



Published in final edited form as:

Nature. 2012 July 5; 487(7405): 51–56. doi:10.1038/nature11129.

Neural population dynamics during reaching

MM Churchland^{1,2,3}, JP Cunningham^{4,5}, MT Kaufman³, JD Foster², P Nuyujukian^{6,7}, SI Ryu^{2,8}, and KV Shenoy^{2,3,6,9}

¹Dept. of Neuroscience, Kavli Institute for Brain Science, David Mahoney Center, Columbia University Medical Center, New York NY, 10032 USA

²Dept. of Electrical Engineering, Stanford University, Stanford CA, 94305, USA

³Neurosciences Program, Stanford University, Stanford CA, 94305, USA

⁴Dept. of Biomedical Engineering, Washington University in St. Louis, St. Louis MO, 63130

⁵Dept. of Engineering, University of Cambridge, Cambridge, CB2 1PZ UK

⁶Dept. of Bioengineering, Stanford University, Stanford CA, 94705, USA

⁷Stanford University School of Medicine, Stanford CA, 94305, USA

⁸Dept. of Neurosurgery, Palo Alto Medical Foundation, Palo Alto CA, 94301, USA

⁹Dept. of Neurobiology, Stanford University School of Medicine, Stanford CA, 94305, USA

Abstract

Most theories of motor cortex have assumed that neural activity represents movement parameters. This view derives from an analogous approach to primary visual cortex, where neural activity represents patterns of light. Yet it is unclear how well that analogy holds. Single-neuron responses in motor cortex appear strikingly complex, and there is marked disagreement regarding which movement parameters are represented. A better analogy might be with other motor systems, where a common principle is rhythmic neural activity. We found that motor cortex responses during reaching contain a brief but strong oscillatory component, something quite unexpected for a non-periodic behavior. Oscillation amplitude and phase followed naturally from the preparatory state, suggesting a mechanistic role for preparatory neural activity. These results demonstrate unexpected yet surprisingly simple structure in the population response. That underlying structure explains many of the confusing features of individual-neuron responses.

Users may view, print, copy, download and text and data- mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Correspondence should be addressed to: Mark Churchland mc3502@columbia.edu.

Author Contributions MMC and JPC contributed equally to this work. The jPCA method was designed by JPC and MMC. MMC and MTK collected data from the reaching monkeys. JDF and PN collected data from the walking monkey. SIR led the array implantation surgeries. KVS contributed to all aspects of the work. All authors discussed the results and commented on the analyses and manuscript.

The authors declare no competing financial interests.

Supplementary material is provided.

Motor and premotor cortex were among the first cortical areas to be extensively studied¹. Yet their basic response properties are poorly understood, and it remains controversial whether neural activity relates to muscles or to abstract movement features^{2–13}. At the heart of this debate is the complexity of individual-neural responses, which exhibit a great variety of multiphasic patterns^{4,14,15}. One explanation is that responses represent many movement parameters:

$$r_n(t) = f_n(\text{param}_1(t), \text{param}_2(t), \text{param}_3(t) \Rightarrow \text{eqn.1}$$

where $r_n(t)$ is the firing rate of neuron n at time t , f_n is a tuning function, and $\text{param}_1(t)$, $\text{param}_2(t)$... are arguments such as hand velocity or target position. Alternatively, motor cortex may constitute a dynamical system that generates and controls movement^{4,8,14–17}. In its simplest, deterministic form this can be expressed as:

$$\dot{\mathbf{r}}(t) = f(\mathbf{r}(t)) + \mathbf{u}(t) \quad \text{eqn.2}$$

where \mathbf{r} is a vector describing the firing rate of all neurons (the ‘population response’ or ‘neural state’), $\dot{\mathbf{r}}$ is its derivative, f is an unknown function, and \mathbf{u} is an external input. In this conception, neural responses reflect underlying dynamics and display ‘tuning’ only incidentally^{18,19}. If so, then dynamical features should be present in the population response. In looking for dynamical structure, we focused on a common principle for movement generation across the animal kingdom: the production of rhythmic, oscillatory activity^{20–22}.

Rhythmic responses in different systems

We first examined neural responses in a context where rhythmic pattern generation is known to occur. The medicinal leech generates rhythmic muscle contractions at ~1.5 Hz during swimming²³, and many single neurons display firing rate oscillations at that frequency (fig. 1a)^{24,25}. Rhythmic structure was also present for cortical responses in the walking monkey: ~1 Hz oscillations matching the ~1 Hz movement of the arm (fig. 1b). If single-neuron oscillations are generated by population-level dynamics, then the population response (the neural state) should rotate with time¹⁵, much as the state of a pendulum rotates in the space defined by velocity and position. We projected the population response onto a two-dimensional state space and found rotations of the neural state for both the swimming leech (fig. 1d; projection of 164 simultaneously recorded neurons) and the walking monkey (fig. 1e; projection of 32 simultaneously recorded channels; also see supplementary movie 1). These observations, while not trivial, are largely expected for a neural dynamical system that generates rhythmic output²².

The projections in figure 1d,e were obtained via two steps. The first was the application of principal component analysis (PCA) to the population response. Inconveniently, PCA does not find dimensions relevant to dynamical structure. We therefore employed a novel method that finds an informative plane within the top PCs. To be conservative, this ‘jPCA’ method was applied only to the top six PCs, which contain the six response patterns most strongly present in the data. The mathematical underpinnings regarding jPCA are described below, but the following is critical. Application of jPCA results in six ‘jPCs’: an orthonormal basis that spans exactly the same space as the first six PCs (supplementary movie 2). The first two

jPCs capture the strongest rotational tendency in the data. The jPC projections are simply linear projections of response patterns that are strongly present in the data; if a given pattern is not present in the top 6 PCs it cannot be present in the jPCs.

The central finding of this study is that quasi-oscillatory neural responses are present during reaches. This is illustrated by the average firing rate of an example motor cortex neuron (fig. 1c) and the corresponding population-level projection (fig. 1f). The rotation of the neural state is short-lived (~1 cycle) but otherwise resembles rotations seen during rhythmic movement. This finding is surprising – the reaches themselves are not rhythmic – yet it agrees with recent theoretical suggestions^{15,22}. One might be concerned that the patterns in figure 1c,f are idiosyncratic. But as shown below, rotations of the neural state are one of the most prominent features of the data.

Quasi-rhythmic responses during reaching

We analyzed 469 single-neuron recordings from motor and premotor cortex of four monkeys (A,B,J,N). We made a further 364 simultaneous recordings (single and multi-unit isolations) from two pairs of implanted 96-electrode arrays (monkeys J,N). Monkeys executed straight reaches (monkeys A,B) or straight and curved reaches (monkeys J,N). An instructed delay paradigm allowed monkeys to prepare their reaches before a go cue. We analyzed 9 datasets, each employing 27–108 reach types ('conditions'). For each neuron/condition we computed and analyzed the average across-trial firing rate.

Most neurons exhibited preparatory and movement-related responses (fig. 2). Responses were typically complex, multiphasic and heterogeneous¹⁴. Yet there appear to be oscillations in many single-neuron responses, beginning just before movement onset and lasting for ~1–1.5 cycles. These quasi-oscillatory patterns are seen for all reach types and all monkeys. Yet interpretational caution is warranted: multiphasic responses might exist for any number of reasons. The critical question is whether there exists orderly rotational structure, across conditions, at the population level.

We have proposed that motor cortex responses reflect the evolution of a neural dynamical system, starting at an initial state set by preparatory activity^{14,15,17,18,26}. If the rotations of the neural state (fig. 1f) reflect straightforward dynamics, then similar rotations should be seen for all conditions. In particular, the neural state should rotate in the same direction for all conditions¹⁵, even when reaches are in opposition.

We projected the population response for all conditions onto the jPC plane. This was done for 200 ms of data, beginning when preparatory activity transitions to movement-related activity (supplementary movie 3 shows a longer span of time). The resulting projections (fig. 3a-f) show four striking features. First, rotations of the neural state are prevalent during reaching. Second, the neural state rotates in the same direction across conditions (*different traces*). Third, the rotation phase follows naturally from the preparatory state. Finally, state-space rotations do not relate directly to reach curvature. Monkeys A and B executed straight reaches; monkeys J and N executed a mixture of straight reaches, clockwise-curving reaches, and counter-clockwise-curving reaches (fig. 2, *insets*). Yet for each dataset the

neural state rotates in the same direction across conditions. Rotations appear to reflect dynamics that are consistent across conditions, rather than the pattern of kinematics *per se*.

