

Supplementary Figure 1: Feedback and spinal cord model variants.

In the main manuscript, we focused on a single recurrent model that directly output EMG. However, the interaction of the motor cortex with other systems is central to the production of movement. To determine whether results depend on the exact form of our simplified model, we tested additional models of two varieties. First, we tested models that incorporate (simulated) sensory feedback. Second, we tested models that incorporate transformations (e.g., spinal circuitry) between motor cortex and muscle activity. (A-D) State space visualizations for monkey J for 4 model variants (same conventions as Figure 6). Each model has a similar architecture to the regularized dynamical model, but with the additions detailed below. (E-F) Summary canonical correlations of model variants and neural data for monkey J. Model details are given below, but the key observation is that the regularized model population response continued to resemble the neural data regardless of the addition of simulated feedback or downstream transformations. This is reflected both in the oscillations about a single fixed point (A-D) and the high canonical correlations. Why does this similarity persist? In the case of feedback, that feedback simply becomes part of the overall dynamics. Once trained in the presence of feedback the network depends on that feedback: its removal would have a profound effect. Yet the overall dynamics remain similar for a straightforward reason: the cost function is still encouraging dynamics that produce the empirical EMG in the simplest way possible. That solution remains similar regardless of whether the dynamics depend on feedback or not. In the case of downstream transformations, the network continues to find a similar solution so long as the ‘basic building blocks’ of EMG remain similar. Consider the case where the model is asked to output muscle synergies (from which EMG can be composed) rather than EMG directly. This constraint changes the model very little: it was already attempting to produce EMG as simply as possible and was thus already encouraged to exploit the redundancy across muscles.

Low-pass Feedback Model Variant We wished to incorporate reasonable sources of feedback into the model. In the absence of afferent recordings, we assumed that feedback is a temporally transformed (filtered) version of the output command. There are many possible such transformations, but we concentrated on two. First, we assumed that sensory feedback is proportional to muscle force / stretch, which is a low-pass filtered version of the electrical activation. To emulate this we filtered the EMG with a low-pass filter having the following parameters: 100-order minimum phase FIR filter with 1Hz pass band with a 2Hz stop-band of 10db attenuation. The filter “group delay” of the resulting filtered EMG was 140ms at 0.5Hz, testing larger delays of peripheral feedback to cortex. The filtered EMG was used as additional input to the model. The remaining parameters were identical to those of the respective regularized dynamical model and EMG fits were just as good (normalized error 7%).

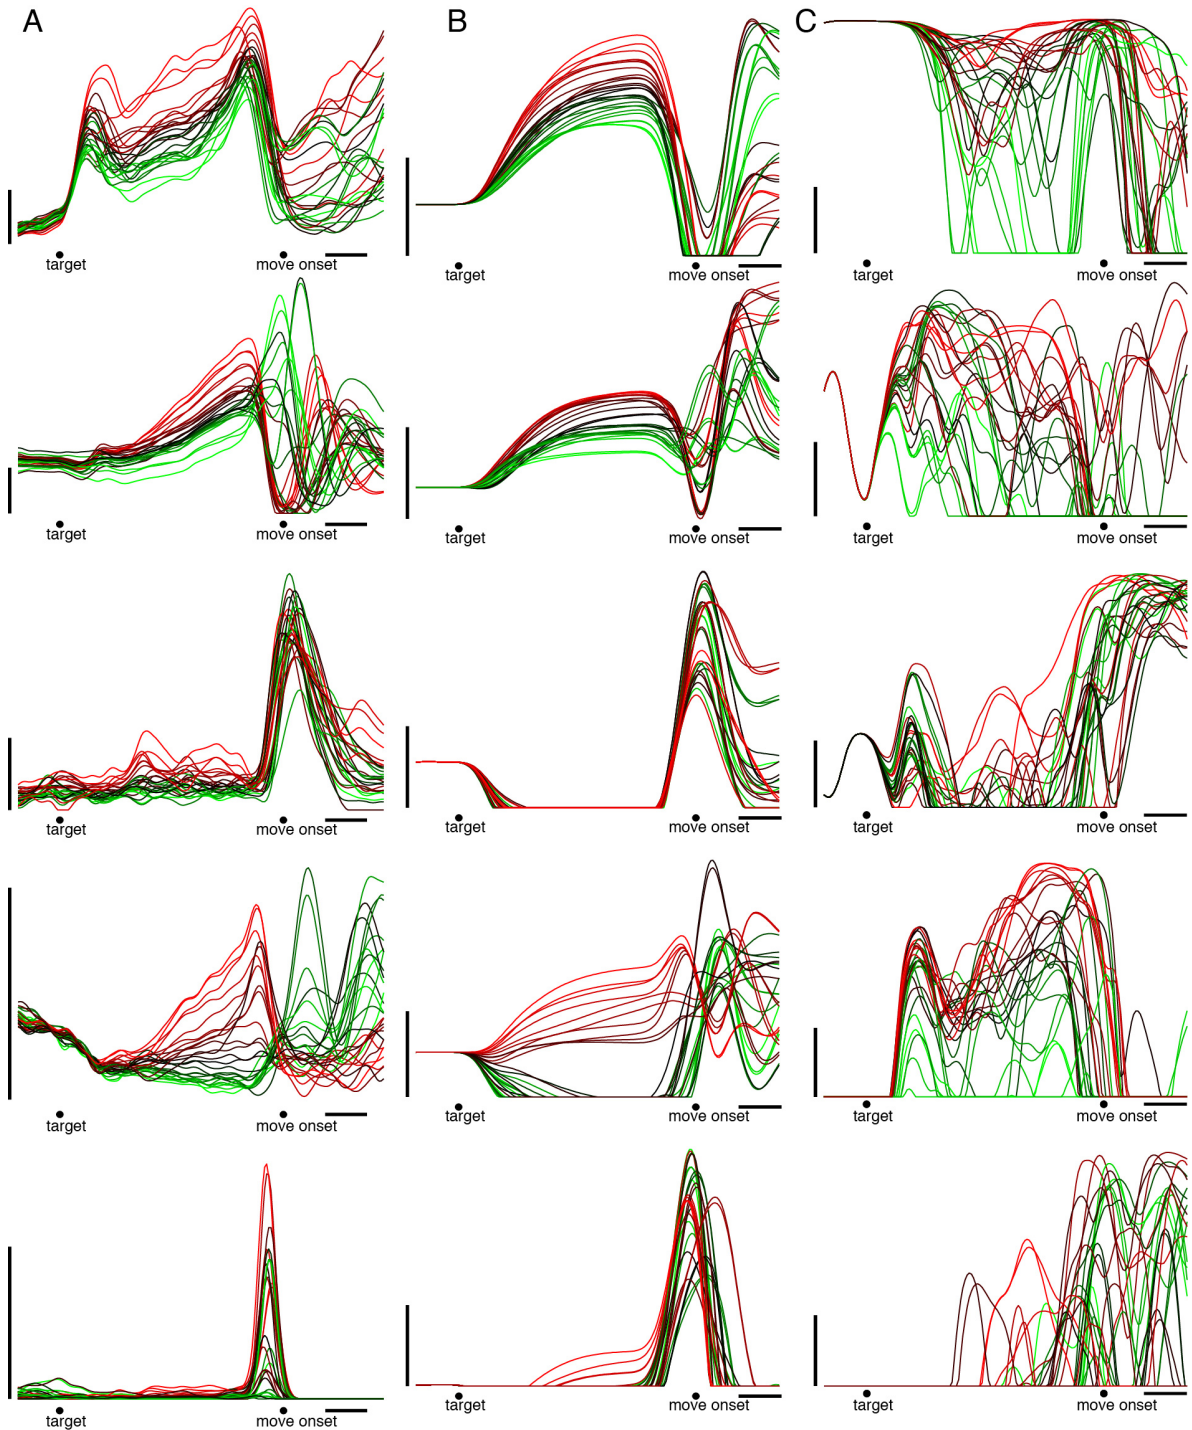
High-pass Feedback Model Variant Another reasonable hypothesis is that sensory feedback is selectively sensitive to *changes* in muscle force, in which case it should be modeled as a high-pass filtered version of the EMG activity. We therefore filtered the EMG with a high-pass filter having the following parameters: 126-order minimum phase FIR filter with a 2Hz pass band and a 1Hz stop-band with 20db attenuation. The filter “group delay” of the resulting filtered EMG was 25ms at 3Hz, testing smaller delays of peripheral feedback to cortex. Simulations were otherwise as described above for the low-pass model. Fits to the EMG remained excellent (normalized error 7%).

