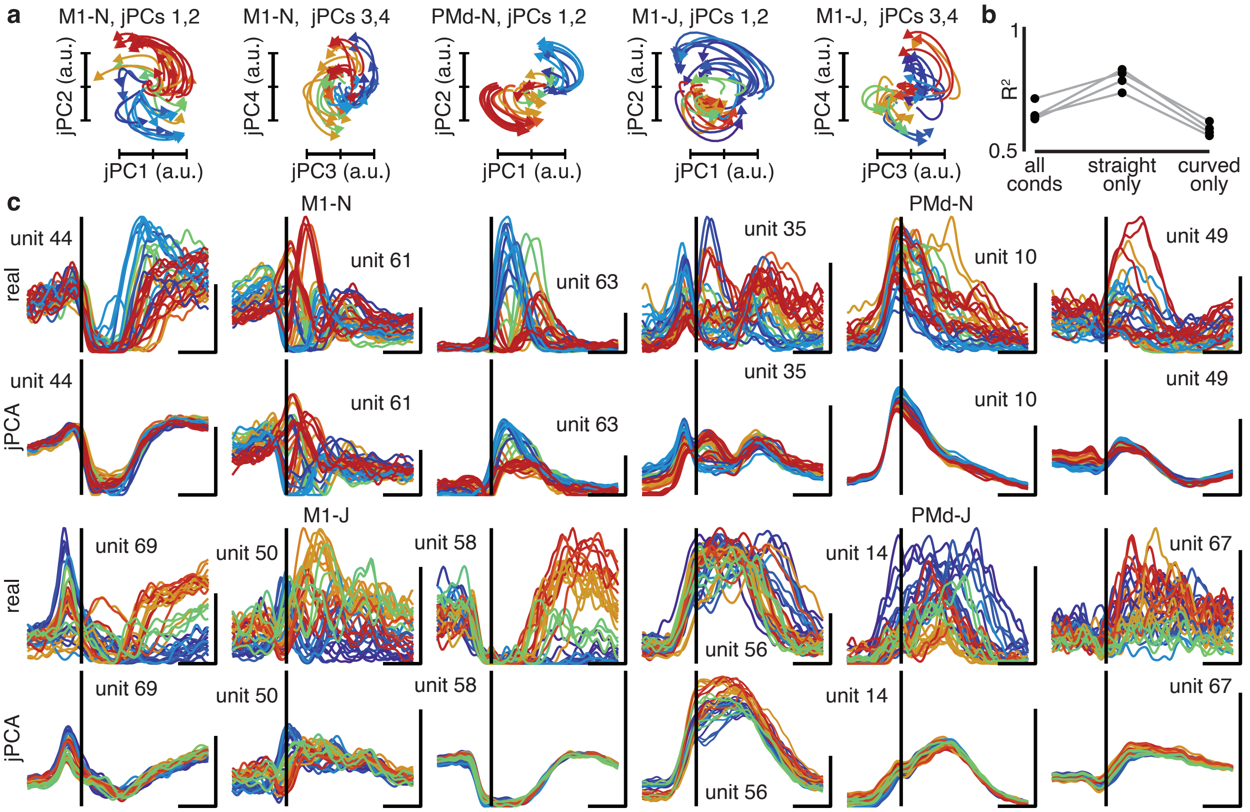

Supplementary Figure 1: jPCA rotations incompletely describe motor cortex activity. a, Rotations uncovered by jPCA across datasets. jPCA identified two rotational planes in M1 activity, and one rotational plane in PMd. **b**, The fit of rotational dynamics identified by jPCA depended heavily on reach complexity. jPCA fit motor cortex activity on straight reaches substantially better than curved reaches ($p < 0.001$, Wilcoxon Signed Rank test). **c**, Multi-unit PSTHs and reconstruction by jPCA rotations and mean activity across conditions. Rotations captured little of the multi-phasic modulation of neural activity during reaching. Traces colored by target angle.

