

A Dynamical Basis Set for Generating Reaches

MARK M. CHURCHLAND¹ AND JOHN P. CUNNINGHAM²

¹*Department of Neuroscience, Grossman Center for the Statistics of Mind, David Mahoney Center for Brain and Behavior Research, Kavli Institute for Brain Science, Columbia University Medical Center, New York, New York 10032*

²*Department of Statistics, Grossman Center for the Statistics of Mind, Center for Theoretical Neuroscience, Institute for Data Science and Engineering, Columbia University, New York, New York 10027*

Correspondence: mc3502@columbia.edu

The motor cortex was the one of the first cortical areas to be explored electrophysiologically, yet little agreement has emerged regarding its basic response properties. Often it is assumed that single-neuron responses reflect a preference for a particular movement or movement variable. It may be further assumed that movement is generated by (or at least accompanied by) a growing population-level preference for the relevant movement. This view has been attractive because it provides a canonical form for the single neuron, a link between preparatory and movement activity, a way of interpreting the population response, and a platform for designing analyses and couching hypotheses. However, this traditional view yields predictions that are at odds with basic features of the data. We discuss an alternative simplified model, in which outgoing commands are produced by dynamics that generate different output patterns as a function of the initial preparatory state. For reaching tasks, we hypothesized simple quasioscillatory dynamics because they provide a natural basis set for the empirical patterns of muscle activity. The predictions of the dynamical model match the data well at both the single-neuron and population levels, and the quasioscillatory patterns explain many of the otherwise odd features of the neural responses.

Since the pioneering work of Evarts (1968), a central goal of motor physiology has been to understand the relationship between neural activity in motor cortex and movement itself. Much of the subsequent literature has leveraged tasks in which monkeys execute reaches (or directed wrist movements) to targets at different spatial locations (Georgopoulos et al. 1986; Moran and Schwartz 1999; Kakei et al. 2001). In many versions of this task (Fig. 1) a target is presented but it is required that movement execution be delayed until after a go cue (Weinrich et al. 1984; Crammond and Kalaska 2000). The resulting “preparatory period” (between target onset and the go cue) is employed because it temporally separates neural activity related to movement preparation from neural activity related to movement itself. The key question is then: How does neural activity unfold during the preparatory and movement periods? Does activity before the go cue prepare the system for the production of movement activity? If so, how? What patterns of activity are present during movement, and what purpose do they serve?

It is widely accepted that neural activity in motor cortex is complex and heterogeneous—across time, across neurons, and across different reach types. A central goal of motor physiology is then to create simple models that capture basic response features and offer explanations for why those features exist. Consider the neural response before and during the two reaches illustrated in Figure 2A. Many models and/or analyses have assumed the framework illustrated in Figure 2B,D. The schematic single-neuron response (Fig. 2B) gives a larger response for the reach down and to the right (*red*); that neuron is thus

said to “prefer” that direction. This preference becomes apparent during preparation and grows stronger during movement. Neural responses consistent with this view are often observed in population averages (Requin et al. 1988; Georgopoulos et al. 1989a; Riehle and Requin 1989; Bastian et al. 1998, 2003; Erlhagen et al. 1999; Cisek 2006a). Accordingly, it is often posited that movement onset is triggered when response strength passes a threshold or reaches a peak.

The single-neuron view in Figure 2B is readily extended to the population level (Georgopoulos et al. 1986, 1989b) by considering the response in neural “state space” (Fig. 2D). Each axis captures a “factor”: a weighted sum of the neural responses. In this case x_1 might be the average activity of rightward-preferring neurons minus the average activity of leftward-preferring neurons. Accordingly, x_2 might be the average activity of upward-preferring neurons minus the average activity of downward-preferring neurons. A reach in a given direction would be generated when the “neural state” expands away from baseline and crosses a threshold. This hypothesis possesses some appealing features. First, it resembles the mechanism by which saccades are prepared and generated (e.g., Hanes and Schall 1996). Second, directions in neural state space correspond to directions in physical space. This provides a concrete interpretation of the neural state as a representation of external factors such as reach direction or velocity. For example, the neural state might be modeled as

$$\mathbf{x}(t, c) = g(\mathbf{v}(t, c)), \quad (1)$$

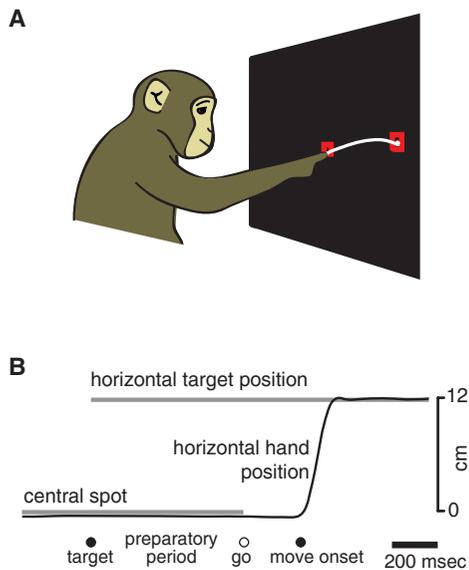


Figure 1. Illustration of behavior for a reaching task, of the type often used to study motor cortex. (A) Reaches were from a central spot to a target. An example trajectory is shown. (B) Task timeline. Target onset was separated from a go cue by a preparatory period. Movement onset occurred shortly (~ 300 msec) after the go cue. Reaches were typically very brisk; in this case lasting < 200 msec.

where $\mathbf{v}(t, c)$ is a vector of hand velocity for time t and condition c . The function $g(\cdot)$ captures “tuning” of neural activity for reach velocity and may be linear or nonlinear. This traditional view remains prevalent in motor physiology (e.g., Pearce and Moran 2012) and virtually all motor prosthetic applications use models that contain relationships similar to that in Equation 1 (e.g., Gilja et al. 2012).

We recently proposed (Churchland et al. 2012; Shenoy et al. 2013) an alternative “dynamical” model. The dynamical model extends the view that motor cortex produces a muscle-like output (Evarts 1968; Cheney and Fetz 1980; Lemon et al. 1986; Todorov 2000; Sergio et al. 2005; Pohlmeier et al. 2007; Schieber and Rivlis 2007; Lillicrap and Scott 2013), yet carries important additional implications. In this dynamical model the dominant motif during reaching is a rotation of the neural state (Fig. 2E) rather than an expansion (Fig. 2D). It is still proposed, as in the traditional model, that the preparatory response involves an expansion away from baseline. However, this preparatory state then “seeds” subsequent oscillatory dynamics, producing a rotation of the neural state that lasts ~ 1 – 1.5 cycles (Churchland et al. 2010). Rotations of the neural state imply quasioscillatory single-neuron response patterns (Fig. 2C), but the converse is not true—the central prediction is more specific than multiphasic single-neuron responses. It is hypothesized that the population neural state evolves according to dynamics of the form

$$\dot{\mathbf{x}}(t, c) = f(\mathbf{x}(t, c)), \quad (2)$$

where $f(\cdot)$ shows a large rotational component. A critical feature of this model is that the dynamics $f(\cdot)$ are the same across all conditions. Thus, each condition is a different neural trajectory in the same dynamical flow field. Although many models imply multiphasic responses at the level of single neurons, many fewer predict population-level rotational dynamics that are consistent across conditions. Below we illustrate why such a model might make sense, given the empirical patterns of muscle activity.

MODEL AND NEURAL RESPONSE PATTERNS

A Basis Set for Muscle Activity? Motivating a Dynamical Hypothesis

The basic features of the traditional model are fairly intuitive: The factors that describe the neural state space correspond to externally measurable parameters such as direction or velocity. In contrast, the appeal of the dynamical model is not immediately obvious; what could be the mechanistic purpose of a neural state governed by internal dynamics (Eq. 2)? A possible answer presents itself if we consider the patterns of activity that can be generated by the traditional and dynamical models. The traditional model generates pulses of activity, whereas the dynamical model generates brief quasioscillatory patterns. Both pulses and oscillatory patterns are potentially useful: Each could constitute a set of building blocks (a “basis set”) for outgoing commands.

To investigate the possible utility of an oscillatory basis set, we simulated a dynamical model in which the state space was greater than two dimensional, allowing for more than one plane of rotation and thus more than one frequency. Such a system produces a small basis set of oscillatory patterns (Fig. 3A), and can do so with any phase and amplitude depending on the preparatory state. If the output is a (fixed) sum of the resulting patterns, one then has a mechanism for generating a range of output patterns.

