

# Neural population dynamics during reaching

Mark M. Churchland<sup>1,2,3\*</sup>, John P. Cunningham<sup>4,5\*</sup>, Matthew T. Kaufman<sup>2,3</sup>, Justin D. Foster<sup>2</sup>, Paul Nuyujukian<sup>6,7</sup>, Stephen I. Ryu<sup>2,8</sup> & Krishna V. Shenoy<sup>2,3,6,9</sup>

**Most theories of motor cortex have assumed that neural activity represents movement parameters. This view derives from what is known about primary visual cortex, where neural activity represents patterns of light. Yet it is unclear how well the analogy between motor and visual cortex holds. Single-neuron responses in motor cortex are 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. Here we find that motor cortex responses during reaching contain a brief but strong oscillatory component, something quite unexpected for a non-periodic behaviour. Oscillation amplitude and phase followed naturally from the preparatory state, suggesting a mechanistic role for preparatory neural activity. These results demonstrate an unexpected yet surprisingly simple structure in the population response. This underlying structure explains many of the confusing features of individual neural responses.**

Motor and premotor cortex were among the first cortical areas to be extensively studied<sup>1</sup>. Yet their basic response properties are poorly understood, and it remains controversial whether neural activity relates to muscles or to abstract movement features<sup>2–13</sup>. At the heart of this debate is the complexity of individual neural responses, which exhibit a great variety of multiphasic patterns<sup>4,14,15</sup>. One explanation is that responses represent many movement parameters:

$$r_n(t) = f_n(\text{param}_1(t), \text{param}_2(t), \text{param}_3(t) \dots) \quad (1)$$

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<sup>4,8,14–17</sup>. In its simplest, deterministic form this can be expressed as:

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

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,  $\mathbf{f}$  is an unknown function, and  $\mathbf{u}$  is an external input. In this conception, neural responses reflect underlying dynamics and display ‘tuning’ only incidentally<sup>18,19</sup>. 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<sup>20–22</sup>.

## 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  $\sim 1.5$  Hz during swimming<sup>23</sup>, and many single neurons display firing rate oscillations at that frequency (Fig. 1a)<sup>24,25</sup>. Rhythmic structure was also present for cortical responses in the walking monkey:  $\sim 1$  Hz oscillations matching the  $\sim 1$  Hz movement of the arm (Fig. 1b). If single-neuron oscillations

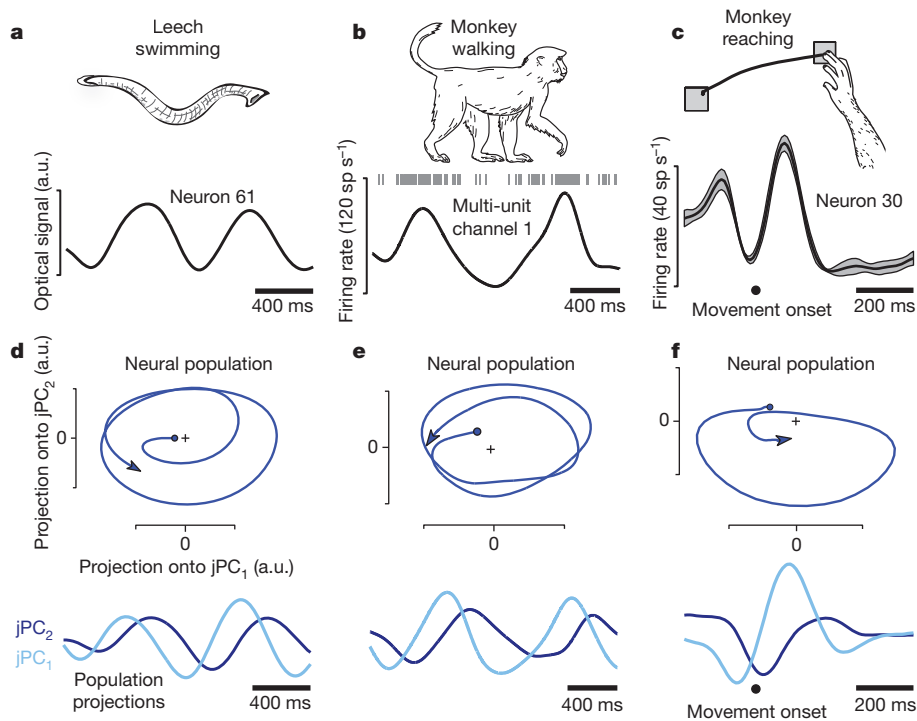
are generated by population-level dynamics, then the population response (the neural state) should rotate with time<sup>15</sup>, much as the state of a pendulum rotates in the space defined by velocity and position. We projected the population response onto a two-dimensional state space and found rotations 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, although not trivial, are largely expected for a neural dynamical system that generates rhythmic output<sup>22</sup>.

The projections in Fig. 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 used a novel method that finds an informative plane within the top principal components (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 six 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 ( $\sim 1$  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<sup>15,22</sup>. There might be a concern that the patterns in

<sup>1</sup>Department of Neuroscience, Kavli Institute for Brain Science, David Mahoney Center, Columbia University Medical Center, New York, New York 10032, USA. <sup>2</sup>Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA. <sup>3</sup>Neurosciences Program, Stanford University, Stanford, California 94305, USA. <sup>4</sup>Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, Missouri 63130, USA. <sup>5</sup>Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK. <sup>6</sup>Department of Bioengineering, Stanford University, Stanford, California 94305, USA. <sup>7</sup>Stanford University School of Medicine, Stanford, California 94305, USA. <sup>8</sup>Department of Neurosurgery, Palo Alto Medical Foundation, Palo Alto, California 94301, USA. <sup>9</sup>Department of Neurobiology, Stanford University School of Medicine, Stanford, California 94305, USA.

\*These authors contributed equally to this work.



**Figure 1 | Oscillation of neural firing rates during three movement types.** **a**, Response of 1 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. a.u., arbitrary units. **b**, Multi-unit response from 1 of 96 electrodes implanted in the arm representation of caudal premotor cortex. Data from 32 such channels were wirelessly transmitted during walking.  $\text{sp s}^{-1}$ , spikes per second. Responses (not averaged across repetitions) were filtered

Fig. 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 analysed 469 single-neuron recordings from motor and premotor cortex of four monkeys (identified as 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 analysed 9 data sets, each using 27–108 reach types ('conditions'). For each neuron and condition we computed and analysed the average across-trial firing rate.

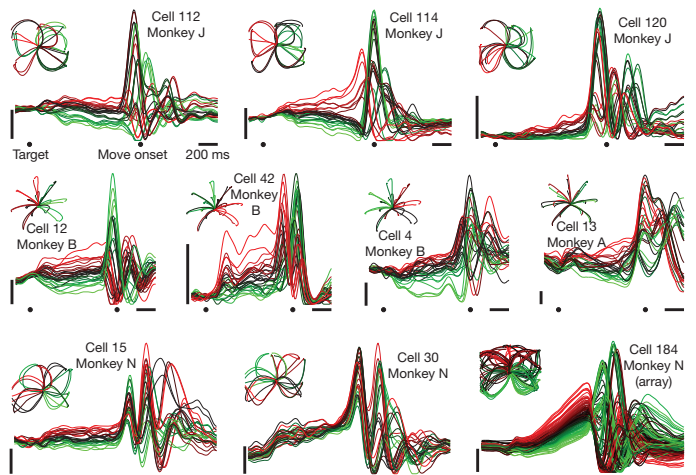
Most neurons exhibited preparatory and movement-related responses (Fig. 2). Responses were typically complex, multiphasic and heterogeneous<sup>14</sup>. Yet there appear to be oscillations in many single-neuron responses, beginning just before movement onset and lasting for  $\sim 1$ –1.5 cycles. These quasi-oscillatory patterns were 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<sup>14,15,17,18,26</sup>. 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 conditions<sup>15</sup>, even when reaches are in opposition.

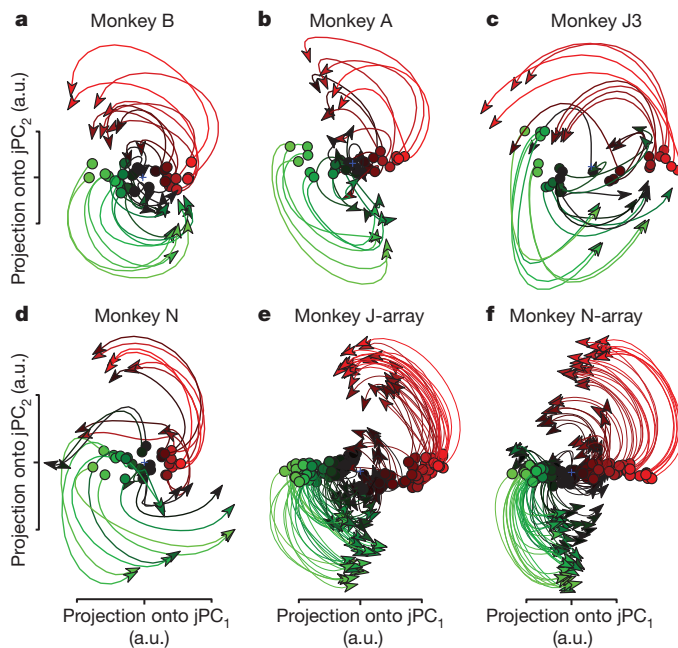
We projected the population response for all conditions onto the jPC plane. This was done for 200 ms of data, beginning when

with a 100 ms Gaussian kernel. **c**, Response of 1 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 s.e.m.). **d**, Projection of the leech population response into the two-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.

preparatory activity transitions to movement-related activity (Supplementary Movie 3 shows a longer span of time). The resulting projections (Fig. 3a–f) show four notable features. First, rotations of



**Figure 2 | Firing rate versus time for ten example neurons, highlighting the multiphasic response patterns.** Each trace plots mean across-trial firing rate for one condition. Traces are coloured red to green based on the level of preparatory activity observed for that neuron. This allows inspection of how the pattern of preparatory tuning changes during the movement. 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  $\text{spikes s}^{-1}$ . Insets plot hand trajectories, which are different for each data set. Traces are coloured using the same code as for the neural data: red traces indicate those conditions with the greatest preparatory response.



**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 coloured on the basis of the preparatory-state projection onto  $jPC_1$ , a.u., arbitrary units. **b**, Projection for monkey A (64 neurons; 28 straight-reach conditions). **c**, Monkey J, data set 3 (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).

the neural state are prevalent during reaching. Second, the neural state rotates in the same direction across conditions (Fig. 3, different traces). Third, the rotation phase follows naturally from the preparatory state. Last, 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 anticlockwise-curving reaches (Fig. 2, insets). Yet for each data set 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 the initial population-level preparatory state is known (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 zero<sup>15</sup>. Nevertheless, the ordered state-space rotations relate directly 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). Single neurons strongly reflect combinations of these underlying patterns. However, that underlying structure is not readily apparent when plotting each pattern alone, or each neuron individually<sup>15</sup>. 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  $\sim 2.5$  Hz) makes them easy to miss unless trial counts are high (data sets averaged 810 trials per 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 Fig. 3. Those projections captured an average of 28% of the total data variance, and thus reveal patterns that are strongly present in the data. However, one must be concerned that

such patterns could have appeared by accident or for trivial reasons. To address this possibility, multiple ‘shuffle’ controls demonstrate that  $jPCA$  does not find rotational structure when such structure is not present (Supplementary Figs 2, 3 and Supplementary Movie 4). Similarly, rotations in the walking monkey were not erroneously found when the monkey was stationary (Supplementary Movie 1).

