Neuron, Volume *54*

Supplemental Data

Motor Learning with Unstable

Neural Representations

Uri Rokni, Andrew G. Richardson, Emilio Bizzi, and H. Sebastian Seung

Changes in Tuning Curves Are Approximated by Changes in Offset and Cosine Components

The directional tuning curve of a neuron in our data is described by the set of 8 mean firing rates for the different movement directions. Changes in the mean firing rates across different directions may be correlated. In this case, changes in tuning curves can be compactly described by a few principal components (PCs), rather than the 8 mean firing rates. To examine this issue, for each neuron we computed a "difference tuning curve", i.e. the changes from block 1 to block 3 in mean firing rate for each direction. We then performed principal component analysis (PCA) on the set of difference tuning curves. We found that the first 3 PCs account for 93% of the variance in the changes in tuning curves (Figure S1A). Figure S1B shows the first three PCs. The first PC (solid line) approximately corresponds to the offset of the tuning curve and the second and third PCs (dashed-dotted and dashed lines, respectively) approximately correspond to the cosine and sine components of the tuning curve. The other PCs have more complicated forms (data not shown). We conclude that changes in tuning curves in the control experiment can be reasonably approximated by changes in the tuning curves' offsets and cosine components.

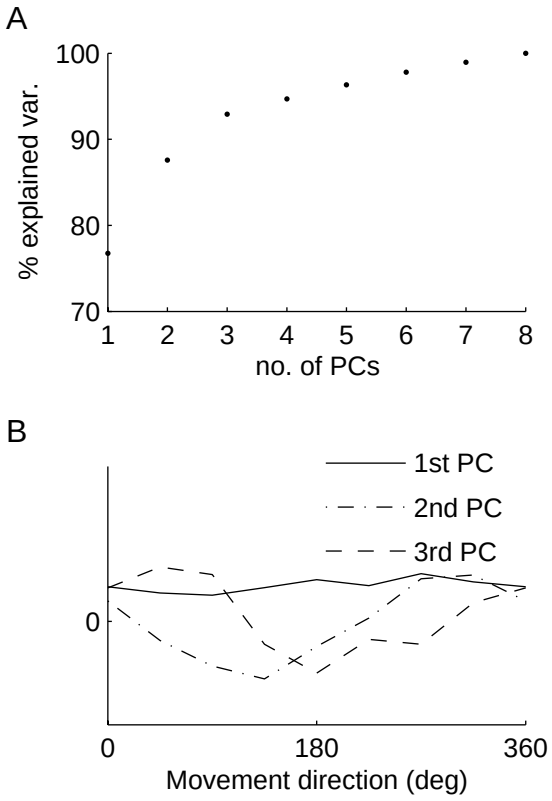

Figure S1. Principal Component Analysis of Changes in Tuning Curves
(A) Percentage of explained variance as a function of number of principal components.
(B) First three principal components. The components have been scaled such that they have the same root mean square, and therefore the y-axis has arbitrary units.

Comparison of Changes in Tuning Curves with Behavioral Changes

We tested whether the drift in directional tuning curves observed in the control experiment is related to drift in behavior. Overall, we had data from 42 control session, 15 for monkey 1 (monkey with SMA data) and 27 for monkey 2 (monkey with MI data). We searched for correlations between changes in tuning curves and changes in hand trajectories across days. We used three measures to quantify changes in tuning curves in a given session: (1) change in PD, (2) change in modulation depth, and (3) change in offset. For each of these measures we averaged the changes from block 1 to block 3 over the cells recorded in each session. Additionally, we selected measures of behavioral change in two ways: (1) we chose behavioral measures which we judged as likely to be related to the neural measures, and (2) we performed principal component analysis (PCA) of the changes in trajectories within control sessions to find the aspects of movement which changed the most.

We performed a total of 18 comparisons between measures of changes in tuning curves and measures of changes in hand trajectories. Table S1 shows the p-value of the correlation found in each comparison, and Table S2 shows for each behavioral measure what fraction of sessions showed a significant change. Out of the 18 comparisons we have made, only a single comparison had a p-value less than 0.05, which is expected by chance. We conclude that we did not find a behavioral correlated of the changes in tuning curves in the control experiment.

While there are many other comparisons that could be done, it is important to keep in mind that adding more comparisons lowers the statistical power of the overall experiment. Additionally, we emphasize that these results still do not rule out entirely the possibility that the neural changes are related to behavioral changes. First, it is possible that more data is required to detect such correlations. Second, it is possible that the neural changes are related to aspects of the behavior unobservable in the hand kinematics, e.g. changes in arm posture or muscle activity. Next, we explain the different comparisons we have made.

PDs vs. Deviation Areas

It is natural to assume that changes in the cells' PDs are related to changes in the direction of movement. Because the deviation area measure captured well the changes in movement direction caused by the forces in the learning experiment (see Figure 1D in main manuscript), we used it also to describe changes in movement direction in the control experiment. We quantified the change in direction of movement on a given session by the change in the mean deviation area from block 1 to block 3. We compared the changes in deviation area with the PD changes, across all control sessions, pooled over both monkeys (see Table S1).

PDs vs. Principal Components of Tangential Velocity Change

Our choice of the deviation area measure for quantifying changes in movement directions is somewhat arbitrary. Another approach for quantifying changes in movement direction is to find the measures that changed the most within the control sessions. For this purpose, we performed the following principal component analysis (PCA). First, we examined the tangential velocity component, i.e. the component of the velocity orthogonal to the axis between the initial hand position and the target (positive tangential velocity corresponds to movement in the counter-clockwise direction). The temporal profile of the tangential velocity is a relatively sensitive gauge of changes in movement direction. Figure S2A shows the tangential velocity profiles averaged over trials in block 1 (black), and averaged over trials in block 3 (grey), in one control session. The change from block 1 to 3 (dashed) reflects changes that occurred in movement direction. We started our analysis by computing the change from block 1 to 3 in the tangential velocity profile for each control session.

Next, we performed PCA on the set of profiles of change in tangential velocity, separately for each monkey. In both monkeys we found that the first two principal components (PCs) account for roughly 50% of the variance of these profiles (Figure S2B), and therefore in subsequent analysis we focused on these two PCs. The profiles of these PCs are shown in Figure S2C. To facilitate the interpretation of these PCs, we computed the mean trajectory of each monkey (Figure S2D, black), and then perturbed these mean trajectories by adding or subtracting each PC (grey). Because the effects on the trajectories were rather small, we artificially magnified the PCs in Figure S2D by a factor of 4, i.e. the typical changes in trajectories were 4 times smaller than shown in the figure. In both monkeys we found that the 1st PC is

related to clockwise/counter-clockwise deviations across the entire trajectory, whereas the 2nd PC is related to deviations mostly at the end of the trajectories.