We examined the patterns of muscle activity across different reaches (Fig. 3B) and found that they could indeed be reconstructed (Fig. 3C) from a basis set of just two underlying oscillations (four total dimensions plus an additional dimension to capture offset). Oscillation frequency and duration were fixed (and thus the dynamics $f(\cdot)$ in Eq. 2 were fixed); only phase and amplitude varied across conditions (implied by the initial state $\mathbf{x}(0, c)$). This simple construction was sufficient to accurately reproduce muscle-like commands (Fig. 3C vs. 3B).

Quasioscillatory patterns are thus an effective basis set, in agreement with a recent modeling study (Rokni and Sompolinsky 2012). There are two further reasons to suspect such a basis set might be used. First, a common feature of large recurrent networks—even randomly connected networks—is the presence of multiple oscillatory “modes.” Network dynamics can, via training, be simplified to produce a small set of well-behaved output activities that are muscle-like (Schaffer et al. 2006; Sussillo et al. 2012). Second, oscillatory activity is common during movement in many species, not only during locomotion but sometimes during brief movements as well.

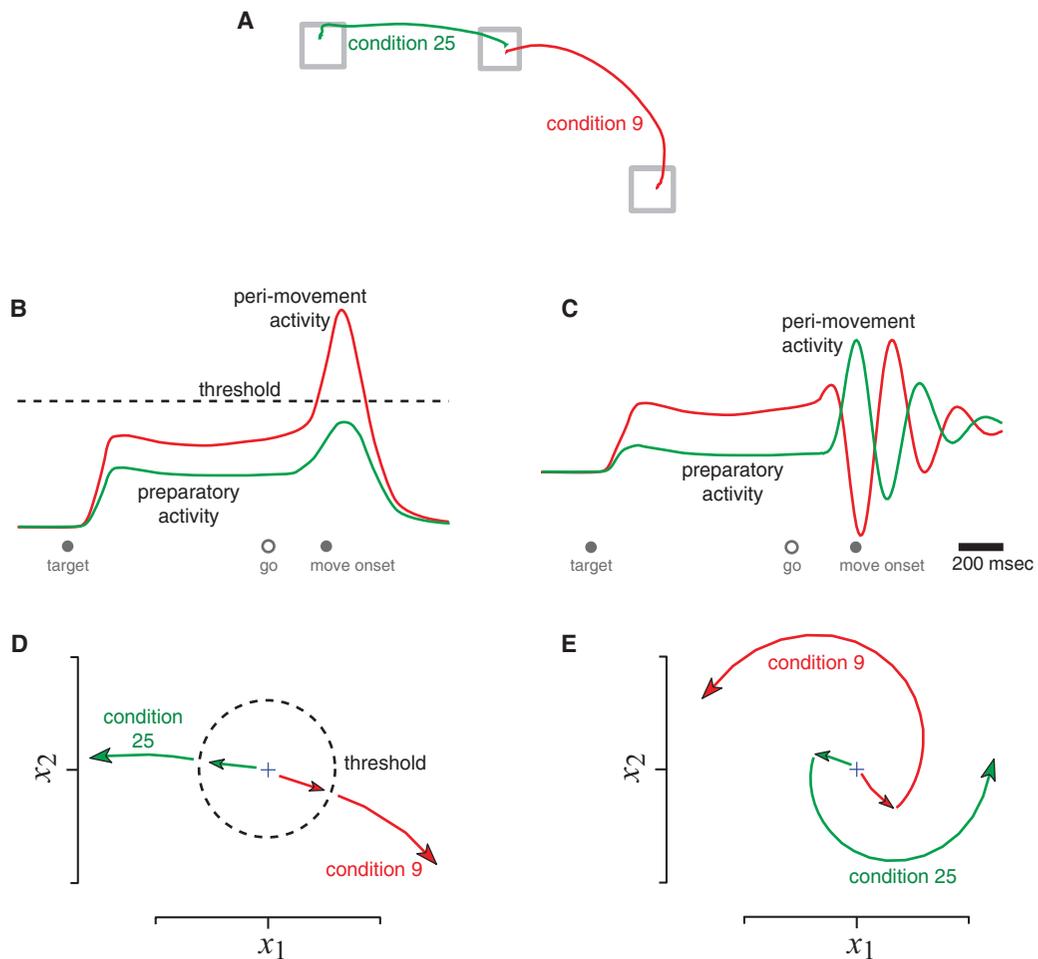


Figure 2. Two simplified models of neural activity during reach generation. (A) Example reach trajectories (2 of 27 total conditions). (B) Cartoon illustration of the traditional model showing the response of one model unit as a function of time for two conditions. The red trace shows the response for the “preferred condition” (condition 9, a roughly rightward reach). (C) Cartoon illustration of the dynamical model, showing the response of one model unit as a function of time for two conditions. (D) Illustration of the traditional model in state space. Each dimension (x_1 and x_2) captures one of the factors represented by the population. In this case x_1 and x_2 correspond to horizontal and vertical reach direction. Motion of the neural state during preparation is away from baseline, and this expansion continues during movement. This expansion corresponds, at the single-unit level, to the development of a preference during preparation that is heightened during movement. For example, a single unit whose response primarily reflects x_1 would show a rightward preference, as in B. (E) Illustration of the dynamical model in state space. The neural state moves away from baseline during preparation, as with the traditional model. During reaching, the neural state rotates with time, tracing a circular trajectory. To connect back to the single-unit level, note that a single unit whose response primarily reflects x_1 would show an oscillatory pattern with an ever-changing “preference,” as in C.

For example, zebrafish larvae produce briefly oscillatory swimming bouts (Orger et al. 2008). Thus, rhythmic activity is a common motif across multiple motor systems that use very different neural architectures. Given the ease with which different neural systems (real and artificial) produce oscillatory patterns, and given the empirical effectiveness of an oscillatory basis set in building muscle activity, a seemingly straightforward hypothesis is that such a mechanism might be at play during reaching.

It should be stressed that a quasioscillatory basis set is not the only (or even necessarily the optimal) solution. The generation of pulsatile patterns—as in the traditional model—is also common across many motor systems and could provide a basis set. And although the above arguments may suggest a quasioscillatory set of basis patterns,

by no means do they necessitate that solution. It was therefore an empirical question whether the data during reaching resembled the traditional model, the dynamical model, or neither.

Dynamical Model Responses Resemble Neural Response

We recorded neural activity from motor and premotor cortex during a variety of reaching tasks (Churchland et al. 2006; Kaufman et al. 2010). Neurons showed a broad range of response patterns (Fig. 4) and rarely resembled the predictions of the traditional model (Churchland and Shenoy 2007; Churchland et al. 2010, 2012). There was little evidence of a “preferred movement.” Indeed, the

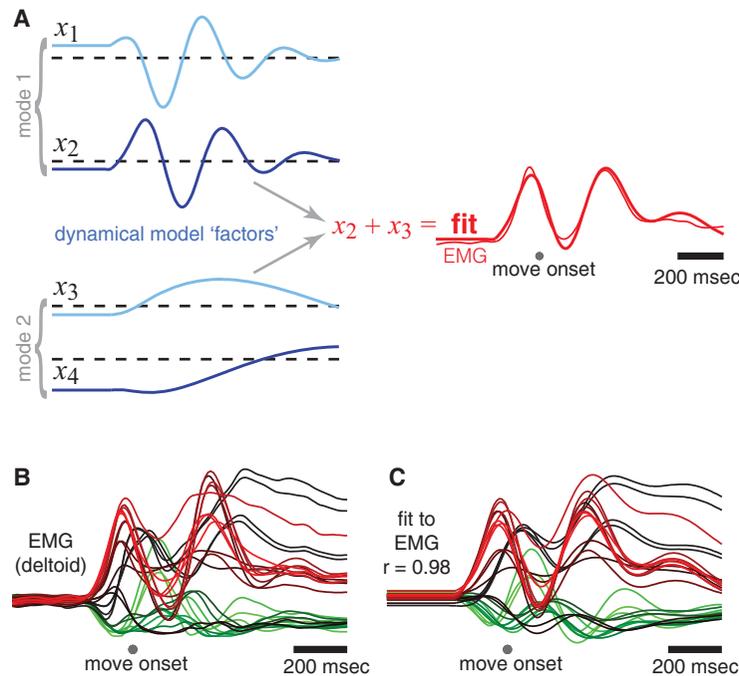


Figure 3. Constructing muscle activity from an oscillatory basis set. (A) The time-evolving factors proposed by the dynamical model. The first pair of factors (x_1 and x_2) capture the two phases of a 2.8-Hz oscillation. The model employs both phases because it is the most natural way for a dynamical system to generate an oscillatory response. The second pair of factors (x_3 and x_4) capture the two phases of a 0.3-Hz oscillation. For both frequencies the oscillations are short-lived; the 2.8-Hz oscillation lasts only a little over one cycle and the 0.3-Hz oscillation lasts less than one cycle. The sum of factors x_2 and x_3 produces a pattern (thick red trace) that approximates the empirical pattern of electromyography (EMG) recorded from the deltoid (thin red trace). Data is shown for one condition (9) of 27 total. Oscillation phase and amplitude were set by the “preparatory state.” The preparatory state is reflected in the fact that the factors are offset from baseline (dashed line) before the onset of pattern generation. For each condition we chose the preparatory state to produce the phases and amplitudes that best fit the EMG. (B) EMG recorded from the deltoid for all 27 conditions. Each trace plots the intensity of the muscle activity for one condition. Red-to-green coloring of the traces is used simply to allow the eye to follow individual traces. (C) Fits to the EMG via the dynamical model, as in A, for all conditions. Correlation between model fit and recorded EMG was 0.98. (B,C, Adapted from Churchland et al. 2012.)

condition that garnered the larger response was typically different at different times during the trial (e.g., sometimes *red traces* are higher and sometimes *green traces* are higher). Thus, individual-neuron responses do not particularly resemble the predictions of the traditional model.

