo et al. 2013, Sussillo & Abbott 2009, Yu et al. 2006) and studying the effects of perturbations applied to the neural activity (e.g., Churchland & Shenoy 2007a, Diester et al. 2011, Gerits et al. 2012, O'Shea et al. 2013). These different investigations are reviewed in greater detail below.

The Population Dynamical State

A dynamical description of cortical function is inherently a description of activity at the

population level. This notion is evident in Equations 2 and 3 above, neither of which can easily be separated into single-neuron components. Unfortunately, obtaining direct empirical access to the relevant scale of population activity is challenging. The full dynamical system is an extensively connected recurrent network of millions of neurons, coupled through input and feedback signals with much of the rest of the brain. The best current measurement technology can record either individual activity of no more than hundreds of neurons (using silicon or microwire arrays or calcium imaging) or aggregate signals that pool over thousands or more neurons at a time (local field potentials, fluorescence changes in voltage-sensitive dyes, or hemodynamic responses). Neither recording scale would seem suited to describing in detail the activity of the whole population. The unreliability of neuronal spiking introduces further challenges (e.g., Faisal et al. 2008, Manwani et al. 2002). Activity cannot be time-averaged on a scale longer than the dynamical time constant of Equation 3 without distorting the resulting dynamics. Similarly, one cannot average over repeated movements to construct a peri-stimulus-time histogram (PSTH) without suppressing intrinsic trial-to-trial variability, which is often of interest (e.g., Afshar et al. 2011, Yu et al. 2009).

These challenges can be addressed using at least two approaches. The first approach avoids the attempt to visualize or describe the dynamical process directly. Instead, hypotheses derived from the dynamical systems viewpoint are tested by assessing related predictions. One example is the prediction that trial-to-trial variance should fall as movement preparation brings the activity of the cortex to a suitable initial point from which appropriate movement activity can be generated (Churchland et al. 2006c, 2010b).

The second approach uses statistical methods to infer the population state from the available data and to examine how that state changes with time. Neurons within a single cortical population do not act alone; instead the coordinated firing of all the neurons

presumably guides the evolution of activity within the population and the evolution of its outputs.

This coordination may be intuitively most clear in the context of representation. Activity in a sensory population that encodes the features of a stimulus will covary as those features change. If the number of features is fewer than the number of neurons in the population, then population-level activity must be confined to a space the dimension of which is lower than the number of neurons. Even if the stimulus were unknown, the relevant aspects of the population activity may still be read out by looking for this lower-dimensional coordination. The same idea applies when the low-dimensional structure derives from the population dynamics rather than from a stimulus representation.

There are at least two reasons to think that the essential dimensions of the dynamical state will be few and will be distributed across many, if not all, of the neurons within a local area. First, the tight recurrent connectivity of the network will naturally tend to spread activity between cells. Second, and more subtly, the need for the network to be robust against the very unreliability that hampers experimental observation favors redundant activation patterns. The vector \mathbf{r} in Equations 2 and 3 spans many neurons, and we assume that independent noise in the activity of those neurons (or, indeed, injury to some of them) has only minimal impact on the output of the map G . Thus both G and the function h that determines the dynamics are likely to pool responses from many neurons, compressing the high-dimensional activity into a smaller set of meaningful degrees of freedom and thereby rejecting noise. These meaningful degrees of freedom may be viewed as defining a restricted space of lower dimension that is embedded within the space of all possible activity patterns. Because only the projection of the activity into this lower-dimensional space matters both to the dynamics of the area and to the influence it exerts on the muscles, the meaningful outputs of h must also be confined to this space. Thus, the projection of \mathbf{r} into this

DEFINITION OF TERMS

Population dynamical state: a set of coordinates, often represented as a vector, describing the instantaneous configuration of a dynamical system and that is sufficient to determine the future evolution of that system and its response to inputs. The population dynamical state of a neuronal network might be the vector of instantaneous firing of all its cells or may incorporate aspects of the neurons' biophysical states. It may also be a lower-dimensional projection of this network-wide description. See Dimensionality Reduction.

Dimensionality reduction: in this context, a technique for mapping the responses of many neurons onto a small number of variables that capture the basic patterns present in those responses. For example, the first variable/dimension might capture the response of a large proportion of neurons that all have very similar responses.

lower-dimensional space defines a population dynamical state (see Definition of Terms).

The dimensions explored by the population dynamical state may depend on the type of movements being performed. Over the full repertoire of movements (e.g., Foster et al. 2012, Gilja et al. 2010, Szuts et al. 2011), the range of dimensions might number in the thousands or more. However, in limited experimental settings with well-controlled motor outputs, the state may be confined to many fewer degrees of freedom. If so, then it should be possible to access the population dynamical state by means of dimensionality reduction techniques applied to the recorded data (e.g., Yu et al. 2009). These methods trace out the trajectory followed by the dynamical system, often on a single trial. Such trajectories make it possible to observe the dynamics more directly—indeed, in some cases the dimensionality reduction itself depends on forming a simultaneous estimate of the dynamical equations—and also to ask qualitative questions about the nature of the dynamics. For example, does the population state observed before the arm moves relate sensibly to the state trajectory traced out during the subsequent movement?

