

The Neuronal Basis of Long-Term Sensorimotor Learning

Yael Mandelblat-Cerf,^{1,5} Itai Novick,^{1,5} Rony Paz,² Yuval Link,³ Sharon Freeman,¹ and Eilon Vaadia^{1,4,5}

¹Department of Medical Neurobiology, Institute for Medical Research Israel–Canada, Faculty of Medicine, Hebrew University, Jerusalem 91120, Israel,

²Department of Neurobiology, The Weizmann Institute of Science, Rehovot, 76100, Israel, ³Department of Physiology, The Faculty of Medicine and

Research Institute, Technion, Haifa 32000, Israel, and ⁴The Interdisciplinary Center for Neural Computation and ⁵The Edmond and Lily Safra Center for Brain Sciences, Hebrew University, Jerusalem 91904, Israel

The brain has a remarkable ability to learn and adjust behavior. For instance, the brain can adjust muscle activation to cope with changes in the environment. However, the neuronal mechanisms behind this adaptation are not clear. To address this fundamental question, this study examines the neuronal basis of long-term sensorimotor learning by recording neuronal activity in the primary motor cortex of monkeys during a long-term adaptation to a force-field perturbation. For 5 consecutive days, the same perturbation was applied to the monkey's hand when reaching to a single target, whereas movements to all other targets were not perturbed. The gradual improvement in performance over these 5 days was correlated to the evolution in the population neuronal signal, with two timescales of changes in single-cell activity. Specifically, one subgroup of cells showed a relatively fast increase in activity, whereas the other showed a gradual, slower decrease. These adapted patterns of neuronal activity did not involve changes in directional tuning of single cells, suggesting that adaptation was the result of adjustments of the required motor plan by a population of neurons rather than changes in single-cell properties. Furthermore, generalization was mostly expressed in the direction of the required compensatory force during adaptation. Altogether, the neuronal activity and its generalization accord with the adapted motor plan.

Introduction

Sensorimotor behavior involves computation of the relation between sensorial guidance and actions. One ingenious feature of the brain is its ability to adapt to changing relations between instructions, actions, and feedback. Psychophysical studies provided evidence for the nature of sensorimotor learning by testing features such as dynamics of acquisition, generalization, and consolidation (Krakauer et al., 1999; Donchin et al., 2003; Caithness et al., 2004). In parallel, physiological studies have demonstrated that neuronal activity in the motor cortex conveys information about the characteristics of the required movement, such as the desired direction (Georgopoulos et al., 1982, 1983) and muscle forces (Evarts, 1968), which can change through learning (Mitz et al., 1991; Chen and Wise, 1995; Wise et al., 1996, 1998; Gandolfo et al., 2000; Paz et al., 2003; Zach et al., 2008).

One of the commonly used learning paradigms in this field is based on force-field perturbations during reaching movements

(Lackner and Dizio, 1994; Shadmehr and Mussa-Ivaldi, 1994; Burdet et al., 2001). The findings for both human and nonhuman primates indicate that subjects learn to predict and compensate for an upcoming force. Bizzi's group studied the neuronal basis of adaptation to a curl force-field (Gandolfo et al., 2000; Li et al., 2001; Padoa-Schioppa et al., 2002) during center-out reaching movements to eight directions. They reported that, late in learning (when trajectories were straight), the preferred directions (PDs) of the cells shifted with force-field direction. When the force-field was removed and aftereffects diminished, some cells, termed "memory cells," still maintained learning-related changes, although the average PD shifts of the population diminished (Li et al., 2001).

Previous studies by our group examined the relation between a specific learned direction and the PDs of cells using single-target force-field adaptation (Arce et al., 2010a). The disadvantage of this design was that it did not allow direct estimation of directional tuning (and PD) of cells during learning, as was done by Bizzi's group.

In our new task design, we track global neuronal and behavioral features during local learning by introducing force-field in movements to a single target, whereas other targets continue to be experienced unperturbed. This enables us to monitor the tuning of cells during local adaptation and study the relation between force-field adaptation and previously reported PD shifts. Furthermore, we study the temporal dynamics of adaptation over several days to test whether learning emerges by processes with multiple timescales (Smith et al., 2006; Fusi et al., 2007; Yamashita and Tani, 2008).

The findings provide evidence for slow and fast timescales of changes in single cells activity during learning, which seem com-

Received Aug. 4, 2010; revised Oct. 10, 2010; accepted Oct. 14, 2010.

This work was supported in part by the Binational Science Foundation (BSF), the Israel Science Foundation, special contributions from the Rosetrees Trust, and the Ida Baruch fund (to the laboratory of E.V.). It was also supported by grants from the Israel Science Foundation and Grant FP7-IRG from the Marie-Curie Actions (to the laboratory of R.P.). E.V. is the Jack H. Skirball Chair of Brain Research. The study is in partial fulfillment of the doctoral thesis of Y.M.C.

The authors declare that they have no competing financial interest.


Y.M.C. and E.V. designed the experiment. Y.M.C. and I.N. conducted the FF experiments. R.P. designed and conducted the VR experiment. Y.M.C. did the data analysis and writing. S.H. assisted in writing. Y.L. provided technical support. E.V. supervised the project, experimentation, data analysis, and write up.

Correspondence should be addressed to Yael Mandelblat-Cerf, The Institute for Medical Research, Faculty of Medicine, The Hebrew University of Jerusalem, Ein Kerem, P.O. Box 12271, Jerusalem, 91120 Israel. E-mail: yaelma@ekmd.huji.ac.il.

DOI:10.1523/JNEUROSCI.4055-10.2011

Copyright © 2011 the authors 0270-6474/11/310300-14\$15.00/0

A Trial flow during a learning block (LRN)



B Sessions (days) flow:




Figure 1. Experimental design. **A**, Example of a trial flow (left to right) during the learning epoch. During the first delay period, the monkey held the robotic arm in the center without moving it. The monkeys maintained their hold at the central circle after target onset for an additional delay and only moved after the go signal. In the figure, the learned target is at 0° and the force-field is clockwise. In the rightmost panels, if the lit target was the selected learned target (bottom row), force-field was applied when movement was initiated, otherwise (top row) the movement was executed under standard conditions. **B**, Recording day flow: all days started with standard trials (center-out reaching movements to 8 directions) followed by a learning epoch. In days 1–4, this was followed by a second standard period and then ended with a second learning period. Days 5 and 6 consisted of three different periods only. Learning was followed by a long standard epoch (WASHOUT).

parable with the shifts in the PDs of the cells as reported previously. However, these apparent shifts can better be explained by the growing discrepancy between target direction and intended hand-movement direction, which leads to a systematic difference in estimations of PDs as a function of these two reference frames. Thus, we suggest that the relation between activity of single cells in motor cortex and the intended arm movement remains the same throughout learning, and adaptation is achieved by adjusting the contributions of motor cortical cells to the population signal that represents the intended movement. Further strengthening this reasoning, the generalization pattern of the population signal reflected the intended movement direction.

Materials and Methods

Animals, recordings, and behavioral task

Two monkeys (one male and one female, *Macaca fascicularis*, 4 kg.) were chronically implanted with a microelectrode array (Cyberkinetics Neurotechnology Systems) in the arm region of M1 contralateral to the performing arm, under anesthesia and aseptic conditions. Animal care and surgical procedures complied with the National Institutes of Health *Guide for the Care and Use of Laboratory Animals* and with guidelines defined by the Institutional Committee for Animal Care and Use at the Hebrew University.

Monkeys used a robotic arm (Phantom Premium 1.5 High Force; SensAble Devices) to control the movements of a cursor on a video screen in a two-dimensional plane. Before surgery, monkeys were well trained to perform a standard eight-target center-out reaching task using this device (“null conditions”). The phantom moved the cursor from the starting point at the center of the screen to a visual target in a delayed go-signal paradigm. The trial sequence and recording day flow is shown in Figure 1. Figure 1A depicts the trial flow from left to right. Each trial began when the monkey positioned the cursor at the central circle. After a variable hold period of 0.85–1.35 s, a target appeared at one of the eight possible positions, which were uniformly distributed in a circle 4 cm from the center (Fig. 1A, second column). After an additional 0.85–1.35 s hold period, the central circle disappeared (go signal, third column), prompting the monkey to move to the target in <0.8 s (fourth column). This

generous time constraint allowed relatively natural reaching movements. After another 0.4 s, a liquid reward was delivered.

A learning week consisted of 6 consecutive days. The same perturbation was applied for the entire learning week: a single target [the “learned target” (LT)] and force-field (FF) direction (clockwise or counterclockwise). Days 1–4 comprised four successive periods (Fig. 1B): (1) a default (standard) eight-target task (STD1) (without FF) of 80 trials; (2) a learning period (LRN) (with force-field) of 240 trials; (3) a second default eight-target task (STD2) of 80 trials, and (4) a second learning period (LRN2). Days 5 and 6 involved only three consecutive periods: (1) STD1, (2) LRN with at least 240 trials, and (3) a long STD (“washout”) of at least 360 trials. No cue was given to the monkey to mark the transitions between standard and learning epochs.

Note that the number of default trials in STD1 and STD2 were introduced with caution; the default trials can show, on one hand, the baseline condition at each day and the post-learning effects. On the other hand, they can interfere with the learning because they have a washout effect. We chose the number of trials in STD1 and STD2 to balance the tradeoff between these two effects: we kept the number of default trials small enough to minimize the interference and large enough to measure the behavior and neural activity in default condition. STD1 served to examine the baseline for each day, before additional learning took place. For example, it made it possible to estimate directional tuning of all cells without perturbation as well as the overnight retention of learning, and STD2 provided a rapid assessment of the learning effect without the perturbation present, with minimal washout. On days 5 and 6, STD2 was replaced by a long STD epoch to negate the learning effect (washout).

In standard periods, the sequence of target appearances was chosen randomly and was executed without any perturbation. During learning periods, targets continued to appear randomly. However, whenever the selected learned target appeared, movement was executed under the curl force-field, which perturbed the hand perpendicular to its direction and proportional to its velocity. Then, on day 6, we assessed the ability to recall the learned motor skill after it was washed out on day 5. Day 6 also ended with a washout to start an additional learning week with a clean slate.

Figure 1A illustrates adaptation to force-field in which the learned target was 0° and force-field was clockwise. Force-field was applied only during the learning period and only to this target. Monkeys were trained for several months with the default eight-target task but were not exposed to force-field before the recordings.

Behavioral analysis

Performance during learning. Behavioral performance was assessed by the directional deviation of the hand from the target direction taken at (1) peak velocity or (2) at movement initiation, 150 ms after movement onset (see the learning curves in Fig. 2B).