If one knows the initial population-level preparatory state (fig. 3, *circles*) subsequent states are reasonably predictable. Such predictability is absent at the individual-neuron level: the correlation between preparatory and movement ‘tuning’ averages nearly zero¹⁵. How do the ordered state-space rotations relate to the seemingly disordered single-neuron responses? Each axis of the jPCA projection captures a time-varying pattern that resembles a single-neuron response (supplementary fig. 1). Indeed, single neurons strongly reflect combinations of these underlying patterns. However, the underlying structure is not readily apparent when plotting each pattern alone, or each neuron individually¹⁵. Furthermore, the rotations during reaching are quasi-oscillatory lasting only 1–1.5 cycles (fig. 1c, also see supplementary movie 3). Their brevity and high frequency (up to ~2.5 Hz) makes them easy to miss unless trial-counts are high (datasets averaged 810 trials/neuron) and data are precisely aligned on movement onset (*methods*).

Controls regarding rotational structure

The central result of this study is the presence of the rotational patterns seen in figure 3. Those projections captured an average of 28% of the total data variance, and thus reveal patterns that are strongly present in the data. Yet might such patterns appear by accident or for trivial reasons? Multiple ‘shuffle’ controls demonstrate that jPCA does not find rotational structure when such structure is not present (supplementary fig. 2,3; supplementary movie 4). Similarly, rotations in the walking monkey were not erroneously found when the monkey was stationary (supplementary movie 1).

The fact that a single plane (two dimensions) captures an average of 28% of the total data variance is notable, given the high dimensionality of the data itself¹⁴. As a comparison, the dimensions defined by PC₂ and PC₃ (which by definition capture the second- and third-most data variance possible) together capture 29% of the total variance. Thus, the jPCA projection simply captures response patterns that were always present in the top PCs, but were difficult to see because they were not axis aligned. In fact, there were typically two or three orthogonal planes that captured rotational structure at different frequencies (supplementary fig. 4). Together these captured 50–70% of the total data variance. Thus, rotations are a dominant feature of the population response. This was true for M1 and PMd independently (supplementary fig. 5).

Rotations, kinematics and EMG

Traditional views posit that motor cortex neurons are tuned for movement parameters such as direction. This perspective does not naturally account for the data in fig. 3. We simulated neural populations that were directionally tuned for velocity with an additional non-directional sensitivity to speed²⁷. Simulated preparatory activity was tuned for reach direction/distance²⁸. We simulated one ‘velocity model’ dataset per recorded dataset, based upon the recorded velocities/endpoints. Firing rates, trial-counts, neuron-counts, and spiking noise were matched to the recorded data. For velocity-model populations, jPCA found no robust or consistent rotations (fig. 4a,d,h). This was true for all datasets (summary analysis

below) including those with curved reaches (e.g., fig. 4*h*). We also simulated a ‘complex kinematic’ model where responses reflected the weighted sum of kinematic parameters (position, velocity, acceleration, and jerk) that correlate with muscle activity⁶. This model produced multiphasic responses but not consistent rotations (fig. 4*b,e,i*). We also recorded EMG from a population of muscles (6–12 recordings per dataset). Although EMG is strongly multiphasic, the population of muscles did not show consistent rotations (fig. 4*c,f,j*, summary data below). This was not due to the smaller size of the muscle population (supplementary fig. 6). In sum, rotations in state space require more than multiphasic responses: they require a pair of multiphasic patterns with phases consistently $\sim 90^\circ$ apart. The neural population contains that complementary pair; the simulated and muscle populations do not. Still, EMG may bear some relationship to the observed rotations – a possibility explored below.

The rotations of the neural state are a robust feature of the physiological data, but how do those rotations relate to the reaches themselves? This question is especially relevant because the reaches were not overtly rhythmic. A possible answer is that muscle activity might be constructed from an oscillatory basis. To test whether this is plausible, we simulated a simple dynamical model possessing two orthogonal rotations in state space: one at a high frequency and one at a low frequency. Muscle activity was fit as the sum of the resulting oscillations in the temporal domain. For example, for monkey J3 the higher-frequency rotation occurred at 2.8 Hz (figure 5*a*). Different conditions (9 and 25 are shown) involved different amplitudes/phases, set by the preparatory state. The vertical, ‘lagging’ dimension impacted simulated muscle activity, and the projections onto that dimension (figure 5*b,c*, *dark blue*) provided key features of the EMG fit. The slower features are provided by a 0.3 Hz oscillation (not shown).

This ‘generator model’ provided excellent EMG fits (fig. 5*b,c*; supplementary fig. 7,8). The fit/EMG correlation ranged from 0.97–0.99 across datasets. Thus, deltoid activity could always be generated from the sum of two underlying rotations whose phases and amplitudes (but not frequencies) vary across conditions. This raises a subtle but key point: although EMG responses do not themselves exhibit state-space rotations, EMG can nevertheless be constructed from underlying rotations. This observation needn’t imply that cortex directly controls muscles. Yet it illustrates the plausibility of direct control⁶, and demonstrates that rotations provide a natural basis set for generating non-rhythmic movement.

One might have expected faster reaches to involve faster rotations, or longer reaches to involve longer rotations. However, EMG could be well fit using two rotations with fixed frequency/duration. This was true even though the 27 conditions differed greatly in reach speed/duration (fig. 5*g*). We return to this point below.

For representational models, individual-unit responses reflect the ‘factors’ for which those units are tuned. For a dynamical model, individual-unit responses should reflect the underlying dynamical factors: the patterns present on each axis of the state space. We simulated a population of ‘generator model’ units whose rates depended, with random weights, on these underlying patterns. ‘Preparatory’ activity was simply the initial state. Simulated units displayed multiphasic response patterns resembling those of real neurons

(fig. 5*d,e*; supplementary fig. 8). Both real and simulated responses exhibited reversals in ‘preferred condition’¹⁴ and a weak correlation between preparatory and movement-period ‘tuning’^{15,29}. Despite such surface complexity, jPCA projections of the simulated populations successfully reveal the simple underlying rotations (fig. 5*f*). These rotations resemble those observed for the neural data (fig. 3).

Population-level quantification

To quantify rotation strength we measured the angle from the neural state in the jPCA plane (x) to its derivative (\dot{x}) for every condition and time (fig. 6*a*). The first jPCA plane is oriented such that the average rotational tendency of the data – however weak or strong – is counterclockwise. Thus, angles near positive $\pi/2$ indicate a strong rotational component. The neural data and the generator model have distributions with peaks near $\pi/2$. In contrast, the velocity-tuned model, the complex-kinematic model, and EMG all had distributions that peaked slightly above zero.

Rotation strength was also quantified via the jPCA computation, in which the data were fit with:

$$\dot{x}(t, c) = M_{skew} x(t, c) \quad \text{eqn.3}$$

where $x(t, c)$ is the k -dimensional ($k=6$ for all figures) population state for time t and condition c . M_{skew} is a skew-symmetric matrix ($M_{skew} = -M_{skew}^T$) that captures rotational dynamics. The first two jPCs were the two eigenvectors associated with the largest-magnitude eigenvalues of M_{skew} , which indicate the strongest rotational plane in the dynamical system fit by eqn. 3. Projecting data onto those jPCs (as in fig. 3,4) reveals the prevalence of rotations. We can also assess rotation prevalence by quantifying how well M_{skew} fits the data relative to an unconstrained matrix M . That unconstrained M provides the best performance of any matrix (skew-symmetric matrices are a subset of unconstrained matrices). For the neural data and generator model, the fit provided by M_{skew} was nearly as good as the fit provided by the best unconstrained M (fig. 6*b*). For the velocity-tuned and complex kinematic models, M_{skew} performed poorly. Results were similar whether we considered all six dimensions or just the plane with the strongest rotations (the plane defined by the first two jPCs). Thus, of those dynamics that can be captured linearly, rotational dynamics dominate only for the generator model and neural data.

