

A neural population mechanism for rapid learning

Matthew G. Perich¹, Juan A. Gallego², Lee E. Miller^{1,2,3,*}

The brain's ability to adapt behavior is crucial to survival. Long-term learning of dexterous motor skills likely requires plastic changes in cortex¹, but we can also learn even from errors in single movements². Such trial-to-trial adaptation likely requires a faster mechanism. Here, we show how the brain can adapt behavior by redeploying existing population activity patterns without altering the functional structure of the cortex. We recorded from both primary motor cortex (M1) and dorsal premotor (PMd) cortex in macaque monkeys during motor learning. We trained computational models to predict single neuron spiking based on the activity of the surrounding neural population in order to study the functional relationships between neurons *within* the two areas³. Intriguingly, the functional structure within each area was preserved throughout learning, suggesting that the underlying neural circuitry remained unaltered^{4,5}. To study the interaction *between* the areas, we separated the PMd activity⁶ into two sets of components: “potent” components that captured activity that mapped onto M1, and “null” components that captured activity patterns with effects only within PMd⁷. As was true within each area, the activity of the potent components consistently predicted M1 spiking throughout its learning-related changes. In stark contrast, the mapping from the null components gradually changed with learning. These results show that, at a population level, PMd develops new motor plans (reflected in the null components) that are transmitted to M1 through an unchanged functional mapping (captured by the potent components) between the two areas. Use of the PMd-to-M1 null space as a neural scratch pad for the gradual development of new motor plans is a powerful mechanism that may explain a variety of rapid learning processes throughout the brain.

In order to make skilled movements, sensory input is combined with internal state variables and transformed into a plan executed by the motor cortices⁸. This transformation may be mediated in part by an “internal inverse model” that maps a desired motor action to the required low-level muscle commands. The dorsal premotor (PMd) and primary motor (M1) cortices, together with the cerebellum, are prime candidate locations for such an inverse model. PMd is involved in movement planning⁹, with diverse inputs and strong connectivity with M1¹⁰, while M1 is the main cortical output to the spinal cord¹¹. Correction of movement errors, such as those caused by external perturbations¹², is thought to lead to alteration of the inverse model, leading to progressively more accurate movements even on a trial-by-trial basis². The rapid speed of these changes seems incompatible with structural changes in synaptic connectivity^{13,14}. Furthermore, monkeys using brain-machine interfaces have great difficulty learning mappings between brain activity and cursor movement that require the normal pattern of covariation among recorded cortical neurons to be altered¹⁵. Given that neural covariance patterns seem to be determined by synaptic connectivity¹⁶, this result further suggests that changes in cortical connectivity may not be the primary mechanism for short-term learning. At the same time, the progressive change in performance over tens of minutes and the performance savings between sessions seem incompatible with a mechanism like the network reverberation that may underlie short-term working memory¹⁷. To reconcile these apparently contradictory observations, we recorded simultaneously from electrode arrays implanted in both M1 and PMd (Figure 1a) as monkeys learned to make accurate reaching movements that were perturbed by a curl field (CF), a velocity-dependent force applied to the hand¹². The CF altered the dynamics such that straight reaches to each target required the monkeys to learn new muscle activation patterns^{2,18}. We investigated whether changes in the functional relationships of the M1 and PMd populations could explain the adapted behavior.

¹Department of Biomedical Engineering, ²Department of Physiology, ³Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL

*Correspondence: lm@northwestern.edu

We trained two rhesus macaque monkeys to perform a standard center-out reaching task (Figure 1a). Each session began with reaches in a null field before the monkeys began to adapt to the CF^{12,18,19}, and progressively straighten their reaches (Figure 1b, S1). Evidence of their learning was revealed by the occurrence of after-effects upon eventual return to the null field (Figure 1b, S1). As has been previously observed, neural activity during adaptation was strikingly heterogeneous across neurons²⁰ (Figure S2). Intriguingly however, the population correlation structure was surprisingly similar before and after learning, despite the large changes in neural firing in both areas (Figures 1d,e, S3).

We used Poisson Generalized Linear Models (GLMs) to predict the spiking of individual neurons based on the activity of the remaining neurons (Figure 2a and Methods)^{3,21}. Using data from late in learning (Figure 2b) when behavior had stabilized, we trained three models: one predicted M1 neurons from the M1 population activity (M1-M1), another predicted PMd neurons from the PMd population (PMd-PMd), and the third predicted M1 neurons from the PMd population (PMd-M1) (Figure 2c). We assessed model performance using a relative pseudo- R^2 (rp R^2) metric, which quantified the improvement in model performance due to the neural inputs above that of reach kinematics alone (Figure S4; Methods). This removed the effect of shared variability due to behavior-related common inputs, and left the unique contributions of individual neurons. We tested generalization of each model from the late CF training data to the early CF trials. Good generalization would indicate that the relationships between neurons were unchanged during learning. The within-area models (M1-M1, PMd-PMd) predicted the complex spiking changes remarkably well throughout learning (Figure 2d-f, Figure S6 shows individual monkeys). However, early in learning, M1 spiking was poorly predicted from PMd. Furthermore, the model's performance improved with a time course very much like that of behavioral adaptation (Figure 2h), suggesting that learning resulted from a change in the functional relationships between neurons in PMd and M1, even though the functional interactions within the two networks remained unchanged^{4,16,5}.

We sought to explain how the relationships between PMd and M1 neurons could change while those within each area remained consistent. Intuitively, PMd population activity at once reflects its inputs, their subsequent processing, and the eventual outputs to M1²² (Figure 3a). We sought to separate these components of population activity by projecting PMd activity patterns onto output-null and output-potent spaces (Figure 3b, see Methods)⁷. We used Principal Component Analysis (PCA) to represent the activity of the M1 and PMd populations as a small number of components that captured mutual covariance patterns across neurons⁶. PMd consistently contained more components than M1 (Figure S5), indicating the existence of a null space containing PMd activity that had no net effect on the M1 components (Figure 3c, see Methods). We hypothesized that these extra components and the resulting null space arise from planning-related computations performed within PMd that did not directly activate M1. Such null-space planning activity could account for the altered overall relationships between PMd and M1, while at the same time, allowing the potent space to maintain a stable mapping from PMd to M1. We repeated the GLM analysis to predict the spiking of individual M1 neurons using either the PMd potent (Pot-M1) or null (Null-M1) components as inputs (Figure 3c). If the null components capture motor planning within PMd that changes with learning, the accuracy of the Null-M1 model should change with behavioral performance, much like the overall PMd-M1 model (Figure 2d-f, purple). However, if the updated motor plans are ultimately sent in a consistent manner to M1, Pot-M1 should remain unchanged. For both monkeys, Pot-M1 predicted M1 spiking consistently, while Null-M1 predictions changed with a time course like that of behavior (Figure 3e). The stability of Pot-M1 shows that, at a population level, there exists a direct mapping between PMd and M1 that persists throughout short-term motor adaptation. The stability of the potent mapping (along with the M1-M1 and PMd-PMd models) supports the conclusion that there were no structural changes in these cortical areas, since the potent

and null spaces are defined simply by different weightings of the same neurons with the same connectivity (Figure S5d,f). Learning, then, results from new activity patterns within PMd, which may be necessary to set a new preparatory state for M1^{7,23}.

We next asked if the observed changes within PMd are a necessary consequence of adapted behavior, or if they are indicative of a more specialized role for PMd in the CF task. On a separate set of sessions, the monkeys learned to reach in the presence of a static rotation of the reach-related visual feedback (visuomotor rotation; VR). Considerable evidence from behavioral studies in humans suggests that the brain areas involved in learning a static visual mapping differ from those required to learn novel effector dynamics (e.g., the CF)^{24,25}. Since VR learning appears to rely heavily on parietal cortex^{25,26}, hierarchically upstream yet of PMd, we hypothesized that VR adaptation would not result in a change in the functional relationship between M1 and PMd. We repeated the above analyses using sessions where the monkeys adapted to a VR of 30 degrees (Figure 4a). There were a number of similarities with the CF sessions: behavioral errors were similar in magnitude and time course (Figures 4b, S1c,d), and there were highly varied, complex changes in neural activity patterns with a preserved correlation structure (Figure S3d,e). However, when we fit GLM models to predict single neuron spiking, all models, including PMd-M1, accurately generalized throughout learning (Figure 4d). Thus, there were no changes in the functional relationships between the PMd and M1 populations, despite diverse changes in single-neuron activity. This result highlights a fundamental difference in the neural adaptation to these two perturbations, and supports the view that VR adaptation occurs upstream of PMd, likely involving parietal cortex^{25,26}. It also strengthens our conclusions about the CF task: the poor generalization of the PMd-M1 GLM model is not a necessary consequence of changing behavior, but rather captures a previously undescribed mechanism by which the motor cortices drive sensorimotor adaptation through population-level activity patterns.

Long-term learning is known to alter connectivity in the motor cortex, resulting in increased horizontal connections²⁷ and synaptogenesis²⁸. Many have proposed that the brain uses similar plastic mechanisms to adapt behavior on shorter timescales^{19,29}. However, structural changes would have impaired predictions of the GLM models^{4,5}. Hence our results suggest that, at least on the time scale of a single experimental session, there were no structural changes within PMd or M1. Our lab has previously found that the relationship between M1 activity and the dynamics of the motor output remains unchanged during CF adaptation¹⁸, with no evidence for adaptive changes in either spatial tuning or firing rates that have a time course like that of learning. Therefore, we hypothesized that CF learning must be mediated by changes in recruitment of M1 by upstream areas, including PMd. Our current results directly support this interpretation: PMd exploits the null space to formulate new motor plans reflecting the modified task demands of the CF, which are then used to recruit M1 without changing the connectivity within either area, or within the potent space from PMd to M1. The lack of any such change in null-space processing during the VR task suggests that VR adaptation occurs upstream of PMd.

Through lesion, computational, and recording studies, the cerebellum has been implicated in a variety of supervised, error-driven motor-learning problems, including both the curl field and visual rotation paradigms explored in this study^{25,30-32}. It is also considered to be a site at which both forward and inverse internal models may be learned^{33,34}. Many forms of cellular plasticity are present in the cerebellum, occurring at multiple sites and over several time scales^{35,36}. The cerebellum also supports rapid, short term memory storage through the bistable properties of Purkinje cells³⁷. The most direct evidence for its role in motor adaptation comes through Purkinje cell recordings during the adaptation of arm³⁸, eye³⁶, and head³⁹ movements. Given the extensive interconnections with PMd⁴⁰ the new motor plans during CF learning may arise from interactions between PMd and an evolving

inverse internal model in the cerebellum^{2,25,34}. Other evidence suggests that while these internal models may depend on the cerebellum for their modification, they may actually be located elsewhere⁴¹. Over a longer time period (hours to days), the motor memory is consolidated, possibly through structural changes in the cerebral cortex^{42,1}. We propose that such structural changes could emerge to support the long-term refinement and recall of skills^{43,42}, while rapid behavioral adaptation is mediated by modified population-wide activity patterns within the existing constrained network structure of the motor cortices. Similar activity patterns have been found in prefrontal cortex for decision-making⁴⁴, working memory⁴⁵, and rule-learning⁴⁶, in the motor cortex for movement planning⁷, and in the parietal cortex for navigation⁴⁷. These widespread observations suggest that the novel coordination mechanism between neuronal populations described here could be exploited throughout the brain for the rapid, flexible adaptation of behavior.

