

Trial-to-Trial Variability of Single Cells in Motor Cortices Is Dynamically Modified during Visuomotor Adaptation

Yael Mandelblat-Cerf,^{1,4} Rony Paz,³ and Eilon Vaadia^{1,2,4}

¹Department of Medical Neurobiology, Institute for Medical Research Israel-Canada, Faculty of Medicine, and ²The Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem 91904, Israel, ³Department of Neurobiology, The Weizmann Institute of Science, Rehovot 76100, Israel, and ⁴The Edmond and Lily Safra Center for Brain Sciences, Hebrew University, Jerusalem 91904, Israel

Neurons in all brain areas exhibit variability in their spiking activity. Although part of this variability can be considered as noise that is detrimental to information processing, recent findings indicate that variability can also be beneficial. In particular, it was suggested that variability in the motor system allows for exploration of possible motor states and therefore can facilitate learning and adaptation to new environments. Here, we provide evidence to support this idea by analyzing the variability of neurons in the primary motor cortex (M1) and in the supplementary motor area (SMA-proper) of monkeys adapting to new rotational visuomotor tasks. We found that trial-to-trial variability increased during learning and exhibited four main characteristics: (1) modulation occurred preferentially during a delay period when the target of movement was already known, but before movement onset; (2) variability returned to its initial levels toward the end of learning; (3) the increase in variability was more apparent in cells with preferred movement directions close to those experienced during learning; and (4) the increase in variability emerged at early phases of learning in the SMA, whereas in M1 behavior reached plateau levels of performance. These results are highly consistent with previous findings that showed similar trends in variability across a population of neurons. Together, the results strengthen the idea that single-cell variability can be much more than mere noise and may be an integral part of the underlying mechanism of sensorimotor learning.

Introduction

Neurons throughout the brain exhibit a high degree of variability in their spiking activity even in seemingly constant task conditions (Bach and Kruger, 1986; Vogels et al., 1989; Arieli et al., 1996a,b; Lee et al., 1998). This variability is probably the result of intrinsic noise from molecular, cellular, and synaptic processes as well as from extrinsic sources such as sensory inputs and muscle activity (Vogels et al., 1989; Arieli et al., 1996a; Stein et al., 2005; Bialek and Setayeshgar, 2008; Faisal et al., 2008). In general, variability is a confounding factor for information transmission and has often been referred to as “noise.” From this standpoint, variability is detrimental to the message carried by the neurons, and the brain should find ways to overcome it to extract the original message (Bialek et al., 1991; Shallden and Newsome, 1998). However, recently it was suggested that variability can also be an important part of the signal itself (Stein et al., 2005; Faisal et al., 2008; Maimon and Assad, 2009).

In the motor system, variability is apparent during movement and force production and correlated with it (Jones et al., 2002;

Hamilton et al., 2004). Recently, a number of specific roles have been attributed to this motor variability. First, as an organizing principle, the motor system may attempt to minimize overall motor and sensory noise (Harris and Wolpert, 1998; van Beers et al., 2004; Todorov, 2005). Second, neuronal variability might be a signature of central mechanisms of movement planning. Variability has been observed before movement onset (MO) when there was no apparent increase in peripheral sensory and motor noise (Churchland et al., 2006a), gradually decreased along the trial, reached its minimum before MO, and was correlated with reaction times (Churchland et al., 2006b).

Computational and modeling studies have put forward the intriguing idea that neuronal variability can be important for learning in general (Kirkpatrick et al., 1983; Fusi, 2002) and for motor learning in particular (Rokni et al., 2007; Fiets et al., 2007; Faisal et al., 2008). In other words, variability enables dynamic representations and continuous exploration of possible motor states and appropriate neuronal configurations that can lead to the desired state by trial and error. Recent research on bird song learning (Kao et al., 2008) and purposeful reaching movements (Paz et al., 2005; Zacksenhouse et al., 2007) lends weight to this hypothesis. However, the direct trial-to-trial variability of neuronal activity as a mechanism that underlies motor learning has received little attention.

To explore this possibility, we analyzed data recorded from M1 and supplementary motor area (SMA) while monkeys adapted daily to rotational visuomotor transformations (Paz et al., 2003). In both humans (Krakauer et al., 2000) and monkeys (Paz et al., 2003), adaptation to such transformations takes sev-

Received June 22, 2009; revised Oct. 12, 2009; accepted Oct. 14, 2009.

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

The authors declare no competing financial interests.

Correspondence should be addressed to Yael Mandelblat-Cerf at the above address. E-mail: yaelma@ekmd.huji.ac.il.

DOI:10.1523/JNEUROSCI.3011-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2915053-10\$15.00/0

eral trials until performance stabilizes. We could therefore compare variability during different stages of learning in both cortical areas. We show that trial-to-trial variability before MO was indeed modulated during learning. Typically, it first increased and later decreased back to its initial levels. The increase was temporally correlated with the progress of behavior and emerged first in SMA and only later in M1, in accordance with the learning hierarchy.

Materials and Methods

Animals, recordings, and behavioral task

The experimental setup, behavioral paradigm, and recording procedures were the same as described in previous studies (Paz et al., 2003, 2005). Briefly, two female rhesus monkeys (*Macaca mulatta*; ~4.5 kg) were trained to perform a center-out task and then were implanted with a recording chamber (27 × 27 mm) above both the right and left hemispheres under anesthesia and aseptic conditions. Animal care and surgical procedures complied with the National Institutes of Health *Guidelines for the Care and Use of Laboratory Animals* (1996) and with guidelines supervised by the Institutional Committee for Animal Care and Use at the Hebrew University.

