

Motor Learning with Unstable Neural Representations

Uri Rokni,^{1,2,*} Andrew G. Richardson,⁴ Emilio Bizzi,^{1,3} and H. Sebastian Seung^{1,2}

¹Department of Brain & Cognitive Sciences

²Howard Hughes Medical Institute

³McGovern Institute for Brain Research

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁴Division of Health Sciences & Technology, Massachusetts Institute of Technology and Harvard Medical School, Cambridge, MA 02142, USA

*Correspondence: rokniu@mit.edu

DOI 10.1016/j.neuron.2007.04.030

SUMMARY

It is often assumed that learning takes place by changing an otherwise stable neural representation. To test this assumption, we studied changes in the directional tuning of primate motor cortical neurons during reaching movements performed in familiar and novel environments. During the familiar task, tuning curves exhibited slow random drift. During learning of the novel task, random drift was accompanied by systematic shifts of tuning curves. Our analysis suggests that motor learning is based on a surprisingly unstable neural representation. To explain these results, we propose that motor cortex is a redundant neural network, i.e., any single behavior can be realized by multiple configurations of synaptic strengths. We further hypothesize that synaptic modifications underlying learning contain a random component, which causes wandering among synaptic configurations with equivalent behaviors but different neural representations. We use a simple model to explore the implications of these assumptions.

INTRODUCTION

Neural recordings in behaving animals have revealed much about the mechanisms underlying motor learning. Changes in single-unit activity have been correlated with learning sensorimotor associations (Mitz et al., 1991; Ojakangas and Ebner, 1992; Paz et al., 2003; Paz and Vaadia, 2004; Wise et al., 1998), learning movement sequences and skills (Cohen and Nicolelis, 2004; Nakamura et al., 1998), and adapting to novel mechanical environments (Gandolfo et al., 2000; Li et al., 2001; Padoa-Schioppa et al., 2002, 2004; Xiao et al., 2006). One assumption implicit in many of these studies is that there is an under-

lying stable neural representation for familiar behavior, and thus changes in the neural representation necessarily reflect motor learning. Empirical support for this assumption is limited to just a few studies (Greenberg and Wilson, 2004; Nicolelis et al., 1997; Schmidt et al., 1976; Taylor et al., 2002; Thompson and Best, 1990; Williams et al., 1999). Interestingly, there are several indications that neural representations may, under some circumstances, change even without obvious learning. For example, when exposed to a fixed environment, hippocampal place fields in mice changed over the course of several hours when attentional demands were low (Kentros et al., 2004). Another study, which motivates the present work, showed that when monkeys performed a familiar reaching task the directional tuning of neurons in the supplementary motor area (SMA) changed substantially (Padoa-Schioppa et al. 2004). We refer to such changes in neural representations, which occur without obvious learning, as *background changes*. The cause of background changes and their function are unknown. Background changes may be related to adaptation to slow changes in the environment, e.g., muscle fatigue. Alternatively, background changes may be unrelated to behavior, and the neural representation of familiar tasks may be truly unstable. The main objective of this work is to study what such instability implies for the plasticity mechanisms underlying motor learning.

In the first half of the paper, we characterize background changes by reanalyzing data from the above mentioned recordings in SMA as well as new data from similar experiments in the primary motor cortex (MI). We study how the directional tuning changes in a “control” experiment, in which the monkey practices a familiar reaching task, and in a “learning” experiment, in which the monkey reaches in the presence of novel forces. In the second half of the paper, we explore the theoretical implications of the possible instability of the motor cortical representation. It has been suggested previously that changes in tuning curves may cancel out at the level of the motor output (Li et al., 2001). Here, we relate this idea to a phenomenon in the theory of neural networks, which we term *redundant networks*. A network is redundant if it uses more neurons

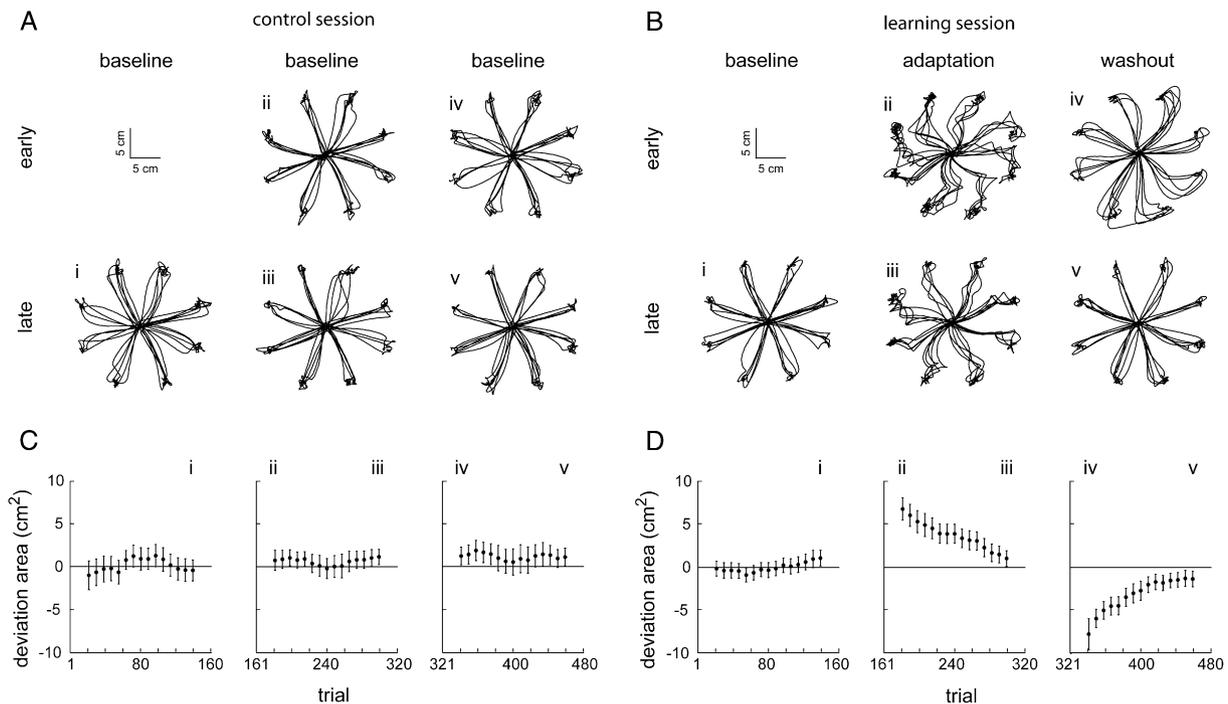


Figure 1. Examples of Behavior in Control and Learning Sessions of the Monkey from Which the M1 Recordings Were Obtained

(i) Trials 121–160, (ii) trials 161–200, (iii) trials 281–320, (iv) trials 321–360, (v) trials 441–480.

(A) Hand trajectories from a control session.

(B) Hand trajectories from a learning session.

(C) Performance during the control session in (A), quantified by the area between the trajectory and a straight line path (40 trial moving average shown with 95% Student's *t* confidence intervals).

(D) Performance during the learning session shown in (B).

than needed to solve its task, such that the neural representation may change without affecting the overall behavior. Using a simple model, we show that noisy learning in a redundant motor cortex produces a background of behaviorally irrelevant changes in tuning curves. Additionally, we examine what further assumptions about the nature of synaptic plasticity are required to explain the observed properties of the background changes.

RESULTS

The Control Experiment: Background Changes Are Random and Slow

In order to characterize background changes, we have analyzed data of a control experiment in which monkeys performed a familiar reaching task, on which they had been trained for several months. On each day of recording, the monkey had to reach to one of eight targets arranged on a circle, 480 times. The hand trajectories showed relatively small changes between different epochs within a practice session (Figure 1A; see also [Supplemental Data](#) available with this article online). We analyzed the movement-related responses of 136 cells, 43 from SMA of one monkey (from [Padoa-Schioppa et al. \[2004\]](#)) and 93 from M1 of a second monkey (new data). Because we found similar results in both brain areas, we

pooled all 136 cells in subsequent analyses. We characterized each cell's movement-related activity by the mean firing rate in a time window from 100 ms prior to movement onset to 300 ms after movement onset. Tuning curves were defined by mean firing rate as a function of the eight reach directions of the task. To examine changes in tuning curves we artificially divided the data from the 480 trials into three consecutive blocks of 160 trials and computed tuning curves for each block separately.