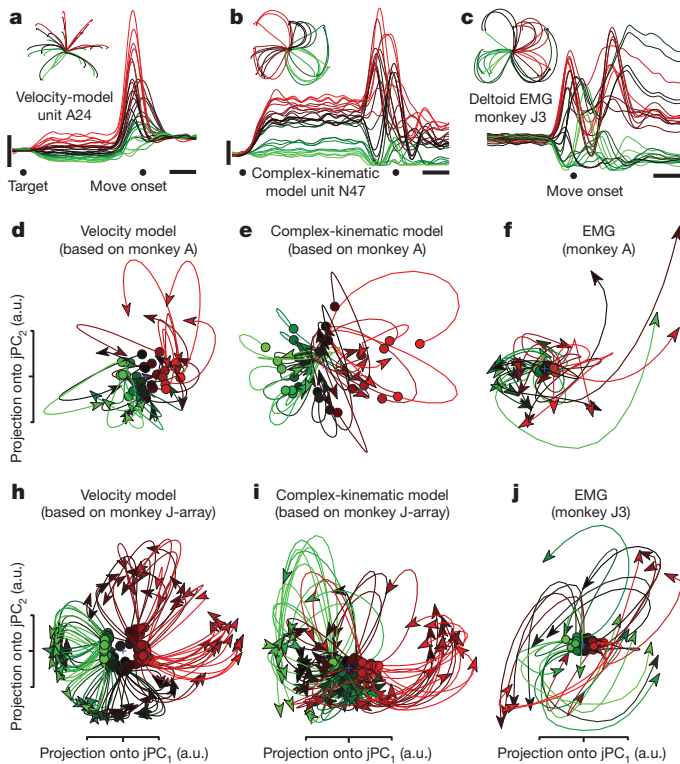
The fact that a single plane (two dimensions) captures an average of 28% of the total data variance is notable, given the high dimensionality of the data itself<sup>14</sup>. As a comparison, the dimensions defined by  $PC_2$  and  $PC_3$  (which by definition capture the second- and third-most data 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 primary motor cortex and dorsal premotor cortex 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 speed<sup>27</sup>. Simulated preparatory activity was tuned for reach direction and distance<sup>28</sup>. We simulated one ‘velocity model’ data set per recorded data set, based upon the recorded velocities and 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 data sets (summary analysis below) including those with curved reaches (for example, Fig. 4h). We also simulated a ‘complex-kinematic’ model in which responses reflected the weighted sum of kinematic parameters (position, velocity, acceleration and jerk) that correlate with muscle activity<sup>6</sup>. This model produced multiphasic responses but not consistent rotations (Fig. 4b, e, i). We also recorded EMG (electromyograms) from a population of muscles (6–12 recordings per data set). Although EMG is strongly multiphasic, the population of muscles did not show consistent rotations (Fig. 4c, 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 it is not immediately apparent how 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, when fitting the deltoid EMG for dataset J3 (the third data set from monkey J) the higher-frequency rotation in the model occurred at 2.8 Hz (Fig. 5a). Different conditions (9 and 25 are shown) involved different amplitudes and phases, set by the preparatory state of the model. The vertical, ‘lagging’ dimension drove simulated muscle activity, and the projections onto that dimension (Fig. 5b, 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. 5b, c and Supplementary Figs 7 and 8). The fit/EMG correlation ranged from

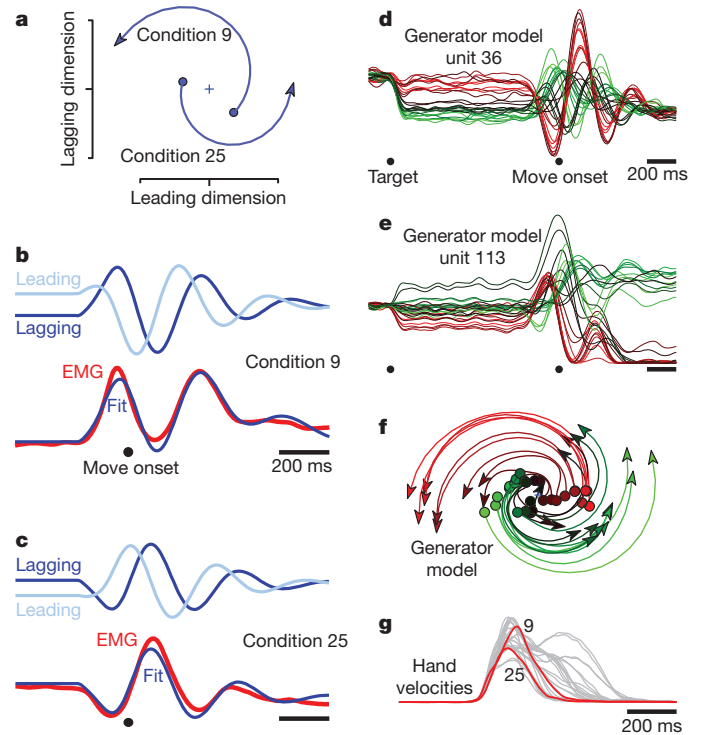


**Figure 4 | Projections of simulated neural and muscle populations.** **a**, Simulated velocity-model unit, based on hand velocities of monkey A. The preferred direction points up and right. Presentation and scale bars as in Fig. 2. **b**, Simulated unit from the complex-kinematic model. **c**, EMG from the deltoid (monkey J, data set 3), coloured 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 Fig. 3. a.u., arbitrary units. **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 J, data set 3.

0.97–0.99 across data sets. 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 need not imply that cortex directly controls muscles. Yet it illustrates the plausibility of direct control<sup>6</sup>, and demonstrates that rotations provide a natural basis for generating non-rhythmic movement.

Faster reaches might have been expected to involve faster rotations, or longer reaches to involve longer rotations. However, EMG could be fit using two rotations with fixed frequency and duration. This was true even though the 27 conditions differed greatly in reach speed and duration (Fig. 5g). 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. 5d, e and Supplementary Fig. 8). Both real and simulated responses exhibited reversals in ‘preferred condition’<sup>14</sup> and a weak correlation between preparatory and movement-period ‘tuning’<sup>15,29</sup>. Despite such surface complexity, jPCA projections of the simulated populations



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

successfully reveal the simple underlying rotations (Fig. 5f). 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 (*ẋ*) for every condition and time (Fig. 6a). The first jPCA plane is oriented such that the average rotational tendency of the data—however weak or strong—is anticlockwise. 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:

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

where  $x(t,c)$  is the *k*-dimensional ( $k = 6$  for all figures) population state for time *t* and condition *c*.  $M_{\text{skew}}$  is a skew-symmetric matrix ( $M_{\text{skew}} = -M_{\text{skew}}^T$ ) that captures rotational dynamics. The first two jPCs were the two eigenvectors associated with the largest-magnitude eigenvalues of  $M_{\text{skew}}$ , which indicate the strongest rotational plane in the dynamical system fit by equation (3). Projecting data onto those jPCs (as in Figs 3 and 4) reveals the prevalence of rotations. We can also assess rotation prevalence by quantifying how well  $M_{\text{skew}}$  fits the data relative to an unconstrained matrix *M*. That unconstrained

matrix  $M$  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  $M_{\text{skew}}$  was nearly as good as the fit provided by the best unconstrained  $M$  (Fig. 6b). For the velocity-tuned and complex-kinematic models,  $M_{\text{skew}}$  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. 6a, b, red), underscoring a central point: state-space rotations result not from a multiphasic

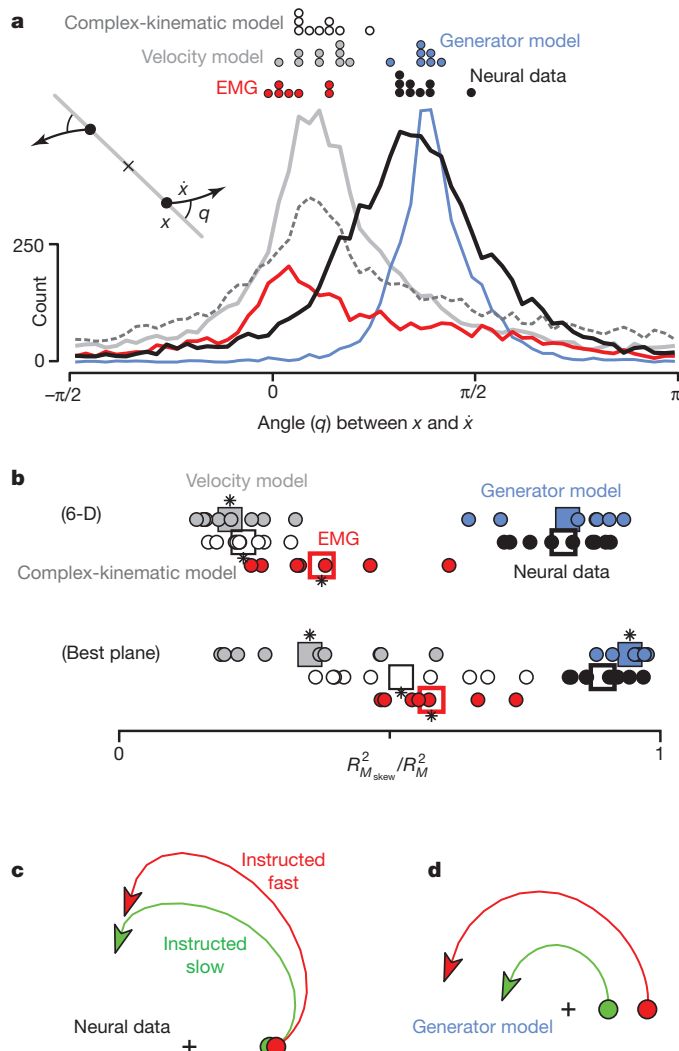
signal, but from how that signal is constructed. For example, the generator model exhibits rotations even though the EMG does not. More generally, many features of the observed rotations make sense in terms of how outputs (EMG, kinematics) might be generated, rather than in terms of the outputs themselves. 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 data sets where target colour 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. 6c, d). The same point can be made by separately applying jPCA for fast and slow reaches: the largest eigenvalues of  $M_{\text{skew}}$  were actually slightly smaller for fast reaches (8.8 versus 9.8 radians  $\text{s}^{-1}$  for monkey A, 12.2 versus 13.5 radians  $\text{s}^{-1}$  for monkey B). Rotation amplitude, rather than frequency, differed between speeds (Fig. 6c, 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 from chance or 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<sup>14,30</sup>, and the lack of correlation between preparatory and movement-period tuning<sup>15,29</sup> (Fig. 5d, e and Supplementary Fig. 8). State-space rotations produce briefly oscillatory temporal patterns that provide an effective basis for producing multiphasic muscle activity, suggesting that non-periodic movements may be generated via neural mechanisms resembling those that generate rhythmic movement<sup>20,22,31–33</sup>.

Recent results suggest that preparatory activity sets the initial state of a dynamical system, whose subsequent evolution produces movement activity<sup>15</sup>. Aspects of these dynamics—a rotation away from the preparatory state—appear straightforward. However, the circuitry that creates these dynamics is unclear; it may be purely local, or may involve recurrent circuitry<sup>34</sup> that links motor cortex with the spinal cord and with subcortical structures<sup>35</sup>. Peripheral feedback is also expected to shape neural dynamics<sup>36</sup>, 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 antisymmetric connectivity: a given neural dimension (for example, jPC<sub>1</sub>) positively influences another (for example, jPC<sub>2</sub>), which negatively influences the first. However, it is unclear whether this population-level pattern directly reflects a circuit-level dominance of antisymmetric connectivity. We also stress that although rotations are a dominant pattern in the data across multiple dimensions (Supplementary Fig. 4), non-rotational components exist as well, perhaps reflecting the nonlinear dynamics necessary for initiating or terminating movement, for stability<sup>37</sup>, and for feedback control<sup>16,38</sup>.

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 population-level dynamical systems perspective. It therefore seems increasingly likely that motor cortex can be understood in relatively straightforward terms: as an engine of movement that uses lawful dynamics.



**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 as illustrated schematically (inset) after projecting the data onto the first jPCA plane. Pure rotation results in angles near  $\pi/2$ ; pure scaling/expansion results in angles near 0. Distributions include all analysed times and conditions. Dots at top show distribution peaks for individual data sets. **b**, Quality of the fit ( $R^2$ ) provided by  $M_{\text{skew}}$  relative to an unconstrained  $M$ . We assessed fit quality for both the rank 6 matrices that capture dynamics in all 6 analysed dimensions (6-D; top row) and the rank 2 matrices that capture dynamics in the first (best) jPCA plane (bottom row). Circles plot performance for individual data sets. 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.

