A curved manifold orients rotational dynamics in motor cortex

Authors: David A. Sabatini^{1,2} and Matthew T. Kaufman^{1,2,*}

¹ Department of Organismal Biology and Anatomy

² Neuroscience Institute, The University of Chicago, Chicago, IL 60637, USA

* Corresponding author, mattkaufman@uchicago.edu

Author contributions: Both authors contributed to all parts of conceptualization, interpretation, writing, and editing. DS performed the analyses with assistance from MK.

Competing interests: The authors declare no competing interests.

Acknowledgements:

This work was supported by funding from The University of Chicago, the Neuroscience Institute at The University of Chicago, and the Alfred P. Sloan Foundation. We thank C. Pandarinath for suggesting controls, and S. Bensmaia and J. Maclean for helpful feedback on the manuscript.

Summary:

Controlling arm movements requires complex, time-varying patterns of muscle activity ^{1,2}. Accordingly, the responses of neurons in motor cortex are complex, time-varying, and heterogeneous during reaching ²⁻⁴. When examined at the population level, patterns of neural activity evolve over time according to dynamical rules ^{5,6}. During reaching, these rules have been argued to be “rotational” ⁷ or variants thereof ^{8,9}, containing coordinated oscillations in the spike rates of individual neurons. While these models capture key aspects of the neural responses, they fail to capture others – accounting for only 20-50% of the neural response variance. Here, we consider a broader class of dynamical models. We find that motor cortex dynamics take an unexpected form: there were 3-4 rotations at fixed frequencies in M1 and PMd explaining more than 90% of neural responses, but these rotations occurred in different portions of state space when movements differ. These rotations appear to reflect a curved manifold of fixed points in state space, around which dynamics are locally rotational. These fixed-frequency rotations obeyed a simple relationship with movement: the orientation of rotations in motor cortex activity were related almost linearly to the movement the animal made, allowing linear decoding of reach kinematic time-courses on single trials. This model constitutes a fundamentally novel way to consider pattern generation: like placing a record player in a large bowl, the frequency of activity is fixed, but the location of motor cortex activity on a curved manifold sets the orientation of locally-rotational dynamics. This system simplifies motor control, helps reconcile conflicting frameworks for interpreting motor cortex, and enables greatly improved neural decoding.

Introduction

Though single neuron activity is complex during motor control, dynamical systems analysis has revealed that population activity obeys relatively simple rules: activity of the neural population determines population activity moments later, consistent with motor cortex acting as a generator^{5,6,10–14}. This approach has been used to argue that motor cortex dynamics are approximately rotational^{7,15–17} or similar^{8,18}; preparatory activity has been shown to set the initial state for future movement-epoch dynamics¹⁹ in dimensions orthogonal to outputs^{20–23}; a large trigger signal, identical across reach types, coincides with the onset of movement-epoch dynamics and may act as a ‘trigger’ for dynamics^{23–26}; and these dynamics produce command signals in special, output-potent dimensions that constitute the muscle readouts of neural activity²⁰.

While this framework has brought a great deal of clarity, it only fits a fraction of the data. In particular, rotational dynamics and similar models fit less than half the movement-specific population variance, which is itself only half the total population variance. Here, we propose a novel feature of motor cortex dynamics: while rotations in motor cortex activity contain consistent frequencies, the location and orientation of rotations in state space vary for different movements, making motor cortex dynamics high-dimensional. We find this variation reflects a curved manifold of fixed points in state space, around which dynamics are locally-rotational, and that this variation smoothly reflects reach parameters, allowing high-fidelity linear decoding of reach kinematics even on single-trials.

Results

Data in this paper were collected for previous studies^{7,20}. Two monkeys, J and N, performed a “maze” variant of a delayed-reach task that evoked various curved and straight reaches (“conditions”; Fig. 1a,b). Simultaneous single- and multi-unit electrophysiological recordings were made from a pair of Utah 96-electrode arrays, one implanted in dorsal premotor (PMd) and one in primary motor cortex (M1). Analyses below used trial-averaged and lightly-smoothed neural data aligned to movement onset unless stated otherwise (Methods).

Figure 1: Units in motor cortex display complex activity during reaching. **a**, Illustration of task. **b**, Straight (*left*) and curved (*right*) reaches performed by monkey J. **c**, Rotations in neural activity uncovered in data from M1-J by jPCA. **d**, *Left*, example trial-averaged activity; *right*, reconstruction by jPCA for unit 61 in M1-N. Traces are conditions. **e**, Distribution of variance explained by jPCA for all units. V.E., variance explained. **f**, *Left*, example trial-averaged activity; *right*, reconstruction by jPCA for unit 75 in M1-N. **g-i**, Models of neural activity where **g**, both the rotational frequency and plane are conserved across conditions; **h**, rotational plane, but not rotational frequency, are conserved across conditions; **i**, rotational frequency, but not rotational plane, are conserved across conditions. n1, neuron 1; dim1, projection dimension 1. In **b**, **c**, **d**, **f**, colors are assigned according to reach angle.

Frequency, but not orientation, of neural dynamics is conserved

Fitting linear dynamical systems can uncover rotational dynamics in motor cortex activity⁷ (Fig. 1c; Extended Data Fig. 1a). These rotational dynamics correspond to ~0.5-4 Hz oscillations in the firing rates of individual neurons, with some key additional constraints described below. Rotational dynamics, together with a condition-invariant signal, fit the activity of some individual M1 and PMd units well (Fig. 1d), but on average explain only 22-49% of the variance in individual unit's firing rates (s.d. < 24% for all datasets; Fig. 1e-f; Extended Data Fig. 1b-e). To determine why rotational dynamics fit only a portion of motor cortex activity, we examined a broader class of dynamical models. Dynamical models can be thought of as specifying two assumptions about the population structure of neural activity. The first assumption specifies the types of trajectories neural activity can follow in state space, i.e., the components that are added together to make up each neuron's time-varying response. The second assumption specifies the location of these trajectories in state space, i.e., how each of these components covaries across neurons. Rotational dynamics assumes that across conditions, neural firing rates will always rotate in state space with the same set of frequencies, and that each distinct-frequency rotation will occur in a consistent plane across conditions (Fig. 1g). At the level of single neurons, rotational dynamics restricts neuron-pairs' oscillations to be locked in phase and amplitude. In a closely-related alternate model, neural activity rotates with different frequencies for different conditions, but remains restricted to the same planes in state space⁹. That model allows firing rates to be modulated at different frequencies across conditions, but keeps neuron-pairs locked in phase and amplitude (Fig. 1h). Alternatively, rotations might maintain the same frequencies but occur in different planes on different conditions, such that the phase and amplitude of oscillations could vary across conditions for the same neuron-pair (Fig. 1i). Finally, activity might form more complex trajectories than rotations that wander through a higher-dimensional space on single conditions²⁷.

To distinguish between these models, we relaxed some constraints by fitting a low-dimensional linear dynamical system to each condition independently (Methods). This model assumes only that motor cortex dynamics are rotational on each condition, but does not constrain the frequency of these rotations or their location in state space to be similar across conditions. We found that condition-specific dynamics on average explained 80-91% (s.d. < 9%) of the population variance, significantly greater than the variance explained by rotational dynamics ($p < 0.01$, Wilcoxon Signed Rank test). To determine the source of this improvement, we first examined the rotational frequencies fit by the dynamics on each condition. Across conditions, we found that rotational frequencies were broadly similar, with frequencies of approximately 0.5, 1.5, and 2.5 Hz appearing in both monkeys in PMd and M1 in all conditions, and with M1 containing an additional rotation at 4 Hz (Fig. 2a; Extended Data Fig. 2a,b). We found that reach kinematics could not be decoded from rotational frequencies (Extended Data Fig. 2c). This indicated that the first assumption of a rotational dynamics model was correct: in each condition, neural activity contained oscillations in the firing rates with approximately the same frequencies on all conditions. Second, we examined the planes in which rotations occurred on each condition, as fit by the condition-specific dynamics. For each condition-pair, we quantified how similar the two planes were for a given rotation frequency using the alignment index^{21,28}. The average alignment index for condition-pairs was 48-74% (s.d. < 16%). This value was less aligned than what would be expected from noisy realizations of a single, low-dimensional subspace (75-84%, Fig. 2b, $p < 0.01$ for all datasets, Wilcoxon Rank Sum test; Extended Data Fig. 2d). This finding indicates that the second assumption of rotational dynamics was incorrect: while different conditions contain firing rate oscillations at consistent frequencies, at the population level neural activity rotates in different planes for different conditions.