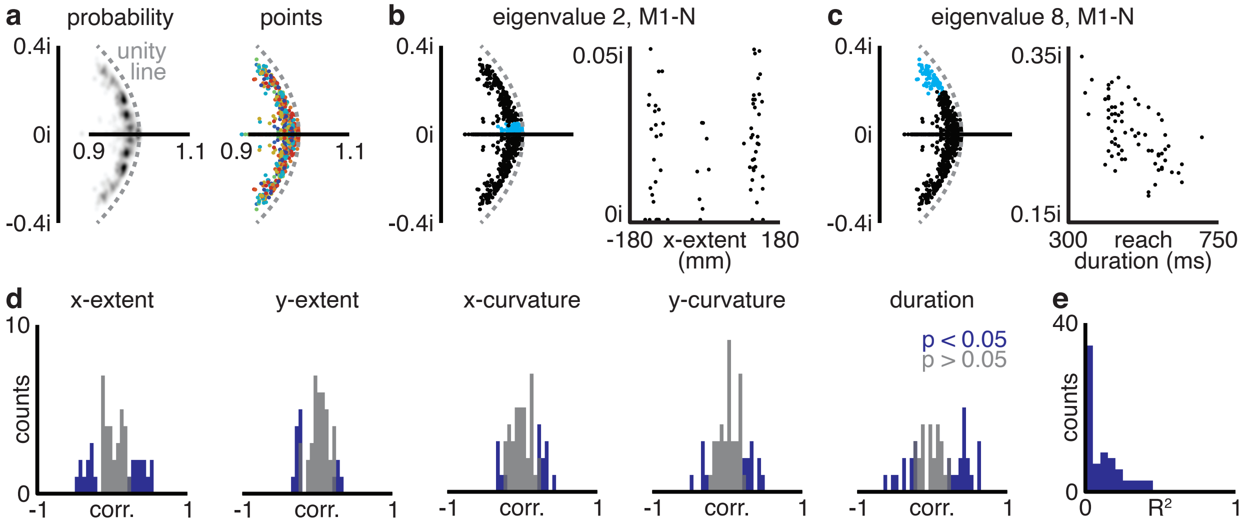

Supplementary Figure 2: Eigenvalues of motor cortical dynamics correlate modestly with reach parameters. **a**, Probability distribution of eigenvalues for single-condition dynamics (*left*) and eigenvalues colored by target angle (*right*). **b**, Second eigenvalues of single-condition dynamics (blue, *left*), and x-extent of reach vs. the imaginary component of the second eigenvalue (*right*). No obvious relationship is present. **c**, Eighth eigenvalues (corresponding to the highest-frequencies that were reliable) of single-condition dynamics (blue, *left*), and reach duration vs. the imaginary component of the eighth eigenvalue (*right*). Data in **a-c** taken from M1-N. **d**, Histogram of correlation coefficients of reach parameters with each real and imaginary component of each eigenvalue. **e**, Histogram of R^2 for linear models predicting real and imaginary components of eigenvalues from reach parameters. Data in **d**, **e** is pooled across datasets.