This study uses data from recordings (with eight microelectrodes in each hemisphere) of single-unit activity in the primary motor cortex (M1) and SMAs (location mapped by magnetic resonance imaging and microstimulation).

The monkeys used a low-friction manipulandum to control the movements of a cursor on a video screen. The manipulandum moved the cursor from the starting point at the center of the screen (origin) to a visual target in a delayed go-signal paradigm. The trial sequence and session flow is shown briefly in Figure 1 and in more detail in the study by Paz et al. (2003). Figure 1A depicts the trial flow from left to right. Note that in each trial, the monkey held its hand in the origin for a random delay of 1000–1500 ms (Hold1) without receiving any information about the type of upcoming trial. Later, after it was presented with a target, it continued to hold its hand in the origin for another random delay of 1000–1500 ms until onset of the go cue (Hold2) and then made a movement. Each recording session (day) involved three consecutive periods (Fig. 1B): (1) the prelearning period, a standard, eight-target task in which the target direction was randomly chosen from eight possible directions uniformly distributed over the circle; (2) the learning period, a transformed, one-target task in which only one target (upward, 90°) was presented and a rotational transformation was introduced between the cursor on the screen and the manipulandum; and (3) the postlearning period, in which the default, eight-target task was presented again. Figure 1A illustrates learning of a -90° visuomotor rotation; when the target is presented at 90° (upward), the hand should move toward 0° (i.e., a rightward hand movement; red) in order for the cursor to move from the origin to the target at 90° (upward cursor movement; green). The term “learned direction” (LD) in this study refers to the direction of hand movement needed to bring the cursor to the target during the visuomotor rotation learning. Rotations were 90, 45, -45, or -90° and were chosen randomly for each session but fixed for the duration of the adaptation period in a session. Monkeys were trained for several months with the default eight-target task but were not exposed to visuomotor rotation before the recordings.

Data analysis

We selected single neurons for analysis that were activated during contralateral arm movements and met two inclusion criteria: (1) well isolated spikes and (2) stable recordings based on firing rates in the first hold period (Hold1) for all trials. For the purposes of our analyses, it was important to obtain as many trials as possible. Therefore, we included only cells that were recorded reliably for N successful trials during learn-

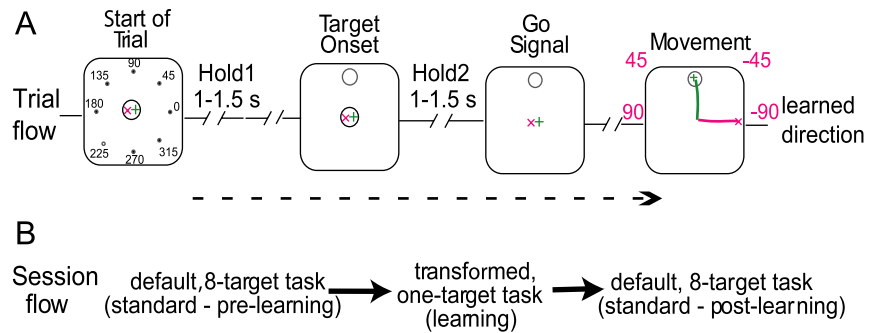


Figure 1. Experimental design. **A**, Single-trial flow (first row, from left to right) and session flow (second row). During period HOLD1, the monkey held the manipulandum in the center (origin) without moving it. The monkey kept holding at the origin after TO for an additional delay (hold2) and moved after the go signal. Shown (rightmost design) is a learning trial with a rotation of 90°. The red line shows arm movement (the LD), and the green line shows the resulting cursor movement. **B**, Session flow. Each day (session) consisted of three different epochs, with standard trials prelearning and postlearning. For details of the experimental design, see the study by Paz et al. (2003).

ing: $N \geq 34$ trials for M1 and $N \geq 38$ for SMA in monkey W; and $N \geq 40$ in monkey X, in which only M1 was sampled.

Behavioral performance during learning. The deviation in trajectories was assessed by a signed normalized deviation (SND), calculated as a directional deviation of the required hand direction minus the actual hand direction (taken at peak velocity), normalized by the rotation size in the session (-45, -90, +45, or +90°). See the learning curve in Figure 2B.

Fano Factor analysis. Fano Factor (FF) is a measure commonly used to estimate process variability (Fano U., 1947). In our case, FF estimates the trial-to-trial variability of spike counts of single cells in a given time window (Lee et al., 1998; Shallden and Newsome, 1998). We calculate the FF of a cell i in a selected group of trials p (a phase) for a time window T , by the following:

$$FF_i(p) = \frac{\text{variance}(r_i(p))}{\text{mean}(r_i(p))}, \quad (1)$$

where r_i is a vector of the spike counts of cell i in the given time interval in the sampled trials p . For a population of N cells, we calculate FF per phase p by the following:

$$FF(p) = \frac{1}{N} \sum_{i=1}^N FF_i(p). \quad (2)$$

The FF (variance/mean) for a Poisson distribution (in which the variance equals the mean) is 1. For processes that are sub- or supra-Poisson (more or less regular than a Poisson process, respectively), the FF depends on the observation length; for a very short observation, the FF tends toward 1, but as the length of the observation increases, the FF of a supra-/sub-Poisson process increases/decreases, respectively (Nawrot et al. 2008). Therefore, the dynamics and the relative change of FF between phases of learning (all having a fixed time window) are more relevant than its absolute values.