Figure 2: Frequency, but not orientation, of rotational dynamics in motor cortex is conserved. **a**, Rotational frequencies found by condition-specific dynamics for M1 and PMd. **b**, Distribution of alignment indices across condition-pairs and the corresponding PCA model for M1 and PMd. **c**, Illustration of factorizing neural activity on a single condition into a condition-specific loading matrix and shared temporal basis functions. **d**, Geometric interpretation of **c**, where shared dynamics are embedded in condition-specific subspaces to produce neural activity. **e**, 2.5 and 4 Hz temporal basis functions for M1-N, plotted as dimensions over time and dimensions plotted against one another. **f**, Distribution of population variance explained in neural activity and corresponding rate-matched homogeneous Poisson processes. **g**, *Left*, example trial-averaged activity; *right*, reconstruction with temporal basis functions. Traces are conditions, colored according to reach angle. **h**, Example of single-condition estimated temporal basis functions from neural activity in M1-N. Shaded regions represent 1 standard deviation across conditions. **i**, Variance explained in the temporal basis functions from the neural activity before and after time-warping.

Conserved temporal patterns in neural activity

This finding argues that the response of any M1 or PMd neuron, for any reaching condition, is the sum of a few conserved oscillations, but that the phases and amplitudes of these oscillations are inconsistent across neurons for different conditions. To capture these relationships, we factorized neural activity into two pieces: a small set of temporal basis functions that are shared across all units and all conditions, and a loadings matrix specific to each condition, which specifies the phase and amplitude of each oscillation for each unit (Fig. 2c; see Methods for details of how this differs from more common approaches). This factorization is mathematically equivalent to fitting a model where rotational frequency is conserved across conditions, but where the planes occupied by these rotations are specific to each condition (Fig. 2d; Methods).

For each dataset independently, we determined the optimal number of temporal basis functions using cross-validation (Methods). For both monkeys, 9 temporal basis functions optimally explained M1 activity, and 7 temporal basis functions optimally explained PMd activity. The temporal basis functions reflect our findings above: in PMd and M1, the basis functions described rotations at 0.5, 1.5, and 2.5 Hz, while the additional pair of basis functions for M1 described rotations at 4 Hz (Fig. 2e; Extended Data Fig. 3a; Extended Data Fig. 4a). Despite the heterogeneity of neural responses during reaching, this factorization captured an average of 90-96% of the population variance (s.d. < 2% for all datasets), far above that expected for similarly smoothed but randomly simulated neural activity (32-42%, $p < 0.01$ for all datasets against smoothed-noise control, Wilcoxon Signed Rank test; Extended Data Fig. 4b), and included complex features of individual units' responses (Fig. 2f-g; Extended Data Fig. 3b).

We then tested whether the shared set of temporal basis functions accurately reflected the frequencies present for each condition, or alternatively whether they merely approximated distinct frequencies for different conditions. To do so, for each condition independently we

recovered the temporal basis functions by multiplying the inverse of the loading matrix by the data (Methods). We then quantified how well this projection explained the variance in the shared temporal basis functions (Fig. 2h). This projection explained 94-98% (s.d. < 2% for all datasets) of the variance in the temporal basis functions, indicating that all the temporal basis functions were present in all conditions and confirming that the frequencies are highly conserved across conditions (Extended Data Fig. 4d).

Of particular concern is that these shared temporal basis functions might comprise an artifact of the task structure. For example, neural activity might have a peak before movement onset and a second peak when stopping, leading to the appearance of a stereotyped activity pattern over time. To test this directly, we warped the neural responses for different conditions to equalize the durations of the reaches (Extended Data Fig. 4b,c). Prior to warping, no significant correlation was observed between a reach being unusually short or long in duration (absolute difference from mean duration) and how well the temporal basis functions could be recovered from neural activity on that condition (Fig. 2i, purple; Pearson's Rho, -0.05 to 0.09; $p > 0.30$ for all datasets). After warping, reaches with particularly short or long durations yielded worse recovery of the temporal basis functions (Fig. 2i, gray; Pearson's Rho, -0.45 to -0.71; $p < 0.01$ for all datasets). This argues that the temporal basis functions did not trivially arise from task structure.

Fixed points govern locally-rotational dynamics

The above findings are inconsistent with the predictions of two possible dynamical models: planar rotations ⁷ and variable-frequency dynamics ⁹. To further understand how the orientation of dynamics varied across conditions, we examined each frequency individually. Each frequency varied its orientation in state space: for a random pair of conditions, same-frequency rotations had an average alignment index of 15-68% (s.d. < 23% for all frequencies

and datasets), significantly below that expected for noisy realizations of planar rotations (Fig. 3a; chance 35-97%, $p < 0.01$ for all frequencies and datasets, Wilcoxon Rank Sum test; Extended Data Fig. 5). To understand the geometry of how the orientations changed, we designed a novel metric we term Subspace Excursion Angles that measures how differently oriented the rotational planes were across all conditions (Methods). We found that the plane for same-frequency rotations varied across conditions in 5-15 dimensions by more than 50° (Fig. 3b). Each frequency's plane therefore pivoted through a high-dimensional space across conditions. We additionally measured the smallest subspace angle between the same-frequency rotation across condition-pairs, and found that it was $> 20^\circ$ for 79-99% of pairs. This indicates that the same-frequency rotation frequently occupied disjoint planes across conditions.

We therefore asked whether this high-dimensionality is compatible with a dynamical systems view of motor cortex activity. Dynamical systems generally segregate different “modes,” such as rotations of different frequencies, to non-overlapping subspaces of state space. To determine whether this segregation was present in the data, we asked whether, across condition-pairs, the planes occupied by different-frequency rotations differed more than same-frequency rotations. They did (Fig. 3c; $p < 0.01$ for all frequencies and datasets, Wilcoxon Rank Sum test; Extended Data Fig. 6), suggesting that the changes in orientation of each rotational plane were large but structured to avoid reusing the same plane for different-frequency rotations.

The behavior of a dynamical system is determined locally by “fixed points” in state space. Examining these fixed points has therefore been used to study the behavior of recurrent neural networks across a variety of tasks^{29,30}. To apply this approach to motor cortex activity, where the system's equations are unknown, on each condition we identified the putative fixed point's location in state space as the point around which activity rotated (Methods). The fixed points for all conditions formed a ring-shaped manifold, with visible correspondence to reach angle, and were spread out in separate dimensions from rotational activity (Fig. 3d). We hypothesized that

the orientation of the rotations for a given condition might depend on the location of the fixed point for that condition. An analogy would be a record player placed in a large bowl: the record always rotates at the same frequency, but the tilt of the record is determined by the placement of the player on the wall of the bowl. We term this model the Location-Dependent Rotations (LDR) model (Fig. 3e, left). We fit an LDR model to motor cortex activity, linearly predicting the orientation of the rotations from the location of the fixed point. The LDR model could explain 73-85% of the population variance (end-to-end 6-fold cross-validated; s.d. < 16%, 31-52% chance, $p < 0.01$ vs. shuffle control for all datasets), indicating that the location of the fixed point on each condition determined the orientation of the planes occupied by rotations (Fig. 3e, right). For a given condition, the orientation of each rotation was informative about every other, but this effect was stronger for low-frequency rotations predicting higher-frequency rotations, and strongest for the fixed point predicting each rotation (Fig. 3f).