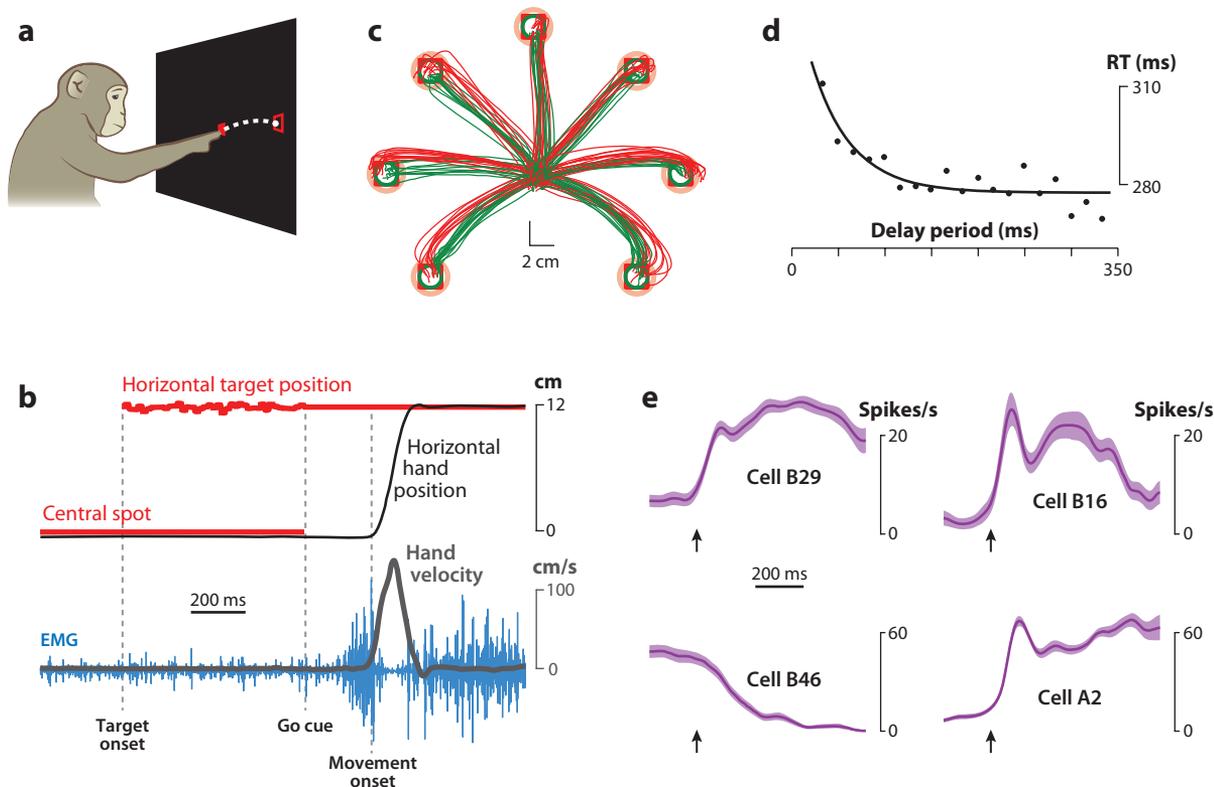


Figure 2

Overview of experimental paradigm, behavioral measurements, muscle measurements, and neural measurements. (a) Illustration of the instructed-delay task. Monkeys sit in a primate chair ~25 cm from a fronto-parallel display. A trial begins by fixating (eye) and touching (hand) a central target (*red filled square*) and holding for a few hundred milliseconds. A peripheral target (*red open square*) then appears, cuing the animal about where a movement must ultimately be made. After a randomized delay period (e.g., 0–1 s) a go cue is given (e.g., extinction of central fixation and touch targets) signaling that an arm movement to the peripheral target may begin. (b) Sample hand measurements and electromyographic (EMG) recordings for the same trial as in panel a. *Top*: Horizontal hand (*black*) and target (*red*) positions are plotted. For this experiment, the target jittered on first appearing and stabilized at the go cue. *Bottom*: Hand velocity superimposed on the voltage recorded from the medial deltoid. (c) Sample reach trajectories and end points in a center-out two-instructed-speed version of the instructed-delay task. Red and green traces/symbols correspond to instructed-fast and instructed-slow conditions. (d) Mean reaction time (RT) plotted versus delay-period duration. The line shows an exponential fit. (e) Examples of typical delay-period firing-rate responses in PMd. Mean \pm Standard Error firing rates for four sample neurons are shown. Figure adapted from Churchland et al. (2006c).

CORTICAL ACTIVITY DURING MOVEMENT PREPARATION

Studies indicate that voluntary movements are prepared before they are executed (Day et al. 1989, Ghez et al. 1997, Keele 1968, Kutas & Donchin 1974, Riehle & Requin 1989, Rosenbaum 1980, Wise 1985). To build intuition, consider the sudden, rapid, and accurate movement needed to swat a fly. An immediate, unpremeditated attack could miss, allowing the fly

to escape. Conversely, a short preparatory delay may permit the accuracy and velocity of movement to be improved, increasing the chances of success. In the laboratory, movement preparation has been studied by instructing a similar, but experimentally controlled, delay prior to a rapid, accurate movement (e.g., Mountcastle et al. 1975). **Figure 2a–c** illustrates the experimental design and task timing, along with sample hand position and EMG measurements.

Evidence that subjects use this instructed delay period to prepare a movement comes in part from the observation that reaction time (RT) is shorter on trials with a delay (e.g., Churchland et al. 2006c, Ghez et al. 1989, Riehle et al. 1997, Riehle & Requin 1989, Rosenbaum 1980). **Figure 2d** illustrates how RT first decreases and then plateaus with delay duration, suggesting that a time-consuming preparatory process has been given a head start during the initial ~ 200 ms of delay (e.g., Crammond & Kalaska 2000, Riehle & Requin 1989, Rosenbaum 1980).

Further evidence for movement preparation comes from neural recordings. Neurons in many cortical areas, including the parietal reach region (e.g., Snyder et al. 1997), PMd (e.g., Weinrich & Wise 1982), and M1 (e.g., Tanji & Evarts 1976), systematically modulate their activity during the delay. Thus, these motor-related areas appear to be engaged in computation prior to the movement (Crammond & Kalaska 2000). **Figure 2e** shows four PMd neurons that exemplify the range of delay-period firing patterns: Some neurons' firing rates increase, some decrease, and some stabilize after an initial transient, whereas others vary throughout. This variety of neural responses contrasts with the simple monotonic decline of behavioral RT (**Figure 2d**) and complicates efforts to understand the role of this activity.

Preparatory Activity as a Subthreshold Representation

Early proposals extended the representational view with the suggestion that preparatory neural activity represents the desired movement at a subthreshold level with the same tuning as that used during movement but with lower overall firing rates (e.g., Tanji & Evarts 1976). This lower-intensity activity is thought not to evoke movement by itself, but instead to reduce the time taken to achieve the correct suprathreshold firing pattern, thus shortening RT. This hypothesis has been assumed by many models of reach generation (e.g., Bastian et al. 1998, Cisek

2006a, Erlhagen & Schöner 2002, Schöner 2004) and agrees with our understanding of the saccadic system (e.g., Hanes & Schall 1996).

Many studies, particularly those exploring summary measures such as the population vector, have indeed reported consistently tuned neural activity before and during movement (e.g., Bastian et al. 1998, 2003; Cisek 2006a; Erlhagen et al. 1999; Georgopoulos et al. 1989; Requin et al. 1988; Riehle & Requin 1989). However, some studies at the single-neuron level have come to the opposite conclusion: that preparatory and movement tuning are often dissimilar (e.g., Wise et al. 1986) and nearly uncorrelated on average (e.g., Churchland et al. 2010a, Crammond & Kalaska 2000, Kaufman et al. 2010, 2013). Attempts to verify a threshold mechanism have also proven inconclusive. Higher premovement firing rates are not consistently associated with shorter RTs (e.g., Bastian et al. 2003, Churchland et al. 2006c). Furthermore, responses of cortical inhibitory interneurons seem inconsistent (Kaufman et al. 2010) with the common hypothesis that subthreshold preparatory activity is released from inhibition to initiate the movement (Bullock & Grossberg 1988, Sawaguchi et al. 1996).

Preparatory Activity as the Initial Dynamical State

The dynamical systems view suggests a different purpose for preparatory activity. Equations 2 and 3 describe the evolution of neural activity and its translation to muscle activity and thus to movement. The population state trajectory, and thus the movement produced, will clearly depend on the dynamics by which the population state evolves, captured by the function $b(\cdot)$. It may also be affected by descending input or feedback $[\mathbf{u}(t)]$ and by any sources of noise (e.g., van Beers et al. 2004). Finally, and crucially for our current purposes, the trajectory will depend on the population state $\mathbf{r}(t_0)$ at the time (t_0) that movement-related activity begins to be generated. Thus, all else remaining equal, different initial states will lead to different movements. This suggests that one role of preparation is

Reaction time (RT):
the time from the go cue until the start of the movement

to bring the population dynamical state to an initial value from which accurate movement-related activity will follow efficiently.

In general, more than one initial population dynamical state may lead to a movement that is sufficiently accurate: for example, a reach adequate to earn a reward. Assuming smoothness in the dynamics, and in the mapping to muscle activity and thus to kinematics, we might expect preparation for each movement to be associated with a compact subregion of the space of all possible population states (**Figure 3a**). State-space trajectories [$\mathbf{r}(t)$ for $t > t_0$] originating from different points in this subregion may vary; however, for the reach to be successful, such variation must (a) be confined to dimensions that are discounted by the mapping to muscles, (b) perturb the movement by too little to affect the desired outcome, or (c) be contained by compensatory changes in the external input provided by other areas, including corrective feedback signals.