Relative FF. The FF during learning was normalized by dividing by the averaged FFs measured during the prelearning standard task for each direction separately:

$$\frac{1}{N_D} \sum_{d \in D} FF(d), \quad (3)$$

where D is the directions that were sampled with at least seven trials each and N_D is the size of D .

Alternative FF calculation. For every FF computation of a given set of trials, we eliminated two of the “suspected outlier” trials by either (1) elimination of the two trials with the maximum and minimum spike counts or (2) elimination of the two trials that deviated the most from the average spike count of the given set of trials (see results in Fig. 5B). An additional analysis compared FF computation for a

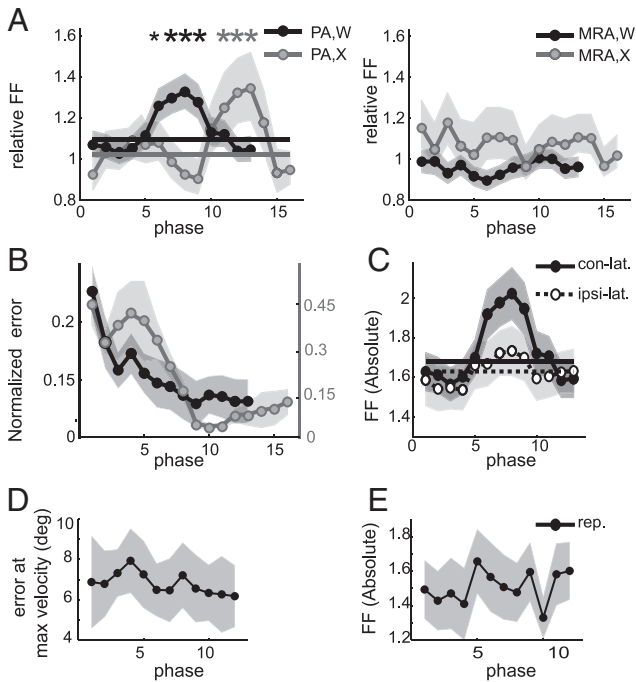


Figure 2. Typical dynamics of averaged FFs of M1 cells is temporally correlated to learning dynamics, shown for monkey X (black; $N_{\text{cells}} = 25$) and monkey W (gray; $N_{\text{cells}} = 114$). **A**, The average of relative FFs over all cells shows significant modulation of PA (left trace) but not of MRA (right trace). Parallel lines denote averaged bootstrap results for FF of PA for each monkey with corresponding colors. Large and small asterisks denote significant difference from chance of 5% and 1%, respectively. **B**, Behavioral improvement during learning. Performance is expressed by averaged SND over all trials of each phase. Note that monkey X was a slow learner and monkey W was much faster. Comparing the graphs in **A** and **B** shows the temporal relationship between behavioral improvement and FF increase. **C**, Comparison of averaged FF dynamics during learning recorded contralateral (con-lat.) to hand movement (monkey W, solid black line, same data as in **A**) to FF dynamics during ipsilateral (ipsi-lat.) recordings, averaged over $n = 107$ cells (dashed black line with white circles). Parallel solid and dashed lines denote averaged bootstrap results for the contralateral and ipsilateral conditions, respectively. Note that here FF is given in absolute values. **D**, Behavioral performance during repetition of movements to the same direction without visuomotor rotation. Performance is expressed by angular error at peak velocity averaged over all trials of each phase. Since there is no rotation, the angular error is not normalized as in **B** above. Note that as expected, no improvement in performance is evident. **E**, Average FFs across cells recorded during repetition sessions (rep.; $N_{\text{cells}} = 39$) shows no modulation along phases.

smaller number of trials (with five, six, eight, or nine trials in each calculation). This later analysis provides better temporal resolution during learning at the expense of the reliability of the variability estimation.

Definition of “phases” and sliding window technique

Phase p with $n = 10$ consecutive trials are in the range of $[j * (p - 1) + 1 \dots j * (p - 1) + n]$, where $j = 2$ is the size, in number of trials, of a step between phases. For instance, phase 1 consists of trials 1–10, phase 2 is trials 3–12, and so forth.

According to this definition, recordings from M1 were included if they consisted of at least 13 phases. A similar computation yielded 15 phases for SMA recordings in monkey W and 16 phases for M1 recordings in monkey X.

Note that neighboring phases include overlapping trials to different degrees and that there are no overlapping data at distances of five phases or more. For example, phase 1 and phase 6 show FF computed from trials 1–10 and 11–20, respectively. Therefore, the fact that neighboring phases at a distance less than five are both significantly different from bootstrap may be attributable to the overlapping data and were not considered as independent variables.

Since the window must be kept large enough for a reliable estimate of FF, the sliding window is a compromise to provide a higher resolution of the process dynamics by estimating FF in smaller steps compared to the window size. Thus, we can observe any triplet of independent phases (such as 1, 6, and 11 or 3, 8, and 13) to examine FF modulation. In short, the sliding window makes it possible to compare independent groups of trials by enabling flexibility when selecting the correct phases to observe the dynamics.