In dynamical systems, the trajectory of the state over time is also determined by the initial state. For motor cortex, that initial state is pre-movement preparatory activity^{7,19,21}. As with fixed points, initial states visibly related to reach type, with similar-angle reaches having similar initial states (Fig. 3g). We fit and evaluated this model of motor cortex identically to an LDR model, but using each condition's initial state rather than fixed point as a predictor. This model explained an average of 75-86% of the population variance (s.d. < 10%, 25-49% chance, $p < 0.01$ vs. shuffle for all datasets), indicating that the location of the initial state on each condition similarly determined the planes occupied by rotations (Fig. 3h).

Given that the fixed point of the system varied across conditions, the dynamics were not consistent with previous models. Nonetheless, if trajectories varied smoothly across conditions and had low tangling¹⁵ (i.e., trajectories never headed in different directions from nearby states; Fig. 3i, left), it remains possible that motor cortex could be well described as a dynamical system with neural activity proceeding deterministically from an initial state to a corresponding fixed

point. To test for these properties, we used nonparametric regression to fit a single smooth dynamical system to motor cortex activity on all conditions (Methods). To test the fit of these dynamics, we used cross-validation, fitting on 5% the conditions and then simulating forward the activity of held-out conditions from their corresponding initial states. We found the fit dynamics could explain an average of 65-78% of the population variance (Fig. 3i right; s.d. < 22%, $p < 0.01$ vs. 13-18% shuffle control for all datasets). This indicates that despite variability in rotational plane orientation, motor cortex activity was well-described as a dynamical system – albeit of a novel form.

Figure 3: Rotations reflect high-dimensional dynamics in motor cortex. **a**, Distribution of alignment indices for same-frequency rotations on pairs of conditions vs. noisy realizations of planar rotations. **b**, Subspace excursion angles for rotations of a single frequency across conditions vs. expected excursion angles for noisy realizations of planar rotations. **c**, Cumulative probability of alignment indices for same-frequency rotations on pairs of conditions vs. different-frequency rotations on pairs of conditions. **d**, Estimated fixed points across conditions in state space, M1-N. Colored according to reach angle. **e**, *Left*, Illustration of LDR model, where the orientation of rotations depends on the fixed point's location in state space; *Right*, distribution of variance explained by LDR model in neural data vs. shuffle control. **f**, Heatmap of mean variance explained predicting the plane occupied by one rotational frequency from another frequency's plane, M1-N. **g**, Initial points across conditions in state space, M1-J. **h**, *Left*, illustration of rotational plane's dependence on initial state, where orientation of rotations depends on initial state location in state space; *right*, distribution of variance explained by the model in neural data vs. shuffle control. **i**, *Left*, illustration of dynamics, where the population derivative changes smoothly as a function of population state; *right*, distribution of variance explained by fit dynamics in neural data vs. shuffle control. **j**, Total rotational variance vs. variance in the first jPCA plane, M1-N and M1-J. **k**, Alignment index between the first jPCA plane and neural activity vs. variance in the first jPCA plane, M1-N and M1-J. **l**, *Left*, distribution of fixed points and initial states along axis of maximal separation; *right*, projection of population activity over time for each condition on to axis of maximal separation between initial state and fixed point. In **d**, **g**, **l**, colors are assigned based on reach angle.

Reconciliation with previously-described observations

To relate high-dimensional motor cortex dynamics to previous rotations models, we used the standard method, jPCA ⁷, to fit planar rotations to our datasets. As previously described, jPCA uncovered rotations in motor cortex activity beginning 150 milliseconds prior to movement onset. However, not all conditions rotated equally in the jPCA planes. One possible explanation is that conditions that rotate less in the jPCA planes have lower-amplitude firing rate oscillations. To measure total rotational activity, for each condition we summed the variance of the rotations in all LDR dimensions excluding the fixed point. For four of six jPCA planes in our datasets we found no significant correlation between total rotational activity and the magnitude of the jPCA-identified rotations (Fig. 3j; Pearson's rho, -0.11-0.1; $p > 0.22$), while for two jPCA planes, one for PMd and another for M1 in monkey J, we found a modest correlation (Pearson's rho, 0.3-0.4; $p < 0.01$). This argues that a given condition having small rotations in the jPCA plane does not mean that condition had weak rotations. Alternatively, we hypothesized that conditions that rotated less in the jPCA planes simply rotated in spaces that were nearly orthogonal to the jPCA planes. To quantify this, we measured the alignment index between each condition's loading matrix and the jPCA planes. We found this measure was strongly correlated with the variance in the jPCA planes (Fig. 3k; Pearson's rho, 0.82-0.96; $p < 0.01$ for all jPCA planes and datasets), indicating that variability in the magnitude of jPCA rotations were not differences in the strength of rotations, but instead reflected misalignment of some conditions' rotations with the jPCA plane.

To relate the geometry of initial states and fixed points to the previously-described condition-invariant trigger signal (CIS) ²⁴, we used Fisher Discriminant Analysis to identify the dimension of maximal discriminability between fixed points and initial states on all conditions for each dataset. We found a dimension in each dataset of strong discriminability between fixed points and initial states (Fig. 3l, left; ROC-AUC = 1 vs. random projection control, $p < 0.01$ for all

datasets). Projections onto this dimension had the initial states tightly clustered, and the fixed points tightly clustered. Activity consisted of a condition-invariant trajectory that started at the initial states and translated over to the fixed points (Fig. 3I, right). This argues that the previously-observed CIS trigger signal reflects a large, consistently-oriented difference in the location of the initial states relative to the fixed points.

Dynamics orientation is linearly related to reach kinematics

Previous work has described correlations between motor cortex activity and numerous kinematic parameters, i.e., tuning³¹⁻³⁷. However, these kinematic-encoding models of motor cortex activity have typically generalized poorly to new conditions³⁸. Dynamical models of motor cortex, in contrast, generalize somewhat better¹¹ but do not address the relationship between motor cortex activity and kinematics. We asked whether the orientation of the planes occupied by rotations might have a simple relationship with reach kinematics. To this end, we fit an LDR-based linear encoding model of motor cortex activity. In this model, the reach kinematics are encoded in the location of the fixed point together with the orientation of the rotational planes. This model was fit and evaluated similarly to the LDR model, but using a kinematic feature-vector describing each condition's reach kinematics as the predictors of that condition's loading matrix, rather than the fixed point or initial state (Fig. 4a; Extended Data Fig. 7a). LDR-based encoding explained an average of 60-81% of the population variance (s.d. < 16% for all datasets), significantly greater than standard linear encoding models of motor cortex activity (Fig. 4b-c; 21-47%, $p < 0.01$ for all datasets, Wilcoxon Signed Rank test).

As encoding and decoding models are inverses, we inverted our LDR-based encoding model to decode reach kinematics from neural activity³⁹. We first evaluated LDR-based decoding on the trial-averaged data using cross-validation: the model was fit to half the conditions, then variance explained for hand position over time was evaluated on the other half.

LDR-based decoding explained an average of 92-97% of the variance in hand position (s.d. < 9% for all datasets), significantly greater than standard regression decoding from (smoothed) motor cortex activity (Fig. 4d-e; 69-91% standard, $p < 0.01$ for all datasets, Wilcoxon Signed Rank test; Extended Data Fig. 7b). We then applied LDR-based decoding to decode reach kinematics from single-trial, non-smoothed spiking data. To do so, we performed a stringent cross-validation: we fit an LDR-based decoder on half the conditions using trial-averaged data, then predicted reach kinematics for the single trials of held-out conditions (Extended Data Fig. 7c). LDR-based decoding explained an average of 83-95% of the variance in hand position (s.d. < 18% for all datasets) on single trials, significantly greater than standard single-trial decoding methods such as regressing from smoothed or binned spiking activity (37-64%), or regression from factors inferred by GPFA⁴⁰ (43-61%; Fig. 4f-g; $p < 0.01$ for all datasets, Wilcoxon Signed Rank test).

