Published in final edited form as:

Prog Brain Res. 2011 ; 192: 33–58. doi:10.1016/B978-0-444-53355-5.00003-8.

A dynamical systems view of motor preparation: Implications for neural prosthetic system design

Krishna V. Shenoy^{†,‡,§,||,*}, Matthew T. Kaufman[‡], Maneesh Sahani^{†,¶}, and Mark M. Churchland^{†,‡}

[†]Department of Electrical Engineering, Stanford University, Stanford, California, USA

[‡]Neurosciences Program, Stanford University, Stanford, California, USA

[§]Department of Bioengineering, Stanford University, Stanford, California, USA

[¶]Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom

^{||}Department of Neurobiology, Stanford University, Stanford, California, USA

Abstract

Neural prosthetic systems aim to help disabled patients suffering from a range of neurological injuries and disease by using neural activity from the brain to directly control assistive devices. This approach in effect bypasses the dysfunctional neural circuitry, such as an injured spinal cord. To do so, neural prostheses depend critically on a scientific understanding of the neural activity that drives them. We review here several recent studies aimed at understanding the neural processes in premotor cortex that precede arm movements and lead to the initiation of movement. These studies were motivated by hypotheses and predictions conceived of within a dynamical systems perspective. This perspective concentrates on describing the neural state using as few degrees of freedom as possible and on inferring the rules that govern the motion of that neural state. Although quite general, this perspective has led to a number of specific predictions that have been addressed experimentally. It is hoped that the resulting picture of the dynamical role of preparatory and movement-related neural activity will be particularly helpful to the development of neural prostheses, which can themselves be viewed as dynamical systems under the control of the larger dynamical system to which they are attached.

Keywords

premotor cortex; motor cortex; motor preparation; state space; dynamical systems; single-trial analysis; neural prostheses; brain machine interface; brain computer interface

Introduction

It is difficult to appreciate just how central movement is to everyday life until this ability is lost due 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 that surround us, and we move our tongues and vocal cords to speak. Movement is not only central to these critical aspects of life, but also to self-image and psychological well-being. In fact, the fundamental reason that tetraplegics wish most for the restored use of their arms is to regain some degree of independence (Anderson, 2004).

Fortunately, it appears that a confluence of knowledge and technology from the fields of (1) systems motor neuroscience, (2) neuroengineering, and (3) electrical engineering and computer science may soon provide a new class of electronic medical systems (termed neural prosthetic systems, brain machine interfaces, or brain computer interfaces) aimed at increasing the quality of life for severely disabled patients.

First, basic neuroscience research across the past several decades has elucidated many of the fundamental principles underlying movement generation and control. A substantial body of knowledge regarding the cortical control of arm movements, particularly in rhesus macaques, now exists (e.g., Evarts, 1964; Georgopoulos et al., 1982, 1986; Schwartz, 1994; Tanji and Evarts, 1976). This literature is reviewed elsewhere (e.g., Kalaska, 2009; Kalaska et al., 1997; Scott, 2004; Wise, 1985). As discussed below, this understanding has been sufficient to help guide the design of first generation prosthetic systems. Yet continued focus on underlying neural mechanisms (in both monkeys and humans), how neural populations behave across timescales, and how neural populations participate in the ongoing control of movement, is essential for creating second generation prostheses capable of higher performance and a greater range of capabilities (e.g., Cunningham et al., 2010; Green and Kalaska, 2010; Truccolo et al., 2008, 2010).

Second, basic neuroengineering research has provided proof-of-concept demonstrations of neural prosthetic systems which translate the electrical activity (action potentials and local field potentials, LFPs) from populations of intracortically recorded neurons into control signals for guiding computer cursors, prosthetic arms, or stimulating the paralyzed musculature. More specifically, a series of designs and demonstrations across the past decade have produced compelling laboratory evidence that intracortical neural signals from rodents (e.g., Chapin et al., 1999), monkeys (e.g., Carmena et al., 2003; Chase et al., 2009; Fetz, 1969; Fraser et al., 2009; Ganguly and Carmena, 2009; Gilja et al., 2010b, c; Heliot et al., 2009; Humphrey et al., 1970; Isaacs et al., 2000; Jackson et al., 2006; Jarosiewicz et al., 2008; Moritz et al., 2008; Mulliken et al., 2008; Musallam et al., 2004; Nuyujukian et al., 2010; Santhanam et al., 2006; Serruya et al., 2002; Shenoy et al., 2003; Taylor et al., 2002; Velliste et al., 2008; Wessberg et al., 2000; Wu et al., 2004), and humans (e.g., Hochberg et al., 2006; Kim et al., 2008) can control prosthetic devices that may provide meaningful quality of life improvement to paralyzed patients. This literature is reviewed elsewhere (e.g., Andersen et al., 2010; Millan and Carmena, 2010; Donoghue, 2008; Donoghue et al., 2007; Fetz, 2007; Hatsopoulos and Donoghue, 2009; Linderman et al., 2008; Nicolelis and Lebedev, 2009; Ryu and Shenoy, 2009; Scherberger, 2009; Schwartz, 2007).

Finally, the semiconductor electronics, optoelectronic telecommunications, micro-electromechanical systems (MEMS), and information technology revolutions over the past four decades have produced extraordinary and relevant technologies. These include low-power and high computational-density circuits and systems, lowpower wireless telemetric systems, advanced light sources and imaging modalities, and small sensor systems that are capable of running sophisticated signal processing algorithms. These technologies have progressed extremely quickly, as described by Moore's Law, and have been leveraged and adapted to create new neurotechnologies for basic neuroscience and neuroengineering applications such as neural prosthetic systems. It is now possible to record from hundreds of neurons simultaneously with bio-MEMS electrode arrays (e.g., Chestek et al., 2009a, 2011; Jackson and Fetz, 2007; Mavoori et al., 2005; Santhanam et al., 2007), filter and "spike sort" all channels in real time (e.g., O'Driscoll et al., 2006; Santhanam et al., 2004, 2006), "decode" the intended arm movement with advanced algorithms (e.g., Achtman et al., 2007; Cunningham et al., 2008; Kemere et al., 2004, 2008; Santhanam et al., 2009; Wu et al., 2006; Yu et al., 2007, 2010), wirelessly telemeter the resulting prosthetic arm control signals with just a few tens of milliwatts of power (e.g., Chestek et al., 2009b; Gilja et al., 2010a;

Harrison et al., 2007, 2009), and soon, this will likely all be possible in fully implantable systems (e.g., Borton et al., 2009; Harrison, 2008; Nurmikko et al., 2010).

While these laboratory proof-of-concept systems and initial FDA phase-I clinical trials are encouraging (e.g., Hochberg, 2008; Hochberg and Taylor, 2007), several barriers remain. If these barriers are unaddressed, they could substantially limit the prospect of intracortically based neural prosthetic systems having a broad and important clinical impact. We recently reviewed what we consider to be three of the most important neuroengineering, bioengineering, electrical engineering, and computer science challenges and opportunities for intracortically based neural prostheses (Gilja et al., 2011). We review here what we consider to be one of the most central and important basic systems-level motor neuroscience questions. The knowledge gained while investigating this question should directly advance our ability to design high-performance neural prostheses. The central question we have been asking is: what are the neural processes that precede movement and lead to the initiation of movement? Neural prostheses will benefit from a deeper and more comprehensive understanding of the neural activity upon which they are based (Green and Kalaska, 2010). This includes activity during both movement preparation and movement generation. We need to understand both because prostheses use both (e.g., Yu et al., 2010), and because the two are presumably causally linked and likely impossible to understand fully if studied in isolation (discussed further in the final section, and Fig. 12). Prostheses should thus benefit from having a firm scientific understanding of how preparatory activity relates to upcoming arm movements, and how this preparatory activity evolves on a millisecond timescale. These are the questions and topics discussed in this review.

Preparing to move the arm

Why should one prepare and then move, as opposed to starting the movement as soon as possible? In some cases, it is critical to move right away, such as when withdrawing a hand from a flame. Animals have evolved low-latency circuits to help in these cases and these circuits underlie a wide range of reflexive movements. However, animals have also evolved circuits to enable voluntary movements which are intentional and purposeful. Voluntary movements require the ability to change, refine, and suppress possible actions before they are actually executed. A simple example is how we swat a fly. One approach would be to see a fly and start moving right away. Unless the nervous system can execute perfectly, this is unlikely to be a good strategy, and if the initial movement is not successful, the fly is likely to depart before a correction can be made. It would thus be beneficial to take slightly more time to initiate the movement, assuming that, in doing so, greater accuracy can be gained. Presumably, we use this slight addition of time to create and refine movement plans until the moment is right and then we initiate the movement. It is this form of deliberate, goal-driven movement that we seek to better understand, both neurally and behaviorally, both out of scientific curiosity and because it could lead to superior prosthetic designs.

There is indeed evidence that voluntary movements are prepared before they are initiated (e.g., Day et al., 1989; Ghez et al., 1997; Keele, 1968; Kutas and Donchin, 1974; Riehle and Requin, 1993; Rosenbaum, 1980; Wise, 1985). An important line of evidence comes from “instructed-delay tasks” where a temporal delay separates an instruction stimulus from a subsequent “go” cue. Figure 1 illustrates the experimental arrangement and task timing, along with example hand position and electromyographic (EMG) measurements. This task is widely employed and is the behavioral task used in the recent studies reviewed here.

At the behavioral level, reaction times (RTs), defined as the time from the go cue until movement onset, are shorter after an instructed-delay period. Figure 2 illustrates how RT decreases and then plateaus as a function of delay period. This RT reduction with delay,

largely occurring during the first 200 ms, suggests that some timeconsuming preparatory process is given a head start by the delay (e.g., Crammond and Kalaska, 2000; Riehle and Requin, 1989; Rosenbaum, 1980). It is straightforward to interpret the importance of this head start on preparation in the context of the fly swatting example offered above. There the goal was not to move instantaneously as soon as the fly landed or was seen. Instead, the goal was to move swiftly and accurately, at a particular speed and along a particular path that perhaps approaches from behind, and to be able to start that movement as quickly as possible when it is decided that the time is right. Thus, a good strategy is to prepare the desired movement as soon as possible, so that one is ready to move as soon as possible when called upon to do so.