## METHODS SUMMARY

Optical recordings from the isolated leech central nervous system were made by K. Briggman and W. Kristan and have been described previously<sup>24,25</sup>. 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 a 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.

**Full Methods** and any associated references are available in the online version of the paper at [www.nature.com/nature](http://www.nature.com/nature).

Received 6 May 2011; accepted 5 April 2012.

Published online 3 June 2012.

- Lemon, R. N. An enduring map of the motor cortex. *Exp. Physiol.* **93**, 798–802 (2008).
- Evarts, E. V. Relation of pyramidal tract activity to force exerted during voluntary movement. *J. Neurophysiol.* **31**, 14–27 (1968).
- Mussa-Ivaldi, F. A. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci. Lett.* **91**, 106–111 (1988).
- Fetz, E. E. Are movement parameters recognizably coded in the activity of single neurons? *Behav. Brain Sci.* **15**, 679–690 (1992).
- Sanger, T. D. Theoretical considerations for the analysis of population coding in motor cortex. *Neural Comput.* **6**, 29–37 (1994).
- Todorov, E. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nature Neurosci.* **3**, 391–398 (2000).
- Hatsopoulos, N. G. Encoding in the motor cortex: was Evarts right after all? Focus on “motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks”. *J. Neurophysiol.* **94**, 2261–2262 (2005).
- Scott, S. H. Inconvenient truths about neural processing in primary motor cortex. *J. Physiol. (Lond.)* **586**, 1217–1224 (2008).
- Morrow, M. M., Pohlmeier, E. A. & Miller, L. E. Control of muscle synergies by cortical ensembles. *Adv. Exp. Med. Biol.* **629**, 179–199 (2009).
- Aflalo, T. N. & Graziano, M. S. A. Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex. *J. Neurosci.* **27**, 2760–2780 (2007).
- Kalaska, J. F. From intention to action: motor cortex and the control of reaching movements. *Adv. Exp. Med. Biol.* **629**, 139–178 (2009).
- Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986).
- Kakei, S., Hoffman, D. S. & Strick, P. L. Muscle and movement representations in the primary motor cortex. *Science* **285**, 2136–2139 (1999).
- Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* **97**, 4235–4257 (2007).
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
- Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nature Neurosci.* **5**, 1226–1235 (2002).
- Graziano, M. S. New insights into motor cortex. *Neuron* **71**, 387–388 (2011).
- Churchland, M. M., Santhanam, G. & Shenoy, K. V. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *J. Neurophysiol.* **96**, 3130–3146 (2006).
- Cisek, P. Preparing for speed. Focus on: “preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach”. *J. Neurophysiol.* **96**, 2842–2843 (2006).
- Georgopoulos, A. P. & Grillner, S. Visuomotor coordination in reaching and locomotion. *Science* **245**, 1209–1210 (1989).
- Grillner, S. Biological pattern generation: the cellular and computational logic of networks in motion. *Neuron* **52**, 751–766 (2006).
- Rokni, U. & Sompolinsky, H. How the brain generates movement. *Neural Comput.* **24**, 289–331 (2011).
- Kristan, W. B. Jr & Calabrese, R. L. Rhythmic swimming activity in neurones of the isolated nerve cord of the leech. *J. Exp. Biol.* **65**, 643–668 (1976).
- Briggman, K. L. & Kristan, W. B. Jr. Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *J. Neurosci.* **26**, 10925–10933 (2006).
- Briggman, K. L., Abarbanel, H. D. & Kristan, W. B. Jr. Optical imaging of neuronal populations during decision-making. *Science* **307**, 896–901 (2005).
- Shenoy, K. V., Kaufman, M. T., Sahani, M. & Churchland, M. M. in *Progress in Brain Research: Enhancing Performance for Action and Perception* (eds Green, A., Chapman, E., Kalaska, J. F. & Lepore, F.) (Elsevier, 2011).
- Moran, D. W. & Schwartz, A. B. Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* **82**, 2676–2692 (1999).
- Messier, J. & Kalaska, J. F. Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *J. Neurophysiol.* **84**, 152–165 (2000).
- Kaufman, M. T. *et al.* Roles of monkey premotor neuron classes in movement preparation and execution. *J. Neurophysiol.* **104**, 799–810 (2010).
- Sergio, L. E., Hamel-Paquet, C. & Kalaska, J. F. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *J. Neurophysiol.* **94**, 2353–2378 (2005).
- Llinas, R. *I of the Vortex* (MIT Press, 2002).
- Yuste, R., MacLean, J. N., Smith, J. & Lansner, A. The cortex as a central pattern generator. *Nature Rev. Neurosci.* **6**, 477–483 (2005).
- 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. *J. Neurophysiol.* **105**, 388–409 (2011).
- Middleton, F. A. & Strick, P. L. Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res. Brain Res. Rev.* **31**, 236–250 (2000).
- Thorn, C. A., Atallah, H., Howe, M. & Graybiel, A. M. Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. *Neuron* **66**, 781–795 (2010).
- Herter, T. M., Korbel, T. & Scott, S. H. Comparison of neural responses in primary motor cortex to transient and continuous loads during posture. *J. Neurophysiol.* **101**, 150–163 (2008).
- Sussillo, D. & Abbott, L. F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).
- Scott, S. H. Optimal feedback control and the neural basis of volitional motor control. *Nature Rev. Neurosci.* **5**, 532–546 (2004).

**Supplementary Information** is linked to the online version of the paper at [www.nature.com/nature](http://www.nature.com/nature).

**Acknowledgements** 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 information technology support. We thank Z. Ghahramani and C. Rasmussen for discussion of jPCA and related methods. We thank D. Sussillo, S. Grossman and M. Sahani for analysis suggestions and commentary on the manuscript. This work was supported by a Helen Hay Whitney postdoctoral fellowship and National Institutes of Health (NIH) postdoctoral training fellowship (M.M.C.), the Burroughs Wellcome Fund Career Awards in the Biomedical Sciences (M.M.C., K.V.S.), Engineering and Physical Sciences Research Council grant EP/H019472/1 and the McDonnell Center (J.P.C.), a National Science Foundation graduate research fellowship (M.T.K.), a Texas Instruments Stanford Graduate Fellowship (J.D.F.), a Paul and Daisy Soros Fellowship (P.N.), the Stanford Medical Scientist Training Program (P.N.), and these awards to K.V.S.: NIH Director’s Pioneer Award (1DP10D006409), 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, the McKnight Foundation, the Sloan Foundation and the Weston Havens Foundation.

**Author Contributions** The jPCA method was designed by J.P.C. and M.M.C. M.M.C. and M.T.K. collected data from the reaching monkeys. J.D.F. and P.N. collected data from the walking monkey. S.I.R. led the array implantation surgeries. K.V.S. contributed to all aspects of the work. All authors discussed the results and commented on the analyses and manuscript.

**Author Information** Reprints and permissions information is available at [www.nature.com/reprints](http://www.nature.com/reprints). The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at [www.nature.com/nature](http://www.nature.com/nature). Correspondence and requests for materials should be addressed to M.M.C. ([mc3502@columbia.edu](mailto:mc3502@columbia.edu)).

## 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<sup>24,25</sup>. Recordings from monkey cortex were made using both a delayed reach task (with head restraint) and from an unrestrained monkey walking on a treadmill<sup>39–41</sup>. 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<sup>15,29,42</sup>. Briefly, four male rhesus monkeys (A, B, J and N) performed delayed reaches on a fronto-parallel screen. Delays ranged from 0–1,000 ms (the exact range varied by monkey). Only trials with delays >400 ms were analysed. Fixation was enforced (at the central spot) during the delay for monkeys J and N. We used two variants of a centre-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 colour (28 total conditions)<sup>18</sup>. 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 data sets (J1–J4). For the monkey J-array and N-array data sets, 108 conditions were presented in the same recording session. Recordings from three of the four monkeys (A, B and J) have been analysed in prior publications (for example, ref. 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 Fig. 5 shows the central analysis divided by area. Seven of the nine data sets (monkey A, B, J1, J2, J3, J4 and N) were recorded with conventional single-electrode techniques. These data sets involved a total of 469 single-unit isolations. The other two data sets (monkey J-array, recorded 18 September 2009; monkey N-array, recorded 23 September 2010) used pairs of chronically implanted 96-electrode arrays (Blackrock Microsystems). These array-based data sets involved, respectively, 146 and 218 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. Although there are important differences in their average response properties (for example, 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 them on the basis of either anatomy or response properties (although see Supplementary Fig. 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<sup>39,41</sup>. Behavioural 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 previously<sup>42</sup>. 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 data sets contained multiple recordings from each muscle (for example, one from each of the three heads of the deltoid). The total number of EMG recordings for some data sets was thus as high as 12. EMG was recorded for all data sets except those recorded using arrays.

**Computing average firing rate as a function of time.** Average trial counts were high (an average of 810 trials per 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 behaviour, and by discarding the few trials for which behaviour was not tightly stereotyped. Average firing rates were further de-noised by filtering with a Gaussian (20 or 24 ms depending on the data set) and using a custom-developed smoothing method that discards idiosyncratic features that are both small and not shared across conditions (see supplementary figure 4 in ref. 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 firing rates, but had essentially no effect on any of the population-level analyses

(Supplementary Fig. 10). EMG recordings and ‘recordings’ of simulated neurons were processed using all the same steps as for the neural 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 data set 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 (for example, Fig. 5a, 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 data set the rotation always had the same frequency regardless of condition, with a rise and decay defined by the same windowing gamma function (for example, 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 data set, 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 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 offset.

**Simulated neural data.** We produced two classes of simulated neural data sets. 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 (for example, see ref. 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<sup>6,30</sup>. 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<sup>-1</sup>) m<sup>-1</sup>, 10 (spikes s<sup>-1</sup>)/(m s<sup>-1</sup>), 1 (spikes s<sup>-1</sup>)/(m s<sup>-2</sup>), 0.05 (spikes s<sup>-1</sup>)/(m s<sup>-3</sup>). These constants are taken from a published model<sup>6</sup>, 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, although 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 (Figs 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’ simulated spikes that were produced via a gamma-interval process (order 2) based on the underlying firing rate. For each neural data set, 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 analysed just as for the actual neural data. The velocity and complex-kinematic models each produced nine simulated data sets (one for every real data set). The generator model produced seven: it could not be simulated for the J-array or N-array data sets, 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 analysed 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 pre-processing 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-centred 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 pre-processing 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 (for example, the population vector) achieve the same end implicitly: non-directional aspects of the response cancel out. The pre-processing steps (and all subsequent analysis steps) were applied in the same way to all data sets, real and simulated.

The most critical pre-processing 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 analysed 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 data<sup>14</sup>. The resulting  $6 \times ct$  matrix,  $X_{\text{red}}$ , defines a six-dimensional neural state for every time and condition. By pre-processing with PCA, we ensure that when jPCA is subsequently applied, it reveals only patterns of activity that are strongly present across neurons. Pre-processing 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 Figs 2 and 3).