Muscle Synergies Model Variant - The muscle synergies model variants were trained to output the top P PCA projections of the EMG for each monkey. Specifically, for both monkeys the following steps were followed: PCA was performed on the $M \times CT$ matrix of the mean-centered EMG. The projections of the EMG onto the resulting top P PCs were used as training data for the regularized models. All inputs were the same between the muscle synergy models and the regularized dynamical models. All other training parameters were the same between the muscles synergy variant and the respective regularized dynamical models. The value for P was set to 4 for monkey J and 5 for monkey N. Shown in C is a model that is optimized to generate muscle synergies instead of EMG (normalized error 5%).

Spinal Cord Model Variant - The spinal cord models used two RNNs. The first RNN simulated the motor cortex and was trained to generate the top 3 muscle synergies. These muscle synergies were computed in the same way as for the muscle synergy model variants, except that the number of synergies was 3 for both monkeys. A second network then used these outputted muscle synergies of the simulated motor cortex models to dynamically generate the full M -dimensional EMG. Aside from producing the 3 muscle synergies as output, the parameters for the simulated motor cortex RNN were identical to those of the regularized dynamical model in the main paper. The parameters for the simulated spinal cord model for both monkeys were as follows: # inputs = 3, $N = 100$, $\alpha = 2e-6$, $\beta = 0.015$, and $\gamma = 2e-6$. The simulated motor cortex RNN is analyzed in Supp. Figure 1D. Shown in D is a model that is optimized with an additional spinal cord module, which receives muscle synergies and outputs EMG. The state-space plot shows the M1 module activity of the M1-spinal cord system (normalized error 2%).

(E) Summary canonical correlations between models shown in panels A-D and neural data. All model variants have as high or higher mean canonical correlation with the data as regularized dynamical model.

(F) The mean CC with the neural data for the respective models is 0.75, 0.76, 0.78, and 0.74.



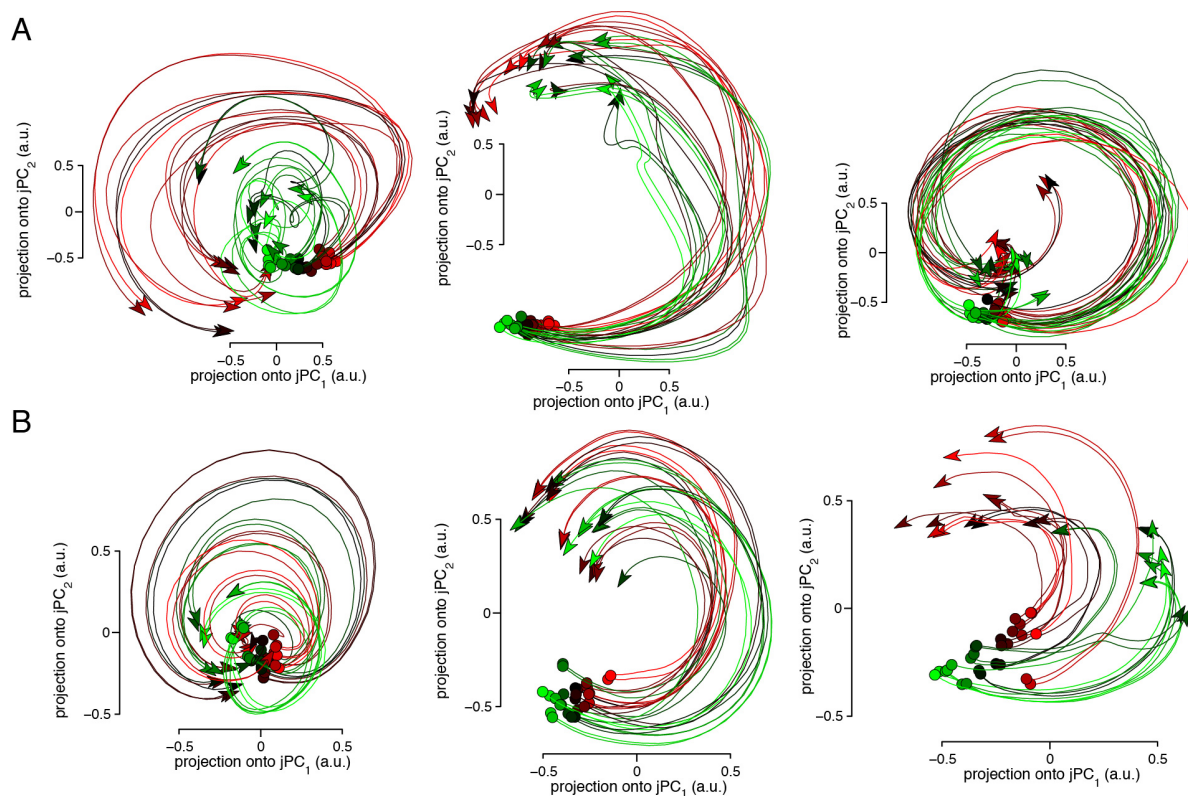
Supplementary Figure 2: Example PSTHs from monkey N, and the regularized and complicated models for monkey N.