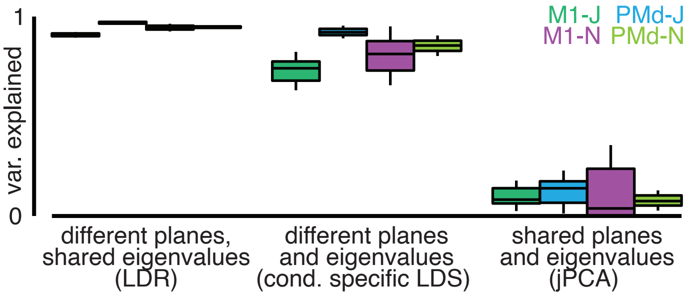

Supplementary Figure 3: Different planes with shared eigenvalues best fit motor cortex activity. We compared models of motor cortex dynamics in which rotational frequencies were held constant between conditions while rotational planes were allowed to vary between conditions (the LDR model; *left*), rotational frequencies and planes were allowed to vary between conditions (*center*), and rotational frequencies and planes were conserved between conditions (*right*). Note that while the “different planes, shared eigenvalues” (LDR) model is a subset of the “different planes and eigenvalues” (condition-specific LDS) model, the “different planes, shared eigenvalues” explains greater neural variance when cross-validated. This improvement has two sources. First, as this model assumes that eigenvalues are shared across conditions, it gets to estimate rotational frequencies using every condition, leading to more stable estimates of eigenvalues. Second, LDS models predict the population state’s derivative from the state, which does not directly optimize variance explained. The “different planes, shared eigenvalues” model, on the other hand, directly optimizes the variance explained.