Tuning Curves Change

The left column in Figure 2 shows four example cells whose tuning curves changed between block 1 (crosses) and block 3 (circles). This tuning instability was not due to recording instability, since the spike waveforms did not change from block 1 to block 3 (Figure 2, right columns), and similar changes in tuning curves were observed in a subpopulation of cells judged as best isolated and most stable (Table 1; see also [Experimental Procedures](#)). In 23% of the 8 directions × 136 neurons, there was a statistically significant change from block 1 to block 3 in the mean firing rate (*t* test, $p < 0.01$). Even more significant changes were seen using aggregate measures of the tuning curves. Seventy-seven percent of the variance of the changes in tuning curves was accounted by changes in their offsets, and 16% was accounted by changes in the cosine components (total of 93%; see [Supplemental](#)

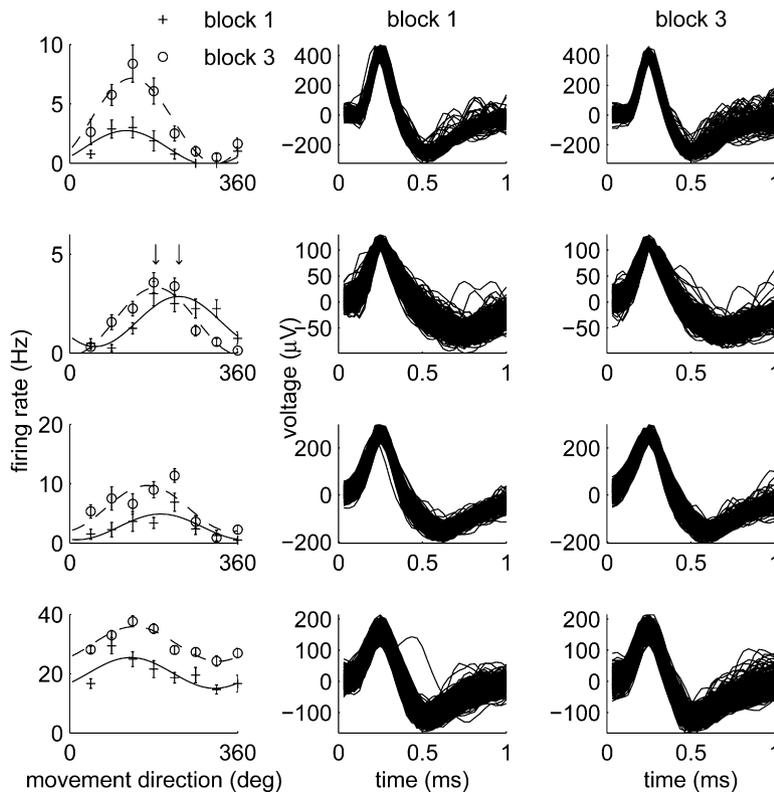


Figure 2. Changes in Tuning Curves in Control Experiment

Each row corresponds to a sample cell. (Left) Mean firing rates in block 1 (crosses) and block 3 (circles) as a function of movement direction, and fitted cosine tuning curves (lines). Error bars correspond to standard errors, and arrows in second row designate the PDs. (Middle and right) Random sample of 1000 spike waveforms in block 1 (middle) and block 3 (right).

Data). Therefore, to quantify the changes in tuning curves, we first fitted the tuning curve of each cell in each block by an offset plus a cosine function (lines in Figure 2, left)

$$r(\theta) = B + A \cos(\theta - \psi), \quad (1)$$

where $r(\theta)$ is firing rate as a function of target direction, B is the offset, A is the modulation depth, and ψ is the preferred direction (PD). Next, we compared the fitted parameters between different blocks. We found changes in offsets (e.g., Figure 2, row 4), modulation depths (e.g., Figure 2, row 1), and PDs (e.g., Figure 2, row 2). In 73% of the neurons, offset changes between blocks 1 and 3 were statistically significant (z test, $p < 0.01$), and in 63% of the neurons, changes between blocks 1 and 3 in the cosine function (i.e., changes in PDs and/or modulation depths) were statistically significant (bivariate z test, $p < 0.01$; see Experimental Procedures). Thus, as observed by Padoa-Schioppa et al. (2004), motor cortical tuning curves may change even when the monkey is performing a familiar task.

Background Changes Are Random across Neurons and Time

We found that changes in offsets, modulation depths, and PDs had qualitatively similar statistical properties. The statistical properties for PD changes of 93 neurons (out of 136) whose tuning curves had a statistically significant cosine component in all blocks (bivariate z test, $p < 0.05$) are shown in Figure 3. Figure 3A presents the distribution of PD changes from block 1 to block 3. The average PD

change was not statistically different from zero (z test, $p > 0.05$). Figure 3A averages across different days, and hence it is possible that on a given day different neurons tend to shift their PDs in the same direction. To test this possibility, for all pairs of cells recorded simultaneously, we plotted the PD change of one neuron against the PD change of the other neuron (Figure 3B). We found no statistically significant correlation among these pairs of PD changes (permutation test, $p > 0.05$). In order to test whether PD changes across different times were correlated, for each cell we plotted the PD change from block 1 to 2 versus its PD change from block 2 to 3 (Figure 3C). Here as well, we found no statistically significant correlation (permutation test, $p > 0.05$). These results show that tuning curve changes in the control experiment were random across neurons and time.

Background Changes Are Slow

We characterized the correlation time of the randomly changing PDs. For this purpose, we binned the data into 12 blocks of 40 trials each. For every pair of bins, we computed the correlation between the populations of PDs at the two time bins and averaged all pairs of bins separated by the same lag. Within the range of measured lags the correlation decayed linearly (Figure 3D; the y intercept is not 1 because of standard errors of the PDs). The slope of the autocorrelation was roughly 1/3000 trials, indicating a slow correlation time of the PDs, on the order of thousands of trials (Figure 3D shows correlations of PDs, which does not contradict the lack of correlation of PD changes

Table 1. Statistics of Changes in Tuning Curves

Statistics	Control Experiment blk 1 to 3	Control Experiment blk 1 to 3 (Best Isolated Cells)	Learning Experiment, Baseline to Washout
Sig. changes in cosine (%)	63 ± 4	61 ± 7	63 ± 3
Mean Δ PD (deg)	-2 ± 3	3 ± 3	-1 ± 2
St. dev. of Δ PD (deg)	29 ± 3	19 ± 3	35 ± 2
Mean Δ mod. depth (Hz)	0.3 ± 0.3	1.1 ± 0.5	0.6 ± 0.2
St. dev. of Δ mod. depth (Hz)	3.3 ± 0.3	3.8 ± 0.5	3.7 ± 0.2
Sig. changes in offset (%)	73 ± 4	69 ± 7	76 ± 2
Mean Δ offset (Hz)	1.5 ± 0.7	3.3 ± 1.4	1.8 ± 0.4
St. dev. of Δ offset (Hz)	7.3 ± 0.7	10.0 ± 1.4	7.0 ± 0.4

The changes are not due to recording instability because similar changes are seen in the best isolated cells (compare columns 1 and 2). The changes in the learning experiment from baseline to washout are not related to adaptation or deadaptation because they have the same statistics as the changes in the control experiment (compare columns 1 and 3). ± indicates standard error.

in Figure 3C). Presently, it is unclear how much the correlation decays over more trials. Additionally, it is unclear whether the time unit relevant for these changes is number of trials or real time.

A similar analysis showed that the offsets and modulation depths also changed slowly and randomly across cells and time (data not shown).

The Learning Experiment: Learning-Related Changes Occur on Top of Background Changes

In this section, we show that learning adds systematic changes on top of the background of random changes described above. The data for this analysis were recorded while the monkeys performed the same reaching task as above, except that novel forces generated by a robotic manipulandum were applied to the arm during the middle 160 trials of each session. Thus, the experiment consisted of three consecutive blocks of 160 trials: (1) a *baseline* block in the absence of forces, (2) an *adaptation* block in the presence of forces, and (3) a *washout* block in the absence of forces. The forces applied during the adaptation block were curl velocity force fields, i.e., proportional to the hand speed and orthogonal to its direction of movement.

Figure 1B shows hand trajectories from one learning session. At the baseline block, the trajectories were fairly straight (Figure 1Bi). At the beginning of the adaptation block, the monkey's hand trajectories were curved by the forces (Figure 1Bii). While practicing in the adaptation block, the monkey learned to compensate for the forces partially, resulting in somewhat straighter hand trajectories (Figure 1Biii). Upon removal of the forces at the beginning of the washout, the monkey's hand trajectories curved in the opposite direction, showing an aftereffect of the forces (Figure 1Biv). Finally, when practicing in the washout block, the monkey relearned the original task, and the trajectories became similar to baseline (Figure 1Bv). To characterize the curvature of the trajec-

ries, we defined the deviation area measure, which is the area between the hand path and a straight path connecting the initial and final hand positions. When integrating this area over the path, counterclockwise deviations are regarded positive and clockwise deviations are regarded negative, and thus the sign of the deviation area indicates the direction of curvature. Figure 1D shows that turning the forces on or off caused abrupt changes in deviation area, followed by gradual adaptation. These learning related changes were large relative to the small changes observed in the control sessions (Figure 1C). Similar behavior was reported previously (Gandolfo et al., 2000).

We analyzed the responses of 172 neurons (67 from M1 and 105 from SMA) recorded during this novel task that had tuning curves with statistically significant cosine components in all blocks (bivariate z test, $p < 0.05$). We constructed three separate tuning curves for each neuron, respectively from the activity of the baseline block, the late adaptation (last 80 trials), and the late washout (last 80 trials). We designate the changes from baseline to late adaptation as *adaptation changes* and the changes from late adaptation to late washout as *washout changes*.