(A) Example PSTHs from five neurons from monkey N (format as in Figure 1B). These examples were chosen to illustrate the range of responses, including neurons with strong preparatory activity (rows one and four), neurons with a broad rise in activation during the movement period (middle row), neurons with oscillatory activity during the movement

period (rows two and four), a neuron with a narrow bump of activation (bottom row). Vertical and horizontal scale bars indicate 20 spikes per second and 200ms.

(B) Example PSTHs chosen from the regularized model for monkey N. Examples were chosen to highlight the similarities between neural and model responses, and are representative of the patterns exhibited by the model units.

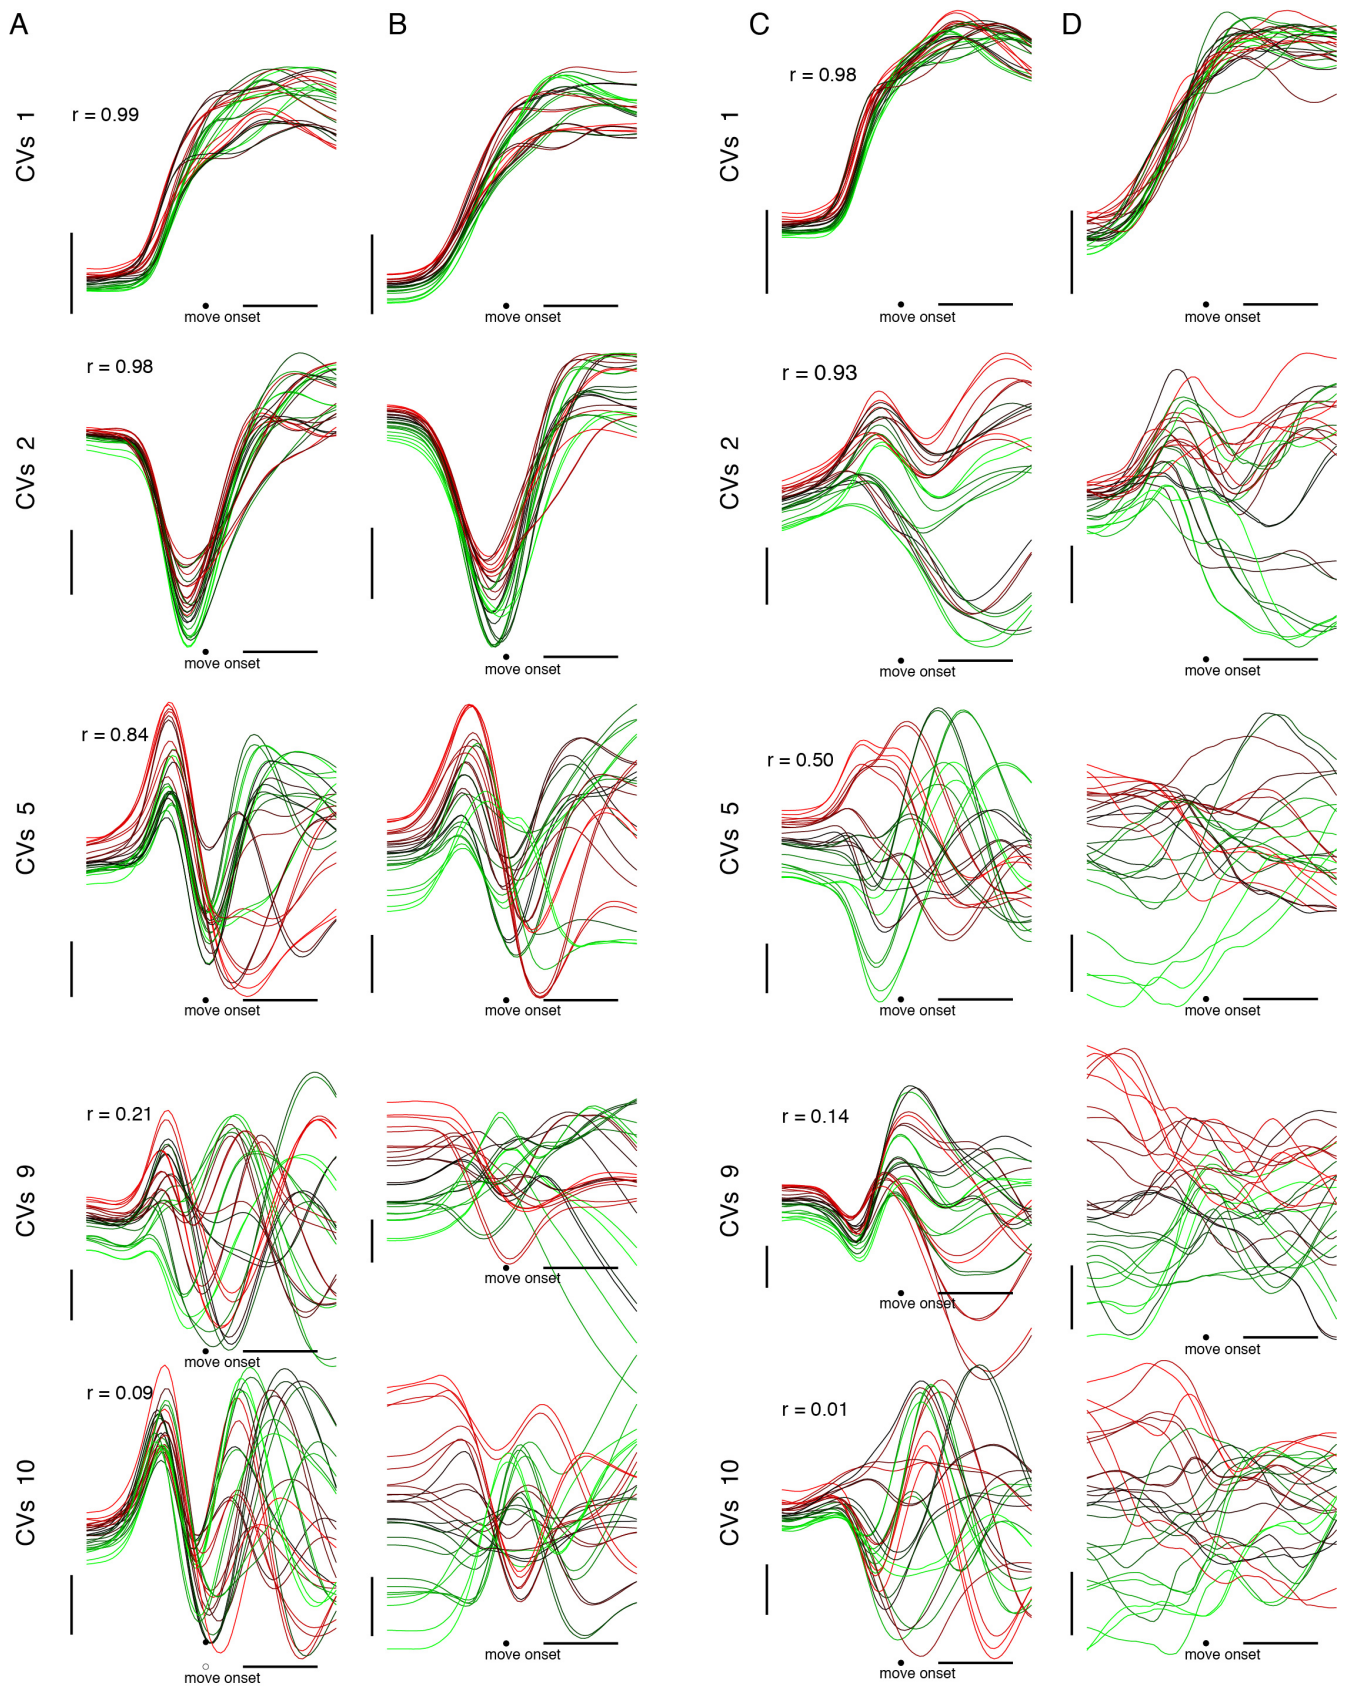
(C) Example PSTHs from five units from the complicated model for monkey N. The PSTHs of the complicated models did not bear any obvious resemblance to those of the neural data. Vertical and horizontal scale bars indicate 20 spikes per second and 200ms.