We then compared neural responses with responses predicted by the dynamical model. Simulated neural responses were based on the model in Figure 3A, which was optimized to generate the empirical electromyography (EMG). The response of each simulated unit was a randomly weighted sum of the patterns present in the model, plus a small amount of noise (Churchland et al. 2012). The level of preparatory activity was determined by the initial plateau (left side of Fig. 3A). This plateau of preparatory activity is critical to the model: It sets the phase and amplitude of the subsequent oscillations ($\mathbf{x}(0, c)$ in Eq. 2). The duration of simulated preparatory activity was set to roughly match that seen for the neural responses.

Dynamical model responses showed a variety of multiphasic patterns that resembled those of cortical neurons (Fig. 4). Preparatory activity was present, yet had an unclear relationship with the subsequent pattern of movement activity. For the model this occurs because the relationship between preparatory state and rotation phase becomes clear only at the level of the population neural

state. In summary, the dynamical model, but not the traditional model, produces single-unit responses that resemble the data. Importantly the dynamical model units in Figure 4 are not fit to recorded neurons; the resemblance naturally arises when the model is fit to the empirical patterns of EMG (Fig. 3A). The next key question was thus: Does the similarity between the data and the dynamical model extend to the population level?

DYNAMICAL STRUCTURE IN THE POPULATION RESPONSE

Model Predictions

The dynamical model makes three key predictions at the population level. First, the population state should evolve according to dynamics described by Equation 2, which can be linearly approximated as

$$\dot{\mathbf{x}}(t, c) = M\mathbf{x}(t, c). \quad (3)$$

It is proposed that dynamics are very “smooth”: They are similar across conditions and can thus be approximated with a single matrix M . The second prediction is that the dynamics in question should have a large oscillatory/

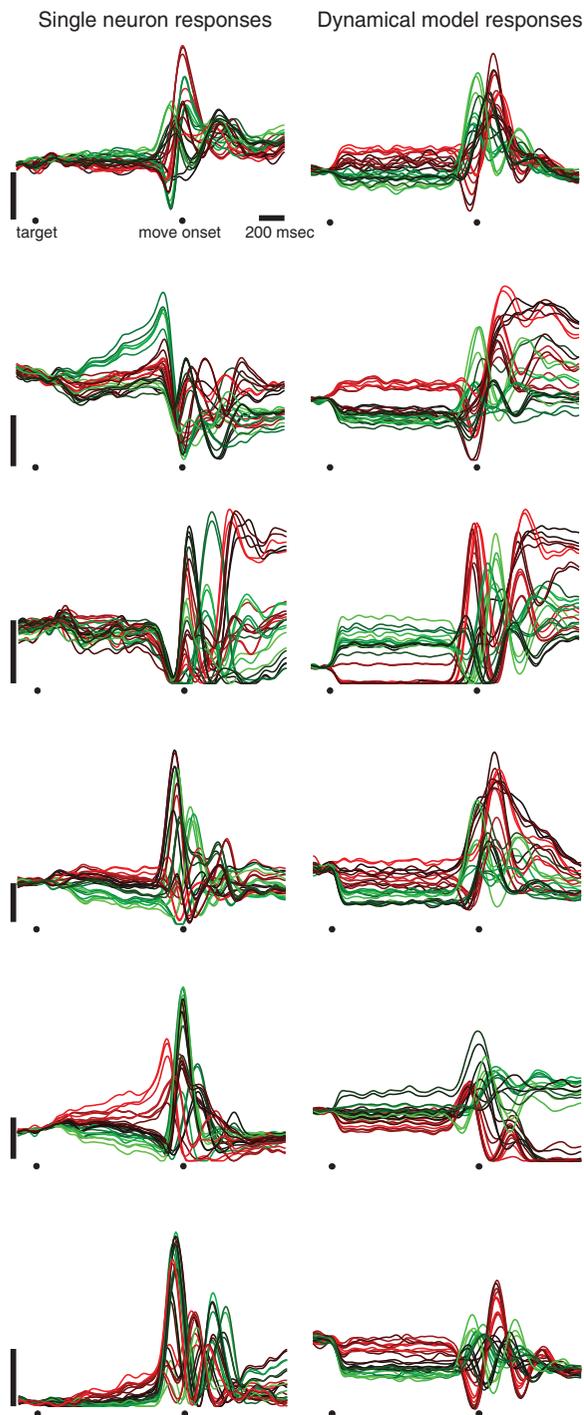


Figure 4. Example responses of real (*left* column) and simulated units from the dynamical model (*right* column). Each subpanel plots the response for one neuron/unit. Each trace plots the response for one reach condition. Vertical scale bars indicate 20 spikes/sec. Traces are colored red to green for visualization. So that the same color coding can be used for all panels, color-coding is based on the initial strength of deltoid EMG, just before movement onset. Note that although the real and simulated data share a number of interesting general features, the individual units/neurons are not meant to (and do not) map directly onto one another. The model responses were never fit to the neural responses—they are a consequence of how the model reproduced the empirical EMG.

rotational component: M should be fairly close to skew-symmetric ($M = -M^T$). Thus, the fit to the data should be almost as good if M is constrained to be skew-symmetric. Third, it should be possible to directly view projections of the data that reveal the rotational structure illustrated in Figure 2E. If there exist factors x_1 and x_2 with the hypothesized structure, then it should be possible to recover those factors by taking an appropriately weighted sum of the activity of all the neurons. Doing so should reveal state-space patterns with clear rotations.

Before examining these predictions, we stress that the use of a weighted sum to yield interpretable “factors” is common in systems neuroscience (Cunningham and Yu 2014)—particularly when decoding represented parameters. For example, in the traditional model, one can recover a representation of reach direction by assigning each neuron a pair of weights based on its preferred direction. Thus, x_1 (x_2) would be the sum of the activities of all neurons, weighted according to the degree to which that neuron preferred rightward (upward) reaches. When examining the predictions of the dynamical model, we wished to accomplish something similar, but to recover projections that might relate to rotational dynamics.

We used principle component analysis as a preprocessing step and reduced the dimensionality of the data to a few factors (typically six). We then fit the temporal evolution of the data with Equation 3. Fit quality (R^2) was assessed using both an unconstrained matrix M , and a matrix M_{skew} that was constrained to be skew-symmetric. Finally, we recovered factors x_1 and x_2 via weights based on the eigenvectors of M_{skew} (the jPCA method; Churchland et al. 2012). We used the eigenvectors that corresponded to the highest-frequency rotations. As a technical aside, we used the eigenvectors of M_{skew} rather than M for three practical reasons. First, we wished to test for the presence of rotations in a subset of dimensions, even if other dimensions contained other patterns. Second, M_{skew} is “normal” and conveniently yields orthogonal eigenvectors. Third, in practice M and M_{skew} were typically very similar for the neural data.