EMG data showed weak rotations (fig 6*a,b, red*), underscoring a central point: state-space rotations result not from a multiphasic signal, but from how that signal is constructed. For example, the generator model exhibits rotations even though the EMG does not. More generally, many features of the observed rotations make sense in terms of how outputs (EMG, kinematics) might be generated, rather than in terms of the outputs themselves. For example, a strong intuition, and a prediction of most hypotheses of motor cortex, is that neural responses should unfold faster for faster movements. However, the generator model makes the opposite prediction; rotation frequencies are fixed. We tested this prediction using two datasets where target color instructed fast versus slow reaches. Both the generator model and the neural data exhibited rotations that were of similar angular velocity for fast

and slow reaches (fig. 6*c,d*). The same point can be made by separately applying jPCA for fast and slow reaches: the largest eigenvalues of M_{skew} were actually slightly smaller for fast reaches (8.8 versus 9.8 radians/s for monkey A, 12.2 versus 13.5 radians/s for monkey B). Rotation amplitude, rather than frequency, differed between speeds (fig. 6*c,d*). This finding is readily interpretable in light of the generator model: larger-amplitude rotations produce more strongly multiphasic responses, a feature of the EMG necessary to drive larger accelerations/decelerations (also see supplementary fig. 9).

Discussion

Rotations of the population state are a prominent feature of the cortical response during reaching. Rotations follow naturally from the preparatory state and are consistent, in direction and angular velocity, across the different reaches that each monkey performed. The rotational structure is much stronger and more consistent than expected by chance or given previous models. These population-level rotations are a relatively simple dynamical feature yet explain seemingly complex features of individual-neuron responses, including frequent reversals of preferred direction^{14,30}, and a lack of correlation between preparatory and movement-period ‘tuning’^{15,29} (fig 5*d,e* and supplementary fig. 8). State-space rotations produce briefly oscillatory temporal patterns that provide an effective basis set for producing multiphasic muscle activity, suggesting that non-periodic movements may be generated via neural mechanisms resembling those that generate rhythmic movement^{20,22,31–33}.

Recent results suggest that preparatory activity sets the initial state of a dynamical system, whose subsequent evolution produces movement activity¹⁵. Aspects of those dynamics – a rotation away from the preparatory state – appear straightforward. However, the circuitry that creates those dynamics is unclear – it may be purely local, or may involve recurrent circuitry³⁴ that links motor cortex with the spinal cord and with critical subcortical structures³⁵. Peripheral feedback is also expected to shape neural dynamics³⁶, although this cannot account for the first ~150 ms of ‘neural motion’ (the hand has yet to move). The finding that dynamics are captured by a skew-symmetric matrix suggests functionally anti-symmetric connectivity: a given neural dimension (e.g., jPC₁) positively influences another (e.g., jPC₂) which negatively influences the first. However, it is unclear whether this population-level pattern directly reflects a circuit-level dominance of anti-symmetric connectivity. We also stress that although rotations are a dominant pattern in the data across multiple dimensions (supp. fig. 4) there exist non-rotational components as well, perhaps reflecting the nonlinear dynamics necessary for initiating/terminating movement, for stability³⁷, and for feedback control^{16,38}.

It is hoped that a focus on the dynamics that generate movement will help transcend the controversy over what single neurons in motor cortex ‘code’ or ‘represent.’ Many of the neural response features that seem most baffling from a representational perspective are natural and straightforward from a dynamical systems perspective. It therefore seems increasingly likely that motor cortex can be understood in relatively straightforward terms: as an engine of movement that employs lawful dynamics.

Methods summary

Optical recordings from the isolated leech central nervous system were made by K. Briggman and W. Kristan and have been described previously^{24,25}. We recorded neural activity from trained monkeys using both single- and multi-electrode techniques. We recorded from the arm representation of premotor cortex using a wireless system while the monkey walked to obtain juice from the front of a treadmill. We recorded from the arm representation of motor and premotor cortex while monkeys reached to targets projected onto a vertically oriented screen, also for juice reward. All surgical and animal care procedures were performed in accordance with National Institutes of Health guidelines and were approved by the Stanford University Institutional Animal Care and Use Committee.

Methods

Recordings and task design

Recordings from the isolated leech central nervous system were made by K. Briggman and W. Kristan and have been described previously^{24,25}. Recordings from monkey cortex were made using both a delayed reach task (with head restraint) and from an unrestrained monkey walking on a treadmill^{39–41}. Animal protocols were approved by the Stanford University Institutional Animal Care and Use Committee.

Most analyses concerned data collected during delayed reach tasks, for which our basic methods have been described previously^{15,29,42}. Briefly, four monkeys (A, B, J, and N) performed delayed reaches on a fronto-parallel screen. Delays ranged from 0–1000 ms (the exact range varied by monkey). Only trials with delays >400 ms were analyzed. Fixation was enforced (at the central spot) during the delay for monkeys J and N. We employed two variants of a center-out reaching task. In the ‘speed task’ monkeys A and B reached to radially arranged targets at two distances. Reach speed was instructed by target color (28 total conditions)¹⁸. In the ‘maze task’ monkeys J and N made both straight reaches and reaches that curved around one or more intervening barriers. This task was beneficial because of the large variety of different reaches that could be evoked. Typically we used 27 conditions: each providing a particular arrangement of target and barriers. Monkey J performed the task for four different sets of 27 conditions, resulting in four datasets (J1 through J4). For the monkey J-array and N-array datasets, 108 conditions were presented in the same recording session. Recordings from three of the four monkeys (A, B and J) have been analyzed in prior publications (e.g.,¹⁵).

Recordings were made from primary motor cortex (M1, both surface and sulcal) and from the adjacent (caudal) aspect of dorsal premotor cortex (PMd). Supplementary figure 5 shows the central analysis divided by area. Seven of the nine datasets (monkey A, B, J1, J2, J3, J4, and N) were recorded with conventional single-electrode techniques. These datasets involved a total of 469 single-unit isolations. The other two datasets (monkey J-array, recorded 09-18-2009; monkey N-array, recorded 09-23-2010) employed a pair of chronically implanted 96-electrode arrays (Blackrock Microsystems). These array-based datasets involved, respectively, 146 and 218 single isolations (each a mix of single and multi-unit isolations).

Arrays were implanted after directly visualizing sulcal landmarks. Single-electrode recordings were guided by stereotaxic criteria, the known response properties of M1 and PMd, and the effects of microstimulation. For all monkeys, at some point the dura was reflected and the sulcal landmarks directly visualized. Recordings were medial to the arcuate spur and lateral to the precentral dimple. Recordings were not made within rostral PMd, near the arcuate sulcus. Sulcal M1, surface M1, and caudal PMd are contiguous. While there are important differences in their average response properties (delay period activity is more common in PMd), these differences are far from absolute: M1-like neurons are frequently found in caudal PMd and vice versa. Most analyses thus considered all neurons without attempting to divide based either on anatomy or on response properties (although see supplementary figure 5).

For the freely walking monkey, data were recorded from an array implanted in the arm representation of PMd. The times of threshold crossings on 32 of the 96 channels were wirelessly transmitted using the HermesD system^{39,41}. Behavioral data were recorded using a commercially available video camera. Juice was dispensed at one end of the treadmill, providing incentive for the monkey to walk continuously.

EMG data were collected as described previously⁴². EMG records were rectified, smoothed and averaged before further analysis. A total of 61 recordings were made from six muscle groups: deltoid, biceps brachii, triceps brachii, trapezius, latissimus dorsi, and pectoralis. Most datasets contained multiple recordings from each muscle (e.g., one from each of the three heads of the deltoid). The total number of EMG recordings for some datasets was thus as high as twelve. EMG was recorded for all datasets except those recorded using arrays.