Supplementary Figure 3: jPCA projections of the population responses for monkey N and the regularized model for monkey N.

(A) jPCA projections of the neural population response for monkey N. Each trace shows the evolution of the neural state over 500ms. Traces start -180ms before movement onset; at the moment the relatively stable preparatory state (circles) transitioned to the movement period trajectory (same conventions as Figure 3). The three projections correspond to the largest magnitude complex eigenvalue pairs of the matrix \mathbf{M}_{skew} , found when fitting the data with $\dot{\mathbf{x}} = \mathbf{M}_{\text{skew}}\mathbf{X}$ (Experimental Procedures). These eigenvalues correspond to frequencies of 2.6Hz, 1.6Hz, and 1.2Hz (left to right) with a quality of fit (R^2) for the optimal purely oscillatory linear system of 0.77.

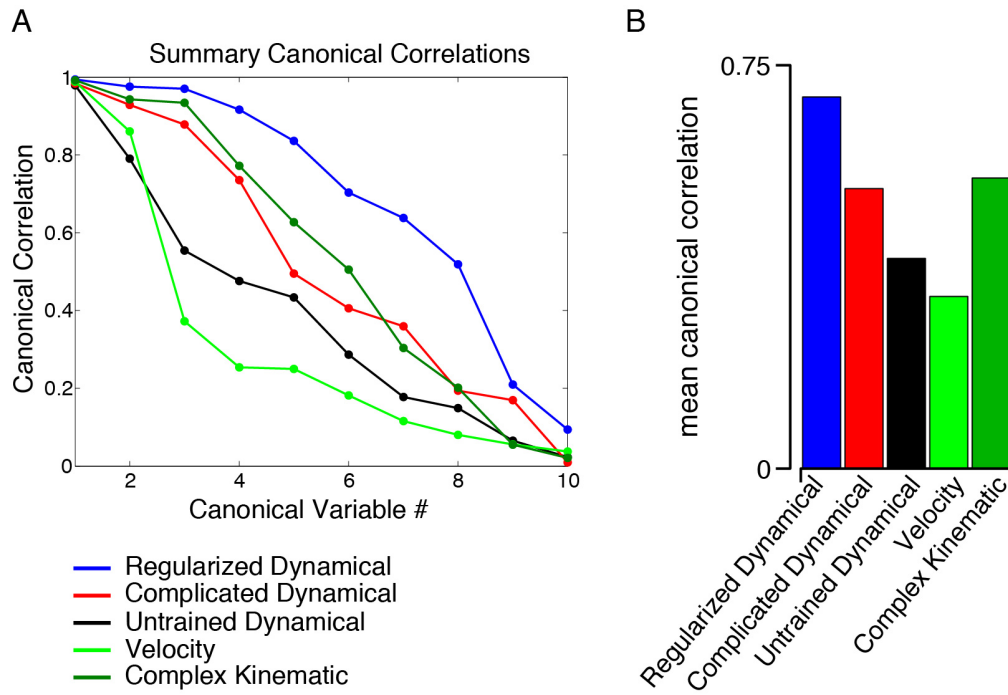
(B) jPCA projections for the regularized model of monkey N. Same conventions as in panel A. Frequencies are 2.2Hz, 1.3Hz, and 0.8 Hz. The linear system, $\dot{\mathbf{x}} = \mathbf{M}_{\text{skew}}\mathbf{X}$ has a quality of fit (R^2) of 0.79.



Supplementary Figure 4: Canonical correlations analysis for monkey N.