Rotations Are a Dominant Motif in the Population Response

When applied to the population of dynamical model responses (e.g., Fig. 4) the jPCA method accurately recovered the true factors (compare Fig. 5A,B). As expected, the quality of the dynamical fit (R^2) was high and was nearly as good when constrained to be purely rotational (R^2_{skew}). For the neural data (Fig. 5C,D) jPCA revealed a similar rotational structure (Churchland et al. 2012). The quality of the dynamical fit was high (although certainly not as high as the idealized model) and was nearly as good when constrained to be purely rotational. Rotation phase and amplitude followed from the preparatory state for both the model and the data. Both model and data contained a second plane with lower-frequency rotations (not shown, although see Fig. 7B,C below). Thus, the three population-level predictions of the dynamical mod-

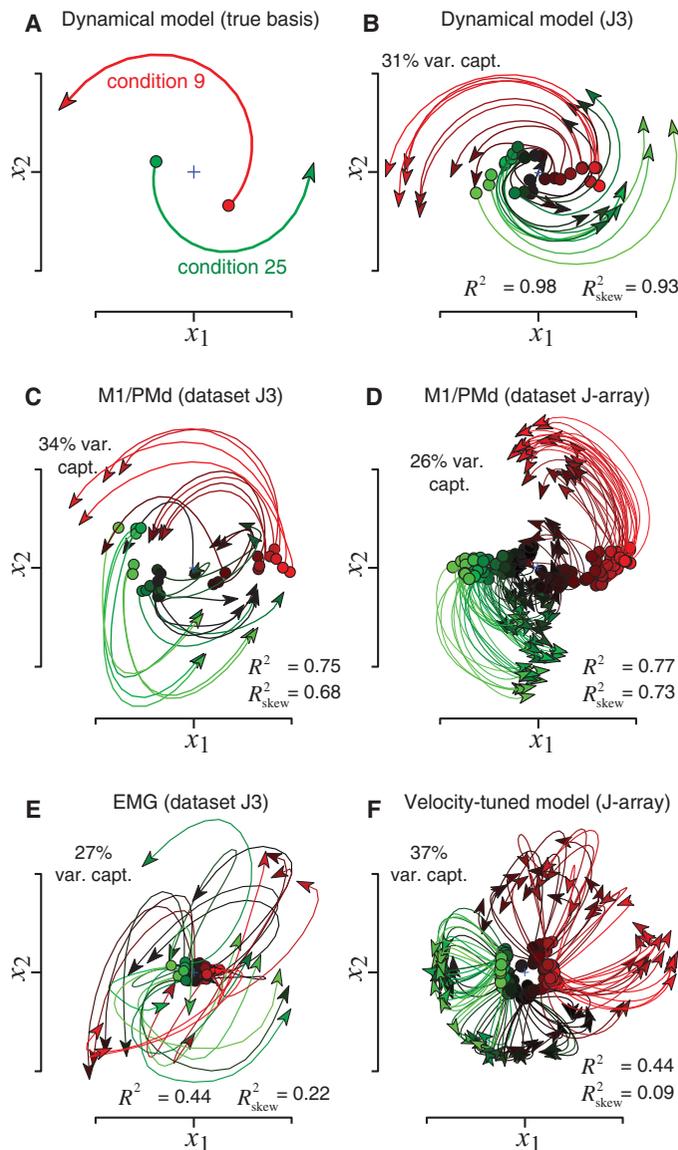


Figure 5. Population structure for real and simulated data. All plots show 200 msec of the evolution of the neural state, beginning just as the preparatory activity (dots) transitions to movement activity (subsequent trajectories). (A) The first two dimensions of the true basis set used by the dynamical model. For simplicity two conditions are shown. (B) Projection, via jPCA, of the simulated population response for the dynamical model. The model was fit to produce the EMG of the deltoid from data set J3, which was recorded across 27 conditions. For visualization, traces are colored based on the projection of the preparatory state onto x_1 . The jPCA method accurately captures the underlying basis from A. Values of R^2 and R^2_{skew} indicate the fit quality provided by a linear model (unconstrained or skew-symmetric respectively) for the two dimensions shown. The variance captured (var. capt.) indicates the total data variance that lies within these two dimensions, as with PCA. (C) jPCA projection of the population response for data set J3 (27 conditions). (D) jPCA projection of the population response for data set J-array (108 conditions). (E) jPCA projection of the muscle population for data set J3. (F) jPCA projection of the simulated population response for a traditional model in which neural responses were tuned for end point position and upcoming peak speed during preparation and tuned for velocity and absolute speed during movement. Traditional model responses were based upon the kinematics recorded for data set J-array. (All panels adapted from Churchland et al. 2012.)

el were observed: (i) the temporal evolution of the data could be fit by a dynamical model, (ii) the fit was almost as good when dynamics were constrained to be rotational, and (iii) the population response could be projected to reveal factors with clear rotations that followed from the preparatory state.

The comparison between neural and model responses in Figure 5 is an instance of “hypothesis-guided dimension-

ality reduction.” We hypothesized the presence of rotational structure and designed a method (jPCA) to seek factors that might display such structure. But to what degree might this constitute a self-fulfilling prophecy: Would rotational structure be seen in any data with certain basic surface-level features (e.g., smoothly changing multiphasic responses)? Might the traditional model also show such rotational structure when analyzed via jPCA?

Rotational Structure Is Absent for Muscles and Traditional Models

For most neural data sets we also recorded the corresponding population of muscle responses. Single-muscle responses were complex and multiphasic (e.g., Fig. 3B) much like single-neuron responses. Single-muscle responses varied smoothly across time and conditions, much like single-neuron responses. Yet the muscle populations did not contain factors with strong rotational structure. Projections of the EMG population (Fig. 5E) often contained looping trajectories, but these did not travel in a consistent direction: The circular flow of some conditions/times was counteracted by a different flow for other conditions/times and was thus inconsistent with the dynamical model. In agreement, the overall dynamical fit was weaker for the muscle population than for the neural populations, ($R^2 = 0.44$ vs. 0.75 and 0.77) and suffered more when constrained to be purely rotational ($R_{\text{skew}}^2 = 0.22$ vs. 0.68 and 0.73). Thus, it is not the case that any population of complex responses will show rotational structure; muscle responses share most surface-level features with the neural data yet did not show consistent rotations. This relates to a key point to which we return below: Although the goal of the dynamical model is to produce muscle activity, the model contains multiple patterns other than the final output, and model neurons are therefore not expected to appear “tuned” for muscle activity.

We also simulated population responses for traditional models where activity is a function of externally measurable parameters (Eq. 1). These population responses did not show strong rotational structure (Fig. 5F), but instead tended to be dominated by the expansion of the neural state. This was true for both simple versions of the traditional model (tuned for reach end point during preparation and reach velocity and speed during movement) and for more complex versions that included tuning for position, acceleration, and jerk (Churchland et al. 2012). This analysis acts as a further control: Again, it is not the case that jPCA reveals rotations for most any data set. More importantly the analyses in Figure 5 reveal that a basic feature of the population response—rotations that follow from the preparatory state—is consistent with the dynamical model but not with a pure muscle-tuning hypothesis or with the traditional model.

In principle, it might be possible to adapt traditional models to produce rotations. For example, a model tuned purely for horizontal position and velocity would produce a partial rotation in state space, if those representations were separable into two roughly orthogonal factors. A more extended rotation could result from a representation of acceleration and jerk. Yet the neural data contain features that are not naturally accounted for by such models. First, rotations during movement follow naturally from the preparatory state, something that would not be true if plotting a “jerk representation” versus an “acceleration representation.” Second, rotation frequency was ~ 1.5 – 2 Hz, considerably slower than the profile of acceleration or jerk. Some movements had durations of < 150 msec,

which would yield a triphasic jerk pattern at ~ 10 Hz—much faster than observed. Third, the observed rotation frequency did not depend on reach speed/movement duration (Churchland et al. 2012) in opposition to the expectation if jerk or acceleration were being represented. Fourth, rotation direction was the same regardless of whether the actual reach was straight, curved clockwise, or curved counterclockwise. These features all emerge naturally from the dynamical model. It might be possible to adapt the traditional model to also account for these features, but it is not clear how to do so.

TRADITIONAL AND DYNAMICAL MODELS SUGGEST DIFFERENT EXPLANATORY APPROACHES

Interpretation of “Represented” Parameters

The traditional model assumes that factors constitute “representations” of parameters such as reach velocity or speed (e.g., Pearce and Moran 2012). A neuron’s response should therefore be related to the parameters with which it covaries most strongly. For example, if a given neuron’s response correlates weakly with muscle activity and strongly with reach velocity, then that would be taken as evidence that the response relates principally to velocity. Under the dynamical model such inferences are fraught. Consider four hypothetical neurons whose responses respectively reflect x_1 through x_4 from the model in Figure 3A. Note that x_3 and x_4 resemble the pattern of EMG rather weakly: x_3 resembles a low-pass filtered version of velocity, and x_4 resembles position. A general feature of the dynamical model is that individual basis elements may often weakly resemble the final output and may appear to represent other incidental factors. This can be seen as the extension, into the temporal domain, of the observation that in a highly redundant system a neuron’s response need not match its causal impact on the final output (Baraduc et al. 2001; Todorov 2002; Schieber and Rivlis 2007).