The force-field pushed the hand perpendicular to its current velocity in a counterclockwise or clockwise direction. Given the components of the observed trajectory in the horizontal plane (x and y), the force-field was generated using the following equation:

$$\begin{bmatrix} \text{FF}x(t) \\ \text{FF}y(t) \end{bmatrix} = k \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} \dot{x}(t) \\ \dot{y}(t) \end{bmatrix},$$

where $\text{FF}x$ and $\text{FF}y$ are the robot-generated forces at time sample t , $k = 8$ Ns/mm, $\theta = \pm 90^\circ$, and \dot{x} and \dot{y} are the components of the hand velocities in the horizontal plane.

“Force-vector” computation. In perturbed movements, the observed hand movement (x and y components) resulted from a vectorial summation of two forces: the force applied by the hand during this movement (“force-vector”) and the robot-generated force-field.

The total of these forces can also be directly derived by Newton’s second law of motion: total force(t) = $m \begin{bmatrix} \ddot{x}(t) \\ \ddot{y}(t) \end{bmatrix}$, where \ddot{x} and \ddot{y} are the hand acceleration components at time t , and m is the approximated mass of the hand. Following morphometry studies (Cheng and Scott, 2000; Graham and Scott, 2003), for ~ 4 kg monkeys, the mass was approximated as 250 g.

Therefore, to extract the force-vector, we assessed the total force and subtracted the external FF: force-vector(t) = $m \begin{bmatrix} \ddot{x}(t) \\ \ddot{y}(t) \end{bmatrix} - \begin{bmatrix} \text{FF}x(t) \\ \text{FF}y(t) \end{bmatrix}$. The direction of the force-vector was assessed for the initial 150 ms of the movement.

Neuronal data analysis

We selected single neurons for analysis that met four inclusion criteria: (1) well isolated spikes; (2) stable recordings based on firing rates in the first hold period before target onset, throughout all trials; (3) the result of a one-way ANOVA showed a significant effect for direction ($p < 0.01$); and (4) a cosine fit [$r(d) = a + b * \cos(d - d_0)$] for directional tuning that exceeded $R^2 = 0.65$.

The neuronal ensemble consists of all the neurons recorded simultaneously during each session.

Firing rate analysis. The firing rate of a cell for each of the eight targets was computed in three time epochs along the trial: (1) post-target onset, 100–600 ms after target onset; (2) a time window of 500 ms before the go signal, in which the target was already known but the monkey still had its hand at the central circle; and (3) movement-related activity, 200 ms before to 300 ms after movement onset.

Preferred direction analysis. Given the eight averaged firing rates for the eight movement directions for each of the epochs, PDs were computed by a cosine fit [$r(d) = a + b * \cos(d - d_0)$] (Georgopoulos et al., 1982) by three different reference frames. (1) Angles were taken as those of the corresponding targets, i.e., 0, 45, 90 . . . (see Fig. 5, red). The PD was accordingly termed “PD_{target}.” (2) For each target direction, the attributed angle was the averaged initial hand-movement directions (see Fig. 5, black). The PD was termed in this case “PD_{hand}.” (3) For each target direction, the attributed angle was the direction of the actively produced force of the hand itself (“force-vector”), which we assumed reflects the motor plan direction. The PD was termed in this case “PD_{plan}.”

For nonperturbed movements, we assumed that the hand direction was the force-vector direction (see Fig. 5, green); therefore, for standard epochs, PD_{plan} is the same as PD_{hand}. However, for the perturbed movements, the observed hand direction was a vectorial summation of the force-vector and the external force-field. Therefore, for the learning epoch, we extracted the force-vector in movements to the learned target (see Materials and Methods).

Differences between estimation of the PD of the cell. We defined a change in PD for a single cell from one epoch to the other as PD_{f1,epoch1} – PD_{f2,epoch2}, which is the difference between the PD in epoch1 computed according to the f1 reference frame and the PD in epoch2 computed according to f2 reference frame, in which f1, f2 can be hand, target, or plan. This PD difference was normalized to the force-field direction such that it was positive if PD_{f1,epoch1} relative to PD_{f2,epoch2} was with force-field direction and negative if in the opposite direction.

PD computations for “repetitive” local learning. In local visuomotor rotation adaptation previously conducted in our laboratory (Paz et al., 2003), only the learned target appeared during learning. The perturbation of visuomotor rotation creates a discrepancy between the cursor and hand direction; to bring the cursor to the learned target and complete the trial successfully, the hand needs to move in a different direction according to the applied angular rotation.

Therefore, to evaluate directional tuning during learning, we assumed that movement-related neuronal activity to all unlearned targets remained as in standard trials, whereas the neuronal activity to the learned target changed according to learning-related modulations. This assumption

is based on the facts that (1) during learning, movement-related activity followed the hand direction, (2) a comparison of prelearning to postlearning trials revealed no changes in this activity to all directions, and (3) the narrow generalization in behavior reported previously (Paz et al., 2005).

Therefore, for each of the seven nonlearned target directions, we computed the averaged movement-related firing rate across trials in the standard epoch and for the learned target, across perturbed trials late in the learning epoch.

PD_{target} was computed by regressing these eight activities against the angles of the targets.

To compute the tuning curve according to the hand reference frame (PD_{hand}), the learned target angle was replaced by the angle to which the hand actually moved under the perturbation (according to the applied visuomotor rotation).

Note that, in this case, because the actual hand movement was unperturbed, PD_{plan} and PD_{hand} are the same.

Error bars indicate SEM. Unless stated otherwise, significance was assessed by t test. Large asterisks show 1% significance, and small asterisks show 5% significance, using the Holm–Bonferroni method to correct for multiple comparisons.

Results

Long-term local adaptation to force-field was used to elucidate the neuronal basis of force-field adaptation, which has been shown previously to elicit changes in directional tuning (Li et al., 2001). A detailed explanation of the behavioral paradigm is described in Materials and Methods and Figure 1. Briefly, the directional tuning of cells was continuously monitored during learning by the task design in which movements in eight directions were experienced in a random order, whereas force-field was applied only when reaching to one of the targets (“learned target”). In this paradigm, learning trials did not appear consecutively but, on average, once every eight trials, resulting in a slower learning rate. Unlike all previous experiments, we followed adaptation not only in a single session (Li et al., 2001; Arce et al., 2010a) but across sessions (5 d). Throughout this period, the same perturbation (force-field direction, clockwise or counterclockwise) was applied to movements toward the same learned target.

To examine the aftereffect of learning, we applied two short periods in which the monkey performed standard reaching to eight targets without perturbation during each day: the first at the beginning of the day (STD1) and the second in the middle of learning (STD2). At the end of day 5, we introduced a longer period of standard eight-target trials to washout the aftereffects. Thus, the session of day 6 could be used to test relearning after washout.

Behavioral findings

Figure 2A demonstrates the progress of behavior in one learning week of 6 d. It shows movements toward the learned target, with center location (white circle), target location (gray), and a clockwise force-field (left to right). Note that (1) movements under force-field became straighter (red) from day to day as monkeys learned to partly compensate for the force-field. As in most studies, the movements never became completely straight. (2) Movements in STD1 (dark blue) and STD2 (cyan) became more curved from day to day, indicating the adaptation aftereffects: the hand pushed against the force-field direction. These aftereffects increased from day to day. (3) The averaged trajectory late in washout of day 5 was considerably straighter (depicted by a yellow line). (4) Trajectories during STD1 on day 6 were almost straight (as predicted from the washout). Nevertheless, the readaptation to force-field on day 6 was not like learning a new task,

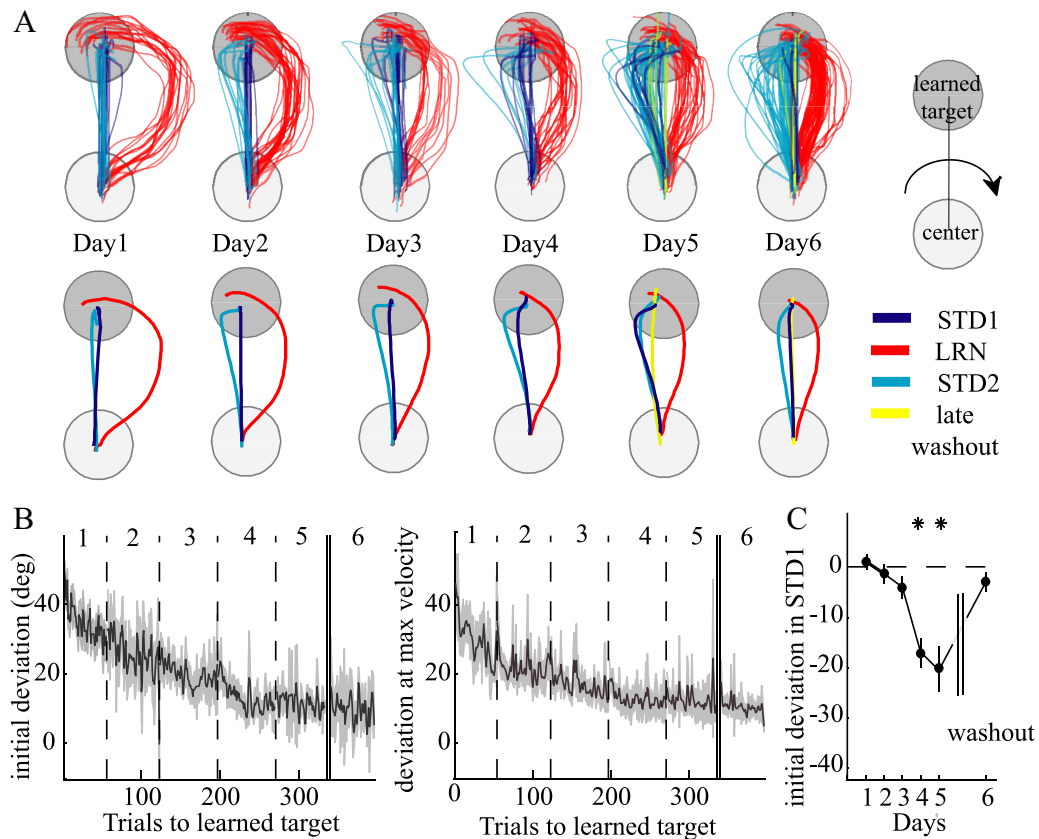


Figure 2. Movement kinematics to the learned target during standard and perturbed trials reflects gradual learning. **A**, Example of trajectories to the learned target in all epochs (top row) and their averages (bottom row) during the 6 learning days (left to right). Movements under force-field are in red. Movements without force-field are in dark blue for STD1, cyan for STD2, and yellow for late in washout. **B**, Gradual behavioral improvement during learning and its retention. The figure shows angular deviations at movement initiation (left) and at maximum velocity (right) with a five-trial moving average as a function of trial number to the learned target along the learning week (5 d and hundreds of trials). Data include 7 weeks (5 weeks from monkey R and 2 weeks from monkey O). **C**, Averaged initial angular deviations across all movements to the learned target during STD1 epoch. Significant overnight aftereffects are evident on days 4 and 5. Washout on day 5 (2 black lines) resulted in small overnight aftereffects on day 6. In all traces, error bars and shaded areas denote SEM. Asterisks denote 1% confidence.

and progress was extremely fast (as shown and discussed below in Fig. 2B, day 6), suggesting recall of existing representation of the learned perturbation.