Low-frequency dynamics explain reach-angle tuning

A large body of work has used kinematic encoding models to describe modulation of average firing rates in motor cortex⁴¹. Though these models poorly describe the fast multiphasic features of motor cortex^{11,14}, they approximate aggregate activity well at long time-scales. In particular, for straight reaches, cosine-tuning has been used to study single-neuron encoding of reach angle^{33,42}. For the straight-reach conditions in our datasets we fit cosine-tuning models to recorded units, explaining an average of 47-57% (s.d. < 26% for all datasets) of the variance in firing rates averaged over both trials and time.

Above, we found that the location of dynamics in state space correlated strongly with reach kinematics. We hypothesized that tuning reflects the large low-frequency components of these dynamics, especially the fixed point location. To test this hypothesis, we projected the fixed point onto each recorded unit for each condition. We found that for three of our four

datasets the fixed point could account for the mean tuning: the projection of the fixed point alone yielded an average correlation of 0.83-0.92 (s.d. < 0.14) with the mean firing rate. For the remaining dataset, the projection of the fixed point and the 0.5 Hz rotation yielded a correlation of 0.83-0.93 (s.d. < 0.14 for all datasets) with the mean firing rate (Fig. 4h). This correlation was significantly greater than the correlation with higher-frequency components ($p < 0.01$ for all datasets, Wilcoxon Signed Rank test). This observation unifies tuning models and dynamical systems models: kinematic tuning in M1 and PMd is synonymous with the low-frequency features of motor cortex dynamics, especially the fixed point, and this location determines the orientation of the higher-frequency dynamics.

Figure 4: Orientation of rotations encodes for kinematic features of reaching movements.

a, Schematic of LDR encoding/decoding as a linear relationship between kinematic features and subspace occupied by neural dynamics. **b**, *Left*, activity of example neuron; *center*, reconstruction by LDR-based encoding; *right*, standard linear-encoding model. **c**, Distribution of population variance explained for LDR-based encoding and standard linear-encoding model. **d**, *Left*, example reach conditions; *center*, reaches predicted from neural activity by LDR-based decoding; *right*, standard regression decoding. **e**, Distribution of variance explained in hand position by LDR-based decoding and standard regression decoding on trial-averaged neural data. **f**, *Left*, single-trial reaches on example reach conditions; *center*, reaches predicted from single-trial spiking activity by LDR-based decoding; *right*, standard regression decoding on smoothed spiking. **g**, Distribution of variance explained in hand position by LDR-based decoding and standard regression decoding on binned and smoothed spiking activity or factors uncovered by GPFA. **h**, *Left*, example unit's mean firing rate as a function of reach angle, along with projection of the fixed point and tuning curve fit to the mean firing rate; *right*, distribution of correlation coefficients of the mean firing rate with the projection of the fixed point low-frequency rotations and high-frequency rotations. In **b**, **d**, **f**, colors are assigned based on reach angle.

Discussion

Here we argue for a new understanding of motor cortex dynamics, of a “record player in a bowl”: preparatory activity selects a fixed point on a curved manifold, and this location determines the orientation of rotational dynamics. Traditional “tuning” corresponds to the fixed point location, while the dynamics focus on the rotations. This model fits the data much more closely than previous dynamical systems models. These findings also agree with previous suggestions that motor cortex is higher dimensional than previously thought ⁴³, and with hints that the rotational plane may vary ^{15,44}. This view appears to challenge notions of fixed output dimensions to the muscles ²⁰. However, it is compatible: many vectors can be orthogonal to one another but have substantial projections onto a common output vector. The LDR model allows the brain to exploit n dynamical dimensions while maintaining an average output projection of $1/\sqrt{n}$ (the projections onto the hypotenuse).

Exploiting manifold curvature ^{45,46} presents a number of advantages for the system. First, the linear relationship between the fixed point location and the reach kinematics solves a major component of the “inverse model problem”: how to produce the needed command signals given a desired outcome. Instead of the brain explicitly solving for the outputs needed, the manifold curvature orients the rotations to produce the needed output signals. Although this manifold curvature must be learned, trial-to-trial this turns the complicated inverse model problem into an easy, linear one. Second, using different dimensions for different reach conditions can allow learning of multiple behaviors without catastrophic interference ⁴⁷, which is otherwise a major challenge for neural networks. Third, the use of distant portions of a curved manifold can reduce interference between conditions ⁴⁸, while reaping the benefit of strong generalization due to local linearity ^{49,50}. Finally, by allowing rotations to pivot in higher-dimensional space, the system can control each output dimension with much greater flexibility than if limited to planar rotational dynamics.

Methods

Pre-processing data

All data have been analyzed in previous studies. In brief, for both monkeys data was collected using two Blackrock “Utah” arrays, one in PMd and one M1. Spiking activity was binned at 1 millisecond, smoothed with a 20 ms Gaussian kernel, trial-averaged, and down-sampled to 10 ms resolution.

Condition-specific linear dynamics systems

For a single condition c 's neural data $\mathbf{X}^{(c)}$, constituting a neurons-by-time points (N -by- T) matrix, fitting linear dynamics can be accomplished by regressing the population discrete-derivative against the current population state, then adding the identity to the fitted matrix. As firing modulation only begins approximately 150 milliseconds prior to movement onset, we fit dynamics on subsequent time points. We eigendecomposed the resulting N -by- N dynamics matrix $\mathbf{M}^{(c)}$ as $\mathbf{M}^{(c)} = \mathbf{E}^{(c)} \text{diag}(\mathbf{d}^{(c)}) \mathbf{E}^{(c)-1}$, where $\mathbf{d}^{(c)}$ is a vector of eigenvalues and $\mathbf{E}^{(c)}$ is a matrix of eigenvectors, both specific to condition c . To avoid overfitting to noise, before fitting dynamics we projected $\mathbf{X}^{(c)}$ into its top k singular vectors, so that after factorizing $\mathbf{X}^{(c)}$ as $\mathbf{X}^{(c)} = \mathbf{U}^{(c)} \mathbf{\Sigma}^{(c)} \mathbf{V}^{(c)*}$, linear dynamics were fit to $\mathbf{U}^{(c)}(1:k, :)^{(c)*} \mathbf{X}^{(c)}$. Based on estimates of the number of temporal basis functions (below), we used $k = 9$ in the case of M1 and $k = 7$ in the case of PMd. We extracted the frequencies and half-lives of the dynamics of condition c as $100/(2\pi/\text{atan2}(\text{imag}(\mathbf{d}), \text{real}(\mathbf{d})))$ for frequency and $\log(0.5)/(\log(|\mathbf{d}|)/100)$ for half-life, where **atan2** is the four-quadrant arctangent function, **real** and **imag** select the real and imaginary portions of their inputs, **log** is the natural logarithm, $||$ denotes the absolute value, and all functions operate element-wise. We further calculated the (non-symmetric) alignment index between two conditions i and j as $1 - \text{var}(\mathbf{U}(:, 1:k)^{(i)} \mathbf{U}(:, 1:k)^{(j)*} \mathbf{X}^{(i)}) / \text{var}(\mathbf{U}(:, 1:k)^{(i)} \mathbf{X}^{(i)})$.