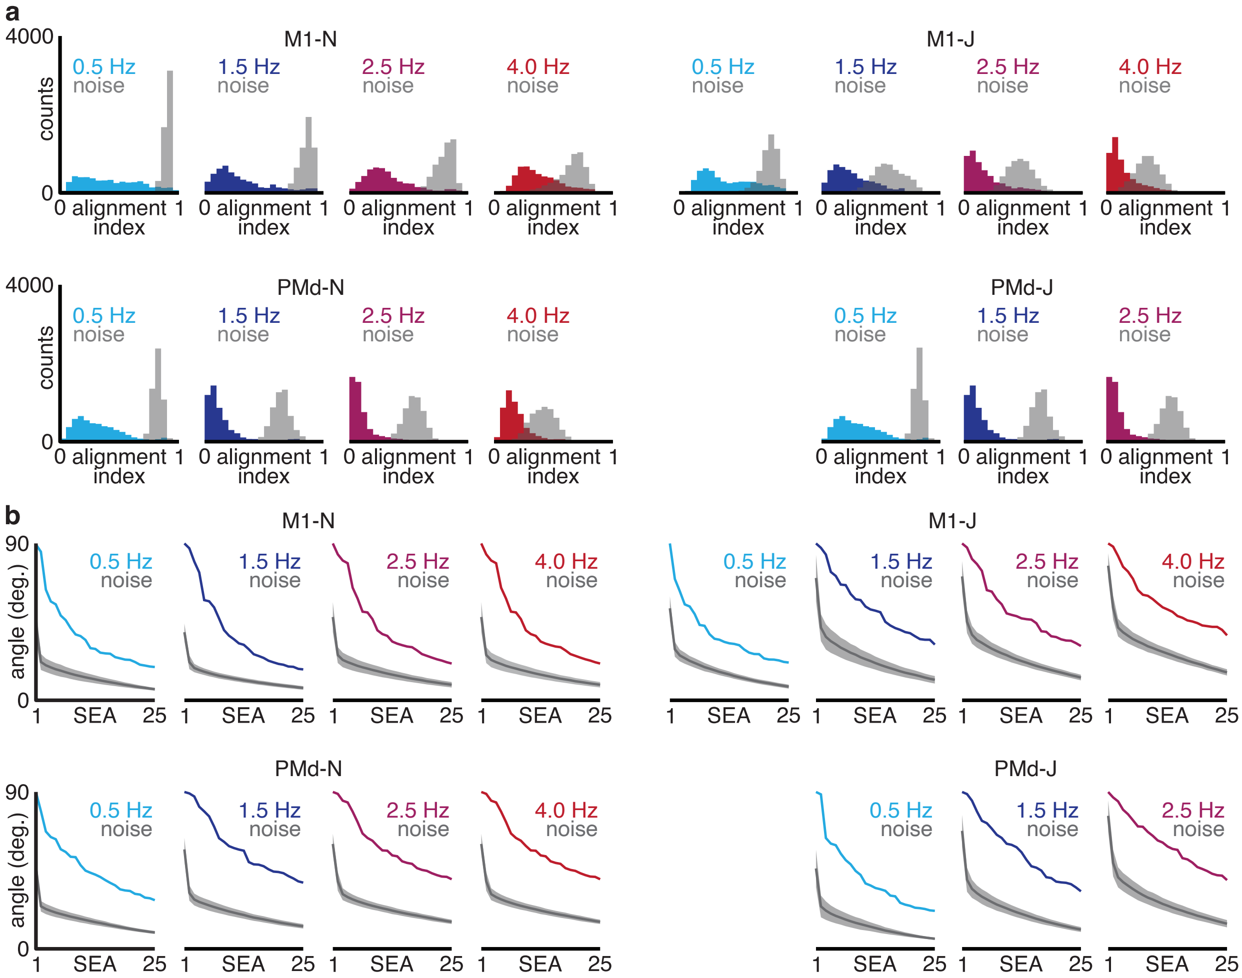

Supplementary Figure 4: Rotational planes differ between conditions. a, Alignment indices between corresponding (same-frequency) rotational planes on pairs of conditions (*colored*) and distribution expected due to estimation noise (*gray*). **b**, Subspace Excursion Angles for corresponding rotational planes across conditions (*colored*), along with angles expected due to estimation noise (*gray*; Line, mean; shaded, 1 standard deviation).

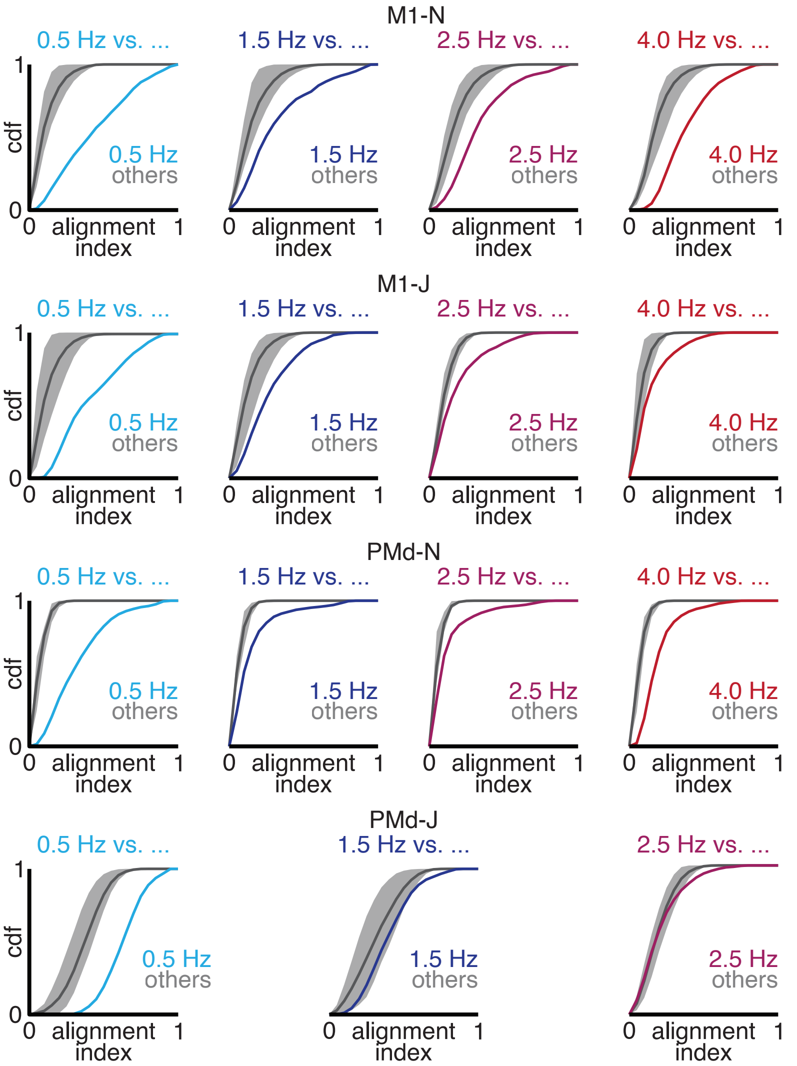

Supplementary Figure 5: Rotations occupied distinct subspaces in neural state space. Cumulative distribution function of alignment index between different-frequency rotational planes (*gray*) and same-frequency rotational planes (*colored*) on pairs of conditions. Same-frequency rotational planes were more aligned than different-frequency rotational planes.