The difference in the interpretations suggested by the traditional and dynamical models is heightened if there are many observed response patterns. Traditionally, a plethora of patterns suggests the representation/control of many parameters (e.g., Ashe and Georgopoulos 1994; Fu et al. 1995; Wang et al. 2007; Pearce and Moran 2012). Under the dynamical hypothesis, a plethora of patterns suggests a rich basis set (Fetz 1992; Churchland and Shenoy 2007; Churchland et al. 2010, 2012; Sussillo et al. 2012; Shenoy et al. 2013; Hennequin et al. 2014).

Interpretation of Why “Tuning” Varies with Other Factors

Neural responses are often analyzed in terms of “tuning”: the manner in which firing rates differ across conditions. Under the traditional framework, tuning is fundamental and reflects the hypothesized relationship between factors and represented parameters. A large literature has sought to identify those parameters to which

tuning relates most reliably, on the assumption that those are the parameters that are truly represented. For example, an extensive series of elegant experiments from multiple laboratories has asked whether tuning is constant with respect to Cartesian reach direction or with respect to other factors such as force and/or muscle activity (for review, see Scott 2008; Kalaska 2009; Reimer and Hatsopoulos 2009). Almost universally, such experiments reveal that tuning is not constant with respect to any of the proposed parameters. Tuning either varies idiosyncratically or undergoes “intermediate” shifts. More generally, tuning for a given factor such as direction is rarely constant with respect to other factors such as reach speed (Churchland et al. 2006; Churchland and Shenoy 2007). For example, Figure 6A shows the mean preparatory response of a neuron whose tuning for each direction de-

pendent on reach distance (top and bottom panels). This might initially seem troubling. Is the underlying representation very complicated? Are responses simply “messy” and unreliable?

Under the dynamical model, inconstant tuning is expected. “Tuning” is merely a consequence of different movement conditions requiring the elements of the basis set to be modulated differently. For example, every trace in Figure 3C was built from the same oscillatory basis set, but with different phases and amplitudes. Because phase and amplitude depend on the preparatory state (Fig. 2E), the preparatory state necessarily varies across conditions. However, such tuning is not constant with respect to other factors. One example is shown in Figure 6B: The model unit changes its preparatory tuning with distance, much as has been regularly observed in the data. The dynamical

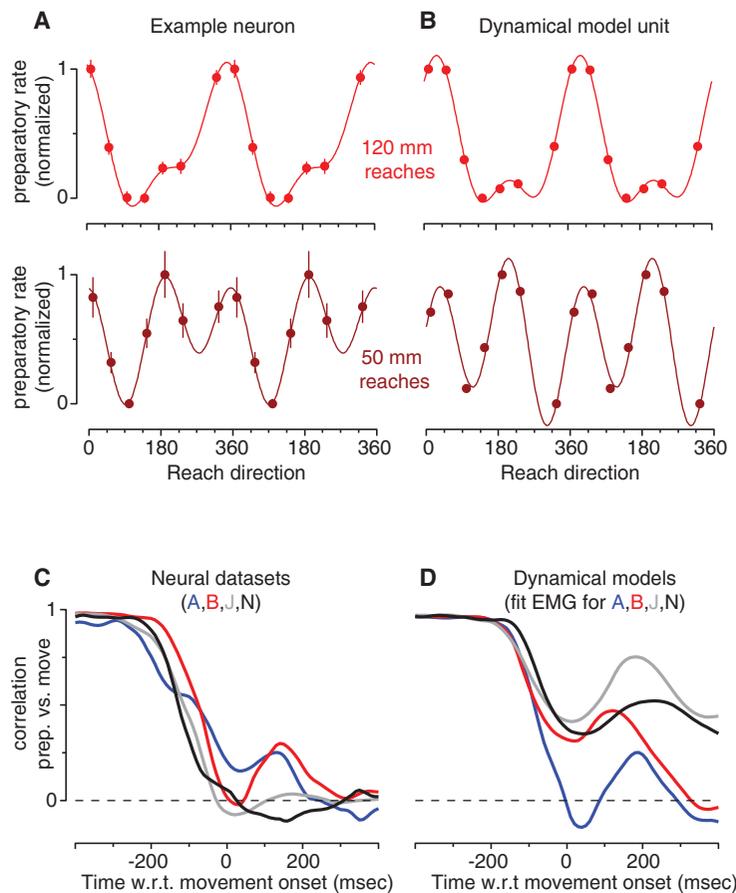


Figure 6. Both real neurons and dynamical model units undergo idiosyncratic changes in the preferred direction as a function of reach distance and as a function of time. (A) The average preparatory firing rate of a single neuron as a function of reach direction, recorded during a center-out reaching task. Error bars give the SEM. Fits use a cosine-tuning model employing both the fundamental (unimodal tuning) and the first harmonic (bimodal tuning). (B) Similar plot but for the preparatory activity of a unit simulated according to the dynamical model, which was trained to produce the EMG of the deltoid. (C) Quantification of the degree to which neurons “preferred” the same conditions during the movement as they did during the preparatory period. The correlation between preparatory and peri-movement tuning is plotted as a function of time. All correlations are between the time shown and a reference time during the middle of the preparatory period. The response of a given neuron at the reference time is a vector containing the firing rate for every condition. For each analyzed time, we correlated the vector of firing rates at that time with the vector at the reference time. This correlation was averaged across all neurons within the data set. Analysis was restricted to neurons robustly tuned during both epochs. Correlations are initially near unity, in agreement with the observation that tuning is stable during the preparatory period. Around 175 msec before movement onset the average correlation began to decrease sharply, indicating that the preference during the movement typically bore little resemblance to that during the preparatory period. (D) Same analysis as above but for data sets simulated using the dynamical model. (C, Adapted from Churchland et al. 2010.)

model thus provides a rather different account from that suggested by the traditional model, where preparatory tuning constitutes a representation of factors such as reach direction or target location (e.g., Shen and Alexander 1997; Pearce and Moran 2012). For the dynamical model “tuning” will rarely have a consistent surface-level relationship with any movement parameter, even those directly controlled by the network. These observations extend an important fact previously noted in the literature: Any system that controls movement will necessarily have activity that varies across conditions, resulting in directional “tuning” that may or may not be consistent (Mussa-Ivaldi 1988; Fetz 1992; Sanger 1994; Todorov 2000, 2002; Cisek 2006b; Ajemian et al. 2008; Kalaska 2009; Reimer and Hatsopoulos 2009; Lillicrap and Scott 2013).

Interpretation of Why “Tuning” Changes with Time

A striking aspect of most single-neuron responses is that tuning changes with time (in addition to the tuning inconstancy across conditions, as noted above). The “preferred condition” during the preparatory period tends not to agree with the preferred condition during movement (Crammond and Kalaska 2000; Churchland et al. 2010). Furthermore, the preferred condition early in the movement tends not to agree with the preferred condition later in the movement (Sergio and Kalaska 1998; Churchland and Shenoy 2007; Hatsopoulos et al. 2007). For example, in Figure 4 the red-to-green ordering of traces is not preserved as time passes. Figure 6C plots the average correlation of “tuning” between each time and a reference time during the preparatory interval. The correlation plummets just before movement begins and approaches zero by the time of movement onset. Thus, the cartoon neuron illustrated in Figure 2B is rarely observed. Under the traditional model these findings suggest that different movement features are represented in different ways at different times (e.g., Fu et al. 1995; Pearce and Moran 2012). An alternative explanation has been that consistent tuning is neither expected nor desired of a system whose function is to generate and control movement (Fetz 1992; Todorov 2000; Cisek 2006b; Kalaska 2009; Reimer and Hatsopoulos 2009; Lillicrap and Scott 2013; Shenoy et al. 2013).

The dynamical model accords with this latter interpretation and supplies a straightforward explanation for tuning that varies with time. The oscillatory basis intrinsically results in a “preference” that is always in flux. For example, in Figure 2C the red trace is sometimes higher and the green trace is sometimes higher. Like the data, the model also shows a declining correlation between preparatory and movement-period “tuning” (Fig. 6D).