The ability to prepare a movement ahead of time is presumably related to the preparatory activity that is widespread in cortex and subcortical structures. Neurons in a number of cortical areas including dorsal premotor cortex (PMd) and primary motor cortex (M1) show changes in activity during the delay period (e.g., Crammond and Kalaska, 2000; Godschalk et al., 1985; Kalaska et al., 1997; Kurata, 1989; Messier and Kalaska, 2000; Riehle and Requin, 1989; Snyder et al., 1997; Tanji and Evarts, 1976; Weinrich et al., 1984). Figure 3 shows four example PMd neurons. While it is typical for the average action potential emission (firing) rate during the delay period to change following target onset, the temporal structure is widely varying across cells: some increase their firing rate, some decrease, some arrive at an approximate plateau level, while others undulate.

This variety of neural responses stands in stark contrast to the simple monotonic decline in behavioral RT as shown in Fig. 2. The central question is, therefore, how does neural activity in the first 200–300 ms of the delay period relate to the decrease in RT? Asked in the context of the fly swatting example, what does this neural activity need to accomplish during the delay so that we are maximally poised to generate the planned movement and, after initiating the movement, successfully hit the fly?

Optimal subspace hypothesis

We have been investigating this question using a dynamical systems perspective (e.g., Briggman et al., 2005; Churchland et al., 2007; Stopfer et al., 2003). What this means in essence is that we wish to understand (1) how the activity of a neural population evolves and achieves the needed preparatory state, (2) how this preparatory state impacts the subsequent arm movement, and (3) what the underlying dynamics (rules) of the neural circuitry are (Churchland et al., 2007; Yu et al., 2006, 2009). We start with as simple an assumption as possible: the arm movement made ($M(t)$) depends upon preparatory activity (P) at the time movement activity begins to be generated (t_0). In other words, $M(t)$ depends on $P(t_0)$.

It is important to note that there are likely sources of variability that impact $M(t)$ but are not accounted for in $P(t_0)$, such as downstream variability in the state of the spinal cord or muscles. Thus, to be strictly true, a noise source should be included, or $P(t_0)$ would need to be the initial state of the entire animal. However, for the moment, we avoid this issue and simply concentrate on the hypothesis that the movement you make is in large part a function of the plan that was present just before movement began. Also, note that the above conception does not rule out a strong (or even dominant) role for feedback. Such feedback could be part of the causal mechanism by which the plan produces the movement.

The central implication of our assumption that $M(t)$ depends on $P(t_0)$ is that motor preparation may be the act of optimizing preparatory activity (i.e., bringing P to the state needed at t_0) so that the generated movement has the desired properties. In the case of monkeys performing reaching movements, the desired movement can be defined as a reach that is accurate enough to result in reward. Consider the space of all possible preparatory

states (all possible P s). For a given reach, there is presumably some small subregion of space containing those values of P that are adequate to produce a successful reach. Although the response of each neuron (i.e., tuning) may not be easily parameterized, there is nonetheless a smooth relationship between firing rate and movement. Therefore, the small subregion of space is conceived of as being contiguous.

Figure 4 illustrates this idea. We conceive of all possible preparatory states as forming a space, with the firing rate of each neuron contributing an axis. Each possible state—each vector of possible firing rates—is then a point in this space. For a given reach (e.g., rightwards), there will be some subset of states (gray region in Fig. 4, referred to as the optimal subspace) that will result in a successful reach that garners a reward. Under this optimal subspace hypothesis, the central goal of motor preparation is to bring the neural state within this subspace before the movement is triggered. This may occur in different ways on different trials (trial 1 and trial 2 in Fig. 4). This framework, though rather general, has provided us with a number of specific and testable predictions, which we review below.

Before doing so, it is worth considering that an almost-trivial prediction of the optimal subspace hypothesis is that different movements require different initial states. If preparatory activity has a strong role in determining movement, then making different movements will require different patterns of preparatory activity. The overall neural state, and thus the state of individual neurons, should therefore vary with different movements. This is of course consistent with the observation that preparatory activity is tuned for reach parameters such as direction and distance (e.g., Messier and Kalaska, 2000). In fact, under the optimal subspace hypothesis, neural activity should appear tuned for essentially every controllable aspect of the upcoming reach (a prediction we will return to shortly).

As a brief aside on the topic of tuning, we note that one could conceive of each axis in Fig. 4 as capturing not the activity of a single neuron, but rather the activity of a population of neurons that are all tuned for the same thing. Thus, the three axes might capture, respectively, the average activity of neurons tuned for direction, distance, and speed. If so, the preparatory state could be thought of as an explicit representation of direction, distance, and speed. However, it has been argued that few individual neurons appear tuned for reach parameters in the straightforward and invariant way that one might hope (e.g., Churchland et al., 2006b; Churchland and Shenoy, 2007b; Cisek, 2006; Fetz, 1992; Scott, 2004, 2008; Todorov, 2000). The optimal subspace hypothesis is largely agnostic to this debate. So long as there is a systematic relationship between preparatory activity and movement, the optimal subspace conception remains viable. Put another way, the space illustrated in Fig. 4 could have axes that capture well-defined parameters, but it need not, and there are reasons to suspect that it does not.

A related and critical point is that the space in which neural activity evolves is certainly larger than the three dimensions illustrated in Fig. 4. Movements vary from one another in more than three different ways. Similarly, neural activity varies across movements in more than three different ways (Churchland and Shenoy, 2007b). Thus, care should be taken when gleaning intuition from illustrations such as that in Fig. 4, to keep in mind that what is illustrated is a projection of a larger and richer space (Churchland et al., 2007; Yu et al., 2009).

We now review a number of specific and testable predictions of the optimal subspace hypothesis.

Prediction 1: Reach-speed modulation

Figure 5a illustrates in state space our first prediction under the optimal subspace hypothesis: preparatory activity should covary with other meaningful aspects of movement, including peak reach speed. Confirming this would be consistent with our assumption above, whereas failing to find this would be consistent with preparatory activity having a higher-level, perhaps more sensory role reflecting the target location but not the more detailed aspects of movement.

To test this prediction, we trained monkeys to reach to targets in a variant of the instructed-delay task. Reaches must be made somewhat faster (1.5 m/s peak speed) when the target was red and somewhat slower (1.0 m/s peak speed) when the target was green (Churchland et al., 2006b). All other movement metrics such as reach path remained similar. Delay-period activity was substantially different ahead of fast and slow arm movements to the same target location. Figure 5b shows the average response of an example neuron, ahead of reaches to a particular target, where the delay-period activity was greater ahead of fast reaches (red) than ahead of slow reaches (green). Figure 5c and d show two more example neurons where this difference in preparatory activity ahead of fast (red) and slow (green) reaches is emphasized by collapsing across all reach target locations. Some neurons had higher average firing rates ahead of fast movements (Fig. 5c), while other neurons had higher average rates ahead of slow movements (Fig. 5d).

In sum, prediction 1 as illustrated in Fig. 5a appears to be correct.

Prediction 2: Reach-speed (trial-by-trial) correlation

Figure 5a also illustrates in state space our second prediction under the optimal subspace hypothesis: preparatory activity should correlate, on a trial-by-trial basis, with the peak reach speed. Our assumption that $M(t)$ depends on $P(t_0)$ predicts that even a slightly different $P(t_0)$ value should lead to a different $M(t)$. If $P(t_0)$ reflects the result of a difficult optimization, then variability in $P(t_0)$ is likely. Therefore, it should be possible to observe a trial-by-trial correlation between $P(t_0)$ and movement metrics $M(t)$.

To test this prediction, we again employed the reach-speed variant of the instructed-delay task. We found trial-by-trial correlations between the firing rate of individual neurons before the go cue and peak reach speed (Churchland et al., 2006a). Consider the instructed-fast condition (red) in Fig. 5c. The horizontal spread of points (one point per trial) reflects the trial-to-trial variance in peak reach speed. The vertical spread of points largely reflects the trial-to-trial variance in estimated firing rate, which is the inevitable result of it being difficult to assess the firing rate of a single neuron on a single trial from a handful of stochastically occurring spikes. Nevertheless a statistically significant correlation was found for most neurons and for both instructed speeds.

Importantly, the state-space illustration (Fig. 5a) further predicts that within the instructed-fast condition, for example, a trial with a slightly slower peak reach speed should have a preparatory state slightly closer to those found in the instructed-slow condition. In other words, if we assume that movement parameters are mapped smoothly from firing rate, the slope of the within-condition correlation (black lines) should agree with the slope of the across-condition mean line fit (gray line) both when the instructed-fast condition had a higher average firing rate (Fig. 5c) and when it had a lower average firing rate (Fig. 5d). We found this to be the case in the majority of neurons (Churchland et al., 2006a).

In sum, prediction 2 as illustrated in Fig. 5a appears to be correct.

Prediction 3: Across-trial firing-rate variance (Fano factor) reduces through time

Figure 6a illustrates in state space our third prediction under the optimal subspace hypothesis: preparatory activity should become, through time, quite accurate and therefore quite similar across trials. Before the target appears, “baseline” neural activity can be somewhat different from trial to trial, leading to some amount of across-trial firing-rate variance (black circles in panel labeled “before target onset” in Fig. 6a, top). After target onset, and for the coming 200–300 ms, preparatory activity on each trial is nominally being optimized and brought to reside within the optimal subspace. The optimal subspace is presumably rather restricted by virtue of the behavioral task constraints and thus should have less across-trial firing-rate variance (black circles within the optimal subspace shaded gray, in the panel labeled “200 ms after target onset” in Fig. 6a, bottom).