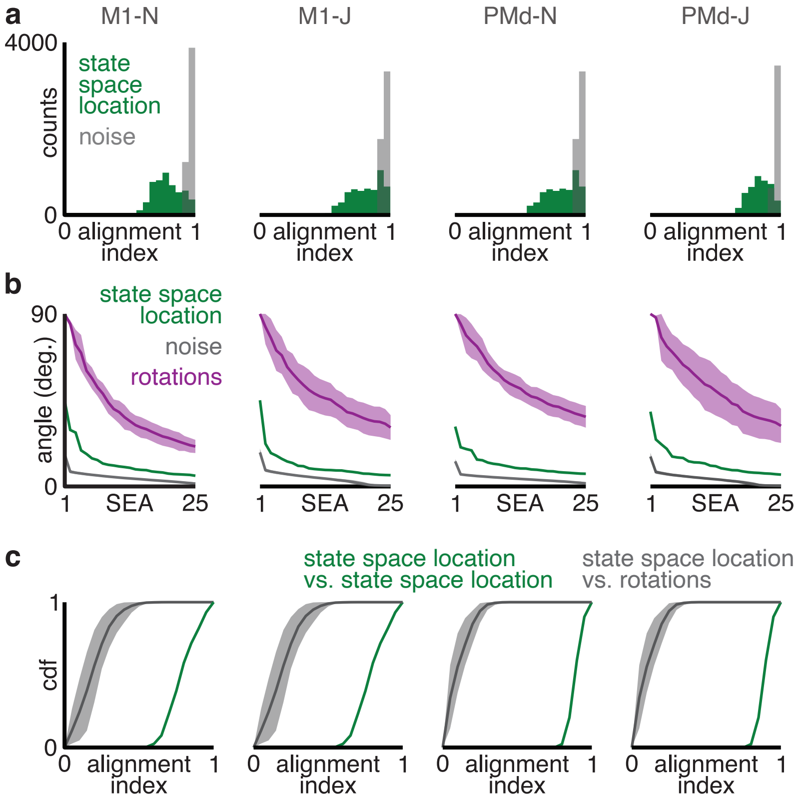

Supplementary Figure 6: State space location differed between conditions. a, Alignment indices between state space locations on pairs of conditions (*green*) and distribution expected due to estimation noise (*gray*). **b**, Subspace excursion angles for state space locations across conditions (*green*), along with angles expected by estimation noise (*gray*) and angles formed by rotations (*purple*). Line, mean; shaded, 1 standard deviation. **c**, Cumulative distribution function of alignment index between rotations and state space locations (*gray*) and state space locations compared to state space locations (*green*) on pairs of conditions. State space locations were more aligned with themselves than with rotations.