Computing average firing rate as a function of time

Average trial counts were high (an average of 810 trials/neuron). To ensure response features were not lost to averaging, a concerted effort was made to compute the average firing rate only over trials with nearly identical reach trajectories. This was done by training to a high level of stereotyped behavior, and by discarding the few trials for which behavior was not tightly stereotyped. Average firing rates were further de-noised by filtering with a Gaussian (20 or 24 ms depending on the dataset) and using a custom-developed smoothing method that discards idiosyncratic features that are both small and not shared across conditions (see supp. fig. 4 of ¹⁵). This method improves the signal-to-noise ratio without over-smoothing in the temporal domain, which was important for preserving high-frequency features of the response. This step aids the visualization of single-neuron PSTHs, but had essentially no effect on any of the population-level analyses (supplementary fig. 10). ‘Recordings’ of simulated neurons and of EMG were processed using all the same steps as for the real data.

Because the delay period and reaction time were variable, firing rates were computed separately locked to target onset, the go cue, and movement onset. For presentation (where one wishes to follow a trace through different epochs) we interpolated over the gaps between the three epochs.

Fitting the generator model to EMG

For the generator model, we directly simulated two state-space rotations. The goal was to start not by simulating the responses of individual units, but by directly simulating the underlying structure of the population data in state space. The two simulated rotations produced patterns that were summed to fit the EMG for the deltoid. For example, the deltoid EMG for dataset J3 was fit using a 2.8 Hz rotation and a 0.3 Hz rotation. Each rotation consisted of leading and lagging sinusoids windowed by a gamma function, with the initial state extended backwards in time to mimic preparatory activity (e.g., fig. 5*a,b,c*). The amplitude and phase of that rotation was different for every condition, to allow the model to fit the different EMG patterns recorded for each condition. Importantly, for a given dataset that rotation always had the same frequency regardless of condition, with a rise and decay defined by the same windowing gamma function (e.g., the 2.8 Hz rotation was always at 2.8 Hz and the 0.3 Hz rotation was always at 0.3 Hz). This mimics a dynamical system that is the same across conditions except for an initial state that determines phase and amplitude. EMG was fit as the sum of the lagging sinusoids, one for each of the two frequencies.

Optimization involved two levels. At the level of each individual condition, the amplitudes and phases that provided the best fit were found via regression. Regression exploited the fact that every possible amplitude/phase of a sinusoid can be constructed via a linear combination of a sine and cosine. This step is thus both fast and guaranteed to find the best fit. Regression involved an offset term, which could be different for each condition. At the level of the whole dataset, we numerically optimized the two frequencies, the mean and shape parameter of the windowing gamma function, and the time when oscillations began. Optimization was started from many initial parameter choices and the best fit was chosen.

Each condition's simulated EMG is simply the sum of two windowed sinusoids and a variable DC offset. However, the central idea of the generator model is that those sinusoids result from rotations in an internal state space. The generator model thus embodied five basic patterns: the pair of leading and lagging dimensions that make up each rotation plus the DC offset.

Simulated neural data

We produced two classes of simulated neural datasets. The first class (the 'velocity model' and 'complex kinematic model') was based on a traditional framework in which units were cosine tuned for kinematic factors. The second class was based on the generator model describe above, which emulates a simple dynamical system. For both classes of model, the firing rates of individual units were assumed to depend upon underlying 'factors.' For the velocity-tuned model, movement-period activity was based upon three underlying factors: horizontal reach velocity, vertical reach velocity, and reach speed (e.g.,²⁷). Each unit thus had a 'preferred direction' in velocity space. Preparatory activity was based upon three additional underlying factors: horizontal reach endpoint, vertical reach endpoint, and peak reach speed. For the complex kinematic model, we assumed that because muscle activity reflects a variety of kinematic factors (position, velocity, acceleration, jerk) neural activity might share this property^{6,30}. As with the velocity model, each unit was cosine tuned with a preferred direction in physical space. Simulated activity depended on motion in that

preferred direction with the following sensitivities: 25 (spikes/s)/m, 10 (spikes/s)/(m/s), 1 (spikes/s)/(m/s²), 0.05 (spikes/s)/(m/s³). These constants are taken from a published model⁶, but have been adjusted as follows. First, the sensitivity to position has been reduced by half, otherwise it tended to dominate to an unrealistic degree. Second, a sensitivity to jerk has been added. This makes for a more stringent control (it increases the multiphasic aspects of the simulated responses) and captures the expectation that cortical activity might be more phasic than muscle activity. Preparatory activity was sensitive to target endpoint.

For the generator model, the underlying factors were the oscillatory patterns captured by the underlying state space. These patterns defined both the movement-period and (via the initial state for each pattern) preparatory activity. Also included as underlying factors were the two gamma functions that defined the oscillation envelopes. These factors were the same across all conditions, and were included to mimic the overall change in excitability that is presumed to cause the waxing and waning of oscillations. The inclusion of these un-tuned factors had a similar impact on the generator-model neurons as did the non-directional speed factor for the velocity-tuned model neurons. However, while the addition of these un-tuned factors served to increase the variety and realism of the simulated responses, it has essentially no impact on the main analyses (fig. 5g and 6). Those analyses are sensitive only to response aspects that differ among conditions.

The firing rate of each simulated unit was a random combination of the underlying factors. For the velocity and complex-kinematic models, the random combinations resulted in a roughly uniform range of preferred directions for both the preparatory and movement periods. For the generator model, the random combinations resulted in simulated units that typically reflected both oscillation frequencies, with a roughly uniform distribution of phases. This is indeed the default expectation for a large network that supports two oscillatory modes.

To produce simulated data with realistic levels of noise, we ‘recorded’ spikes that were produced via a gamma-interval process (order 2) based on the underlying firing rate. For each neural dataset, we simulated one unit for every recorded neuron, and matched the overall firing rates and trial-counts of each simulated unit to those of the respective recorded neuron. The simulated spiking data was then analyzed just as for the actual neural data. The velocity-model and complex-kinematic models each produced 9 simulated datasets (one for every real dataset). The generator-model produced 7: it could not be simulated for the J-array or N-array datasets, as we did not attempt to record EMG for those 108-condition experiments.

Projections that capture rotational structure

We produced projections of the population data using a novel dimensionality reduction method, jPCA, designed for the present application. For most analyses we analyzed 200 ms of time, sampled every 10 ms, starting just before the rapid change in neural activity that precedes movement onset. Before applying jPCA, a number of preprocessing steps were applied to the data (these same steps were also applied to the simulated data and EMG). Responses were normalized to have a similar firing-rate range for all neurons. ‘Soft’ normalization was used, so that neurons with very strong responses were reduced to

approximately unity range, but neurons with weak responses had less than unity range. For each neuron, the data were mean-centered at every time: the average across-condition response was subtracted from the response for each condition. Thus, all subsequent analysis focused on those aspects of the neural response that differ across conditions. This preprocessing step can be skipped (see supplementary fig. 11) but the resulting projections often capture rotations that are similar for all conditions. In such cases one fails to gain multiple views of the underlying process, making it difficult to infer whether rotations are due to dynamics or to more trivial possibilities. It was thus deemed more conservative to only interpret projections where activity unfolds differently across conditions. Related population analyses (e.g., the population vector) achieve the same end implicitly: non-directional aspects of the response cancel out. The preprocessing steps (and all subsequent analysis steps) were applied in the same way to all datasets, real and simulated.

The most critical preprocessing step was the use of traditional PCA. We compiled a data matrix, X , of size $n \times ct$, where n is the number of neurons, c is the number of conditions, and t is the number of time-points. This matrix simply contains the firing rates of every neuron for every condition and every analyzed time. We then used PCA to reduce the dimensionality of X from $n \times ct$ to $k \times ct$. $k = 6$ for all analyses in the main text, which is conservative given the true dimensionality of the data¹⁴. The resulting $6 \times ct$ matrix, X_{red} , defines a 6-dimensional neural state for every time and condition. By preprocessing with PCA, we ensure that when jPCA is subsequently applied, it reveals only patterns of activity that are strongly present across neurons. Preprocessing with PCA greatly reduces any potential concern that the observed rotations were found simply by looking in a very high-dimensional space (also see shuffle controls in supplementary fig. 2 and 3).

jPCA is a method for finding projections (onto an orthonormal basis) that captures rotational structure in the data. jPCA is based upon a comparison of the neural state with its derivative. We computed \dot{X}_{red} , of size $6 \times c(t-1)$ by taking the difference in the state between adjacent time-points within each row of X_{red} . We then fit using:

$$\dot{X}_{red} = M X_{red} \quad \text{eqn.4}$$

and,

$$\dot{X}_{red} = M_{skew} X_{red} \quad \text{eqn.5}$$

(To keep the dimensions appropriate, the final time-point for each condition in X_{red} was removed so that it was also $6 \times c(t-1)$). Thus, we are attempting to find matrices M and M_{skew} that take the state at each time in X_{red} , and transform it into the derivative of the state in \dot{X}_{red} . M can be found using linear regression. Finding M_{skew} requires more complex (but still linear) operations (see supplementary derivation for more detail). The quality of the above fits was assessed using R^2 (e.g., fig. 6b). R^2 captures the ability to linearly predict the data in \dot{X}_{red} (across all times and conditions) from the data in X_{red} .