To test this prediction, we analyzed data from an instructed-delay task using the Fano factor: the across-trial spike-count variance divided by the mean (Churchland et al., 2006c, 2007, 2010c). Normalization, and an additional set of controls, is necessary to ensure that the measured changes in variance are not simply due to the well-known scaling of spike-count variance with spike-count mean (as happens, e.g., for a Poisson process; Churchland et al., 2007, 2010c; Rickert et al., 2009). As shown in Figure 6b, we found that the Fano factor declines over the course of approximately 200 ms and then approximately plateaus (Churchland et al., 2006c). This is somewhat remarkable, as it so closely resembles the decline and plateau seen in the behavioral curves (RT versus delay, Fig. 2).

In sum, prediction 3 as illustrated in Fig. 6a appears to be correct. Moreover, it appears that the across-trial firing-rate variance (as measured by the Fano factor) parallels the reduction in RT: both drop over the course of approximately 200 ms and then hold at that level. This possibility is explored below, as predictions 4-I and 4-II.

As a brief aside, it could be the case that this reduction in across-trial firing-rate variance is principally a motor phenomenon and is specific to the preparation of arm movements. However, we found that this same general structure of a reduction in across-trial firing-rate variance following a stimulus onset is present across much if not all of cerebral cortex (Churchland et al., 2010c). Figure 7 shows a substantial reduction in Fano factor following stimulus onset in numerous areas, across all four cortical lobes, and in a variety of behaviors. This reduction seems to be a general property of the nervous system responding to an input, much as the mean (across-trial) firing-rate changing is a general property of cortical neurons. This reduction in across-trial firing-rate variance in each area may be correlated with the relevant functions performed therein (e.g., sensation, cognition, behavior), again just as the mean firing rate is well known to correlate with the function of each area.

Predictions 4-I and 4-II: Lower Fano factor and lower RTs

Figure 6c illustrates the first part of our fourth prediction (prediction 4-I) under the optimal subspace hypothesis: the lower the across-trial firing-rate variance at the time of the go cue, the lower too should be the RT. Having seen the similarity between how the Fano factor descends and holds as a function of delay duration (Fig. 6b), and how RT descends and holds as a function of RT (Fig. 2), it is natural to predict that there should exist a positive correlation between Fano factor and RT. For example, one expects that short delays should lead to high Fano factors and high RTs, while long delays should lead to low Fano factors and low RTs.

To test this prediction, we analyzed short delay duration trials from the instructed-delay task. Figure 2 shows representative RT data from monkey G when 30, 130, and 230 ms delay durations were used. Figure 6c shows the Fano factor at the three critical times: 30, 130, and 230 ms after target onset. Figure 6d shows RT data plotted against Fano factor data, from the same trials in Monkey G, and a clear relationship is seen. The lower the across-trial firing-rate variance at the time of the go cue (as measured by the Fano factor), the lower the RT (Churchland et al., 2006c).

Figure 8a illustrates in state space the second part of our fourth prediction (prediction 4-II) under the optimal subspace hypothesis: the lower the across-trial firing-rate variance at the time of the go cue, the lower too should be the RT, even in long delay-duration trials where sufficient time has elapsed for “complete” motor preparation to result. For long delay durations (e.g., >200–300 ms), the Fano factor has nominally plateaued, as has the RT, at a low level. But, as depicted in Fig. 8a, there could still remain some variability. On trials that “wander outside” the optimal subspace (red circles), some additional time (i.e., increased RT) should be required to complete preparatory optimization following a go cue. In contrast, on trials where the preparatory state is within the optimal subspace (green dots) and therefore motor preparation is complete and ready for execution, movement can begin with a minimum of latency following the go cue (i.e., low RT).

To test this prediction, we started with all trials with 200 ms or longer delay durations, across 7 days of experiments using a 96-channel electrode array. This helped assure sufficient data. Second, we sorted trials according to whether the RT was shorter than or longer than the median RT. Third, we calculated the across-trial firing-rate variance (Fano factor) for the half of trials with shorter than median RT, and the same for the half of trials with longer than median RT. We did so for times ranging from 200 ms before the go cue until 200 ms after the go cue in order to assess the robustness of the result. Figure 8b plots the Fano factor curve for shorter than median trials (green curve) and longer than median trials (red curve). These curves are statistically significantly different (not shown), and as predicted, the lower across-trial firing-rate variance trials (green curve) are associated with lower RTs (Churchland et al., 2006c).

In sum, prediction 4-II as illustrated in Fig. 8a appears to be correct. When combined with the experiments and results associated with prediction 4-I, it appears clear that there is a close relationship between the across-trial firing-rate variability at the time of the go cue and the resulting RT. Recently, similar results have been found in area V4 ahead of saccadic eye movements, suggesting that this relationship is not limited to the arm movement system alone (Steinmetz and Moore, 2010).

Prediction 5: Perturbing neural activity increases RT

The inset of Fig. 9 illustrates in state space the fifth prediction under the optimal subspace hypothesis: perturbing the preparatory state out of the optimal subspace should result in an increased RT. But it should not reduce movement accuracy.

The first four predictions of the optimal subspace hypothesis were correlative, and their affirmation provides important evidence supporting the optimal subspace hypothesis. The inset of Fig. 9 illustrates a causal prediction, wherein a preparatory state within the optimal subspace is deliberately perturbed (curly line with displaced preparatory state, black circle) and it is predicted that the RT should increase. This follows from the reasoning that if the goal of motor preparation is to help make accurate movements, then the brain must somehow be able to monitor preparatory activity and determine when it is accurate enough to initiate movement. If preparatory activity were optimized and within the optimal subspace, but were then perturbed away from the optimal subspace, the brain should wait for

the plan to reoptimize to the optimal subspace (i.e., recover) before initiating movement (red dashed arrow labeled “reoptimization”). Importantly, after taking time to reoptimize preparatory activity, the resulting movements should be as accurate as on nonstimulated trials.

To test this prediction, we delivered subthreshold electrical microstimulation to PMd on a subset of trials and did so at various times relative to the go cue (Churchland and Shenoy, 2007a). Figure 9 plots experimental results from all (30) stimulation sites in PMd in one monkey. RTs are increased when microstimulation is delivered around the time of the go cue. This is seen as a rightward shift in the red hand-speed curve (stimulation around time of go cue, as indicated by the red bar) relative to the black hand-speed curve (no stimulation). This is consistent with time having been consumed (increased RT) to reoptimize preparatory activity. Importantly, aside from delaying the onset of movement, all other movement metrics were extremely similar to the nonstimulated trials (consistent with prediction 5). Note that the red and green averaged curves in Fig. 9 have lower peak hand speed due to staggered RTs, but individual trials do achieve the same, higher peak hand speed (see Churchland and Shenoy, 2007a for details).

As it is critical to establish effect specificity when conducting causal perturbation experiments, we performed several additional control experiments (Churchland and Shenoy, 2007a). Four are briefly summarized here. First, we found that stimulating well before the go cue (Fig. 9, green bar) had little impact on the RT. This can be seen in Fig. 9 by noting that the green curve largely overlaps with the black curve. This result is consistent with there being sufficient time for reoptimization to occur before the go cue appears. This is an important temporal control and indicates that the effect of subthreshold microstimulation exerts its influence just when the preparatory state is most needed (consistent with prediction 5). Second, stimulating on zero delay-duration trials where there was presumably no optimized preparatory activity present to perturb did not alter RT. This is an important control as it confirms the necessity of there first existing a preparatory state near the optimal subspace (consistent with prediction 5). Third, stimulating in M1 where there is relatively less preparatory activity resulted in little RT increase. The importance of this control is twofold. (i) It confirms that perturbing motor preparation is easier in an area where preparatory activity is prevalent (PMd) than in an area where it is not (M1). (ii) It confirms area specificity. M1 is just a few millimeters from PMd, but the effect is dramatically reduced and thereby helps assure that subthreshold microstimulation is not just a generalized distraction. Both findings are consistent with prediction 5. Fourth and finally, the effect of microstimulation was specific to arm movements and produced little increase in saccadic eye movement RT. This is an important control for the possibility that microstimulation is altering attention, which should impact both effectors equivalently.

In sum, prediction 5 as illustrated in Fig. 9 is borne out. This contributes causal supportive evidence for the optimal subspace hypothesis, which complements predictions 1–4 as well as 6 and 7 (below).

Prediction 6: Single-trial neural trajectories

The conceptual sketch in Fig. 4 illustrates in state space the sixth prediction under the optimal subspace hypothesis: it should be possible to construct single-trial state-space neural trajectories and use them to directly confirm that across-trial firing-rate variability decreases through time.

To test this prediction, we must begin by measuring many neurons simultaneously. This is essential as we seek an accurate estimate of the preparatory state on each individual trial and on a fine time scale. Both require data from many neurons, instead of the more traditional

technique of trial averaging, in order to mathematically reduce the deleterious effects of spiking noise (Churchland et al., 2007; Yu et al., 2009). These measurements can be made with electrode arrays, which have been developed substantially as part of neural prosthesis research. The analyses can be performed using modern dimensionality reduction and visualization methods such as Gaussian Process Factor Analysis (GPFA; Yu et al., 2009). Dimensionality reduction is needed for two reasons. First, reducing the dimensionality of the data from its original ~100 D space (e.g., 100 neurons measured simultaneously constitutes a 100 D space) down to 10–15 D appears to be possible without significant loss of information and has the benefit of effectively denoising the data (Yu et al., 2009). This can be thought of as essentially performing a weighted average to combine the responses of neurons that share some important aspect of their response. Second, while reducing the dimensionality further (below 10–15 D) does result in a loss of information, it can be quite useful for visualization purposes. This is because the two or three dimensions used in drawings can be the two or three dimensions that capture the greatest variance in the data, and the resulting plots are still sufficient to spur on hypotheses and predictions as described above.