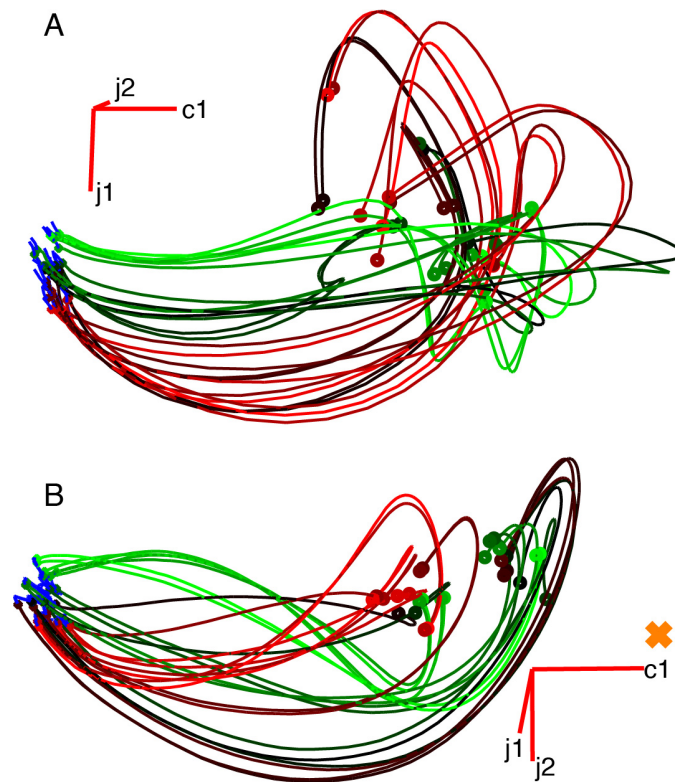
(A, B) CCA projections (canonical variables) of the neural population response (A) and the regularized model for monkey N (B). In each row is shown the canonical variables 1, 2, 5, 9, and 10 out of 10 (conventions as in Figure 4). The vertical scale bar indicates 1 arbitrary unit while the horizontal scale bar indicates 200ms.

(C, D) Canonical variables of the neural population response (C) and the complicated model for monkey N (D) (format as in A and B).



Supplementary Figure 5: Comparison of simulated and neural population responses for monkey N.

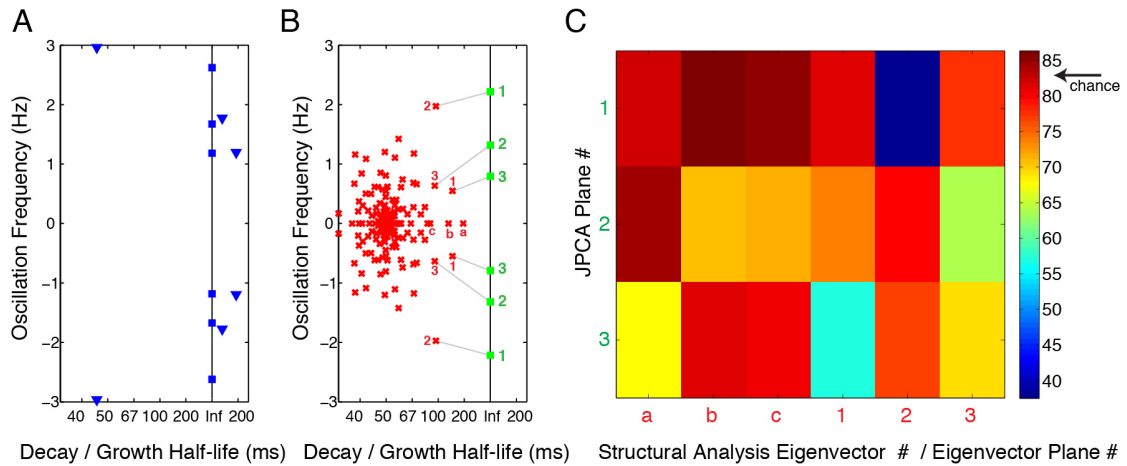
(A) Summary Canonical Correlations (same conventions as in Figure 5.) (B) The average of the canonical correlations between the models and the data.



Supplementary Figure 6: Monkey N and regularized model state-space visualizations.

(A) 3D visualization of the neural data during the movement period for monkey N. The projection is comprised of the first j PC plane (Supp. Figure 3A, left panel) and an additional dimension that captures variance from the cross-condition mean. Each trace is color-coded to show one of the 27 reach conditions. For all conditions the trajectory during the preparatory period is colored blue. Time shown is 400ms before to 220ms after movement onset.

(B) Analogous 3D visualization of the regularized model for monkey N (same conventions as in panel A). Additionally, the single, condition-independent fixed point of the model, which organizes the dynamics of movement generation, is shown with an orange 'x'. Time shown is 800ms before to 220ms after movement onset.

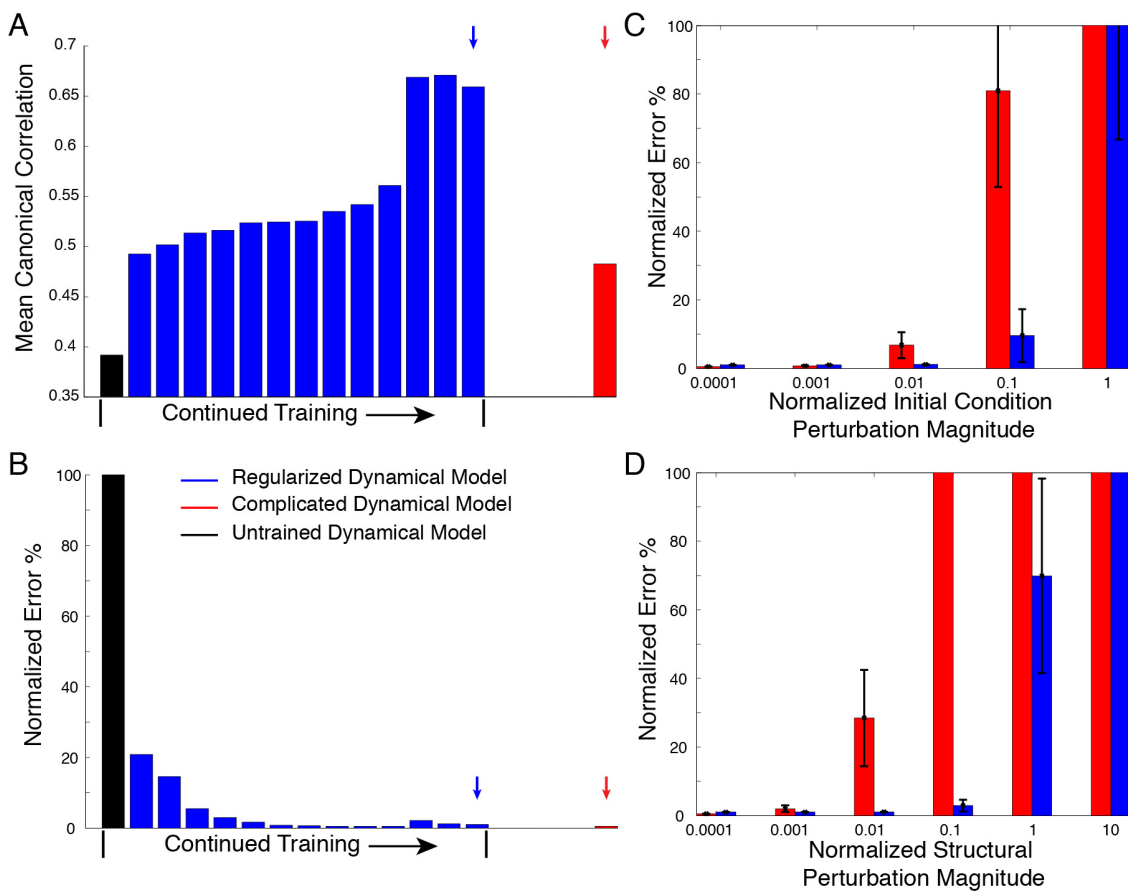


Supplementary Figure 7: Frequency analysis of neural data and regularized model for monkey N.