In contrast with the control experiment, in which PD changes at different times were uncorrelated across cells (Figure 3C), in the learning experiment adaptation and washout changes were anticorrelated and distributed along the $y = -x$ diagonal (Figure 3E; see also Padoa-Schioppa et al. [2004]). This indicates that on average, adaptation changes were reversed by washout. However, there were also deviations from the diagonal, indicating that after washout PDs of individual cells did not return to their baseline values. Using similar methods as we used for the control experiment, we found that for many cells differences in tuning curves between baseline and washout were statistically significant (Table 1; see also Padoa-Schioppa et al. [2004]). It was previously proposed that the baseline-to-washout changes underlie

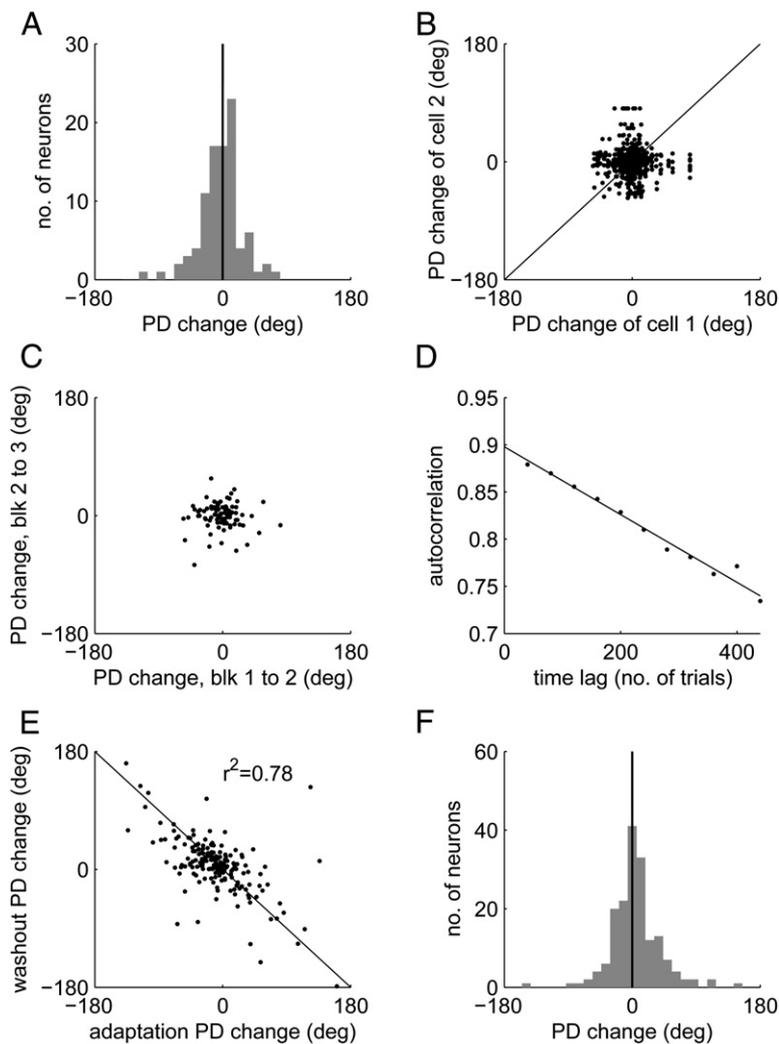


Figure 3. Statistics of Changes in PDs in Control and Learning Experiments

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

(B) PD change from block 1 to 3 of one cell versus PD change from block 1 to 3 of another cell recorded simultaneously, across all simultaneously recorded pairs. Each pair is represented by two points symmetrically positioned around the $y = x$ diagonal (solid line).

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

(D) Autocorrelation of population of PDs and linear fit (solid line).

(E) Adaptation PD changes versus washout PD changes, across cells in learning experiment. Solid line represents the $y = -x$ diagonal.

(F) Distribution across cells of baseline-to-washout PD changes in learning experiment.

learning of the force task. Additionally, it is possible that these changes are related to the monkeys not fully dead-adapting in the washout. To challenge these interpretations, we compared the statistics of the baseline-to-washout changes in the learning experiment with changes in the control experiment over a similar number of trials. We found that the distribution of baseline-to-washout PD changes (Figure 3F) was similar to the distribution of PD changes in the control experiment from block 1 to block 3 (Figure 3A). Furthermore, every statistic we have examined—of the changes in PDs, modulation depths, and offsets of the tuning curves—showed no statistically significant difference between the learning and control experiments (z test, $p > 0.05$; Table 1). This result suggests that the changes from baseline to washout are unrelated to either the force adaptation or deadadaptation processes.

As an alternative interpretation, we suggest that changes in tuning curves in the learning experiment are a sum of two components: systematic learning-related changes and random background changes which exist

regardless of learning. The learning-related changes reverse at washout and are therefore responsible for the anticorrelation observed between adaptation and washout changes. The background changes are responsible for the changes from baseline to washout. The fact that the statistics of these background changes were so similar in the control and learning experiments implies that the learning-related and background changes do not interact.

Theory: Background Changes Are Caused by Noisy Learning in a Redundant Motor Cortex

What is the interpretation of the background changes? Perhaps background changes reflect subtle behavioral changes (although we did not find evidence for this possibility—see Supplemental Data). Alternatively, the background changes may be behaviorally irrelevant. We have constructed a theory which suggests why behaviorally irrelevant changes in the neural representation might occur. The theory is based on three assumptions: (1) motor

cortex is redundant in the sense that it uses more neurons than required to produce the desired sensorimotor transformation, (2) when practicing a task, sensory feedback about motor errors is translated to synaptic changes which reduce the errors, and (3) this plasticity mechanism is noisy. We found that under these conditions a background of behaviorally irrelevant changes in tuning curves is produced.

Our assumption that motor cortex is redundant allows it to achieve the same sensorimotor transformation with different neural representations. In terms of synaptic weights, this implies a continuum of configurations which produce the desired sensorimotor transformation, which we term the *optimal manifold*. The synaptic configurations within this manifold are minima of the motor error. Therefore, one way to imagine this optimal manifold is by a flat valley in the landscape of the motor error as a function of synaptic weights (shown schematically in Figure 4A). Synaptic learning can be described as going down the error landscape. If learning is noisy and ongoing, then even after reaching the valley and mastering the task, synaptic strengths continue to wander along the valley. Thus, a background of behaviorally irrelevant changes in the neural representation is produced.

A Model of the Background Changes in Motor Cortical Tuning Curves

To demonstrate our theory, we constructed a simple model of a redundant cortical network which generates reaching movements in the horizontal plane. Following the approach of Salinas and Abbott (1995), our model generates reaching by the following stages (Figure 4B): (1) the appearance of the target activates two sensory units, in proportion to the x-y coordinates of the target, (2) the sensory units activate a large number of motor cortical neurons, (3) the motor cortical neurons generate a two-dimensional endpoint force on the hand, and (4) the force moves the hand to a new position in the plane. Each of these stages is modeled as a linear static mapping. In this static framework, we cannot represent dynamic force perturbations, so we used a perturbation of a static rotation, which similar to the curl velocity force field requires rotation of endpoint force.

The tuning curves of our model cells are defined by the firing rates as a function of target direction. Notice that if the network is wired properly, such that movement direction equals target direction, then these tuning curves also describe tuning to movement direction. In our model, tuning curves are cosine shaped (Figure 4B, inset), resembling the broad unimodal tuning curves observed in motor cortex. The cosine-shaped directional tuning stems from our assumption of a linear relation between firing rates and Cartesian position coordinates (Mussa-Ivaldi, 1988; Todorov, 2000). The modulation depths and PDs of the tuning curves are determined by the cells' input connections. When these connections are modified by synaptic plasticity, the tuning curves change. In this work, we did not model the offsets of the tuning curves.

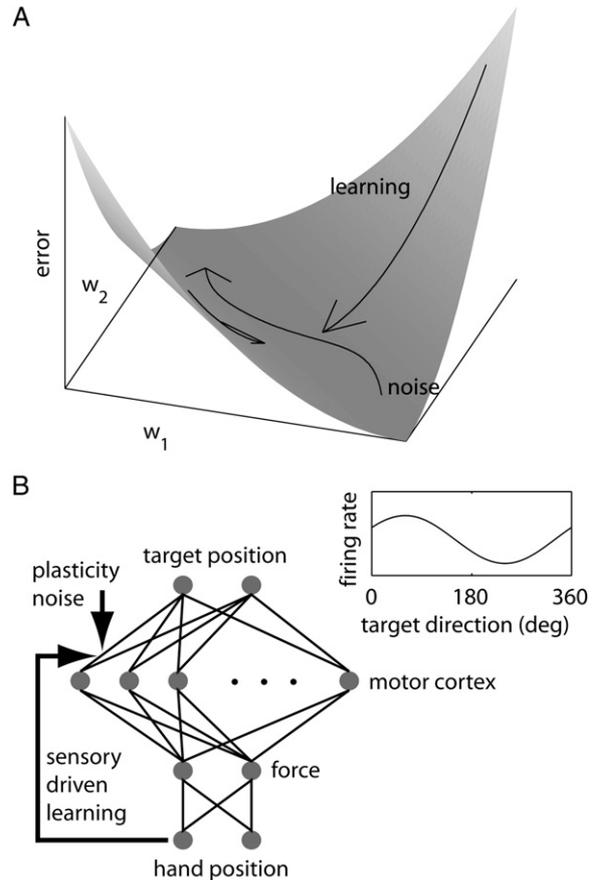


Figure 4. Theory for Cause of Background Changes

(A) Learning pushes the synaptic strengths down an error landscape which has a valley of minima at the optimal manifold, and noise causes the synaptic strengths to drift along this valley.