Subsequently, for each control session we computed the coefficients of the two PCs. These coefficients represent to what extent each PC contributed to the change in the tangential velocity profile. Finally, we used these coefficients as our new measures of change in movement direction, and compared them with the changes in PDs, across sessions, separately for the two monkeys (2 PCs x 2 monkeys = 4 comparisons; see Table S1). To judge the statistical significance of the coefficient of a given PC on a given session, we computed the coefficients of that PC for all individual trials in blocks 1 and 3 of that session, and used a z-test to test the statistical significance of the difference in the means between blocks 1 and 3.

Modulation Depths and Offsets vs. Movement Speed

Because it has been shown that the modulation depths and offsets of directional tuning curves of motor cortical cells scale with hand speed (Moran and Schwartz, 1999), we compared the changes in modulation depths and offsets to changes in hand speed. To quantify the change in hand speed we first computed the average hand speed for each trial by dividing distance from starting point to target by total movement time. Movement onset was defined as time when speed exceeds 4cm/s, and movement offset was defined as time when speed becomes less than 4cm/s. For each control session, we averaged the mean hand speeds within trials, across trials in block 1 and across trials in block 3, separately in each block. The difference between the mean speeds in blocks 1 and 3 was our measure of speed change. We compared this measure of speed change with the offset changes and modulation depth changes, across all control sessions, pooled over both monkeys (2 comparisons; see Table S1).

Modulation Depths and Offsets vs. Principal Components of Radial Velocity Changes

Above we supplemented the measure of change in deviation area with measures of change in movement direction derived by PCA of the changes in tangential velocity. Following a similar approach, we supplemented the measure of change in hand speed with measures derived by PCA of the changes in radial velocity (the component along the axis from the starting point to the target). While changes in PDs are more likely related to changes in tangential velocity, changes in modulation depths and offsets are more likely related to changes in radial velocity. Figure S3A shows for one example session of each monkey the mean radial velocity profiles in block 1 (black), block 3 (grey), and the change from block 1 to 3 (dashed). We performed PCA on the set of profiles of radial velocity change across all sessions, separately for each monkey. We found that the first two PCs account for roughly 60% of the variance of the changes in the radial velocity (see Figure S3B). Figure S3C shows the profiles of change in radial velocity which correspond to these PCs. To visualize the effect of these PCs on the movement we computed the mean radial velocity profile of each monkey (Figure S3D, black), and perturbed this mean profile by adding or subtracting each PC (grey). Because the effects on the radial velocity profiles were small, we artificially magnified the PCs in Figure S3D by a factor of 4, i.e. the typical changes in trajectories were 4 times smaller than shown in the figure. The two monkeys showed qualitatively different PCs. Monkey 1 showed changes in movement time (PC 1), and increases/decreases in the duration of the acceleration phase coupled with overshooting/undershooting the target (PC 2). Monkey 2 showed changes in the shape of the velocity profile.

We computed the coefficients of these PCs for each session, and used them as our new measures of movement changes. We compared these coefficients with the changes in modulation depths and offsets across sessions, separately for each monkey (2 PCs x 2 neural measures x 2 monkeys = 8 comparisons; see Table S1). The statistical significance of the coefficients was judged by the same method used for the coefficients of the PCs of the tangential velocity (see above).

PDs, Modulation Depths, and Offsets vs. Hit Rates

The above tests included only successful trials. It could be argued that the changes in the neuronal tuning are related to the changes in the probability of success. After all, getting the reward is probably the aspect of the task that is most important to the monkeys. Within each block in each control session we computed the percentage of successful trials, which we term the hit rate. Hit rates across control sessions were 95%±3% for monkey 1, and 84%±6% for monkey 2 (mean±s.d.). Changes in hit rates from blocks 1 to 3 were 1.5%±3.6% for monkey 1 and 0.3±6.9% for monkey 2 (mean±s.d.). We compared the changes in hit

rates from blocks 1 to 3 with the changes in PDs, modulation depths and offsets, pooled over both monkeys (3 comparisons; see Table S1).

Table S1. p Values of Comparisons (Permutation Test)

	PDs	modulation depths	offsets
deviation areas (both monkeys)	0.08		
monkey 1, tangential PC 1	0.31		
monkey 1, tangential PC 2	0.02		
monkey 2, tangential PC 1	0.43		
monkey 2, tangential PC 2	0.35		
mean speed (both monkeys)		0.86	0.74
monkey 1, radial PC 1		0.79	0.60
monkey 1, radial PC 2		0.78	0.55
monkey 2, radial PC 1		0.56	0.94
monkey 2, radial PC 2		0.12	0.12
hit rates (both monkeys)	0.69	0.36	0.99

Table S2. Fraction of Control Sessions with Significant Behavioral Change from Block 1 to 3 (z Test, $p < 0.01$)

measure	fraction of significant changes
deviation areas (both monkeys)	21%
monkey 1, tangential PC 1	20%
monkey 1, tangential PC 2	27%
monkey 2, tangential PC 1	18%
monkey 2, tangential PC 2	18%
mean speed (both monkeys)	62%
monkey 1, radial PC 1	60%
monkey 1, radial PC 2	40%
monkey 2, radial PC 1	86%
monkey 2, radial PC 2	36%
hit rates (both monkeys)	14%