Tuning properties of single cells. We divided the cells into subgroups according to their tuning properties.

(1) A cell was said to be “tuned” if (1) the result of a one-way ANOVA showed a significant effect for direction ($p < 0.05$) and (2) a cosine fit $[r(d) = a + b * \cos(d - d_0)]$ that exceeded $R^2 = 0.55$. Otherwise, a cell was defined as “untuned.”

(2) The peak of the fitted cosine tuning curve was defined as the preferred direction (PD) of the cell. Each cell was also characterized by the absolute distance of its PD from the LD, the direction used during the learning session in which the cell was recorded.

Learning completion phase and maximal FF phase (“Maxphase”). Learning completion phase (LCP) is a metric for learning completion that is computed according to the SND as follows:

$\text{Errs}(p) = \text{mean}_{\text{tr}}(\text{SND}(\text{tr}))$ $\text{tr} = \{\text{trials of phase } p\}$ is the average error for phase p .

$\text{Limit} = \text{Errs}(1) - 0.8 * (\text{Errs}(1) - \text{Errs}(N_{\text{phases}}))$, where N_{phases} is the number of phases, is the limit for learning completeness, defined by 80% of the difference between the initial error and the last error.

$\text{LCP} = \min_p(\text{Errs}(p) < \text{limit})$ is the index of the first phase that crosses the limit.

LCP was evaluated for all sessions in which the learning curves showed consistent improvement (80% of the sessions in M1 recordings from monkey W, 75% of the sessions in M1 recordings of monkey X, and 71% of the sessions in SMA recordings from monkey W).

We examined the relationship between LCP and FF modulation. Cell c was included in the analysis if (1) its averaged FF along learning was larger than the averaged FFs in standard over all cells; (2) FFs showed a considerable modulation, namely the following:

$$\frac{\max(\text{FFs}) - \min(\text{FFs})}{\max(\text{FFs}) + \min(\text{FFs})} \geq 0.4, \quad (4)$$

and (3) LCP was defined.

Maxphase_c is the phase in which FFs for cell c peaked during learning, denoted as $\text{FF}_{\text{Maxphase}_c}$, and LCP_c is the LCP for the session in which cell c was recorded.

Cells were divided into two groups with sizes N_1 and N_2 according to their relationships between Maxphase and LCP for each cell: $\text{Maxphase}_c < \text{LCP}_c$ or $\text{Maxphase}_c \geq \text{LCP}_c$. For each group, we evaluated the variability index (VI):

$$\text{variability index, VI}_i = \frac{1}{N_1 + N_2} \sum_{c=1}^{N_i} \text{FF}_{\text{Maxphase}_c} \quad (5)$$

where $i = 1, 2$.

To test for a significant difference between these two groups, we used a bootstrap technique as follows: in each simulation, we randomly divided the cells into two groups and evaluated the VI for each group. The differences $\text{VI}_1 - \text{VI}_2$ and $N_1 - N_2$ were tested against the distributions of 1000 differences between simulated pairs of VIs and N_s

Significance of modulation. We defined the FFs discussed above as the “observed FFs”: $\text{obsFF}_{p,c}$, $p = 1 \dots N_{\text{phases}}$, $c = 1 \dots N_{\text{cells}}$, where N_{cells} is the number of cells and obsFF_p is a vector of all FFs in phase p .

To test the significance of obsFF modulations along phases, we used a bootstrap technique across trials, computing 1000 simulated FFs per cell. In each simulation, for each cell, we randomly chose 10 trials (length of a phase) from all the learning trials (from the session in which the cell was recorded) and computed the FF for the spike counts in these trials, denoted as $\text{simFF}_{s,c}$ [$c = 1 \dots \text{number of cells}$, $s = 1 \dots 1000$].

To evaluate obsFF_p , we used two criteria: (1) a Z test to check whether obsFF_p was significantly different from the distribution of 1000 simulated averages calculated as follows:

$$\overline{\text{simFF1}}(s) = \frac{1}{N_{\text{cells}}} \sum_{c=1}^{N_{\text{cells}}} \text{simFF}_{s,c} \quad (6)$$

where $s = 1 \dots 1000$ and where for each simulation we averaged its simulated FFs over all cells; and (2) a Z test to check whether obsFF_p was significantly different from the distribution of the N_{cells} simulated averages of N_{cells} calculated as follows:

$$\overline{\text{simFF2}}(c) = \frac{1}{1000} \sum_{s=1}^{1000} \text{simFF}_{s,c} \quad (7)$$

where $c = 1 \dots N_{\text{cells}}$ and where for each cell we averaged over all its simulated FFs.

We considered a phase to be significantly different from chance if p values were smaller than 1% under these two criteria. When the p value was between 1% and 5%, significance is denoted by a small asterisk (even if it was 1% for one of the two).

Significance of the temporal difference between two FF curves. We calculated the temporal difference (in phases) between two curves (averaged FF values per phase) by the difference between their peaks. The peak for each curve was obtained by two methods: (1) fitting a cosine function. A good cosine fit ensures that both simulated groups approximately had the same FF pattern of modulation as the real data, therefore making the computation of a temporal shift meaningful; and (2) a generic peak-finding algorithm: the phase with a maximum FF was detected (P). The peak was detected by a quadratic least-squares curve-fitting of phases [$P - 1$ P $P + 1$].