Controlling for alignment index of randomly-oriented planes

To control for the expected alignment index of randomly oriented planes in state space, we performed a Monte Carlo estimate of the distribution of the expected alignment index for each dataset independently. To do so, we first used a Singular Value Decomposition on a N -by- CT matrix to identify the weighted singular vectors of the dataset. We then drew two random matrices of sizes N -by- k ($k = 9$ for M1 and $k = 7$ for PMd) and multiplied them by the weighted singular values of the dataset. Finally, we quantified the alignment index between these two samples. We repeated this procedure 11664 times to produce a distribution of the expected alignment index of randomly-oriented planes.

Extracting temporal basis functions from neural data

Typically, when neural data is factored (such as with PCA), the basis provided is patterns of covariance across neurons. The result of the factorization is therefore a weighted sum of these basis patterns at each time point for each condition. Here, we factorize activity differently: the basis provided is patterns of activity over time. The response of a given neuron for a given condition (over an entire trial) is a weighted sum of these temporal basis functions. Our method is shown schematically in Figure 2c. Formally, given C recordings of a neural population with N units across T time-points, neural recordings are naturally organized into a data tensor $\mathbf{X} \in \mathbb{R}^{N \times T \times C}$, so that element $\mathbf{X}(n, t, c)$ of the tensor is the firing rate of neuron n on time-point t on condition c ¹¹. Given conserved temporal structure across neurons and conditions, this data tensor can be approximated as the product of a matrix of k temporal basis functions $\mathbf{B} \in \mathbb{R}^{T \times k}$ with a loadings

tensor $\mathbf{L} \in \mathbb{R}^{N \times k \times C}$, providing a low-rank approximation to $\mathbf{X}$. In this approximation, $\mathbf{B}$ specifies the patterns of firing rate modulation over time present in the neural population, and $\mathbf{L}$ specifies how the temporal basis functions in $\mathbf{B}$ are used to reconstruct the population activity in $\mathbf{X}$. In particular, element $\mathbf{L}(n,k,c)$ specifies how temporal basis function k is used to reconstruct the activity of neuron n on condition c . We extracted $\mathbf{B}$ and $\mathbf{L}$ to minimize the mean-squared error between $\mathbf{X}$ and its approximation via a Singular Value Decomposition on the unfolding of $\mathbf{X}$ into $\mathbf{X}' \in \mathbb{R}^{NC \times T}$, which decomposes $\mathbf{X}'$ into $\mathbf{X}' = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$. This can be used for approximation by selecting the first k singular values and vectors. Of these terms, we defined the unfolding of $\mathbf{L}$, $\mathbf{L}' \in \mathbb{R}^{NC \times k}$ as $\mathbf{U}(:,1:k)\mathbf{\Sigma}(1:k,1:k)$, while $\mathbf{B}$ was defined as $\mathbf{V}(:,1:k)$. As neural activity on separate conditions in $\mathbf{X}$ is separately reconstructed, we sliced $\mathbf{L}$ into loadings matrices $\mathbf{L}^{(c)}$ so that $\mathbf{X}(:, :, c)$ is approximated by $\mathbf{L}^{(c)}\mathbf{B}^*$. We refer to this decomposition as a “temporal basis factorization”.

Here, we show that in the case where motor cortex dynamics have rotations of conserved frequencies and half-lives, but rotational planes are not conserved across conditions, temporal basis factorization correctly identifies the conserved rotations and planes on each condition in which rotations occur. Let the dynamics on condition c be $\mathbf{x}^{(c)}_{t+1} = \mathbf{M}^{(c)}\mathbf{x}^{(c)}_t$, where $\mathbf{x}^{(c)}_t \in \mathbb{R}^N$ is the population state at time t on condition c and $\mathbf{M}^{(c)}$ is the dynamics matrix. Eigendecompose $\mathbf{M}^{(c)}$ to yield $\mathbf{x}^{(c)}_{t+1} = \mathbf{E}^{(c)}\text{diag}(\mathbf{d})\mathbf{E}^{(c),-1}\mathbf{x}^{(c)}_t$, where $\mathbf{d}$ is a vector of eigenvalues determining the frequency content and half-lives of the patterns of activity produced by the population, and $\mathbf{E}^{(c)}$ contains the eigenvectors of $\mathbf{M}^{(c)}$. Critically, $\mathbf{d}$ is constructed as shared across conditions. By iteratively applying the dynamics matrix to run forward additional time steps, $\mathbf{x}^{(c)}_{t+p} = \mathbf{E}^{(c)}\text{diag}(\mathbf{d}^p)\mathbf{E}^{(c),-1}\mathbf{x}^{(c)}_t$ for all integers p , and we see that the dynamics are controlled by $\mathbf{d}$. The half-life of decay (or expansion) is dictated by the magnitude of the entries $\mathbf{d}$, and the frequencies of rotation are dictated by the angles of the entries of $\mathbf{d}$ in the complex plane. We then reformulate the original problem as $\mathbf{X}^{(c)} = \mathbf{E}^{(c)}\mathbf{D}^{(c)*}$, where $\mathbf{X}^{(c)} \in \mathbb{R}^{N \times T}$ and $\mathbf{X}(:, t) = \mathbf{x}^{(c)}_t$, $*$ indicates transposition, $\mathbf{D}^{(c)} \in \mathbb{R}^{N \times T}$ and $\mathbf{D}^{(c)}(t, :) = \text{diag}(\mathbf{d}^t)\mathbf{E}^{(c),-1}\mathbf{x}^{(c)}_0$, and $\mathbf{E}^{(c),-1}\mathbf{x}^{(c)}_0$ is the initial condition. To remove the dependence of $\mathbf{D}^{(c)}$ on the condition, we rewrite the factorization as $\mathbf{X}^{(c)} = \mathbf{L}^{(c)}\mathbf{B}^*$, where the unscaled eigenvectors are defined $\mathbf{L}^{(c)} = \mathbf{E}^{(c)}\text{diag}(\mathbf{E}^{(c),-1}\mathbf{x}^{(c)}_0)$ and the temporal basis functions are $\mathbf{B}(t, :) = \mathbf{d}^t$, weighting the eigenvectors rather than the patterns of activity over time. Note that eigenvalues in $\mathbf{d}$ that are nearly-zero in magnitude will decay away very rapidly, and therefore can be truncated to produce a low-rank model of motor cortex dynamics on each condition.

Cross-validating the temporal basis functions

To avoid overfitting, we used cross-validation to determine the number of temporal basis functions. For a subset of trials (the training set), temporal basis functions and loading matrices were computed as usual. For a given number of temporal basis functions, we then reconstructed the activity of each neuron over time for each condition, and quantified the variance explained in the PSTHs made from the test set. This was repeated for 10 splits, and sweeping the number of temporal basis functions. We averaged the variance explained across splits for each number of temporal basis functions, then chose the minimum of the curve as the number of temporal basis functions. In cases where the number of basis functions chosen was even, we rounded down to the nearest odd number in order to have each basis function paired with another phase-shifted, same-frequency basis function.

Testing the ability to recover temporal basis functions from neural data

One concern about the temporal basis factorization is that good fits might not actually indicate a “conserved eigenvalues, variable eigenvectors” model, but instead may act as a generalized

Fourier transform and fit any smooth data reasonably. The validity of the LDR dynamics model rests on two requirements. First, all the temporal basis functions must be used in each condition, so that no columns of a loading matrix can be zero for any condition. Second, for the temporal basis functions to describe dynamics, all the temporal basis functions must be used to reconstruct population activity along linearly-independent dimensions. The validity of these requirements was tested for each condition by quantifying the invertibility of each condition's loading matrix, by quantifying the variance explained in $\mathbf{B}^*$ by $\mathbf{L}^{(c),-1}\mathbf{X}^{(c)}$.