(A) Eigenvalue analysis of the neural data. Shown on the line of stability ('Inf' – neither decaying nor growing) are the purely imaginary eigenvalues associated with the jPCA analysis of the neural data in Figure 3A (blue squares). Also shown are the top eigenvalues of an unconstrained linear fit to the neural data (blue triangles).

(B) The complex eigenvalue spectrum of the linearized system around the fixed point in the regularized model for monkey N (red 'x') based on a structural analysis of the weight matrix. Highlighted with red numbers / letters are those modes of the linearized system that have a slow decay. Shown along the line of stability ('Inf') are the purely imaginary eigenvalues associated with the jPCA analysis of the regularized model data (green squares). Grey lines show the connection between the jPCA analysis and the structural analysis, as given by a subspace angle analysis of eigenvectors in panel C.

(C) On the vertical axis are listed the 3 oscillatory planes found by jPCA (corresponding to the green numbered modes in panel B). Color indicates the minimum subspace angle (the minimum angle between the corresponding planes). For comparison, the minimum subspace angle between two randomly chosen planes in a $N = 300D$ space is 84 ± 2 degrees (mean and std dev., black arrow labeled 'chance'). Thus, a minimum subspace angle of 30 to 40 degrees indicates highly overlapping subspaces. In the present case jPC plane 1 overlaps heavily with mode 2 (the fastest frequency), jPC plane 2 overlaps most strongly with oscillatory mode 3 and jPC plane 3 overlaps most strongly with oscillatory mode 1.



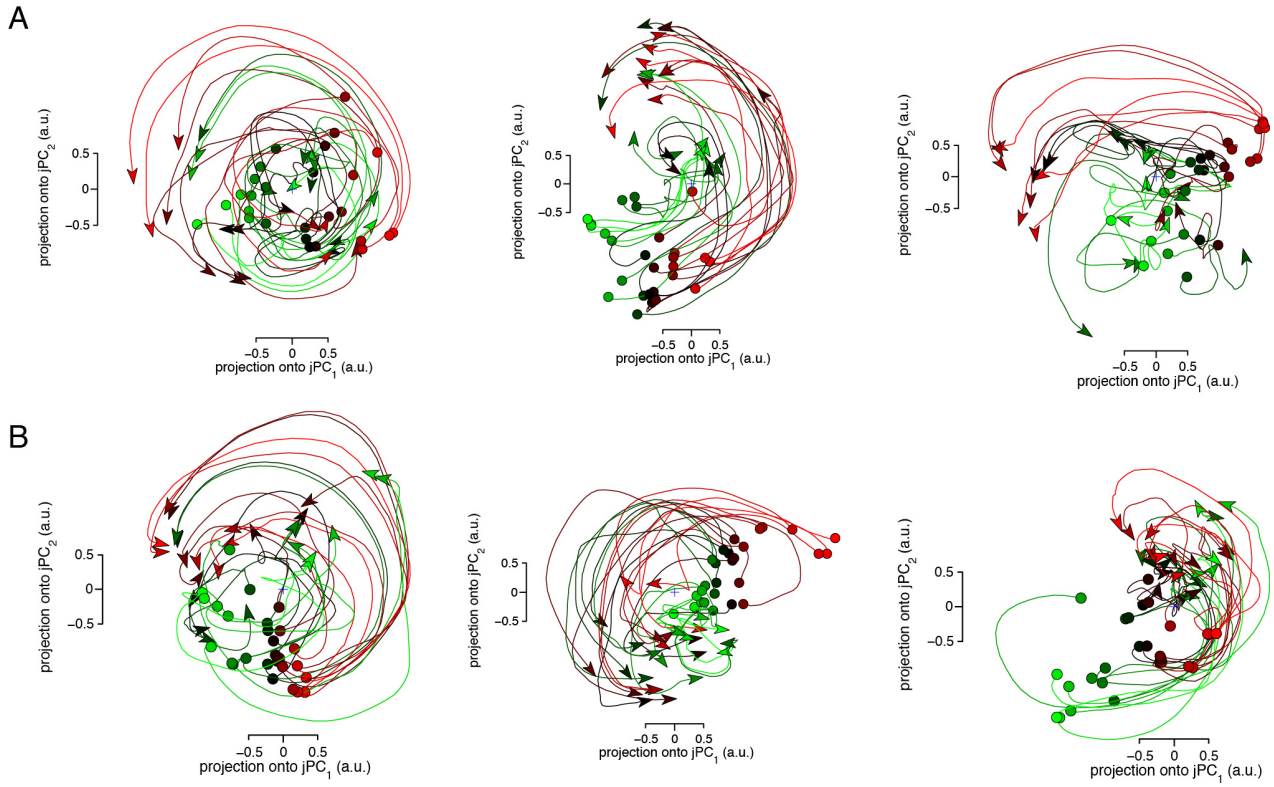
Supplementary Figure 8: Regularization impacts similarity to data and model robustness (monkey N).

(A) Average canonical correlation, as training progresses, between the regularized model and the neural data from monkey N. Black shows the mean canonical correlation of 0.39 between the neural data and the data from the untrained model with correct inputs injected. This provides a baseline for comparing mean canonical correlation. As training goes forward with heavy regularization (blue) the model becomes more and more similar to the neural data, ending with a mean canonical correlation of 0.65 for this model (blue arrow). When trained to generate EMG without any regularization, the model has a mean canonical correlation with the data of 0.48 (red arrow).

(B) The normalized error of the network output as a function of training for the untrained regularized and complicated models (same conventions as in panel A). Note that the error decreases very quickly for the regularized model, while in panel A, the mean canonical correlation continues to increase over a much longer period of training, demonstrating that simplicity also makes the model more similar to the data, and not simply fitting the EMG.