Figure 10a shows multiple single-trial neural trajectories in a 2D state space created with GPFA (Churchland et al., 2010c; Yu et al., 2009). This is the first time that true single-trial neural trajectories (gray lines in Fig. 10a), as opposed to the cartoon depictions in Fig. 4, are plotted in this review. It is reassuring to see in Fig. 10a that the scatter in across-trial preparatory states at each point in time (black dots) reduces as the trial progresses. As the trial progresses from before target onset (100 ms pretarget) to just after target onset when the preparatory state is evolving toward the optimal subspace (100 ms post-target), and on to when the neural state on each trial is presumably within the optimal subspace (200 ms post-target), the variance (scatter) of the preparatory state reduces. This is consistent with the results presented above, inferred with Fano factor analyses.

Figure 10b again shows multiple single-trial neural trajectories in a 2D state space created with GPFA but now goes on to show data until the time of movement onset (Churchland et al., 2010c). This reveals for the first time that neural trajectories (gray lines) follow a largely stereotyped path through state space, after the initial convergence following target onset. They start in the baseline pretarget state (blue circles), progress, and slow-down (if an extended delay period) in the optimal subspace until the time of the go cue (green), and then arch around and arrive at a small region where the arm movement is first detected (black). Highlighted in red is one outlier trial which had a substantially longer RT than typical. With single-trial visualization of even (entirely) internal neural processing, it is now possible to ask, for the first time, what the reason might be. On this trial, the preparatory process appears to have completed normally. The green circle is within the (presumed) optimal subspace and surrounded by other trials that had normal RTs. While we cannot conclude what the cause was from this data alone, we appear to be able to rule out incomplete motor preparation. Figure 10c plots data from a different data set. Again, one trial with a particularly long RT is highlighted in red. This trial's neural trajectory undergoes an entire loop between the go cue and movement onset.

In sum, prediction 6 as illustrated in Fig. 4 is borne out. It is possible to construct single-trial neural trajectories and use these trajectories to directly see key features that can only be inferred less directly with single-neuron recordings. It also appears to now be possible to begin to investigate the reasons for outlier and other types of unique trials. Intriguingly, it should now also be possible to design experiments aimed at creating inherently single-trial phenomenon such as single-trial decision making, which could shed considerable insight on internal cognitive processing and neural dynamics (Kalmar et al., 2010; Rivera-Alvidrez et al., 2008). Further, as tasks become more complex (e.g., Churchland et al., 2008) and

naturalistic (e.g., Chestek et al., 2009b; Gilja et al., 2010a; Jackson et al., 2006; Santhanam et al., 2007), both behavior and the preceding neural processes will likely become less stereotyped and may therefore often require single-trial analyses.

Prediction 7: Farther and faster along loop reduces RT

We have posited that the preparatory state has a large impact on the subsequent movement. We have also seen several predictions that stem from the optimal subspace hypothesis along with evidence supporting these predictions. It does seem to be the case that the preparatory state at the time of the go cue has a substantial influence on the subsequent movement. But why should this be?

One possibility is illustrated in Fig. 11a. It could be that the preparatory state at the time of the go cue (green circle) acts as the initial state of a subsequent dynamical system that serves to generate muscle activity and create movement (green and blue arrows; Churchland et al., 2010a). As such, it is important that the preparatory state be within the optimal subspace in order to help create the desired movement. Thus some region or regions of the brain appear to monitor and wait for this to be true before “pulling the trigger” to initiate movement. After the movement trigger has been pulled, if the preparatory state happens to be farther along the “loop” and moving in the standard direction around the loop (see arrows), then RT may be further lowered.

Figure 11b illustrates in state space the seventh prediction under the optimal subspace hypothesis: the farther the preparatory state is along the loop when the movement trigger is pulled, and the faster it is moving along the loop in the standard direction (not shown in Fig. 11b), the shorter the RT should be. Figure 11b depicts the “loop” structure seen in Fig. 10b, where most individual trials follow a stereotyped path in state space. The two single-trial neural trajectories in Fig. 11b exit the optimal subspace in a particular direction, as was the case for all single-trial neural trajectories (gray lines) in Fig. 10b and c. Prediction 7 states that a trial like trial 1 in Fig. 11b should have a shorter RT than trial 2, because (i) the preparatory state at the time of the go cue is within the optimal subspace for both trials (and thus the movement trigger is presumably pulled at the same time) and (ii) the preparatory state for trial 1 (green circle) is nearer the exit-edge of the optimal subspace and is thus farther along the stereotyped path that it will need to take to generate movement. If the preparatory state also happened to be moving along the stereotyped path in the standard direction, as opposed to not moving or moving in the opposite direction, then the RT ought to be shorter still.

To test this hypothesis, we correlated, on a trial-by-trial basis, how far along the loop the preparatory state was (at the time of the go cue) with RT. As predicted, we found a statistically significant negative correlation, and primarily in just the exit-edge direction (Afshar et al., 2011). Also as predicted, we found a statistically significant correlation between the direction of movement of the preparatory state at the time of the go cue and RT: preparatory states that were moving in the direction of (subsequent) loop travel had lower RTs than comparably positioned preparatory states moving in the opposite direction (Afshar et al., 2011).

This appears to suggest that a trial with preparatory activity at the time of the go cue (green circle) that is (i) within the optimal subspace and (ii) farther along, and moving in, the standard “loop” direction is in some sense “doubly advantaged” because it is both (i) a well-optimized preparatory state (i.e., within the optimal subspace so the movement trigger can be pulled straight away) and (ii) fortuitously positioned and already headed along the path it will need to take to generate movement. In sum, prediction 7 as illustrated in Fig. 11b is borne out.

Future directions

The above predictions were derived from a dynamical systems perspective, and to some degree their confirmation argues for that perspective. Yet the most central questions remain largely unaddressed. What is the nature of the relevant dynamics (e.g., Yu et al., 2006)? Do they relate to the dynamics of movement-generating circuits in simpler organisms (e.g., Grillner, 2006; Kristan and Calabrese, 1976)? How and why do dynamics change as a function of overall state (e.g., resting vs. planning vs. moving)? What is the nature of the circuitry, both local and feedback, that produces those dynamics? Answering such questions will likely depend on progress in three domains: (1) the ability to better perturb and probe dynamics, (2) the ability to resolve dynamical structure in neural data, and (3) the ability to relate the recorded “neural trajectories” to externally measurable parameters such as muscle activity and hand movement. We consider these in turn.

First, when reverse-engineering any system, the ability to perturb and observe is critical. As described above, we used intracortical electrical microstimulation to ask how a perturbation of neural activity influenced RT. Pharmacological manipulations are also possible and would offer cell-type specific manipulation of the system. Recent advances in optogenetic stimulation of neurons in rhesus monkeys may also provide important new insights due to the ability to excite and inhibit neurons in a cell-type specific manner (unlike electrical microstimulation), do so on a millisecond timescale (unlike pharmacological manipulations), and not interfere with simultaneous electrical recordings (unlike electrical microstimulation; Diester et al., 2011; Han et al., 2009). This could enable the direct visualization of neural trajectories throughout trials where the neural state is optically perturbed at various times. Recall the Fig. 9 inset, where the curly black line and dashed red line could only illustrate our speculation about how the neural trajectory might evolve during and directly following stimulation, because electrical microstimulation interferes with electrical array recordings. Also, while permanently altering the underlying neural circuitry transgenically is not currently possible in rhesus monkeys, it is possible to reversibly lesion brain regions, cell-types, and neuronal projections pharmacologically and optogenetically. Understanding how the subsequent alteration of neural trajectories relates to altered behavior would deepen our understanding of motor preparation and generation, and the role of specific cells and connectivity (Kaufman et al., 2009, 2010a; Lerchner et al., 2011).

Second, any real understanding of dynamics will hinge upon the ability to go beyond merely plotting state-space trajectories. One wishes to take those seemingly complex neural trajectories, which evolve in many dimensions, and infer meaningful and parsimonious underlying dynamics (Yu et al., 2006). Indeed, if this cannot be done—if the proposed dynamics are not simpler than the data they seek to explain—then the dynamical systems perspective may have little to offer. Fortunately, it appears that simple dynamics may well be able to explain a considerable amount of the structure of the data (Churchland et al., 2010b, 2011; Cunningham et al., 2011; Macke et al., 2011; Petreska et al., 2011), but further progress in this realm will depend on the continued development of analysis methods that can capture how one neural state leads to the next.

Finally, while the observed state-space trajectories often appear rather abstract, they must exist for a concrete purpose: producing movement. That is, there must be some direct and causal relationship between the neural trajectory and some externally measurable quantity such as muscle activation or arm kinematics. Historically, the relationship between movement- period activity and the parameters of movement has been contentious (e.g., Kalaska, 2009). The dynamical systems perspective will not on its own resolve this debate, but there are a number of contributions it can make. The dimensionality- reduction techniques that produce the statespace trajectories force the experimenter to focus on those

patterns that are most strongly present in the data (e.g., Rivera-Alvidrez et al., 2009, 2010a, b). Also, the relatively high dimensionality of the state space makes it clear that not all aspects of neural activity can or should be related directly to external factors: some dimensions may be important to the overall dynamics but may not exert any direct influence on the periphery (Kaufman et al., 2010b, 2011).

Progress in the above domains should also increase the breadth of questions addressable under the dynamical systems perspective. Already it has been possible to ask whether neural variability decreases during learning (Mandelblat-Cerf et al., 2009). More generally, we wish to know what the “state-space” correlate of motor learning might look like. For example, does the location of the optimal subspace change following learning? Or does learning change the dynamics that determine the trajectory away from that planned state? Of course, one suspects that both such mechanisms might be at play, perhaps depending on the type of learning (e.g., the former strategy might be more rapid but less flexible). The further development of critical tools, as described above, could open the door to many experiments of this type.