(B) Model of motor cortical network which generates reaching movements. Tuning curves of model cells are cosine shaped (inset).

The goal of the network is to minimize the error between hand position and target position. For simplicity, we assumed that in order to achieve this task, only the input weights of the motor cortical cells can be modified, whereas the cells' output weights are fixed. In this sense, the sensorimotor transformation is stored in the input weights of motor cortex. We assumed that after each trial, i.e., a single run on the network, sensory feedback about the motor error is used to modify the input weights in order to reduce subsequent motor error. Our major assumption regarding this plasticity process is that noise is added to the learning signal, independently at different synapses, and that this plasticity is operative even when the network has already mastered its task. In addition to the noise and learning signal we also added a decay term which limits the degree of wandering of synaptic weights.

Simulation of the Control Experiment

In order to show how background changes are generated and explain why they are random and slow, we simulated

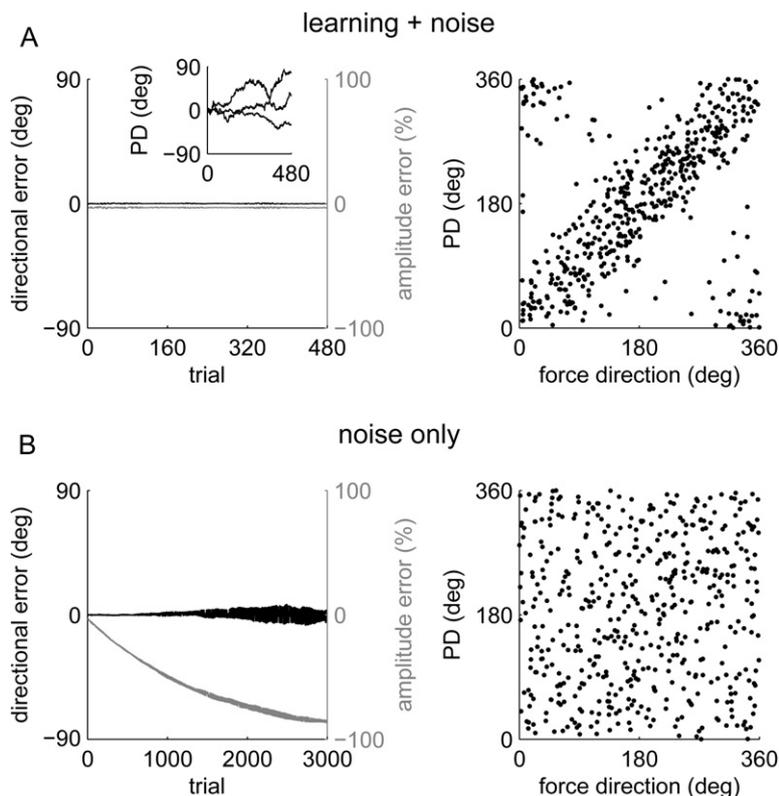


Figure 5. Behavior and Neural Representation in Simulations of Control Experiment

(A) Simulation with noisy learning rule, $\sigma = 0.025$, $\tau_{\text{forget}} = 1500$, $\tau_{\text{learn}} = 50$. (Left) Error in movement direction (black), error in movement amplitude as percentage of desired movement amplitude (gray), and PDs of three sample cells whose PDs started close to zero (inset). (Right) PD (computed from the firing rate r_i) versus force direction (denoted α_i in the [Experimental Procedures](#)) from last trial of simulation shown on left.

(B) Same as (A) but without a learning signal $\sigma = 0.025$, $\tau_{\text{forget}} = 1500$, $\tau_{\text{learn}} = 10^6$. In both simulations $N = 10000$, but only 500 randomly sampled cells are shown on right panels.

the control experiment. First, we pretrained the model for many trials to mimic the excessive pretraining the monkeys had experienced. Next, we simulated 480 trials of the task, where at each trial a target was chosen randomly from eight targets arranged uniformly on a circle.

Model Generates Background Changes

In the simulation of the control experiment, the model maintained good performance ([Figure 5A](#), left; there was a small bias because of the weight-decay term), and yet the PDs of the model cells changed considerably ([Figure 5A](#), inset). Thus, the neural representation wandered in a manifold of configurations which produce the same behavior. To understand this redundancy, we first consider one simple configuration of tuning curves within this manifold, in which the PD of each cell equals the direction of force it generates (and all cells have the same modulation depth). This configuration generates a motor output in the correct direction because cells which produce force directions close to the desired direction are preferentially recruited, and the force components orthogonal to the desired direction cancel out. We refer to these tuning curves as the *relevant tuning curves*. Because of the vast convergence from cells to motor outputs, it is possible to add *irrelevant components* to these tuning curves, whose effects on the motor output cancel out. Thus, a generic configuration of tuning curves in the manifold can be decomposed into relevant components which produce the desired output and irrelevant components which do not contribute to the outputs. In such configura-

tions, PDs are correlated with, rather than equal to, the force directions ([Figure 5A](#), right). During noisy plasticity, as the neural representation wanders in the manifold, the relevant components remains fixed and the irrelevant components change randomly. The typical size of the irrelevant components is determined by the amplitude of the plasticity noise. The stronger the plasticity noise, the larger the irrelevant components, and therefore the weaker is the correlation between PDs and force directions. In our simulations, the noise amplitude was tuned to reproduce the magnitude of the observed PD changes.

To emphasize the active role of the learning signal in maintaining the performance, we also performed a simulation with the learning signal turned off. In this case, the noise randomized the PDs ([Figure 5B](#), right). Consequently, cells generated forces more or less equally in all directions and the net output diminished ([Figure 5B](#), left; this does not necessarily imply that prolonged sensory deprivation causes immobilization because other sources of drive may take over). These random changes had relatively little effect on movement direction because they tended to averaged out. The time constant of this forgetting process was set by the time constant of the decay term in the weight update rule, denoted τ_{forget} .

We compared the properties of the model-generated background changes (with the learning signal on) and the experimentally observed background changes. For this purpose, we replicated our analysis of the experimental data on the simulation data. We divided the simulation

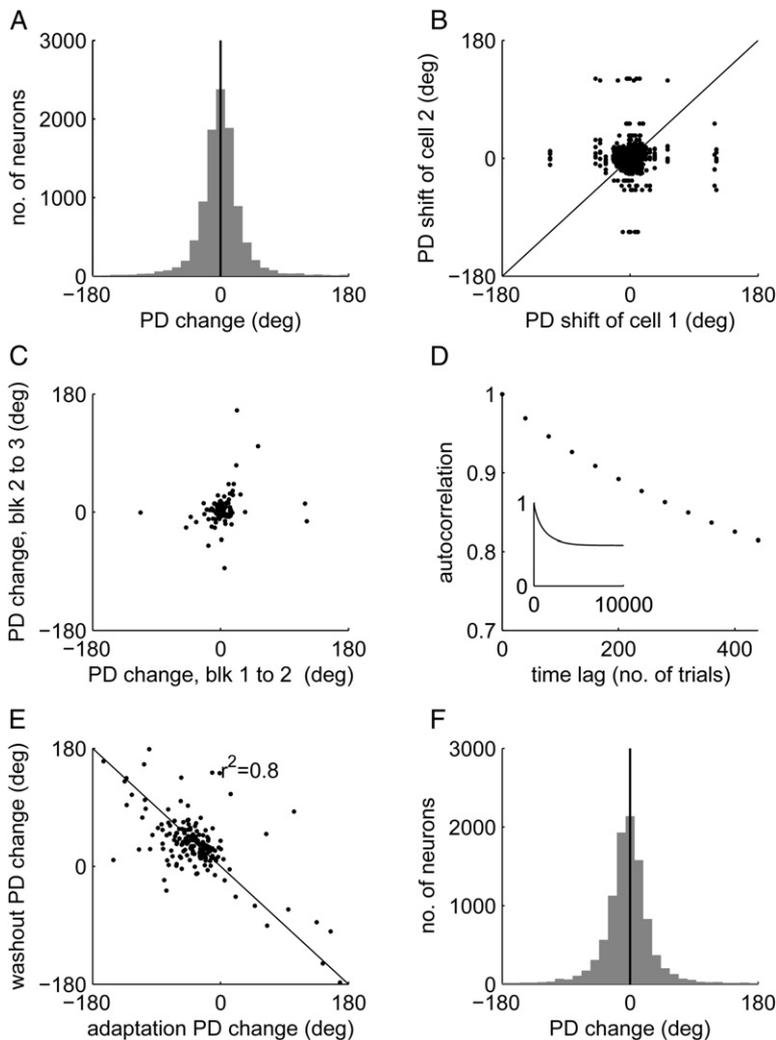


Figure 6. Statistics of Changes in Tuning Curves in Simulations of Control and Learning Experiments

(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 versus 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 versus 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. Model parameter values are $\tau_{learn} = 50$, $\tau_{forget} = 1500$, $\sigma = 0.025$, $\varphi = 60^\circ$, $N = 10000$.

data into three equal blocks and used the neural activities within each block to construct directional tuning curves.