Acknowledgements

The authors would like to thank Dr. Sara Solla for her helpful discussions in preparing these analyses. J.A.G. received funding from the European Commission (FP7-PEOPLE-2013-IOF-627384). M.G.P. received funding from NIH NINDS T32 HD07418 and NIH NINDS F31 NS092356. This project was additionally funded by NIH NINDS NS053603 and NIH NINDS NS074044.

Ethical Statement

All procedures involving animals in this study were performed in accordance with the ethical standards of the Northwestern University Institutional Animal Care and Use Committee and are consistent with Federal guidelines.

Author Contributions

M.G.P. and L.E.M. conceived and designed experiments. M.G.P., J.A.G., and L.E.M. prepared figures, interpreted results, and wrote the manuscript. M.G.P. conducted experiments and analyzed data.

References

1. Nudo, R. J., Milliken, G. W., Jenkins, W. M. & Merzenich, M. M. Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J. Neurosci.* **16**, 785–807 (1996).
2. Thoroughman, K. a & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–7 (2000).
3. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13**, 105–11 (2010).
4. Ahissar, E. *et al.* Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science (80-.)*. **257**, 1412–5. (1992).
5. Gerhard, F. *et al.* Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone. *PLoS Comput. Biol.* **9**, 32–34 (2013).
6. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* (2014). doi:10.1038/nn.3776
7. 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**, 440–8 (2014).
8. Kalaska, J. F., Scott, S. H., Cisek, P. & Sergio, L. E. Cortical control of reaching movements. *Curr. Opin. Neurobiol.* **7**, 849–859 (1997).
9. Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–14 (2005).
10. Dum, R. P. & Strick, P. L. Motor areas in the frontal lobe of the primate. *Physiol. Behav.* **77**, 677–82 (2002).
11. Rathelot, J.-A. & Strick, P. L. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 918–23 (2009).
12. Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–24 (1994).
13. Bailey, C. H. & Chen, M. Morphological basis of short-term habituation in Aplysia. *J. Neurosci.* **8**, 2452–2459 (1988).
14. Bailey, C. H. & Chen, M. Morphological basis of long-term habituation and sensitization in Aplysia. *Science* **220**, 91–3 (1983).

15. Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**, 423–426 (2014).
16. Okun, M. *et al.* Diverse coupling of neurons to populations in sensory cortex. *Nature* (2015). doi:10.1038/nature14273
17. Major, G. & Tank, D. Persistent neural activity: Prevalence and mechanisms. *Curr. Opin. Neurobiol.* **14**, 675–684 (2004).
18. Cherian, A., Fernandes, H. L. & Miller, L. E. Primary motor cortical discharge during force field adaptation reflects muscle-like dynamics. *J. Neurophysiol.* **110**, 768–83 (2013).
19. Li, C. S., Padoa-Schioppa, C. & Bizzi, E. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* **30**, 593–607 (2001).
20. Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* **97**, 4235–57 (2007).
21. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–9 (2008).
22. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, (2015).
23. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* (2012). doi:10.1038/nature11129
24. Krakauer, J. W., Ghilardi, M. F. & Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat. Neurosci.* **2**, 1026–31 (1999).
25. Diedrichsen, J., Hashambhoy, Y., Rane, T. & Shadmehr, R. Neural correlates of reach errors. *J. Neurosci.* **25**, 9919–9931 (2005).
26. Tanaka, H., Sejnowski, T. J. & Krakauer, J. W. Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *J. Neurophysiol.* **102**, 2921–32 (2009).
27. Rioult-Pedotti, M. S., Friedman, D., Hess, G. & Donoghue, J. P. Strengthening of horizontal cortical connections following skill learning. *Nat. Neurosci.* **1**, 230–4 (1998).
28. Kleim, J. A. *et al.* Motor Learning-Dependent Synaptogenesis Is Localized to Functionally Reorganized Motor Cortex. *Neurobiol. Learn. Mem.* **77**, 63–77 (2002).
29. Classen, J., Liepert, J., Wise, S. P., Hallett, M. & Cohen, L. G. Rapid plasticity of human cortical movement representation induced by practice. *J. Neurophysiol.* **79**, 1117–23 (1998).
30. Izawa, J., Criscimagna-Hemminger, S. E. & Shadmehr, R. Cerebellar contributions to reach adaptation and learning sensory consequences of action. *J. Neurosci.* **32**, 4230–9 (2012).
31. Herzfeld, D. J. *et al.* Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *Neuroimage* (2014). doi:10.1016/j.neuroimage.2014.04.076
32. Galea, J. M., Vazquez, A., Pasricha, N., de Xivry, J.-J. O. & Celnik, P. Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cereb. Cortex* **21**, 1761–70 (2011).
33. Imamizu, H. *et al.* Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* **403**, 192–5 (2000).
34. Wolpert, D. M., Miall, R. C. & Kawato, M. Internal models in the cerebellum. *Trends Cogn. Sci.* **2**, 338–47 (1998).
35. Zheng, N. & Raman, I. M. Synaptic inhibition, excitation, and plasticity in neurons of the cerebellar nuclei. *Cerebellum* **9**, 56–66 (2010).
36. Yang, Y. & Lisberger, S. G. Interaction of plasticity and circuit organization during the acquisition of cerebellum-dependent motor learning. *Elife* **2013**, 1–19 (2013).
37. Loewenstein, Y. *et al.* Bistability of cerebellar Purkinje cells modulated by sensory stimulation. *Nat. Neurosci.* **8**, 202–11 (2005).
38. Gilbert, P. F. & Thach, W. T. Purkinje cell activity during motor learning. *Brain Res.* **128**, 309–28 (1977).
39. Brooks, J. X., Carriot, J. & Cullen, K. E. Learning to expect the unexpected: rapid updating in primate cerebellum during voluntary self-motion. *Nat. Neurosci.* **18**, 1–10 (2015).
40. Dum, R. P. & Strick, P. L. An unfolded map of the cerebellar dentate nucleus and its projections to the cerebral cortex. *J. Neurophysiol.* **89**, 634–639 (2003).
41. Shadmehr, R. & Holcomb, H. H. Neural correlates of motor memory consolidation. *Science* **277**, 821–5 (1997).
42. Peters, A. J., Chen, S. X. & Komiyama, T. Emergence of reproducible spatiotemporal activity during motor learning. *Nature* (2014). doi:10.1038/nature13235
43. Bailey, C. H. & Kandel, E. R. Structural Changes Accompanying Memory Storage. *Annu. Rev. Physiol.* **55**, 397–426 (1993).
44. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
45. Machens, C. K., Romo, R. & Brody, C. D. Functional, but not anatomical, separation of ‘what’ and ‘when’ in prefrontal cortex. *J. Neurosci.* **30**, 350–60 (2010).
46. Durstewitz, D., Vittoz, N. M., Floresco, S. B. & Seamans, J. K. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* **66**, 438–448 (2010).
47. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–8 (2012).

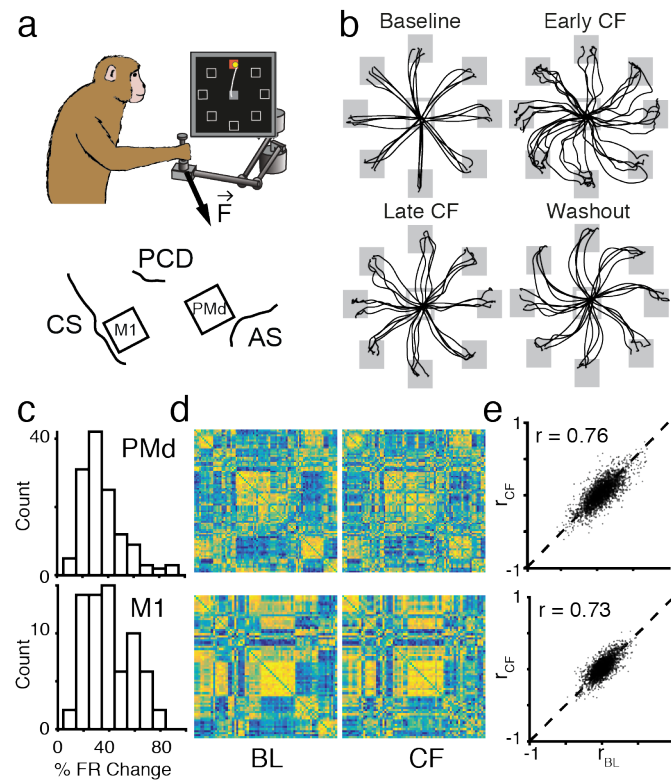


Figure 1 | Curl field task. a) Monkeys performed a standard center-out task with a variable instructed delay period following cue presentation (top). We recorded from single neurons in M1 and PMd (bottom; CS: central sulcus, PCD: pre-central dimple, AS: arcuate sulcus). b) Example position traces for the first reaches to each target for four representative CF sessions. Curvature increased when the CF was imposed (top right), but straightened during learning (bottom left). We observed oppositely directed after-effects in Washout (bottom right). c) Summary of percent of firing rate change for all cells recorded on a single session. d) Normalized pairwise correlations between all cells recorded on the same session as Panel c. Clustering was performed in Baseline (BL, left) as a means to visualize the correlation structure, and the same ordering was kept for late CF (right). e) Summary of pairwise correlations between BL and late CF for all combinations of neurons recorded in each of the nine CF sessions. A random subsample of 5,000 pairs is plotted. The value of ‘r’ for each plot indicates the Pearson’s correlation coefficient to assess similarity between the null-field and CF conditions for all pairwise-correlations from all sessions.