Projections of the Population

Different perspectives on tuning lead to different perspectives regarding how to summarize the population response. In the traditional view, underlying factors relate to represented movement parameters and should be

recovered via reference to those parameters. As one example, a two-dimensional population vector can be constructed by weighting each neuron’s preferred direction by its response and taking the vector sum. More formally, assume the encoding model:

$$\mathbf{r} = D\mathbf{x},$$

where \mathbf{r} is an $n \times 1$ vector of neural responses, \mathbf{x} is a 2×1 vector of factors describing horizontal and vertical velocity, and D is an $n \times 2$ matrix of preferred directions. The underlying factors in \mathbf{x} can be extracted via

$$\mathbf{x} = D^\dagger \mathbf{r},$$

where D^\dagger is the pseudoinverse of D . When D has columns that are close to orthonormal, this can be approximated using the transpose:

$$\mathbf{x} = D^T \mathbf{r}.$$

A population vector produced via this method is shown in Figure 7A for 28 reach conditions (7 directions \times 2 distances \times 2 reach speeds). This view captures a considerable proportion of the overall variance (20%). Still, although an expansive component is present, it is somewhat disorganized and there is little clear relationship between preparatory and movement activity across conditions. Thus, although Figure 7A is a valid view of the population response, one wonders whether other reasonable views also exist and might be at least as informative.

Under the dynamical perspective, the most revealing projections of the population will be those that capture the dynamical relationship between factors. One particularly desires projections in which a flow field describes the evolution of the neural state. jPCA pursues this end by fitting with a dynamical model. The data are then projected via

$$\mathbf{x} = V\mathbf{r},$$

where each row of the matrix V is an eigenvector of the matrix that fit the dynamics. The number of rows of V can be chosen based on the number of eigenvectors with large corresponding eigenvalues. We found that the data typically contained two to three planes of rotations (4–6 total dimensions) as illustrated in Figure 7B,C (same data set as in Fig. 7A). The first two planes capture 56% of the variance and reveal rotations at different frequencies. Such structure accords with the dynamical model in which a basis set is built from a small handful of oscillatory modes

Pitfalls of the “Population PSTH”

It should be stressed that the three views in Figure 7A–C are all linear transformations of the underlying population response. One might prefer one view to

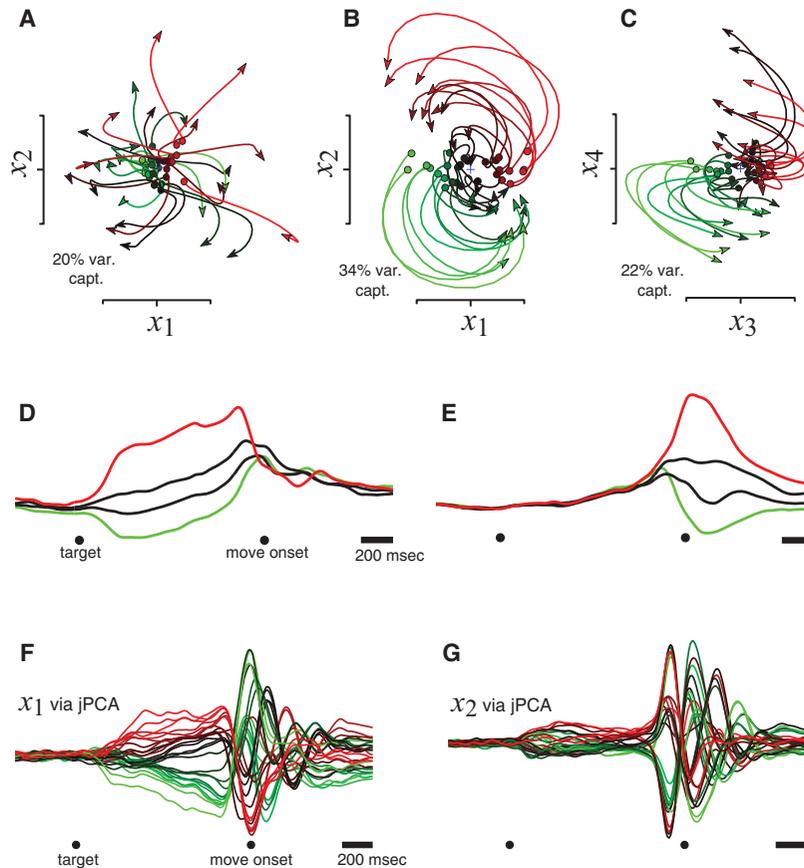


Figure 7. Various views of the population response. For state-space plots (A–C) all conventions are as in Figure 5. (A) The population vector for a data set which was recorded during straight reaches to seven radially arranged targets at two different reach speeds (data from monkey B). (B) The first pair of dimensions of the jPCA projection for that same data set. (C) The second pair of dimensions of that jPCA projection. (D) Population peristimulus time histograms (PSTHs) based on the preferred condition during the preparatory period (monkey J). Traces plot the average firing rate (across all neurons) for the condition with the strongest preparatory response (red), the condition with the weakest preparatory response (green), and two intermediate conditions (intermediate shades). Such conditions were found separately for each neuron. Responses were then averaged after normalizing each neuron’s activity by its maximum rate (results were similar without normalization). The resulting population PSTH displays very little movement period activity, even though such activity was strongly present. Indeed, analysis was restricted to include only neurons with robust preparatory tuning and robust movement-period tuning. (E) A similar population PSTH but with the preferred condition found during the movement period. Despite being based on exactly the same data as in D, the population PSTH now looks very different: There appears to be little preparatory activity. (F) First dimension (x_1) of the projection found via jPCA for data set J3, plotted as a function of time. Each trace corresponds to one condition. (G) As in F, but for the second jPCA dimension (x_2). (B, Adapted from Churchland et al. 2012.)

another depending on one’s hypothesis and on the structure and variance captured by that view. Yet all are valid views of the data, and each reveals response aspects not captured by the others. However, there exist common analysis methods, often inspired by the traditional model, that will not in general produce valid projections of the population response. For example, it is common to summarize the temporal evolution of the population response by computing the “population PSTH (peristimulus time histogram)”: the average activity of every neuron, each for its own “preferred” condition. Thus, the activity of a rightward-preferring neuron for a rightward reach is averaged with the activity of an upward-preferring neuron for an upward reach. This procedure is justified only if the tuning model is accurate and behavior is symmetric (e.g., if every reach trajectory is similar excepting the direction itself). Otherwise the resulting view is not a

projection of the population response and may be difficult to interpret.

To illustrate, Figure 7D plots the population PSTH based on the preferred condition during the preparatory period. The red trace is the average response of all neurons for their most-preferred condition. The green trace corresponds to the least-preferred condition, and dark traces show intermediate preferences. The population PSTH shows a small peak before movement onset, after which the four traces become similar and relatively flat. The plot thus fails to reflect the fact that most neurons are strongly “tuned” and responsive at this time. That structure is lost in the population PSTH because the preferred condition during movement and preparation are weakly correlated (Fig. 6C). The population PSTH looks very different if the preferred condition is measured later, during the movement (Fig. 7E). Indeed, peaks in tuning can be produced at

different times depending on exactly when the preferred condition was assessed. Thus, the seemingly sensible and oft-used approach of averaging responses based on their preferences can produce tautological conclusions: A population will appear strongly tuned at a particular time simply because it was sorted accordingly.

The dynamical perspective suggests alternative approaches for examining the population response versus time. For example, one can produce “factor PSTHs” in which each factor is plotted versus time across all conditions (Fig. 7F,G). These projections of the neural data reveal response structure that is lost in the traditional population PSTH. In particular, these projections reveal that a rough plateau of preparatory activity leads to a strongly oscillatory response during the movement. This pattern—also often observed at the single-neuron level—is lost in the population PSTH.

The Dynamical Model: Toward a Mechanistic Hypothesis

Under the dynamical perspective, the “meaning” of the neural state can seem abstract because most factors do not

relate to externally measurable parameters. However, the dynamical perspective provides a natural language for expressing high-level mechanistic hypotheses. The model in Figure 3A is not a circuit-level model, but does propose a mechanism by which the circuit as a whole might perform its function. Dynamical hypotheses can often be succinctly expressed in state space, where they become hypotheses regarding the flow field (Fig. 8B)—that is, the influence of factors on each other. Because many of those factors have no external meaning, they can be interpreted in terms of their mechanistic roles in producing the final output. This style of explanation is often useful even when the underlying connectivity is known, as is the case for many model networks (Sussillo and Barak 2013). Furthermore, this approach promotes the generation of hypotheses that might otherwise lack a clear language. Two examples follow.