Local Noise and High Redundancy Explain Randomness of Background Changes

Figures 6A and 6B show that PD changes in the model are random across cells, similar to the randomness observed in the experimental data (compare with Figures 3A and 3B). The randomness across cells in the model results from our assumptions of local synaptic noise sources and a high degree of redundancy. An alternative to local noise sources is noise which comes from the environment, e.g., muscle noise, which through sensory feedback contaminates the learning signal. We found that such environmental noise creates changes in tuning curves which are correlated across cells (Supplemental Data). Additionally, if redundancy is not high, changes in different cells may be coupled. When both local noise and high redundancy are assumed, PDs of different cells change nearly independently. Figure 6C shows that PD changes in the model are also random across time, similar to the randomness observed in the experimental data (compare

with Figure 3C). This temporal randomness results from the temporal randomness of the plasticity noise.

Background Changes Must Be Slow to Allow Learning

The control experiment showed that the background changes are slow, in the sense that PDs have a correlation time on the order of thousands of trials. In our model, the correlation time of the background changes was determined by τ_{forget} . We set $\tau_{forget} = 1500$ trials, so that the autocorrelation of the PDs decayed slowly (Figure 6D), similarly to the experimental autocorrelation (Figure 3D). We found that in order to obtain good performance of the model, τ_{forget} must be much greater than the learning time constant which was set by another parameter, τ_{learn} . There is continual competition between the learning signal which stores motor memories and plasticity noise which erases the memory. When τ_{forget} is much larger than τ_{learn} , the erasure causes only a slight bias of the model's output (Figure 5A, left). However, when τ_{forget} is comparable to τ_{learn} , this bias becomes large (data not shown).

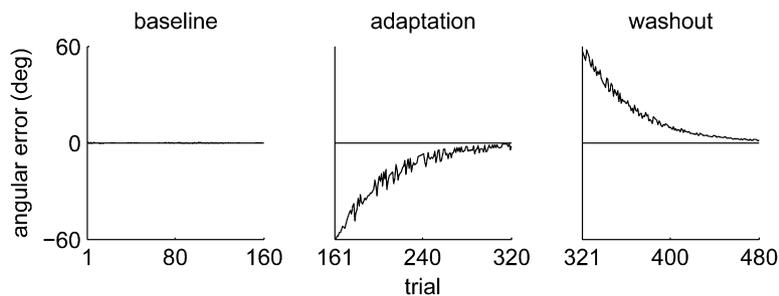


Figure 7. Model's Error in Direction of Movement in Simulation of Learning Experiment

$T_{learn} = 50$, $T_{forget} = 1500$, $\sigma = 0.025$, $\phi = 60^\circ$, $N = 10000$.

At Long Times, PDs Are Not Completely Randomized

Even after 10,000 trials of the control simulation, the auto-correlation of the PDs did not vanish, but rather decayed to a positive baseline (Figure 6D, inset). This baseline correlation reflects the fixed relevant components of the tuning curves. In other words, the tuning curves do not change completely arbitrarily, but are rather confined to configurations which produce the correct behavior. The value of this baseline correlation depends on the relative magnitude of the relevant and irrelevant components, which in turn depends on the amplitude of plasticity noise.

Simulation of the Learning Experiment

In order to explain how learning related changes in tuning curves combine with background changes, we simulated the learning experiment. We modeled the effect of the forces as a rotation of the outputs by 60° . We first pretrained the model for many trials and then trained the model on (1) a baseline block of 160 trials without the perturbation, (2) an adaptation block of 160 trials with the rotation perturbation, and (3) a washout block of 160 trials without the perturbation. When the perturbation was turned on or off, the model produced a large error which was subsequently reduced by learning (Figure 7).

We repeated the analysis we had performed on the data of the learning experiment on our simulation data. We computed average tuning curves for the baseline, late adaptation (last 80 trials), and late washout (last 80 trials). As in the experiment, we designate changes from baseline to late adaptation as *adaptation changes* and changes from late adaptation to late washout as *washout changes*.

Model Generates a Combination of Learning-Related Changes and Background Changes

The learning experiment showed that adaptation and washout changes were anti-correlated, albeit with a considerable spread (Figure 3E). We interpreted this result as indicating that changes in tuning curves are a sum of learning-related changes and background changes. Similarly, in the learning simulation, the adaptation changes and washout changes were anticorrelated with considerable spread (Figure 6E). In the model, the learning-related and background changes are caused by the learning signal and plasticity noise, respectively. The learning signal causes behaviorally relevant changes in synaptic strengths in order to improve performance. At the same time, plasticity noise changes synapses randomly and

causes behaviorally irrelevant changes. Because of these irrelevant changes, after washout synapses are in a configuration which is different from their baseline configuration, and thus tuning curves change from baseline to washout (Figure 8A).

Similarity of Baseline-to-Washout Changes and Control Changes Is Explained by Additive Plasticity Noise

Our experiments showed that changes from baseline to washout in the learning experiment had similar statistics as the changes in the control experiment over a similar number of trials. This also holds for our model, e.g., the distribution of PD changes from baseline to washout (Figure 6F) is very similar to the distribution of PD changes in the control simulation from block 1 to 3 (Figure 6A; small differences between the two distributions are caused by the fact that learning is not entirely complete at the late washout). This similarity results from our assumption of additive plasticity noise. Because the noise is additive, the relevant changes caused by the learning signal and the irrelevant changes caused by the noise do not interact. Thus, learning a new task does not affect the statistics of the irrelevant background changes. If the noise were multiplicative, i.e., scaling with the motor error, learning a novel task would have increased the background changes. Linearity of neurons is not necessary to make the statistics of background changes independent of learning. Even with nonlinear neurons (and additive noise), as long as learning does not change the statistics of the gains between synaptic changes and firing rate changes, there is no effect on the statistics of the background changes.

DISCUSSION

In experiments on motor learning, it is often assumed that there is an underlying neural representation that is stable and that adaptation takes place on top of this stable background. Our experimental and theoretical results suggest a radically different picture. The experiments show that tuning curves of motor cortical cells are constantly changing even when performing a familiar task. Furthermore, when learning a new task, learning-related changes occur on top of this background of changing tuning curves. To explain these results, we proposed a theory which is based on the following assumptions: (1) motor cortex is

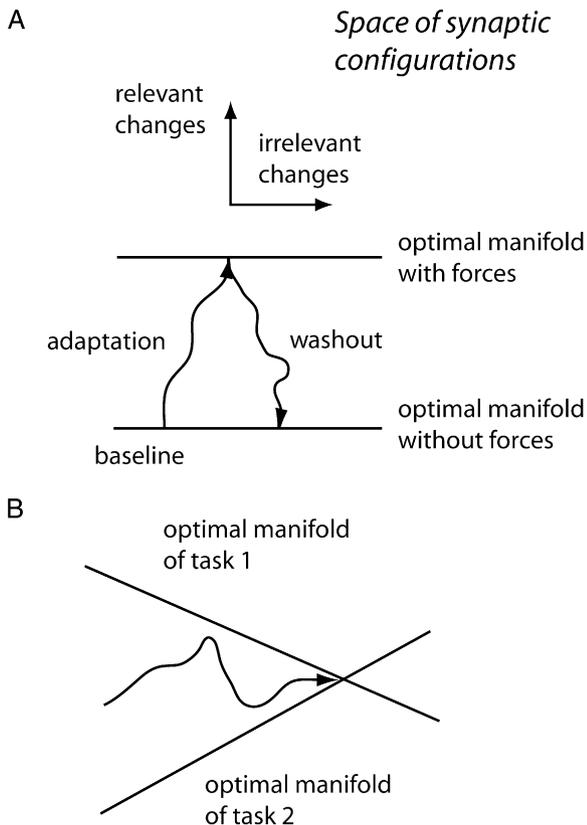


Figure 8. Optimal Manifolds of Multiple Tasks

(A) Changes in tuning curves in the learning experiment are a combination of behaviorally relevant changes created by the learning signal and irrelevant changes created by plasticity noise. After washout, synapses return to the manifold optimal for the no-force task at a configuration different from baseline.