Thus, the representational and dynamical perspectives both suggest that different movements should require different preparatory activity. Indeed, premovement firing is found to vary with every movement parameter studied so far (Cisek 2006b), including direction (e.g., Kurata 1989, Wise 1985), distance (e.g., Crammond & Kalaska 2000), speed (Churchland et al. 2006b), and curvature (Hocherman & Wise 1991); even apparently random variability in the preparatory state correlates with variability in the subsequent movement (Churchland et al. 2006a). However, short of a rapid de- and re-coding of activity between preparation and movement, the representational view predicts that preparatory and movement tuning should be congruent, which contrasts with the single-unit data as reviewed above. If the link between pre- and peri-movement activity were simply dynamical, on the other hand, then there would be no reason to necessarily expect such congruence.

Two recent studies have extended further support for the dynamical view. First, Churchland et al. (2010a) showed that although preparatory activity does indeed covary with

movement parameters such as direction, distance, speed, and curvature, it is more closely related to the pattern of cortical neural activity during the movement—as would be expected if the premovement population state led directly to the subsequent trajectory of movement-period neural activity, and only indirectly to the movement. Second, Kaufman et al. (2011) observed that preparatory states associated with different reaches were arranged along dimensions orthogonal to the dimensions of activity that correlate with changes in muscle force during movement. This result is consistent with a view in which preparatory activity does not itself engage changes in muscle output through the mapping $G[\]$ but nonetheless leads to movement control signals that do. In a representational picture, where preparatory and movement activity are similarly tuned, such orthogonality would be unexpected.

The Dynamics of Preparation

The end point of motor preparation is hypothesized to be an initial population dynamical state, from which the movement-period neural activity evolves to generate the desired movement. How is the correct initial state achieved between the times when the subject first sees the target and subsequently is told to move? Clearly the dynamics of movement preparation cannot be the same as the dynamics of movement activity. During movement preparation, the dynamical system must bring the population state toward the optimal preparatory region (as in **Figure 3a**) not away from it. Is it possible to detect signatures of this convergent dynamical process?

Activity in the experimental premovement period starts from a baseline condition, in which the only behavioral constraints are that the eyes remain fixated and the hand remains still (**Figure 2a**). There is little to prevent motor cortical activity in this state varying substantially across trials. During preparation, the activity then approaches the preparatory state, while avoiding the premature generation of movement. Again, because the intervening

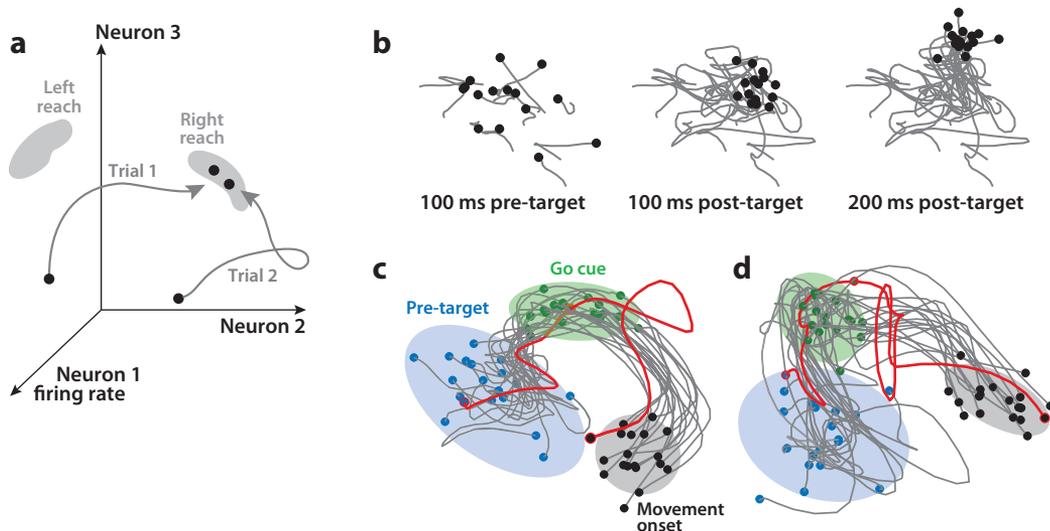


Figure 3

Schematic illustration of the optimal subspace hypothesis and single-trial neural trajectories computed using Gaussian process factor analysis (GPFA). (a) The configuration of firing rates is represented in a state-space, with the firing rate of each neuron contributing an axis, only three of which are drawn. For each possible movement, we hypothesize that there exists a subspace/subregion of states that are optimal in the sense that they will produce the desired result, with a minimal reaction time, when the movement is triggered. Different movements will have different optimal subspaces (*shaded areas*). The goal of motor preparation would be to optimize the configuration of firing rates so that it lies within the optimal subspace for the desired movement. For different trials (*arrows*), this process may take place at different rates, along different paths, and from different starting points. Figure from Churchland et al. (2006c). (b) Projections of PMd activity into a two-dimensional state-space. Each black point represents the location of neural activity on one trial. Gray traces show trajectories from 200 ms before target onset until the indicated time. The stimulus was a reach target (135° , 60 mm distant), with no reach allowed until a subsequent go cue; 15 (of 47) randomly selected trials are shown. (c) Trajectories were plotted until movement onset. Blue dots indicate 100 ms before stimulus (reach target) onset. No reach was allowed until after the go cue (*green dots*), 400–900 ms later. Activity between the blue and green dots thus relates to movement planning. Movement onset (*black dots*) was ~ 300 ms after the go cue. For display, 18 randomly selected trials are plotted, plus one hand-selected outlier trial (*red*, trial ID 211). Covariance ellipses were computed across all 47 trials. This is a two-dimensional projection of a ten-dimensional latent space. In the full space, the black ellipse is far from the edge of the blue ellipse. This projection was chosen to preserve accurate relative sizes (on a per-dimension basis) of the true ten-dimensional volumes of the ellipsoids. Data are from the G20040123 dataset. (d) Data are presented as in panel c, with the same target location, but for data from another day's data set (G20040122; *red* outlier trial: ID 793). Figure panels b–d adapted from Churchland et al. (2010b).

states do not themselves engage muscles, they may well be less constrained than those traversed during the movement's active phase. (See schematic trials 1 and 2 in **Figure 3a**.) The final preparatory state, however, is constrained by the need to generate the correct movement. Thus, we might expect that as preparation progresses, the relative variability across different trials should fall. Such a decrease has indeed been identified in the Fano factor of individual neurons in both the premotor cortex and motor cortex (Churchland et al. 2006c, Mandelblat-Cerf et al. 2009, Rickert et al. 2009). The

decline in variability is also apparent in the population dynamical state directly visualized via dimensionality reduction (Churchland et al. 2010b), as shown in **Figure 3b**. As predicted, the reduction in variance comes primarily from convergence in the low-dimensional population dynamical state rather than in the spiking noise of each cell (Churchland et al. 2010b). Finally, variability is only partially reduced when incomplete information about the target is provided (Rickert et al. 2009). These findings support the hypothesis that motor preparation requires network activity to converge to a

relatively tight set of population dynamical states. As an aside, a similar decline in neural variability is present in a variety of different cortical areas whenever a relevant stimulus is presented (Churchland et al. 2010b). These findings suggest that many different cortical computations may involve attractor-like dynamics. Nonetheless, the significance of such computation must depend on the function of the area. In visual cortex, the decline in variability may reflect the formation of a more consistent representation of the visual stimulus. In the motor areas considered here, the evidence (discussed below) indicates that the decline in variability reflects convergence to a preparatory state that has motor consequences.