To test for a significant temporal difference between the two curves, we ran a bootstrap analysis as follows: we pooled the data of the two curves (FFs of all cells from all phases). In each simulation, we randomly split the data into two groups and computed the temporal difference between them according to both peak-finding methods. We ran the simulation 1000 times and tested whether the temporal difference obtained according to method 1(2) could come from the distribution of temporal differences computed by method 1(2) we obtained by bootstrap.

Results

Behavior

The behavioral task was reported by Paz et al. (2003) and is summarized here in Materials and Methods and illustrated in Figure 1. Analyses of the monkeys' behavior were also reported by Paz et al. (2003). In brief, trajectory deviations were taken at peak velocity and normalized to the magnitude of transformation (see Materials and Methods). The deviations appeared to be the largest for the target used in the learning epoch and decreased as a function of angular distance from this target and as a function of trial number after learning (Paz et al., 2003) (Fig. 2A). Thus, the aftereffects showed limited generalization across the work space.

During the learning epoch itself, more than 10 trials, on average, were required to reach a behavioral performance plateau with small and stable errors (Paz et al., 2003) (Fig. 2A). As in other studies, learning consisted of two phases: a transient phase during the first 10 trials, followed by a slower phase. The next recording day, no aftereffects of the previous learning day were observed. Furthermore, a different LD (see Fig. 1 legend) from the previous day was always chosen to further ensure that learning occurred on a daily basis, making it possible to pool neurons across recording days.

Neuronal data

We analyzed neuronal activity in two epochs of the trial: (1) activity from 100 to 600 ms after target onset (TO; before the go signal, when the animal is instructed to avoid any movement) was

termed "preparatory activity" (PA); and (2) the activity around MO (from 100 ms before to 400 ms after) was termed "movement-related activity" (MRA). Note that we use these traditional terms for convenience. The first period is distinctly different from the second only by the design of animal training (in the first, the monkeys are not allowed to move, and in the second, they should move to reach success). This temporal separation allows some separation of the process choosing the action (from TO and on) and initiation of the response (from go signal and on).

The sample totaled 139 cells (114 in monkey W and 25 in monkey X) from M1 and 78 cells (all in monkey W) from the SMA. Criteria for selecting cells are described in Materials and Methods.

The analysis focused on evaluating variability of trial-to-trial neuronal activity during adaptation to visuomotor rotation. Variability was assessed for spike counts of single cells, expressed by the FF (see Materials and Methods). We show the dynamics of variability modulations by calculating averaged FF across populations of cells (except for Fig. 4, which shows the FFs for single cells).

Analysis of neuronal activity in M1 and SMA-proper revealed the specific dynamics of this variability. The sections below describe the dynamics, characterize the group of cells that were primarily involved in the generation of these dynamics, and compare the two cortical areas.

FF analysis of single-cell activity in M1

FF during PA versus MRA

Figure 2A compares the averaged variability across the sampled population of cells during different epochs of the trial shown separately for each of the two monkeys (black, monkey W; gray, monkey X). For each monkey, it shows the relative FF (see Materials and Methods), calculated separately for PA (left) and MRA (right).

The figure shows that FF for PA increased in a series of consecutive phases. For monkey W, a significant increase was already visible at phase 6 until phase 9, whereas for monkey X, the increase was significant from phase 11 to phase 13. The phases with a statistically significant increase are marked by asterisks (bootstrap analysis; large asterisks, $p < 0.01$; small asterisks, $p < 0.05$; see Materials and Methods). The figure also shows that increased FFs were only observed during the PA epoch and not during MRA (similarly, there were no significant FF changes during the Hold1 period; data not shown). FF curves of monkeys W and X peaked at phases 7.9 and 11.8, respectively, according to a cosine fit and at phases 8.1 and 12.6 according to a generic peak-finding algorithm (see Materials and Methods). A temporal difference of 3.9 and 4.5 phases in the first and second methods were significant ($p < 0.01$).

The increase emerged at different phases in the two monkeys. To better understand this difference, we looked at the learning curves of the two monkeys and computed the averaged directional error at peak velocity for each phase. These learning curves are compared in Figure 2B. Note that monkey W (black line) was a faster learner; monkey X (gray line) showed little reduction in directional error in the first five phases, whereas monkey W had almost completed learning at this time.

To quantify learning dynamics, we calculated for each monkey and for each session the phase number at which the directional error was reduced by 80%. We called this phase the LCP (see Materials and Methods). Monkey W showed an overessions averaged LCP of 6.2, and monkey X had an LCP of 9. This difference of 2.8 phases between the learning curves of the two monkeys is statistically significant (ANOVA, $p < 0.05$).

Table 1. Relation between FF increase and learning completeness

Recording area	Monkey, number of cells	Percentage of cells ^a			VI		
		Maxphase > LCP	Maxphase < LCP	<i>p</i> value	Maxphase > LCP	Maxphase < LCP	<i>p</i> value
M1	Monkey W, <i>n</i> = 33	76%	24%	1	3.13	0.88	1
	Monkey X, <i>n</i> = 12	75%	25%	5	3.67	0.83	5
	Both monkeys, <i>n</i> = 45	75%	25%	1	3.27	0.95	1
SMA	Monkey W, <i>n</i> = 42	23%	77%	1	0.75	2.9	1

Cells are divided to two groups according to the relationship between Maxphase and LCP for each cell: Maxphase > LCP or Maxphase < LCP. For each monkey and cortical area, the data show the percentage of cells and VI of each group. The *p* value is the significance of the difference between the percentage of cells (on the left) and VIs (on the right) of the two groups.