Importance to neural prosthetic system design

Neural prosthetic system design depends critically on a fundamental and comprehensive scientific understanding of how populations of neurons prepare and generate natural movements (Green and Kalaska, 2010). How neural populations evolve on a fast timescale is of particular interest, as neural prostheses must operate rapidly to ensure accurate and stable control. Figure 12 highlights perhaps the most fundamental problem that basic neuroscience (reviewed above) and neural prosthetic systems have in common: how to understand noisy electrical activity from a population of neurons on a millisecond timescale and on a single-trial basis. Figure 12 shows the now familiar instructed-delay reach task and, for the first time in this review, reasonably raw and unprocessed electrode-array neural data. The most striking feature is how noisy spiking data really are. Staring at this figure for a few moments reveals the subtle difference in response pattern between the preparatory period following target onset and the baseline period preceding it. More obvious is the difference between the movement period and preparatory period. While the eye is often poor at discerning signal from noise, particularly without the visual benefit of trial averaging, this is a useful exercise as it helps one appreciate the challenge before us as neuroscientists and neuroengineers. We seek to understand how these neural responses arise, how they support behavior, and how we can use these volitional signals as an information source for controlling neural prostheses. While several important design insights for prostheses have already resulted from a deeper scientific understanding, some of which are discussed briefly below, there is little doubt that the most important leaps forward in prostheses will result from future scientific discovery. This has historically been the case, with science deeply informing engineering.

In addition to how the basic scientific understanding of brain organization and movement control has already informed neural prosthetic system design, as described in the Introduction, there are two additional points to briefly note here. First, communication prostheses rely on preparatory activity, as can motor prostheses (e.g., Musallam et al., 2004; Santhanam et al., 2006; Shenoy et al., 2003; Yu et al., 2007, 2010). These systems can make use of new discoveries such as preparatory activity in PMd reflecting the speed of the upcoming movement. Second, at the heart of the dynamical systems perspective and the associated quest for single-trial neural trajectories is time. How long does it take for the neural trajectory to actually traverse from baseline to the optimal subspace where, once there, the neural activity can be fruitfully decoded and used to guide a prosthesis? This is precisely the “transit time” we needed to know as part of our recent prosthesis research, so

that we could skip this transition period to avoid inadvertently decoding neural activity that is still in flux. This time (T_{skip} ; Santhanam et al., 2006) is approximately 200 ms as measured with single-trial neural trajectories in scientific experiments (as described above) and agrees with measurements from neural prosthetic experiments. Similarly, for prosthetics designs, it is important to know how long neural activity should be integrated (T_{int}), so as to best estimate the parameters of interest, and this is related to how stable preparatory neural activity is while in and around the optimal subspace. Single-trial neural trajectories can, and have, revealed important features which will continue to inform the design of neural prosthetic systems.

Summary

The ability to move voluntarily is central to the human experience. By pursuing a deeper scientific understanding of the neural control of natural movement, it should be possible to advance the design of neural prostheses, with the goal of helping patients who have lost their ability to move. A potentially underappreciated part of controlling movement is preparing movement. Preparation is, after all, how each movement begins. Motor preparation can be studied in many different ways. We have elected to adopt a dynamical systems perspective in order to facilitate the construction of hypotheses, and set about testing their predictions. The optimal subspace hypothesis has led to seven tested predictions. It appears that this dynamical systems perspective, and closely associated state-space diagrams, is helping to generate an ongoing series of testable predictions. As a result of the recent studies reviewed above, numerous questions are now more apparent and remain to be addressed as described in the Future Directions. We believe that the dynamical systems perspective should continue to help generate new and testable ideas and lead to deeper insights for both basic and applied neuroscience.

Acknowledgments

This work was supported by Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (K. V. S and M. M. C.), DARPA REPAIR N66001-10-C-2010 and NIH-NINDS CRCNS R01-NS-054283 (K. V. S. and M. S.), an NIH Director's Pioneer Award 1DP1OD006409 (K. V. S.), a National Science Foundation graduate research fellowship (M. T. K.), and the Gatsby Charitable Foundation (M. S.).

References

- Achtman N, Afshar A, Santhanam G, Yu BM, Ryu SI, Shenoy KV. Free-paced high-performance braincomputer interfaces. *Journal of Neuroengineering*. 2007; 4:336–347.
- Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV. Single-trial neural correlates of arm movement preparation. *Neuron*. 2011 In press.
- Andersen RA, Hwang EJ, Mulliken GH. Cognitive neural prosthetics. *Annual Review of Psychology*. 2010; 61:169–190.
- Anderson KD. Targeting recovery: Priorities of the spinal cord-injured population. *Journal of Neurotrauma*. 2004; 21:1371–1383. [PubMed: 15672628]
- Borton, DA.; Song, Y-K.; Patterson, WR.; Bull, CW.; Park, S.; Laiwalla, F.; Donoghue, JP.; Nurmikko, AV. Wireless, high-bandwidth recordings from nonhuman primate motor cortex using a scalable 16-ch implantable microsystem. *Conference proceedings: IEEE engineering in medicine and biology society*; 2009. p. 5531-5534.
- Briggman KL, Abarbanel HD, Kristan WB. Optical imaging of neuronal populations during decisionmaking. *Science*. 2005; 307:896–901. [PubMed: 15705844]
- Carmena JM, Lebedev MA, Crist RE, O'Doherty JE, Santucci DM, Dimitrov DF, et al. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*. 2003; 1:193–208.

- Chapin JK, Moxon KA, Markowitz RS, Nicolelis MAL. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*. 1999; 2:664–670.
- Chase SM, Schwartz AB, Kass RE. Bias, optimal linear estimation, and the differences between openloop simulation and closed-loop performance of spikingbased brain-computer interface algorithms. *Neural Networks*. 2009; 22:1203–1213. [PubMed: 19502004]
- Chestek, CA.; Cunningham, JP.; Gilja, V.; Nuyujukian, P.; Ryu, SI.; Shenoy, KV. Neural prosthetic systems: Current problems and future directions. *Proceedings of the 31st annual conference of IEEE EMBS*; 2009a. p. 3369-3375.
- Chestek CA, Gilja V, Nuyujukian P, Foster JD, Fan JM, Kaufman MT, et al. Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *Journal of Neural Engineering*. 2011
- Chestek CA, Gilja V, Nuyujukian P, Kier R, Solzbacher F, Ryu SI, et al. HermesC: Lowpower wireless neural recording system for freely moving primates. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2009b; 17:330–338. [PubMed: 19497829]
- Churchland MM, Afshar A, Shenoy KV. A central source of movement variability. *Neuron*. 2006a; 52:1085–1096. [PubMed: 17178410]
- Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron*. 2010a; 68:387–400. [PubMed: 21040842]
- Churchland, MM.; Cunningham, JP.; Kaufman, MT.; Ryu, SI.; Shenoy, KV. Firing rate oscillations underlie motor cortex responses during reaching in monkey. *Frontiers in Neuroscience*. Conference Abstract. *Computational and Systems Neuroscience (COSYNE)*; Salt Lake City, UT. 2011. p. III-32
- Churchland, MM.; Kaufman, MT.; Cunningham, JP.; Ryu, SI.; Shenoy, KV. Some basic features of the neural response in motor and premotor cortex. Program No. 382.2. *Neuroscience meeting planner*; San Diego, CA: Society for Neuroscience Online; 2010b.
- Churchland, MM.; Kaufman, MT.; Shenoy, KV. A novel maze task for the study of motor preparation. Program No. 182.14. *Neuroscience meeting planner*; Washington, DC: Society for Neuroscience; 2008.
- Churchland MM, Santhanam G, Shenoy KV. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of Neurophysiology*. 2006b; 96:3130–3146. [PubMed: 16855111]
- Churchland MM, Shenoy KV. Delay of movement caused by disruption of cortical preparatory activity. *Journal of Neurophysiology*. 2007a; 97:348–359. [PubMed: 17005608]
- Churchland MM, Shenoy KV. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of Neurophysiology*. 2007b; 97:4235–4257. [PubMed: 17376854]
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, et al. Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*. 2010c; 13:369–378.
- Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. Neural variability in premotor cortex provides a signature of motor preparation. *The Journal of Neuroscience*. 2006c; 26:3697–3712. [PubMed: 16597724]
- Churchland MM, Yu BM, Sahani M, Shenoy KV. Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Current Opinion in Neurobiology*. 2007; 17:609–618. [PubMed: 18093826]
- Cisek P. Preparing for speed. Focus on “preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach”. *Journal of Neurophysiology*. 2006; 96:2842–2843. [PubMed: 16928790]
- Crammond DJ, Kalaska JF. Prior information in motor and premotor cortex: Activity during the delay period and effect on pre-movement activity. *Journal of Neurophysiology*. 2000; 84:986–1005. [PubMed: 10938322]