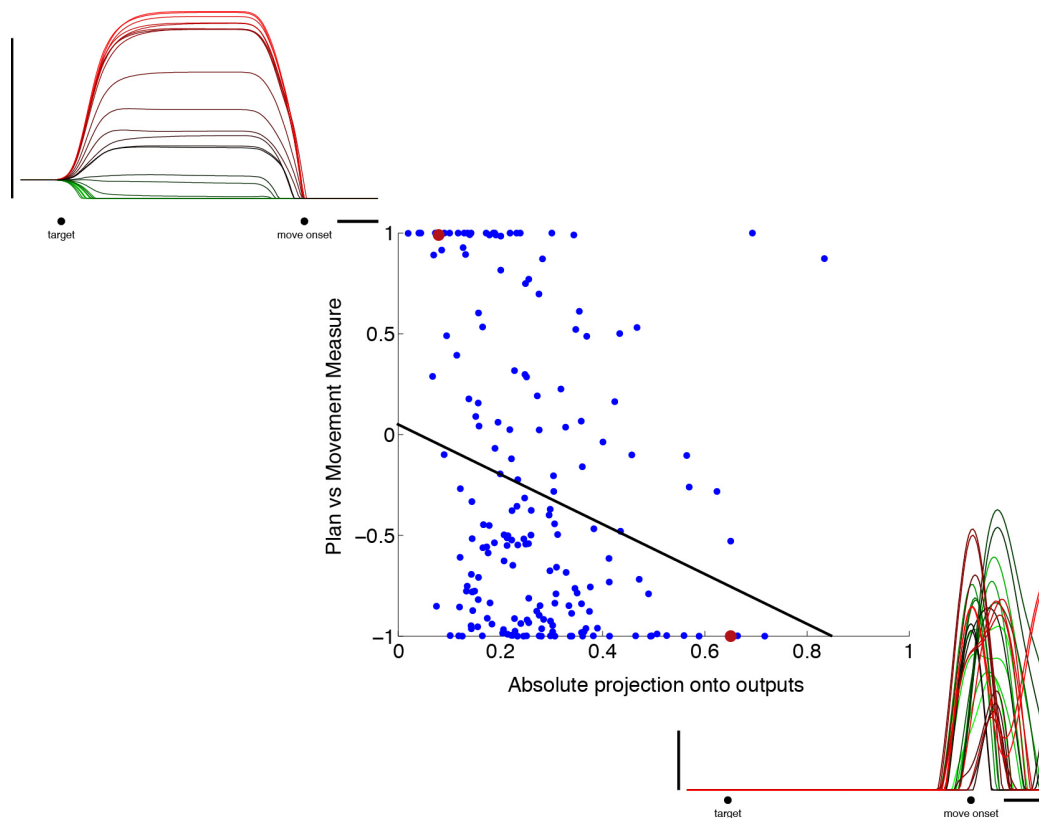
(C) Perturbation test of the initial conditions for the regularized and complicated models analyzed in panel A (blue and red arrows, respectively). The inputs were randomly perturbed according to a normalized percentage of the input strength (as given on horizontal axis) 50 times. The network was then run and the mean normalized EMG error of the outputs was averaged (vertical axis). Error bars show standard deviation. The vertical axis is truncated at 100% error.

(D) A structural perturbation test of the recurrent connectivity matrix in equation (1) for the regularized and complicated models analyzed in panel A (blue and red arrows, respectively). The connectivity matrix was randomly perturbed 50 times according to a normalized percentage of the mean absolute connection strength (as given on horizontal axis). The perturbed network was then run and the mean normalized EMG error of the outputs was averaged (vertical axis). Error bars show standard deviation. The vertical axis is truncated at 100% error.



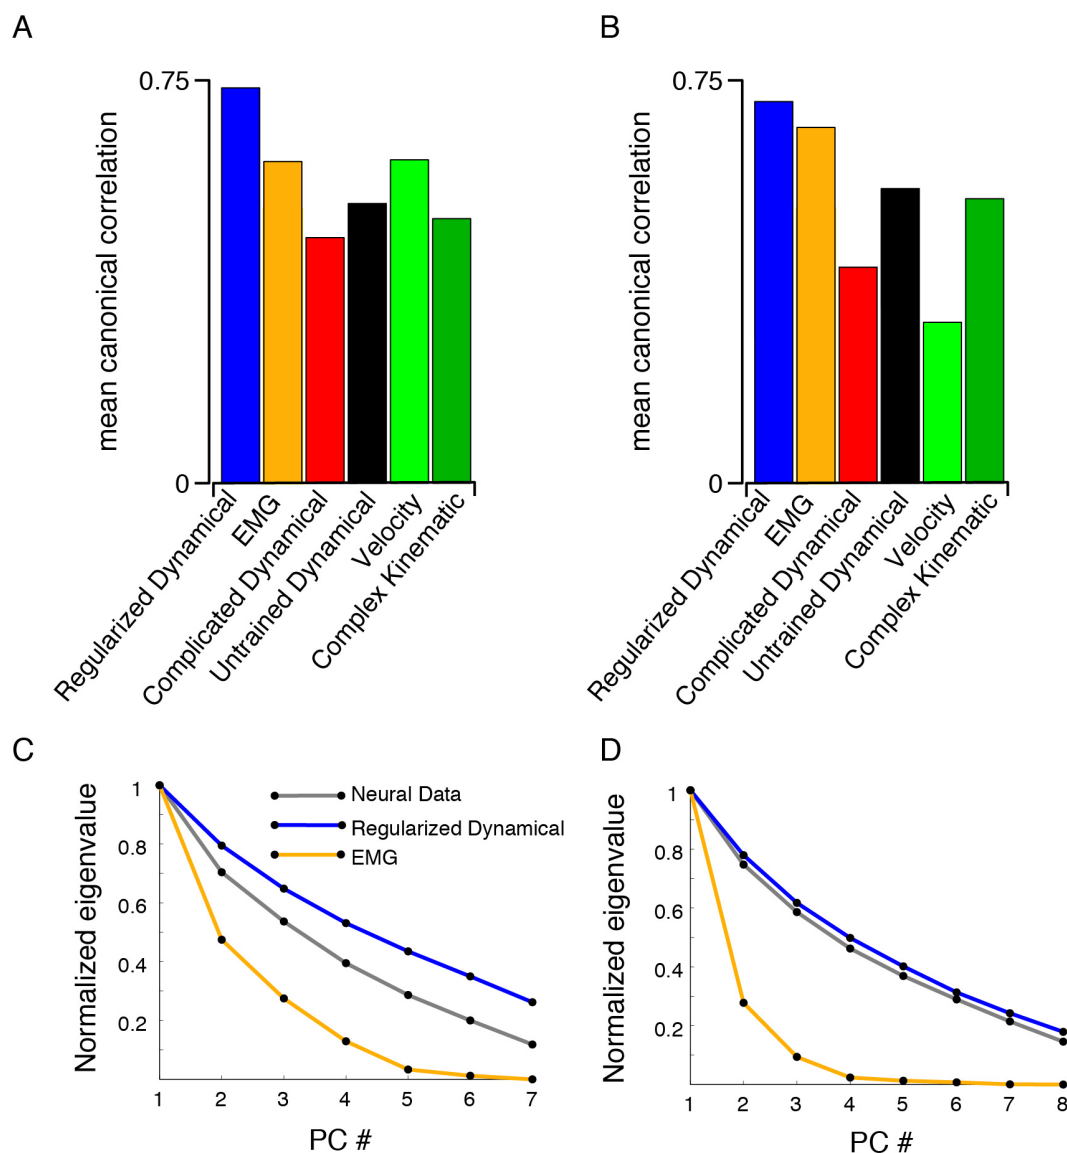
Supplementary Figure 9: jPCA projections for the complicated models of monkey J and monkey N.