To quantify the phenomena demonstrated above, we analyzed all the learning sessions (7 weeks) and computed the average directional deviations of trajectories to the learned target, taken at movement initiation (Fig. 2B, left trace) and at peak velocity (Fig. 2B, right trace) during the learning. Deviations were normalized to force-field direction such that positive/negative errors were with/counter to force-field direction, respectively. Previous studies of local adaptation to force-field in our laboratory (Arce et al., 2010a), in which the learned target appeared repetitively, showed that adaptation was completed very fast in a single session, reaching plateau of $\sim 15^\circ$ in a few dozen trials. In our study, 5 d and >300 trials to the learned target were necessary to reach a similar behavioral plateau.

As in previous studies (Karni et al., 1998; Krakauer et al., 2000), learning was composed of a fast stage, characterized by substantial improvements, followed by a slower stage with much smaller improvements. In this study we found that the fast stage lasted for the two first days (called here “early in learning”) and the slow stage followed on days 4 and 5 (called here “late in learning”). Day 3 was not included to avoid the gray zone between the stages.

A clear evidence for day-to-day progress of learning was provided by the overnight aftereffect of learning in standard trials that began each day. Observing the STD1 trials (Fig. 2C) showed

that hand movements to the learned target were straight only in the first day (before learning started) and gradually, from day to day, became more curved in a direction counter to the force-field that had been experienced in the previous day.

As mentioned above and demonstrated in Figure 2B, the performance under force-field reached a plateau at day 5. At the end of this day, we introduced a long epoch of standard trials of at least 360 trials. During this washout epoch, aftereffects became substantially smaller but still significantly different from zero (t test, $p < 0.01$), with an average deviation of $\sim 9^\circ$. On the following day (day 6), the aftereffect during STD1 was even smaller (curvature of 3° ; t test, $p > 0.05$), indicating that during the night the washout continued. However, as learning started (the same perturbation), deviations of trajectories were again $\sim 10^\circ$ with force-field direction, similar to performance on day 5. Therefore, although aftereffects were small in STD1, showing little retention, the learned task was successfully recalled almost immediately. Washout of additional 360 trials at the end of this day was even more complete, with minimal aftereffects not significantly different from 0 (t test, $p > 0.1$).

Examination of movement velocities revealed that movements were significantly slower in learning relative to standard trials; however, they were stable across days (supplemental Fig. 1A,B, available at www.jneurosci.org as supplemental material). One could argue that the relatively permissive parameters that

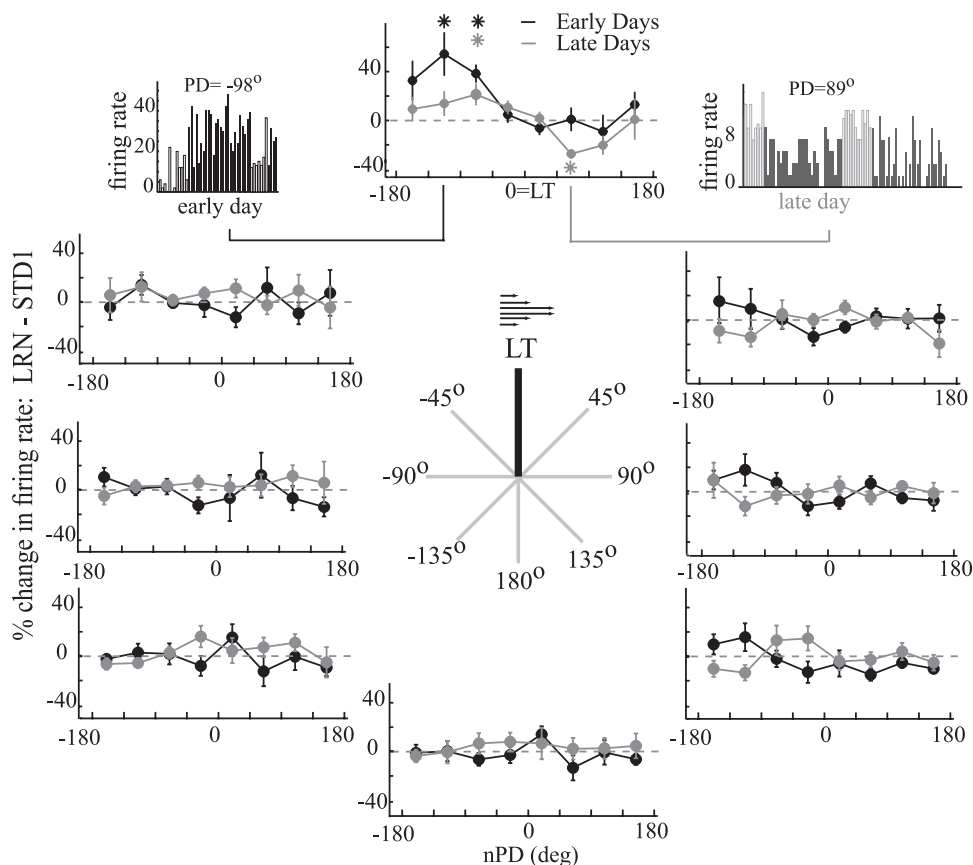


Figure 3. Changes in movement-related firing rates are only evident in movements to the learned target (LT), dependent on the preferred directions of the cells and on the learning stage. The figure shows the averaged percentage of daily changes of firing rates between standard and learning epochs, as a function of the nPDs of cells in bins in a range of 45° . Analysis was done separately for movements to each of the eight targets around the learned target. Data are divided into early (black) and late (gray) days of learning. Note that cells with nPDs around -90° from the LT (counter-FF cells) increased their activity mostly early in learning, whereas cells with nPDs around $+90^\circ$ from the LT (co-FF cells) showed some decrease late in learning. Insets provide examples of firing rates (blank bars, standard trials; full bars, learning trials), for two single cells: a cell with PD counter to force-field direction early in learning showed increased activity (left), and a cell with PD with force-field direction late in learning showed decreased activity (right). In all traces, asterisks denote 1% confidence, and error bars denote SEM. $n = 346$ in early days and $n = 264$ in late days.

allowed monkeys to successfully reach the learned target with slower movements and somewhat curved trajectories allowed the subjects to avoid adaptation to force-field using the strategy of reduced speed to reduce the force. However, the data clearly indicate that this was not the case. The relatively stable velocity during learning trials produced a similar force-field. Under this stable force-field, the curvatures were reduced (Fig. 2B) and the aftereffects gradually increased (Fig. 2C).

Neuronal findings

We analyzed neuronal activity during three different time epochs: (1) post-target onset, (2) before the go signal, and (3) movement-related activity.

The dataset includes directionally tuned cells that met the inclusion criteria as described in Materials and Methods. Of the total sample, 304 cells were directionally tuned during post-target onset (8 from monkey O, 296 from monkey R), 542 before the go signal (43 from monkey O, 499 from monkey R), and 830 during the movement-related activity (214 from monkey O, 616 from monkey R). All analyses were done for each period using only the tuned cells. Note that we did not attempt to record the same neurons on different days. Therefore, comparisons over days were done by averaging daily the changes in cell activity over the population.

Firing rate modulations

Cells were classified according to the angular distance between their PDs and the learned target, signed by the force-field direction (“nPD”). Specifically, zero nPDs were assigned to cells with PDs located at the learned target, positive values, to cells with PDs “pushing” in the direction of the force-field, and negative values pushing against the force-field. Cells with nPDs in the range of 45° to 135° were defined as “co-FF,” and cells with nPDs in the range of -135° to -45° were defined as “counter-FF.”

To study the distribution of the changes in firing rates of the population of cells in the movement period and its relation to the learning process, the cells were divided into eight subgroups according to their nPD, in bins of 45° (-180° to -135° , -135° to -90° , ..., 135° to 180°). Daily changes in firing rates were evaluated separately for movements to each of the targets in the early (Fig. 3, black) and late (Fig. 3, gray) stages of learning. The figure shows the changes in firing rates as a function of nPD, in which nonlearned targets were aligned around the learned target. The plot clearly shows that neuronal activity to each of the nonlearned targets did not change significantly during either the early or the late days of learning. Consistent with our previous results (Arce et al., 2010a), the changes in activity to the learned target depended on the nPD. On top of that, we found that it also depended on the learning stage: early in learning (black), co-FF cells did not show

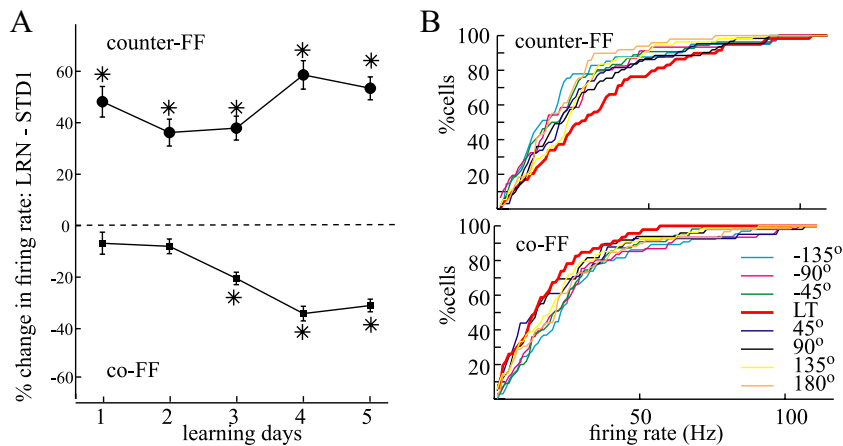


Figure 4. Dynamics of evolving learning-related activity during perturbations depend on the PD of the cell. **A**, Percentage of change in the averaged firing rates across cells with PDs in the range $-135^\circ < \text{nPD} < -45^\circ$ (counter-FF, top) and the range $45^\circ < \text{nPD} < 135^\circ$ (co-FF, bottom) compared with the averaged firing rate of cells with PDs in the same range during prelearning standard trials of day 1. The plot shows that counter-FF cells significantly increase their activity in day 1 and maintain the elevated activity throughout learning, whereas co-FF cells show a significant decrease but only on day 3. Asterisks denote 1% confidence, and error bars denote SEM. **B**, The cumulative distributions of firing rates to the learned target (LT) late in learning (thick red line) for counter-FF (left) and co-FF (right) cells. For each of the nonlearned targets (different colors, excluding the thick red line), the plots show the cumulative distribution of the firing rates of the population of cells with nPDs located in a counter-FF (left) or co-FF (right) direction relative to each target. Note that all curves lie above (left) or below (right) the red curve, which marks the respective cumulative distributions of the counter- and co-FF activities in movements to the learned target. $n = 264$.

significant changes, whereas counter-FF cells increased their firing rate during learning significantly (t test, $p < 0.01$). Late in learning (gray), the changes from STD1 to LRN of counter-FF cells were smaller than in the early days ($p < 0.01$), and the co-FF cells showed some significant ($p < 0.01$ for $45^\circ < \text{nPD} < 90^\circ$) decreases in activity. Examples of trial-by-trial activity to the learned target of two cells are shown in the insets. On the left, a counter-FF cell ($\text{nPD} = -98^\circ$), recorded at an early stage of learning (day 1), shows increased activity during learning ($p < 0.01$). On the right, a co-FF cell ($\text{nPD} = 89^\circ$) from day 4 (late stage) shows decreased activity during learning ($p < 0.05$). Each bar represents the firing rate in a single trial to the learned target; blank bars denote standard trials, and filled bars denote learning trials.