M_{skew} has imaginary eigenvalues, and thus captures rotational dynamics. The strongest, most rapid rotations in the dynamical system occur in the plane defined by the eigenvectors

associated with the largest two (complex-conjugate) imaginary eigenvalues. These eigenvectors (V_1 and V_2) are complex, but the associated real plane can be found by: $jPC_1 = V_1 + V_2$, and $jPC_2 = j(V_1 - V_2)$ (after resolving the ambiguity in the polarity of V_1 and V_2 such that their real components have the same sign). The first jPCA projection is then $X_{jPCA} = [jPC_1; jPC_2] \times X_{red}$. X_{jPCA} is thus a $2 \times ct$ matrix describing the neural state, projected onto two dimensions, for every time and condition. For a given jPCA plane, the choice of orthogonal vectors (jPC_1 and jPC_2) within that plane is arbitrary. We therefore selected jPC_1 and jPC_2 so that any net rotation was counter-clockwise (the same choice was of course used across all conditions for a given dataset) and so that the spread of preparatory states was strongest along jPC_1 . We also computed the proportion of the total data variance captured by the jPCA plane, in a manner exactly analogous to that for PCA.

It is worth stressing that the six jPCs form an orthonormal basis that spans exactly the same space as the first 6 PCs. Thus, all patterns seen in the jPCA projections are also present in the PCA projections (the rotational patterns are simply not axis aligned in the latter case, and are thus less obvious to the eye; see supplementary movie 2).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are deeply grateful to K. Briggman and W. Kristan for providing data recorded from the leech. We thank M. Risch for animal care, S. Eisensee for administrative support, and D. Haven and B. Oskotsky for IT support. We thank Z. Ghahramani and C. Rasmussen for discussion of jPCA and related methods. We thank both David Sussillo and Sanford Grossman for analysis suggestions and commentary on the manuscript. This work was supported by a Helen Hay Whitney postdoctoral fellowship and NIH postdoctoral training fellowship (MMC), the Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (MMC, KVS), a UK EPSRC EP/H019472/1 (JPC), a National Science Foundation graduate research fellowship (MTK), a Texas Instruments Stanford Graduate Fellowship (JDF), a Paul and Daisy Soros Fellowship (PN), the Stanford Medical Scientist Training Program (PN), and these awards to KVS: NIH Director's Pioneer Award (1DP1OD006409), NIH NINDS EUREKA Award (R01-NS066311), NIH NINDS BRP (R01-NS064318), NIH NINDS CRCNS (R01-NS054283), DARPA-DSO REPAIR (N66001-10-C-2010), Stanford Center for Integrated Systems, NSF Center for Neuromorphic Systems Engineering at Caltech, Office of Naval Research, and the Whitaker Foundation.

References

1. Lemon RN. An enduring map of the motor cortex. *Exp Physiol.* 2008; 93:798–802. 93/7/798 [pii]. 10.1113/expphysiol.2007.039081 [PubMed: 18562475]
2. Evarts EV. Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol.* 1968; 31:14–27. [PubMed: 4966614]
3. Mussa-Ivaldi FA. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci Lett.* 1988; 91:106–111. [PubMed: 3173781]
4. Fetz EE. Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences.* 1992; 15:679–690.
5. Sanger TD. Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation.* 1994; 6:29–37.
6. Todorov E. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci.* 2000; 3:391–398. [PubMed: 10725930]
7. Hatsopoulos NG. Encoding in the motor cortex: was evarts right after all? Focus on “motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks”. *J Neurophysiol.* 2005; 94:2261–2262. [PubMed: 16160087]

8. Scott SH. Inconvenient truths about neural processing in primary motor cortex. *J Physiol.* 2008; 586:1217–1224.10.1113/jphysiol.2007.146068 [PubMed: 18187462]
9. Morrow MM, Pohlmeier EA, Miller LE. Control of Muscle Synergies by Cortical Ensembles. 2009; 629:179–199.10.1007/978-0-387-77064-2_9
10. Afshar TN, Graziano MSA. Relationship between Unconstrained Arm Movements and Single-Neuron Firing in the Macaque Motor Cortex. *Journal of Neuroscience.* 2007; 27:2760–2780.10.1523/jneurosci.3147-06.2007 [PubMed: 17360898]
11. Kalaska JF. From intention to action: motor cortex and the control of reaching movements. *Adv Exp Med Biol.* 2009; 629:139–178.10.1007/978-0-387-77064-2_8 [PubMed: 19227499]
12. Georgopoulos AP, Schwartz AB, Kettner RE. Neuronal population coding of movement direction. *Science.* 1986; 233:1416–1419. [PubMed: 3749885]
13. Kakei S, Hoffman DS, Strick PL. Muscle and movement representations in the primary motor cortex. *Science.* 1999; 285:2136–2139. [PubMed: 10497133]
14. Churchland MM, Shenoy KV. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J Neurophysiol.* 2007; 97:4235–4257. [PubMed: 17376854]
15. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron.* 2010; 68:387–400.10.1016/j.neuron.2010.09.015 [PubMed: 21040842]
16. Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nat Neurosci.* 2002; 5:1226–1235. [PubMed: 12404008]
17. Graziano MS. New insights into motor cortex. *Neuron.* 2011; 71:387–388.10.1016/j.neuron.2011.07.014 [PubMed: 21835337]
18. Churchland MM, Santhanam G, Shenoy KV. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J Neurophysiol.* 2006; 96:3130–3146. [PubMed: 16855111]
19. Cisek P. Preparing for speed. Focus on: “Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach”. *J Neurophysiol.* 2006
20. Georgopoulos AP, Grillner S. Visuomotor coordination in reaching and locomotion. *Science.* 1989; 245:1209–1210. [PubMed: 2675307]
21. Grillner S. Biological Pattern Generation: The Cellular and Computational Logic of Networks in Motion. *Neuron.* 2006; 52:751–766.10.1016/j.neuron.2006.11.008 [PubMed: 17145498]
22. Rokni U, Sompolinsky H. How the brain generates movement. *Neural Comput.* 2011
23. Kristan WB Jr, Calabrese RL. Rhythmic swimming activity in neurones of the isolated nerve cord of the leech. *J Exp Biol.* 1976; 65:643–668. [PubMed: 1018167]
24. Briggman KL, Kristan WB. Imaging Dedicated and Multifunctional Neural Circuits Generating Distinct Behaviors. *Journal of Neuroscience.* 2006; 26:10925–10933.10.1523/jneurosci.3265-06.2006 [PubMed: 17050731]
25. Briggman KL, Abarbanel HD, Kristan WB Jr. Optical imaging of neuronal populations during decision-making. *Science.* 2005; 307:896–901. [PubMed: 15705844]
26. Shenoy, KV.; Kaufman, MT.; Sahani, M.; Churchland, MM. *Progress in Brain Research: Enhancing Performance for Action and Perception.* Green, A.; Chapman, E.; Kalaska, JF.; Lepore, F., editors. Elsevier; 2011.
27. Moran DW, Schwartz AB. Motor cortical representation of speed and direction during reaching. *J Neurophysiol.* 1999; 82:2676–2692. [PubMed: 10561437]
28. Messier J, Kalaska JF. Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *J Neurophysiol.* 2000; 84:152–165. [PubMed: 10899193]
29. Kaufman MT, et al. Roles of Monkey Premotor Neuron Classes in Movement Preparation and Execution. *Journal of Neurophysiology.* 2010; 104:799–810.10.1152/jn.00231.2009 [PubMed: 20538784]
30. Sergio LE, Hamel-Paquet C, Kalaska JF. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *J Neurophysiol.* 2005; 94:2353–2378. [PubMed: 15888522]