- Cunningham, JP.; Churchland, MM.; Kaufman, MT.; Shenoy, KV. Extracting rotational structure from motor cortical data. *Frontiers in Neuroscience. Conference Abstract Computational and Systems Neuroscience (COSYNE)*; Salt Lake City, UT. 2011. p. II-33
- Cunningham JP, Nuyujukian P, Gilja V, Chestek CA, Ryu SI, Shenoy KV. A closed-loop human simulator for investigating the role of feedback-control in brain-machine interfaces. *Journal of Neurophysiology*. 2010; 105:1932–1949. [PubMed: 20943945]
- Cunningham J, Yu B, Gilja V, Ryu S, Shenoy K. Toward optimal target placement for neural prosthetic devices. *Journal of Neurophysiology*. 2008; 100:3445–3457. [PubMed: 18829845]
- Day BL, Rothwell JC, Thompson PD, de Noordhout AM, Nakashima K, Shannon K, et al. Delay in the execution of voluntary movement by electrical or magnetic brain stimulation in intact man. *Brain*. 1989; 112:649–663. [PubMed: 2731025]
- Millan J, del R, Carmena JM. Invasive or noninvasive: Understanding brain-machine interface technology. *IEEE Engineering in Medicine and Biology Magazine*. 2010; 29:16–22. [PubMed: 20209672]
- Diester I, Kaufman MT, Mogri M, Pashaie R, Goo W, Yizhar O, et al. An optogenetic toolbox designed for primates. *Nature Neuroscience*. 2011; 14:387–397.
- Donoghue JP. Bridging the brain to the world: A perspective on neural interface systems. *Neuron*. 2008; 60:511–521. [PubMed: 18995827]
- Donoghue JP, Nurmikko A, Black MJ, Hochberg LR. Assistive technology and robotic control using motor cortex ensemble-based neural interface systems in humans with tetraplegia. *Journal of Physiology*. 2007; 579:603–611. [PubMed: 17272345]
- Evarts EV. Temporal patterns of discharge of pyramidal tract neurons during sleep and waking in the monkey. *Journal of Neurophysiology*. 1964; 27:152–171. [PubMed: 14129768]
- Fetz EE. Operant conditioning of cortical unit activity. *Science*. 1969; 163:955–957. [PubMed: 4974291]
- Fetz EE. Are movement parameters recognizably coded in the activity of single neurons? *The Behavioral and Brain Sciences*. 1992; 15:679–690.
- Fetz EE. Volitional control of neural activity: Implications for brain-computer interfaces. *Journal of Physiology*. 2007; 579:571–579. [PubMed: 17234689]
- Fraser GW, Chase SM, Whitford A, Schwartz AB. Control of a brain-computer interface without spike sorting. *Journal of Neural Engineering*. 2009; 6:055004. [PubMed: 19721186]
- Ganguly K, Carmena JM. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biology*. 2009; 7:e1000153. [PubMed: 19621062]
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of Neuroscience*. 1982; 2:1527–1537. [PubMed: 7143039]
- Georgopoulos AP, Schwartz AB, Kettner RE. Neuronal population coding of movement direction. *Science*. 1986; 233:1416–1419. [PubMed: 3749885]
- Ghez C, Favilla M, Ghilardi MF, Gordon J, Bermejo R, Pullman S. Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research*. 1997; 115:217–233.
- Gilja V, Chestek CA, Diester I, Henderson JM, Deisseroth K, Shenoy KV. Challenges and opportunities for next generation intra-cortically based neural prostheses. *IEEE Transactions on Biomedical Engineering*. 2011; 10.1109/TBME.2011.2107553
- Gilja V, Chestek CA, Nuyujukian P, Foster J, Shenoy KV. Autonomous head-mounted electrophysiology systems for freely-behaving primates. *Current Opinion in Neurobiology*. 2010a; 20:676–686. [PubMed: 20655733]
- Gilja, V.; Nuyujukian, P.; Chestek, CA.; Cunningham, JP.; Yu, BM.; Ryu, SI., et al. A high-performance continuous cortically-controlled prosthesis enabled by feedback control design. Program No. 20.6. Neuroscience meeting planner; San Diego, CA: Society for Neuroscience Online; 2010b.
- Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Ryu SI, et al. High-performance continuous neural cursor control enabled by a feedback control perspective. *Frontiers in Neuroscience*. 2010c; 10.3389/conf.fnins.2010.03.00249

- Godschalk M, Lemon RN, Kuypers HG, van der Steen J. The involvement of monkey premotor cortex neurones in preparation of visually cued arm movements. *Behavioural Brain Research*. 1985; 18:143–157. [PubMed: 4096829]
- Green AM, Kalaska JF. Learning to move machines with the mind. *Trends in Neurosciences*. 2010; 34:61–75. [PubMed: 21176975]
- Grillner S. Biological pattern generation: The cellular and computational logic of networks in motion. *Neuron*. 2006; 52:751–766. [PubMed: 17145498]
- Han X, Qian X, Bernstein JG, Zhou HH, Franzesi GT, Stern P, et al. Millisecond-timescale optical control of neural dynamics in the nonhuman primate brain. *Neuron*. 2009; 62:191–198. [PubMed: 19409264]
- Harrison RR. The design of integrated circuits to observe brain activity. *Proceedings of the IEEE*. 2008; 96:1203–1216.
- Harrison RR, Kier RJ, Chestek CA, Gilja V, Nuyujukian P, Ryu SI, et al. Wireless neural recording with single low-power integrated circuit. *IEEE Transactions on Neural and Rehabilitation Systems Engineering*. 2009; 17:322–329.
- Harrison RR, Watkins PT, Kier RJ, Lovejoy RO, Black DJ, Greger B, et al. A low-power integrated circuit for a wireless 100-electrode neural recording system. *IEEE Journal of Solid-State Circuits*. 2007; 42:123–133.
- Hatsopoulos NG, Donoghue JP. The science of neural interface systems. *Annual Review of Neuroscience*. 2009; 32:249–266.
- Heliot R, Ganguly K, Jimenez J, Carmenta JM. Learning in closed-loop brain-machine interfaces: Modeling and experimental validation. *IEEE Transactions on Systems, Man, and Cybernetics. Part B: Cybernetics*. 2009; 6:3710–3715.
- Hochberg LR. Turning thought into action. *The New England Journal of Medicine*. 2008; 359:1175–1177. [PubMed: 18784110]
- Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, et al. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*. 2006; 442:164–171. [PubMed: 16838014]
- Hochberg LR, Taylor DM. Intuitive prosthetic limb control. *Lancet*. 2007; 369:345–346. [PubMed: 17276755]
- Humphrey DR, Schmidt EM, Thompson WD. Predicting measures of motor performance from multiple cortical spike trains. *Science*. 1970; 170:758–762. [PubMed: 4991377]
- Isaacs RE, Weber DJ, Schwartz AB. Work toward real-time control of a cortical neural prosthesis. *IEEE Transactions on Rehabilitation Engineering*. 2000; 8:196–198. [PubMed: 10896185]
- Jackson A, Fetz EE. Compact movable microwire array for long-term chronic unit recording in cerebral cortex of primates. *Journal of Neurophysiology*. 2007; 98:3109–3118. [PubMed: 17855584]
- Jackson A, Mavoori J, Fetz EE. Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*. 2006; 444:56–60. [PubMed: 17057705]
- Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB. Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences of the United States of America*. 2008; 105:19486–19491. [PubMed: 19047633]
- Kalaska JF. From intention to action: Motor cortex and the control of reaching movements. *Advances in Experimental Medicine and Biology*. 2009; 629:139–178. [PubMed: 19227499]
- Kalaska JF, Scott SH, Cisek P, Sergio LE. Cortical control of reaching movements. *Current Opinion in Neurobiology*. 1997; 7:849–859. [PubMed: 9464979]
- Kalmar, R.; Reppas, J.; Ryu, SI.; Shenoy, KV.; Newsome, WT. Ensemble activity underlying movement preparation in prearcuate cortex. *Frontiers in neuroscience. Conference abstract: Computational and systems neuroscience (COSYNE); Salt Lake City, UT*. 2010.
- Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, et al. The roles of monkey premotor neuron classes in movement preparation and execution. *Journal of Neurophysiology*. 2010a; 104:799–810. [PubMed: 20538784]

- Kaufman, MT.; Churchland, MM.; Shenoy, KV. Program No. 463.11. Neuroscience Meeting Planner. Chicago, IL: Society for Neuroscience; 2009. Putative interneurons respond more rapidly than pyramidal cells in monkey premotor cortex.
- Kaufman, MT.; Churchland, MM.; Shenoy, KV. nanosymposium Lack of evidence for inhibitory gating in monkey M1. Program No. 732.2. 2010. Neuroscience Meeting Planner; San Diego, CA: Society for Neuroscience. Online; 2010b.
- Kaufman, MT.; Churchland, MM.; Shenoy, KV. Cortical preparatory activity avoids causing movement by remaining in a muscle-neutral space. *Frontiers in Neuroscience*. Conference Abstract. Computational and Systems Neuroscience. (COSYNE); Salt Lake City, UT. 2011. p. II-61
- Keele SW. Movement control in skilled motor performance. *Psychological Bulletin*. 1968; 70:387–403.
- Kemere C, Santhanam G, Yu B, Afshar A, Ryu S, Meng T, et al. Detecting neural-state transitions using hidden Markov models for motor cortical prostheses. *Journal of Neurophysiology*. 2008; 100:2441–2452. [PubMed: 18614757]
- Kemere CT, Shenoy KV, Meng TH. Modelbased neural decoding of reaching movements: A maximum likelihood approach. *IEEE Transactions on Biomedical Engineering*. 2004; 51:925–932. [PubMed: 15188860]
- Kim SP, Simeral JD, Hochberg LR, Donoghue JP, Black MJ. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *Journal of NeuroEngineering*. 2008; 5:455–476.
- Kristan WB, Calabrese RL. Rhythmic swimming activity in neurones of the isolated nerve cord of the leech. *The Journal of Experimental Biology*. 1976; 65:643–686. [PubMed: 1018167]
- Kurata K. Distribution of neurons with set- and movement- related activity before hand and foot movements in the premotor cortex of rhesus monkeys. *Experimental Brain Research*. 1989; 77:245–256.
- Kutas M, Donchin E. Studies of squeezing: Handedness, responding hand, response force, and asymmetry of readiness potential. *Science*. 1974; 186:545–548. [PubMed: 4469679]
- Lerchner, A.; Shenoy, KV.; Sahani, M. The role of horizontal long-range connections in shaping the dynamics of multi-electrode array data. *Frontiers in Neuroscience*. Conference Abstract. Computational and Systems Neuroscience (COSYNE); Salt Lake City, UT. 2011. p. 1-27.
- Linderman MD, Santhanam G, Kemere CT, Gilja V, O'Driscoll S, Yu BM, et al. Signal processing challenges for neural prostheses. *IEEE Signal Processing Magazine*. 2008; 25:18–28.
- Macke, J.; Busing, L.; Cunningham, JP.; Yu, BM.; Shenoy, KV.; Sahani, M. Modelling lowdimensional dynamics in recorded spiking populations. In: *Frontiers in Neuroscience*. Conference Abstract. Computational and Systems Neuroscience, (COSYNE 2011); Salt Lake City, UT. 2011. p. I-34.
- Mandelblat-Cerf Y, Paz R, Vaadia E. Trial-totrial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *The Journal of Neuroscience*. 2009; 29:15053–15062. [PubMed: 19955356]
- Mavoori J, Jackson A, Diorio C, Fetz E. An autonomous implantable computer for neural recording and stimulation in unrestrained primates. *Journal of Neuroscience Methods*. 2005; 148:71–77. [PubMed: 16102841]
- Messier J, Kalaska JF. Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *Journal of Neurophysiology*. 2000; 84:152–165. [PubMed: 10899193]
- Moritz CT, Perlmuter SI, Fetz EE. Direct control of paralysed muscles by cortical neurons. *Nature*. 2008; 456:639–642. [PubMed: 18923392]
- Mulliken GH, Musallam S, Andersen RA. Decoding trajectories from posterior parietal cortex ensembles. *The Journal of Neuroscience*. 2008; 28:12913–12926. [PubMed: 19036985]
- Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA. Cognitive control signals for neural prosthetics. *Science*. 2004; 305:258–262. [PubMed: 15247483]
- Nicolelis MAL, Lebedev MA. Principles of neural ensemble physiology underlying the operation of brainmachine interfaces. *Nature Reviews Neuroscience*. 2009; 10:530–540.