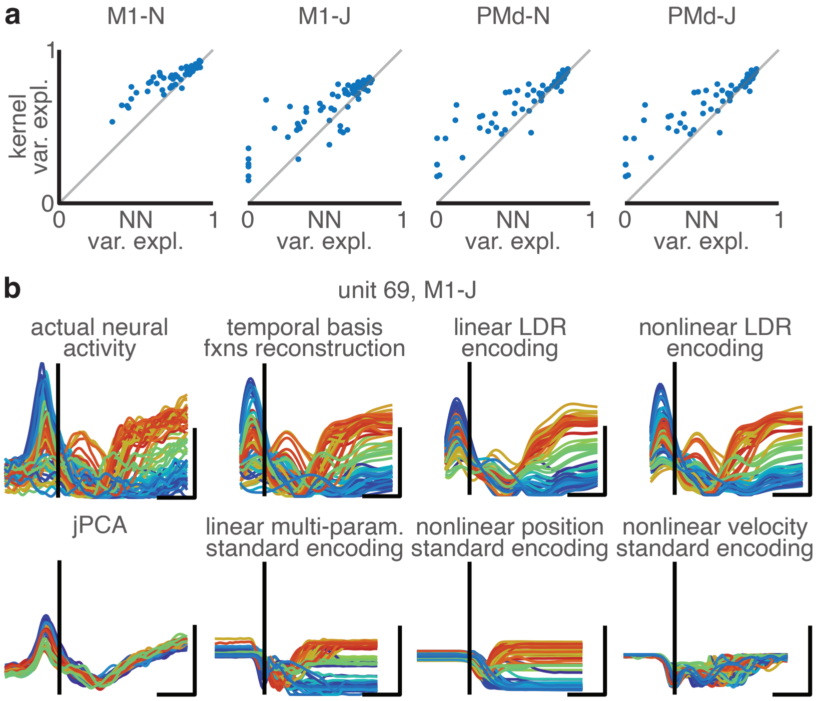

Supplementary Figure 7: Reach kinematics are sufficient to predict motor cortex dynamics. a, To demonstrate that LDR encoding used continuous features to determine test condition dynamics, rather than simply memorizing the training set, we compared kernel regression (a continuous model) with nearest-neighbor regression (which memorizes the training set). Conditions were sorted by target angle. To make the (six-fold) cross-validation more challenging, for each fold we held out conditions with targets in the same sextant. The higher performance of kernel regression indicates that nonlinear LDR encoding was not merely memorization of the training set. **b**, Example unit for the 7 models of neural activity fit in this study.

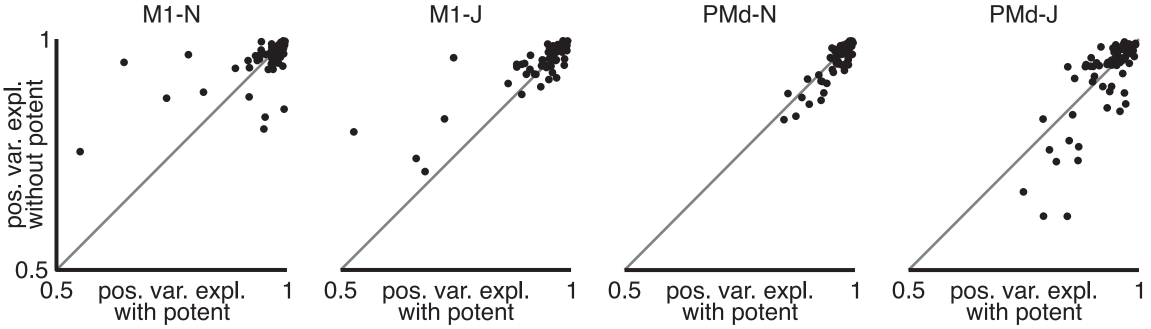

Supplementary Figure 8: LDR-based decoding does not rely on output-potent dimensions exclusively. Variance explained in hand position using LDR-based decoding with and without output-potent activity. Each dot is a condition. LDR-based decoding still performed well even after excluding output-potent dimensions.

Supplement

Basis function representation of a diagonalizable linear dynamical system

Linear dynamical systems are defined by a transfer matrix, **M**, that maps the state of the dynamical forward in time. This can be defined in continuous time as

$$(1)\quad d\mathbf{x}/dt = \mathbf{M}\mathbf{x}$$

so that the derivative of the state is just a linear transformation on the current state. From some initial state $\mathbf{x}_0$, the state at time t can be expressed as

$$(2)\quad \mathbf{x}(t) = \mathrm{mexp}(\mathbf{M}t)\mathbf{x}_0$$

where mexp() refers to the matrix exponential. In this section, we will work in continuous time, though our argument generalizes to discrete time naturally. Under a few assumptions, the transfer matrix **M** can be diagonalized