What are the consequences if the convergence of the preparatory state is not complete at the time of the go cue? Instructed-delay experiments in which accuracy was emphasized have gathered some data to address this question. First, Churchland et al. (2006c) found that neural variability was lower among trials with short RT, in which motor preparation was likely to have been complete at the time of the go cue, than among trials with longer RT, in which motor preparation may have been incomplete or not quite accurate. This result is consistent with lower variability indicating greater preparatory accuracy (i.e., closer to the optimal preparatory state). Second, when subthreshold electrical microstimulation disrupted the preparatory state, RT was increased (Churchland & Shenoy 2007a). This effect could reflect additional time taken to recover the appropriate preparatory state (see also Ames et al. 2012 for another possibility). The effect was specific (Churchland & Shenoy 2007a). First, RT was more strongly affected when the microstimulation targeted the premotor cortex (where preparatory activity is more common) rather than the motor cortex (where preparatory activity is less common). Second, effects were seen only when the preparatory state was disrupted around the time of the go cue; disruption of the preparatory state before it was needed had little impact on RT. Third, the impact on reach RT was much greater than the

impact on saccadic RT, consistent with the role of the premotor cortex in preparing reaches rather than saccades and inconsistent with the possibility that microstimulation simply distracted the animal. Finally, O'Shea et al. (2013) recently found that optical stimulation of optogenetically transfected PMd neurons during the preparatory period similarly increases RT.

THE TRANSITION TO MOVEMENT

By itself, the idea that motor cortical activity represents or codes movement parameters (Equation 1) does not constrain the relationship between preparatory and movement activity. Nonetheless, this transition has frequently been thought to depend on the strengthening of a representation until it crosses a firing-rate threshold, by analogy to the oculomotor system. However, direct evidence for such a threshold has been lacking in the case of reaches (Bastian et al. 2003, Churchland et al. 2006c). By contrast, under the dynamical systems perspective (Equations 2 and 3), the transition to movement is a transition between two different types of network dynamics, most likely to be mediated by a change in the external input term $\mathbf{u}(t)$ of Equation 3. Preparatory dynamics, which brings the population to a suitable state of readiness, then gives way to the dynamics that generate movement. As the movement is triggered, the population dynamical state departs from the prepared initial state and follows a trajectory through state-space. It is that state-space trajectory—determined by the initial state, the neural dynamics, and any feedback—that drives the movement.

The transition from preparatory to movement dynamics may be directly observable. Petreska et al. (2011) used an unsupervised machine-learning technique to study changes in dynamics within multielectrode neuronal data gathered while animals performed instructed-delay reaches. They observed stereotyped dynamical transitions occurring at times shortly after target presentation as well as between the go cue and the beginning

of movement. The timing of the dynamical transition that followed the go cue was correlated with the timing of subsequent movement onset. Indeed, this transition—the identity and timing of which were determined by the neural data alone—predicted trial-to-trial variation in RT much better than did alternatives based on a threshold applied to overall firing or to the length of the population vector, even when that threshold was chosen with direct reference to the behavior.

Afshar et al. (2011) addressed a related issue. They reasoned that natural variability might occasionally displace the population dynamical state from the average point of preparation toward the direction in which that state will need to evolve when the movement is to be initiated. Any such variability could actually be beneficial to the initiation of movement and might reduce RT. Indeed, Afshar et al. found that the displacement of the dynamical state at the time of the go cue in the direction defined by the movement-period activity was negatively correlated with RT. Furthermore, RT was even lower if the preparatory state happened to be moving in that direction at the time of the go cue. Thus, although previous studies (Churchland et al. 2006c, Rickert et al. 2009) stressed the importance of an accurate and consistent preparatory state (RTs being lower on average when neural activity is near that state), displacement from the preparatory state can, in fact, result in a lower RT when the displacement is in the direction that is to be traversed during movement.

CORTICAL ACTIVITY DURING MOVEMENT

The dynamical systems perspective focuses on the population dynamical state and its evolution. Testing specific hypotheses therefore often requires direct visualization of that state. The traces in **Figure 3c,d** illustrate the trajectories of the population dynamical state on 19 trials from just before target onset to the moment when movement begins. After target onset, the dynamical state approaches

a preparatory region and its variability falls. Then, following the go cue, the neural state moves rapidly away from the preparatory state in a curved trajectory. Some trial-to-trial variability is evident even after the go cue. In particular, for two outlier trials, the neural state wanders before falling back on track. On these trials, the monkey hesitated for an abnormally long few hundred milliseconds before beginning to reach. These observations underscore the ability of dimensionality reduction methods, when applied to data collected from multielectrode arrays, to reveal single-trial (and potentially rare) phenomena that would normally have been lost to averaging or discarded owing to abnormal behavior. That said, for this highly practiced task, such trials were rare (about 0.1%). On the vast majority of trials, the population state evolved along a stereotyped curved trajectory. How can we characterize that trajectory: its shape, its time evolution, and the principles that give rise to it?

This relates to Fetz's original question, "Are movement parameters recognizably coded in the activity of single neurons?" If they are, then the neural state-space trajectory should reflect the trajectory of the represented parameters. For example, consider the model in which neural activity is cosine-tuned for reach velocity (Moran & Schwartz 1999). This relationship can be written in matrix form as $\mathbf{r}(t - \tau) = \mathbf{M}\mathbf{v}(t)$, where $\mathbf{r}(t)$ is an $n \times 1$ vector describing the firing rate of each neuron, $\mathbf{v}(t)$ is the three-dimensional reach velocity vector at time t , τ is the lag by which neural activity leads movements, and \mathbf{M} is an $n \times 3$ matrix describing (in each row) each neuron's preferred direction. Under this model, the population state would be three dimensional, with those dimensions capturing the neural representation of velocity. The population vector is a dimensionality reduction method made specifically for just such a situation.

However, a simple velocity-tuned model is inadequate to fully capture the richness of the neural responses. **Figure 4** illustrates the PSTH responses of two typical neurons recorded from the motor cortex

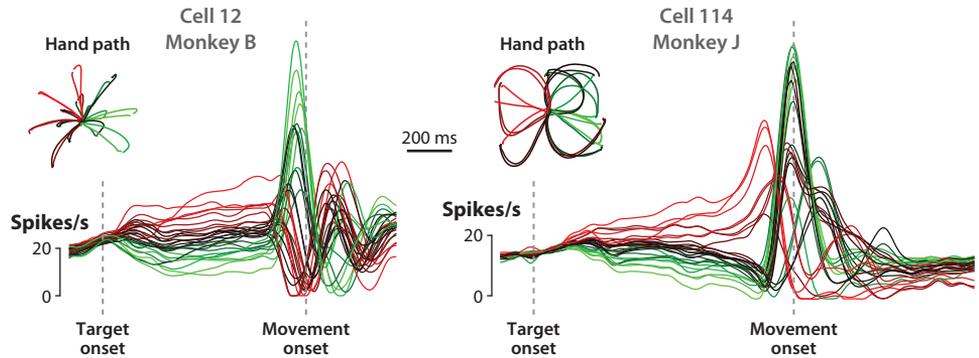


Figure 4

Peri-stimulus-time histogram (PSTHs) and arm-movement kinematics. PSTH from two sample neurons (*left and right panels*), including hand paths for sample reaches (*insets*). Monkey B performed a standard center-out reaching task, with two distances and two instructed speeds. Monkey J performed a more complex version of this task where some reaches were required to curve around a barrier. Traces are colored red to green on the basis of the relative level of preparatory activity for each condition. Insets show hand trajectories and are color-coded to reveal the directional nature of preparatory tuning. Note that what was preferred during the preparatory period was typically not preferred during the movement period; e.g., the first neuron shows a preference for left and up during the preparatory period but a preference for right and down by movement onset. Figure adapted from Churchland et al. (2010a).

of two monkeys. Monkey B performed a standard center-out reaching task, with two distances and two instructed speeds. Monkey J performed a similar task, but in it some reaches were required to curve around a barrier.