Next, we focus on the two subgroups of cells with nPDs in the co-FF and counter-FF ranges, which are the ones that showed learning-related effects.

We first assessed the cumulative changes of activity of these cells day by day. Because the cells recorded each day were not necessarily the same cells, we computed the cumulative change by averaging the population activity separately for each learning trial on each day and comparing these activities with the averaged activity across all cells and all trials to the learned target in the first epoch of standard trials before learning started (the activity on the first day of the week, STD1, day 1). This day-by-day analysis was performed separately for each of the weeks (7 weeks) and then averaged across weeks. Figure 4A shows for each numbered day (days 1–5, x -axis) the average across all learning trials for all weeks (y -axis). It shows that counter-FF cells (circles) increased their activity in early stages and became significant (50% increase on day 1, $p < 0.01$) even on the first day of the week. Thus, the fast increased firing rate in movements to the learned target of cells with nPDs in the counter-FF range shown in Figure 3 and the example of the single unit (left) represent the tendency of this group of cells to show relatively rapid increases in activity, which was maintained throughout the week. As shown in Figure 3, co-FF cells showed a tendency to decrease their firing rates in

movement to the learned target only late in learning. Figure 4A extracts the cumulative effect of this tendency and presents it day by day (squares). The slow evolution of decreased activity in these cells is clear. This was a much slower process than the one seen in counter-FF cells, which became significant much later in learning.

Second, we performed an analysis to support the notion that the observed learning-related changes were not a chance phenomenon. To do so, we tested whether increases and decreases of firing rates occurred more specifically and consistently in movements to the learned target. We studied the firing rates of cells during movements to the learned target compared with activity during movements to nonlearned targets and evaluated their cumulative distributions. Figure 4B shows the cumulative percentages of cells as a function of firing rate for each of the seven nonlearned targets and for the learned target. The figure shows the distributions for co-FF cells (right plot, thick red line) and counter-FF cells

(left thick red line) separately. The other curves (thinner lines) show, for each of the nonlearned targets, the cumulative distribution for the firing rates of cells with PDs between -135° and -45° away from that target (left) and cells with PDs between 45° and 135° away from that target (right).

The top plot shows that the percentage of cells firing above a given rate was always larger in the counter-FF subgroup during movements to the learned target (thick red line) compared with the percentage of cells above the same firing rate value in movements to any of the nonlearned targets (thin lines). The co-FF cells (bottom plot) show the opposite relation, having a lower percentage of cells accumulated for a given firing rate value.

To conclude, the analyses show opposing modulations in activity of counter-FF and co-FF cells and their different time dynamics; during learning, the activity of counter-FF cells increased significantly early in learning, resulting in activity above chance (higher than firing rates to unlearned targets, $p < 0.01$), whereas the activity of co-FF gradually decreased along the week and eventually was below chance ($p < 0.05$).

Are there learning-related changes in preferred directions?

We investigated how the changes in firing rates were reflected in the tuning properties of the cells. We assume that, in reaching tasks, subjects apply a motor plan reflecting the intention to move the arm in a desired trajectory toward the target. When the force-field perturbed the movement by pushing the hand, it introduces a difference between the direction of the target (Fig. 5A, left panels, early in learning, red dashed line) and the observed hand trajectory (black arrow). At this stage, the motor plan (early in learning, green arrow) is no longer appropriate. Then, when subject adapts, the observed hand direction again points in the target direction (Fig. 5A, left panels, late in learning, black arrow), which can only be achieved by changing the motor plan to compensate for the force-field (Fig. 5A, left panels, late in learning, green arrow). As shown in Figures 3 and 4, the activity of counter-FF cells increased and the activity of co-FF cells gradually

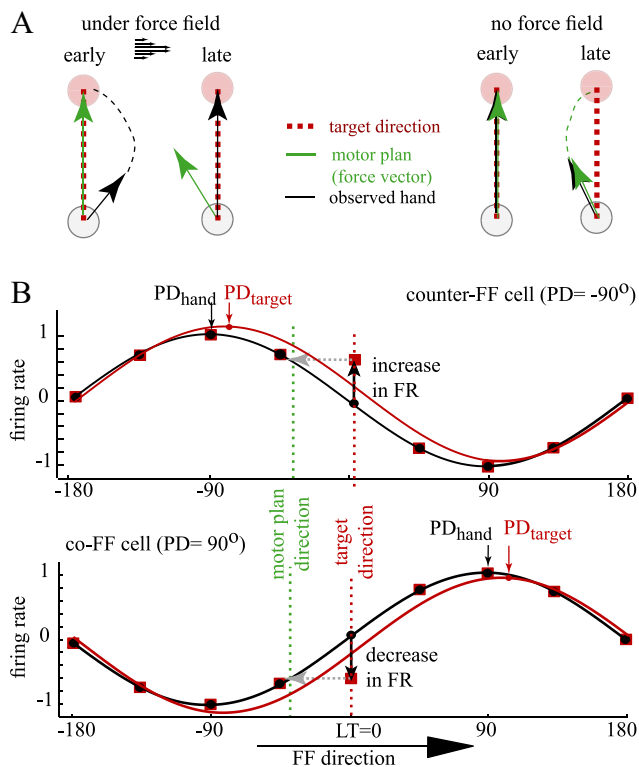


Figure 5. Directional tuning of cells under perturbation depends on the selected reference frame: schematics. **A**, Schematic chart directions of the learned target (red), the observed hand (black), and the force-vector applied by the hand (green) for early and late stages of learning. The scheme shows, during adaptation to the perturbing force-field (left), the observed hand direction gradually points to the target, whereas the motor plan becomes deviated in the opposite direction. When the force-field is removed (right), the observed hand direction always follows the new motor plan. Therefore, early in learning, it points to the target, and, later, it points counter to the force-field direction. **B**, A model of PD computation. Black curves are the directional tuning of counter-FF (up) and co-FF (down) cells. A new cosine fit for the counter-FF (or co-FF) cell with increased (or decreased) activity for the LT results in red curves, which are shifted with force-field direction. FR, Firing rate. Dotted green lines denote the direction of a motor plan that can produce such changes in activity.

decreased during adaptation. In the following analysis, we assume that these changes reflect the new motor plan and that the goal at all stages remained to move the hand to the target.

Figure 5B shows a simulated demonstration of the expected changes in tuning by the increases and decreases of firing rates of counter-FF and co-FF cells, respectively.

The top plot shows tuning of a counter-FF cell ($PD = -90^\circ$, learned target = 0°) before adaptation (black circles are firing rates and black line is the cosine fit of these data). Assuming that the cell codes for the target location, the increased firing rate of the counter-FF cell during movements to the learned target is assigned to the learned target direction (red square). As shown, this introduces a shift of the tuning curve with force-field direction (PD shifted from -90° to -81°). We term the PD when computed in this reference frame PD_{target} . Alternatively, the cell activity may reflect the adapted motor plan; the increased firing rate reflects in this case the intention to push in direction against the force-field, in a new motor plan direction (green dotted line). Therefore, the resulted tuning remained as the original one (black line). How can we test experimentally the direction of the new motor plan? Under force-field, the hand direction is always a vectorial sum of the force-vector (motor plan) and the external force applied on the hand. Therefore, from the observed hand-

movement direction, we assessed the force-vector that the hand applied (see Materials and Methods) and termed the PD when computed in this reference frame as PD_{plan} . The intended movement direction can also be estimated by measuring the initial observed hand direction in movements to the learned target immediately after removal of the force-field perturbation (Fig. 5A, right panels). When we calculated the tuning according to the initial hand direction, the preferred direction was termed PD_{hand} . Note that, during standard trials, PD_{hand} and PD_{plan} are the same because the hand is not perturbed and moves in the direction of the motor plan.

Similar to the effect of the increased activity of a counter-FF cell, the decreased firing rate of co-FF cell (Fig. 5B, bottom plot) introduced a difference in PD estimations with force-field direction, from 90° to 99° .

An example of a calculation of PD_{target} and PD_{plan} for the activity of a single cell (supplemental Fig. 2, available at www.jneurosci.org as supplemental material) late in learning shows a difference of 8° between these PD estimations.

For each cell, we computed directional tuning curves in each of the reference frames and determined its PD_{target} , PD_{hand} , and PD_{plan} . This was calculated separately for epochs STD1, LRN, and STD2.

We calculated the averaged differences between PD estimations in various epochs and the PD in the STD1 epoch on the same recording day. Specifically, for each epoch, we compared PD_{target} , PD_{hand} , or PD_{plan} to the PD_{plan} of STD1, to analyze the learning-related changes in neuronal activity under the different reference frames. We examined the PD differences as a function of the nPD of the cell, grouped in bins of 45° .

Figure 6A shows the differences between PD_{plan} in STD1 and PDs late in learning. The differences for PD_{target} are shown in red, for PD_{hand} in black, and for PD_{plan} in green. Similarly, Figure 6B shows the difference between STD1 and STD2. The PD differences in Figure 6A demonstrates a double-peak pattern, which indicates that both co-FF and counter-FF cells showed consistent apparent shifts in PDs. This pattern was not observed early in learning (supplemental Fig. 3A, available at www.jneurosci.org as supplemental material) when only cells in the counter-FF range were shifted in the force-field direction. This reinforces the finding that co-FF firing decreased late in learning, and the counter-FF cells continued to show increased activity (as depicted in Fig. 4). These apparent PD shifts can be explained by the combined effect of firing rate modulations in co-FF and counter-FF cells as detailed in Figure 5B.