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

$$\dot{X}_{\text{red}} = M X_{\text{red}} \quad (4)$$

and

$$\dot{X}_{\text{red}} = M_{\text{skew}} X_{\text{red}} \quad (5)$$

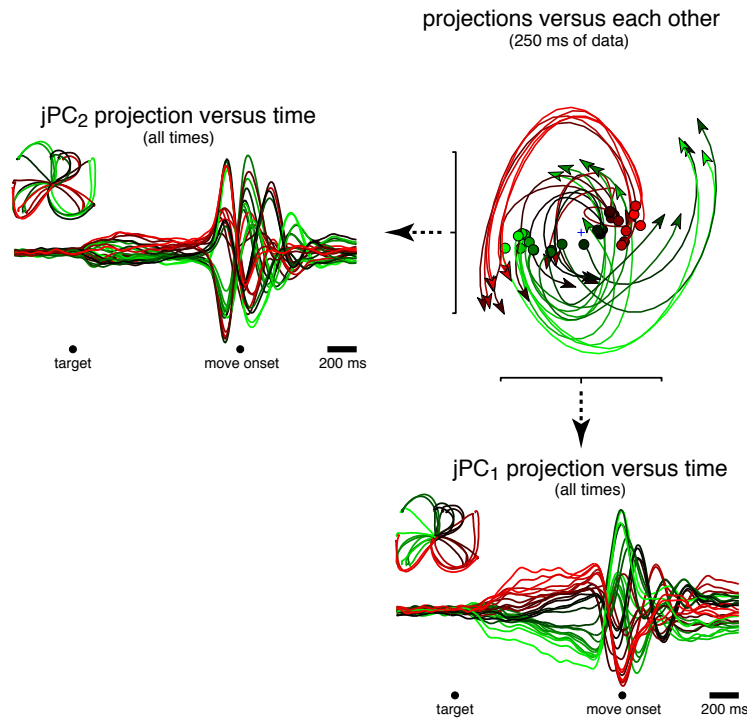
(To keep the dimensions appropriate, the final time point for each condition in  $X_{\text{red}}$  was removed so that it was also  $6 \times c(t - 1)$ ). Thus, we are attempting to find matrices  $M$  and  $M_{\text{skew}}$  that take the state at each time in  $X_{\text{red}}$ , and transform it into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

$M_{\text{skew}}$  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:  $\text{jPC}_1 = V_1 + V_2$ , and  $\text{jPC}_2 = j(V_1 - V_2)$  (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  $X_{\text{jPCA}} = (\text{jPC}_1; \text{jPC}_2) \times X_{\text{red}}$ . The matrix  $X_{\text{jPCA}}$  is thus of size  $2 \times ct$ , and describes the neural state, projected onto two dimensions, for every time and condition. For a given jPCA plane, the choice of orthogonal vectors ( $\text{jPC}_1$  and  $\text{jPC}_2$ ) within that plane is arbitrary. We therefore selected  $\text{jPC}_1$  and  $\text{jPC}_2$  so that any net rotation was anticlockwise (the same choice was of course used across all conditions for a given data set) and so that the spread of preparatory states was strongest along  $\text{jPC}_1$ . 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 six 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).

39. Gilja, V., Chestek, C. A., Nuyujukian, P., Foster, J. D. & Shenoy, K. V. Autonomous head-mounted electrophysiology systems for freely behaving primates. *Curr. Opin. Neurobiol.* **20**, 676–686 (2010).
40. Foster, J. D. *et al.* in *Proc. of the 5th International IEEE EMBS Conference on Neural Engineering* 613–615 (IEEE, 2011).
41. Miranda, H. *et al.* A high-rate long-range wireless transmission system for simultaneous multichannel neural recording applications. *IEEE Trans. Biomed. Circ. Syst.* **4**, 181–191 (2010).
42. Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G. & Shenoy, K. V. Neural variability in premotor cortex provides a signature of motor preparation. *J. Neurosci.* **26**, 3697–3712 (2006).

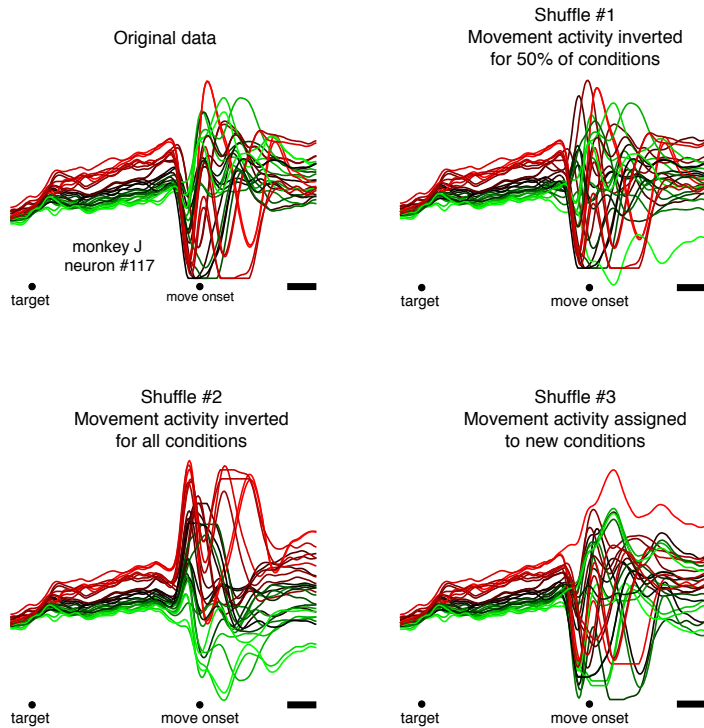




**Supplementary figure 1.** Projection onto each jPCA axis as a function of time. Data is shown for monkey J3. The plot of the jPCA plane uses the same format as in figure 3 of the main text. The plots versus time use the same format as for the individual-neuron PSTHs in figure 2 of the main text (though here the vertical units are arbitrary). Traces are colored red to green based on the level of preparatory activity for that projection. This allows visualization of ‘tuning’ with respect to the reach trajectories (inset). Direction ‘tuning’ is present but imperfect in the projections, much as it is for the neurons upon which the projections are based.

Indeed, in many ways the projections versus time look as if they could be the responses of single neurons. This is not accidental: the jPCA projections capture responses that are strongly present in the responses of individual neurons. Conversely, the jPCA projections are simply weighted sums of individual-neuron responses. Each jPCA projection can thus be interpreted much as with a traditional ‘population average’. The key difference is that the weights used for jPCA are chosen by the algorithm, rather than set by hand according to the ‘preferred direction’.

As with the supplementary movies and other plots that show an extended period of time, these projections were based on the top ten PCs rather than the top six. This is not critical but aids in finding a projection that works well across a wide range of times.



**Supplementary figure 2.** Construction of shuffle controls, illustrated for an example neuron. A potential concern is that the jPCA method might be powerful enough to find state-space rotations for *any* population response that contains diverse and multiphasic responses. This would be a large concern were one analyzing a few conditions in a very high-dimensional space. However, our analyses involved 27-108 conditions, and were applied only after the data dimensionality was reduced using traditional PCA. Nevertheless, it is prudent to empirically evaluate the degree to which rotational patterns might be found by chance simply because responses are diverse and complex.

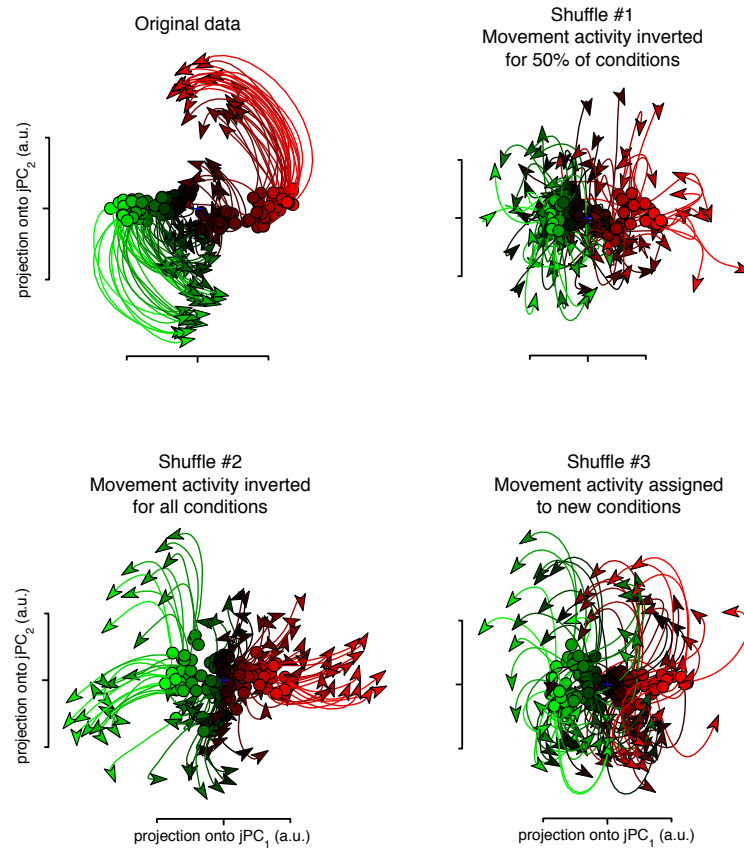
To evaluate this possibility we performed three shuffle controls that preserve response diversity and complexity, but disrupt the deep structure of the data. If robust rotational structure can be seen in the shuffled controls, then this would be a cause for serious concern. It would indicate that rotations can be found by our methods even when not truly present.

All shuffle controls are based on the distinction between preparatory activity (which is left intact) and peri-movement activity (which was shuffled in three different ways). We picked a time-point 50 ms after the go cue as the dividing point between preparatory and peri-movement activity. For the first shuffled control, the pattern of peri-movement activity was inverted for half of the conditions, selected at random. The inversion was performed around the dividing time-point, such that continuity with preparatory activity was preserved. This procedure was performed separately for each neuron.

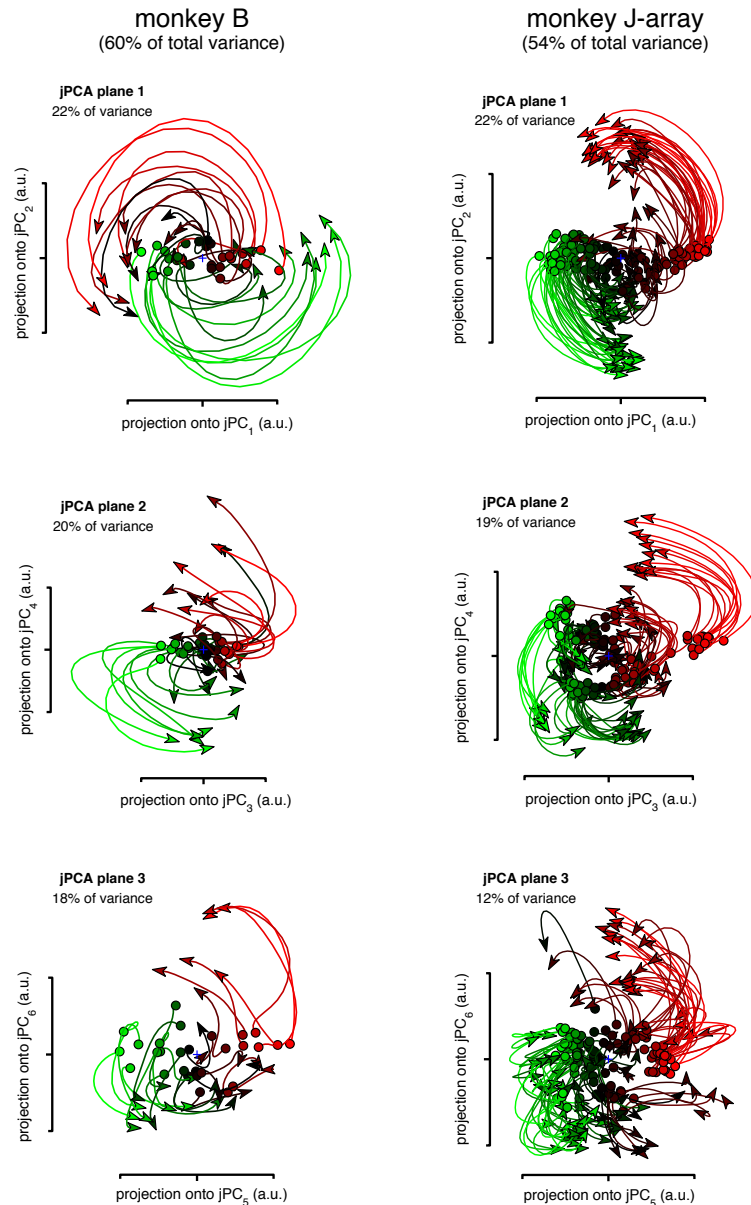
The second shuffle control was similar to the first, but inverted the peri-movement activity pattern for all conditions. This manipulation is not expected to remove all rotational structure (most such structure is merely sign-inverted). However, this manipulation is expected to largely remove any consistent relationship (assuming there is one) between the preparatory state and the phase of subsequent oscillations. Thus, for the time period of interest, this shuffle control is expected to greatly reduce the consistency of any true rotations, especially the relationship between phase and initial state.