31. Llinas, R. *I of the Vortex*. MIT Press; 2002.
32. Yuste R, MacLean JN, Smith J, Lansner A. The cortex as a central pattern generator. *Nat Rev Neurosci*. 2005; 6:477–483.10.1038/nrn1686 [PubMed: 15928717]
33. Yakovenko S, Krouchev N, Drew T. Sequential Activation of Motor Cortical Neurons Contributes to Intralimb Coordination During Reaching in the Cat by Modulating Muscle Synergies. *Journal of Neurophysiology*. 2010; 105:388–409.10.1152/jn.00469.2010 [PubMed: 21068260]
34. Middleton FA, Strick PL. Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res Brain Res Rev*. 2000; 31:236–250. S0165017399000405 [pii]. [PubMed: 10719151]
35. Thorn CA, Atallah H, Howe M, Graybiel AM. Differential Dynamics of Activity Changes in Dorsolateral and Dorsomedial Striatal Loops during Learning. *Neuron*. 2010; 66:781–795.10.1016/j.neuron.2010.04.036 [PubMed: 20547134]
36. Herter TM, Korb T, Scott SH. Comparison of Neural Responses in Primary Motor Cortex to Transient and Continuous Loads During Posture. *Journal of Neurophysiology*. 2008; 101:150–163.10.1152/jn.90230.2008 [PubMed: 19005005]
37. Sussillo D, Abbott LF. Generating coherent patterns of activity from chaotic neural networks. *Neuron*. 2009; 63:544–557. [PubMed: 19709635]
38. Scott SH. Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci*. 2004; 5:532–546. nrn1427 [pii]. 10.1038/nrn1427 [PubMed: 15208695]
39. Gilja V, Chestek CA, Nuyujukian P, Foster JD, Shenoy KV. Autonomous head-mounted electrophysiology systems for freely behaving primates. *Curr Opin Neurobiol*. 2010; 20:676–686. S0959-4388(10)00104-2 [pii]. 10.1016/j.conb.2010.06.007 [PubMed: 20655733]
40. Foster, JD., et al. Proc. of the 5th International IEEE EMBS Conference on Neural Engineering; p. 613-615.
41. Miranda H, Gilja V, Chestek CA, Shenoy KV, Meng TH. HermesD: A high-rate long-range wireless transmission system for simultaneous multichannel neural recording applications. *IEEE Transactions on Biomedical Circuits and Systems*. 2010; 4:181–191. [PubMed: 23853342]
42. Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV. Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci*. 2006; 26:3697–3712. [PubMed: 16597724]

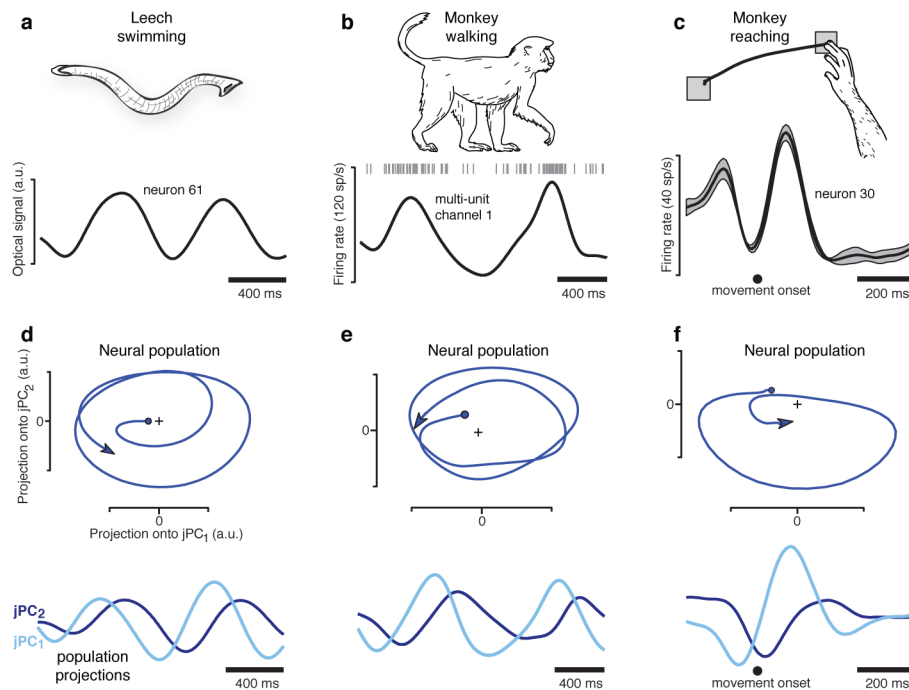


Figure 1.

Oscillation of neural firing rates during three movement types. **a.** Response of one of 164 neurons (simultaneously recorded using voltage-sensitive dye) in the isolated leech central nervous system during a swimming motor pattern. Responses (not averaged across repetitions) were filtered with a 100 ms Gaussian kernel. **b.** Multi-unit response from one of 96 electrodes implanted in the arm representation of caudal premotor cortex. Data from 32 such channels was wirelessly transmitted during walking. Responses (not averaged across repetitions) were filtered with a 100 ms Gaussian kernel. **c.** Response of one of 118 neurons recorded from motor cortex of a reaching monkey (N) using single-electrode techniques. Firing rates were smoothed with a 24 ms Gaussian and averaged across 9 repetitions of the illustrated leftwards reach (flanking traces show SEM). **d.** Projection of the leech population response into the 2-dimensional jPCA space. The two dimensions are plotted versus each other (*top*) and versus time (*bottom*). Units are arbitrary but fixed between axes. **e.** Similar projection for the walking monkey. **f.** Similar projection for the reaching monkey.

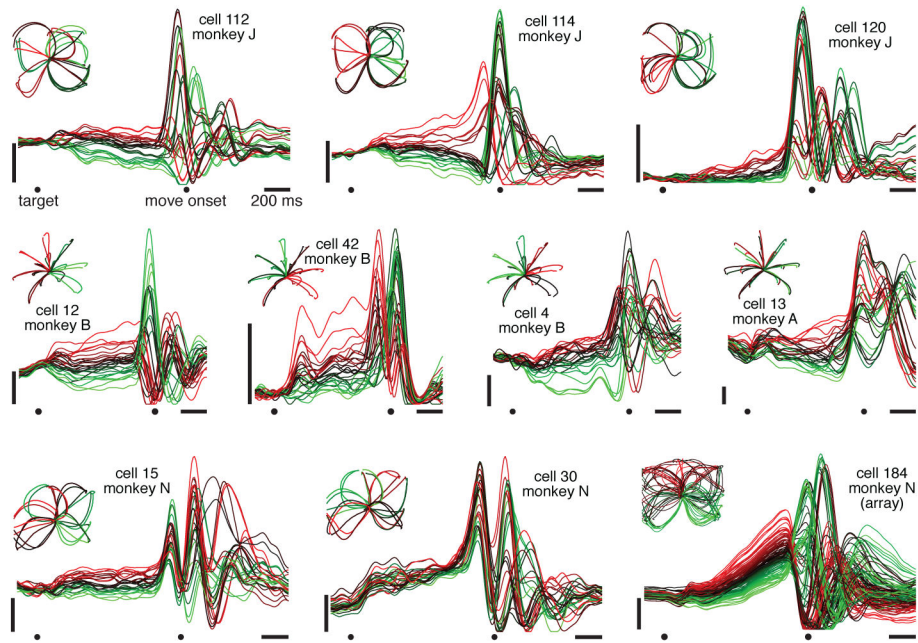


Figure 2.

Firing rate versus time for 10 example neurons, highlighting the multiphasic response patterns. Each trace plots mean across-trial firing rate for one condition, colored based on the firing rate at the end of the preparatory period. Data were averaged separately locked to target onset, the go cue, and movement onset. To aid viewing, traces are interpolated across the gaps between epochs. Vertical scale bars indicate 20 spikes/s. Insets plot hand trajectories, which are different for each dataset.