However, the PD_{plan} during learning (Fig. 6A, green line) was not different from STD1, suggesting that the relation between firing rates and force-vectors (directional tuning) do not change. Assuming the force-vector reflects the adapted motor plan (as illustrated in Fig. 5A, left panels), the observed changes in firing rates reflect the change in the motor plan (Kalaska et al., 1989). Note that the similarity between PD_{plan} in STD1 and learning suggests that a comparison between PD_{target} and PD_{plan} , both computed for the learning epoch, should result in the same double-peak pattern of PD shifts as shown in Figure 6A.

In summary, late in learning, when the hand moved almost straight to the target, PD_{hand} and PD_{target} were similar on average but different from PD_{plan} . This difference reflects the effect of learning, i.e., modification of the motor plan. Note that some cells showed changes in the PD_{plan} , but with no consistent relations to the FF direction, suggesting that they could originate from inaccuracies in PD estimations attributable to short stan-

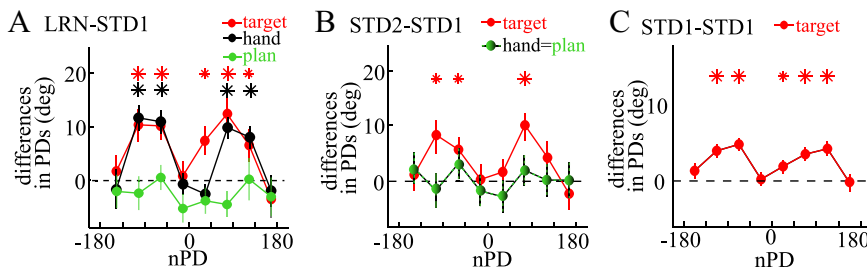


Figure 6. Directional tuning of cells under perturbation depends on the selected reference frame: neuronal data. The *x*-axis in all plots is the nPD of the cells, in bins of 45°, with learned target at 0. The *y*-axis in **A–C** shows differences between PDs late in learning. The title of each plot shows the compared epochs. **A**, The PD differences between learning and STD1 are positive, showing apparent shifts with force-field direction. The plots show that the differences depend on (1) the nPD of the cell and (2) the reference frame during learning (PD_{target} (red), PD_{hand} (black), or PD_{plan} (green)). **B**, PD_{target} (red), but not PD_{hand} and PD_{plan} (black–green), show in STD2 apparent shifts relative to STD1 with force-field direction in a similar pattern as in learning (**A**). **C**, Comparing PD_{target} and PD_{plan} in STD1 (red) reveals a similar pattern of apparent PD shifts. Note that the comparison was made between the same cells with the same activity for the same trials and differed only by the reference frame. $n = 264$. In all traces, error bars denote SEM. Large and small asterisks denote 1 and 5% confidence, respectively.




Figure 7. Control analyses for the relation between directional tuning, selected reference frames, and adaptation. **A**, The relation between apparent shifts and discrepancy between hand and target directions is demonstrated by separating STD1 trials to the learned target to those with large (gray) and small (black) overnight aftereffects. **B**, The observed (obs) firing rate during learning matches the expected (exp) activity from movements in the force-vector direction (black) and the observed activity when perturbation is removed (gray). $n = 264$. In all traces, error bars denote SEM.

standard periods and/or from ongoing variability in neuronal activity (Rokni et al., 2007).

Comparing the first and the second standard epochs (STD1 to STD2, Fig. 6B) showed similar but smaller differences between PD_{target} in STD2 relative to STD1. However, unlike the apparent PD shifts of PD_{hand} during learning (Fig. 6A, black line), in STD2, PD_{hand} were not significantly different from STD1, as shown by the black–green line of Figure 6B. The result suggests that the relation between firing rates and initial observed hand direction (or motor plan) did not change. Because PD_{hand} and PD_{plan} are assumed to be the same under standard conditions, the black–green line is effectively showing the same comparison as the green line in Figure 6A. Therefore, the difference between the PD_{hand} (or PD_{plan}) and PD_{target} in STD2 further strengthens our previous results and shows that the modified motor plan remains after force-field removal.

We noticed that the movements in the standard trials of the late days were significantly curved throughout the STD1 epoch. Therefore, we calculated the difference between PD_{target} and PD_{plan} for the very same cells and the same trials of STD1. Obviously, any differences between these PDs only reflect the difference of reference frame. Indeed, early in learning, there was no significant difference between the hand direction and the target direction and therefore no significant differences between the PDs (supplemental Fig. 3C, available at www.jneurosci.org as supplemental material). Late in adaptation, when the trajectories expressed overnight aftereffects and their initial direction deviated from the target direction, PD_{target} deviated from PD_{plan} . Cal-

culating the differences between PDs for all cells, the same double-peak pattern of apparent PD shifts in the co-FF and counter-FF ranges emerged (Fig. 6C).

To strengthen the conclusion that the apparent PD shifts were related to deviations of the motor plan rather than to changes in tuning properties of the cells, we performed three additional control analyses. First, trials to the learned target during STD1 in late days were divided into two groups by the size of the overnight aftereffects on movements to this target: one with initial directional deviations larger than the median and one with the smaller deviations. Thus, the calculation of PDs only differed in the activity and hand direction to the learned target. Figure 7A shows $PD_{\text{target,STD1}} - PD_{\text{plan,STD1}}$ for the two groups, with small (black) and large (gray) overnight aftereffects. The figure shows that the larger the deviations between the target and hand directions, the larger the difference between PD_{target} and PD_{plan} of the co-FF and counter-FF cells.

Second, for each perturbed movement, we estimated the direction of the force applied by the hand (force-vector). For each cell, we compared its observed firing rate in each of the perturbed movements to the one expected from the force-vector direction, according to the tuning curve of the cell on standard trials (Fig. 7B, black). Furthermore, we compared the observed activity with the activity in the first trial to the learned target, after removal of force-field (in STD2 epoch), in which the hand clearly showed the learning-related aftereffect (Fig. 7B, gray). In both comparisons, on average, we did not find significant differences, suggesting that the observed behavioral changes are likely to be sufficient to account for the neuronal effects.

Finally, we examined days in which no force-field learning took place. On these days, the monkeys performed the highly familiar default eight-target nonperturbed reaching movements, and no changes were expected in motor plans and neuronal activity. We compared the early trials with late trials for the same day and did not find any consistent differences between PDs (data not shown).

The results strongly support the conclusion that, during force-field adaptation, the subjects learned to push their arm in a new direction to compensate for the force-field, and the emerging PD shifts only reflect the difference in reference frames that we use to compute them. Namely, the PDs remain the same, and the population generates a new motor plan by the increases and decreases of the counter-FF and co-FF cells.

Washout and relearning

Days 5 and 6 were used to test for washout and recall of the learned task. The behavioral effects of washout (day 5) and relearning (day 6) evolved faster than the learning, which spanned >200 trials to the learned target. Therefore, we examined the related activity changes in the previously defined two populations of cells with fast and slow dynamics.

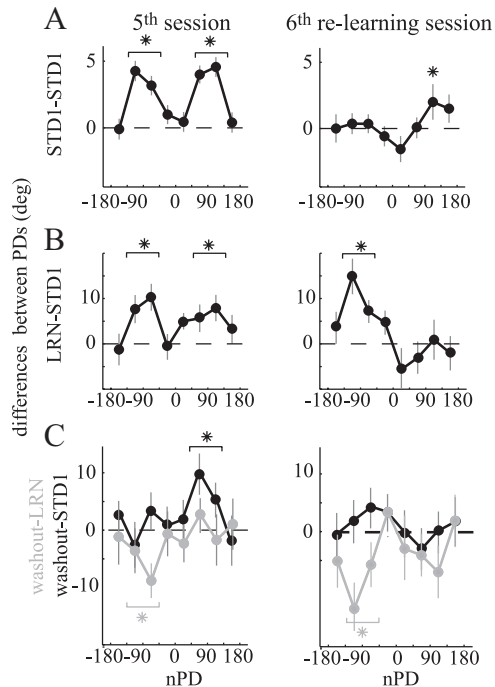


Figure 8. Apparent PD shifts are reversed during washout and relearned during relearning; this effect is seen only in counter-FF cells, reflecting only the fast process. The analysis shows day 5 (left) and day 6 (right). The y-axes show differences in PDs between reference frames, as a function of the nPD of the cell. The solid black lines show comparisons of the PD_{plan} in STD1 to PD_{target} in STD1 (A), learning (B), and late in washout (C). The differences between PD_{target} late in washout and PD_{target} in learning are shown as gray lines. $n = 130$ on day 5 and $n = 126$ on day 6. In all traces, asterisks show 1% confidence, and error bars denote SEM. Square brackets indicate that significance was calculated on cells from two adjacent bins together, in a total range of 90° .

Figure 8 depicts in black the PD_{target} in standard (A), learning (B), and late in washout (C) compared with PD_{plan} in standard. Figure 8C specifically depicts (gray) the effect of washout by computing the difference between PD_{target} late in washout and PD_{target} in learning.

Figure 8, A and B (left), first shows the typical double-peak pattern of apparent PD shifts on day 5, as shown in Figure 6. Then, at the end of learning on day 5, we applied a long washout epoch of at least 360 standard trials in all directions (45 to the learned target). Comparing the activity of co-FF and counter-FF cells late in washout to their activity during the preceding learning epoch showed that counter-FF cells decreased their firing rates ($p < 0.01$, data not shown), whereas co-FF cells maintained their learning-related activity and did not change significantly. Figure 8C (left, gray) shows that, as a result, PDs of counter-FF cells “shifted” counter to the force-field direction. Because the co-FF cells did not change their firing rates during washout, they kept the same PDs as in the learning period (gray) and still differed from the PDs in STD1 (black).

Note that, late in washout, the trajectories still showed an aftereffect. This incomplete washout may have resulted from the learning-related activity of co-FF cells, which was maintained at the end of the washout period, whereas the learning effects in counter-FF cells washed out relatively fast.

Next, we tested relearning of the perturbation on day 6, after the washout on day 5 and the STD1 epoch of day 6.

Trajectories to the learned target in STD1 epoch of day 6 were almost straight (Fig. 2A, C), and therefore there were almost no differences between PD_{target} and PD_{plan} in STD1 epoch (Fig. 8A,

right). Relearning was expressed only by an increased activity of the counter-FF cells. As a result, their PD_{target} differed from their PD_{plan} in standard trials (Fig. 8B, right), similar to the PD differences observed in the early days of learning (supplemental Fig. 3A, available at www.jneurosci.org as supplemental material). These differences in PDs were reversed during the successive washout epoch (Fig. 8C, right, gray), resulting in PDs that were similar to those in standard (Fig. 8C, right, black). Note that, late in this washout, the hand moved in the direction of the target (no significant aftereffects), suggesting that the lack of differences between PDs in washout and standard are in accordance with the complete washout.