Four features of the neural responses are relevant to the controversy over what is being coded. First, the same neurons exhibit both preparatory- and movement-period activity, yet tuning during the preparatory period often differs from that during the movement period (e.g., cell 12 prefers up-left during the preparation, but down-right by movement onset) (Churchland et al. 2010a). Second, the movement-period responses are complex and multiphasic (Churchland et al. 2010a, 2012; Churchland & Shenoy 2007b). Third, the responses of different neurons are heterogeneous, even in the same animal and the same local region of the cortex (Churchland et al. 2010a, Churchland & Shenoy 2007b, Fetz 1992). In dynamical systems terms, the neural responses occupy a relatively high-dimensional space, on the order of 15–30 dimensions (Churchland & Shenoy 2007b). Thus, if neurons are to represent movement parameters, there must be

many such parameters (e.g., Pearce & Moran 2012). Finally, neural firing fluctuates over 400–800 ms, even when the reaches themselves are quite brief (e.g., the reaches for Monkey B lasted ~150–300 ms). Thus if movement parameters are represented directly, there must be some unexpected temporal multiplexing (e.g., Fu et al. 1995).

In large part because the neural responses are complex, our field does not yet agree on the relationship between movement-period neural responses and movement itself; the nature of G in Equation 2 remains mostly unresolved (although, see, e.g., Fetz et al. 2000). Yet some recent progress has been made in characterizing the nature of b in Equation 3 and in describing the dynamics that generate movement-period neural responses. First, the collective activity of motor cortical neurons is better described by a model in which activity is driven by a low-dimensional dynamical model, relative to a model in which firing coordination emerges from direct connections between recorded cells (Macke et al. 2011).

Second, the dynamics at play during movement appear to have some simple aspects. In

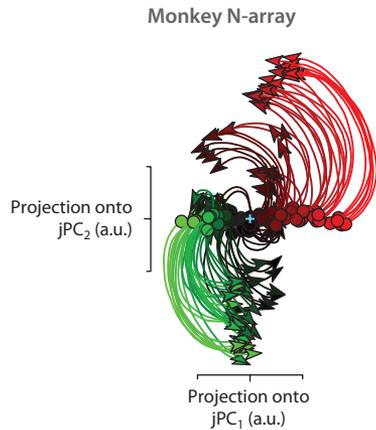


Figure 5

Projections of the neural population response, produced by applying jPCA to the first six principal components of the data. Two-dimensional projection using 218 single- and multiunit isolations, 108 conditions employing straight and curved reaches, from the N-array data set. Each trace is the average trajectory for one of 108 conditions. Trial-averaged neural trajectories are colored red to green on the basis of the level of preparatory activity for that projection. Each trace (one condition) plots the first 200 ms of movement-related activity away from the preparatory state (*circles*). Figure adapted from Churchland et al. (2012).

many lower-dimensionality projections, one of which is shown in **Figure 5**, the neural trajectory simply rotates with a phase and amplitude set by the preparatory state (Churchland et al. 2012). The rotational trajectories of the neural state resemble what is seen during rhythmic movement, even though the reaches were not overtly rhythmic. These trajectories suggest that the role of motor cortex may be most naturally thought of in the context of pattern generation. Consistent with this idea, EMG activity was well fit by the sum of two rhythmic components that were fixed in frequency but varied in phase and amplitude across the different movements (Churchland et al. 2012).

It should be stressed that the neural dynamics have aspects that are not well fit by a simple linear model (Churchland et al. 2012, Sussillo et al. 2013). However, a simple linear time-invariant system accounted for a large proportion of the variance (48.5% of

variance explained over the nine data sets tested) (Elsayed et al. 2013). Furthermore, the linear component was almost entirely normal and rotational (linear systems constrained to be normal and rotational performed 93.2% and 91.3% as well as did an unconstrained linear system) (Elsayed et al. 2013).

Although purely rotational dynamics are only an approximation to the true nonlinear dynamics, the observation that neural pattern generation involves rotations of the neural state illustrates two key points of the dynamical systems perspective. First, the goal of the preparatory state is not to act as an overt representation, but rather to set the amplitude and phase of the subsequent rotations. Second, those state-space rotations produce, in the temporal domain, brief sinusoidal oscillations that form a natural basis set for building more complex patterns (e.g., the blue traces in **Figure 1** can be linearly combined to fit the red EMG trace quite well). Thus, the neural dynamics and the resulting patterns can be understood in simple terms, even though they do not constitute an overt representation of movement parameters.

The observed similarity between the average trajectories during rhythmic and nonrhythmic movement might come about because the nervous system has redeployed old principles for a new purpose. Alternatively, rotations are a common dynamical motif and produce a natural basis set (Rokni & Sompolinsky 2012). Thus we see more than one potential explanation for the key features of the observed dynamics. These key features are that dynamics are similar across different reach types, have a strong rotational component, and have their phase and amplitude determined by the neural state achieved during movement preparation. A number of dynamical models, possibly including control-theory style models, may be able to account for these features (e.g., Scott 2004, Todorov & Jordan 2002). That said, the data are inconsistent with many classes of dynamics (Churchland et al. 2012). Because the observed dynamics are similar across reach types, they are not consistent with a system that is dominated by reach-specific inputs (e.g., a dynamical system that

converts velocity commands to muscle activity). The rotational patterns are also not consistent with rise-to-threshold or burst-generator models. This discrepancy is important because many previous intuitions regarding single-neuron activity derived from such models. Most centrally, these intuitions included the expectation that the preparatory- and movement-period preferred direction should be closely related. In contrast, the empirical preferred direction is in a constant state of flux, a natural consequence of the underlying rotations (Churchland et al. 2010a, 2012). This observation illustrates how the complex responses of individual neurons can often hide simple underlying structure—structure that is readily interpretable from a state-space, dynamical systems perspective.

DISCUSSION

The past 50 years have witnessed remarkable progress in our understanding of the cortical control of arm movements. Many fundamental principles have been discovered, and this basic scientific knowledge has led to rapid advances, including early clinical trials of cortically controlled neural prostheses for paralyzed participants (Collinger et al. 2013; Hochberg et al. 2006, 2012). Despite this considerable progress, it remains arguable whether an adequate conceptual framework has yet been identified, around which experimental, computational, and theoretical research can be oriented. Indeed without an adequate conceptual framework, it is unclear how a unified and comprehensive understanding for cortical motor control could be assembled or even recognized. This sentiment, initially expressed by Fetz in his 1992 article, appears to be of at least as great a concern today as it was 20 years ago. The following excerpts from recent articles serve as examples.

A shift in how to examine the motor system occurred in the 1980s from a problem of control back to a problem of what variables were coded in the activity of neurons. . . . [P]erhaps it is time to re-evaluate what we are learning

about M1 function from continuing to ask what coordinate frames or neural representations can be found in M1. Perhaps it is time to stop pursuing the penultimate goal of identifying the coordinate frame(s) represented in the discharge patterns of M1 and again move back to the question of control. (Scott 2008, p. 1220)

Neurophysiological experiments have revealed neural correlates of many arm movement parameters, ranging from the spatial kinematics of hand path trajectories to muscle activation patterns. However, there is still no broad consensus on the role of the motor cortex in the control of voluntary movement. The answer to that question will depend as much on further theoretical insights into the computational architecture of the motor system as on the design of the definitive neurophysiological experiment. (Kalaska 2009, p. 172)

An epic, twenty-year battle was fought over the cortical representation of movement. Do motor cortex neurons represent the direction of the hand during reaching, or do they represent other features of movement such as joint rotation or muscle output? As vigorous as this debate may have been, it still did not address the nature of the network within the motor cortex. Indeed, it tended to emphasize the properties of individual neurons rather than network properties. . . . The battles over the cortical representation of movement never satisfactorily addressed those questions. (Graziano 2011b, p. 388)

It appears that the field has reached a point where a new way of conceptualizing cortical motor control is needed. We have reviewed here a relatively new dynamical systems framework that appears to have several of the desired attributes. As such, the dynamical systems framework may help deepen our understanding of the neural control of movement. It may help do so by (a) making relatively few assumptions (e.g., being agnostic to tuning curves, specifically

their lack of invariance and generalization); (b) observing and documenting the dominant, and perhaps essential, features of neural state-space trajectories; (c) providing single-trial neural correlates of behavior that offer insights beyond those available using average relationships alone; (d) generating hypotheses that can be answered in the quantitative terms of population dynamical states and evolution rules (equations of motion), without the need to ascribe representational meaning to the detailed response of single neurons; and (e) being entirely open to the nature of the dynamics uncovered (e.g., ranging from pattern generators through sophisticated feedback controllers).