^aCells were included if: (1) variability criterion: averaged FF along learning > averaged FF of STD over all cells; (2) modulation criterion: (maxFF – minFF)/(maxFF + minFF) > 0.4; and (3) LCP was defined.

To further explore the apparent relationship between FF modulation and improvement in performance, we investigated the relationship between the phase in which FFs of single cells reached its peak (Maxphase) and the LCP. Although the overall correlation was weak, a clear relationship was found when dividing the cells into two groups, one with Maxphase < LCP (cells that showed peak of FF modulation before LCP) and the other with Maxphase ≥ LCP. For each of the monkeys, we found that the first group was significantly smaller than the second group (monkey W, *p* < 0.01; monkey X, *p* < 0.05). Namely, in fewer cells, Maxphase preceded LCP, and it had a significantly lower VI (see Materials and Methods; statistical significance was *p* < 0.01 for monkey W and *p* < 0.05 for monkey X) than the second group (Table 1). Thus overall, this analysis supports the notion that the dynamics of neuronal variability are related to the dynamics of learning. In both monkeys, variability in M1 was low in the early learning trials when improvement was fast (Fig. 2*B*). The increase in FF did not emerge, in either monkey, before the end of this fast-learning stage. However, the increase was not tightly coupled to LCP epoch and could also occur in later phases of learning.

At the very end of the learning block, the FFs decreased again to their initial values (monkey W, phase 10; monkey X, phase 14). This profile of transient increased variability during the learning block is termed here “typical FF dynamics during learning.”

In the sliding window technique, as stated previously, phases that are at least five phases apart do not overlap; for example, phases 3, 8, and 13 correspond to groups of trials 5–14, 15–24, and 25–34. For every phase *p* (marked by an asterisk in the figures), we can construct a triplet of phases [*p* – 5 *p* + 5]. In all such triplets, only phase *p* significantly increased while the other two did not. Therefore, although the sliding window yields an overlap between phases, each such triplet exhibits the typical dynamics without any overlapping. In the example of the triplet [3,8,13] in monkey W, FF did not increase above chance in the first and third groups of trials but showed a significant increase for the second group (Fig. 2*A*).

We tested FF of the PA in several conditions against the FF of the cells shown in Figure 2*A* for monkey W. First, we compared the FF modulations of these cells recorded in the contralateral hemisphere to 107 cells that were recorded in the ipsilateral hemisphere. Figure 2*C* (solid vs dashed black lines; absolute FF values without normalization that was used in 2*A*) shows that ipsilateral activity failed to reach a significant level of modulation (bootstrap, *p* > 0.1).

Second, some sessions served as a “repetition control,” which entailed a one-target task without visuomotor rotation. These were helpful to preclude the possibilities that the observed effect was merely attributable to repetition of a given movement during the learning epoch and that FF modulation can occur without an improvement in performance taking place. Our findings clearly show that although directional error did not change along the repetitive movement executions (Fig. 2*D*), values of FF (~1.5,




Figure 3. Averaged FFs (top row) and PETHs (bottom row) during epochs around (A) TO and (B) MO for phases 3, 8, and 13. No difference between the phases was found around MO for either PETHs or FFs (right column). Around TO, PETHs did not show any difference between the phases, whereas FFs of phase 8 after TO deviated significantly above phases 3 and 13 (left column, bottom row). The blue dashed line marks the time of event occurrence, and black lines mark the time windows of PA and MRA, as defined in this work. Large and small asterisks denote significant difference from chance of 5% and 1%, respectively.

computed from 39 cells) never deviated from chance during any of the phases (Fig. 2*E*) and did not differ from FF of standard trials (data not shown; bootstrap, *p* > 0.1).

To obtain a better understanding of the origin of FF dynamics, we depict an additional analysis of the neuronal data from monkey W in Figure 3. The figure shows the averaged FFs in the top row and perievent time histograms (PETHs) in the bottom row for all sampled neurons of this monkey in 100 ms bins, around TO (A) and around MO (B). The FFs and PETHs are shown for three distinct phases: phases 3, 8, and 13. Only phase 8 (red) showed high FF (see Fig. 2), whereas phases 3 (blue) and 13 (green) did not. Clearly, there was a significant FF increase only during PA, and only in phase 8 in comparison to phases 3 and 13. Note that this difference was not correlated with the firing rate modulation as shown by the PETHs for the same time interval, as there was no significant difference between the responses in phases 3, 8, and 13 (ANOVA, *p* > 0.5). Figures 2 and 3 therefore show that FF modulations of the population of cells occurred only during the preparatory epoch (PA). Figure 4 shows examples of these modulations in three single cells. For cell 1, it shows the raster plots around TO (first row) and MO (second row) and spike counts in each trial for PA and MRA, respectively (Fig. 4*A*). Two red dashed lines delimit the trials of phase 8, and the red circles mark spike counts. Figure 4*B* shows the FFs during PA for this cell (black line), depicting the peak modulation in phase 8, whereas in MRA (gray line), no significant changes in FF were observed. Red circles mark the FFs of phase 8 in both cases.