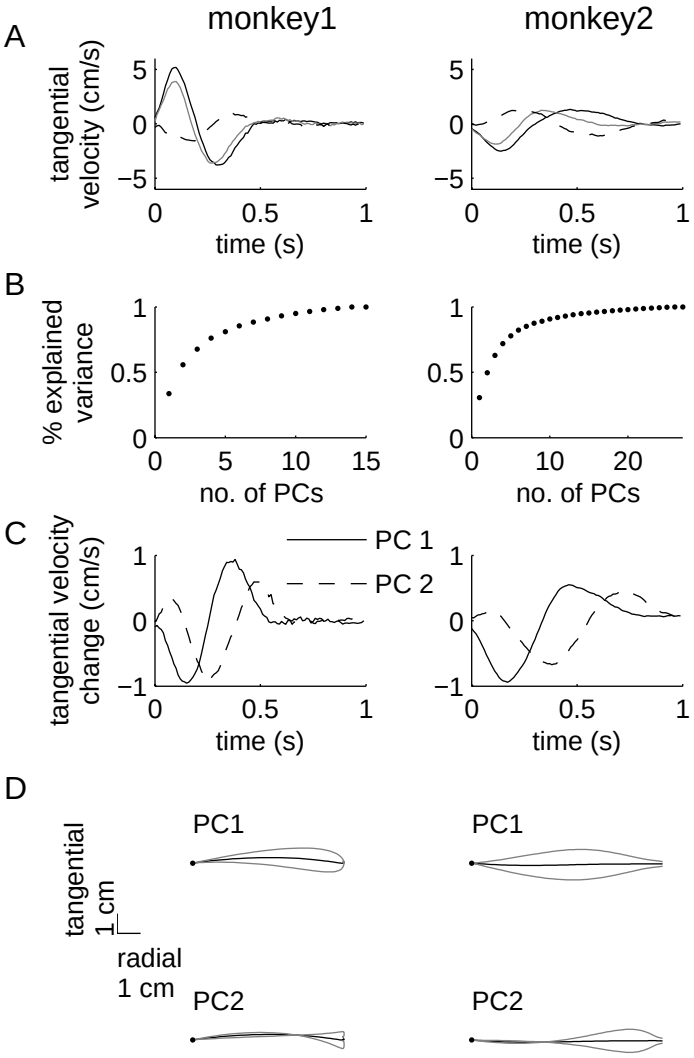

Figure S2. PCA of Tangential Velocity Profiles
Left column corresponds to monkey 1 and right column corresponds to monkey 2.
(A) Mean tangential velocity profile, averaged over block 1 of an example session (black), averaged over block 3 of the same session (grey), and the difference between blocks 1 and 3 (dashed).
(B) Percentage of explained variance of the changes from block 1 to 3, in the tangential velocity profiles, as a function of the number of PCs.
(C) Velocity profiles of first two PCs.
(D) Mean trajectories (averaged over blocks 1 of all sessions) (black), and the effect of adding/subtracting each PC, magnified by 4 (grey). The dot represents the starting point, the x-axis represents the radial axis from starting point to target, and the y-axis represents the tangential axis.

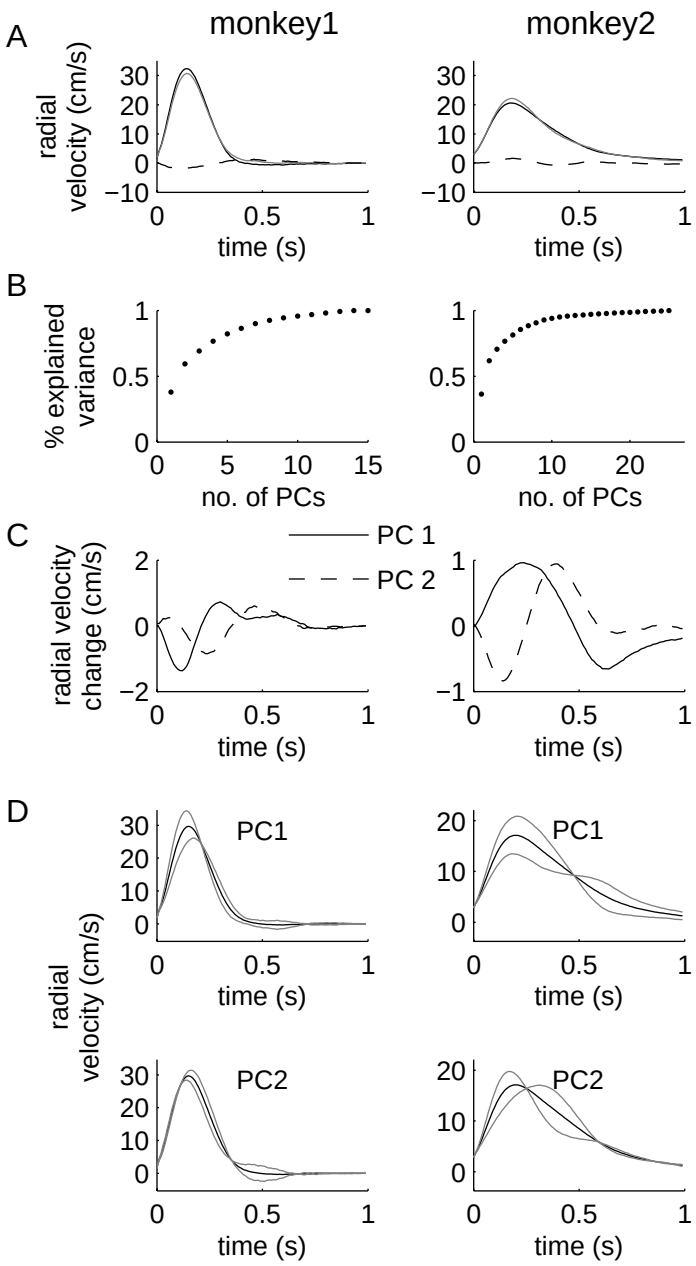

Figure S3. PCA of Radial Velocity Profiles
Left column corresponds to monkey 1 and right column corresponds to monkey 2.
(A) Mean radial velocity profile, averaged over block 1 of an example session (black), averaged over block 3 of the same session (grey), and the difference between blocks 1 and 3 (dashed).
(B) Percentage of explained across-session variance of the changes from block 1 to 3 in the mean radial velocity profiles as a function of the number of PCs.
(C) Velocity profiles of first two PCs.
(D) Mean radial velocity profiles averaged over blocks 1 of all sessions (black), and the effect of adding/subtracting each PC, magnified by 4 (grey).

Global Noise from the Environment Causes Correlated Changes across Cells

In the manuscript, we have shown that the changes in tuning curves in the control experiment are random across cells. Additionally, we have shown that randomness across cells can be reproduced by a model in which the source of variability in plasticity is local, i.e. independent across synapses. Here we examine the effect of global plasticity noise, i.e. noise from the environment which through sensory feedback contaminates the learning signal. For this purpose, we modify the plasticity rule (Equation 10 in main text) to

$$\Delta W_{ij} = -\frac{N}{\tau_{learn}} \frac{\partial E}{\partial W_{ij}} \left(x + \sigma n, x^t \right) \tag{S1}$$

where n is Gaussian noise that is added on top of the model output vector x, and σ is the noise amplitude, which we set to 1. n may be interpreted as either sensory noise, i.e. noise in the estimation of the hand position, or as muscle noise. We simulated the control experiment with this modified plasticity rule. Figure S4 shows that in this simulation the PDs of different cells change in a correlated manner. Thus, global noise from the environment causes changes in tuning curves which are correlated across cells.