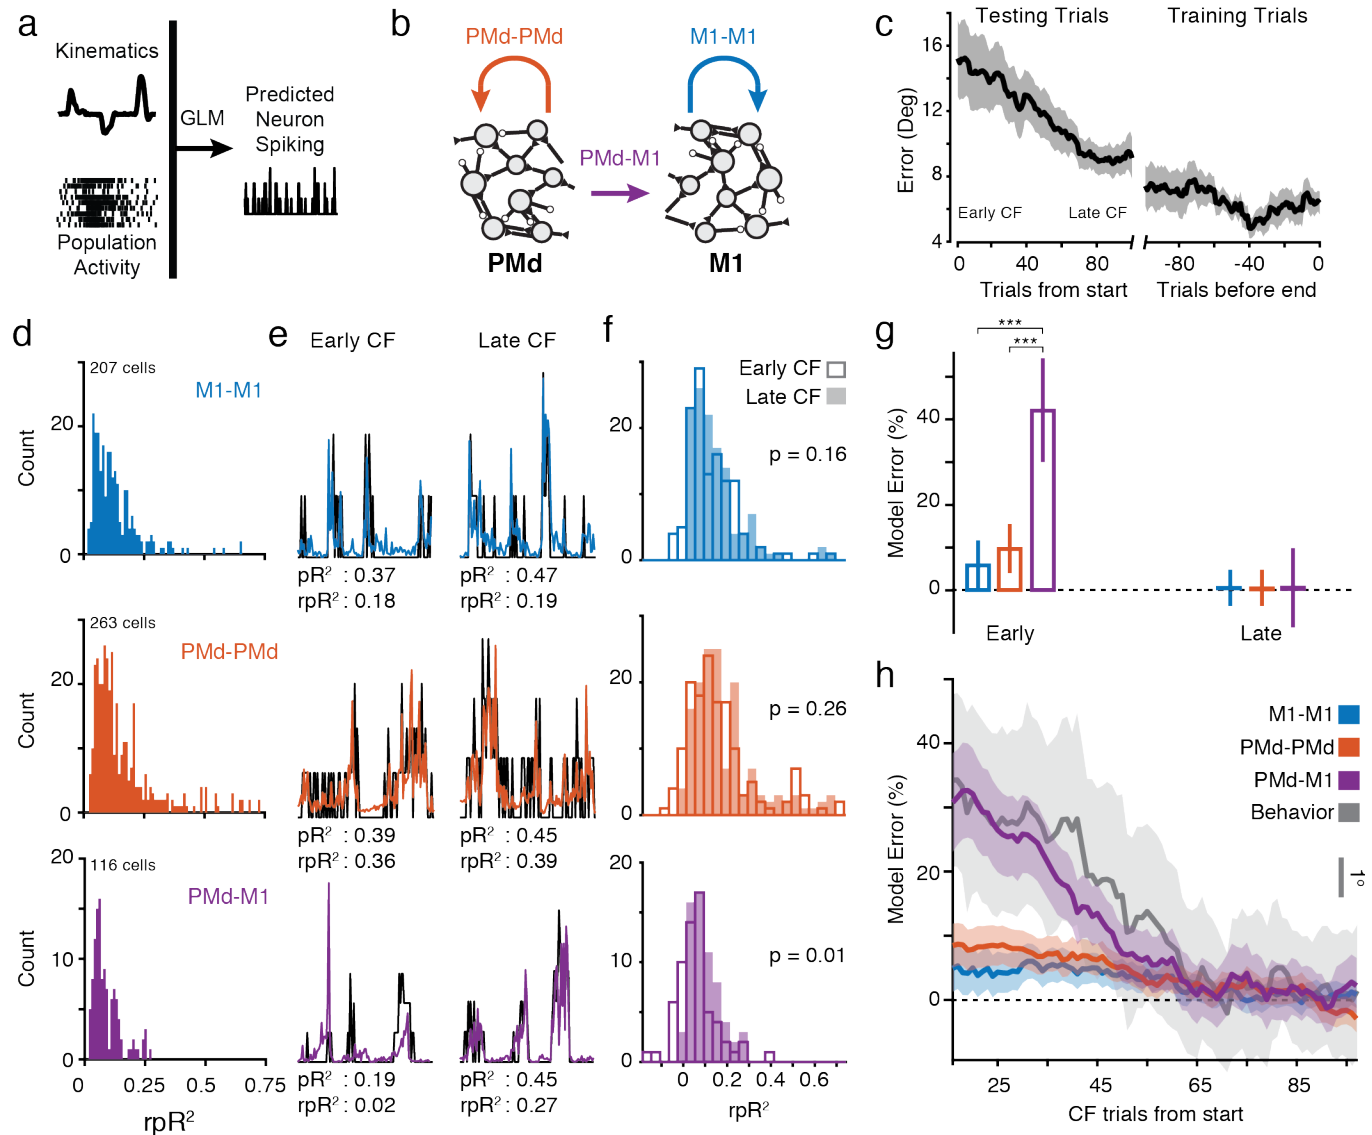


Figure 2 | Predicting neural spiking with GLMs. a) We used Poisson GLMs to predict single-neuron spiking using the kinematics and the activity of the remaining population. b) We trained three models: two within an area (M1-M1, and PMd-PMd) and one between the areas (PMd-M1). c) Angular error (mean \pm st.e. across sessions) during CF learning. We trained GLMs using data recorded late in learning, after behavior had stabilized (right) and tested them for generalization throughout the initial phase of learning, beginning at the first CF trial (left). In Panels e, f, and g, we compared Early CF (highest error) and Late CF (lowest error) trials within the testing block. d) Cross-validated rpR^2 for all cells with significant fits for the three GLM models. Cells were pooled across two monkeys and nine total sessions. e) Spiking of three representative neurons (black) and model predictions (colors) during three early and three late learning trials. f) Summary histograms of rpR^2 values for predictions of a block of 5 trials in early CF (hollow) and late CF (solid). M1-M1 and PMd-PMd had similar distributions during early and late CF, but Early CF predictions by the PMd-M1 models were significantly lower than Late CF ($p = 0.01$, two-sample t-test). g) Percent error in model performance during early and late learning. Significant differences were observed between PMd-M1 during Early and Late, and PMd-M1 compared with M1-M1 and PMd-PMd during late, with a significance level defined at $p < 0.01$ (two-sample t-test, N is given by cell counts in Panel d). h) Time course of model performance changes. Predictions were made for individual trials, and then smoothed with a 30 trial moving average (see Methods). Plotted data are mean and standard error across predicted neurons for each model. Behavioral error processed with the same methods is overlaid in gray. The decrease in PMd-M1 model error followed the time course of behavior.

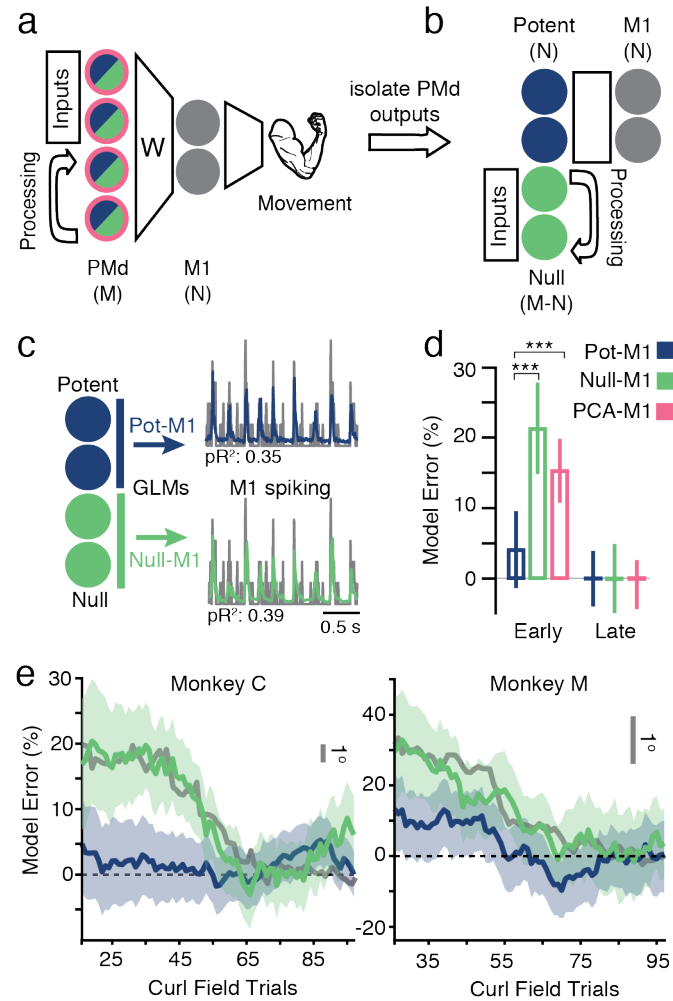


Figure 3 | Predictions from potent and null components. a) Hierarchical schematic of motor planning in PMd and M1. b) We devised an analysis to demix the PMd outputs to M1 from the other functions of the population. The former comprises the potent space, while the latter resides in the null space. c) The time-varying projections of PMd activity onto these potent and null axes were used as the inputs to a GLM model to predict M1 spiking. d) Bar plot comparing model error performance during early and late trials with potent (Pot-M1), null (Null-M1), and all PMd PCs (PCA-M1). Pot-M1 performed significantly better than the other models ($p < 0.01$, two-sample t-test). e) The time course of model performance for Pot-M1 and Null-M1 for all sessions with the two subjects. Gray line is the mean behavioral error corresponding to those trials.

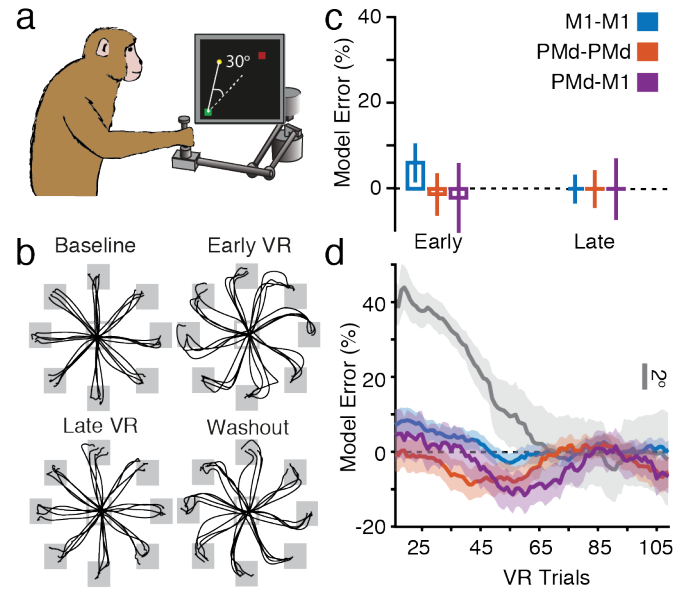


Figure 4 | Model performance during visuomotor rotation. a) The monkeys also adapted to a visuomotor rotation (VR) using the same center-out task as Figure 1. The VR rotated the visual cursor feedback on the screen by 30 degrees. b) Position traces for the first (or last) reach to each target for four representative sessions with the VR, shown as in Figure 1b. The monkeys exhibited behavioral errors that were similar to those of the CF condition (see Figure S1). c) Model prediction error during early and late VR trials. There were no significant changes throughout learning. d) Same format as Figure 2h. All GLM models, including PMd-M1 generalized well from late VR to early VR trials, despite clear behavioral adaptation (gray).