The dynamical systems framework is not single-neuron nihilistic: It does not ignore or attempt to average away the complex features of single-neuron responses. Indeed, we hope that by capturing the underlying dynamics it will become possible to explain the many seemingly odd aspects of individual-neuron responses. The dynamical systems perspective also provides a clear road map and goal, which is to quantify the dynamical systems instantiated by neural circuits. This mathematical quantification comes in two inter-related parts: the state-space neural trajectories (a focus of this review article) and the form and meaning of the evolution rule or equations of motion (a focus of ongoing research; e.g., Abeles et al. 1995, Churchland et al. 2012, Petreska et al. 2011, Rabinovich et al. 2008, Seidemann et al. 1996, Smith et al. 2004, Vaadia et al. 1995, Yu et al. 2006). It appears possible that three primary dynamical systems underlie reaching arm movements: one to prepare the neural state in an appropriate manner; a second system to use this computationally optimized starting point to generate movement activity, muscle contraction, and thus movement itself; and a third system that uses feedback for control. Much future research is certainly needed to explore this possibility, to extend and relate it to numerous behaviors beyond the instructed-delay point-to-point reaching task, and to see whether the dynamical systems perspective can

ultimately help provide a more comprehensive understanding of cortical motor control.

FUTURE ISSUES

The predictions, experiments, and analyses described above stem from a dynamical systems perspective, and to some degree their confirmation argues for that perspective. Yet many central questions remain largely unaddressed. What is the nature of the relevant dynamics (b in Equation 3), and why are they what they are? Do they relate to the dynamics of movement-generating circuits in simpler organisms? Do they reflect sophisticated mechanisms of online control and feedback (e.g., Scott 2004, Todorov & Jordan 2002)? How and why do dynamics change as a function of input from other brain areas (e.g., resting versus planning versus moving)? What is the nature of the circuitry, both local and feedback, that produces these dynamics? What is the mapping between the population dynamical state and muscular activity (G in Equation 2)? Answering such questions will likely depend on progress in three domains of research: first, the increased ability to resolve dynamical structure in neural data; second, the increased ability to perturb the population dynamical state while observing dynamical structure; and third, the increased ability to relate state-space trajectories to externally measurable parameters.

For example, it is becoming possible to employ optogenetic techniques to briefly increase (or decrease) the firing rate of excitatory or inhibitory neurons in the cortex of the rhesus monkey (e.g., Diester et al. 2011, O'Shea et al. 2013). This can be accomplished at various times relative to withholding, preparing, or generating arm movements while simultaneously observing the resulting perturbation and recovery of the population dynamical state. This class of pump-probe experiment should enable more quantitative measurement of the neural dynamics in operation during various phases of the behavioral task and should help illuminate the nature and operation of the neural circuitry underlying these neural dynamics.

SUMMARY POINTS

1. Movement preparation has long been thought to be a critical step in generating movement. Recent work supports this idea and argues that achieving the correct preparatory state is important for producing the desired movement.
2. Measurements of the preparatory state predict both reaction time and trial-to-trial movement variability. Disruption of the preparatory state increases reaction time.
3. Preparatory activity is not a subthreshold version of movement activity but instead appears to serve as an initial state for dynamics that engage shortly before movement onset.
4. The onset of these dynamics is tied to movement onset rather than to the go cue and is predictive of trial-by-trial reaction time.
5. Neural responses during the movement appear complex but have at least some simple aspects: Dynamics can be approximated by a linear differential equation in which the same dynamics govern many reaching movements.
6. Because dynamics are similar across conditions, the pattern of movement-related activity is determined largely by the preparatory state.
7. The best linear approximation (to the true nonlinear dynamics) is dominated by rotational dynamics. Preparatory activity sets the amplitude and phase of the movement-period rotations.
8. The resulting firing rate patterns form a natural basis set for building more complex patterns such as muscle activity.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We apologize in advance to all the investigators whose research could not be appropriately cited owing to space limitations. This work was supported by Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (to K.V.S. and M.M.C.), DARPA REPAIR N66001-10-C-2010 (to K.V.S. and M.S.), NSF-NIH CRCNS NINDS R01NS054283 (to K.V.S. and M.S.), an NIH Director's Pioneer Award 1DP1OD006409 (to K.V.S.), the Gatsby Charitable Foundation (to M.S.), Searle Scholars Program (to M.M.C.), and NIH Director's New Innovator Award DP2NS083037 (to M.M.C.).

LITERATURE CITED

- Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, et al. 1995. Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. USA* 92:8616–20
- Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. 2011. Single-trial neural correlates of arm movement preparation. *Neuron* 71:555–64
- Ames KC, Ryu SI, Shenoy KV. 2012. Neural dynamics of reaching following incomplete or incorrect planning. *Front. Neurosci. Conf.: Comput. Syst. Neurosci. (COSYNE)*. Abstr. T-5
- Andersen RA, Buneo CA. 2002. Intentional maps in posterior parietal cortex. *Annu. Rev. Neurosci.* 25:189–220

- Bastian A, Riehle A, Erlhagen W, Schöner G. 1998. Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport* 9:315–19
- Bastian A, Schöner G, Riehle A. 2003. Preshaping and continuous evolution of motor cortical representations during movement preparation. *Eur. J. Neurosci.* 18:2047–58
- Batista AP, Santhanam G, Yu BM, Ryu SI, Afshar A, Shenoy KV. 2007. Reference frames for reach planning in macaque dorsal premotor cortex. *J. Neurophysiol.* 98:966–83
- Briggman KL, Abarbanel HD, Kristan WB Jr. 2005. Optical imaging of neuronal populations during decision-making. *Science* 307:896–901
- Briggman KL, Abarbanel HD, Kristan WB Jr. 2006. From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Curr. Opin. Neurobiol.* 16:135–44
- Broome BM, Jayaraman V, Laurent G. 2006. Encoding and decoding of overlapping odor sequences. *Neuron* 51:467–82
- Brown TG. 1914. On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. *J. Physiol.* 48:18–46
- Buesing L, Macke JH, Sahani M. 2012. Learning stable, regularised latent models of neural population dynamics. *Network* 23:24–47
- Bullock D, Grossberg S. 1988. Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychol. Rev.* 95:49–90
- Caminiti R, Johnson PB, Galli C, Ferraina S, 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.* 11:1182–97
- Churchland MM, Afshar A, Shenoy KV. 2006a. A central source of movement variability. *Neuron* 52:1085–96
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, et al. 2012. Neural population dynamics during reaching. *Nature* 487:51–56
- Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. 2010a. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* 68:387–400
- Churchland MM, Santhanam G, Shenoy KV. 2006b. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J. Neurophysiol.* 96:3130–46
- Churchland MM, Shenoy KV. 2007a. Delay of movement caused by disruption of cortical preparatory activity. *J. Neurophysiol.* 97:348–59
- Churchland MM, Shenoy KV. 2007b. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* 97:4235–57
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, et al. 2010b. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13:369–78
- Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. 2006c. Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* 26:3697–712
- Churchland MM, Yu BM, Sahani M, Shenoy KV. 2007. Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* 17:609–18
- Cisek P. 2006a. Integrated neural processes for defining potential actions and deciding between them: a computational model. *J. Neurosci.* 26:9761–70
- Cisek P. 2006b. Preparing for speed. Focus on “Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach.” *J. Neurophysiol.* 96:2842–43
- Collinger JL, Wodlinger B, Downey JE, Wang W, Tyler-Kabara EC, et al. 2013. High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* 381:557–64
- Crammond DJ, Kalaska JF. 2000. Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *J. Neurophysiol.* 84:986–1005
- Day BL, Rothwell JC, Thompson PD, Maertens de Noordhout A, Nakashima K, et al. 1989. Delay in the execution of voluntary movement by electrical or magnetic brain stimulation in intact man. Evidence for the storage of motor programs in the brain. *Brain* 112(Pt. 3):649–63
- Diester I, Kaufman MT, Mogri M, Pashaie R, Goo W, et al. 2011. An optogenetic toolbox designed for primates. *Nat. Neurosci.* 14:387–97