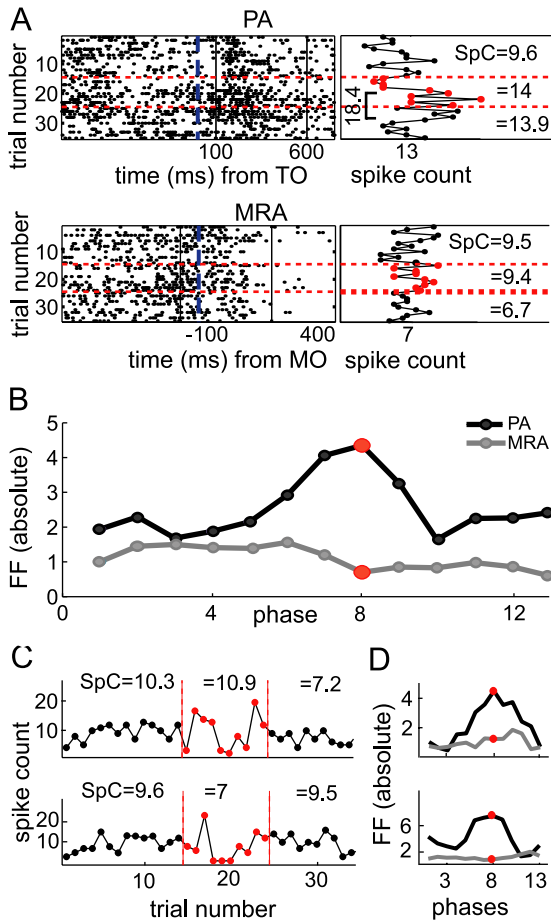


Figure 4. FF modulations of M1 single cells. **A**, First and second rows show raster plots around TO and MO, respectively. Spike counts in the windows of PA and MRA for each trial are denoted on the right. **B**, FFs along the phases for PA (black) and MRA (gray). Phase 8 is highlighted by red markers both for PA and MRA computations; dashed red lines indicate the bounds of phase 8 in the PETHs, and red circles indicate their corresponding spike counts and FF. Averaged spike counts before, in, and after phase 8 are denoted. Black and blue lines are as in Figure 3. **C, D**, Activity modulations in two other cells (top and bottom rows). **C**, Spike counts of PA along learning trials. **D**, FFs along phases for PA and MRA. Markers are as before. SpC, Spike count.

To explore possible relationships between spike counts and FFs for this single cell, we examined phases 8 and 10 (delimited by a square bracket in the top right plot of Fig. 4A). The spike counts increased along the trials of phase 8 (to 14). The transition from low counts to high counts in this phase is reflected by the high FF. It could be claimed that the high FF is merely a result of the high counts. However, this does not reconcile with the relationship between spike counts and FF in phase 10, where the average spike count was even higher (18.4) than in phase 8 (14) but the FF was much lower (as seen in Fig. 4B). In fact, FF in phase 10 was close to FFs in the first few phases, when the spike counts were the lowest (9.6). Therefore, the increase in FF in this cell was not attributable to an increased spike count but rather caused by the modulation from low to high spike count. Figure 4, C and D, depict spike counts (Fig. 4C) and FFs (Fig. 4D) of two other cells that illustrate two additional ways for FF to increase because of spike count modulations. FF may increase without a change in spike count, but merely because of an increase in variability (Fig. 4C,D, top row) or a sudden and temporary decrease in spike counts (Fig. 4C,D, bottom row). Altogether, the observed FF in a given phase could not be predicted by the average spike counts in that phase.

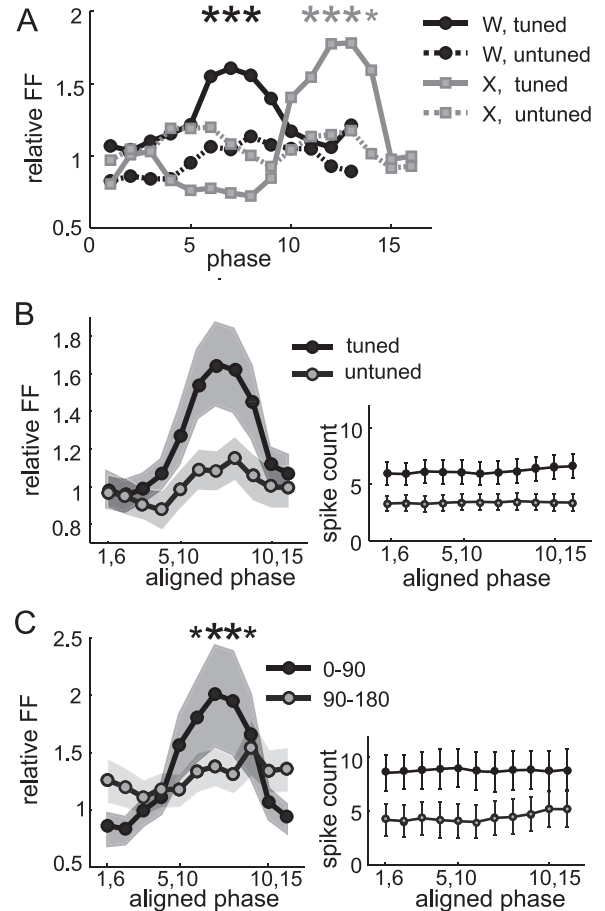


Figure 5. FF typical dynamics emerge primarily in directionally tuned cells with a PD close to the LD. **A**, FF dynamics for monkey W (black, circle) and monkey X (gray, square), split into tuned (solid lines) and untuned (dotted lines) cells. Tuned cells (monkey W, $N_{\text{cells}} = 17$; monkey X, $N_{\text{cells}} = 7$) show the typical dynamics, whereas untuned do not. **B**, Left, Alignment of both monkeys' FF dynamics to pool their cells. Aligned phases are denoted for monkeys W and X. Right, Spike counts for tuned and untuned cells. **C**, Left, Division of tuned cells according to distance of the cell's PD from the LD. Cells with PD $< 90^\circ$ away from the LD ($N_{\text{cells}} = 9$) showed the typical FF modulation. Right, Spike counts for closely and distantly tuned cells. Apparently, FF typical dynamics cannot be explained by changes in spike counts. Large and small asterisks denote significant difference from chance of 5% and 1%, respectively.