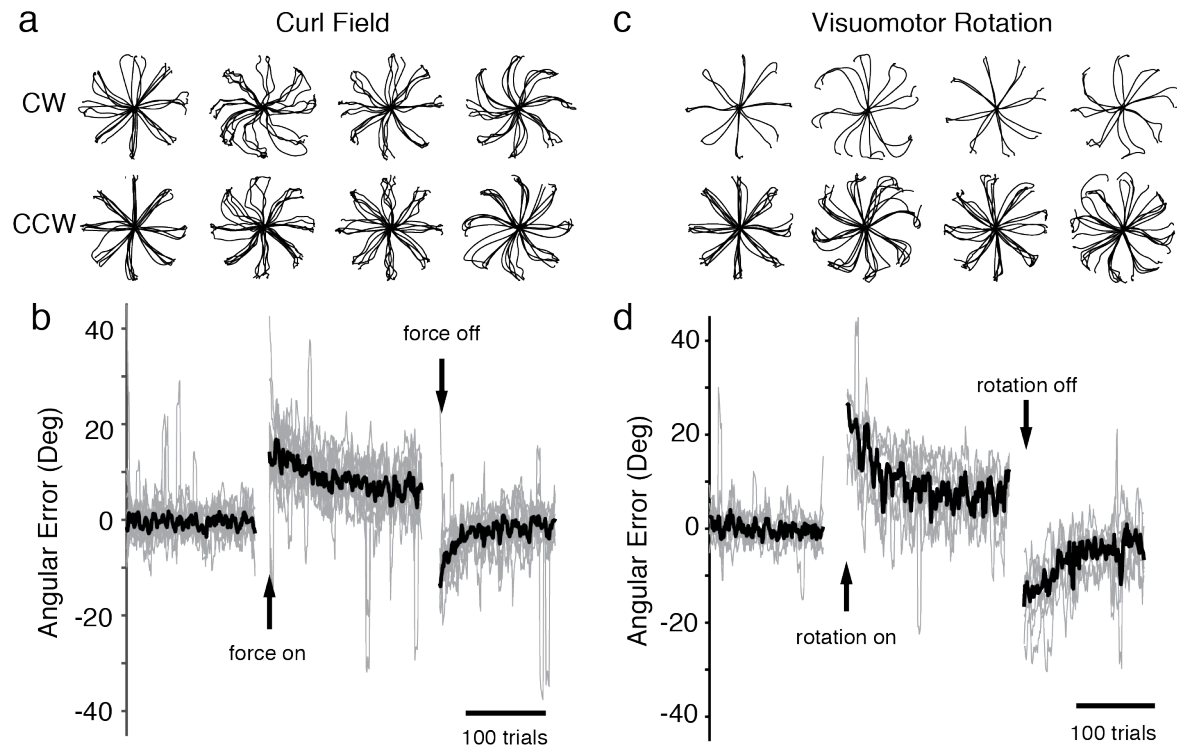


Figure S1 | Behavioral adaptation. a) Position traces for the first reach to each target from four sessions with a clockwise CF (top row) and five sessions with a counter-clockwise CF (bottom row). Sessions from both monkeys are included. Data from the three sessions with Monkey C with shorter reaches were plotted on a different scale to provide uniform length for visualization purposes. b) Error in the takeoff angle for all sessions (light gray lines), with the median across sessions shown in black. Gray traces were smoothed with a 4 trial moving average to reduce noise while preserving the time course of adaptation. c) Same as Panel a, but for the VR sessions (two clockwise and 5 counter-clockwise). d) Same as Panel b for the VR sessions.

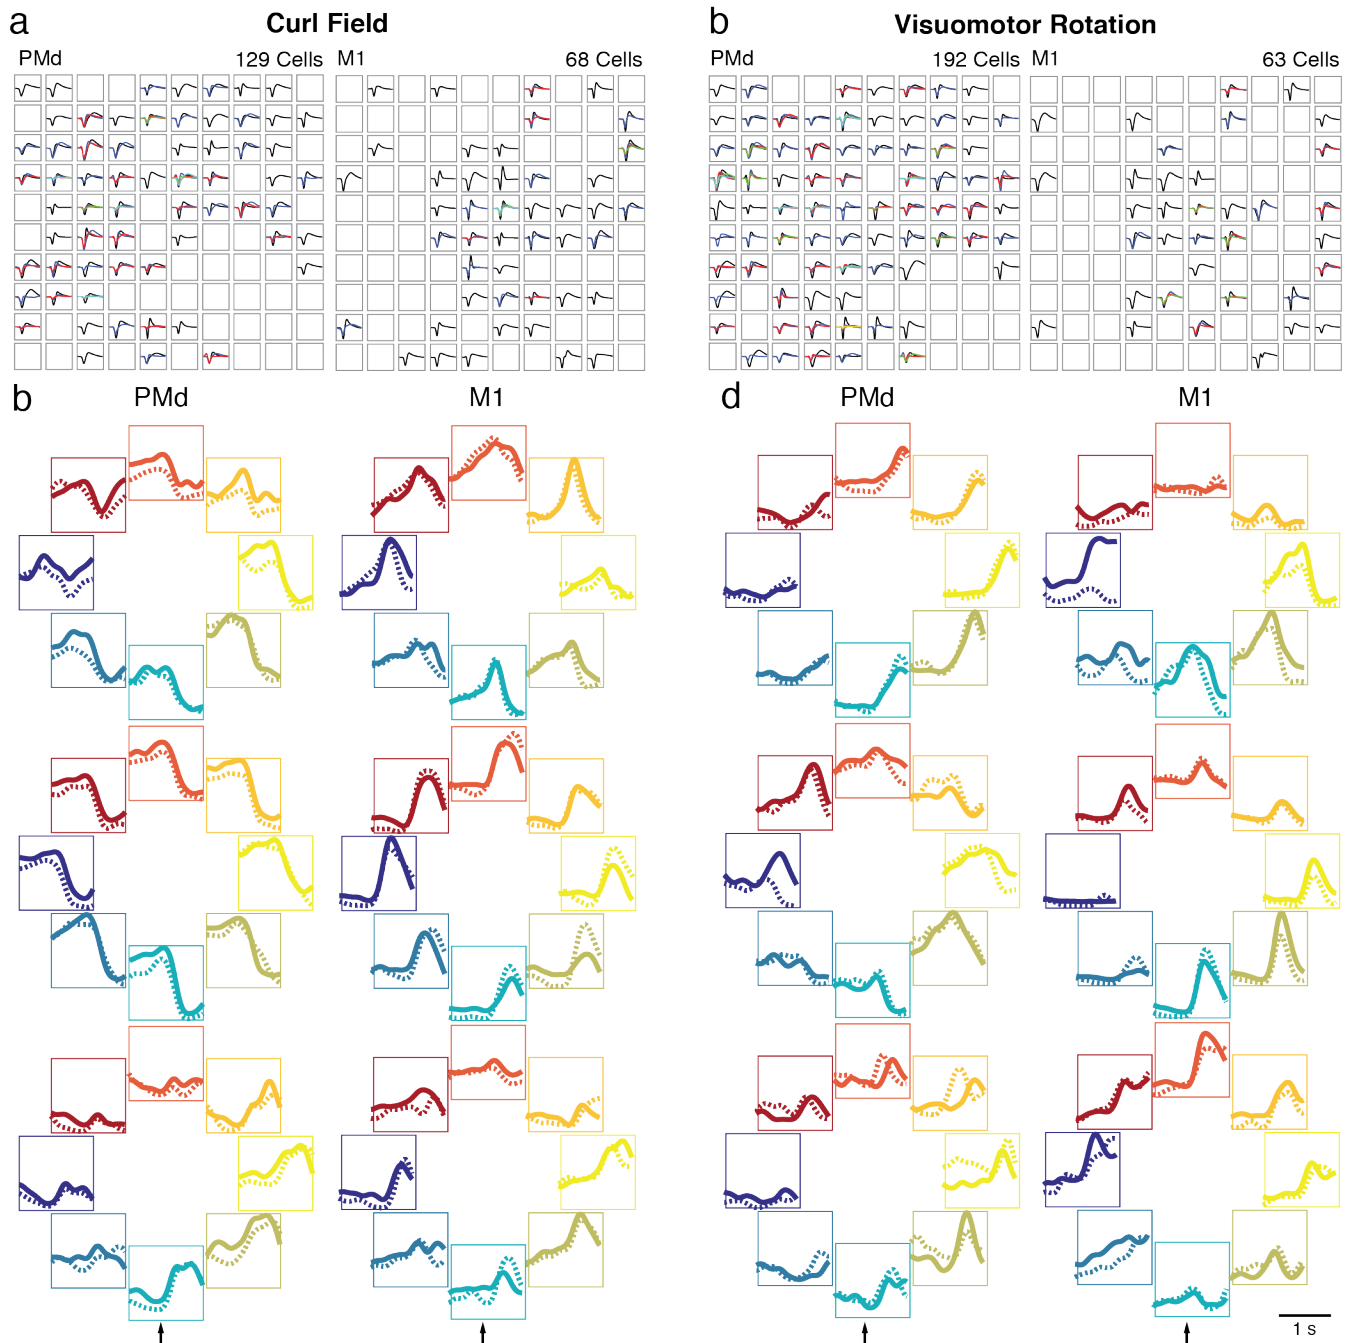


Figure S2 | Example neural activity before and after CF learning. a) Array recordings for a single session for PMd (left) and M1 (right). The mean waveform is plotted in the spatial location on the array, with each sorted unit having a different color. b-d) Activity patterns for three example neurons from the same CF session. Neural activity for each trial was smoothed with a Gaussian kernel, aligned to the go cue (indicated by the arrow at the bottom of each panel), and averaged across reaches to each target (colored directions) for Baseline (solid lines) and late CF (dashed lines). e-g) Same as Panels a-c, but for three M1 neurons.