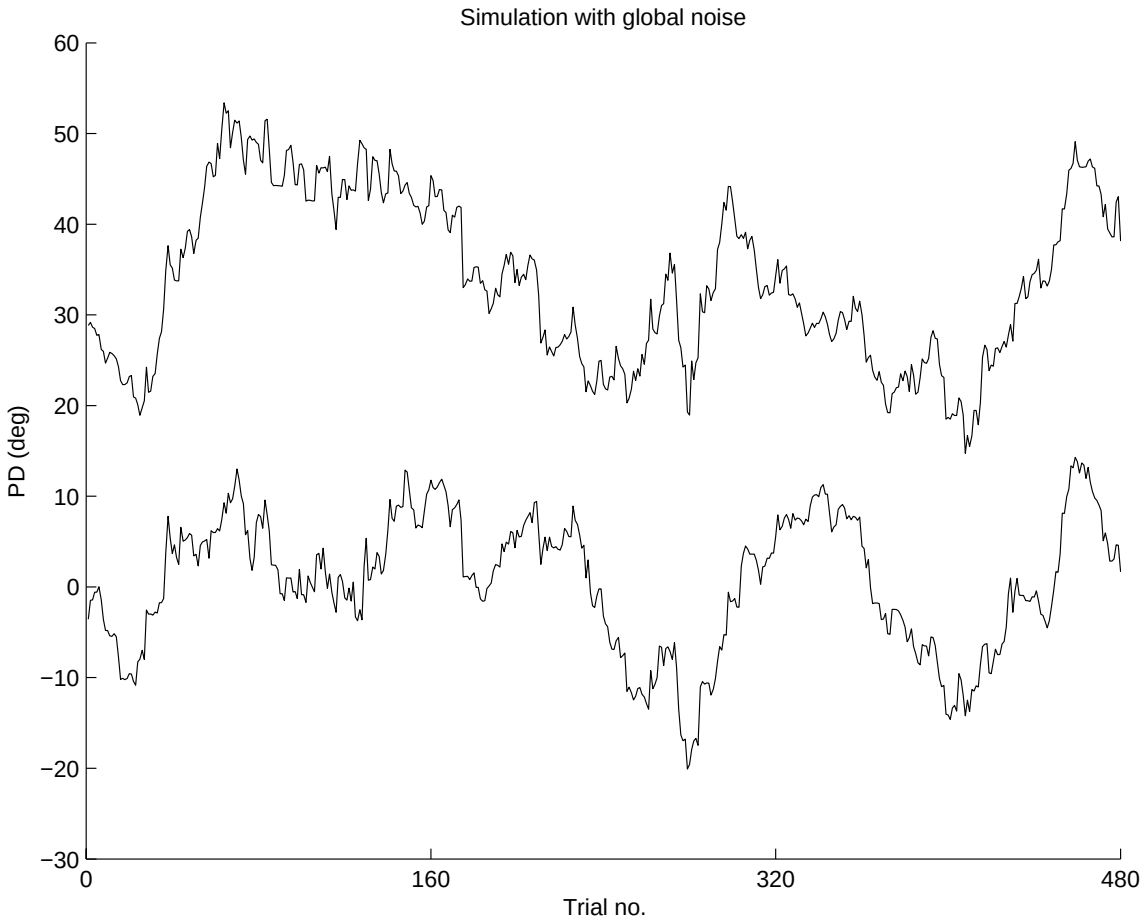

Figure S4. Simulation of Model with Global Noise from the Environment
PDs of two model cells as a function of trial number.

Statistics of Modulation Depths Are Accounted by Including Plasticity of Neuronal Excitability

Problems with Original Model

The model we have presented in the main text accounts reasonably well for the statistics of the PDs, but does not account well for the statistics of the modulation depths. In order to compare the predicted statistics of the modulation depths to the experimental data, first we scaled the firing rates of the model cells to match the empirical mean modulation depth in the control condition (~5 Hz). Next, we compared the predicted distribution of modulation depths in the control condition to the empirical distribution (Figure S5A, middle and left, respectively). The model predicted the distribution is peaked around modulation depths of intermediate size, whereas the true distribution is peaked at small modulation depths. Additionally, we examined the distribution of changes in modulation depths from block 1 to 3 in the control experiment (Figure S5B, middle and left, respectively). We found that the model underestimates the changes in modulation depths. Similarly, we found that the model underestimates the changes in modulation depths during adaptation and washout in the learning experiment, although it does account for the anti-correlation between these changes (Figure S5C, middle and left, respectively). Finally, in all the distributions shown in Figure S5, the data has much heavier tails than the predicted distributions.

The Modified Model

We found that the model's discrepancies in the statistics of the modulation depths can be remedied by assuming an additional plasticity mechanism that alters neuronal excitability. In the original model, changes in tuning curves occur by synaptic modifications. The level of noise in these synaptic modifications was tuned to reproduce the degree of PD changes observed experimentally. Under these conditions, the model underestimated the changes in modulation depths. One possible explanation for this discrepancy is the existence of an additional plasticity mechanism which changes modulation depths but not PDs. In order to modify a cell's modulation depth without altering its PD the cell may change its excitability, and thus modify uniformly the throughput of all its synapses. Persistent changes in neuronal excitability have been reported previously, and were found to be correlated with learning in several experimental preparations (Zhang and Linden, 2003). These experience-dependent excitability changes are a result of the modification of voltage-gated ion channels responsible for action potential generation and propagation. Thus, our modified model may be viewed as storing task-related information through both synaptic and intrinsic modifications.

To test whether noisy learning in both synapses and neuronal excitability can account for the experimental data we extended the original model to include also plasticity of neuronal excitability. In the original model the gain between the sensory input x_j^t and the firing rate of a cortical cell r_i was determined by the synaptic weight W_{ij} (see Equation 6 in main text). In the modified model, we substitute W_{ij} with $g_i W_{ij}$, where g_i is the excitability of neuron *i*

$$r_i = \sum_{j=1}^{2} g_i W_{ij} x_j^t \tag{S2}$$

Similar to the original model, we use noisy gradient descent learning to train the model parameters (see Equation 10 in main text). However in the modified model we train both W_{ij} and g_i

$$\begin{aligned} \Delta W_{ij} &= -\frac{W_{ij}}{\tau_{forget}} + \sigma n_{ij} - \frac{N}{\tau_{learn}} \frac{\partial E\left(x, x^t\right)}{\partial W_{ij}} \\ \Delta g_i &= -\frac{g_i}{\tau_{forget}^g} + \sigma^g n_i^g - \frac{N}{\tau_{learn}} \frac{\partial E\left(x, x^t\right)}{\partial g_i} \end{aligned} \tag{S3}$$

where g_i are constrained to be positive. n_{ij}, n_i^g are independent normalized unbiased Gaussian noise terms. The modified model has the same parameters as the original model and two additional parameters:

the amplitude of the noise in the plasticity of the excitabilities σ^g, and the decay time of the excitabilities τ^g_{forget}. The learning time constants of the excitabilities and the synapses are assumed the same (denoted τ_{learn}) and equal to the typical learning time observed in the monkeys' behavior (50 trials). As in the original model the cost that is optimized is the squared output error

$$E\left(x,x^t\right)=\frac{1}{2}\sum_{i=1}^{2}\left(x_i-x_i^t\right)^2 \tag{S4}$$

The gradients of this cost with respect to W_{ij} and g_i are related to the error by

$$\begin{aligned}\frac{\partial E\left(x,x^t\right)}{\partial W_{ij}}&=\sum_{k=1}^{2}\left(x_k-x_k^t\right)Z_{ki}g_i x_j^t\\ \frac{\partial E\left(x,x^t\right)}{\partial g_i}&=\sum_{k=1}^{2}\left(x_k-x_k^t\right)Z_{ki}\sum_{j=1}^{2}W_{ij}x_j^t\end{aligned} \tag{S5}$$

As in the simulations of the original model, before simulating the experiments we trained the model for 10000 trials, thus mimicking the monkeys' previous experience. The performance of the model was as good as that of the original model (not shown). Adding the excitability changes had a relatively small effect on the changes in PDs and therefore the parameters of the original model have not been changed much[1] (Table S3). We tuned the parameters σ^g, τ^g_{forget} to match the distributions of changes in modulation depths in the control and learning simulations with the empirical distributions.