The third shuffle control randomly reassigned the peri-movement activity from one condition to the preparatory activity from another. The beginning of the peri-movement pattern was simply appended to the final firing rate during the preparatory state, such that there was no discontinuity. The same reassignment was performed for all neurons. As with the second control, this third shuffle control is not expected to remove all rotational structure (many of the peri-movement activity patterns are altered only modestly by this manipulation). However, any true relationship between rotation phase and initial state is expected to be disrupted. Thus, for the time period of interest, this shuffle control is expected to greatly reduce the consistency of any true rotations, especially the relationship between phase and initial state.

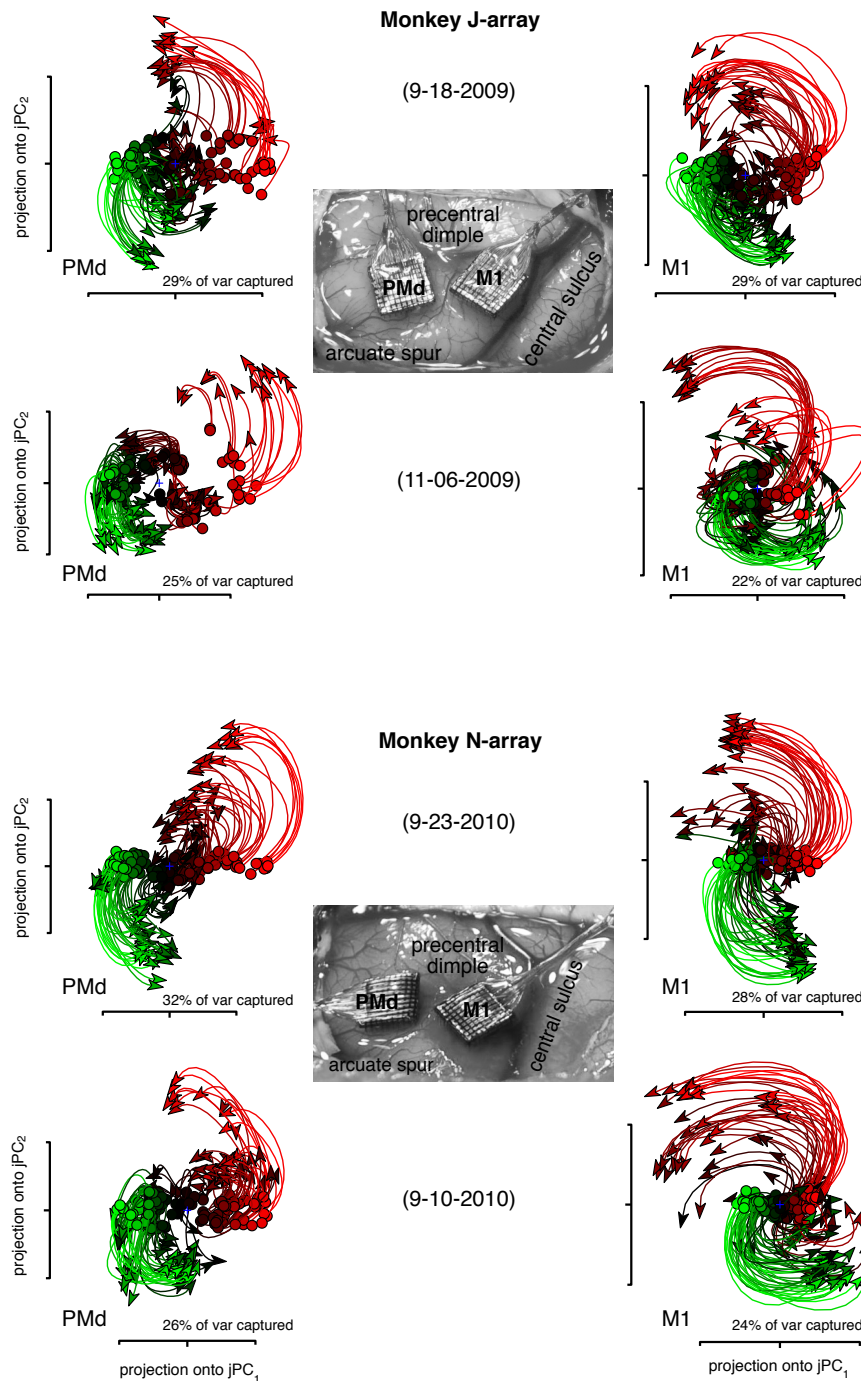
For each control, if strong rotational structure survives the shuffle procedure, then that will be taken as evidence that the jPCA method can erroneously extract such structure even when it is absent or weakly present. This would make the central results in figure 3 impossible to interpret. If strong rotational structure is lost following the shuffle, then that will be taken as evidence that such structure was in fact present in the original data to a much greater degree than expected by chance.



**Supplementary figure 3.** Effect of the shuffle controls illustrated in supplementary figure 2. Each panel plots the jPCA projection of the population response for the J-array dataset (108 reach conditions). The top-left panel plots the projection for the original, un-shuffled data (same as figure 3e of the main text). The other panels plot the projections when jPCA was applied following the three shuffle controls. While many individual trajectories remain curved, the overall robustness of the rotational structure, and the relationship between phase and initial state, is largely lost following shuffling. This indicates that the pattern seen in the top-left panel reflects real underlying structure in the population response, rather than the ability of our method to find such structure by chance. Shuffle controls had similar effects across all datasets.



**Supplementary figure 4.** The data contain multiple planes with rotations. Data are shown for two datasets (monkey B and monkey J-array). Very similar findings were obtained for all datasets. Each column plots the first three jPCA planes (the top six jPCs) found within the top 10 PCs. All three planes contained rotational structure that was coherent (in the same direction and at a similar angular velocity) across conditions. However, rotations were slower (as expected) and less orderly for the higher-numbered jPCA planes. The top three jPCA planes (spanning six dimensions) captured 60% (monkey B) and 54% (monkey J-array) of the total variance in the data. For comparison, the top six PCs (which by definition capture the most variance possible) captured 72% and 64% of the variance. These findings were typical: all datasets contained two, and usually three, planes that captured rotations. Together these planes captured 50-70% of the data variance (range is across datasets). The planes shown were found by applying jPCA to the top 10 PCs, rather than the top 6 as in the main text. This is necessary; it is unlikely that the six dimensions with the strongest rotations would fortuitously align with the top 6 PCs.

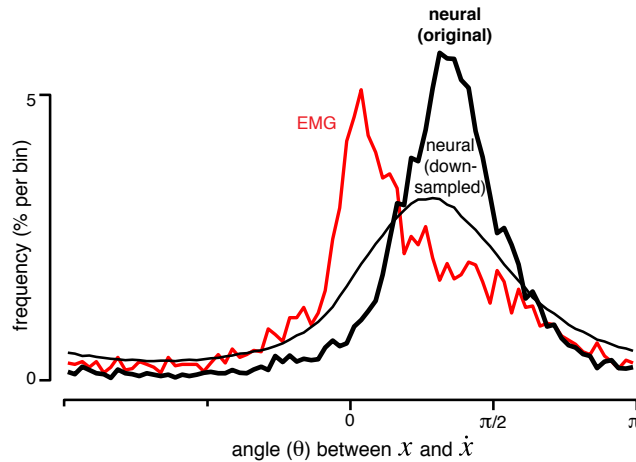


**Supplementary figure 5.** Analysis segregated by area (caudal PMd versus M1). Analysis employed data from two monkeys, each implanted with a pair of electrode arrays. Each array spanned ~4-6 mm of anterior-posterior distance, with ~1-2 mm of separation at the closest point. We refer to the two arrays as the ‘PMd array’ and the ‘M1 array’, though it should be kept in mind that border between these two areas is not sharp. Shown are the J-9-18-2009 and N-9-23-2010 datasets employed in the main text, plus two additional datasets collected on different dates: J-11-06-2009 and N-9-10-2010. The additional datasets employed the same design as those shown in the main text, and were added to give a sense of the consistency of the effects described below.

Both PMd and M1 exhibited clear rotations of the neural state. However, there were two subtle but noticeable differences between the projections for the PMd and M1 arrays. First, for PMd the preparatory state tended to be better separated across conditions. This is consistent with the long-reported fact that PMd tends to exhibit stronger preparatory activity. Second, the angular velocity of the rotations was in each case slightly higher for M1. We suspect that this effect may be real. It is consistent with our informal observations, notable in every dataset we have

inspected so far, that neurons recorded in posterior sites are more likely to exhibit high-frequency response features. Still, we wish to stress that PMd-like neurons are frequently found in M1 and *vice versa*. And of course both arrays exhibited overall patterns that were qualitatively very similar: the neural state rotated away from the initial preparatory state, with a direction and angular velocity that was reasonably consistent across all 108 reach conditions.

Variance captured refers to the proportion of total data variance that lies within the jPCA plane. This was typically 20-35%. Capturing more than this using two dimensions is unlikely, given the high-dimensional nature of the neural data.

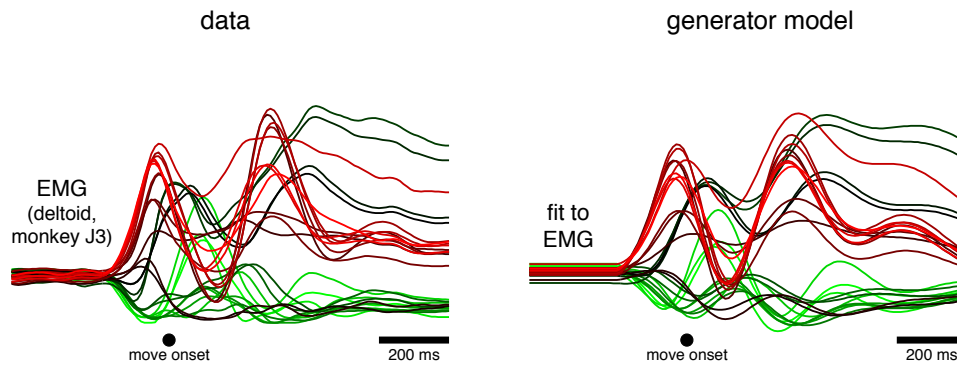


**Supplementary figure 6.** Effect of down-sampling the population of recorded neurons to the same size as the population of recorded muscles. As reported in the main text (fig. 4*f,j* and fig. 6*a,b*) the recorded populations of muscles did not exhibit strong rotations in the jPCA plane. A potential concern is that the population of muscle recordings might have shown weak rotations (relative to the neural population) simply because fewer muscles were recorded than neurons. Across datasets, 6–12 muscles were recorded, compared with 50–218 neural isolations.