Altogether, these results hint that the stability of learning is dependent on the relative fraction of fast and slow neuronal changes in the learning process. The incomplete washout on day 5 was arguably a reflection of the slow dynamics of co-FF cells that maintained their learning-related changes until late in washout. These changes were probably gradually washed out overnight, because on day 6 the overnight aftereffects were much smaller. Thus, the fast relearning was mediated mainly by changes of the “fast” counter-FF cells, whereas the changes of co-FF cells were most likely too slow to accumulate. Hence, after “fast relearning” on day 6, the washout could be completed because it resulted mostly from the rapidly reversed changes of counter-FF cells.

Population-vector analysis

Population-vectors are thought to relate neuronal activity in the motor cortex to the direction of arm movements (Georgopoulos et al., 1986). We computed population-vectors in movements to the learned target with and without force-field as an additional way to explore possible changes in the motor plan during force-field adaptation.

For each day and cell, we computed the population-vectors based on PD_{plan} in STD1 using the optimization linear estimator (Salinas and Abbott, 1994). Figure 9A presents an example of population-vectors for trials to the learned target along 2 learning weeks comprising 5 d. Both had the same learned target (at 0° , horizontal line) but opposing force-field directions; week 1 with a counterclockwise (CCW) field and week 2 with a clockwise (CW) field. Each dot reflects the direction of a population-vector for a single trial. Dots appear in the order of trials along days, where blue/cyan dots are population-vectors of STD1 trials and red/magenta are population-vectors of LRN trials, for week 1 and week 2, respectively. The figure shows that, as learning progressed, population-vectors clearly deviated in different directions from the learned target in the 2 weeks, each counter to its force-field direction (compensating the perturbation). Note that population-vector deviations were apparent not only in LRN trials but also in STD1 trials (blue/cyan), in close concurrence with the overnight aftereffects (Fig. 2C).

Averaging over the 7 learning weeks (Fig. 9B) shows the gradual divergence of population-vectors from the learned target. Because population-vectors were computed by the PD_{plan} in the STD1 epoch, we expected them to point in the initial hand direction during unperturbed movements. Indeed, for movements in STD1 and STD2, the population-vectors and initial hand direction were not significantly different (data not shown, t test, $p > 0.5$) throughout the learning week. During learning trials, when the force-field pushed the hand, population-vector (Fig. 9B, red line) and initial hand direction (Fig. 9B, black line, taken from Fig. 2B for illustration purposes only) pointed in different directions, with a difference that remained at approximately the same level of 40° , throughout the week (ANOVA, $p > 0.1$), whereas the

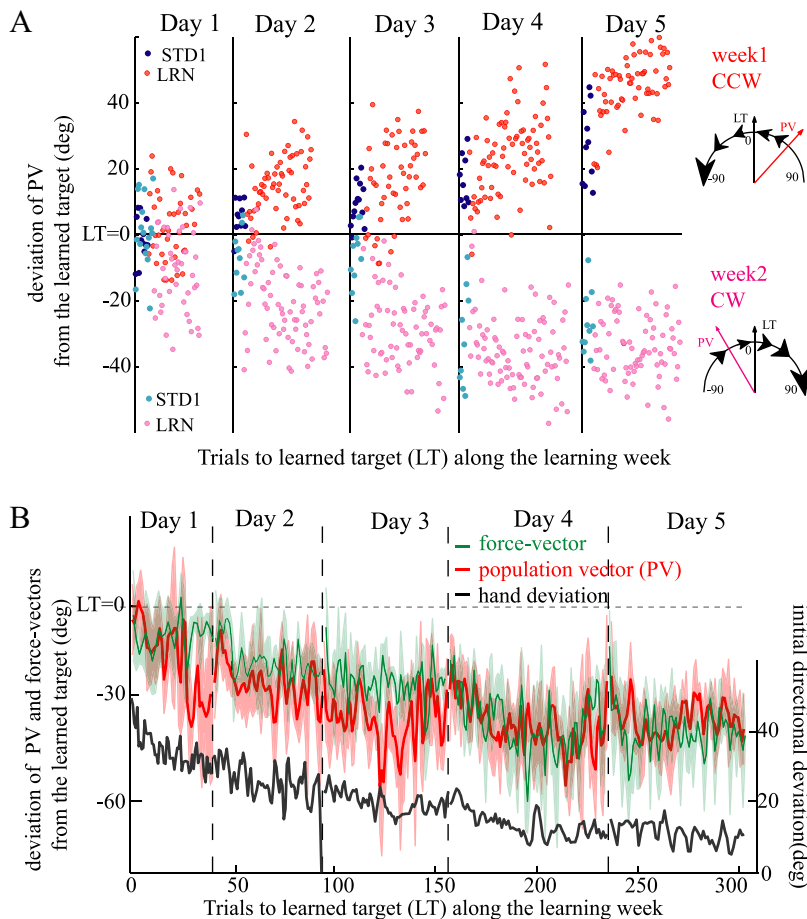


Figure 9. Adaptation of the population-vectors (PVs). The PVs during movements to the learned target (LT) gradually deviate in a direction against the force-field, with the force-vector direction. **A**, Population-vector estimations for single trials to learned target (dots) during 2 learning weeks that had the same learned target (at 0°) but opposing force-fields. STD1 trials are denoted in blue (week 1) and cyan (week 2) and learning trials in red (week 1) and magenta (week 2). The opposing force-fields induced opposing deviations of population-vectors from the learned target, each against its force-field direction. Note that, on days 4 and 5, population-vectors in STD1 are deviated and coincide with the overnight aftereffects that are described in Figure 2C. CW, Clockwise; CCW, counterclockwise. **B**, Population-vectors (red) and force-vectors (green) averaged over 7 weeks show a gradual and similar deviation from the learned target, counter to force-field direction. The average learning curve (copied from Fig. 2B) is shown (black line) with its y -axis on the right, for comparison of temporal evolution. Note that, when the population-vectors and force-vector point to the learned target (early in learning), the hand (black) deviates away from it. $n = 704$. Shaded areas denote SEM.

population-vector and force-vector (green line) directions were the same ($p > 0.15$).

Because the initial hand direction deviated less from the target direction (straighter movements), population-vectors deviated more. Namely, under the relatively constant force-field during learning, there was a tight relation between the changes in motor plan as reflected by changes of population-vectors and the observed hand movement.

Premovement neuronal activity

In each trial, the monkeys were informed on the direction of the upcoming movements 850–1350 ms before the go signal, and the monkey was required to hold the robotic arm at the central circle during this time window. Studies have suggested that neuronal activity in M1 during such delay period reflected attributes of the intended reaching movement (Wise et al., 1998) and could show learning-related modulation (Paz et al., 2003; Mandelblat-Cerf et al., 2009). Bizzi's group (Padoa-Schioppa et al., 2004; Xiao et al.,

2006) showed during force-field adaptation gradual shifts in PDs in premotor areas during such a delay period.

We examined the premovement delay period by comparing PD_{target} in several epochs with PD_{plan} in STD1 (i.e., $PD_{\text{target,epoch1}} - PD_{\text{plan,STD1}}$, where epoch 1 presents STD1, LRN, or STD2) during late days of learning. We performed the analysis separately for two premovement segments: one that immediately followed the target onset and the other a period of 500 ms before go signal. Although neuronal activity after target onset did not show significant differences in PDs (data not shown), the activity before the go signal (Fig. 10A) showed the double-peak pattern of apparent PD shifts of the co-FF and counter-FF cells in each of the epochs (STD1, LRN, STD2) compared with STD1. The population-vector of this activity (Fig. 10B) diverged from the learned target in a similar (but smaller and variable) way as the population-vectors that were estimated for the movement-related activity (Fig. 9B). These apparent PD shifts may reflect a change of the motor plan. The changes of population-vectors during the delay period support the notion that the motor cortex can plan a movement ahead of its execution.

PDs analyses in local learning of kinematics task

A previous experiment in our laboratory investigated neuronal activity during local learning of visuomotor rotation (Paz et al., 2003), which introduces a constant rotation between the hand and cursor directions without applying any force on the hand. Therefore, monkeys learned to move their hand in a new direction to bring the cursor to the learned target. The visuomotor rotation task requires a change of the kinematics parameters, whereas force-field perturbation requires a change of the dynamics parameters to maintain similar kinematics. In previous studies by Paz et al. and others (Ghahramani et al., 1996; Krakauer et al., 2000), only the learned target was repetitively presented during adaptation, resulting in a fast learning and performance that reached a plateau within dozens of trials on a single day.

In local adaptation, one can only record neuronal activity to the single learned target. Paz et al. showed that movement-related activity during adaptation was as expected from the firing rates of cells on standard trials when the hand moved in the new adapted direction (Paz et al., 2003, their Fig. 5). In addition, it was shown that the behavioral generalization after such adaptation was limited. Based on these two findings, we assume that we can estimate the directional tuning of cells during local visuomotor adaptation by using firing rates to the seven nonlearned targets from the standard trials with the activity to the learned target from the

adaptation period and test whether movement-related activity shows the same pattern of double-peak apparent PD shifts and the dependency on reference frame (hand direction vs target direction). The visuomotor rotation task is highly useful to examine the effect of reference frame because the adapted motor plan is clearly observed by the updated hand direction, which moves in the rotated angle. Therefore, unlike force-field adaptation, we can directly examine PD- $_{plan}$ (and PD- $_{hand}$, which is the same) during learning.

Furthermore, note that, in this short-term learning task, the prelearning movements during standard epoch were straight toward the learned target, and therefore the PDs computed in different reference frames (plan, target, and hand) for these movements were all practically the same.

Figure 11 shows that the differences between PD- $_{target}$ in learning and PD- $_{plan}$ in standard as a function of the cells PD (red solid lines) are very similar to the double-peak pattern shown in Figure 6 (red curves); namely, apparent PD shifts are evident for cells with PD at a distance of approximately $\pm 90^\circ$ from the learned target.

Interestingly, as in force-field adaptation, we found no systematic differences in PDs when PD- $_{plan}$ during visuomotor rotation learning were compared with PD- $_{plan}$ in standard (Fig. 11, black line). This result is in line with Paz et al. (2003), showing that movement-related firing rates during adaptation followed the hand direction. Thus, all cells can contribute to the new hand-movement direction, with increased activity in cells with PD in the new movement direction and decreased activity in cells in the opposite direction. However, the activity of cells with PDs around $\pm 90^\circ$ from the learned target direction changed the most. This is not surprising given the cosine-like directional tuning. As shown in Figure 5B, for these cells, the target is located at the range of maximal slope of the cosine; hence, the sensitivity to changes in the direction of the motor plan is maximal. Cells with PDs near the learned target or 180° away showed little change in activity because of the small slope. These changes were probably smaller than background variability (Lee et al., 1998; Stein et al., 2005; Rokni et al., 2007; Faisal et al., 2008) and escaped detection. Therefore, although all cells could be involved in generation of the new movement, we only observed the larger modulations in the two PD ranges around $\pm 90^\circ$ from the learned target.