(B) Learning several tasks with the same neural circuitry can be described as moving synaptic strengths to a configuration in the intersection of the manifolds optimal for these tasks.

redundant in the sense that it uses more neurons than required to generate the desired sensorimotor transformation, (2) when practicing a task, sensory feedback is transformed into synaptic changes which reduce the motor error, and (3) this plasticity mechanism is noisy. The redundancy of the system allows changes in the neural representations that do not affect behavior. The noise changes tuning curves randomly, and the learning signal shapes these changes so they do not harm task performance. As a result, tuning curves wander randomly between different configurations which are behaviorally equivalent.

Alternative Interpretations

While our theory provides an explanation for why tuning curves changed in the control experiment, there are a number of alternative interpretations which at this point cannot be ruled out. Changes in tuning curves may be related to behavioral changes which we have overlooked, e.g., postural changes that are not reflected in our hand

kinematics data. Additionally, it is possible that the recording electrodes injured the cells and consequently affected their tuning curves. Finally, perhaps neuromodulation underlies the changes in tuning curves, rather than synaptic changes.

How Does Our Interpretation of the Data Differ from Previous Interpretations?

In previous work, the significance of background changes was not fully appreciated, and consequently the data were interpreted differently. Specifically, in previous work cells were classified by how they changed their PDs in the learning experiment. Cells were classified as *kinematic* cells whose PDs changed very little, *dynamic* cells whose PDs changed during adaptation and changed back during washout, and *memory* cells whose PDs changed without returning to their baseline values (Li et al., 2001; Padoa-Schioppa et al., 2004). However, the data do not show clear clusters corresponding to these cell classes, but rather a continuum of response types (e.g., Figure 3E). According to our interpretation, such diversity of response types does not reflect specialized cell classes, but rather the randomness inherent in plasticity. If our interpretation is correct, then recordings across days will show that cells switch randomly between the different classes.

The previous studies proposed that changes from baseline to washout in memory cells reflect memory of the adaptation. In contrast, according to our theory, changes in tuning curves from baseline to washout are behaviorally irrelevant changes caused by plasticity noise. This interpretation is supported by our result that the statistics of the changes from baseline to washout are very similar to the statistics of the changes in the control experiment over a similar number of trials. According to our interpretation, recordings across days will show changes that are uncorrelated, whereas if changes are learning related, they are more likely to be consistent across days.

What Additional Evidence Is There for the Theory?

According to our theory, even when practicing a familiar task, sensory feedback is continually used to learn and prevent noise from erasing motor memories. Thus, our theory predicts that in the absence of sensory feedback familiar tasks are forgotten (Figure 5B). This prediction is confirmed by experiments that show that interfering with auditory feedback in adult finches or adult humans causes their well-learned vocalizations to slowly deteriorate (reviewed in Brainard and Doupe [2000]). Additionally, our theory predicts that as a task becomes more demanding the neural representations become more stable. This is predicted to occur because when more task constraints are added the dimension of the optimal manifold reduces, thus reducing the drift in synaptic strengths (for this effect to be appreciable redundancy should be low). This prediction is confirmed by an experiment which shows that as the requirements on spatial navigation of mice increases the spatial representation in their hippocampus becomes more stable (Kentros et al., 2004).

How Could the Theory Be Further Tested?

Besides the above mentioned recordings across days, one may use brain computer interface (BCI) experiments, in which a population of cortical cells is used to control a computer device (for review, see [Schwartz \[2004\]](#)). The advantage of BCI experiments is that the mapping from neural activity to motor output is fully known. Knowledge of this mapping can be used to test directly our assertion that changes in the neural representation are shaped so they would not affect the motor output.

What Have We Learned about Plasticity Underlying Motor Learning?

First, from the existence of background changes, we concluded that this plasticity process is considerably variable. Second, from the spatial randomness of background changes, we inferred that the source of variability is local, i.e., independent in different synapses, rather than noise from the environment, e.g., muscle noise, which through sensory feedback contaminates the learning signal. Third, from the fact that baseline-to-washout changes in the learning experiment have similar statistics to changes in the control experiment, we concluded that plasticity noise is additive. Finally, from the long correlation time of the background changes, we concluded that noise changes synapses very slowly. According to our theory, this slowness is necessary to prevent the noise from erasing motor memories. Notice that all these conclusions are based on the assumption that the observed changes in tuning curves are caused by synaptic changes.

The Meaning of Tuning Curves

It is commonly assumed that the tuning of a neuron's activity to a movement parameter directly reflects its effect on movement. For example, a cell's PD is thought to represent the direction of force it generates. However, our model shows that the cells' PDs deviate randomly from the force directions ([Figure 5A](#), right). For the parameter values we used, the mean absolute difference between PDs and force directions was about 40° . We conclude that the tuning of cells to motor parameters does not uniquely specify their effect on movement, but rather specifies how the cells are recruited to produce the movement.

Doesn't the Theory Imply that Neural Representations Have No Spatial Order?

Recent studies report that nearby motor cortical cells tend to have similar directional tuning, more than expected by a completely random arrangement ([Amirikian and Georgopoulos, 2003](#); [Ben-Shaul et al., 2003](#); [Cheney and Fetz, 1985](#)). In our model, PDs are correlated with the directions of endpoint forces generated by the cells ([Figure 5A](#)). Therefore, if the force directions are spatially organized within cortex, then PDs should also be spatially organized. In our simulations, PDs of cells which produce similar force directions are on average 55° apart, and thus we predict that PDs of nearby cells differ on average at least by 55° (this would be the case if nearby cells pro-

duced exactly the same force directions). This is consistent with the finding of [Ben-Shaul et al. \(2003\)](#) that during movement time PDs of nearby cells differ on average by 75° .

What Is the Function of Plasticity Noise?

One possibility is that plasticity noise hinders learning but has not been eliminated over the course of evolution because its effects are small. However, it is also possible that plasticity noise is useful for learning. It is well known in learning theory that adding noise to the learning process may prevent settling in local minima of the performance ([Kirkpatrick et al., 1983](#)). Additionally, it was proposed that stochastic plasticity is useful for preventing newly formed memories from overriding existing memories ([Fusi, 2002](#)). Finally, there is a whole class of learning methods, known as reinforcement learning, which is based on noise. In reinforcement learning, noise is injected into the system in order to probe different possible outputs and evaluate their success ([Sutton and Barto, 1998](#)). Recent studies suggest how reinforcement learning algorithms can be implemented in biophysically realistic, spiking neural circuits ([Fiete and Seung, 2006](#); [Seung, 2003](#); [Xie and Seung, 2004](#)).

What Is the Function of Redundancy?

Redundancy provides robustness to damage and noise. Additionally, motor cortex may be highly redundant with respect to a given task because it needs to store in the same neural circuit motor memories related to other tasks. This scenario can be visualized with the concept of the optimal manifold, which is the continuum of all synaptic configurations appropriate for a given task. For example, teaching a neural circuit two tasks can be described as moving the synaptic strengths to a configuration in the intersection of the two manifolds optimal for these tasks ([Figure 8B](#)).

How Would Our Results Generalize to More Complicated Networks?

Our linear model network tunes $2N$ synapses to perform a 2×2 linear transformation. Consequently, it has a linear optimal manifold of dimension $2N-4$. Generally, in a linear network the manifold dimension is the difference between the number of synapses and the number of constraints imposed by the task. In a nonlinear network, the optimal manifold is curved, and we speculate that its dimension is roughly equal to the difference between the total number of synapses and the number of synapses actually needed to solve the task. At present, there is no theory describing the nature of these manifolds.

Statistics of Modulation Depths Can Be Explained by Assuming Plasticity of Neuronal Excitability

While our model accounts reasonably well for the statistics of the PDs, it does not describe well the statistics of the modulation depths. The model predicts that the distribution of modulation depths should peak at intermediate values, whereas the empirical distribution is peaked at

low values. Additionally, the model underestimates the degree of modulation depth changes. We found that these problems are remedied if we assume that neuronal excitability is also changed plastically by learning (see [Supplemental Data](#)). Previous studies show that neuronal excitability is plastic and that changes in excitability are correlated with learning (Zhang and Linden, 2003). Thus, task related information may be stored in both synapses and intrinsic cellular properties.

What Are the Limitations of Our Model?

In this work, we chose the simplest model that illustrates that the neural representation of a redundant system may be inherently unstable. Because of its simplicity, our model did not capture accurately certain aspects of the data. First, the model readapted to the baseline condition as fast as it learned the novel task (Figure 7), whereas the monkeys usually readapted to the baseline faster and more completely than they adapted to the forces (Figure 1). Second, the distributions of PD changes generated by the model tended to have heavier tails than the empirical distributions (compare Figures 3B and 6B). Third, the autocorrelation of the PDs in the model had a slight curvature which was not observed in the data (compare Figures 3C and 6C). Fourth, the distribution of learning related PD changes in the model was more biased than the empirical distribution (compare Figures 3E and 6E).