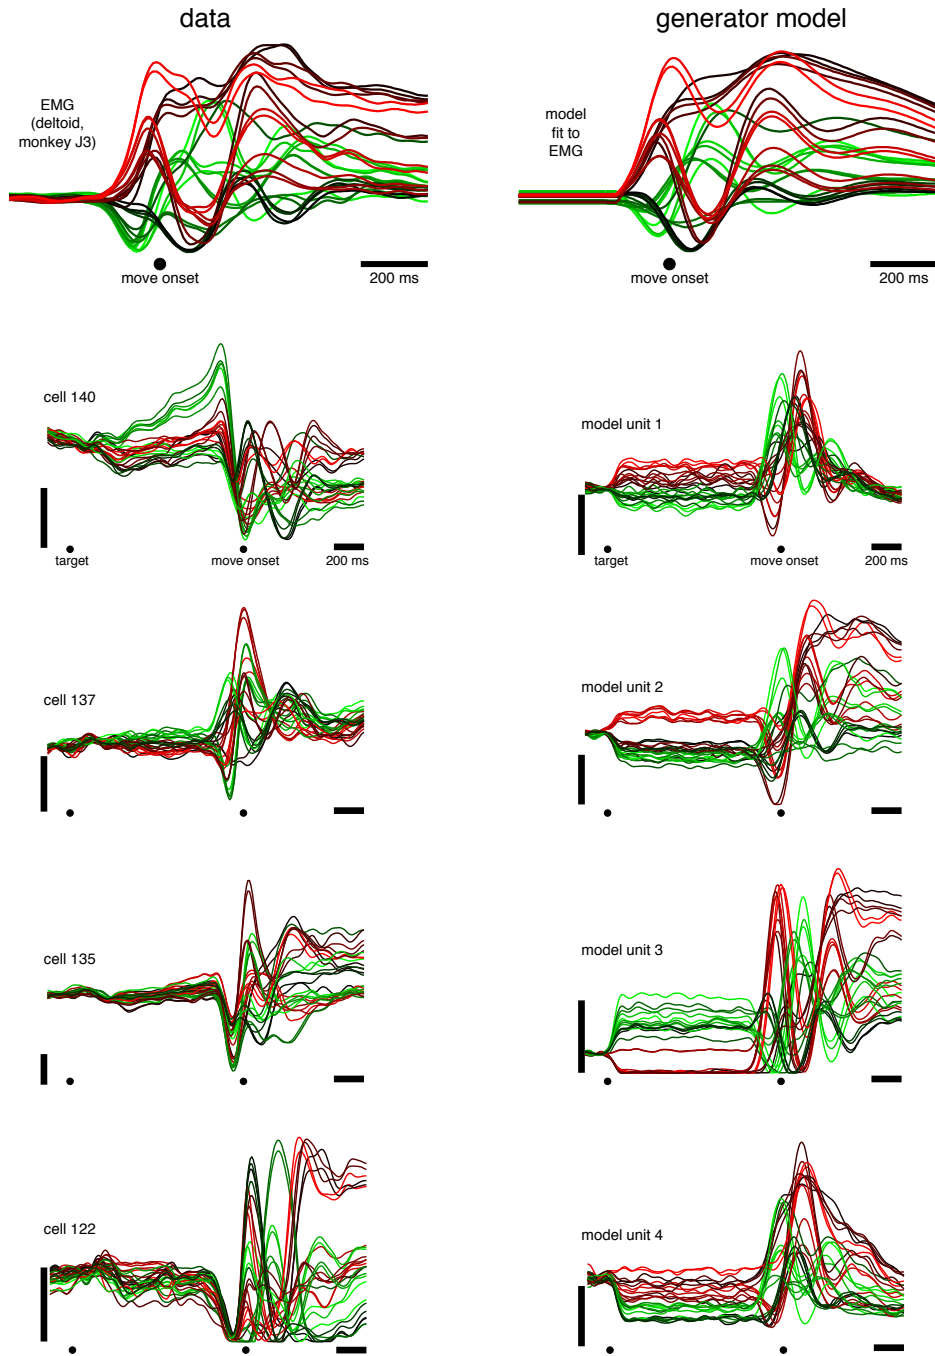
Before describing the control shown here, two points are worth making. First, for all analyses in the main text, the dimensionality of both the muscle and neural populations was reduced to six before applying jPCA. Thus, for both muscles and neurons, the jPCA method attempts to find a plane with rotations within a six-dimensional space. Specifically, it was not the case that jPCA was allowed to ‘search’ within a larger space for the neural data. Second, the central hypothesis of this study holds that the output of cortex reflects a subset of the patterns present in the neural data. This is consistent with the finding that muscle activity is typically found to be low dimensional<sup>43,44</sup>. Thus, it *should* in some sense be true that the muscle population shows weaker rotations because it is lower dimensional. This is certainly true of the generator model, where the single output is constructed from underlying rotations, but does not itself contain any rotations (main text; fig. 5). Nevertheless, it is important to rule out the statistical concern that the higher-dimensional neural space is somehow given an ‘unfair’ advantage.

To control for this possibility, for each dataset we down-sampled the neural population to equal the size of the muscle population. This was repeated 100 times for each dataset. Individual-neuron recordings are often much noisier than individual EMG recordings; we therefore restricted analysis to neurons where the strength of peri-movement activity was better than average. The down-sampled data was then analyzed as in fig. 6 of the main text. The distribution for the down-sampled data is broader and shifted slightly to the left of the distribution for the original data. This is expected: the increase in sampling error will broaden the distribution. Furthermore, rotations will be weakened on those draws where the sample of neurons does not contain both relevant phases to roughly equal degrees. Nevertheless, the distribution for the down-sampled neural data peaks well to the right of the distribution for the muscle recordings.

In summary, populations of neurons contain considerably stronger rotations of their state than do populations of muscles, even when the two populations are matched in size.

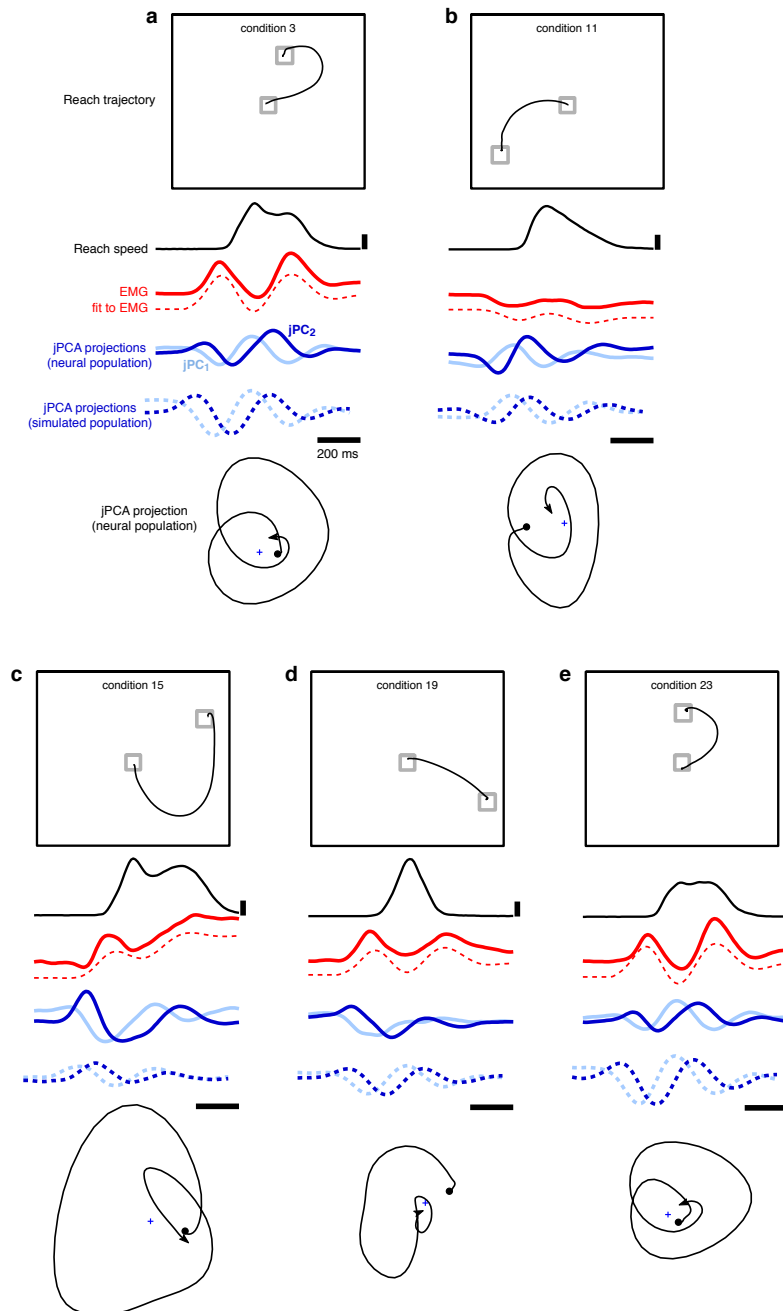


**Supplementary figure 7.** Fits of the generator model (right) to deltoid EMG (left). Fits are the same as those in figure 5*b,c* of the main text, but are shown for all 27 conditions (dataset J3). Color-coding is based on the initial strength of deltoid EMG, just before movement onset.



**Supplementary figure 8.** Further examples (from dataset J2) of how the generator model fits EMG (*top*). Also shown are example responses of both real (*left column*) and simulated units from the generator model (*right column*). 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. Vertical scale bars indicate 20 spikes/s. 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 for the model, neural responses directly reflect the underlying rotational patterns used to generate EMG. Nevertheless the responses of most model neurons do not match the EMG profile.





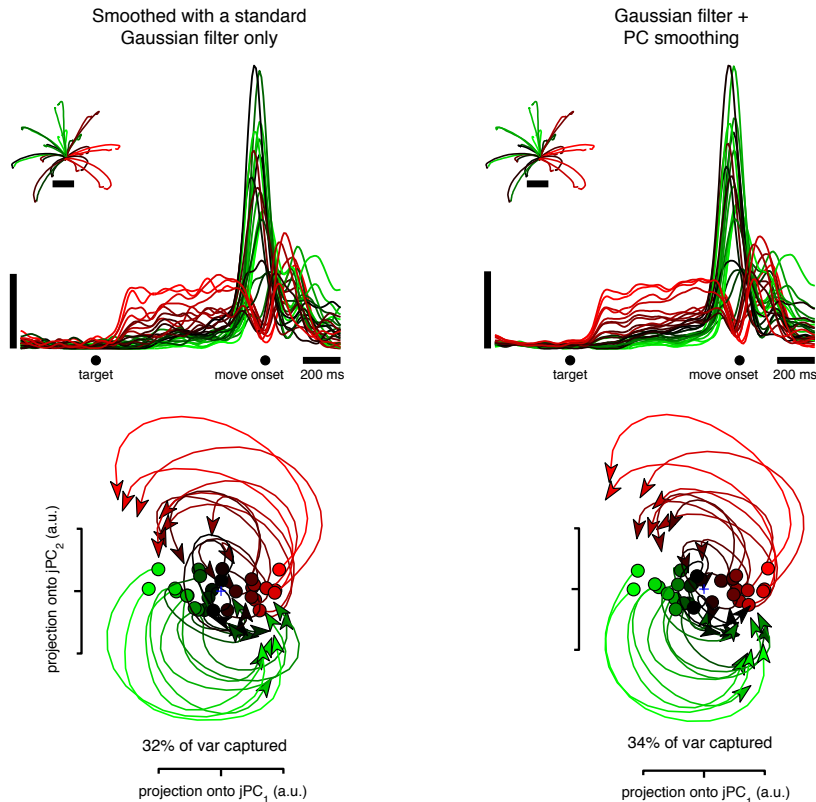
**Supplementary figure 9.** Example data for individual conditions, allowing one to compare kinematics, EMG, jPCA projections of neural data, and jPCA projections of generator model ‘data’ (monkey J3 dataset). The goal is not to make direct quantitative comparisons, but to inform intuition regarding how oscillatory patterns are expected to change across conditions. Data are from monkey J3. Each panel (*a-e*) plots data for one condition. Shown are hand trajectory (*black, top*), reach speed (*black*), deltoid EMG (*red*), EMG fit produced by the generator model (*dashed red*), projections onto the jPCs versus time for real (*blue*) and generator model data (*dashed blue*), and the jPCA plane for the neural data (*black, bottom; jPC1 and jPC2 are the horizontal and vertical axes*). The jPCA projections for the generator model have been advanced in time by 50 ms to emulate a lag between neural and EMG signals.

A comparison of kinematics, EMG, and neural data makes two points. First, there is no straightforward relationship between reach duration and oscillation duration. For example, compare panel *d* with panel *a*. Although reach durations differ, the multiphasic muscle and neural patterns have a similar duration in both panels. Second, there is no obvious relationship between reach duration and oscillation frequency. These two observations were true by design for the generator model, but were not guaranteed to be true for the neural data or for the EMG.