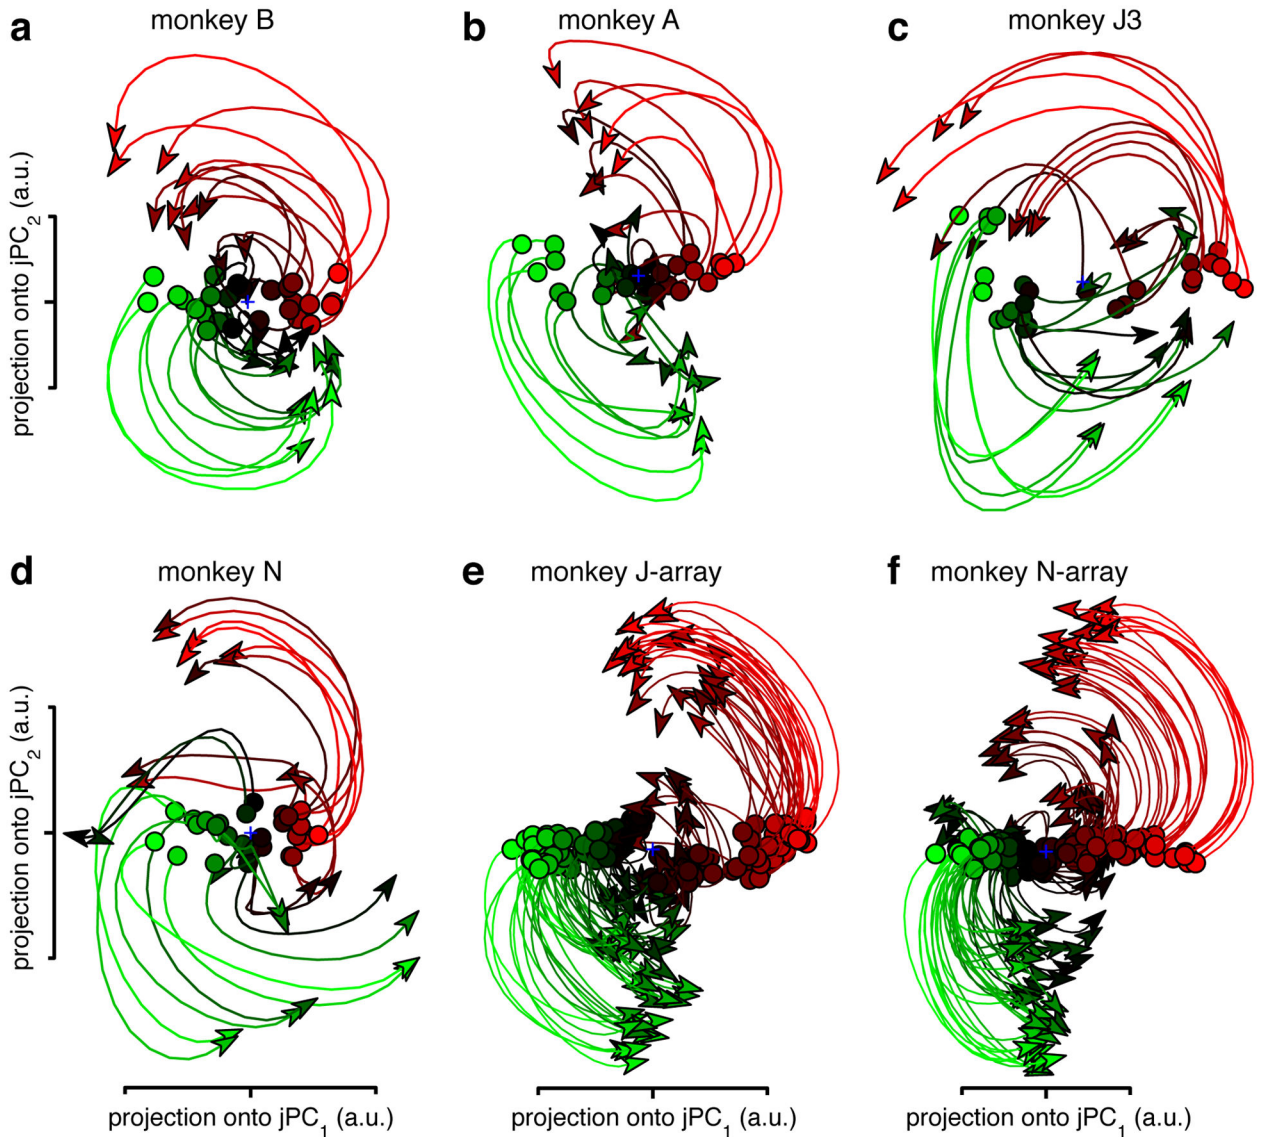


Figure 3.

Projections of the neural population response. *a.* Projection for monkey B (74 neurons; 28 straight-reach conditions). Each trace (one condition) plots the first 200 ms of movement-related activity away from the preparatory state (circles). Traces are colored based on the preparatory-state projection onto jPC_1 . *b.* Projection for monkey A (64 neurons; 28 straight-reach conditions). *c.* Monkey J3 (55 neurons; 27 straight- and curved-reach conditions). *d.* Monkey N (118 neurons; 27 straight- and curved-reach conditions). *e.* Monkey J-array (146 isolations; 108 straight- and curved-reach conditions). *f.* Monkey N-array (218 isolations; 108 straight- and curved-reach conditions).

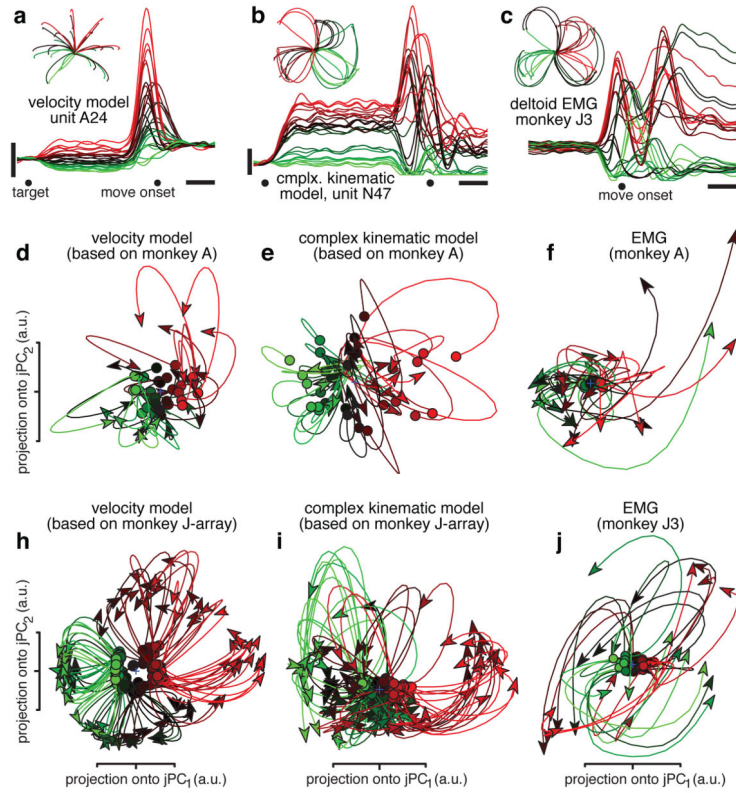


Figure 4. Projections of simulated neural and muscle populations. *a.* Simulated ‘velocity model’ unit, based on hand velocities of monkey A. The preferred direction points up and right. Presentation and scale bars as in figure 2. *b.* Simulated unit from the ‘complex kinematic model’. *c.* EMG from the deltoid (monkey J3), colored by initial response. EMG was rectified, smoothed, and averaged across trials. *d.* Projection of the ‘velocity model’ population response (64 simulated neurons) for monkey A. Identical presentation and analysis as figure 3. *e.* Projection of the ‘complex-kinematic model’ population response for monkey A. *f.* Projection of the recorded muscle population response for monkey A. *h.* Same as *d* but for monkey J-array. *i.* Same as *e* but for monkey J-array. *j.* Same as *f* but for monkey J3.