To conclude, the results strongly suggest that the differences in PDs that we observe in adaptation to force-field and in visuomotor rotation do not reflect PD shifts. Instead, they reflect a discrepancy between the target location and the motor plan direction. Given this discrepancy, new weights are assigned to the cells (changes in firing rates) and the population signal is adjusted accordingly.

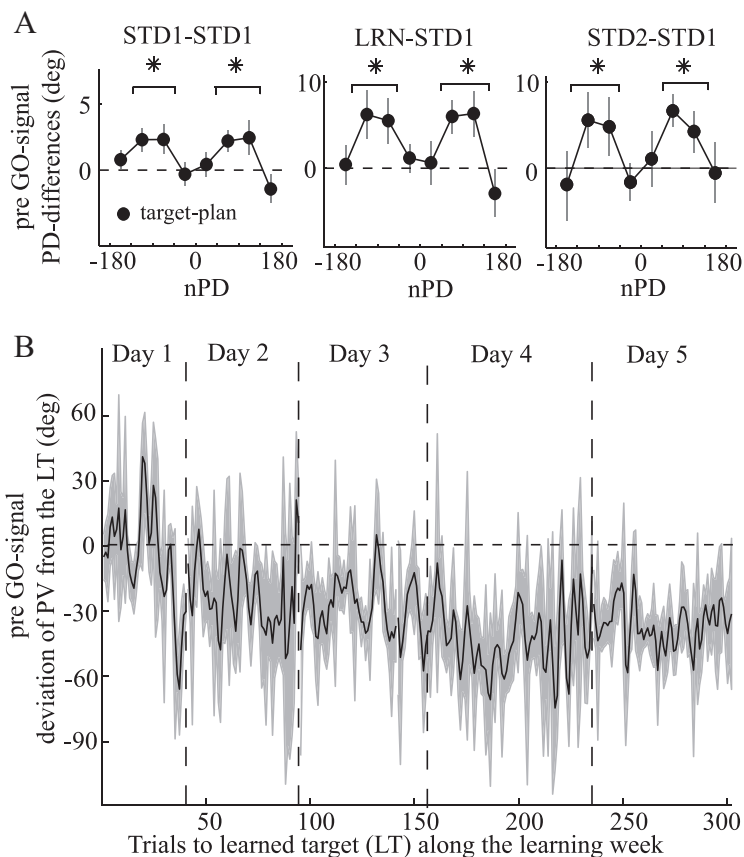


Figure 10. The neuronal changes are apparent before movement onset. **A**, The figure depicts PD differences computed for neuronal activity before the go signal (late learning days). The comparisons of PD- $_{plan}$ in STD1 to PD- $_{target}$ in STD1 (left), LRN (middle), and STD2 (right) show the double-peak pattern of apparent PD shifts in all three cases. Asterisks denote 1% confidence, and square brackets indicate that significance was calculated on cells from two adjacent bins together, in a total range of 90° . **B**, Population-vectors (PV) estimated for activity before the go signal, averaged over 7 weeks, show a gradual deviation from the LT counter to force-field direction. $n = 486$. Shaded areas and error bars denote SEM.

Generalization: behavioral and neuronal data

To test generalization, we studied the effect on movements to the seven nonlearned targets (without force-field) that were executed throughout the learning epochs in the late days of learning.

The results are depicted in Figure 12 describing the initial directional deviations of trajectories (“TRJ-aftereffects”) and the corresponding deviations of the population-vectors from the target (“PV-deviations”) using optimal linear estimator; Salinas and Abbott, 1994) for unperturbed trials that immediately followed force-field trials. Because trials appeared in a random sequence during the learning epoch, each target occasionally followed a force-field trial. Figure 12 presents the averaged aftereffect (black) and population-vector deviation (gray) for each target as a function of its angular distance from the learned target. The figure shows, as expected, that the strongest aftereffects were observed in the learned target direction (-37°) and the corresponding deviations of population-vector (-30°). Interestingly, aftereffects on trajectories and deviations of population-vectors were also observed in other movement directions, reflecting generalization. Generalization was mostly expressed in movements to the -45° target (target adjacent to the learned target in a direction counter to the force-field) with aftereffects of -9° (t test, $p < 0.01$). Similarly, the corresponding population-vectors deviated from the target direction, with directional errors of approximately -7° . Significant but considerably smaller effects ($p < 0.01$) were evident in movements to $+45^\circ$ target ($\sim -3^\circ$). Move-

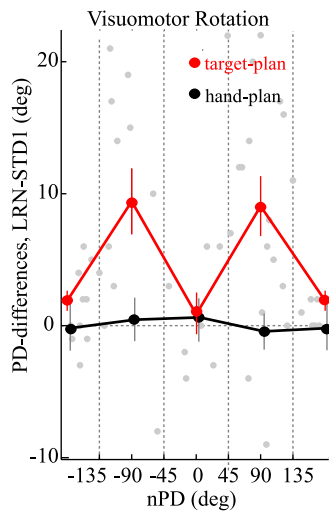


Figure 11. PD_{target} and not PD_{hand} show apparent PD shifts in a kinematic task (visuomotor rotation). Gray dots indicate the differences for single cells between PD_{target} in the learning epoch compared with the PD in standard epoch. The red line indicates the averaged PD differences across cells, binned in 90° . For purposes of illustration, the average of subgroup with $|PD| > 135$ is depicted twice at both ends of the x-axis. Comparing PD in standard with PD_{hand} in learning (black) did not show any systematic differences. $n = 52$. Error bars denote SEM.

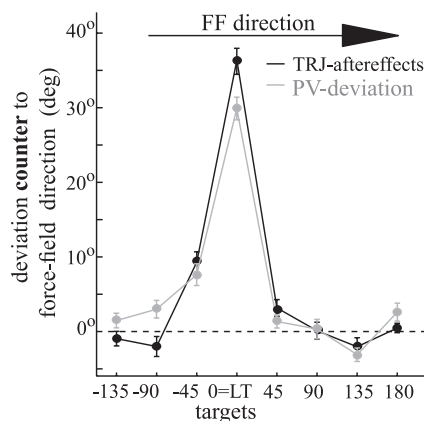


Figure 12. Generalization of local adaptation to force-field is affected by the motor plan. Generalization is measured by initial directional deviations of trajectories (TRJ-aftereffects, black line) and the corresponding population-vectors (PV-deviation, gray line). Note that trajectories aftereffects to the learned target were only measured in STD2 after the first epoch of learning trials, but all other directions were measured during learning. The figure shows the trajectories aftereffects and population-vector deviations from each of the targets around the learned target, taken at late stages of adaptation. Note the elevated effect to target -45° . $n = 264$. Error bars denote SEM.

ments to other targets did not show any systematic effects and on average did not differ from zero. Note that we did not find significant daily changes in firing rates during movement to any of the nonlearned targets (Fig. 3), although some generalization effects are evident, probably attributable to a smaller overall effect relative to changes in movements to the learned target.

In conclusion, as expected from previous studies, generalization was narrow (limited to $\pm 45^\circ$). Interestingly, we found that the effect of generalization was increased in the direction to which the hand was pushing to compensate for the perturbing force-field. Namely, generalization followed the motor plan direction (counter to force-field direction) rather than the target location or the observed hand direction (with force-field direction).

Discussion

This paper presents a new approach to study the neuronal basis of adaptation, using a modified version of the curl force-field that leads to long-term local adaptation. The design of experiment allowed for monitoring the local effects and at the same time measuring the directional tuning of single cells and the population signals. We found that adaptation involved processes of different (fast and slow) timescales that were expressed in modulations of firing rates. Analysis of these changes revealed that they do not reflect consistent modifications of PDs of single cells (Li et al., 2001). Rather, they reveal the change of intention and motor plan without changing properties of single cells.

Specifically, we found that (1) during local, long-term adaptation the activity of the neuronal population was modified and correlated with improvement in performance from day to day. (2) Modulations in single-cell activity were primarily expressed in two subgroups of cells with fast and slow timescales. (3) In both timescales, the cells showed consistent apparent shifts of their PDs with force-field direction. However, (4) these calculated “PD shifts” may be explained by the discrepancy between the motor plan direction and target location rather than intrinsic changes in the tuning of the cells. (5) The generalization pattern, with similar neuronal and behavioral signatures, reflected the effect of the motor plan direction.

Generally, our findings agree and extend predictions of psychophysical studies in humans and monkeys showing that gradual adaptation to force-field is expressed by reduction of trajectory curvature under the perturbation and by increased aftereffects (opposite curvature) when perturbation is removed (Shadmehr and Mussa-Ivaldi, 1994; Li et al., 2001; Arce et al., 2009). As in previous studies, we also found day-to-day retention of learning (Shadmehr and Brashers-Krug, 1997) with aftereffects that were maintained overnight. These behavioral signatures are commonly considered to reflect modifications of internal models (Kawato and Wolpert, 1998; Kawato, 1999; Hwang et al., 2006).

Dynamics of neuronal modulations reflect fast and slow adaptation processes

The modification of the internal model allows for better prediction of the amplitude and direction of the perturbation. We looked for neuronal changes that could reflect a change of an intention of action, i.e., a change in motor plan. It has been suggested that these changes are performed by multiple timescale processes, which theoretically could be advantageous for flexibility in synaptic plasticity (Fusi et al., 2007). It was also suggested that fast and slow processes could account for several phenomena in short-term motor adaptation (Smith et al., 2006) when the slow process is more robust. Namely, the slower the process, the longer the memory; the fast process that is expressed in behavior early in learning leaves only short-term traces, and the slow process that emerges only later in learning leaves longer traces. We studied long-term rather than trial-by-trial modulations. Therefore, the “fast” and “slow” dynamics we report are at longer timescales. Nevertheless, they may serve similar functions as proposed by Smith et al. Interestingly, our data also support the notion that the slow process that evolves much slower than the fast also decays more slowly (Figs. 3, 4).

The fast and slow changes that we see in motor cortex may depend on activity in various brain structures in addition to the motor cortex. A major candidate is the cerebellum (Medina et al., 2000, 2001), in which it was suggested that error rapidly teaches the cerebellar cortex whereas the cerebellar cortex slowly teaches

the cerebellar nuclei. Karni et al. (1998) demonstrated (in a functional magnetic resonance imaging study) fast and slow modulations of M1 activity and related the fast stage to the actual execution of movements and the slower to the formation of a “unitary motor plan and a rather ‘automatic’ movement.”