Comparing EMG with the generator model illustrates why the generator model is able to provide

such good fits. The EMG contains multiphasic features that differ, across conditions, in their amplitude and phase. Yet those multiphasic features are reasonably consistent in their frequency. The generator model is thus able to contribute those features using a short-lived oscillatory pattern whose phase and amplitude differs across conditions. The lower-frequency features are contributed by the second (and slower) of the two oscillations provided by the model (not shown).

These data are inadequate to address the possibility that neural activity directly drives muscle activity. Indeed, one presumes that reasonably strong dynamics are contributed by the spinal cord. Yet the examples shown illustrate that there is nothing paradoxical regarding some of the salient features of the neural state-space rotations. Such rotations have a relationship with kinematics that is sometimes counter-intuitive, but they form a natural basis for driving the muscle activity that results in those kinematics.



**Supplementary figure 10.** Effect of the pre-processing step where the PSTH of each neuron was smoothed using the ‘PC smoothing’ method<sup>45</sup>. This step acts only to remove small amounts of noise from the PSTHs of individual neurons. PC smoothing is done at a single-neuron level and has no relationship to the use of PCA or jPCA at the population level.

*Top panels:* response of one neuron (B16). The data in the left column were smoothed using a 20 ms Gaussian kernel in the traditional fashion. The data in the right column were further smoothed using the PC smoothing method.

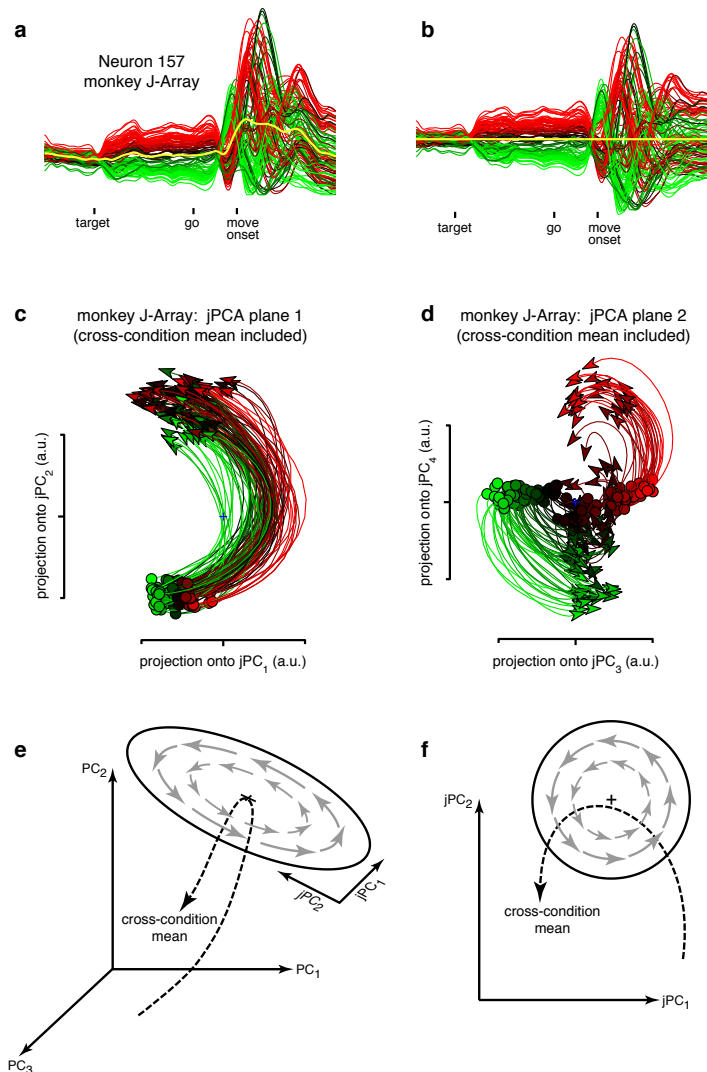
*Bottom panels:* jPCA projections of the population for unsmoothed and smoothed data. The right panel is identical to that in fig. 3 of the main text. PC smoothing has almost no impact at the population level, for reasons described below.

PC smoothing was used to pre-process all data presented in the main text. The details of this method are provided below. There

are two key high-level points. First, PC smoothing reduces the noise present in the PSTH of a given neuron. This provides a better visual estimate of that neuron’s response. Second, PC smoothing has essentially no impact on the jPCA projections. The removal of small amounts of individual-neuron noise is largely irrelevant at the population level. The jPCA projections are thus virtually identical regardless of whether smoothing was applied. This was true of all datasets, including the EMG recordings.

Although it PC smoothing had essentially no impact on the central results, for completeness we describe below the motivation and methodology. First note that a temporal filter exploits the fact that different times cannot be arbitrarily different from one another. In contrast, PC smoothing exploits the fact that different conditions cannot be arbitrarily different from one another. PC smoothing exploits this fact by using PCA at the level of an *individual neuron* (note that this is very different from the more usual approach of using PCA at the population level). To apply PC smoothing, we compiled a *cxt* data matrix, where each row contained the response of that neuron for one condition across all times. We then decomposed this matrix into its principal components (PCs), and reconstructed the data using the first six principal components. This procedure preferentially discards small high-frequency events that are unique to one of the 27 conditions (and thus likely to result from sampling error). Something similar could of course be obtained by using a broader temporal filter, but that could come at the cost of losing real high-frequency aspects of the response. An advantage of PC smoothing is that it does not remove high-frequency aspects of the response if they are shared among a number of conditions.

To reiterate a technical point: this use of PCA to smooth individual-neuron PSTHs contrasts with the more typical use of PCA at the population level. In the current study we use PCA twice: once for PC smoothing as described above (a minor component of the overall analysis) and a second time at the population level during the computation of the jPCA plane (a key component of our analyses).



**Supplementary figure 11.** Effects of the pre-processing step where each neuron's response was centered by removing the cross-condition mean. The goal of this pre-processing step was to focus all further analyses on dimensions where activity differed strongly across conditions. The panels below show how the cross-condition mean was removed (*top row*), the consequences of not subtracting the cross-condition mean (*middle row*), and possible relationships between the cross-condition mean and the rotational patterns (*bottom row*). **a**. Firing rate as a function of time for one example neuron (monkey J-array dataset; 108 conditions). Each green/red trace plots the average firing rate for one condition, shaded based on the level of preparatory activity. The yellow trace plots the cross-condition mean (the mean of all the other traces). Before subsequent analysis (the application of PCA and jPCA) this cross-condition mean was subtracted. This was done independently for each neuron. **b**. Response of the same example neuron after mean subtraction. The cross-condition mean (yellow) is now zero at all times. **c**. Application of jPCA to the monkey J-array dataset, with *no subtraction of the cross-condition mean* (i.e., the pre-processing step in *a, b* was *not* applied). This panel can be contrasted with that in figure 3*d* of the main text (for which the cross-condition mean was subtracted). When the cross-condition mean is not subtracted, the projection onto the first jPCA plane captures response patterns that vary only weakly across conditions. This is

not surprising: many neurons display strong response features that are similar across conditions (the well-known ‘non-directional’ component of the response<sup>46</sup>). Thus, the top PCs (via standard PCA) often contain large features that are very similar across conditions. The jPCs are simply a projection of what is already contained in the PCs. It was therefore common (as in this example) for the first jPCA plane to contain patterns that were largely condition-independent. Such patterns are consistent with our general hypothesis (they still involve rotations) but are difficult to interpret because there are many potentially trivial explanations for the observation of curved trajectories that are similar across all conditions. **d**. The second jPCA plane (jPC<sub>4</sub> versus jPC<sub>3</sub>) for the same analysis as in *c* (again, the cross-condition mean was *not* removed). This pattern is similar to that seen in the original analysis in figure 3*d* (and exists in a plane orthogonal to that in panel *c* of this figure). Thus, for this dataset, the first jPCA plane largely captures the condition-independent response, and the second plane largely captures the condition-dependent response, and is thus very similar to that in the original analysis. However, across the 8 datasets it was not uncommon for both planes to capture a mixture of condition-independent and condition-dependent components. Thus, to allow the first jPC plane to always be the primary focus of analysis, and to allow that plane to capture structure across conditions, all jPCA analyses in the main text were performed after subtracting the cross-condition mean. **e, f**. Schematic illustration of possible trajectories for the condition-independent aspect of the response (the cross-condition mean) relative to the pattern of rotations. Those two aspects might be orthogonal (as in *e*) or might lie within the same plane (as in *f*). In both cases, rotations will be best isolated if the cross-condition mean is first removed. These schematics also illustrate a possible role for the condition-independent component of the response.

(As a technical point, the analyses in this figure consider the top 10 PCs, to allow sufficient dimensions to capture both condition-independent and condition-dependent features of the response).

### Supplementary Derivation

jPCA is a dynamical variant of PCA that finds planes of significant rotational structure within data. We consider high dimensional time series data  $\mathbf{x}(t) = [x_1(t), \dots, x_n(t)]$ . We can represent this data as a matrix  $X \in \mathbb{R}^{ct \times n}$ , where  $ct$  is the number of time points in the time series across all conditions, and  $n$  is the dimensionality of the data at any time point. This is as in the main text/methods, except here it is convenient to orient  $X$  as  $ct \times n$  rather than  $n \times ct$ . Also note that, in practice, jPCA was always applied *after* traditional PCA to  $X_{red} \in \mathbb{R}^{ct \times k}$  for  $k = 6$  or similar. For the purposes of understanding jPCA at a mathematical level, that preprocessing is a distraction (and has no bearing on the validity of the algorithm itself), so we ignore it here. Hence, hereafter we consider data of the form  $X \in \mathbb{R}^{ct \times n}$ .

Traditional PCA begins by calculating the data covariance  $\Sigma = X^T X$  (assuming  $X$  is mean-centered). However, since we are particularly interested in dynamical structure, we seek a different  $n \times n$  matrix summarizing the data. The simplest dynamical system that we can fit to that data is a time-invariant linear dynamical system, which has the form  $\dot{\mathbf{x}}(t) = \mathbf{x}(t)M$  for any matrix  $M \in \mathbb{R}^{n \times n}$ . Solving for such an  $M$ , given our data  $X$ , reduces to a simple least squares problem. We can write  $\dot{X} = XM$ , and then the least squares solution solves the problem  $M^* = \operatorname{argmin}_{M \in \mathbb{R}^{n \times n}} \|\dot{X} - XM\|_F$  (where the subscript  $F$  denotes the Frobenius norm). This is solved in simple closed form as  $M^* = (X^T X)^{-1} X^T \dot{X}$  (often written as  $M^* = X \setminus \dot{X}$ ). Importantly, this dynamics matrix  $M^*$  is a valid summary matrix: whereas the data covariance  $\Sigma = X^T X$  describes the ellipsoid that best fits the data (without regard to temporal information),  $M^*$  describes the linear dynamical system that best fits the data  $X$ .

General linear dynamical systems describe both expansions/contractions and rotations, but they make no distinction between these aspects of the data. In this work we seek to investigate the role that rotations play in the dynamics of motor cortical neurons, and hence we are specifically interested in *rotational* linear dynamical systems. Every linear transformation  $M$  is some mixture of a symmetric matrix and a skew-symmetric matrix, which we write  $M = M_{symm} + M_{skew}$ . The symmetric part is, by analogy to functions, the “even” part of the matrix, and is defined as  $M_{symm} = (M + M^T)/2$ . Such matrices satisfy  $M_{symm} = M_{symm}^T$ . Accordingly, the skew-symmetric matrix (odd part of the matrix, also called anti-symmetric) is  $M_{skew} = (M - M^T)/2$ . Such matrices satisfy  $M_{skew} = -M_{skew}^T$ . Adding  $M_{symm}$  and  $M_{skew}$  returns the matrix  $M$ , so indeed these are a general description of any square matrix  $M$ . Additionally, the symmetric component  $M_{symm}$  has purely real eigenvalues and thus describes only expansions and contractions of data. So too, the skew-symmetric component has purely imaginary eigenvalues (in complex conjugate pairs) and describes rotations in the data.