Results of Modified Model

We found that the modified model accounts much better than the original model for the observed statistics of the modulation depths. Remarkably, without fine tuning of parameters the modified model produced a distribution of modulation depths which is peaked at small values and is very similar to the empirical distribution (Figure S5A). This distribution of modulation depths was established in the pre-training stage of the model, prior to the simulation of the control experiment, and regardless of initial conditions. In the original model cells with weak modulation depths are rare. This property of the original model is related to the structure of the inputs to the cortical cells. Each cortical cell sums sine and cosine functions of direction, weighted by two synaptic weights (see Experimental Procedures in main text). Under these conditions, a cell's modulation depth is weak when a rare coincidence occurs in which both input synapses are weak simultaneously. More generally, when a cell receives multiple synaptic inputs with different directional tuning, the cell has a weak modulation depth when a rare configuration of synaptic weights occurs, such that the input tuning curves nearly cancel out. In contrast, in the modified model a neuron may have a weak modulation depth whenever its excitability is weak. In other words, in the original model a weak modulation depth requires several (two) variables to be weak simultaneously, whereas in the modified model it is sufficient that one variable, namely excitability, is weak. This occurs often as noise changes excitabilities, and consequently weak modulation depths are abundant.

Additionally, the modified model accounts much better for the distribution of the modulation depth changes in the control experiment (Figure S5B). As expected, changes in modulation depths became larger than in the original model because of the added contribution of the excitability changes. To account well for the size of the modulation depth changes in the control condition we had to assume that random excitability changes occur on a faster time constant than synaptic changes ($\tau^g_{forget}=500, \tau_{forget}=1500$). The reason

[1] σ has been slightly reduced because the excitability changes cause the PD changes to increase a little bit, indirectly. In the original model the random PD changes are slowed down by the learning signal as it makes corrections to the tuning curves in order to preserve the performance. In the modified model with the excitability changes, part of these corrections is done by excitability changes, thus less corrections are done in the synapses, and the PDs are somewhat more free to change.

for this difference in time constants is that empirically, modulation depths had larger fractional changes than PDs. Changes in modulation depths had a standard deviation of ~3 Hz, which is roughly 60% of the mean modulation depth (~5 Hz). In contrast, PD changes had a standard deviation of ~30˚, which is roughly 15% of the maximal possible change of 180˚.

Adding the excitability changes also produced heavy tails in the distribution of changes in modulation depths, similar to the heavy tails observed in the experimental data (Figure S5B). These heavy tails result from the product of synaptic changes and excitability changes. In general, taking the product of two random variables tends to produce heavy tails, because occasionally the two variables simultaneously deviate from the average moderately, which causes the product to deviate from the average strongly. Thus, according to our model the heavy tails in the distribution of modulation depth changes reflect coincidences between excitability changes and synaptic changes. For similar reasons, our model produces heavy tails also in the distribution of the modulation depth changes in the learning experiment (Figure S5C) and in the distribution of the modulation depths (Figure S5A).

We found that in order to account for the changes in modulation depths in the learning condition (Figure S5C) we must assume that the excitabilities are tuned to improve performance. When we simulated the learning experiment without the excitability gradient term in Equation S3, but still included the random changes in g_i, the model underestimated the size of the changes in modulation depths and the predicted anticorrelation between the adaptation and washout changes was much weaker than the empirical anticorrelation (not shown). Thus, according to our model excitability changes are not purely random, but are in fact tuned to improve the behavior.

The modified model still accounts reasonably well for the statistics of the changes in PDs (Figure S6). The only aspect in which the modified model is slightly worse than the original model, is that the predicted anticorrelation between the adaptation and washout PD changes is somewhat weaker than the observed anticorrelation and the anticorrelation predicted by the original model (compare Figure S6E with Figures 3E, 6E in main text). Because in the modified model part of the learning is achieved by excitability changes, it has somewhat less learning related changes in PDs, and therefore less anticorrelation between adaptation and washout PD changes.