$$(7)\quad \mathbf{M} = \mathbf{E}\mathrm{diag}(\mathbf{d})\mathbf{E}^{-1}$$

where **E** is the eigenvectors of the dynamics, **d** is a vector of the eigenvalues, and diag() returns a diagonal matrix with the argument along the diagonal. If **M** is diagonalizable, then the state at time t can be expressed more simply as

$$(8)\quad \mathbf{x}(t) = \mathbf{E}\exp(\mathrm{diag}(\mathbf{d})t)\mathbf{E}^{-1}\mathbf{x}_0$$

where exp() is the exponential function applied to each element in its argument. Define the vector

$$(9)\quad \mathbf{x}_{\mathbf{e}0} = \mathbf{E}^{-1}\mathbf{x}_0$$

as the initial state in the eigenvector basis. Using that, for two vectors **z** and **y** of identical dimensions,

$$(10)\quad \mathrm{diag}(\mathbf{z})\mathbf{y} = \mathrm{diag}(\mathbf{y})\mathbf{z}$$

then (8) can be re-written as

$$(11)\quad \mathbf{x}(t) = \mathbf{E}\mathrm{diag}(\mathbf{x}_{\mathbf{e}0})\exp(\mathbf{d}t)$$

so that exp($\mathbf{d}t$) is now a column-vector. By further defining the terms

$$(12)\quad \mathbf{b}(t) = \exp(\mathbf{d}t)$$

$$(13)\quad \mathbf{L} = \mathbf{E}\mathrm{diag}(\mathbf{x}_{\mathbf{e}0})$$

(11) can be re-written as

$$(14)\quad \mathbf{x}(t) = \mathbf{L}\mathbf{b}(t)$$

so that the state at time t can be expressed as the product of loading matrix and a set of temporal basis functions. This factorization isolates the effect of the eigenvalues in determining the “temporal” characteristics of the population into **b**. **L**, by contrast, expresses both the effect of the eigenvectors of determining “where” in state space dynamics occur, and the initial state in setting the phases and magnitudes of the dynamics.

Fitting motor cortex activity with conserved rotations, but varying rotational planes

In this study, we required a method to identify a model of motor cortex activity where rotations in motor cortex activity were conserved in their temporal characteristics (identical eigenvalues), but occupied different planes on different conditions (different eigenvectors). In this model, motor cortex activity on single conditions is an LDS. This system is, however, nonlinear across conditions, thus disallowing simply regressing the state’s derivatives against the state. We exploit the basis-function representation of LDSs to work around this issue. As conditions share eigenvalues, they share temporal basis functions. This means that two conditions, *i* and *j*, have states as a function of time *t*

$$(15)\ \mathbf{x}^{(i)}(t) = \mathbf{L}^{(i)}\mathbf{b}(t)$$

$$(16)\ \mathbf{x}^{(j)}(t) = \mathbf{L}^{(j)}\mathbf{b}(t)$$

$$(17)\ \mathbf{L}^{(j)} = \mathbf{E}^{(j)}\mathrm{diag}(\mathbf{E}^{(j),-1}\mathbf{x}^{(j)}{}_{0})$$

where $\mathbf{x}^{(j)}(t)$ and $\mathbf{x}^{(i)}(t)$ is the firing rates at time *t* on conditions *j* and *i* respectively, $\mathbf{E}^{(j)}$ is the matrix of eigenvectors of the transfer matrix for condition *j*, and $\mathbf{x}^{(j)}{}_{0}$ is the initial state on condition *j*. Usefully, while the loading matrices in (15) and (16) differ, the temporal basis functions are conserved between conditions because the eigenvalues are conserved. This means that activity on each condition is just a different linear transformation on the same temporal basis functions. Fitting this model therefore reduces to identifying an optimal set of temporal basis functions that explain neural activity. These can be obtained by the temporal singular vectors of the neural activity concatenated into a (*neurons* x *conditions*)-by-*time* matrix, or equivalently the eigenvectors of the Gramian matrix over time. This is comparable to similar operations over neurons that are used to extract the dimensions of highest variance in neural state space.