Another limitation of our model is that synapses have a single forgetting time constant on the order of thousands of trials. We chose this time constant to fit the rate of background changes observed in our data. Since some motor tasks are retained over many years without practice, it is more plausible that there are multiple synaptic forgetting time constants, some of which are very long (Fusi, 2002). To address this issue, we have extended our model to include several synaptic forgetting time constants. We found that the model can reproduce the observed rate of background changes and yet in the absence of sensory feedback partially retain motor memories for indefinitely long times (see [Supplemental Data](#)). Finally, in the future, our model should be extended to allow storage of multiple sensorimotor transformations, by including contextual inputs (e.g., Salinas, 2004).

EXPERIMENTAL PROCEDURES

Task

All experimental procedures adhered to NIH guidelines on the use of animals and were approved by the MIT Committee for Animal Care. Two rhesus macaques (*Macaca mulatta*) were trained for at least 4 months on the visuomotor reaching paradigm described in Li et al. (2001). The animals sat in a chair and with their right arm held onto a handle at the end of a two-link robotic manipulandum, whose endpoint mapped to a cursor on a monitor. On each trial, they moved the handle in order to move the cursor from a center target displayed on the monitor to one of eight peripheral targets, uniformly located around a circle. Each trial began with a 1 s hold time at the center target, followed by the presentation of a pseudorandomly chosen peripheral target (i.e., the cue). The center target remained on for a variable 0.5 to 1.5 s after the cue to indicate the instructed delay time. Upon dis-

appearance of the center target (i.e., the go signal), the monkey made an 8–10 cm reaching movement to place the cursor in the peripheral target, where it had to remain for 1 s to receive a juice reward. Movements had to be less than 3 s and confined to $\pm 60^\circ$ about a line connecting the center and peripheral targets. Any error resulted in abortion of the trial without reward. The hand trajectory on each trial was recorded and saved for analysis.

In the control experiment, which typically lasted 1–2 hr, the monkeys performed 480 correct trials with no external forces. In the learning experiment, the monkeys performed 160 correct trials with no external forces (baseline block), followed immediately by another 160 correct trials during which the robotic manipulandum applied forces on the hand that were proportional and perpendicular to its velocity vector (adaptation block), and finally another 160 correct trials with no external forces (washout block). The magnitude of the velocity-dependent force field was 6 Ns/m.

Neural Recordings

After sufficient training, a head-restraining device was fixed to the skull and a craniotomy (28 mm diameter) was performed under aseptic conditions. The craniotomies were centered, relative to the interaural line and midline, at anterior 22 mm, lateral 0 mm for one monkey (SMA recordings) and anterior 20 mm, lateral (left hemisphere) 15 mm for the other (M1 recordings). Intracortical microstimulation (50 ms trains of biphasic pulses at 330 Hz, with 0.2 ms pulse duration and 10–120 μ A pulse amplitude) was used to map out the proximal arm representation in each cortical area. Extracellular recordings were made from these locations using epoxyite-insulated tungsten microelectrodes (1–3 M Ω impedance). Up to eight recording electrodes were used in each session, each lowered with a manual microdrive with the goal of having one well-isolated cell per electrode. The recordings were preamplified at the headstage (unity gain), amplified (10,000 gain), and filtered (300 Hz to 10 kHz, passband). Action potentials were detected by a manually determined threshold crossing, and the spike times and behavioral task event times were saved for off-line analysis. Spike waveforms were digitized (1.00–1.75 ms duration) and saved for subsequent spike sorting.

Spike sorting was done manually, with the aid of software packages (Autocut 3, DataWave Technologies; MClust 3.3, A. David Redish, University of Minnesota), by detecting clusters in spike waveform feature space. Clusters of spikes were assumed to come from one neuron if they were (1) reasonably separated from other clusters and noise spikes in feature space, (2) had temporally continuous, if not constant, waveform features, and (3) exhibited at least a 1 ms refractory period. To assess how the quality of spike sorting impacted our results, some of the analyses described below were repeated on a subset of neurons which were judged subjectively to be (1) the best isolated, by having no overlap between their clusters and other clusters or noise spikes in at least one projection of feature space, and (2) the most stable, by having temporally constant waveform features. The results of our analysis were similar whether we included all cells or just the best isolated, most stable cells (Table 1). As further evidence that unstable tuning was not due to unstable recordings, we show several examples of tuning curve instabilities in stably recorded cells (Figure 2).

Data Analysis

We analyzed 136 cells (93 from M1 and 43 from SMA) recorded in the control experiment and 304 cells (105 from M1 and 199 from SMA) recorded in the learning experiment. For each cell and each trial, we computed the average firing rate between 100 ms prior to movement onset and 300 ms after movement onset. We identified movement onset as the last time at which hand speed crossed a 4 cm/s threshold prior to the time of peak speed. For cells recorded during control sessions, we divided the trials into three consecutive blocks of 160 trials. For cells recorded during learning sessions, we divided the data into entire baseline block, last 80 trials of adaptation block, and last 80 trials of washout block. The first 80 trials of the adaptation and washout

were excluded to focus only on the postadaptation phase (cf. Li et al. [2001]).

We estimated the tuning curve of each cell in each block by eight mean firing rates corresponding to the different movement directions. We fitted each tuning curve by a sum of an offset (mean of the eight firing rates), and a cosine function (Equation 1). To fit the cosine function, we defined the two-dimensional AC vector

$$AC = \frac{1}{4} \sum_{k=1}^8 \begin{pmatrix} \cos \theta_k \\ \sin \theta_k \end{pmatrix} r_k \quad (2)$$

where θ_k are the movement directions and r_k are the mean firing rates. The amplitude A and phase ψ of the cosine were set to the magnitude and direction of AC, respectively. This commonly used method minimizes the squared error between the cosine function and the mean firing rates. We used a t test to estimate the significance of changes across blocks of mean firing rates at individual movement directions. Because the offsets and cosine components are averages over a large number of trials (160), to test the significance of their changes we used a z test (t statistic with a Gaussian null distribution; see Montgomery and Runger [1999]). To test the changes in the cosine functions, we used a bivariate z test on the two-dimensional AC vectors, thus testing for changes in PD and/or modulation depth (see Christensen [2001] on Hotelling's t statistic for multivariate data). To decide whether a tuning curve has a significant cosine component, we tested whether the AC vector is significantly different from zero by a bivariate z test with the assumption of isotropic noise. For the analysis of PDs, we chose only cells which had significant cosine components in all three blocks with $p < 0.05$, including 93 cells in the control experiment (59 from M1 and 34 from SMA) and 172 cells in the learning experiment (67 from M1 and 105 from SMA).

To test whether PD changes have a nonzero mean, we used a z test. The correlation between PD changes across cells and across time was estimated by Pearson's correlation coefficient. To estimate the significance of the correlation coefficient, we used a nonparametric permutation test. For example, for data $(x_1, y_1), \dots, (x_n, y_n)$, we randomized the y data with respect to the x data 1000 times and recompute the correlation coefficient for each iteration. The p value of the correlation coefficient is the fraction of times the simulated correlation coefficient has an absolute value larger than the real correlation coefficient.

In order to estimate the autocorrelation of the population of PDs in the control experiment, we binned the data into 12 blocks of 40 trials. We computed the PD of each cell within each block as described above. We defined the correlation between PDs in two bins numbered k and m as

$$c(k, m) = \frac{1}{N_{\text{cells}}} \sum_{i=1}^{N_{\text{cells}}} \cos[\psi_i(k) - \psi_i(m)] \quad (3)$$

where $\psi_i(k)$ is the PD of cell i in bin k . The autocorrelation function with lag k was estimated by

$$ACF(k) = \frac{1}{N_{\text{bins}} - k} \sum_{m=1}^{N_{\text{bins}} - k} c(m, m+k) \quad (4)$$

where $N_{\text{bins}} = 12$.

Model Equations

Model of Reaching

When presented with a target in direction θ , the two sensory inputs are activated proportionally to the target coordinates

$$\begin{aligned} x_1^t &= \cos \theta \\ x_2^t &= \sin \theta. \end{aligned} \quad (5)$$

The two sensory inputs activate N motor cortical cells

$$r_i = \sum_{j=1}^2 W_{ij} x_j^t \quad (6)$$

where i is the cell number and W_{ij} are the input weights of the cortical cells. r_i is interpreted as firing rate averaged over movement time relative to baseline firing rate before movement and therefore may be negative. The motor cortical cells generate an endpoint force

$$f_i = \sum_{j=1}^N Z_{ij} r_j \quad (7)$$

where $i = 1, 2$ indexes the force components and Z_{ij} are the cells' output weights. Z_{ij} are fixed to

$$\begin{aligned} Z_{1,j} &= \frac{2}{N} \cos \alpha_j \\ Z_{2,j} &= \frac{2}{N} \sin \alpha_j \end{aligned} \quad (8)$$

where α_j are the directions of forces generated by the neurons, which we distribute uniformly. We normalize Z_{ij} by $1/N$ such that firing rates of order 1 produce a force of order 1. The final hand coordinates are

$$x_i = \sum_{j=1}^2 R_{ij} f_j \quad (9)$$

where R_{ij} is the 2×2 identity matrix without the perturbation and a rotation of angle φ with the perturbation.