- Nurmikko AV, Donoghue JP, Hochberg L, Patterson WR, Song YK, Bull C, et al. Listening to brain microcircuits for interfacing with external world—Progress in wireless implantable microelectronic neuroengineering devices. *Proceedings of the IEEE*. 2010; 98:375. [PubMed: 21654935]
- Nuyujukian, P.; Gilja, V.; Chestek, CA.; Cunningham, JP.; Fan, JM.; Yu, BM., et al. Generalization and robustness of a continuous cortically-controlled prosthesis enabled by feedback control design. Program No. 20.7. Neuroscience Meeting Planner; San Diego, CA: Society for Neuroscience Online; 2010.
- O'Driscoll, S.; Meng, TH.; Shenoy, KV.; Kemere, C. Neurons to silicon: Implantable prosthesis processor. *International solid state circuits conference (ISSCC)*; 2006. p. 552-553.p. 672
- Petreska, B.; Cunningham, JP.; Santhanam, G.; Yu, BM.; Ryu, SI.; Shenoy, KV.; Sahani, M. Detecting changes in neural dynamics within single trials. *Frontiers in Neuroscience. Conference Abstract Computational and Systems Neuroscience (COSYNE)*; Salt Lake City, UT. 2011. p. I-33
- Rickert J, Riehle A, Aertsen A, Rotter S, Nawrot MP. Dynamic encoding of movement direction in motor cortical neurons. *The Journal of Neuroscience*. 2009; 29:13870–13882. [PubMed: 19889998]
- Riehle A, Requin J. Monkey primary motor and premotor cortex: Single-cell activity related to prior information about direction and extent of an intended movement. *Journal of Neurophysiology*. 1989; 61:534–549. [PubMed: 2709098]
- Riehle A, Requin J. The predictive value for performance speed of preparatory changes in neuronal activity of the monkey motor and premotor cortex. *Behavioural Brain Research*. 1993; 53:35–49. [PubMed: 8466666]
- Rivera-Alvidrez, Z.; Kalmar, R.; Afshar, A.; Santhanam, G.; Yu, BM.; Ryu, SI., et al. Single-trial representation of uncertainty about reach goals in macaque PMd. Presentation III-52. *Computational and Systems Neuroscience (COSYNE) meeting*; Salt Lake City, UT. 2008.
- Rivera-Alvidrez, Z.; Kalmar, R.; Ryu, SI.; Shenoy, KV. Low-dimensional neural features predict muscle EMG signals. *Proceedings of the 32nd annual international conference of the IEEE EMBS*; 2010a. p. 6027-6033.
- Rivera-Alvidrez, Z.; Kalmar, R.; Ryu, SI.; Shenoy, KV. Low dimensional neural features predict specific muscle EMG signals. Program No. 732.4. Neuroscience Meeting Planner; San Diego, CA: Society for Neuroscience; 2010b.
- Program No. 463.4, Neuroscience Meeting Planner; Chicago, IL: Society for Neuroscience; 2009.
- Rosenbaum DA. Human movement initiation: Specification of arm, direction, and extent. *Journal of Experimental Psychology General*. 1980; 109:444–474. [PubMed: 6449531]
- Ryu SI, Shenoy KV. Human cortical prostheses: Lost in translation? *Neurosurgical Focus*. 2009; 27:E5. [PubMed: 19569893]
- Santhanam G, Linderman MD, Gilja V, Afshar A, Ryu SI, Meng TH, et al. HermesB: A continuous neural recording system for freely behaving primates. *IEEE Transactions on Biomedical Engineering*. 2007; 54:2037–2050. [PubMed: 18018699]
- Santhanam G, Ryu SI, Yu BM, Afshar A, Shenoy KV. A high-performance brain-computer interface. *Nature*. 2006; 442:195–198. [PubMed: 16838020]
- Santhanam, G.; Sahani, M.; Ryu, SI.; Shenoy, KV. An extensible infrastructure for fully automated spike sorting during online experiments. *Proceedings of the 26th annual international conference of IEEE EMBS*; San Francisco, CA. 2004. p. 4380-4384.
- Santhanam G, Yu BM, Gilja V, Ryu SI, Afshar A, Sahani M, et al. Factor-analysis methods for higher-performance neural prostheses. *Journal of Neurophysiology*. 2009; 102:1315–1330. [PubMed: 19297518]
- Scherberger H. Neural control of motor prostheses. *Current Opinion in Neurobiology*. 2009; 19:629–633. [PubMed: 19896364]
- Schwartz AB. Direct cortical representation of drawing. *Science*. 1994; 265:540–542. [PubMed: 8036499]
- Schwartz AB. Useful signals from motor cortex. *Journal of Physiology*. 2007; 579:581–601. [PubMed: 17255162]

- Scott SH. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*. 2004; 7:524–546.
- Scott SH. Inconvenient truths about neural processing in primary motor cortex. *Journal of Physiology*. 2008; 586:1217–1224. [PubMed: 18187462]
- Serruya MD, Hatsopoulos NG, Paninski L, Fellows MR, Donoghue J. Instant neural control of a movement signal. *Nature*. 2002; 416:141–142. [PubMed: 11894084]
- Shenoy KV, Meeker D, Cao S, Kureshi SA, Pesaran B, Mitra P, et al. Neural prosthetic control signals from plan activity. *Neuroreport*. 2003; 14:591–596. [PubMed: 12657892]
- Snyder LH, Batista AP, Andersen RA. Coding of intention in the posterior parietal cortex. *Nature*. 1997; 386:167–170. [PubMed: 9062187]
- Steinmetz NA, Moore T. Changes in the response rate and response variability of area v4 neurons during the preparation of saccadic eye movements. *Journal of Neurophysiology*. 2010; 103:1171–1178. [PubMed: 20018834]
- Stopfer M, Jayaraman V, Laurent G. Intensity versus identity coding in an olfactory system. *Neuron*. 2003; 39:991–1004. [PubMed: 12971898]
- Tanji J, Evarts EV. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *Journal of Neurophysiology*. 1976; 39:1062–1068. [PubMed: 824409]
- Taylor DM, Tillery SIH, Schwartz AB. Direct cortical control of 3D neuroprosthetic devices. *Science*. 2002; 296:1829–1832. [PubMed: 12052948]
- Todorov E. Direct cortical control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience*. 2000; 3:391–398.
- Truccolo W, Friehs GM, Donoghue JP, Hochberg LR. Primary motor cortex tuning to intended movement kinematics in humans with tetraplegia. *The Journal of Neuroscience*. 2008; 28:1163–1178. [PubMed: 18234894]
- Truccolo W, Hochberg LR, Donoghue JP. Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nature Neuroscience*. 2010; 13:105–111.
- Velliste M, Perel S, Spalding MC, Whitford AS, Schwartz AB. Cortical control of a prosthetic arm for self-feeding. *Nature*. 2008; 453:1098–1101. [PubMed: 18509337]
- Weinrich M, Wise SP, Mauritz KH. A neurophysiological study of the premotor cortex in the rhesus monkey. *Brain*. 1984; 107:385–414. [PubMed: 6722510]
- Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, et al. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*. 2000; 408:361–365. [PubMed: 11099043]
- Wise SP. The primate premotor cortex: Past, present and preparatory. *The Annual Review of Neuroscience*. 1985; 8:1–19.
- Wu W, Gao Y, Bienenstock E, Donoghue JP, Black MJ. Bayesian population decoding of motor cortical activity using a kalman filter. *Neural Computation*. 2006; 18:80–118. [PubMed: 16354382]
- Wu, W.; Shaikhouni, A.; Donoghue, JP.; Black, MJ. Closed-loop neural control of cursor motion using a kalman filter. *Proceedings of the 26th annual international conference of the IEEE, engineering in medicine and biology society*; 2004. p. 4126-4129.
- Yu, BM.; Afshar, A.; Santhanam, G.; Ryu, SI.; Shenoy, KV.; Sahani, M. Extracting dynamical structure embedded in neural activity. In: Weiss, Y.; Schölkopf, B.; Platt, J., editors. *Advances in neural information processing systems*. Vol. 18. Cambridge, MA: MIT Press; 2006. p. 1545-1552.
- Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*. 2009; 102:614–635. [PubMed: 19357332]
- Yu BM, Kemere C, Santhanam G, Afshar A, Ryu SI, Meng TH, et al. Mixture of trajectory models for neural decoding of goal-directed movements. *Journal of Neurophysiology*. 2007; 97:3763–3780. [PubMed: 17329627]
- Yu, BM.; Santhanam, G.; Sahani, M.; Shenoy, KV. Neural decoding for motor and communication prostheses. In: Oweiss, KG., editor. *Statistical signal processing for neuroscience*. Elsevier; 2010. p. 219-263. p. 58

Fig. 1. Illustration of the instructed-delay task, hand measurements, and EMG recordings. (a) Monkeys sit in a primate chair approximately 25 cm from a fronto-parallel display. Movements begin and end with the hand touching the display. The hand is a few millimeters from the screen while in flight. The white trace shows the reach trajectory for one trial. (b) Time line of the task and behavior for the same trial. T, target onset; G, go cue; and M, movement onset. Horizontal hand (black) and target (red) position is plotted (top). The target jittered on first appearing and ceased at the go cue. Bottom: Gray trace plots hand velocity (computed in the direction of the target), superimposed on the voltage recorded from the medial deltoid (arbitrary vertical scale). Traces end at the time of the reward. Data are from monkey A in a session focused on EMG recordings. Figure adapted from Churchland et al. (2006b).