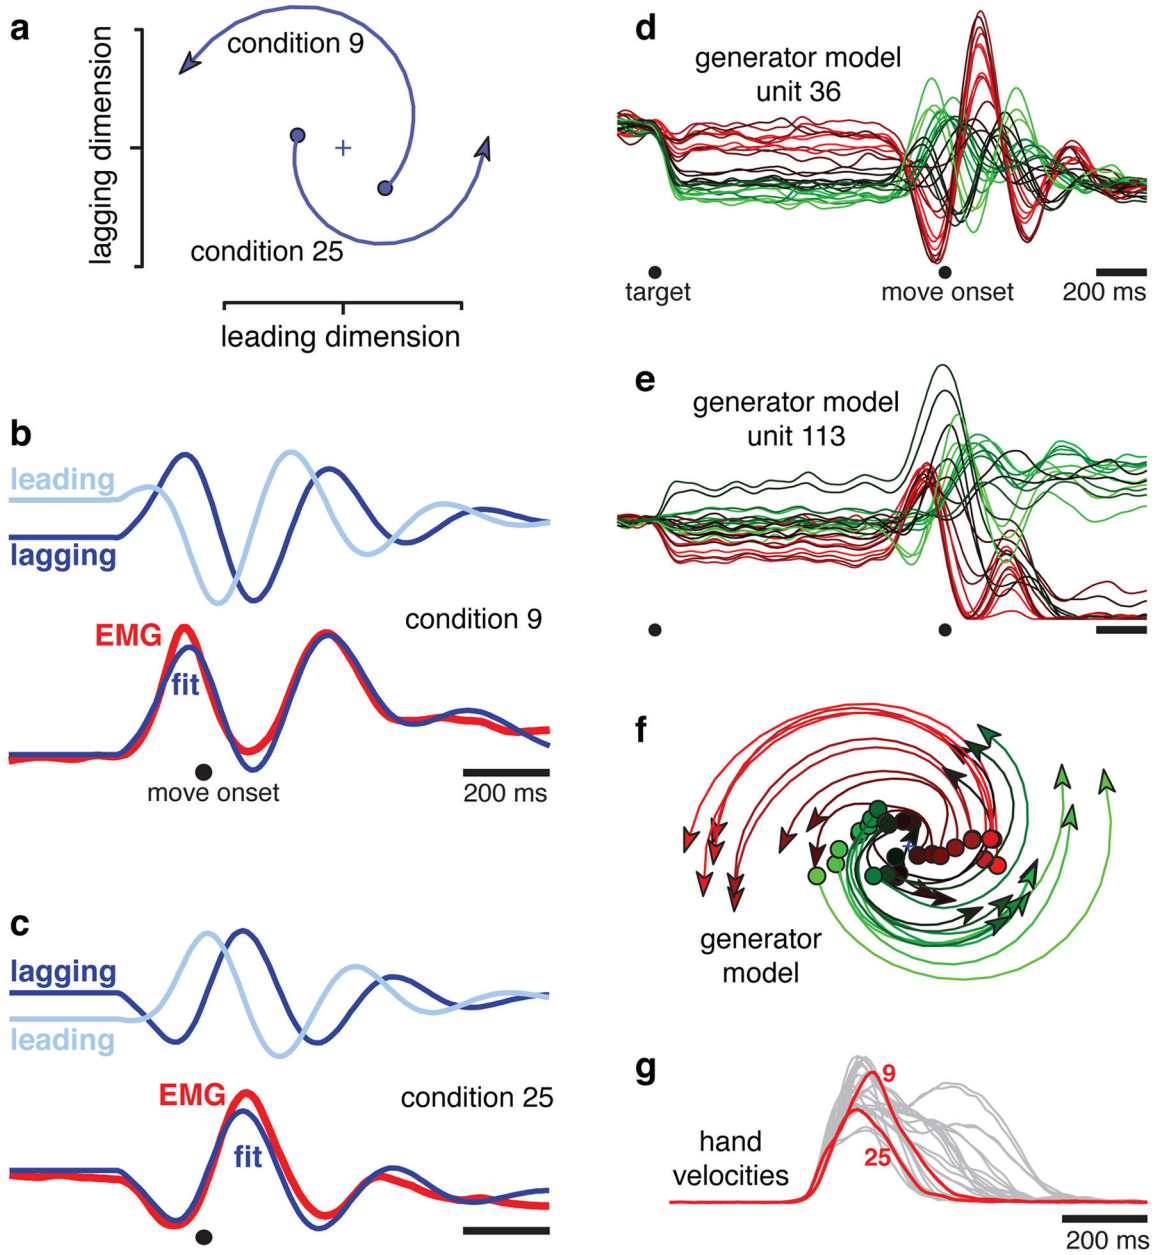


Figure 5. Illustration of how a simple model generates fits to EMG using a pair of brief rotations. **a.** The higher frequency rotation (2.8 Hz) for conditions 9 and 25, plotting the first 200 ms of evolution away from the preparatory state (filled circles). The preparatory state determines rotation amplitude and phase. One state dimension (‘leading’, on the x-axis) is always 90° ahead of the other (‘lagging’, on the y-axis). **b.** Projections onto the leading and lagging dimensions (light and dark blue) versus time (condition 9). The fit to the EMG is the sum of lagging components from the 2.8 Hz rotation (shown) and a lower frequency rotation (0.3 Hz, not shown). **c.** Similar to **b**, but for condition 25. The 2.8 Hz rotation is ~180° out of phase with that in **b**. **d.** Simulated response of a unit from the generator model (methods).

Presentation as in figure 2. *e*. A second simulated unit. *f*. jPCA projection of the simulated population; compare with the neural data in figure 3*c*. *g*. Hand velocities for the 27 conditions in *d*, *e*, *f*. Red traces show conditions 9 and 25.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

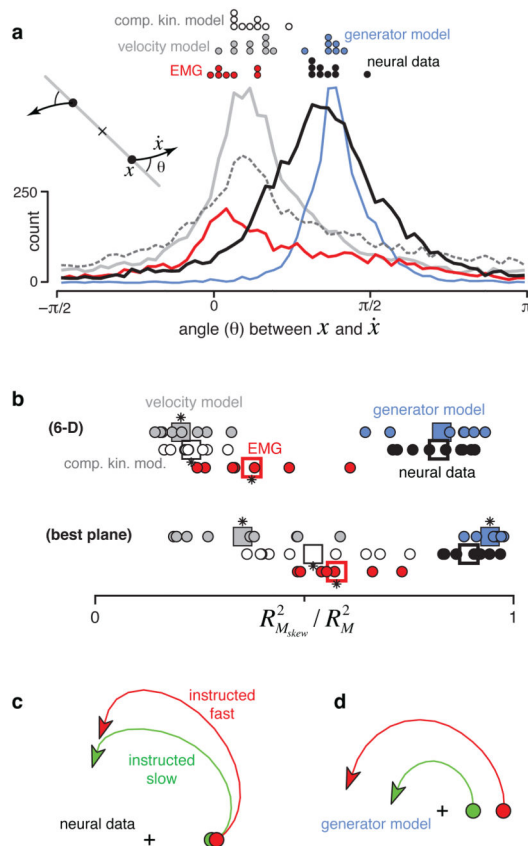


Figure 6.

Consistency of rotational dynamics for real and simulated data. **a.** Histograms of the angle between the neural state, x , and its derivative, \dot{x} , for real and simulated data. The angle was measured after projecting the data onto the first jPCA plane as illustrated schematically (*inset*). Pure rotation results in angles near $\pi/2$; pure scaling/expansion results in angles near 0. Distributions include all analyzed times and conditions. Dots at top show distribution peaks for individual datasets. **b.** Quality of the fit (R^2) provided by M_{skew} relative to an unconstrained M . We assessed fit quality for both the rank 6 matrices that capture dynamics in all 6 analyzed PCs/jPCs (*bottom row*) and the rank 2 matrices that capture dynamics in the first jPCA plane (*top row*). Circles plot performance for individual datasets. Squares give overall averages. Asterisks indicate a significant difference (t-test, $p < 0.05$) from neural data. **c.** Average (across monkeys A and B) neural trajectory for all instructed-slow conditions (*green*) and all instructed-fast conditions (*red*). **d.** Similar to *c* but for the generator model.