Since we are interested only in rotational linear dynamical systems, we must solve the original least squares problem ( $\dot{X} = XM$ ) not over all matrices  $M$  (that describe any linear dynamical system), but instead we should solve the original problem constrained only to the set of rotational linear systems. We call this set of skew-symmetric matrices  $\mathcal{S}^{n \times n}$ , and we are interested in solutions to the original least squares problem of the form  $M_{skew} \in \mathcal{S}^{n \times n}$ , namely  $M^* = \operatorname{argmin}_{M \in \mathcal{S}^{n \times n}} \|\dot{X} - XM\|_F$ .

To solve this constrained optimization, we use an equivalent but slightly more cumbersome notation. We begin by noting that in the original unconstrained least squares solution  $X \setminus \dot{X}$ , the columns of  $M$  are solved *independently* of each other. If we write  $M^* = X \setminus \dot{X}$  in the expanded least squares form

$M^* = (X^T X)^{-1} X^T \dot{X}$ , we see that each column of  $\dot{X}$  determines the corresponding column of  $M$  independently of other columns. Thus, we can rewrite the original problem as a vector problem instead of a matrix problem. We introduce the vector  $\mathbf{m} \in \mathbb{R}^{n^2}$ , which is simply the matrix  $M \in \mathbb{R}^{n \times n}$  unrolled to a vector (we denote this using the common notation  $\mathbf{m} = M(:)$ ). We can then rewrite the unconstrained least squares problem as  $\mathbf{m}^* = \operatorname{argmin}_{\mathbf{m} \in \mathbb{R}^{n^2}} \|\dot{\mathbf{x}} - \tilde{X} \mathbf{m}\|_2$ , where  $\dot{\mathbf{x}} = \dot{X}(:)$ , and  $\tilde{X}$  is a block diagonal matrix with the matrix  $X$  repeated on the  $n$  diagonal blocks. These two forms of the least squares problem give identical solutions (it is just a formatting choice: do we write the solution as a vector  $\mathbf{m}$  or a matrix  $M$ ).

Now we rewrite the rotational optimization problem in this vector notation, as that will allow us to naturally incorporate the constraint that  $M^*$  be a member of  $\mathcal{S}^{n \times n}$ . Whereas previously we sought to solve the unconstrained  $M^* = \operatorname{argmin}_{M \in \mathbb{R}^{n \times n}} \|\dot{X} - XM\|_F$ , here we want to solve the constrained  $M^* = \operatorname{argmin}_{M \in \mathcal{S}^{n \times n}} \|\dot{X} - XM\|_F$ . For this set of skew-symmetric matrices  $\mathcal{S}^{n \times n}$ , by virtue of the constraint  $M_{skew} = -M_{skew}^T$ , these matrices only have  $n(n-1)/2$  free parameters (not the  $n^2$  of general matrices  $\mathbb{R}^{n \times n}$ ).

The key step is to note that we can represent skew-symmetric matrices as a vector of  $n(n-1)/2$  free parameters: we call these vectors  $\mathbf{k} \in \mathbb{R}^{n(n-1)/2}$ . Further, we can specify a linear map from these vectors onto the space of our vectors  $\mathbf{m} \in \mathbb{R}^{n \times n}$ . This matrix, which we call  $H$ , simply places each of the elements of  $\mathbf{k}$  into two indices in  $\mathbf{m}$ . We can consider  $\mathbf{k}$  to be the lower triangle of a skew-symmetric matrix. Suppose we have an element of  $\mathbf{k}$ ,  $k_r$ , that corresponds to the  $(i, j)$  element of the matrix for  $i > j$  (the lower triangle).  $H$  then takes this element  $k_r$  and places: (1) the element  $k_r$  in the element of  $\mathbf{m}$  corresponding to the  $(i, j)$  index of a matrix  $M$ , and (2) the negated element  $-k_r$  in the element of  $\mathbf{m}$  corresponding to the  $(j, i)$  index of  $M$ . By this construction we define a matrix  $H$  that maps from  $\mathbb{R}^{n(n-1)/2}$  (the number of free parameters in a skew-symmetric matrix) to  $\mathbb{R}^{n^2}$ .

Finally, we return to our problem of interest. Since, by our construction, the quantity  $H\mathbf{k}$  is a vector in  $\mathbb{R}^{n^2}$  and is equivalent to the set of all skew-symmetric matrices  $\mathcal{S}^{n \times n}$ , our original problem has the equivalence:

$$M^* = \operatorname{argmin}_{M \in \mathcal{S}^{n \times n}} \|\dot{X} - XM\|_F \iff \mathbf{k}^* = \operatorname{argmin}_{\mathbf{k} \in \mathbb{R}^{\frac{1}{2}n(n-1)}} \|\dot{\mathbf{x}} - \tilde{X} H \mathbf{k}\|_2.$$

This form can then simply be solved in closed form by grouping  $H$  with  $\tilde{X}$  and solving  $\mathbf{k}^* = (\tilde{X} H) \setminus \dot{\mathbf{x}}$ . Since  $\mathbf{k}^*$  is a vector of length  $n(n-1)/2$  that defines a skew-symmetric matrix  $M_{skew}^*$  (via  $H$ ), and since skew-symmetric matrices describe rotational dynamics, this solution is the matrix defining the best-fitting rotational linear dynamical system.

As an implementation note, one might prefer to write the Lagrangian for the constrained optimization problem. Doing so will result in, depending on programming choices, a very similar implementation as above. In any event, this optimization problem has a unique global optimum, so any valid derivation will produce exactly the same result. As a second note, representing the large matrices  $H$  and  $\tilde{X}$  can be computationally burdensome. Instead, we can avoid the explicit creation of these matrices by solving  $\mathbf{k}^* = (\tilde{X} H) \setminus \dot{\mathbf{x}}$  with a gradient based-method, which allows us to manipulate vectors and matrices in the most computationally convenient way. The result is an extremely fast and accurate method for solving

the skew-symmetric least squares problem (only slightly slower than regular unconstrained least squares).

We now have successfully calculated  $M_{skew}^*$ , the summary matrix that describes rotational dynamics in the data. Traditional PCA takes the summary matrix (data covariance) and does an eigenvalue decomposition, which creates a ranked set of orthonormal vectors that we can use to project our data. Decomposition of a skew-symmetric matrix also produces a ranked set of orthonormal vectors. The eigenvectors produced by the decomposition of  $M_{skew}$  come in orthogonal, complex conjugate pairs (these vectors form a unitary matrix), and thus we can naturally think of the decomposition of  $M_{skew}$  as producing orthogonal planes (the eigenvalues are purely imaginary). To find a plane - any pair of real-valued projection vectors  $\{\mathbf{u}_{i,1}, \mathbf{u}_{i,2}\}$  - each complex conjugate pair of vectors  $\{\mathbf{v}_{i,1}, \mathbf{v}_{i,2}\}$  is combined as  $\mathbf{u}_{i,1} = \mathbf{v}_{i,1} + \mathbf{v}_{i,2}$  and  $\mathbf{u}_{i,2} = j(\mathbf{v}_{i,1} - \mathbf{v}_{i,2})$ , which are then suitably normalized. These  $\mathbf{u}$  vectors can be used exactly as in traditional PCA to project the high dimensional data to lower dimensionality (e.g., by projecting onto just the first plane, after ordering from largest to smallest pair of imaginary eigenvalues).

The magnitude of the eigenvalues allow us to select the plane (or planes) with the highest frequency and consistency in the *rotational* linear dynamical system. Thus, projecting the data down onto these planes allows us to visualize planes in the data with significant rotations. By connection to the imaginary eigenvalues of  $M_{skew}$ , we call this algorithm jPCA. To be concrete, all figures showing projections (except those noted otherwise) show the top jPC plane, which is the plane of strongest activity within the rotational dynamical system that best fits the data.

It is important to revisit the fact that jPCA and PCA solve different objective functions and thus produce different results: while PCA finds directions of maximal variance, jPCA finds directions (planes) of significant rotational dynamics. Indeed, the plane with the strongest rotations could in principle capture very little data variance. Thus, to prevent finding significant rotations that capture little variance, jPCA was always applied after first using PCA to find the handful of dimensions ( $k = 6$  typically) with the most variance. In this application, jPCA simply rotates the dimensions found by PCA to better reveal dynamical structure (see supplementary movie 2).

As a technical note, we are taking the plane of strongest rotations within the rotational dynamical system that best fits the data. As an alternative approach, one might try to directly fit the plane of strongest rotations from the data, which would in principle do at least as well as the first jPCA plane. Thus our jPC projections (cf. Figure 3 and 4) are theoretically conservative, though our empirical investigation indicates this is an insignificant effect. These two approaches are equivalent when the data is white ( $X^T X = I$ ), and we note that whitening the data also produces very little change to the resulting projections.

As a final point, the jPCA algorithm can be readily modified to a symmetric version 'symmPCA' to focus on directions of largest expansion and contraction. The convenience of being able to eigendecompose a summary matrix and yield orthogonal vectors belongs to the class of normal matrices, which by definition are diagonalizable by a unitary matrix. The class of normal matrices includes symmetric and skew-symmetric matrices, among others. This fact suggests a broader class of PCA variants that are a subject of future work.

### **Supplementary Movies**

**Supplementary movie 1.** Neural trajectory in the walking monkey (2× real time). The movie begins with the monkey stationary. After ~8 seconds (16 seconds of real time), the monkey begins walking. After ~25 cycles, the monkey pauses, then begins walking again. The monkey ceases walking once more just before the movie ends. Clear rotations are seen only during epochs of walking (QuickTime 376 KB).

**Supplementary movie 2.** Illustration of how the PCA axes were rotated to find the jPCA projection. The movie contains clips for three datasets: monkey B, monkey J-array, and monkey N-array. For each clip the opening frame plots the PCA-based projection (PC<sub>2</sub> versus PC<sub>1</sub>). On each subsequent frame the projection is rotated until the jPCA projection is reached. Compare with figure 3*a,e,f* of the main text (QuickTime 3.1 MB).

**Supplementary movie 3.** The jPCA projections as a function of time. The movie contains clips for four datasets: monkey B, monkey J3, monkey N, and monkey N-array. Time is 1/8<sup>th</sup> real time and starts just as preparatory activity is giving way to movement-epoch activity. The projections differ slightly from those in figure 3 of the main text for two reasons. First, they are based on more time. Second, to allow jPCA to best isolate the plane that worked well across a broad range of times, analysis was based on the top 10 PCs. (The true dimensionality of the data is higher still, but we wished to remain modestly conservative). The shuffled controls in *supplementary movie 4*, and the analysis of EMG in *supplementary movie 5* similarly employ the greater range of times and projections based on the top 10 PCs (QuickTime 2.4 MB).

**Supplementary movie 4.** Just as for *supplementary movie 3* but after applying the shuffle control. The movie contains clips for the same four datasets: monkey B, monkey J3, monkey N, and monkey N-array. The shuffle that was applied was version #1 (see *supplementary figure 2* and 3). (QuickTime 1.8 MB).

**Supplementary movie 5.** The jPCA projections as a function of time for populations of muscle recordings. The movie contains clips for three datasets: monkey A, monkey J3, and monkey N. (EMG data were not recorded for the full 108-condition task used for the array-based datasets). (QuickTime 1.2 MB)

### **Supplementary References**

- 43 Bizzi, E., Cheung, V. C., d'Avella, A., Saltiel, P. & Tresch, M. Combining modules for movement. *Brain Res Rev* **57**, 125-133 (2008).
- 44 Tresch, M. C. & Jarc, A. The case for and against muscle synergies. *Curr Opin Neurobiol* **19**, 601-607 (2009).
- 45 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine? *Neuron* **68**, 387-400 (2010).
- 46 Moran, D. W. & Schwartz, A. B. Motor cortical representation of speed and direction during reaching. *J Neurophysiol* **82**, 2676-2692 (1999).