A long-standing question has been how preparatory activity avoids causing movement. For the traditional model, it is usually suggested that preparatory activity lies below threshold or is held at bay by an inhibitory gate. The dynamical model suggests a different hypothesis: that preparatory activity lives in the null space of the

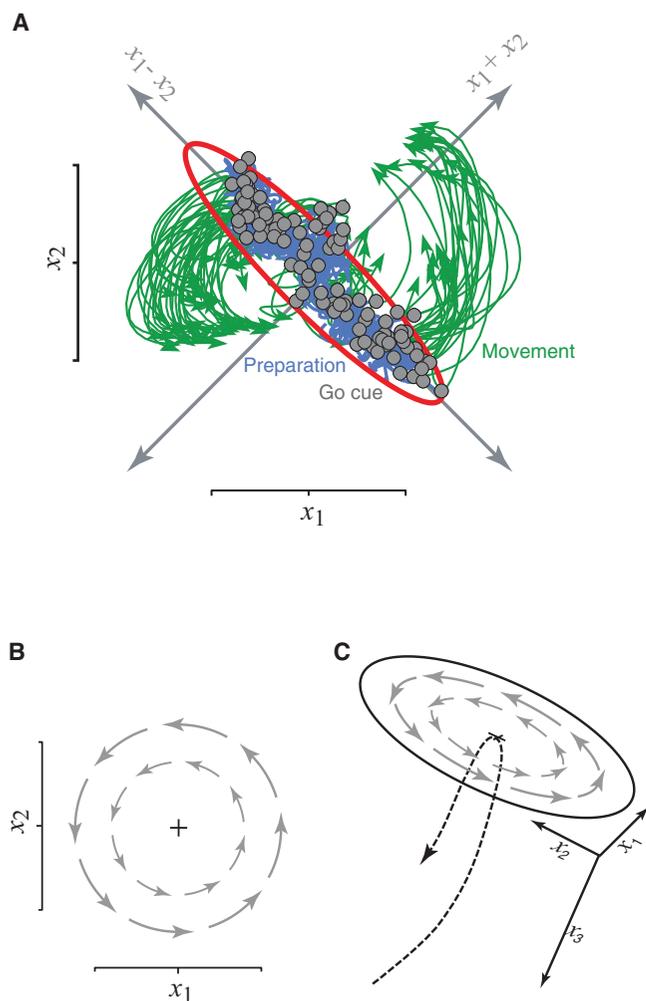


Figure 8. Dynamical hypotheses expressed in state space. (A) Illustration, using data, of the “null space hypothesis.” The projection is of the population from data set J-array (as in Fig. 5D but with slightly different preprocessing as it was computed for a different study). Preparatory activity is shown in blue; the go cue was given at the time of the gray circles, after which the green traces reveal the trajectory of movement-period activity. The red ellipse shows two standard deviations of the covariance ellipse describing the spread of preparatory activity. (B) Hypotheses regarding the interactions between factors can be expressed as flow fields in state space. (C) Illustration of how a factor might play a “higher-level” role. In this case, a translation in direction x_3 could bring the neural state into a region of strong oscillatory dynamics, and thus initiate pattern generation. In this case, one dimension (x_3) would be used to recruit pattern generation in a space spanned by other factors (x_1 and x_2). (A, Adapted from Kaufman et al. 2014.)

dimensions that map to output muscle activity. This is indeed true of the dynamical model outlined in Figure 3. Preparatory activity is present in the model—no factor is at baseline before the oscillation begins—and plays the role of setting the phase and amplitude of the subsequent oscillation. Yet preparatory activity cancels when the factors are summed to produce the output. Thus, the output is virtually identical across conditions during the preparatory period (Fig. 3C) even though most units have preparatory activity (Fig. 4). Figure 8A illustrates a simple view of this proposal overlaid on a projection of neural data. If preparatory activity (gray dots) is primarily extended in the $x_1 - x_2$ direction then it has little projection onto the $x_1 + x_2$ direction. A downstream target driven by both x_1 and x_2 will mostly ignore preparatory activity. Such an effect is indeed present: Those dimensions that most resemble muscle activity have recently been shown to have the least preparatory activity (Kaufman et al. 2014).

A second outstanding question regards the neural events that trigger movement onset. If activity does not rise past a threshold (as in Fig. 2B), then what triggers movement onset? The model in Figure 2C,E assumes—and the data confirms—the sudden onset of rotational dynamics just before movement onset. We recently hypothesized (Kaufman et al. 2013) that this sudden change in dynamics may be driven by a large overall translation in state space (Fig. 8C). This hypothesis is plausible because local dynamics could be very different in different regions in state space, as suggested by recent modeling (Sussillo et al. 2012). By exploiting such a mechanism, the circuit could control the neural dynamics that drive movement. Hypotheses such as this are readily conceived in a dynamical framework, but difficult to express otherwise.

CONCLUSIONS

We compared predictions of two idealized models. We considered a traditional model, in which a movement preference appears during the preparatory period and becomes stronger as movement is generated. The basic response features of this model were at odds with the patterns observed in single-neuron responses (Fig. 4). In contrast, a dynamical model produced reasonably realistic single-neuron responses. At the population level, the dynamical model predicted the existence of rotations in state space. Such rotations were also prevalent in the neural data. This prevalence of rotations was not a consequence of the application of the jPCA method: That same method revealed weak rotational tendencies in the population of muscle responses and in simulated population responses of traditional models. More generally, the linear dynamical fit to the neural data was dominated by a rotational component, something not true for the traditional model or the muscle population.

The dynamical model naturally accounts for some seemingly confusing response properties, including the tendency of neurons to display “tuning” that changes with

time and with other parameters such as distance and speed (Fig. 6). That said, the dynamical model is currently quite incomplete, and the instantiation shown in Figure 2E and Figure 3A is highly idealized. Future work will have to explain how similar dynamics can be instantiated in neural networks (Schaffer et al. 2006; Sussillo et al. 2012; Hennequin et al. 2014), how dynamics are “turned on” at the right time, the degree to which dynamical motifs other than rotations are important, and the degree to which dynamics are driven by internal versus feedback connections. It also remains an open question why rotations are a dominant motif when other basis sets are possible. One possibility is that neural mechanisms for locomotion were co-opted, on an evolutionary timescale, to produce other kinds of motion (Rokni and Sompolinsky 2012). Alternatively, an oscillatory basis set may be used simply because it is easily generated and provides a robust solution (Sussillo et al. 2012).

In summary, although the dynamical model might initially seem less intuitive, it accounts for the data more successfully than the traditional model. The dynamical model explains a number of features of the data that otherwise seem troublesome or “messy.” The model provides a high-level mechanistic explanation for how the system might generate the desired outputs: The output is hypothesized to be the sum of a basis set of patterns that can be generated via internal dynamics. In this view, preparatory activity is truly “preparatory”: Its goal is to set the initial state of those dynamics, so that the basis set is recruited appropriately. The subsequent patterns of movement activity may appear complex, but they have a straightforward goal: to provide the building blocks of a temporally structured output.

The above hypothesis is readily extensible. For example, the true basis set is surely more complex than two simple oscillations. Indeed, the data often contain three planes of rotation and significant nonrotational components as well. Thus, rotations may be a dominant motif but they are not the only motif. Furthermore, although an oscillatory basis set will be appropriate for many movements (e.g., those involving multiphasic EMG), other bases will presumably be appropriate for other classes of movement. Indeed, the heart of the dynamical hypothesis is not the particular basis set being used, but the broader strategy at play. This includes three key elements: the presence of more neural dimensions than output dimensions, dynamics that govern the evolution of the neural state within those dimensions, and preparatory activity that plays the role of initializing the dynamics. The dynamical view provides the language and tools for exploring how variations upon this theme may allow motor cortex to generate different classes of movement.

ACKNOWLEDGMENTS

We are very grateful to our colleagues with whom the original research was performed: Krishna Shenoy, Matthew Kaufman, Maneesh Sahani, Stephen Ryu, David Sussillo, Justin Foster, and Paul Nuyujukian. We

are deeply grateful to our theoretical colleagues whose work formed much of the inspiration for the dynamical model, in particular Larry Abbott, David Sussillo, and Evan Schaffer. We would also like to acknowledge recent funding sources that have allowed this work to continue: The Grossman Center for the Statistics of Mind, The Burroughs Wellcome Fund, The Searle Scholars Program, The Simons Foundation (SCGB#325171 and SCGB#325233), The Sloan Foundation, The McKnight Foundation, and The National Institutes of Health (NIH) Director's New Innovator Award Program (DP2 NS083037), the Kavli Foundation, and the Gatsby Charitable Trust.