(A) jPCA projections for the complicated model of monkey J. Same conventions as in Figure 3. Frequencies are 1.3Hz, 0.8Hz, and 0.6Hz. The linear system, $\dot{\mathbf{x}} = \mathbf{M}_{\text{skew}}\mathbf{X}$ has a quality of fit (R^2) of 0.35.
 (B) jPCA projections for the complicated model of monkey N. Frequencies are 1.2Hz, 0.9Hz, and 0.8Hz. The linear system, $\dot{\mathbf{x}} = \mathbf{M}_{\text{skew}}\mathbf{X}$ has a quality of fit (R^2) of 0.37.



Supplementary Figure 10: The relative strength of preparatory activity to movement activity correlates with model unit output strength.

There is a well known anterior to posterior gradient in PMd / M1: preparatory activity is strongest in more anterior regions, while more posterior regions are more strongly dominated by movement-period activity. This gradient runs opposite another well-known gradient: more posterior regions have stronger and more direct connections to the spinal cord. While the model was not designed with such an anatomical gradient, it is nevertheless reasonable to inquire whether the units that are more strongly 'preparatory tuned' tend to project less strongly onto the output while the units that are more 'movement tuned' tend to project more strongly onto the output. To do so we measured the ratio of preparatory activity to movement activity in each unit. Specifically, for all units and all times we calculated the total variance of preparatory activity (-600 to -400 before movement onset) across conditions and the total variance of perimovement period activity (200 to 400 ms after movement onset) across conditions. We then took the time point for each unit when these numbers were maximal, subtracted the movement variance from the preparatory variance, and normalized over the total variance. This gave a measure of the units' activity that varied between preparatory-activity dominated (+1) or movement-activity dominated (-1). We then took this measure and correlated it with the absolute strength of each unit's projection onto the output units (i.e. how strongly that unit influence the muscle activity directly). The correlation for the regularized model for monkey J was -0.25 (p-value of 0.007) and for the regularized model for monkey N was -0.27 (p-value of 0.0002). A scatter plot of these two measures is shown (blue dots are individual units). For illustration, PSTHs are shown for two example units near the extremes (red dots in scatter, conventions as in Figure 2). The black line is the linear fit. We also performed an additional analysis examining possible differences between those units with relatively stronger versus weaker output connections. We found that the population of neurons with stronger than average connections exhibited slightly higher-frequency oscillations (by 14% and 7% for the two monkeys) relative to the population with weaker than average connections. While it is hard to draw any firm conclusions, this result is in agreement with our recent report that the M1 population has slightly higher-frequency oscillations than the PMd population.



Supplementary Figure 11: Further comparisons between models and data, that include EMG.

(A, B) Average canonical correlations between the neural data and various datasets, as in Figure 5B (same conventions, EMG in orange). In this analysis the number of canonical correlations was first reduced to a number of principal components equal to the number of muscles in our recorded population (7 for monkey J and 8 for monkey N). In this way, we could also include the EMG versus neural data canonical correlations as another benchmark against which we could measure the similarity of the regularized model to the neural data. Note that this analysis is thus based on fewer principle components than the analysis in Figure 5 and Supplementary Figure 5. In both monkeys, (A) monkey J, (B) monkey N, the average canonical correlation between the regularized model and the neural data was higher than all other models studied, including also the EMG, thus showing that the similarity between the neural data and the regularized models does not arise simply by generating EMG.

(C, D) Comparison of the eigenvalue spectra between the neural data, the regularized dynamical model, and the EMG. PCA was run on the EMG and the eigenvalue spectrum plotted (orange). For the neural data and model data. For both the neural data and the regularized dynamical model, the PC eigenvalue spectra were generated on a number randomly selected neurons / units equal to the number of EMG each model output. These eigenvalue spectra were then averaged

over 100 random neuron/unit selections. Finally, the eigenvalue spectra were normalized so that the eigenvalue for the first PC had a value of 1.0. In this way, the dimensionality of the datasets can be seen in the falloff of the spectra. This analysis shows that even though the regularized dynamical model was trained to generate EMG, it contains more dynamical content than just the EMG. (C) monkey J, (D) monkey N.

Table 1 Model hyperparameters that varied between models

Model	g	h	M	Nonlinearity	α	β	γ	Multiple Delays
Regularized, monkey J (Figs. 1-7 , Supp. Figures 1, 10, 11)	1.5	1.0	7	rectified tanh	1e-4	0.03	1e-4	Yes
Regularized, monkey N (Supp. Figures 1-7, 10, 11)	1.5	1.0	8	rectified tanh	1e-4	0.01	5e-5	Yes
Complicated, monkey J (Figs. 1-7 , Supp. Figure 9)	3.75	1.0	7	rectified tanh	0	0	0	No
Complicated, monkey N (Supp. Figures 1-7, 9)	4.0	1.0	8	rectified tanh	0	0	0	No
Complicated to Regularized Study, monkey J (Regularized) (Fig. 8)	2.7	1.5	7	tanh	4e-4	0	4e-5	No
Complicated to Regularized Study, monkey N (Regularized) (Supp. Figure 8)	2.8	1.5	8	tanh	4e-4	0	4e-5	No
Complicated to Regularized Study, monkey J (Complicated) (Fig. 8)	2.7	1.5	7	tanh	0	0	0	No
Complicated to Regularized Study, monkey N (Complicated) (Supp. Figure 8)	2.8	1.5	8	tanh	0	0	0	No