Uncovering dynamical structure in temporal basis functions

The loadings matrix $\mathbf{L}^{(c)}$ from the temporal basis factorization may be used to uncover the coordinates in which dynamics of conserved frequencies and half-lives occur on condition c . In general, however, the SVD-based approach for uncovering temporal basis functions will yield functions that are linear combinations of the eigenmodes of the system, and so do not contain pure frequencies or half-lives. To purify the dynamical modes, we fit a dynamics matrix to the temporal evolution of $\mathbf{B}$, so that $\mathbf{B}(t+1,:) = \mathbf{M}\mathbf{B}(t,:)$, where $\mathbf{M}$ can be identified via regressing the state of the temporal basis functions against the discrete-derivative, then adding the identity to the fitted matrix. Because firing modulation only begins approximately 150 milliseconds prior to movement onset, we fit dynamics only to subsequent timepoints in the basis functions. We then eigendecomposed $\mathbf{M}$ into $\mathbf{M} = \mathbf{E}\mathbf{diag}(\mathbf{d})\mathbf{E}^{-1}$, where $\mathbf{d}$ is a vector of eigenvalues and $\mathbf{E}$ is a matrix of eigenvectors. To purify $\mathbf{B}$, we replaced it with $\mathbf{B}\mathbf{E}^{-1}$, and to preserve the factorization, replaced $\mathbf{L}^{(c)}$ with $\mathbf{L}^{(c)}\mathbf{E}$. To avoid introducing complex numbers into the temporal basis itself, for complex-conjugate pairs of eigenvalues $\mathbf{d}(i)$, $\mathbf{d}(\bar{i})$, with corresponding eigenvectors $\mathbf{E}(:,i)$, $\mathbf{E}(:,\bar{i})$, we extracted the corresponding eigenplane as $\mathbf{E}(:,i) + \mathbf{E}(:,\bar{i})$ and $i(\mathbf{E}(:,i) - \mathbf{E}(:,\bar{i}))$, where i is the complex unit⁷. To avoid exponential growth due to noise in the fitting procedure, we capped eigenvalue magnitude at 0.99. In practice, this only required a single real eigenvalue with magnitude slightly greater than one to be altered. Each eigenvalue can then be used to calculate the half-lives and frequencies of the corresponding column of $\mathbf{B}$.

Subspace Excursion Angles

In this analysis we sought to understand whether, for a set of low-dimensional subspaces $\mathbf{S} = \{\mathbf{S}^{(1)}, \dots, \mathbf{S}^{(c)}\}$ of a high-dimensional space, these subspaces only “wobbled” slightly relative to one another, or whether the objects truly “pivoted” far into many different dimensions and formed large angles with one another. Here, each $\mathbf{S}^{(i)}$ was the plane of rotations for a given frequency on a given condition (taken from the columns of $\mathbf{L}^{(i)}$), embedded in high-dimensional neural space. Existing methods of estimating dimensionality do not distinguish wobbling from pivoting, as long as the use of each additional dimension is above the noise level. To better characterize the geometry of this set of subspaces, we took one of the observed subspaces $\mathbf{S}^{(i)}$ as a seed and found the greatest subspace angle (the principal angle) to any other observed subspace $\mathbf{S}^{(j)}$. We then replaced $\mathbf{S}^{(i)}$ with the union of $\mathbf{S}^{(i)}$ and the vector in $\mathbf{S}^{(j)}$ that formed the greatest angle with $\mathbf{S}^{(i)}$. We repeated the process to find the next greatest angle to the augmented $\mathbf{S}^{(i)}$, and so forth. We optimized over all possible sequences of angles, including a tolerance for the maximum angle so that we could maximize the overall sum of angles. For computational efficiency, the tolerance used was 1% of the maximum subspace angle at any given depth in the angle sequence. To break ties between sequence permutations, we identified the unique sequence in the tied set that produced a monotonically decreasing series of angles. To stretch terminology, this method produces a sequence of angles which characterize how thoroughly a set of subspaces spans a given dimensionality.

Controlling for noisy realizations of a single rotational plane

To estimate the expected alignment index for noisy realizations of a single rotational plane, for each identified rotational frequency in each dataset we performed a bootstrap control on each condition independently. For each condition, we randomly partitioned the trials in half, then smoothed and trial-averaged each half of the partition independently. In each half of the partition, we then used regression to fit a loading matrix to neural activity. After sub-selecting the columns corresponding to the plane occupied by rotations of that frequency, we quantified the alignment index for each half of the partition with the other half. For each frequency and each condition, we performed 500 bootstraps to estimate the distribution of alignment indices, then pooled distributions across conditions. We also estimated the noise floor for Subspace Excursion Angles. To do so, for each condition we randomly selected 108 of the partitions, pseudo-inverted $\mathbf{B}$ to estimate a loading matrix from the neural activity, selected the columns corresponding to the plane occupied by rotations of the appropriate frequency, then measured the subspace excursion angles between these 108 realizations of a single plane. We pooled subspace excursion angles among conditions to estimate the distribution of excursion angles for noisy realizations of a single plane for each frequency. Note that because these procedures half the amount of data used, the resulting noise floor is an overestimate.

Predicting and decoding from loading matrices

Let $\mathbf{L} \in \mathbb{R}^{Nk \times p}$ be the matrix of vectorized loadings matrices in the training set of p conditions, where $\mathbf{L}(:,c) = \text{vec}(\mathbf{L}^{(c)})$. Let $\mathbf{Q}' \in \mathbb{R}^{d' \times p}$ be the matrix of explanatory variables across the training set of p conditions, where $\mathbf{Q}'(:,c)$ are the d' explanatory variables on condition c . Using a Singular Value Decomposition, factorize $\mathbf{Q}'$ as $\mathbf{Q}' = \mathbf{U}\mathbf{\Sigma}\mathbf{V}$. Define the low-dimensional approximation to $\mathbf{Q}'$ as $\mathbf{Q} = \mathbf{\Sigma}(1:z,1:z)\mathbf{V}(:,1:z)^*$, where z is a parameter chosen to capture some specified percentage of the sum of the singular values. Let $\mathbf{X}^{(i)} \in \mathbb{R}^{N \times T}$ and $\mathbf{q}^{(i)} \in \mathbb{R}^{d'}$ respectively denote the matrix of motor cortex population activity and explanatory variables on condition i , a test set condition that was held out from forming the temporal basis functions $\mathbf{B}$ and the matrix $\mathbf{Q}'$. Population activity on condition i can then be inferred as $\text{reshape}(\mathbf{L}\mathbf{Q}^*(\mathbf{Q}\mathbf{Q}^*)^{-1}\mathbf{U}(:,1:z)^*\mathbf{q}^{(i)}, N, k)\mathbf{B}^*$, where reshape reverses the vectorization operator to form an $N \times k$ matrix. The explanatory variables on condition i can be predicted as $\mathbf{U}(:,1:z)\mathbf{Q}\mathbf{L}^*(\mathbf{L}\mathbf{L}^*)^{-1}\text{vec}(\mathbf{X}^{(i)}\mathbf{B}^{-1})$, where the term $\mathbf{X}^{(i)}\mathbf{B}^{-1}$ is a regression-fit loading matrix for condition i .