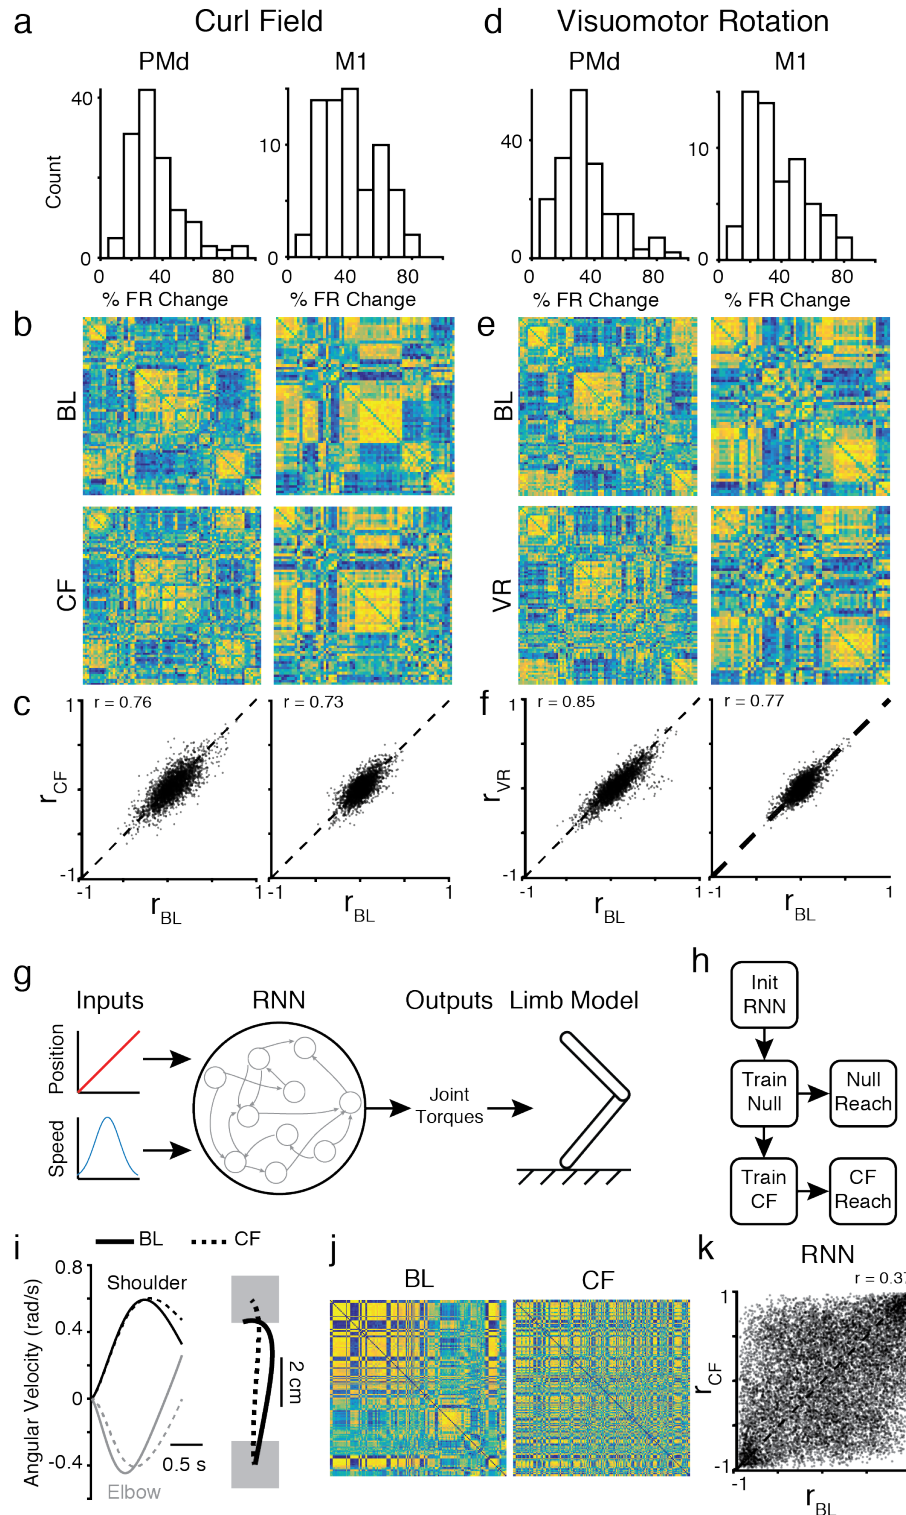


Figure S3 | Correlation structure during learning. a) Summary of percent firing rate changes for M1 and PMd neurons in a CF session, reproduced from Figure 1c. b) Correlation structure for all M1 (left) and PMd (right) in a CF session, for Baseline (BL, top) and CF (bottom). Data reproduced from Figure 1d. c) Scatter plot of BL and CF pairwise correlations, reproduced from Figure 1e. d-f) Same as panels a-c, but for a representative VR session. The two perturbations had similar amounts of activity changes, and both had preserved correlation structure. g) A simple Recurrent Neural Network (RNN) was trained to provide joint torques to drive a model of the limb (see Methods). h) The RNN was trained to perform a single reach

direction, then the curl field was applied to the endpoint and the model was further trained to apply compensatory torques. In this case, learning is known to occur through structural changes in the network, and provides a reference to compare against the recorded neural data. i) Angular velocities for the elbow and shoulder joints (left) and endpoint position (right) for BL (solid) and CF (dashed) reaches to the target. j) Panel b was replicated using pairwise correlations between units in the RNN during BL and CF. Much of the structure changed during the CF. k) Panel c was replicated using units in the RNN. There is much less correlation between activity of the units in BL and CF than in the recorded neurons, suggesting that the real CF and VR learning does not result in structural changes.

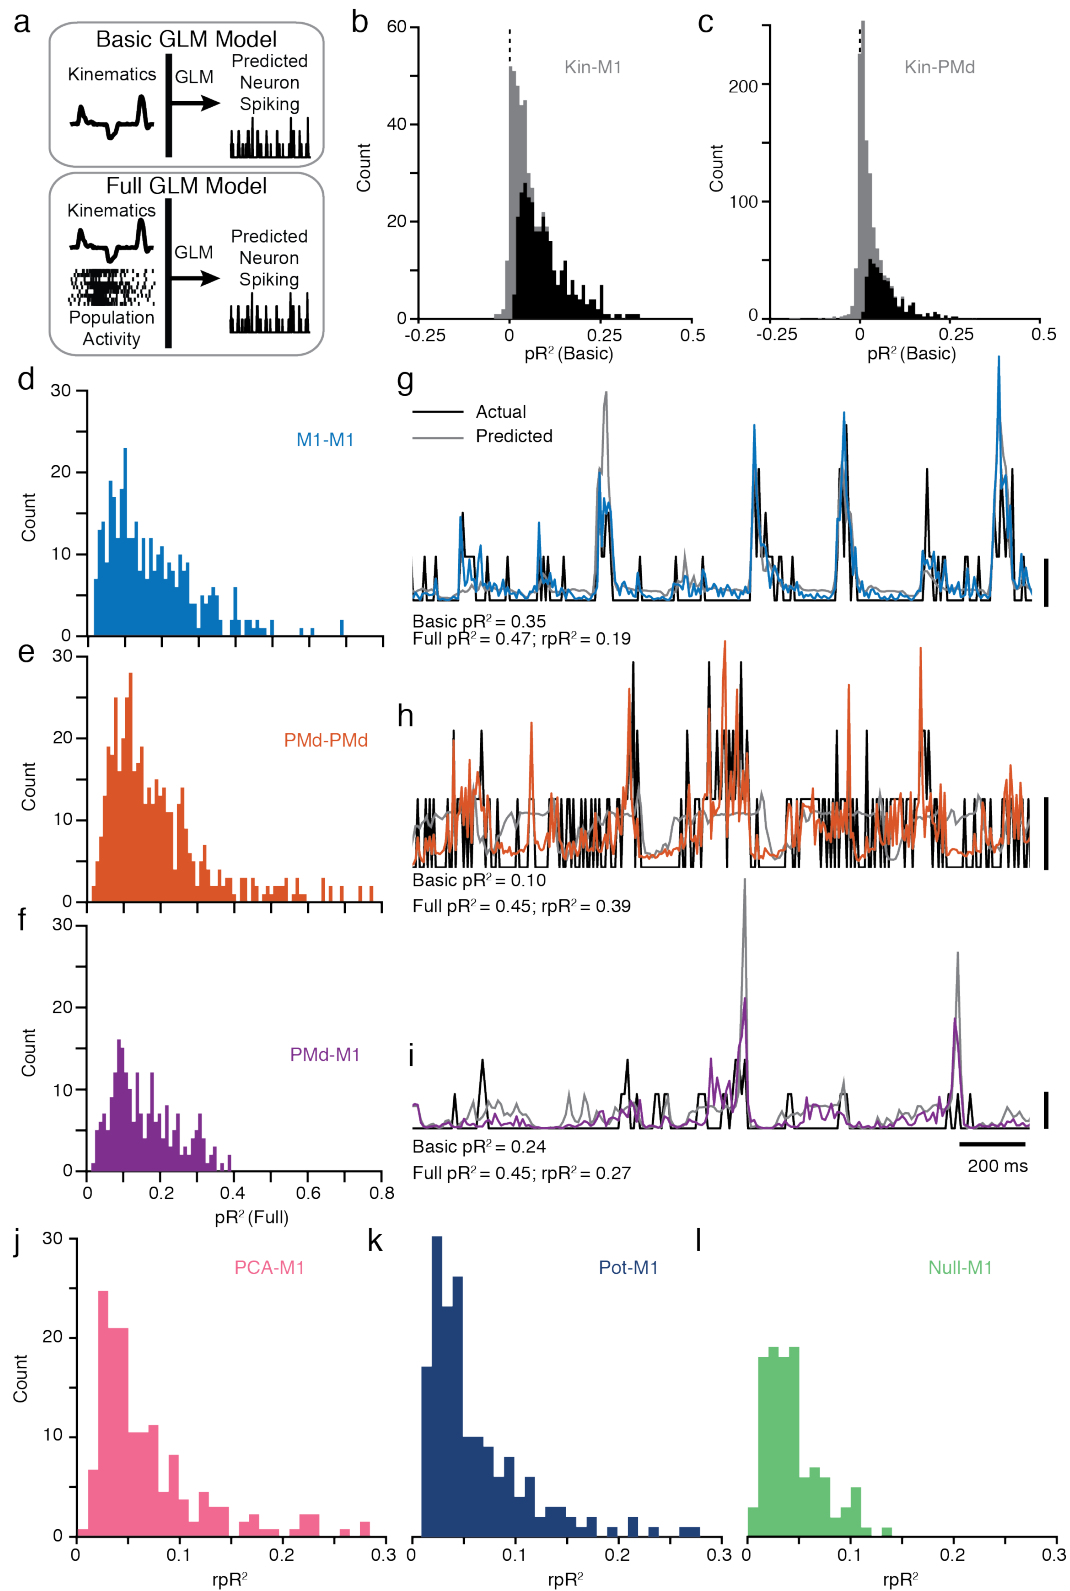


Figure S4 | GLM Model Performance a) Schematic representation of the GLM models. The Basic model included only kinematic covariates (see Methods), while the Full model included both kinematics and neural activity. The relative pseudo- R^2 metric was a comparison between these two models. b) Distribution of cross-validated pseudo- R^2 values for predictions of

all M1 neurons from all sessions with the Basic model (gray). Black overlaid distribution shows cells with significant model fits (see Methods). c) Same as Panel b, but for predictions of PMd neurons. d) Distribution of cross-validated pseudo- R^2 values for the Full M1-M1 model. All neurons from all sessions with significant fits are included. e-f) Same as Panel d, but for the PMd-PMd and PMd-M1 models. g) Predictions of M1 spiking (black) for a single example neuron. The Basic model prediction (gray) is compared to the M1-M1 model prediction (blue). The black scale bar on the right indicates one spike, and pseudo- R^2 values are shown below the predictions. h-i) Same as Panel g, but for PMd-PMd and PMd-M1. j-l) Distribution of cross-validated relative pseudo- R^2 values for the PCA-M1 (j), Pot-M1 (k), and Null-M1 (l) models. The three models had similar performance.