Table S3. Parameter Values Used in Simulations of Modified Model

τ_{learn}	τ_{forget}	σ	τ^{g}_{forget}	σ^{g}	φ	N
50	1500	0.02	500	0.018	60˚	10000

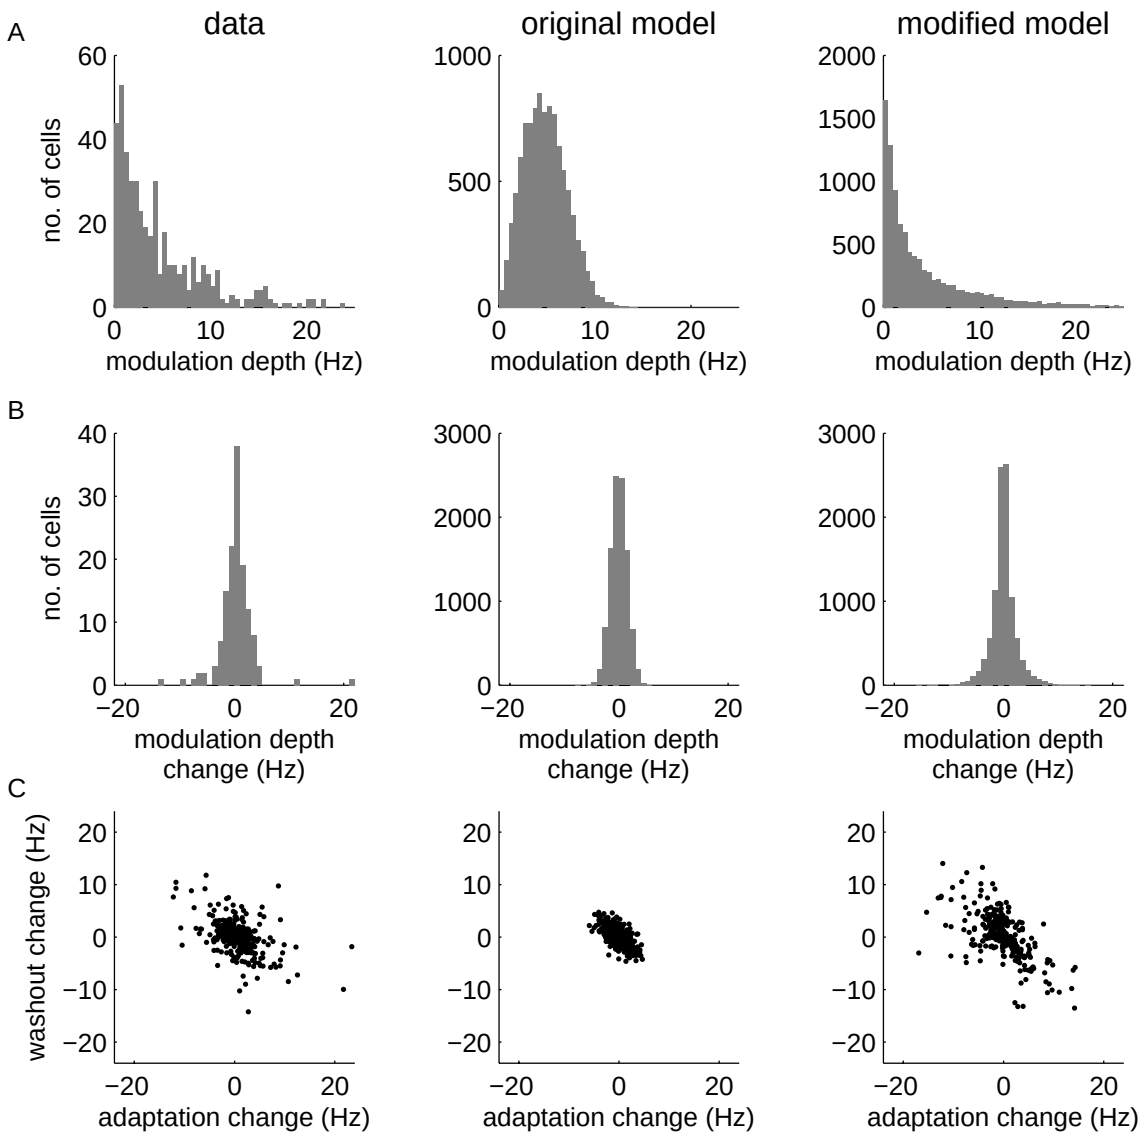

Figure S5. Statistics of Modulation Depths in Experimental Data (Left), Original Model (Middle), and Modified Model (Right)
(A) Distribution of modulation depths across cells measured from activity in block 1 of control experiment/simulation.
(B) Distribution of changes in modulation depths across cells from block 1 to 3 in control experiment/simulation.
(C) Washout change in modulation depth (last 80 trials of adaptation block to last 80 trials of washout block) vs. adaptation change in modulation depth (baseline block to last 80 trials of adaptation block) across cells in learning experiment/simulation.

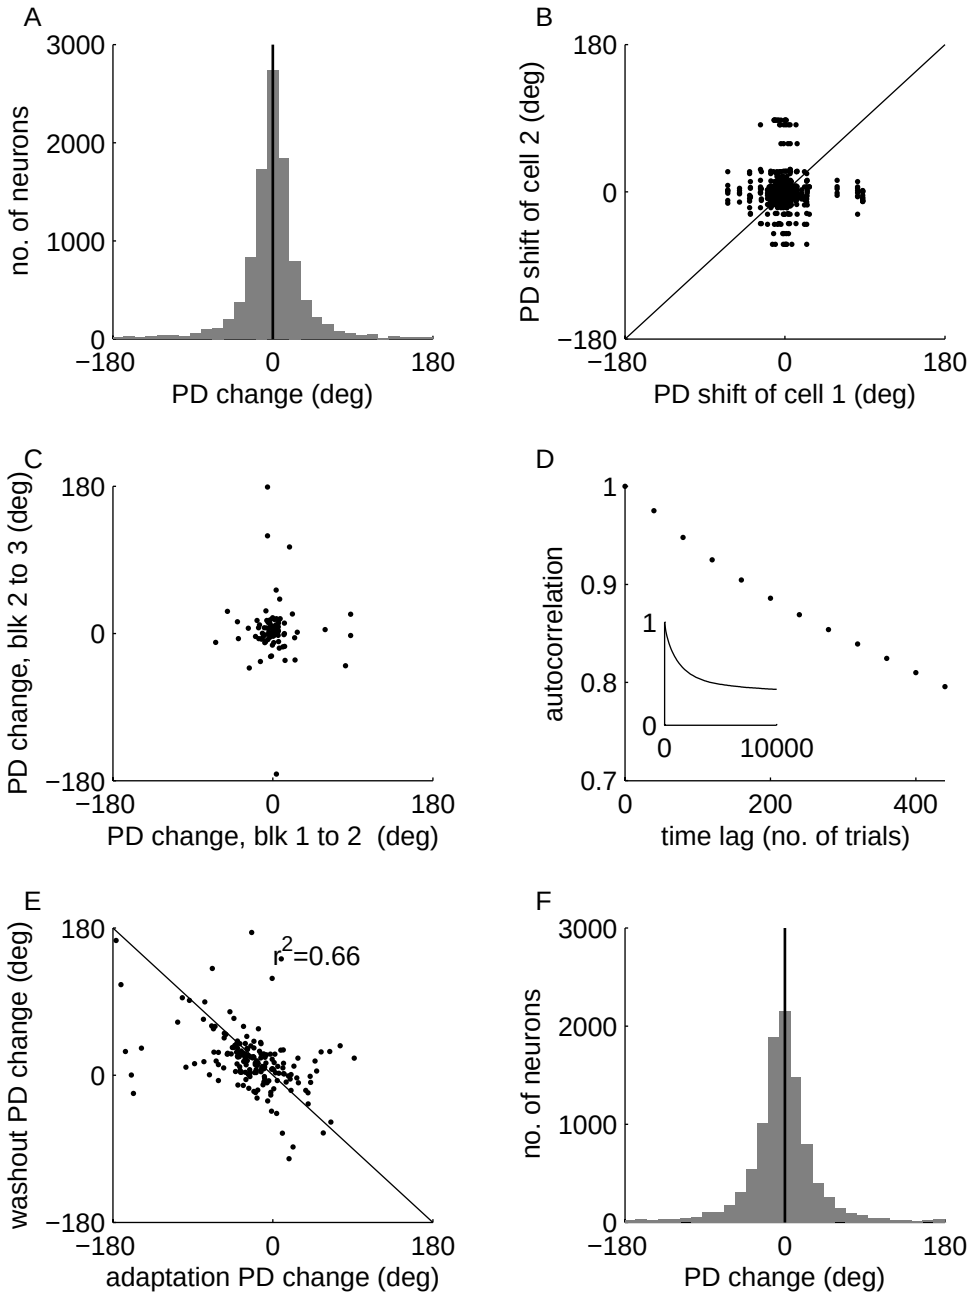

Figure S6. Statistics of Changes in Tuning Curves in Simulations of Control and Learning Experiments with the Modified Model (for Parameter Values See Table S3)

(A) Distribution across cells of PD changes from block 1 to 3 in control simulation.

(B) PD changes from block 1 to 3 of pairs of cells. Each pair is represented by two points symmetrically positioned around the $y = x$ diagonal (solid line).

(C) PD change from block 1 to 2 vs. PD change from block 2 to 3, across cells in control simulation.