- Elsayed G, Kaufman MT, Ryu SI, Shenoy KV, Churchland MM, Cunningham JP. 2013. Characterization of dynamical activity in motor cortex. *Front. Neurosci. Conf.: Comput. Syst. Neurosci. (COSYNE)*. Abstr. III-61
- Erlhagen W, Bastian A, Jancke D, Riehle A, Schöner G. 1999. The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *J. Neurosci. Methods* 94:53–66
- Erlhagen W, Schöner G. 2002. Dynamic field theory of movement preparation. *Psychol. Rev.* 109:545–72
- Evarts EV. 1964. Temporal patterns of discharge of pyramidal tract neurons during sleep and waking in the monkey. *J. Neurophysiol.* 27:152–71
- Evarts EV. 1968. Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophysiol.* 31:14–27
- Faisal AA, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9:292–303
- Fetz EE. 1992. Are movement parameters recognizably coded in the activity of single neurons? *Behav. Brain Sci.* 15:679–90
- Fetz EE, Perlmutter SI, Prut Y. 2000. Functions of mammalian spinal interneurons during movement. *Curr. Opin. Neurobiol.* 10:699–707
- Foster JD, Nuyujukian P, Freifeld O, Ryu SI, Black MJ, Shenoy KV. 2012. A framework for relating neural activity to freely moving behavior. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* 2012:2736–39
- Fu QG, Flament D, Coltz JD, Ebner T. 1995. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J. Neurophysiol.* 73:836–54
- Ganguli S, Bisley JW, Roitman JD, Shadlen MN, Goldberg ME, Miller KD. 2008. One-dimensional dynamics of attention and decision making in LIP. *Neuron* 58:15–25
- Georgopoulos AP, Crutcher MD, Schwartz AB. 1989. Cognitive spatial-motor processes 3. Motor cortical prediction of movement direction during an instructed delay period. *Exp. Brain Res.* 75:183–94
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2:1527–37
- Georgopoulos AP, Schwartz AB, Kettner RE. 1986. Neuronal population coding of movement direction. *Science* 233:1416–19
- Gerits A, Farivar R, Rosen BR, Wald LL, Boyden ES, Vanduffel W. 2012. Optogenetically induced behavioral and functional network changes in primates. *Curr. Biol.* 22:1722–26
- Ghez C, Favilla M, Ghilardi MF, Gordon J, Bernejo R, Pullman S. 1997. Discrete and continuous planning of hand movements and isometric force trajectories. *Exp. Brain Res.* 115:217–33
- Ghez C, Hening W, Favilla M. 1989. Gradual specification of response amplitude in human tracking performance. *Brain Behav. Evol.* 33:69–74
- Gilja V, Chestek CA, Diester I, Henderson JM, Deisseroth K, Shenoy KV. 2011. Challenges and opportunities for next-generation intracortically based neural prostheses. *IEEE Trans. Biomed. Eng.* 58:1891–99
- Gilja V, Chestek CA, Nuyujukian P, Foster J, Shenoy KV. 2010. Autonomous head-mounted electrophysiology systems for freely behaving primates. *Curr. Opin. Neurobiol.* 20:676–86
- Gilja V, Nuyujukian P, Chestek C, Cunningham J, Yu B, et al. 2012. A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* 15:1752–57
- Graziano MSA. 2008. *The Intelligent Movement Machine: An Ethological Perspective on the Primate Motor System*, Vol. 224. Oxford/New York: Oxford Univ. Press
- Graziano MSA. 2011a. Cables vs. networks: old and new views on the function of motor cortex. *J. Physiol.* 589:2439
- Graziano MSA. 2011b. New insights into motor cortex. *Neuron* 71:387–88
- Green AM, Kalaska JF. 2011. Learning to move machines with the mind. *Trends Neurosci.* 34:61–75
- Grillner S. 2006. Biological pattern generation: the cellular and computational logic of networks in motion. *Neuron* 52:751–66
- Hanes DP, Schall JD. 1996. Neural control of voluntary movement initiation. *Science* 274:427–30
- Harvey CD, Coen P, Tank DW. 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68
- Hatsopoulos NG. 2005. Encoding in the motor cortex: Was Evarts right after all? Focus on “motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks.” *J. Neurophysiol.* 94:2261–62

- Hatsopoulos NG, Amit Y. 2012. Synthesizing complex movement fragment representations from motor cortical ensembles. *J. Physiol. Paris* 106:112–19
- Hatsopoulos NG, Donoghue JP. 2009. The science of neural interface systems. *Annu. Rev. Neurosci.* 32:249–66
- Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, et al. 2012. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485:372–75
- Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, et al. 2006. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164–71
- Hocherman S, Wise SP. 1991. Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Exp. Brain Res.* 83:285–302
- Kakei S, Hoffman DS, Strick PL. 1999. Muscle and movement representations in the primary motor cortex. *Science* 285:2136–39
- Kalaska JF. 2009. From intention to action: motor cortex and the control of reaching movements. *Adv. Exp. Med. Biol.* 629:139–78
- Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, et al. 2010. Roles of monkey premotor neuron classes in movement preparation and execution. *J. Neurophysiol.* 104:799–810
- Kaufman MT, Churchland MM, Shenoy KV. 2011. Cortical preparatory activity avoids causing movement by remaining in a muscle-neutral space. *Front. Neurosci. Conf.: Comput. Syst. Neurosci. (COSYNE)*. Abstr. II-61
- Keele SW. 1968. Movement control in skilled motor performance. *Psychol. Bull.* 70:387–403
- Kopell N, Ermentrout GB. 2002. Mechanisms of phase-locking and frequency control in pairs of coupled neural oscillators. In *Handbook on Dynamical Systems, Vol. 2: Toward Applications*, ed. B Fiedler, pp. 3–54. Philadelphia: Elsevier
- Kurata K. 1989. Distribution of neurons with set- and movement-related activity before hand and foot movements in the premotor cortex of rhesus monkeys. *Exp. Brain Res.* 77:245–56
- Kutas M, Donchin E. 1974. Studies of squeezing: handedness, responding hand, response force, and asymmetry of readiness potential. *Science* 186:545–48
- Lisberger SG, Sejnowski TJ. 1992. Motor learning in a recurrent network model based on the vestibulo-ocular reflex. *Nature* 360:159–61
- Machens CK, Romo R, Brody CD. 2010. Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *J. Neurosci.* 30:350–60
- Macke JH, Büsing L, Cunningham JP, Yu BM, Shenoy KV, and Sahani M. 2011. Empirical models of spiking in neural populations. See Shawe-Taylor et al. 2011, pp. 1350–58
- Mandelblat-Cerf Y, Paz R, Vaadia E. 2009. Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J. Neurosci.* 29:15053–62
- Manwani A, Steinmetz PN, Koch C. 2002. The impact of spike timing variability on the signal-encoding performance of neural spiking models. *Neural Comput.* 14:347–67
- Maynard EM, Hatsopoulos NG, Ojakangas CL, Acuna BD, Sanes JN, et al. 1999. Neuronal interactions improve cortical population coding of movement direction. *J. Neurosci.* 19:8083–93
- Mazor O, Laurent G. 2005. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–73
- McGuire LM, Sabes PN. 2011. Heterogeneous representations in the superior parietal lobule are common across reaches to visual and proprioceptive targets. *J. Neurosci.* 31:6661–73
- Miles FA, Evars EV. 1979. Concepts of motor organization. *Annu. Rev. Psychol.* 30:327–62
- Moran DW, Schwartz AB. 1999. Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* 82:2676–92
- Mountcastle VB, Lynch JC, Georgopoulos A, Sakata H, Acuna C. 1975. Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. *J. Neurophysiol.* 38:871–908
- Mullette-Gillman OA, Cohen YE, Groh JM. 2009. Motor-related signals in the intraparietal cortex encode locations in a hybrid, rather than eye-centered reference frame. *Cereb. Cortex* 19:1761–75
- Mussa-Ivaldi FA. 1988. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci. Lett.* 91:106–11
- O’Shea D, Goo W, Kalanithi P, Diester I, Ramakrishnan C, et al. 2013. Neural dynamics following optogenetic disruption of motor preparation. *Front. Neurosci. Conf.: Comput. Syst. Neurosci. (COSYNE)*. Abstr. III-60