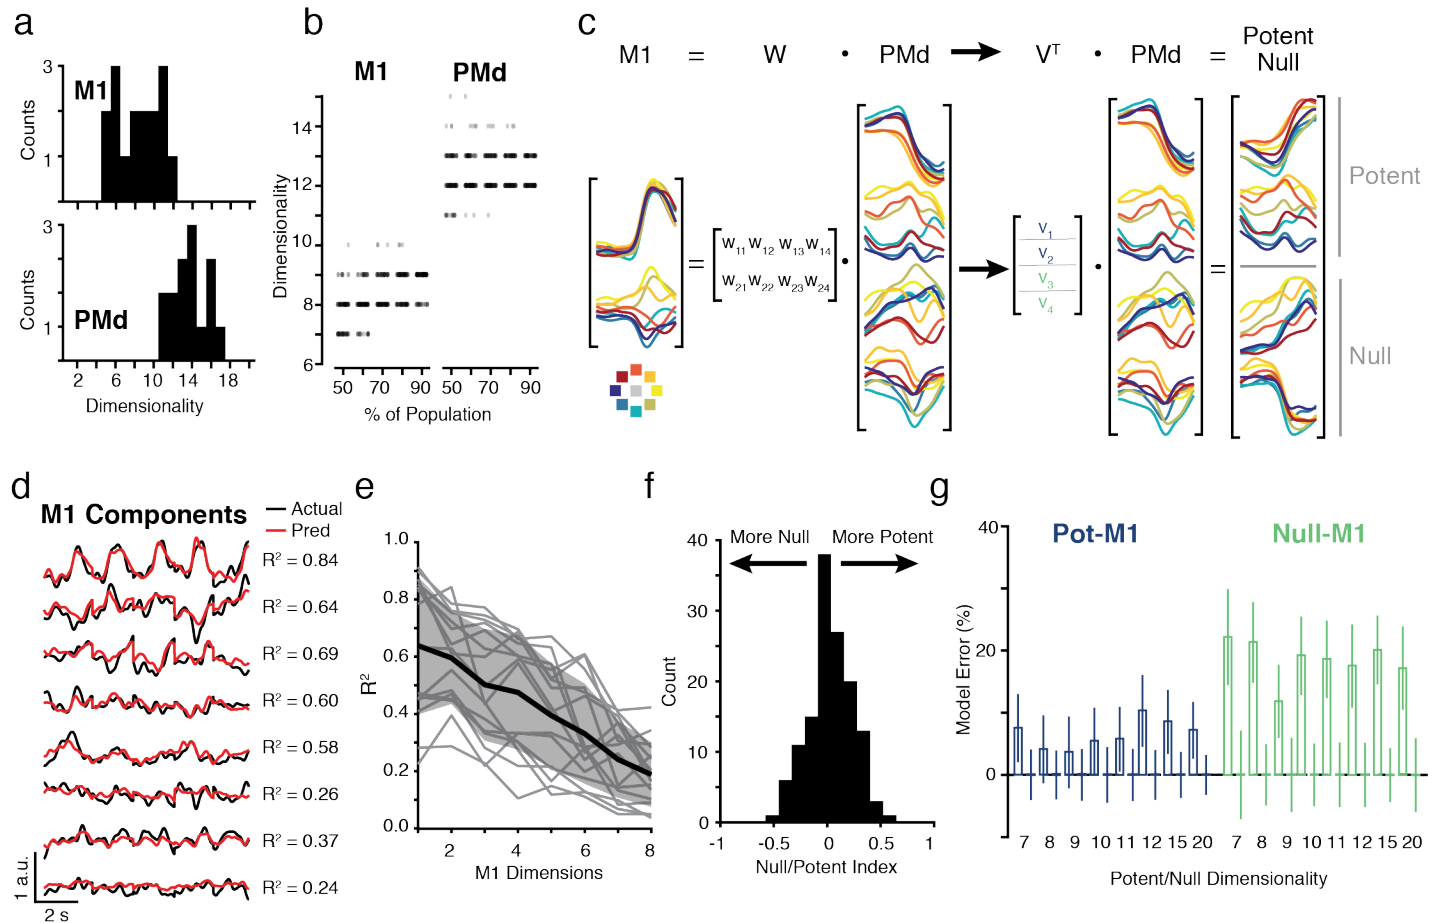


Figure S5 | Details on Potent and Null Spaces a) Summary histograms across all sessions from both monkeys of dimensionality for M1 (top) and PMd (bottom). b) For one example session from Monkey C, the effect of population size on dimensionality. We randomly subsampled the neural populations 100 times at each percentage and repeated the dimensionality analysis. The result of each repetition is plotted as a single point with a random jitter on the horizontal axis to show the density. PMd (right) was consistently higher dimensional than M1 (left). c) Schematic representation of the method to identify output-potent and output-null spaces using trial-averaged data from a single session. Traces show the activity of components found by PCA on a representative session, with each color corresponding to one of the eight target directions. We used multi-linear regression to build a matrix W relating the activity of the M1 PCs (here, dimensionality of two) to the PMd PCs (here, dimensionality of four). Thus, W is a 2x4 matrix. Using Singular Value Decomposition (see Methods), we identified a matrix V^T , the first two rows of which contained the basis vectors of the potent space, while the last two rows defined the null space. We multiplied the PMd PCs by this matrix to get the time-varying potent (top) and null (bottom) projections. d) Example predictions (red) of the first eight M1 PCs (black) from the first sixteen PMd PCs, with R^2 quantifying quality of fit for a single session. e) Summary of R^2 for M1 PC predictions across sessions (gray lines). Black line and gray shading indicate mean and st.dev. across sessions. f) We attempted to identify potent or null subpopulations using an index that quantified the relative weights of each neuron onto the potent and null axes (see Methods). Values of 1 indicate the cell was exclusively potent, and values of -1 indicate the cell was exclusively null. The distribution of cells was centered around zero, indicating that the potent and null spaces captured population-wide activity patterns. g) Comparison of GLM performance error between early CF (left bars for each dimension) and late CF trials (right bars) for Pot-M1 and Null-M1 as a function of the selected dimensionality. Values for eight dimensions plotted here are those included in Figure 3d. Our primary effect that Pot-M1 generalizes to early CF trials better than Null-M1 was consistent for a range of dimensionalities.

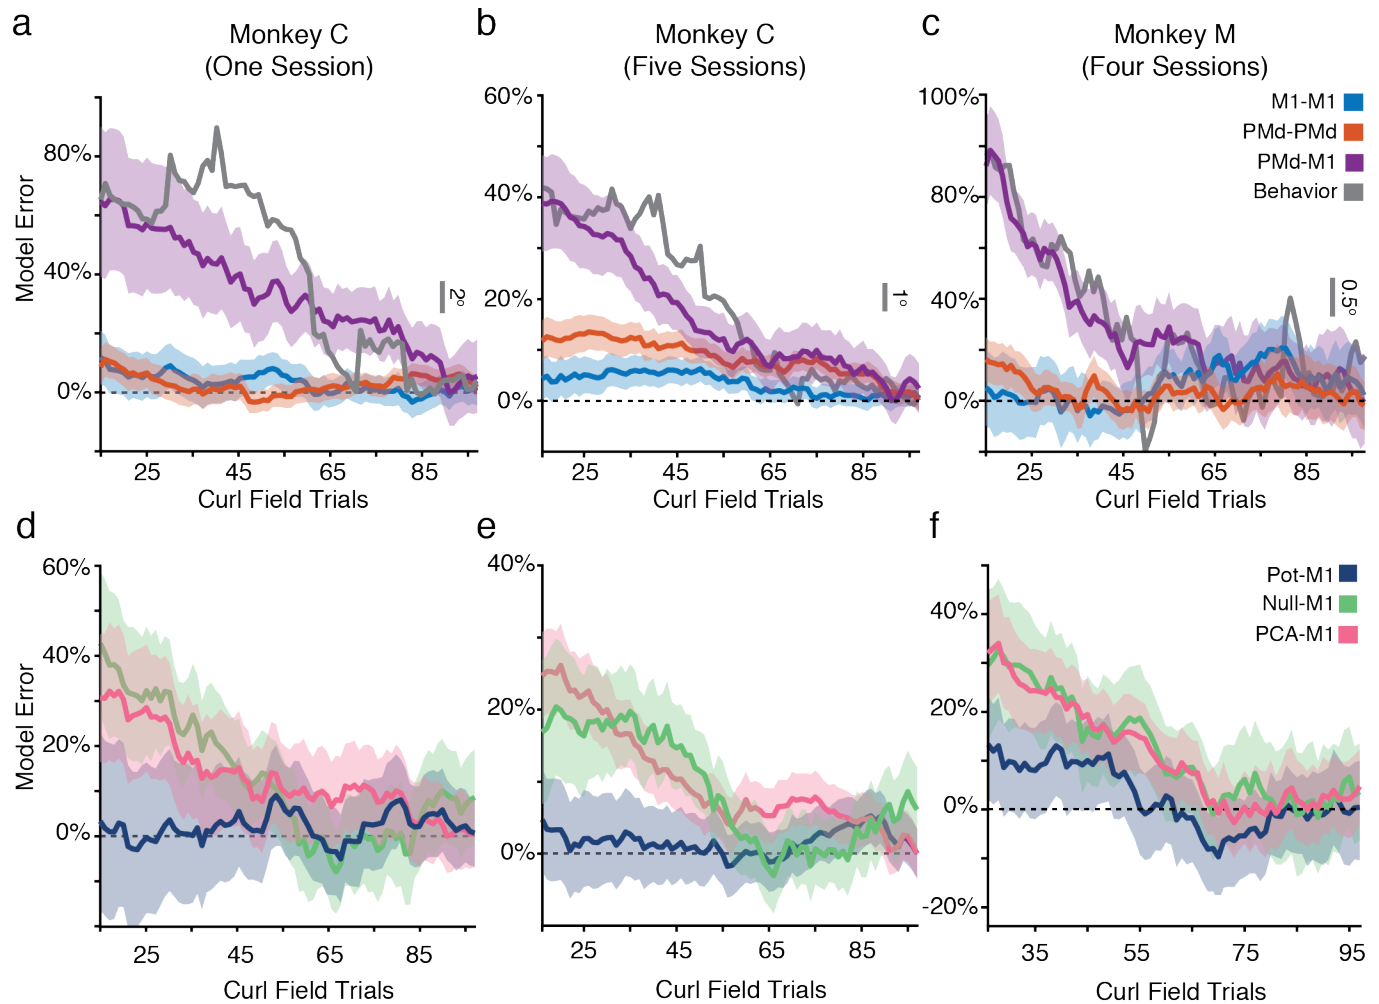


Figure S6 | Subject-specific GLM performance. a) Identical to Figure 2h, but for all neurons recorded on a single session from Monkey C. The gray line shows a moving average of the mean behavioral error of the monkey on that session for the same trials. b-c) Same as Panel a, but for all sessions from Monkey C (b) and Monkey M (c). d) Identical to Figure 3e, but for a single session from Monkey C. e) Identical to Panel d, but for all sessions from Monkey C. f) Identical to Panels d, but for a session from Monkey M. Since Monkey M contained fewer neurons and predictions were considerably noisier, we extended the moving average window size to 50 trials.