(D) Autocorrelation of population of PDs. (D, inset) Autocorrelation over long times (E) Adaptation PD changes vs. washout PD changes, across cells in learning simulation. Solid line represents the $y = -x$ diagonal. (F) Distribution across cells of baseline-to-washout PD changes in learning simulation. To facilitate the comparison with the experimental results, we show in (B), (C) and (E) samples of cells of the same size as in the corresponding subplots in Figure 3 in the main text.

Extension of Model to Several Forgetting Time Constants

In our model, noise driven changes in tuning curves are shaped by sensory driven learning, so they will not harm performance. In our simulations, the correlation time of the noise driven changes τ_{forget} was set to 1500 trials, in order to reproduce the observed correlation time of the tuning curves. In the absence of sensory feedback, the noise erases the motor memory stored in the network, within several time constants τ_{forget}, i.e. within several thousands of trials. This consequence of the model seems implausible, since some motor skills are retained, at least partially, over many years without practice.

Here we show that by extending the model, such that it includes several forgetting time constants, it is possible to reproduce the observed correlation time of changes in tuning curves and yet maintain, at least partially, memory of motor skills for very long times. We assume a model with the same architecture as the original model, except that each sensory neuron makes a pair of synapses onto each motor cortical cell, including one unstable synapse and one stable synapse. Thus, instead of Equation 6 in the main text, the firing rates of the cortical cells are computed by

$$r_i = \sum_{j=1}^{2} \left(W_{ij}^{U} + W_{ij}^{S} \right) x_j^t \qquad \text{(S6)}$$

where W_{ij}^{U}, W_{ij}^{S} are the weights of the unstable and stable synapses, respectively. W_{ij}^{U}, W_{ij}^{S} are modified by the same rule as in the original model (see Equations 10 and 12 in main text), except that for W_{ij}^{U} we use τ_{forget}=750 trials, and for W_{ij}^{S} we use $\tau_{forget}=10^6$ trials. Additionally, we set $\sigma = 0.018$ for both types of synapses, and the number of cells N=1000. A similar model with two forgetting and learning time constants has been proposed recently to explain interesting behavioral findings in learning experiments (Smith et al., 2006). In contrast, we are interested in how the multiple time constant affect the changes in the neural representation.

Figure S7 compares the original model with the extended model. In both models, simulation of the control experiment produces a similar amount of change in tuning curves (Figures S7A,B). On the one hand, in the extended model only half of the synapses are unstable, whereas in the original model all synapses are unstable. On the other hand, the forgetting time constant of the unstable synapses is shorter in the extended model compared with the original model. These two effects cancel out such that tuning curves change similarly in both models. However, the performance of the models is very different when simulating without sensory feedback over many trials. The original model completely forgets what it has learned (Figure S7C), whereas the extended model partially retains its performance (Figure S7D). The extended model maintains its memory in the stable synapses. We conclude that by having several time constants, it is possible to account for the observed degree of change in tuning curves and have motor skills remembered over very long times.

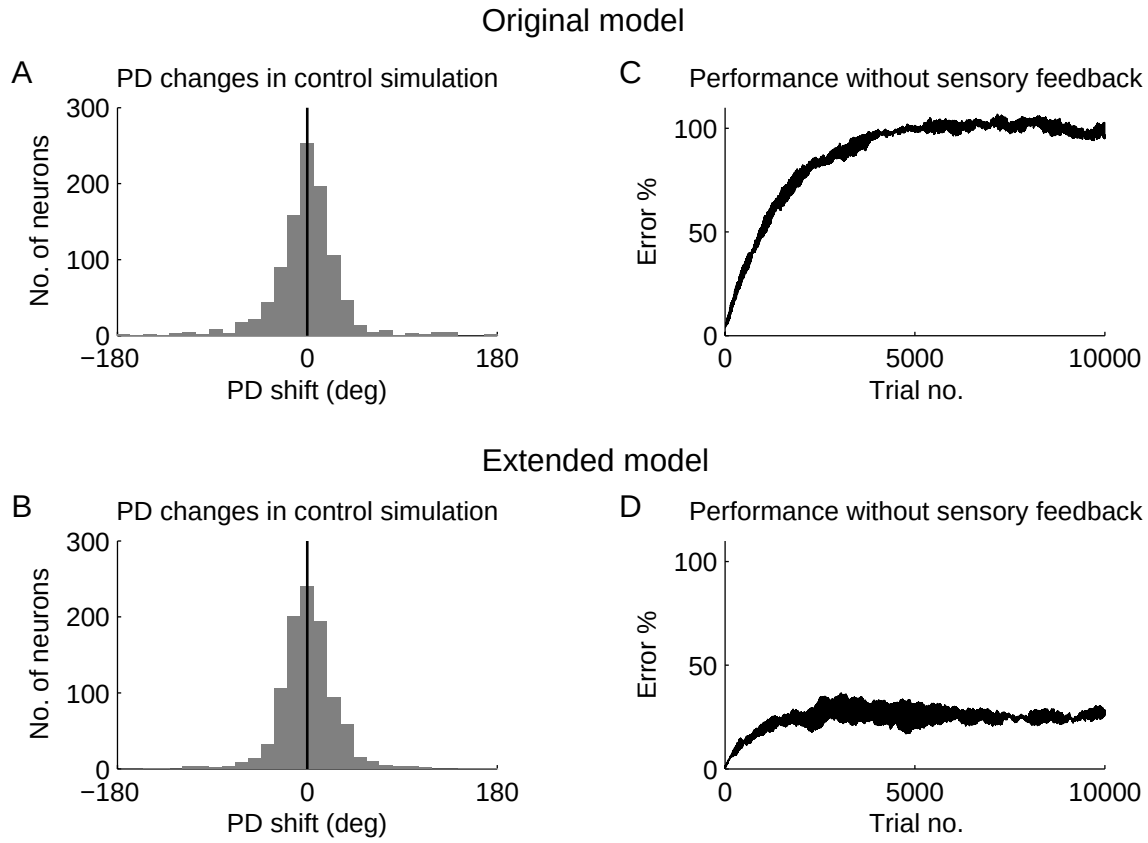

Figure S7. Extension of Model to Several Forgetting Time Constants
(A) Distribution of block 1 to 3 PD changes in control simulation of original model.
(B) Same as A for extended model.
(C) Error in hand position, in % of desired movement amplitude, as function of trial number in simulation of original model without sensory feedback.
(D) Same as C for extended model.

Dependence of Model Performance on Number of Neurons

In our simulations, the model's output slightly deviated from the desired output. Furthermore, the output for the same target changed over trials, with a typical time constant of the learning time constant τ_{learn}. In terms of the synaptic weights, these errors imply that the system fluctuates around the optimal manifold, i.e. the valley in the error landscape (Figure 4A in main text). In this section we examine the dependence of these errors on the number of neurons N. To asses the overall performance of the model we computed the mean squared error (MSE) in hand position

$$MSE = \left\langle \left\| x - x^t \right\|^2 \right\rangle \tag{S7}$$

where x is the hand position vector, x^t is the target position vector, and $\langle \ldots \rangle$ denotes averaging over many trials (we used 100000 trials). We estimated the MSE for different values of N, and found that it asymptotes at roughly N~100 (Figure S8A, dots).