Model of Plasticity

The task of the network is to have $x_i = x_i^t$. In order to learn this task, W_{ij} are incremented after each trial by

$$\Delta W_{ij} = -\frac{W_{ij}}{\tau_{\text{forget}}} + \sigma n_{ij} - \frac{N}{\tau_{\text{learn}}} \frac{\partial E(x, x^t)}{\partial W_{ij}}. \quad (10)$$

The second term is noisy synaptic changes, where n_{ij} are unbiased normalized i.i.d. Gaussian noise components and σ is the noise amplitude. The first term is a decay term which prevents W_{ij} from drifting without bound. When only the first two terms are present (Figure 5B), the synaptic weights perform a leaky random walk process, with a time constant τ_{forget} and a variance which scales as $\sigma^2 / \tau_{\text{forget}}$. The third term is a gradient descent learning signal, which is a method commonly used for teaching artificial neural networks (Rumelhart et al., 1986). This method optimizes a cost function E with respect to the network weights W_{ij} by making small steps of W_{ij} in the direction which decreases E the most. We use the squared error cost

$$E(x, x^t) = \frac{1}{2} \sum_{i=1}^2 (x_i - x_i^t)^2. \quad (11)$$

The gradient of this cost is related to the error by

$$\frac{\partial E(x, x^t)}{\partial W_{ij}} = \sum_{k=1}^2 (x_k - x_k^t) Z_{ki} x_j^t. \quad (12)$$

Because the gradient scales as $1/N$, we introduce a prefactor N in front of the learning signal in Equation 10. The gradient with respect to a synapse depends on information not local to that synapse, e.g., Z_{ki} in Equation 12. However, previous work has shown that the gradient can be computed by correlating noise in synaptic transmission with a global reward signal (Seung, 2003). Such a learning rule produces a noisy estimate of the gradient, similar to our noisy gradient learning rule (Equation 10).

Model Parameters

For Figure 6, we set $\tau_{\text{learn}} = 50$, $\tau_{\text{forget}} = 1500$, $\sigma = 0.025$, $\varphi = 60^\circ$, $N = 10000$. The model performs well at $N \sim 100$ (Supplemental Data), but for better statistics we set $N = 10000$. τ_{learn} was set according to the observed learning time constants. We set σ to a value which in the control simulation produced PD changes of a magnitude similar to the observed magnitude. τ_{forget} was set to reproduce the rate of the experimentally observed background changes. φ was fit to reproduce the observed anticorrelation between adaptation and washout changes. All simulations started with 10,000 trials of pretraining, which

is considerably longer than τ_{forget} (1500 trials), and thus our results do not depend on the initial W_{ij} (which is zero).

Supplemental Data

The Supplemental Data for this article can be found online at <http://www.neuron.org/cgi/content/full/54/4/653/DC1/>.

ACKNOWLEDGMENTS

We thank Camillo Padoa-Schioppa for collecting the SMA data, Michale Fee, Michael Long, Beata Jarosiewicz, Simon Overduin, Sen Song, Srinivas Turaga, Olivia White, and Dan Rokni for commenting on this manuscript and Yonatan Loewenstein for helpful discussions.

Received: September 21, 2006

Revised: January 26, 2007

Accepted: April 30, 2007

Published: May 23, 2007

REFERENCES

- Amirikian, B., and Georgopoulos, A. (2003). Modular organization of directionally tuned cells in the motor cortex: is there a short-range order? *Proc. Natl. Acad. Sci. USA* *100*, 12474–12479.
- Ben-Shaul, Y., Stark, E., Asher, I., Drori, R., Nadasdy, Z., and Abeles, M. (2003). Dynamical organization of directional tuning in the primate premotor and primary motor cortex. *J. Neurophysiol.* *89*, 1136–1142.
- Brainard, M.S., and Doupe, A.J. (2000). Auditory feedback in learning and maintenance of vocal behaviour. *Nat. Rev. Neurosci.* *1*, 31–40.
- Cheney, P.D., and Fetz, E.E. (1985). Comparable patterns of muscle facilitation evoked by individual corticomotoneuronal (CM) cells and by single intracortical microstimuli in primates: evidence for functional groups of CM cells. *J. Neurophysiol.* *53*, 786–804.
- Christensen, R. (2001). Multivariate linear models. In *Advanced Linear Modeling*, G. Casella, S. Fienberg, and I. Olkin, eds. (New York: Springer-Verlag), pp. 1–73.
- Cohen, D., and Nicolelis, M.A. (2004). Reduction of single-neuron firing uncertainty by cortical ensembles during motor skill learning. *J. Neurosci.* *24*, 3574–3582.
- Fiete, I.R., and Seung, H.S. (2006). Gradient learning in spiking neural networks by dynamic perturbation of conductances. *Phys. Rev. Lett.* *97*, 048104.
- Fusi, S. (2002). Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cybern.* *87*, 459–470.
- Gandolfo, F., Li, C., Benda, B., Padoa-Schioppa, C., and Bizzi, E. (2000). Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl. Acad. Sci. USA* *97*, 2259–2263.
- Greenberg, P., and Wilson, F. (2004). Functional stability of dorsolateral prefrontal neurons. *J. Neurophysiol.* *92*, 1042–1055.
- Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* *42*, 283–295.
- Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983). Optimization by simulated annealing. *Science* *220*, 671–680.
- Li, C.S., Padoa-Schioppa, C., and Bizzi, E. (2001). 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.
- Mitz, A.R., Godschalk, M., and Wise, S.P. (1991). Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. *J. Neurosci.* *11*, 1855–1872.
- Montgomery, D.C., and Runger, G.C. (1999). *Applied Statistics and Probability for Engineers*, 2nd Edition (New York: John Wiley & Sons Inc.).
- Mussa-Ivaldi, F.A. (1988). Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci. Lett.* *97*, 106–111.
- Nakamura, K., Sakai, K., and Hikosaka, O. (1998). Neuronal activity in medial frontal cortex during learning of sequential procedures. *J. Neurophysiol.* *80*, 2671–2687.
- Nicolelis, M.A., Ghazanfar, A.A., Faggin, B.M., Votaw, S., and Oliveira, L.M. (1997). Reconstructing the engram: simultaneous, multisite, many single neuron recordings. *Neuron* *18*, 529–537.
- Ojakangas, C.L., and Ebner, T.J. (1992). Purkinje cell complex and simple spike changes during a voluntary arm movement learning task in the monkey. *J. Neurophysiol.* *68*, 2222–2236.
- Padoa-Schioppa, C., Li, C.S., and Bizzi, E. (2002). Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. *Neuron* *36*, 751–765.
- Padoa-Schioppa, C., Li, C.S., and Bizzi, E. (2004). Neuronal activity in the supplementary motor area of monkeys adapting to a new dynamic environment. *J. Neurophysiol.* *91*, 449–473.
- Paz, R., and Vaadia, E. (2004). Specificity of sensorimotor learning and the neural code: Neuronal representations in the primary motor cortex. *J. Physiol. (Paris)* *98*, 331–348.
- Paz, R., Boraud, T., Natan, C., Bergman, H., and Vaadia, E. (2003). Preparatory activity in motor cortex reflects learning of local visuomotor skills. *Nat. Neurosci.* *6*, 882–890.
- Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature* *323*, 533–536.
- Salinas, E. (2004). Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *J. Neurosci.* *24*, 1113–1118.
- Salinas, E., and Abbott, L.F. (1995). Transfer of coded information from sensory to motor networks. *J. Neurosci.* *15*, 6461–6474.
- Schmidt, E.M., Bak, M.J., and McIntosh, J.S. (1976). Long-term chronic recording from cortical neurons. *Exp. Neurol.* *52*, 496–506.
- Schwartz, A.B. (2004). Cortical neural prosthetics. *Annu. Rev. Neurosci.* *27*, 487–507.
- Seung, H.S. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* *40*, 1063–1073.
- Sutton, R., and Barto, G. (1998). *Reinforcement Learning: An Introduction* (Cambridge, MA: The MIT Press).
- Taylor, D.M., Tillery, S.I., and Schwartz, A.B. (2002). Direct cortical control of 3D neuroprosthetic devices. *Science* *296*, 1829–1832.
- Thompson, L.T., and Best, P.J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res.* *509*, 299–308.
- Todorov, E. (2000). Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* *3*, 391–398.
- Williams, J.C., Rennaker, R.L., and Kipke, D.R. (1999). Stability of chronic multichannel neural recordings: implications for a long-term neural interface. *Neurocomputing* *26–27*, 1069–1076.
- Wise, S.P., Moody, S.L., Blomstrom, K.J., and Mitz, A.R. (1998). Changes in motor cortical activity during visuomotor adaptation. *Exp. Brain Res.* *121*, 285–299.
- Xiao, J., Padoa-Schioppa, C., and Bizzi, E. (2006). Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. *Exp. Brain Res.* *168*, 106–119.
- Xie, X., and Seung, H. (2004). Learning in neural networks by reinforcement of irregular spiking. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* *69*, 041909.
- 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.