Methods

Behavioral task

Two monkeys (male, *mucaca mulatta*; Monkey C: 11.7 kg, Monkey M: 10.5 kg) were seated in a primate chair and made reaching movements with a custom 2-D planar manipulandum to control a cursor displayed on a computer screen. We recorded the position of the handle at a sampling frequency of 1kHz using encoders. The monkeys performed a standard center-out reaching task with eight outer targets evenly distributed around a circle at a radius of 8cm. All targets were 2cm squares. The first three sessions with Monkey C used a radius of 6 cm. However, we observed no qualitative difference in the behavioral or neural results for the shorter reach distance, and all sessions were thus treated equally. Each trial began when the monkey moved to a center target. After a variable hold period (0.5 – 1.5 s), one of the eight outer targets appeared. The monkey had a variable instructed delay period (0.5 – 1.5 s) which allowed us to study neural activity during explicit movement planning and preparation, in addition to movement execution. The monkeys then received an auditory go cue, and the center target disappeared. The monkeys had one second to reach the target, where they had to hold for 0.5 s.

In the curl field (CF) task, two motors applied torques to the elbow and shoulder joints of the manipulandum in order to achieve the desired endpoint force. The magnitude and direction of the force depended on the velocity of hand movement according to Equation 1, where \vec{F} is the endpoint force, $\dot{\vec{p}}$ is the derivative of the hand position \vec{p} , θ_c is the angle of curl field application (85°), and k is a constant ($0.15 \text{ N}\cdot\text{s}/\text{cm}$):

$$\vec{F} = \begin{bmatrix} F_x \\ F_y \end{bmatrix} = k \begin{bmatrix} \cos \theta_c & -\sin \theta_c \\ \sin \theta_c & \cos \theta_c \end{bmatrix} \begin{bmatrix} \dot{p}_x \\ \dot{p}_y \end{bmatrix} \quad (1)$$

In the visuomotor rotation (VR) task, hand position p was rotated by θ_r (here, chosen to be 30°) to provide altered cursor feedback \vec{C} on the screen. The rotation was position-dependent so that the cursor would return to the center target with the return reach:

$$\vec{C} = \begin{bmatrix} C_x \\ C_y \end{bmatrix} = \begin{bmatrix} \cos \theta_r & -\sin \theta_r \\ \sin \theta_r & \cos \theta_r \end{bmatrix} \begin{bmatrix} p_x \\ p_y \end{bmatrix} \quad (2)$$

Both the CF and VR perturbations were imposed continuously throughout the block of learning trials, including the return to center and outer target hold periods.

Each session was of variable length since we allowed the monkeys to reach as long as possible to ensure that behavior had sufficient time to stabilize, and allow for large testing and training sets for the GLM. For the CF sessions, the monkeys performed a block of unperturbed Baseline trials (range across sessions: 170 – 225 rewards) followed by an Adaptation block with the CF perturbations (201 – 337 rewards). The session concluded with a Washout block, where the perturbation was removed and the monkeys readapted to making normal reaches (153 – 404 rewards). The curl field was applied in both clockwise (CW) and counter-clockwise (CCW) directions, though we saw no qualitative difference between the sessions. Monkey C had three CW sessions and two CCW sessions, while Monkey M had four CCW sessions. For the VR sessions, the monkeys performed 154 – 217 successful trials in Baseline, 219 – 316 during VR (either CW or CCW), and then 162 – 348 in Washout. Monkey C performed two CW VR sessions and two CCW sessions, while monkey M performed three CCW sessions. There is considerable evidence that learning can be consolidated, resulting in savings across sessions¹. In this study, we minimized the effect of savings to focus on single-session learning. The monkeys typically: 1) received different perturbations day-to-day, as we alternated between CF and VR sessions, 2) received opposing directions of the perturbation on subsequent days, and 3) had multiple days between successive perturbation exposures.

Behavioral adaptation analysis

For a quantitative summary of behavioral adaptation, we used the errors in the angle of the initial hand trajectory. We measured the angular deviation of the hand from the true target direction 150 ms after movement onset. To account for the natural biases of the monkeys, we found the difference on each trial from the average deviation for that target in Baseline trials. Sessions with the CW and CCW perturbations were similar except for the sign of the effects. Thus, for the behavioral data in Figures 2h, 4f, and S1, we pooled all perturbation directions together and simply flipped the sign of the CW errors. Figures 1b and 4b show the position traces for example CF and VR sessions, respectively. Since the target size was 2cm, there could be some deviation in the starting and ending positions, and subsequently some deviation in the total length of the reaches. For visualization purposes only, we normalized the length of each reach to begin in the center of the workspace and have a total linear distance of 8cm between the starting and ending points.

Neural recordings

After extensive training in the unperturbed center-out reaching task, we surgically implanted chronic multi-electrode arrays (Blackrock Microsystems, Salt Lake City, UT) in M1 and PMd. From each array, we recorded 96 channels of neural activity using a Blackrock Cerebus system (Blackrock Microsystems, Salt Lake City, UT). The snippet data was manually processed offline using spike sorting software to identify single neurons (Offline Sorter v3, Plexon, Inc, Dallas, TX). We sorted data from all three task epochs (Baseline, CF/VR learning, and Washout) simultaneously to ensure we reliably identified the same neurons throughout the sessions. With such array recordings, there is a small possibility that duplicate neurons can appear on different channels as a result of electrode shunting, which would influence our GLM models by providing perfectly correlated inputs for these cells. While such duplicate channels are often easily identifiable during recording, we took two precautionary steps to ensure our data included only independent channels. First, we used the electrode crosstalk utility in the Blackrock Cerebus system to identify and disable any potential candidates with high crosstalk. Second, offline we computed the percent of coincident spikes between any two channels, and compared this percentage against an empirical probability distribution from all sessions of data. We excluded any cells whose coincidence was above a 95% probability threshold (in practice, this was approximately 15-20% coincidence, which excluded no more than one or two low-firing cells per session).

Across all sessions, we isolated between 137 – 256 PMd and 55 – 93 M1 neurons for Monkey C, and 66 – 121 PMd and 26 – 51 M1 neurons for Monkey M. For the pairwise correlation analysis, we excluded cells with a trial-averaged firing rates of less than 1 Hz. Our GLM models were by necessity poorly fit for neurons with low firing rates. Thus, for the GLM analyses, we only considered neurons with a trial-averaged mean firing rate greater than 5 Hz. Pooled across all monkeys and CF and VR sessions, this gave a population of 918 M1 and 2221 PMd neurons. Given the chronic nature of these recordings, it is certain that some individual neurons appeared in multiple sessions. However, our analyses primarily focus on the population-level relationships which we found to be robust to changes in the exact cells recorded, so we do not expect our results to be biased by partial resampling.

Dimensionality reduction

We counted spikes in 10 ms bins and square root transformed the raw counts to stabilize the variance². We then convolved the spike train of each neuron for each trial with a Gaussian kernel of width 100 ms to compute a smooth firing rate. We used Principal Component Analysis (PCA) to reduce the smoothed firing rates of the neurons in each session to a small number of components². PCA finds the dominant covariation patterns in the population and provides a set of orthogonal basis vectors that captures most of the population variance. Importantly, the axes of PCA capture population-wide interactions, with nearly all neurons contributing to the dominant components.

For the null and potent space analysis described below, we needed to select dimensionalities for M1 and PMd. We adapted a method developed by Machens et al³ to estimate the dimensionality of our recorded populations. In brief, PCA provides an orthogonal basis set with the same dimensionality as the neural input. However, the variance captured by many of the higher dimensions (with the smallest eigenvalues) is typically quite small. We estimated the noise in the neural activity patterns using the trial-to-trial variation in the activity of each neuron.

We sampled a random pair of trials for each reach direction and subtracted the activity of each neuron. This gave an estimate of the variance of each neuron across two different reaches to each target. We then ran PCA on the neural “noise” space provided by this difference for all targets. We repeated this 1000 times, giving a distribution of eigenvalues for each of these noise dimensions. We used the 99% limit of these distributions to estimate the amount of noise variance explained for each dimension. This allowed us to put a ceiling on the amount of variance that could be explained by noise. The dimensionality was thus defined by the number of dimensions needed to explain 95% of the remaining variance (Figure S5d shows all sessions for M1 and PMd). Importantly, the dimensionality we estimated was robust to the number of recorded neurons since it reflected population-level patterns. We performed a control where we repeated the above analysis with random subsamples of neurons, taking 50-90% of the M1 or PMd populations (Figure S5e) and observed no change in the estimated dimensionality.

Potent and null space calculation

Using the above method, we estimated the dimensionality of the M1 and PMd populations on each session. Since we identified a larger dimensionality for PMd than M1, there existed a “null space” in PMd, which encompasses PMd activity that has no net effect on M1⁴. To identify the geometry of the null and potent spaces, we constructed multi-input multi-output (MIMO) linear models relating the N-dimensional PMd space to the O-dimensional M1 space (with $N > O$):

$$M = WP \tag{3}$$

M ($O \times t$) and P ($N \times t$) are matrices whose rows contain the activity of each PC for M1 and PMd, respectively, and whose columns contain the time points (t). We then performed singular value decomposition (SVD) of the matrix W ($O \times N$) that maps PMd onto M1:

$$W = USV' \tag{4}$$

SVD decomposes the rank-deficient rectangular matrix W into a set of orthonormal basis vectors that allows us to define the null and potent spaces. For our purposes, the matrix V' defines the vectors that define the potent and null spaces, with the first N rows corresponding to the potent space, and the remaining $M - N$ rows defining the null space (Figure S5a, Equation 5):

$$V = \begin{bmatrix} v_{11} & \cdots & v_{1M} \\ \vdots & \ddots & \vdots \\ v_{N1} & \cdots & v_{NM} \end{bmatrix} \tag{5}$$

We used only trials from the Baseline period of each session to find the axes for PCA, as well as the null and potent spaces. The Baseline trials were independent of the CF/VR trials used for both testing and training the GLM models, ensuring that we did not bias our results to find any specific solutions. However, we obtained nearly identical results if we used all of the data, or data only from the CF/VR trials, indicating that the null and potent spaces identified through this analysis did not change throughout the session. It is also important to note that the null and potent spaces, as with the PCA axes, typically comprised population-wide activity patterns, rather than sub-groups of neurons (Figure S5f).

Single neuron correlation analysis

We studied the correlations between individual neurons using the same smoothed firing rates we used for PCA. We then aligned each trial at the time of movement onset and isolated a window beginning 700 ms before and ending 800 ms after movement onset. We averaged across trials for each target direction, during both the pre-learning Baseline and the learning epochs. We excluded the first 50% of CF or VR trials to look at neural activity when adaptation was most complete. Examples of these trial-averaged activity profiles were used for the plots in Figures 1, 4, S2, and S7. We then performed pairwise cross-correlations between all neurons recorded on each

session during Baseline and late CF/VR. The coefficient of correlation values shown in Figures 1e and S3 were computed using the pairwise correlation values using all pairs of neurons from each session, as a means to quantify the similarity between the two conditions. For the heat maps shown in Figures 1e and S3, we normalized the range of each row to scale from -1 to 1 to enhance visualization. We then used a simple hierarchical clustering algorithm to sort the neurons in the Baseline condition. This same sorting order was used for the late CF heat map as a means of visually assessing the consistency in the correlation structure.