To understand this behavior, it is useful to separate the squared error into the squared bias and variance contributions

$$MSE = bias^2 + variance \tag{S8}$$

The bias is the systematic deviation of the hand from the target position

$$bias^2 = \left\langle \left\| \langle x \rangle_{x|\theta} - x^t(\theta) \right\|^2 \right\rangle_\theta \tag{S9}$$

Here $\langle \ldots \rangle_{x|\theta}$ denotes averaging over trials for which the target was in direction θ, $x^t(\theta)$ is the position of the target in direction θ, and $\langle \ldots \rangle_\theta$ denotes averaging over movement directions. The bias is caused by the forgetting term in our synaptic weight update rule (Equation 10 in main text), which biases the synaptic weights slightly from the optimal manifold towards zero. This bias is independent of N (Figure S8A, circles; small changes are caused by sampling noise). The size of the bias depends on the ratio of time constants $\tau_{learn}/\tau_{forget}$ which determines the relative sizes of the forgetting and learning terms. When $\tau_{learn}/\tau_{forget}$ increases the forgetting term is more effective in biasing the synaptic weights from the optimal manifold, and the bias increase. At the regime where τ_{forget} is much longer than τ_{learn} this relation is approximately linear (Figure S8B)

$$bias \sim \frac{\tau_{\text{learn}}}{\tau_{\text{forget}}} \tag{S10}$$

The second contribution to the MSE is the variance, i.e. the randomness in hand position

$$variance = \left\langle \left\langle \left\| x - \langle x \rangle_{x|\theta} \right\|^2 \right\rangle_{x|\theta} \right\rangle_\theta \tag{S11}$$

The variability in hand position is caused by the synaptic plasticity noise which pushes the synaptic weights around the optimal manifold. The size of the variance depends on a number of factors. First, the variance decreases with N (Figure S8A, crosses), proportionally to $1/N$ (not shown), because there are on the order of N independent noise sources which average out at the output. Second, the variance scales with the variance of the noise σ^2 (Figure S8C). Lastly, the variance depends on τ_{learn}. When τ_{learn} increases, the learning signal becomes weaker, and thus the force driving the synaptic weights back to the optimal manifold weakens (Equation 10 in main text). Consequently, the variance in synaptic weights and the motor output increases, in proportion to τ_{learn} (Figure S8D). Putting these different factors together, we have approximately

$$variance \sim \frac{\tau_{\text{learn}} \sigma^2}{N} \tag{S12}$$

Now we are in a position to explain why the MSE asymptotes at roughly N~100. The MSE asymptotes when the variance term becomes small relative to the bias term, which according to Equations S10 and S12 implies

$$\left(\frac{\tau_{\text{learn}}}{\tau_{\text{forget}}}\right)^2 >> \frac{\tau_{\text{learn}}\sigma^2}{N} \tag{S13}$$

Rearranging terms we find

$$N >> \frac{\left(\tau_{\text{forget}}\sigma\right)^2}{\tau_{\text{learn}}} \sim 30 \tag{S14}$$

For the parameter values that we used in our simulations (τ_{learn}=50, σ=0.025, τ_{forget}=1500), the right hand side of Equation S14 is approximately 30. This explains why the error does not decrease much beyond N~100.

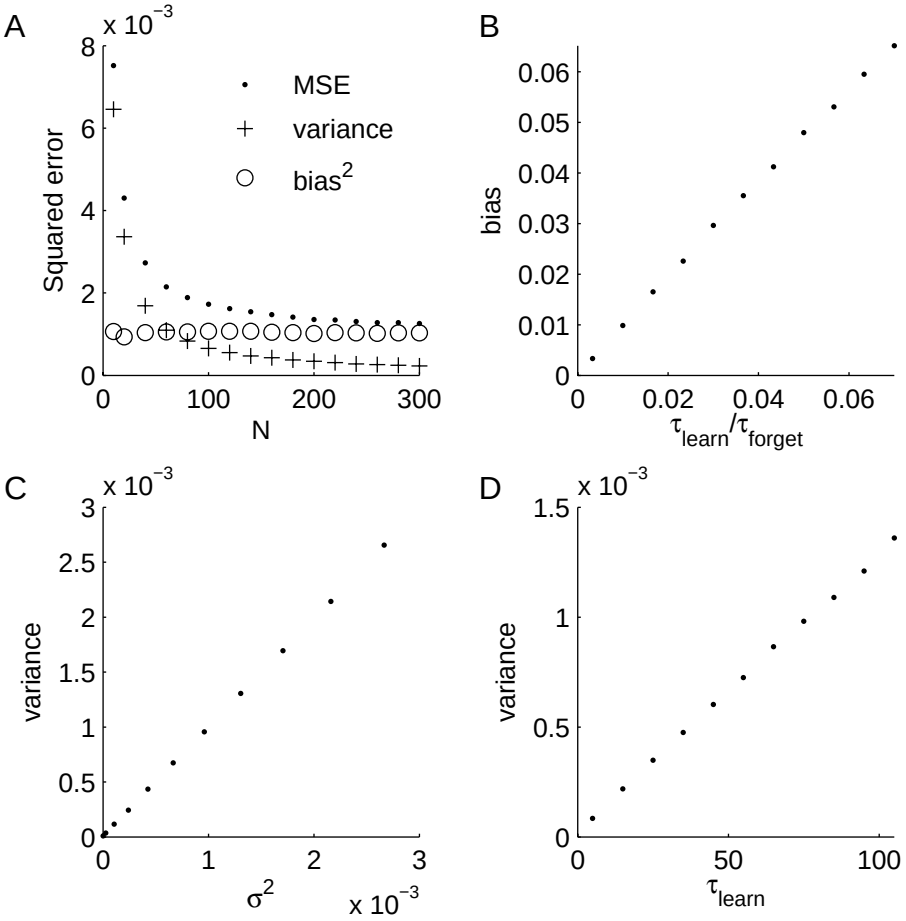

Figure S8. Dependence of Model's Error in Position on Parameters
Baseline parameter values are τ_{learn}=50, σ=0.025, τ_{forget}=1500, N=100. 1 unit of error corresponds to the desired movement amplitude.
(A) Simulations with different N. Shown are MSE (dots), squared bias (circles) and a variance (crosses).
(B) Simulations with different ratios of τ_{forget} and τ_{learn} (we modified τ_{learn}).
(C) Simulations with different σ.
(D) Simulations with different τ_{learn}.

Supplemental References

Moran, D. W., and Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. J Neurophysiol *82*, 2676-2692.

Smith, M. A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol *4*, e179.

Zhang, W., and Linden, D. J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. Nat Rev Neurosci *4*, 885-900.