Extracting fixed-point location and prediction of loading matrices

Given that fixed points determine the behavior of a nonlinear dynamical system within a local neighborhood, we considered whether the planes occupied by conserved dynamics are predictable from the location of the fixed point. The loadings matrix $\mathbf{L}^{(c)}$ describes the subspace which the dynamics in $\mathbf{B}$ occupy on condition c , so this question can be framed as inquiring whether there exists shared structure between the loadings matrices and identified fixed points. The putative fixed point on a given condition can be identified using the loading matrix $\mathbf{L}^{(c)}$. For each dataset, the first temporal basis function $\mathbf{B}(:,1)$ described exponential decay with a half-life of 3.2-69.3 seconds, 4.5-99x longer than the longest reach duration. Given this mismatch in timescales, at the timescale of a reach these temporal basis functions can be approximated as stationary. With this approximation, $\mathbf{L}(:,1)^{(c)}$ then describes a stationary dimension in state space on condition c , where $\mathbf{L}(:,1)^{(c)}$ is fit via regression on held-out conditions. Weighting this vector by the mean of the corresponding temporal basis function, $\mathbf{L}(:,1)^{(c)}(\mathbf{B}(1,1) + \mathbf{B}(2,1) + \dots + \mathbf{B}(T,1))/T$, then identifies the approximate fixed point on condition c into which rotational neural activity decays. For predicting the loading matrices, we chose z to capture 80% of the magnitude of the

summed singular values. As a column of the loading matrix was used as a predictor, only the rotation-corresponding columns, $\mathbf{L}(:,2:\text{end})^{(c)}$, were predicted, replacing the inferred quantity as $\mathbf{L}(:,1)^{(c)}\mathbf{B}(:,1)^* + \text{reshape}(\mathbf{LQ}^*(\mathbf{QQ}^*)^{-1}\mathbf{U}(:,1:z)^*\mathbf{L}(:,1)^{(c)}, N, k-1)\mathbf{B}(:,2:\text{end})^*$. Controls for prediction quality were performed by shuffling columns of the loading matrices independently between conditions.

Extracting initial state location and prediction of loading matrices

Given that initial states determine the subsequent behavior of a dynamical system, we considered whether the eventual rotational planes were predictable from the corresponding preparatory activity. The loadings matrix $\mathbf{L}^{(c)}$ describes the subspace which the dynamics in $\mathbf{B}$ occupy on condition c , so this question can be framed as inquiring whether there exists shared structure between the loadings matrices and preparatory activity. As a quantification of preparatory activity, for each condition we used the N -dimensional vector of mean firing rates 350 to 250 milliseconds prior to movement onset. For predicting the loading matrices, we chose z to capture 80% of the magnitude of the summed singular values. For the initial state, we used the pooled initial state of M1 and PMd to predict the loading matrices of either M1 or PMd separately. Controls for prediction quality were performed by shuffling the initial states and loading matrices between conditions.

Fitting dynamics to motor cortex activity using non-parametric regression

Substantial evidence has argued that during reaching motor cortex acts as a largely autonomous dynamical system, with the current population state determining the future population state^{5,6}. To test this generally and directly, we utilized a variant of nearest-neighbor regression to fit a function mapping from the current population state to the discrete-derivative. To fit this function, we used Softmax-Weighted Nearest Neighbor Regression, identical to standard inverse-distance weighted KNN regression but with $k = N$ and the inverse distances passed through a softmax function. We hand-tuned the temperature parameter B of the softmax function, setting $B = 0.5$. To reduce noise, we performed prediction after approximating neural activity using temporal basis functions. Controls for prediction quality were performed by independently shuffling each column of the loading matrices between conditions.

Linking orientation of motor cortex dynamics with reach kinematics

Given that motor cortex dynamics acts causally in producing movement, we considered whether the planes occupied by conserved dynamics encode for associated reach kinematics. As the loadings matrix $\mathbf{L}^{(c)}$ corresponds to the subspace which the dynamics in $\mathbf{B}$ occupy on condition c , this question can be framed as inquiring whether there exists shared structure between the loadings matrices and reach kinematics. Reach kinematics were quantified on each condition c as a matrix $\mathbf{K}^{(c)} \in \mathbb{R}^{T \times 2}$, where $\mathbf{K}^{(c)}(t,:)$ is the x - and y -coordinates of the hand at time t . As a predictor for loading matrices on each condition c , we used $\text{vec}(\mathbf{K}^{(c)})$. For predicting the loading matrices, As the kinematics were essentially noiseless, we chose z to capture 98% of the magnitude of the summed singular values. Note that to improve model performance by avoiding decreasing the SNR, we did not multiply $\mathbf{B}$ by the eigenvectors of its dynamics prior to fitting encoding or decoding models, so that $\mathbf{B}$ remains orthogonal. For single-trial decoding of kinematics, $\mathbf{X}^{(c)}$ is replaced by a $100\mathbf{T}^{(c)}$, where $\mathbf{T}^{(c)}$ is an N -by- T matrix of spike counts binned at 10 ms, with the multiplication by 100 converting the spiking to instantaneous firing rates in units of spikes/second. We both quantified the quality of predicting reach kinematics from M1 activity alone, PMd activity alone, or from both areas by averaging over the predictions of M1 and PMd. Controls for predicting neural activity from kinematics were performed by using linear regression to fit neural activity as encoding position, velocity and acceleration in the X and Y dimensions,

and the sine and cosine of the hand's angle from the origin. Controls for decoding kinematics from trial-averaged data were performed by using linear regression to read out the position of the hand from the neural state smoothed over time with a 100 ms Gaussian kernel, with a 100 ms lag between the neural state and the kinematics. Controls for decoding kinematics from single-trial data were performed by using linear regression to read out the position of the hand from the spiking binned at 20 or 100 ms, spiking smoothed with a 100 ms Gaussian kernel, or from 20 factors inferred from Gaussian Process Factor Analysis, with all decoding methods incorporating a 100 ms lag between the neural state and the kinematics.

Data availability

Data will be made available via DANDI.

Code availability

Novel tests (Subspace Excursion Angles) will be made publicly available. Other code will be made available by request.

References

1. Scott, S. H. Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* **5**, 532–46 (2004).
2. 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).
3. Crammond, D. J. & Kalaska, J. F. Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *J Neurophysiol* **84**, 986–1005 (2000).
4. Fu, Q. G., Flament, D., Coltz, J. D. & Ebner, T. J. Temporal encoding of movement kinematics in the discharge of primate primary motor and premotor neurons. *J Neurophysiol* **73**, 836–54 (1995).
5. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation Through Neural Population Dynamics. *Annu Rev Neurosci* **43**, 249–275 (2020).
6. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a dynamical systems perspective. *Annu Rev Neurosci* **36**, 337–59 (2013).
7. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
8. Hennequin, G., Vogels, T. P. & Gerstner, W. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* **82**, 1394–1406 (2014).
9. Logiaco, L., Abbott, L. F. & Escola, S. Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell Rep* **35**, 109090 (2021).
10. Omrani, M., Kaufman, M. T., Hatsopoulos, N. G. & Cheney, P. D. Perspectives on classical controversies about the motor cortex. *J. Neurophysiol.* **118**, 1828–1848 (2017).
11. Seely, J. S. *et al.* Tensor Analysis Reveals Distinct Population Structure that Parallels the Different Computational Roles of Areas M1 and V1. *PLoS Comput. Biol.* **12**, e1005164 (2016).
12. Perich, M. G. *et al.* Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior. <http://biorxiv.org/lookup/doi/10.1101/2020.07.30.228767> (2020) doi:10.1101/2020.07.30.228767.
13. Gallego, J. A., Perich, M. G., Miller, L. E. & Solla, S. A. Neural Manifolds for the Control of Movement. *Neuron* **94**, 978–984 (2017).
14. Michaels, J. A., Dann, B. & Scherberger, H. Neural Population Dynamics during Reaching Are Better Explained by a Dynamical System than Representational Tuning. *PLoS Comput Biol* **12**, e1005175 (2016).
15. Russo, A. A. *et al.* Motor Cortex Embeds Muscle-like Commands in an Untangled Population Response. *Neuron* **97**, 953–966.e8 (2018).
16. Pandarinath, C. *et al.* Neural population dynamics in human motor cortex during movements in people with ALS. *Elife* **4**, e07436 (2015).
17. Stavisky, S. D. *et al.* Neural ensemble dynamics in dorsal motor cortex during speech in people with paralysis. *Elife* **8**, e46015 (2019).
18. Kao, T.-C., Sadabadi, M. S. & Hennequin, G. Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. *Neuron* **109**, 1567–1581.e12 (2021).
19. 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).

20. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
21. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun* **7**, 13239 (2016).
22. Kohn, A. *et al.* Principles of Corticocortical Communication: Proposed Schemes and Design Considerations. *Trends Neurosci* **43**, 725–737 (2020).
23. Zimnik, A. J. & Churchland, M. M. Independent generation of sequence elements by motor cortex. *Nat Neurosci* **24**, 412–424 (2021).
24. Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects Movement Timing but Not Movement Type. *eNeuro* **3**, (2016).
25. Michaels, J. A., Dann, B., Intveld, R. W. & Scherberger, H. Predicting Reaction Time from the Neural State Space of the Premotor and Parietal Grasping Network. *J Neurosci* **35**, 11415–11432 (2015).
26. Lara, A. H., Cunningham, J. P. & Churchland, M. M. Different population dynamics in the supplementary motor area and motor cortex during reaching. *Nat Commun* **9**, 2754 (2018).
27. Gao, P. *et al.* A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv* 214262 (2017) doi:10.1101/214262.
28. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).
29. Golub, M. & Sussillo, D. FixedPointFinder: A Tensorflow toolbox for identifying and characterizing fixed points in recurrent neural networks. *JOSS* **3**, 1003 (2018).
30. Sussillo, D. & Barak, O. Opening the Black Box: Low-Dimensional Dynamics in High-Dimensional Recurrent Neural Networks. *Neural Computation* **25**, 626–649 (2013).
31. Tanji, J. & Evarts, E. V. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J Neurophysiol* **39**, 1062–1068 (1976).
32. Evarts, E. V. Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology* **31**, 14–27 (1968).
33. Georgopoulos, A. P., Kalaska, J. F., Caminiti, R. & Massey, J. T. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* **2**, 1527–1537 (1982).
34. 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**, 152–65 (2000).
35. Riehle, A., MacKay, W. A. & Requin, J. Are extent and force independent movement parameters? Preparation- and movement-related neuronal activity in the monkey cortex. *Exp Brain Res* **99**, 56–74 (1994).
36. 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**, 3130–46 (2006).
37. Hocherman, S. & Wise, S. P. Effects of hand movement path on motor cortical activity in awake, behaving rhesus monkeys. *Exp Brain Res* **83**, 285–302 (1991).
38. Oby, E. R., Ethier, C. & Miller, L. E. Movement representation in the primary motor cortex and its contribution to generalizable EMG predictions. *J Neurophysiol* **109**, 666–678 (2013).
39. Salinas, E. & Abbott, L. F. Vector reconstruction from firing rates. *J Comput Neurosci* **1**, 89–107 (1994).
40. Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J Neurophysiol* **102**, 614–35 (2009).

41. Branco, M. P., de Boer, L. M., Ramsey, N. F. & Vansteensel, M. J. Encoding of kinetic and kinematic movement parameters in the sensorimotor cortex: A Brain-Computer Interface perspective. *Eur J Neurosci* **50**, 2755–2772 (2019).
42. Schwartz, A. B. & Moran, D. W. Arm trajectory and representation of movement processing in motor cortical activity. *Eur J Neurosci* **12**, 1851–6 (2000).
43. Marshall, N. J. *et al.* *Flexible neural control of motor units*.
<http://biorxiv.org/lookup/doi/10.1101/2021.05.05.442653> (2021)
doi:10.1101/2021.05.05.442653.
44. Michaels, J. A., Schaffelhofer, S., Agudelo-Toro, A. & Scherberger, H. *A modular neural network model of grasp movement generation*.
<http://biorxiv.org/lookup/doi/10.1101/742189> (2019) doi:10.1101/742189.
45. Sohn, H., Narain, D., Meirhaeghe, N. & Jazayeri, M. Bayesian Computation through Cortical Latent Dynamics. *Neuron* **103**, 934–947.e5 (2019).
46. Balasubramaniam, R. *et al.* Neural Encoding and Representation of Time for Sensorimotor Control and Learning. *J Neurosci* **41**, 866–872 (2021).
47. Duncker, L., Driscoll, L., Shenoy, K. V., Sahani, M. & Sussillo, D. Organizing recurrent network dynamics by task-computation to enable continual learning. *Advances in Neural Information Processing Systems* **33**, (2020).
48. Sun, X. *et al.* *Skill-specific changes in cortical preparatory activity during motor learning*.
<http://biorxiv.org/lookup/doi/10.1101/2020.01.30.919894> (2020)
doi:10.1101/2020.01.30.919894.
49. Recanatesi, S. *et al.* Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nat Commun* **12**, 1417 (2021).
50. Chung, S., Lee, D. D. & Sompolinsky, H. Classification and Geometry of General Perceptual Manifolds. *Phys. Rev. X* **8**, 031003 (2018).

Extended data figure 1: Individual units in motor cortex are not uniformly well-described by rotational dynamics. **a**, Rotational activity uncovered by jPCA for each dataset independently. For M1 in both datasets, two of the identified planes contained activity that was strongly rotational, while for PMd in both datasets one plane contained strongly rotational activity. **b-e**, Recorded activity of individual units (real) and activity reconstructed from a “trigger signal and rotations” (CIS and jPCA) model for M1-N (**b**), PMd-N (**c**), M1-J (**d**), and PMd-J (**e**). In all subplots, colors are assigned based on reach angle.

Extended data figure 2: Frequency, but not orientation, of motor cortex dynamics is conserved. **a**, Rotational frequencies found by condition-specific dynamics for M1 and PMd for monkeys N and J. **b**, Density plot in the complex plane of eigenvalues for condition-specific dynamics for M1 and PMd for monkey N. **c**, Distribution of variance explained in hand position by linearly-decoding reach kinematics from condition-specific rotational frequencies and shuffle control. Frequency variation explained kinematics poorly. **d**, Distribution of alignment indices across condition-pairs and the corresponding PCA model.

Extended data figure 3: Factorizing motor cortex activity with temporal basis functions. a, Temporal basis functions for activity recorded from M1, monkey N. **b,** Example unit activity along with reconstructions provided by temporal basis functions, M1-N. In b, traces colored according to reach angle.

Extended data figure 4: Conserved rotational frequencies do not reflect artifacts of task or preprocessing. **a**, Original temporal basis functions (*left*), frequency-purified temporal basis functions (*center*), and temporal basis functions simulated from fit dynamics (*right*). **b**, Single-trial and trial-averaged activity for unit 61 in M1-N, along with corresponding time-warped activity and a rate-matched homogeneous Poisson process. **c**, Speed profiles for reaches performed by monkey J before and after time-warping. **d**, Temporal basis functions recovered from example condition's neural activity. In **b**, **c**, **d**, colors are assigned based on reach angle.

Extended data figure 5: Rotations in motor cortex activity pivot through a high-dimensional space across conditions. a-d, Distribution of alignment indices for same-frequency rotations on pairs of conditions and noisy realizations of a dimensionality-matched subspace for M1-N (a), M1-J (b), PMd-N (c), and PMd-J (d).

Extended data figure 6: Motor cortex segregates different-frequency rotations to different subspaces. a-d, Cumulative density function of alignment indices for same-frequency rotations and different-frequency rotations on pairs of conditions for M1-N (a), M1-J (b), PMd-N (c), and PMd-J (d).

Extended data figure 7: LDR allows decoding of kinematics from neural activity. **a**, 6 of 10-12 reach components used as a basis for kinematics during LDR-based decoding, M1-N. **b**, Neural activity smoothed with a 20 ms (*left*) and 100 ms (*right*) Gaussian kernel. Neural activity was smoothed with a 100 ms Gaussian kernel prior to standard regression decoding. **c**, Single-trial spiking, trial-averaged activity, and firing rates fitted to single trials using the temporal basis functions for example units. In **b** and **c**, colors assigned according to reach angle.