REFERENCES

- Ajemian R, Green A, Bullock D, Sergio L, Kalaska J, Grossberg S. 2008. Assessing the function of motor cortex: Single-neuron models of how neural response is modulated by limb biomechanics. *Neuron* **58**: 414–428.
- Ashe J, Georgopoulos AP. 1994. Movement parameters and neural activity in motor cortex and area 5. *Cereb Cortex* **4**: 590–600.
- Baraduc P, Guigon E, Burnod Y. 2001. Recoding arm position to learn visuomotor transformations. *Cereb Cortex* **11**: 906–917.
- Bastian A, Riehle A, Erlhagen W, Schöner G. 1998. Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport* **9**: 315–319.
- Bastian A, Schöner G, Riehle A. 2003. Preshaping and continuous evolution of motor cortical representations during movement preparation. *Eur J Neurosci* **18**: 2047–2058.
- Cheney PD, Fetz EE. 1980. Functional classes of primate corticomotoneuronal cells and their relation to active force. *J Neurophysiol* **44**: 773–791.
- Churchland MM, Shenoy KV. 2007. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol* **97**: 4235–4257.
- Churchland MM, Santhanam G, Shenoy KV. 2006. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J Neurophysiol* **96**: 3130–3146.
- Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. 2010. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron* **68**: 387–400.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012. Neural population dynamics during reaching. *Nature* **487**: 51–56.
- Cisek P. 2006a. Integrated neural processes for defining potential actions and deciding between them: A computational model. *J Neurosci* **26**: 9761–9770.
- Cisek P. 2006b. Preparing for speed. Focus on “Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach”. *J Neurophysiol* **96**: 2842–2843.
- Crammond DJ, Kalaska JF. 2000. Prior information in motor and premotor cortex: Activity during the delay period and effect on pre-movement activity. *J Neurophysiol* **84**: 986–1005.
- Cunningham JP, Yu BM. 2014. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* **17**: 1500–1509.
- Erlhagen W, Bastian A, Jancke D, Riehle A, Schöner G. 1999. The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *J Neurosci Methods* **94**: 53–66.
- Evarts EV. 1968. Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol* **31**: 14–27.
- Fetz EE. 1992. Are movement parameters recognizably coded in the activity of single neurons? *Behav Brain Sci* **15**: 679–690.
- Fu QG, Flament D, Coltz JD, Ebner TJ. 1995. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J Neurophysiol* **73**: 836–854.
- Georgopoulos AP, Schwartz AB, Kettner RE. 1986. Neuronal population coding of movement direction. *Science* **233**: 1416–1419.
- Georgopoulos AP, Crutcher MD, Schwartz AB. 1989a. Cognitive spatial-motor processes. 3. Motor cortical prediction of movement direction during an instructed delay period. *Exp Brain Res* **75**: 183–194.
- Georgopoulos AP, Lurito JT, Petrides M, Schwartz AB, Massey JT. 1989b. Mental rotation of the neuronal population vector. *Science* **243**: 234–236.
- Gilja V, Nuyujukian P, Chestek CA, Cunningham JP, Yu BM, Fan JM, Churchland MM, Kaufman MT, Kao JC, Ryu SI, et al. 2012. A high-performance neural prosthesis enabled by control algorithm design. *Nat Neurosci* **15**: 1752–1757.
- Hanes DP, Schall JD. 1996. Neural control of voluntary movement initiation. *Science* **274**: 427–430.
- Hatsopoulos NG, Xu Q, Amit Y. 2007. Encoding of movement fragments in the motor cortex. *J Neurosci* **27**: 5105–5114.
- Hennequin G, Vogels TP, Gerstner W. 2014. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* **82**: 1394–1406.
- Takei S, Hoffman DS, Strick PL. 2001. Direction of action is represented in the ventral premotor cortex. *Nat Neurosci* **4**: 1020–1025.
- Kalaska JF. 2009. From intention to action: Motor cortex and the control of reaching movements. *Adv Exp Med Biol* **629**: 139–178.
- Kaufman MT, Churchland MM, Santhanam G, Yu BM, Afshar A, Ryu SI, Shenoy KV. 2010. Roles of monkey premotor neuron classes in movement preparation and execution. *J Neurophysiol* **104**: 799–810.
- Kaufman MT, Seely JS, Shenoy KV, Churchland MM. 2013. A large signal in motor cortex that predicts movement onset. In *Neural Control of Movement Annual Meeting*, San Juan, Puerto Rico.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2014. Cortical activity in the null space: Permitting preparation without movement. *Nat Neurosci* **17**: 440–448.
- Lemon RN, Mantel GW, Muir RB. 1986. Corticospinal facilitation of hand muscles during voluntary movement in the conscious monkey. *J Physiol* **381**: 497–527.
- Lillicrap TP, Scott SH. 2013. Preference distributions of primary motor cortex neurons reflect control solutions optimized for limb biomechanics. *Neuron* **77**: 168–179.
- Moran DW, Schwartz AB. 1999. Motor cortical representation of speed and direction during reaching. *J Neurophysiol* **82**: 2676–2692.
- Mussa-Ivaldi FA. 1988. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci Lett* **91**: 106–111.
- Orger MB, Kampff AR, Severi KE, Bollmann JH, Engert F. 2008. Control of visually guided behavior by distinct populations of spinal projection neurons. *Nat Neurosci* **11**: 327–333.
- Pearce TM, Moran DW. 2012. Strategy-dependent encoding of planned arm movements in the dorsal premotor cortex. *Science* **337**: 984–988.
- Pohlmeier EA, Solla SA, Perreault EJ, Miller LE. 2007. Prediction of upper limb muscle activity from motor cortical discharge during reaching. *J Neural Eng* **4**: 369–379.
- Reimer J, Hatsopoulos NG. 2009. The problem of parametric neural coding in the motor system. *Adv Exp Med Biol* **629**: 243–259.
- Requin J, Riehle A, Seal J. 1988. Neuronal activity and information processing in motor control: From stages to continuous flow. *Biol Psychol* **26**: 179–198.
- Riehle A, Requin J. 1989. Monkey primary motor and premotor cortex: Single-cell activity related to prior information about

- direction and extent of an intended movement. *J Neurophysiol* **61**: 534–549.
- Rokni U, Sompolinsky H. 2012. How the brain generates movement. *Neural Comput* **24**: 289–331.
- Sanger TD. 1994. Theoretical considerations for the analysis of population coding in motor cortex. *Neural Comput* **6**: 29–37.
- Schaffer ES, Rajan K, Churchland MM, Shenoy KV, Abbott LF. 2006. Generating complex repeatable patterns of activity by gain modulating network neurons. In *Society for Neuroscience Annual Meeting*, Atlanta, GA.
- Schieber MH, Rivlis G. 2007. Partial reconstruction of muscle activity from a pruned network of diverse motor cortex neurons. *J Neurophysiol* **97**: 70–82.
- Scott SH. 2008. Inconvenient truths about neural processing in primary motor cortex. *J Physiol* **586**: 1217–1224.
- Sergio LE, Kalaska JF. 1998. Changes in the temporal pattern of primary motor cortex activity in a directional isometric force versus limb movement task. *J Neurophysiol* **80**: 1577–1583.
- Sergio LE, Hamel-Pâquet C, Kalaska JF. 2005. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *J Neurophysiol* **94**: 2353–2378.
- Shen L, Alexander GE. 1997. Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. *J Neurophysiol* **77**: 1195–1212.
- Shenoy KV, Sahani M, Churchland MM. 2013. Cortical control of arm movements: A dynamical systems perspective. *Annu Rev Neurosci* **36**: 337–359.
- Sussillo D, Barak O. 2013. Opening the black box: Low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Comput* **25**: 626–649.
- Sussillo D, Churchland MM, Kaufman M, Shenoy KV. 2012. A recurrent neural network that produces EMG from rhythmic dynamics. In *Translational and Computational Motor Control (TCMC)*, New Orleans, LA.
- Todorov E. 2000. Direct cortical control of muscle activation in voluntary arm movements: A model. *Nat Neurosci* **3**: 391–398.
- Todorov E. 2002. Cosine tuning minimizes motor errors. *Neural Comput* **14**: 1233–1260.
- Wang W, Chan SS, Heldman DA, Moran DW. 2007. Motor cortical representation of position and velocity during reaching. *J Neurophysiol* **97**: 4258–4270.
- Weinrich M, Wise SP, Mauritz KH. 1984. A neurophysiological study of the premotor cortex in the rhesus monkey. *Brain* **107** (Pt 2): 385–414.



Cold Spring Harbor Symposia on Quantitative Biology

A Dynamical Basis Set for Generating Reaches

Mark M. Churchland and John P. Cunningham

Cold Spring Harb Symp Quant Biol published online April 7, 2015

Access the most recent version at doi:[10.1101/sqb.2014.79.024703](https://doi.org/10.1101/sqb.2014.79.024703)

P<P	Published online April 7, 2015 in advance of the print volume.
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here

Advance online articles have not yet appeared in the print volume. Citations to Advance online articles must include the digital object identifier (DOI) and date of initial publication.

To subscribe to *Cold Spring Harbor Symposia on Quantitative Biology* go to:
<http://symposium.cshlp.org/subscriptions>