Recurrent Neural Network model

We sought to develop a reference to interpret our pairwise correlation models in a scenario where there are known structural changes. We trained a Recurrent Neural Network (RNN) to drive torques applied to the elbow and shoulder joints of a 2-D planar arm (Figure S3g). The RNN was trained using the HebbRNN toolkit⁵. In brief, we initialized a 300 unit neural network randomly, and trained over subsequent iterations to reach the desired endpoint. No constraints were placed on the hand path or velocity. We trained the initial model to control the limb in the null field, then retrained the same network to make the same reach with a modeled curl field imposed on the endpoint (Figure S3h,i). We compared the activation profiles for these two reaches using pairwise correlations, as was done with the real neurons (Figure S3j,k).

Generalized Linear Models

We trained Poisson Generalized Linear Models⁶ (GLMs) to predict the spiking activity of individual neurons on a single-trial basis⁷. GLMs extend Gaussian multilinear regression approaches for the Poisson statistics of neural spiking. We take weighted linear combinations of the desired covariates, x_i , such as limb kinematics:

$$\sum_i \theta_i x_i = X\theta \quad (6)$$

The weighted covariates were passed through an exponential inverse link function. The exponential provides a non-negative conditional intensity function λ , analogous to the firing rate of the predicted neuron:

$$\lambda|X, \theta = \exp(X\theta) \quad (7)$$

The number of observed spikes, n , in any given time bin is assumed to a Poisson process with an instantaneous firing rate mean of λ :

$$n|\lambda \sim \text{Poisson}(\lambda dt) \quad (8)$$

Covariate inputs to the GLMs

In our analyses, we used GLMs to predict the spiking activity of single neurons based on the activity of the remaining population and kinematic signals. We binned the neural spikes at 50 ms intervals and downsampled the continuous kinematic signals to 20 Hz to match the binned spikes. We shifted the kinematic signals backwards in time by three spiking bins (150 ms) to account for transmission delays between cortical activity and the motor output. Previous studies have observed a broad range of delays⁸, so we convolved the kinematic signals with raised cosine basis functions centered at 0 ms and -100 ms, adapting the method of Pillow et al., where bases further back in history become wider⁹. By including these convolved signals as inputs to our GLM models, we allowed the neurons to have more flexible temporal relationships with the kinematics. Note that all GLM models included the same convolved endpoint position, velocity, and acceleration signals as covariates.

We trained two types of models: the Basic models included only kinematic covariates, while the Full models included both the kinematic covariates and the spiking activity of the single-neuron populations (Figure S3). For the GLMs with single neuron inputs (Figure 2), we trained three different types of Full models. M1-M1 models predicted the spiking activity of each M1 neuron from the activity of all other M1 neurons recorded on the same session, PMd-PMd models predicted the spiking of each PMd neuron from all other PMd neurons, and PMd-M1

models predicted M1 neurons using the activity of all PMd cells. For the GLM analysis with potent and null components (Figure 3), we used low-dimensional summaries of PMd population activity as inputs to the GLMs, rather than single neurons. For PCA-M1, we projected PMd activity into the PCA space (see above) and selected only the first 16 dimensions as input to the GLM. Since PCA captures population-wide covariance patterns, we expected that this approach would provide nearly identical results to the single neuron models of PMd-M1, and it was included primarily as a control. For Pot-M1 and Null-M1, we projected the time-varying PMd signals onto the basis vectors for the potent and null space, respectively (see above). We then used these time-varying signals as inputs to GLMs to predict the spiking of M1 neurons. It is worth noting that although we defined the Null space as activity which produced zero output in the low-dimensional M1 components, we could still predict M1 spiking quite well from the Null space in the cross-validated training data (Figure S3). Although potentially unintuitive, it worked well for a number of reasons. First, we identified the potent and null spaces using population-wide components, and used these activity patterns to predict the spiking single neurons. Additionally, for a given reach direction within a condition, the stereotyped activity in the null space could be well-correlated with activity in the potent space (and subsequently M1) due to the lawful relationship between them^{4,10}. It is under the condition of changing behavior that these correlations can begin to break.

Training the GLMs

We trained the models using the last 50% of CF or VR trials when behavior was most stable, including only trials where the monkeys made successful reaches to acquire the outer target (reward trials). This allowed us to test the generalization of the GLMs during the early adaptation trials. For the CF, it was important to both train and test the GLMs using trials from the CF epoch to avoid extrapolating between the Null and CF conditions. When we imposed the CF, it changed the relationship between the kinematics and dynamics of limb movement. Thus, if we trained the GLM on Baseline trials, the relationship between kinematics and neural activity changed immediately on CF trials¹¹, leading to poor GLM generalization for all models. By both training and testing within the block of CF trials, we avoided the problem of extrapolating to new dynamics conditions. Although the VR sessions did not have this problem, we adopted this same approach the sake of consistency.

We trained the models using a maximum likelihood method (glmfit in Matlab, The Mathworks Inc). In the case of our full population spiking models, we had dozens to hundreds of covariate inputs for a single predicted output. Although we had very large numbers of training datapoints (typically on the order of 10,000 samples), there is the possibility our models were impaired by overfitting. We guarded against overfitting using ten-fold cross-validation of our training dataset. We also repeated our analyses using Lasso GLM for regularization and observed nearly identical results (data not shown). We thus chose to use the non-regularized GLM for simplicity and to reduce the computational load, since it did not impact our results.

Evaluating GLM performance

We evaluated GLM performance using a particular formulation of the pseudo- R^2 (pR^2). The pR^2 is analogous to the R^2 commonly used in model-fitting with Gaussian statistics, but it is generalized to incorporate the assumed Poisson statistics of the neural spiking data:

$$pR^2 = 1 - \frac{\log L(n) - \log L(\hat{\lambda})}{\log L(n) - \log L(\bar{n})} \quad (9)$$

The pR^2 finds the difference in log-likelihood between the observed spiking data (n) and the model predictions ($\hat{\lambda}$). This value is compared against the difference in log-likelihood for the mean of the dataset (\bar{n}). We used the Likelihood (L) for Poisson data according to:

$$L = \prod_{t=1}^T \text{Poisson}(n_t | \lambda_t) = \prod_{t=1}^T \frac{\lambda_t^{n_t} \exp(-\lambda_t)}{n_t!} \quad (10)$$

And thus, the log-likelihood ($\log L$) across all time bins (t) of a given spike train is:

$$\log L = \sum_{t=1}^T (n_t \log \lambda_t - \lambda_t - \log n_t!) \quad (11)$$

Although the upper bound for pR² is one, poor model fits can be less than zero. A pR² of one indicates a perfect model fit, a value of zero indicates that the model prediction performs as well as finding the mean of the data, while values less than zero indicate that the model performed worse than merely fitting the mean. Typical pR² values are smaller in magnitude than those typically found with the Gaussian R². When evaluating GLM fits, we used a bootstrapping procedure with 1000 iterations to obtain 95% confidence bounds on the pR² value. We considered a model fit to be significant if this bootstrapped confidence interval was above zero, indicating that the model helped to explain the spiking activity. For many analyses, we used the relative pseudo-R² (rpR²), which directly compares two separate GLM models. While pR² compared the log-likelihood of the model predictions to the mean of the data, the rpR² compares the predictions of a Full model to a Basic model with fewer covariates.

$$\text{rpR}^2(\text{Basic, Full}) = 1 - \frac{\log L(n) - \log L(\hat{\lambda}_F)}{\log L(n) - \log L(\hat{\lambda}_B)} \quad (12)$$

Here, $\hat{\lambda}_F$, the Full model prediction, which includes both the kinematics and the population spiking, is compared to $\hat{\lambda}_B$, the prediction of the Basic model, which includes only kinematics. This metric thus quantifies the improvement in performance afforded by the additional neuronal inputs. Positive values indicate that the Full model performed better than the Basic model, while negative values indicate that predictions were better with kinematics alone. As with the pR², we obtained confidence bounds with a bootstrapping procedure and assessed significance by determining if the lower bound was above zero. This indicated that the addition of population spiking added information over the kinematics alone, and thus could be capturing meaningful functional relationships between the population and the predicted cell.

For the time course plots, such as Figure 2h, we predicted neural spiking on individual trials. However, predictions could be quite noisy with such small numbers of datapoints. For example, if a cell fired very few spikes on a particular trial, the pR² may be quite low, even though the model generally performed quite well. To remove some of this variability, we smoothed the trial-to-trial predictions for each neuron (as well as the overlaid behavior) with a moving average. We chose a window of 30 trials, though we observed similar (but slightly more variable) traces even down to window sizes of 5-10 trials. Since there were rapid behavioral improvements in the early trials, we padded the beginning and end with NaNs, each of a length of half of the window size. This helped to prevent averaging out the changing behavioral effects, with the tradeoff of slightly increasing noise. In practice, our results were similar without this padding.

Selecting cells with significant population relationships

For most of our analyses, we studied cells that were well-predicted by our GLMs. We determined this by two main criteria using ten-fold cross-validation on the training data. First, we required that the Basic pR² was significantly above zero. This reduced the pool of candidate cells to 522/918 (57%) in M1 and 612/2221 (28%) in PMd, but was necessary so that the rpR² would be well defined. Qualitatively, we obtained similar results when we relaxed this criterion to include more cells. We also required that the rpR² was significantly above zero. We only included cells that were significantly above zero for all ten of the folds for all pR² and rpR². This method was very conservative, but ensured that we only studied cells that were reliably predicted.

Statistical tests

For the GLM models, we assessed the significance of model fits empirically using a bootstrapping procedure on cross-validated data as described above. We additionally used two-sample Student's t-tests to compare the distributions of pseudo-R² changes in Early and late learning. For Figure 2f, this was done using the raw rpR² values. For Figures 2g, 3d, 4c, and Extended Data Figure 5g, the t-test was done using the normalized change in rpR².

Additional References

1. Huang, V. S., Haith, A., Mazzoni, P. & Krakauer, J. W. Rethinking Motor Learning and Savings in Adaptation Paradigms: Model-Free Memory for Successful Actions Combines with Internal Models. *Neuron* **70**, 787–801 (2011).
2. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* (2014). doi:10.1038/nn.3776
3. Machens, C. K., Romo, R. & Brody, C. D. Functional, but not anatomical, separation of ‘what’ and ‘when’ in prefrontal cortex. *J. Neurosci.* **30**, 350–60 (2010).
4. 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**, 440–8 (2014).
5. Michaels, J. A. & Scherberger, H. hebbRNN: A Reward-Modulated Hebbian Learning Rule for Recurrent Neural Networks. *J. Open Source Softw.* (2016). doi:http://dx.doi.org/10.21105/joss.00060
6. JA, N. & R, B. Generalized linear models. *Encyclopedia of Statistical Sciences* (1972).
7. Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13**, 105–11 (2010).
8. Moran, D. W. & Schwartz, A. B. Motor cortical representation of speed and direction during reaching. *J. Neurophysiol.* **82**, 2676–92 (1999).
9. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–9 (2008).
10. 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.* 13239 (2016). doi:10.1038/ncomms13239
11. Cherian, A., Fernandes, H. L. & Miller, L. E. Primary motor cortical discharge during force field adaptation reflects muscle-like dynamics. *J. Neurophysiol.* **110**, 768–83 (2013).