- Pearce TM, Moran DW. 2012. Strategy-dependent encoding of planned arm movements in the dorsal premotor cortex. *Science* 337:984–88
- Pesaran B, Nelson MJ, Andersen RA. 2006. Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron* 51:125–34
- Petreska B, Yu BM, Cunningham JP, Santhanam G, Ryu SI, et al. 2011. Dynamical segmentation of single trials from population neural data. See Shawe-Taylor et al. 2011, pp. 756–64
- Rabinovich M, Huerta R, Laurent G. 2008. Transient dynamics for neural processing. *Science* 321:48–50
- Reimer J, Hatsopoulos NG. 2009. The problem of parametric neural coding in the motor system. *Adv. Exp. Med. Biol.* 629:243–59
- Requin J, Riehle A, Seal J. 1988. Neuronal activity and information processing in motor control: from stages to continuous flow. *Biol. Psychol.* 26:179–98
- Rickert J, Riehle A, Aertsen A, Rotter S, Nawrot MP. 2009. Dynamic encoding of movement direction in motor cortical neurons. *J. Neurosci.* 29:13870–82
- Riehle A, Grün S, Diesmann M, Aertsen A. 1997. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–53
- Riehle A, Requin J. 1989. Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *J. Neurophysiol.* 61:534–49
- Rokni U, Sompolinsky H. 2012. How the brain generates movement. *Neural Comput.* 24:289–331
- Rosenbaum DA. 1980. Human movement initiation: specification of arm, direction, and extent. *J. Exp. Psychol. Gen.* 109:444–74
- Sawaguchi T, Yamane I, Kubota K. 1996. Application of the GABA antagonist bicuculline to the premotor cortex reduces the ability to withhold reaching movements by well-trained monkeys in visually guided reaching task. *J. Neurophysiol.* 75:2150–56
- Schöner G. 2004. Dynamical systems approaches to understanding the generation of movement by the nervous system. In *Progress in Motor Control*, Vol. 3, ed. ML Latash, MF Levin. Champaign, IL: Hum. Kinet.
- Scott SH. 2000. Population vectors and motor cortex: neural coding or epiphenomenon? *Nat. Neurosci.* 3:307–8
- Scott SH. 2004. Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* 5:532–46
- Scott SH. 2008. Inconvenient truths about neural processing in primary motor cortex. *J. Physiol.* 586:1217–24
- Scott SH, Kalaska JF. 1995. Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures. *J. Neurophysiol.* 73:2563–67
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E. 1996. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J. Neurosci.* 16:752–68
- Shawe-Taylor J, Zemel RS, Bartlett P, Pereira F, Weinberger KQ, eds. 2011. *Advances in Neural Information Processing Systems (NIPS)*, Vol. 24. Red Hook, NY: Curran
- Shenoy KV, Kaufman MT, Sahani M, Churchland MM. 2011. A dynamical systems view of motor preparation: implications for neural prosthetic system design. *Prog. Brain Res.* 192:33–58
- Skavenski AA, Robinson DA. 1973. Role of abducens neurons in vestibuloocular reflex. *J. Neurophysiol.* 36:724–38
- Smith AC, Frank LM, Wirth S, Yanike M, Hu D, et al. 2004. Dynamic analysis of learning in behavioral experiments. *J. Neurosci.* 24:447–61
- Snyder LH, Batista AP, Andersen RA. 1997. Coding of intention in the posterior parietal cortex. *Nature* 386:167–70
- Stopfer M, Jayaraman V, Laurent G. 2003. Intensity versus identity coding in an olfactory system. *Neuron* 39:991–1004
- Sussillo D, Abbott LF. 2009. Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63:544–57
- Sussillo D, Churchland MM, Kaufman MT, Shenoy KV. 2013. A recurrent neural network that produces EMG from rhythmic dynamics. *Front. Neurosci. Conf.: Comput. Syst. Neurosci. (COSYNE)*. Abstr. III-67
- Szuts TA, Faddeyev V, Kachiguine S, Sher A, Grivich MV, et al. 2011. A wireless multi-channel neural amplifier for freely moving animals. *Nat. Neurosci.* 14:263–69

- Tanji J, Evarts EV. 1976. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J. Neurophysiol.* 39:1062–68
- Todorov E. 2000. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3:391–98
- Todorov E, Jordan MI. 2002. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5:1226–35
- Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, et al. 1995. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373:515–18
- van Beers RJ, Haggard P, Wolpert DM. 2004. The role of execution noise in movement variability. *J. Neurophysiol.* 91:1050–63
- Wang W, Chan SS, Heldman DA, Moran DW. 2010. Motor cortical representation of hand translation and rotation during reaching. *J. Neurosci.* 30:958–62
- Weinrich M, Wise SP. 1982. The premotor cortex of the monkey. *J. Neurosci.* 2:1329–45
- Wise SP. 1985. The primate premotor cortex: past, present, and preparatory. *Annu. Rev. Neurosci.* 8:1–19
- Wise SP, Weinrich M, Mauritz KH. 1986. Movement-related activity in the premotor cortex of rhesus macaques. *Prog. Brain Res.* 64:117–31
- Yu BM, Afshar A, Santhanam G, Ryu SI, Shenoy KV, Sahani M. 2006. Extracting dynamical structure embedded in neural activity. In *Advances in Neural Information Processing Systems (NIPS), Vol. 18*, ed. Y Weiss, B Scholkopf, J Platt, pp. 1545–52. Cambridge, MA: MIT Press
- Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. 2009. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* 102:614–35
- Yuste R, MacLean JN, Smith J, Lansner A. 2005. The cortex as a central pattern generator. *Nat. Rev. Neurosci.* 6:477–83