Fig. 2.

Mean RT (in milliseconds) is plotted versus delay period duration. For monkeys A and B, this was for the catch trials with short delays. Although the delay period was selected from a continuum, in practice, delay periods were integer multiples of 16 ms because of video presentation, and this binning is used in the plots. Lines show exponential fits. For monkey G, we did not use catch trials (the minimum delay for most experiments was already quite short, at 200 ms). The plotted data are therefore from one experiment using three discrete delay durations (30, 130, and 230 ms; black symbols) and another (performed the previous day) using a continuous range (200–700 ms; white symbols). For the latter, data have been binned (ranges shown in parentheses). From Churchland et al. (2006c).

Fig. 3.

Examples of typical delay-period responses in PMd. (a) Mean $\pm$ SE firing rates for four example neurons. Three of these showed increases in firing rate after target onset, whereas one showed a decrease. Data are from experiments using a continuous range of delay periods (500–900 for monkey B and 400–800 for monkey A). For each time point, mean firing rate was computed from only those trials with a delay period at least that long. Labels give the monkey initial and cell number. Details (direction, distance, instructed speed, and trials/condition) were as follows: cell B29, 45°, 85 mm, fast, 23 trials; cell B16, 135°, 60 mm, fast, 20 trials; cell B46, 335°, 85 mm, fast, 41 trials; cell A2, 185°, 120 mm, slow, 42 trials. From Churchland et al. (2006c).

Fig. 4.

Illustration of the optimal subspace hypothesis. 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 of states that are optimal in the sense that they will produce the desired result 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. From Churchland et al. (2006c).

Fig. 5.

Predicted and measured relationships between neural activity and reach velocity. (a) The prediction that preparatory activity should covary with instructed reach speed (prediction 1) is visualized in the state-space framework. Two optimal subspaces are illustrated: one shaded red for the fast instruction and another shaded green for the slow instruction. The prediction that preparatory activity should correlate, on a trial-by-trial basis, with peak reach speed (prediction 2) can also be visualized in this state space. For example, an instructed-fast trial with a slower-than-typical actual reach speed should have a preparatory state toward one end of the “fast” optimal subspace, nearer to the “slow” optimal subspace. (b) Examination of the first prediction: the mean firing rate is plotted as a function of time for one neuron, one target location, and both instructed speeds. For this neuron, the mean firing rate was highest when preparing a fast reach. Other neurons showed the opposite pattern. “T,” “G,” and “M” indicate target onset, the go cue, and the median time of movement onset. (c) Trial-by-trial correlation between preparatory firing rate and peak reach speed. Data are from one neuron (B24). Each dot plots the mean delay-period firing rate versus peak reach speed for one trial. Trials have been pooled across target locations (for this neuron all locations involved a preference for the fast instruction). To allow pooling, firing rates and peak speeds are expressed relative to the mean for the relevant condition. The offset between the left subpanel (instructed-slow reaches) and the right subpanel (instructed-fast reaches) indicates the degree to which firing rates were on average higher for the instructed-fast condition. This defines a “predicted slope” (gray line) with which one can compare the slopes computed from the trial-by-trial variability (black lines). (d) Similar plot but for a second neuron (A06) for which the average firing rate was higher for instructed-slow reaches. In agreement, the trial-by-trial correlations show negative slopes. Figure adapted from Churchland et al. (2006a).

Fig. 6.

State-space view of prediction 3, and Fano factor relationship with RT in the instructed-delay task. (a) State-space illustration, as in Fig. 4, showing the state of several trials (black circles) from before target onset until converging within the optimal subspace (gray shaded region) approximately 200 ms later. (b) Fano factor reduces from the time of target onset, approximately holds at a plateau level throughout delay periods longer than 200 ms, and then reduces further following the go cue. (c) By employing shorter delay periods, where the go cue comes at 30, 130, or 230 ms after target onset (colored arrows), it is possible to ask if the resulting RTs are longer when the Fano factor is higher (prediction 4-I). (d) Mean $\pm$ SEM measured RT plotted against mean $\pm$ measured Fano factor for the three delay durations (monkey G). A clear correlation is observed, with longer delay durations leading to both lower Fano factors and lower RTs as predicted. As a technical aside, when measuring across-trial variability our later publications (e.g., Churchland et al. 2007, Churchland et al. 2010c) employed the Fano factor while our original publication (Churchland et al. 2006c) employed the closely related ‘normalized variance’. The above plots are reproduced from that original manuscript, and it should thus be kept in mind that the vertical axis is not technically the Fano factor, because spiking was assessed in a Gaussian window rather than a square counting window. That said, results are very similar regardless of the exact window shape (e.g., Churchland et al. 2007, Fig. 4).

Fig. 7.

Changes in firing-rate variability for 10 datasets (one per panel). Insets indicate stimulus type. Data are aligned on stimulus onset (arrow). For the two bottom panels (MT area/direction and MT speed), the dot pattern appeared at time zero (first arrow) and began moving at the second arrow. The mean rate (gray) and the Fano factor (black with flanking SEM) were computed using a 50-ms sliding window. For OFC, where response amplitudes were small, a 100-ms window was used to gain statistical power. Analysis included all conditions, including nonpreferred. The Fano factor was computed after mean matching (Churchland et al., 2010c). The resulting stabilized means are shown in black. The mean number of trials per condition was 100 (V1), 24 (V4), 15 (MT plaids), 88 (MT dots), 35 (LIP), 10 (PRR), 31 (PMd), 106 (OFC), 125 (MT direction and area), and 14 (MT speed). Figure from Churchland et al. (2010c).

Fig. 8. State-space view of prediction 4-II, and relationship of the Fano factor to natural RT variability. (a) State-space view of prediction 4-II. The shaded area represents the optimal subspace for the movement being prepared, as in Fig. 4. Each dot corresponds to one trial and represents the configuration of firing rates around the time of the go cue. For some trials, that configuration may lie within the optimal subspace (green dots), leading to a short RT. For other trials, the configuration may lie outside (red dots), leading to a longer RT. (b) Red and green traces show the Fano factor, around the time of the go cue, for trials with RTs longer and shorter than the median. Data were pooled across the recordings from 7 days (monkey G), including all trials with delay periods >200 ms. Figure adapted from Churchland et al. (2006c).

Fig. 9.

State-space view of prediction 5 and influence of subthreshold microstimulation in PMd on arm movement RT. Inset: Statespace view of the predicted effect of subthreshold microstimulation on the preparatory state (curly arrow displacing the state, shown as a black circle), and the time consuming process of reoptimizing the preparatory state so that it is returned to within the optimal subspace (red dashed line). Black curve: nonstimulated trials. Red curve: microstimulation occurring just after the go cue, when preparatory activity is most needed. Green curve: stimulation occurring well before the go cue, when time exists following stimulation offset for preparatory activity to recover before the go cue appears. Red and green bars indicate when stimulation was delivered. Data are from 30 experiments in PMd, and curves plot mean $\pm$ SE.

Fig. 10.

Single-trial neural trajectories computed using GPFA. (a) 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. Fifteen (of 47) randomly selected trials are shown. (b) 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 approximately 300 ms after the go cue. For display, 18 randomly selected trials are plotted, plus one hand-selected trial (red, trialID 211). Covariance ellipses were computed across all 47 trials. This is a two-dimensional projection of a 10-dimensional latent space. In the full space, the black ellipse is far from the edge of the blue ellipse. This projection was chosen to accurately preserve the relative sizes (on a per-dimension basis) of the true 10-dimensional volumes of the ellipsoids. Data are from the G20040123 dataset. (c) Data are presented as in (b), with the same target location, but for data from another day's dataset (G20040122; red trial, trialID 793). From Churchland et al. (2010c).

Fig. 11.

Low-dimensional state space (as in Fig. 4) illustrating prediction 7. (a) A single-trial neural trajectory is illustrated with activity from the time of target onset (red circle) shown in red, from the time the go cue is presented (green circle) shown in green, and from the time of movement onset (blue circle) shown in blue. One or more regions of the brain appear to monitor the preparatory state and initiate movement (“pull the movement trigger”) only if it is within the optimal subspace, so as to assure that the desired movement results. (b) Illustration of prediction 7. After the movement trigger is pulled, a preparatory state that happens to be farther along the loop in the standard direction should have a shorter RT. Trial 1 should have a shorter RT than trial 2. Not shown is the related prediction that a preparatory state that is moving faster in the standard direction of loop travel should also have a shorter RT.

Fig. 12.

Instructed-delay reach task and neural activity from a single trial, highlighting the need to understand all phases of volitional arm movement (holding, preparing, and moving) on a trial-by-trial basis as this is the fundamental information source and (millisecond) timescale on which neural prosthetic systems depend. Top: Schematic illustration and timeline of the instructed-delay reach task introduced in Fig. 1 and discussed throughout this review. Middle: measured hand speed. Photo: 100 electrode array used to record from many tens of neurons simultaneously. Bottom: action potential (spike) raster from 97 simultaneously recorded single and multi-units, with the time of an action potential indicated by a black tick mark. Red vertical lines indicate the time of target onset, go cue presentation, and movement onset.