The faster process in our study was expressed by an increased firing rate in a selected population of cells with nPDs counter to the force-field direction. This increase can account for most of the required compensation of the force-field, pushing the hand against its direction. As learning progressed, the slow process gradually accumulated, as expressed by the decreased firing rate in cells showing nPDs with the force-field direction (co-FF cells). The contribution of these cells to the correction of direction was small, yet the decreased firing rate of these cells apparently acted together with the counter-FF cells to overcome the perturbation. Assuming that the brain attempts to minimize some cost function during adaptation (Todorov and Jordan, 2002; Todorov, 2004), it is likely that these changes, which reduce the total muscle activation, are cost effective and facilitate optimization of the motor plan.

Note that, when force-field is applied on several directions (“global learning” as in the study by Bizzi’s group) counter-FF cells are also the co-FF cells for targets in the opposite direction. Thus, it is more accurate to talk about a “counter-FF process” and “co-FF process” that act together to optimize the movements under force-field. Therefore, our results predict that, during global learning, a directionally tuned cell should increase its firing rate in a “fast process” and decrease its firing rate in a “slow process” depending on the target location.

Neuronal modulations reflect the adapted motor plan

Our primary findings show that, late in learning, the fast and slow processes generated a profile of firing rate modulations that depend on the PD of the cell, similar to the cosine-like pattern observed in our previous study (Arce et al., 2010a). However, then, we could not measure the directional tuning during learning, which we did here. Tuning of cells during learning was computed in three reference frames: target location, observed initial hand-movement direction, and the computed force-vector that the hand applied. For nonlearned targets, these three frames were practically identical. However, it changed for the learned target: when the perturbation was first introduced, the force-vector pointed to the target, whereas the hand deviated with force-field direction. Late in learning, the force-vector pointed counter to force-field direction, and its vectorial summation with the force-field generated hand movements toward the target (Fig. 5A). Comparisons of directional tuning and the resulting PDs in the three reference frames indicated that, during learning, neuronal activity was always consistent with the force-vector reference frame (Kalaska et al., 1989) but not with target or hand reference frames (Fig. 6A,B) in which the PDs seemed to shift with force-field direction as shown previously (Li et al., 2001; Arce et al., 2010b). When perturbation was removed, we observed apparent PD shifts on the population level in the opposite direction, as found in these previous studies. However, these studies also reported specific memory effects in selected groups of cells. We could not test this feature because in our task design of long-term learning (without daily washout), prelearning and postlearning standard movements were not recorded on the same day.

Our results provide a new explanation for the PD shift phenomena. Because PDs do not change in the force-vector reference frame, the apparent shifts may express firing rate modulations in the selected reference frame, whereas PDs do not shift. When the

force-field pushes the hand, subjects modify the motor plan, and the same cells with same PDs are recruited to perform this new plan.

This conclusion is supported by the analysis of kinematics adaptation (visuomotor rotation task) (Fig. 11). In that task, there was an instructed discrepancy between the hand and target directions. Thus, the hand moved unperturbed in the planned direction, whereas the cursor moved to the target. The analysis showed that learning-related changes were reflected by apparent PD shifts in the target reference frame but not in the hand/plan reference frame. These results clearly suggest that kinematics adaptation, like dynamics adaptation, can evolve without PD shifts.

Although PDs did not shift systematically, the population-vectors showed a day-to-day gradual shift counter to force-field direction, pointing in the same direction as the estimated force-vector (the assumed motor plan direction) and tightly related to improvement in performance. Therefore, the findings further support the notion that the neuronal tuning always reflects the motor plan.

Late in learning, the adapted plan was evident before the go signal (Fig. 10) (in line with Padoa-Schioppa et al., 2002). This adapted activity in motor cortex may serve as the basis for observations showing that the brain can predict the perturbation and compute in advance how to compensate to move straight to the target. The new computation could be performed by either “re-aiming” to a new direction (as suggested previously for rotation by Jarosiewicz et al., 2008) or aiming at the target but computing another force-vector direction. Our data do not distinguish between these strategies; although they are conceptually different, they result in the same directional motor plan that produces the same profile of increases and decreases of activity in the motor cortex.

The generalization pattern of learning described in Figure 12 is intriguing. During adaptation, the hand was pushed in the force-field direction, and the actual hand movements were curved in this direction. However, generalization showed an increased effect in the direction of the adapted motor plan (counter to force-field direction) rather than in the perturbed hand direction. Therefore, we suggest that the adapted motor plan influences responses to nonlearned targets, inducing generalization in its direction.

To conclude, we tracked neuronal activity during long-term sensorimotor adaptation and distinguished between two different timescales of changes in single-cell firing rates. These changes occurred without systematically changing the preferred directions of the cells but rather reflected the change in the motor plan direction.

References

- Arce F, Novick I, Shahar M, Link Y, Ghez C, Vaadia E (2009) Differences in context and feedback result in different trajectories and adaptation strategies in reaching. *PLoS One* 4:e4214.
- Arce F, Novick I, Mandelblat-Cerf Y, Israel Z, Ghez C, Vaadia E (2010a) Combined adaptiveness of specific motor-cortical ensembles underlies learning. *J Neurosci* 30:5415–5425.
- Arce F, Novick I, Mandelblat-Cerf Y, Vaadia E (2010b) Neuronal correlates of memory formation in motor cortex after adaptation to force field. *J Neurosci* 30:9189–9198.
- Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414:446–449.
- Caithness G, Osu R, Bays P, Chase H, Klassen J, Kawato M, Wolpert DM, Flanagan JR (2004) Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J Neurosci* 24:8662–8671.
- Chen LL, Wise SP (1995) Supplementary eye field contrasted with the fron-

- tal eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 73:1122–1134.
- Cheng EJ, Scott SH (2000) Morphometry of *Macaca mulatta* forelimb. I. Shoulder and elbow muscles and segment inertial parameters. *J Morphol* 245:206–224.
- Donchin O, Francis JT, Shadmehr R (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J Neurosci* 23:9032–9045.
- Evarts EV (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. *J Neurophysiol* 31:14–27.
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303.
- Fusi S, Asaad WF, Miller EK, Wang XJ (2007) A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple time-scales. *Neuron* 54:319–333.
- Gandolfo F, Li C, Benda BJ, Schioppa CP, Bizzi E (2000) Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc Natl Acad Sci U S A* 97:2259–2263.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537.
- Georgopoulos AP, Caminiti R, Kalaska JF, Massey JT (1983) Spatial coding of movement: a hypothesis concerning the coding of movement direction by motor cortical populations. *Exp Brain Res [Suppl]* 7:327–336.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
- Ghahramani Z, Wolpert DM, Jordan MI (1996) Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci* 16:7085–7096.
- Graham KM, Scott SH (2003) Morphometry of *Macaca mulatta* forelimb. III. Moment arm of shoulder and elbow muscles. *J Morphol* 255:301–314.
- Hwang EJ, Smith MA, Shadmehr R (2006) Dissociable effects of the implicit and explicit memory systems on learning control of reaching. *Exp Brain Res* 173:425–437.
- Jarosiewicz B, Chase SM, Fraser GW, Velliste M, Kass RE, Schwartz AB (2008) Functional network reorganization during learning in a brain-computer interface paradigm. *Proc Natl Acad Sci U S A* 105:19486–19491.
- Kalaska JF, Cohen DA, Hyde ML, Prud'homme M (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *J Neurosci* 9:2080–2102.
- Karni A, Meyer G, Rey-Hipolito C, Jezzard P, Adams MM, Turner R, Ungerleider LG (1998) The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex. *Proc Natl Acad Sci U S A* 95:861–868.
- Kawato M (1999) Internal models for motor control and trajectory planning. *Curr Opin Neurobiol* 9:718–727.
- Kawato M, Wolpert D (1998) Internal models for motor control. *Novartis Found Symp* 218:291–304; discussion 304–307.
- Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2:1026–1031.
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20:8916–8924.
- Lackner JR, Dizio P (1994) Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* 72:299–313.
- Lee D, Port NL, Kruse W, Georgopoulos AP (1998) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18:1161–1170.
- Li CS, Padoa-Schioppa C, 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.
- Mandelblat-Cerf Y, Paz R, Vaadia E (2009) Trial-to-trial variability of single cells in motor cortices is dynamically modified during visuomotor adaptation. *J Neurosci* 29:15053–15062.
- Medina JF, Noes WL, Ohyama T, Mauk MD (2000) Mechanisms of cerebellar learning suggested by eyelid conditioning. *Curr Opin Neurobiol* 10:717–724.
- Medina JF, Garcia KS, Mauk MD (2001) A mechanism for savings in the cerebellum. *J Neurosci* 21:4081–4089.
- Mitz AR, Godschalk M, Wise SP (1991) Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. *J Neurosci* 11:1855–1872.
- Padoa-Schioppa C, Li CS, Bizzi E (2002) Neuronal correlates of kinematics-to-dynamics transformation in the supplementary motor area. *Neuron* 36:751–765.
- Padoa-Schioppa C, Li CS, 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, Boraud T, Natan C, Bergman H, Vaadia E (2003) Preparatory activity in motor cortex reflects learning of local visuomotor skills. *Nat Neurosci* 6:882–890.
- Paz R, Nathan C, Boraud T, Bergman H, Vaadia E (2005) Acquisition and generalization of visuomotor transformations by nonhuman primates. *Exp Brain Res* 161:209–219.
- Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representation. *Neuron* 54:653–666.
- Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. *J Comput Neurosci* 1:89–107.
- Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. *J Neurosci* 17:409–419.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.
- Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4:e179.
- Stein RB, Gossen ER, Jones KE (2005) Neuronal variability: noise or part of the signal? *Nat Rev Neurosci* 6:389–397.
- Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7:907–915.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- Wise SP, di Pellegrino G, Boussaoud D (1996) The premotor cortex and nonstandard sensorimotor mapping. *Can J Physiol Pharmacol* 74:469–482.
- Wise SP, Moody SL, Blomstrom KJ, Mitz AR (1998) Changes in motor cortical activity during visuomotor adaptation. *Exp Brain Res* 121:285–299.
- Xiao J, Padoa-Schioppa C, Bizzi E (2006) Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. *Exp Brain Res* 168:106–119.
- Yamashita Y, Tani J (2008) Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS Comput Biol* 4:e1000220.
- Zach N, Inbar D, Grinvald Y, Bergman H, Vaadia E (2008) Emergence of novel representations in primary motor cortex and premotor neurons during associative learning. *J Neurosci* 28:9545–9556.