The remainder of the results focuses only on the PA (i.e., the epoch that showed significant modulations of the FF).

Relationship of single-cell directional tuning and its FF dynamics during learning

Figure 5A shows plots of FF values as a function of phase (as in Fig. 2) calculated separately for tuned (solid lines) and untuned (dotted lines) cells (see Materials and Methods), for each of the two monkeys (monkey W, black with circles; monkey X, gray with squares). The FF of tuned and untuned cells was not significantly different in the first and last phases (ANOVA, $p > 0.1$). During learning, the tuned cells exhibited significantly more pronounced typical FF dynamics than the untuned cells [i.e., an ANOVA showed ($p < 0.01$) a significant difference between tuned and untuned cells only in phases 6–8 in monkey W and phases 11–13 in monkey X] (Fig. 5A).

To examine the relationship between FF modulation and directional tuning more closely, we divided the tuned cells into two groups: the first had PDs close to the LD ("closely tuned"; $0-90^\circ$ from LD: note that LD is the actual direction of hand movement used during learning), and the second had PD away from the LD

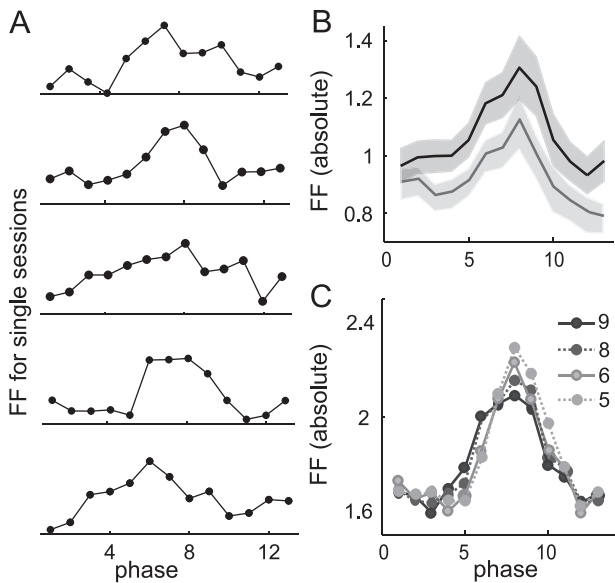


Figure 6. FF typical dynamics emerged under several control conditions ($N_{\text{cells}} = 114$, all from monkey W). **A**, Single sessions. Each trace corresponds to a single session, showing the averaged FFs of tuned cells recorded along it. **B**, Eliminating, for each phase, the maximal and minimal spike counts (black) or those two that deviated most from the average (gray). In both cases, FF modulation remained significant, suggesting it did not result from mere outliers. **C**, FF computations for five, six, eight, or nine trials (from light to dark shades of gray) in a phase. All show the typical FF modulation.

(“distantly tuned”; -90 to 180°). To calculate with a sufficient number of cells, we pooled all neurons from the two monkeys. As mentioned above, we found that the temporal difference in dynamics of FF modulation between the two monkeys was related to the dynamics of learning (Fig. 2). To compare the typical FF dynamics in the two monkeys, we aligned the phases in which modulation increased in the two monkeys according to the first phase, which showed a significantly increased FF (“starred phase”). Thus, we used the phases starting from five phases before the starred phase until five phases after it. These phases are denoted 1,6–11,16 in Figure 4, *B* and *C*, for monkeys W and X, respectively. As expected from Figure 4*A*, there was a clear difference between the FF modulations for tuned cells (black) and untuned cells (gray) after alignment (Fig. 5*B*, left trace). Furthermore, closely tuned cells (black) showed the typical FF modulation, whereas the distantly tuned cells (gray) showed no modulation. The clear difference as seen in Figure 5*C* (left trace) is statistically significant ($p < 0.01$).

Although the group of tuned cells exhibited a higher firing rate, on average, than the untuned ones (Fig. 5*B*, right trace), as was the case for the closely tuned cells compared with the distantly tuned cells (Fig. 5*C*, right trace), the results indicate that the different firing rates could not be the direct source of FF modulation during learning. This is apparent from the analysis that shows that the firing rates do not show any significant modulation in different phases, including the relevant ones (phases 6–8).

Controls for FF modulation

To further test the validity of FF modulations, we conducted additional controls. First, we tested whether this typical modulation of FFs during PA could also be observed in single sessions (Fig. 6*A*). The modulation of the tuned cells (while naturally less significant for fewer data) was evident even in single sessions. Second, we controlled for the possibility that the increase in variability was not merely attributable to a change in the spike count

of a single trial, but rather to an overall propensity for an increase in variability during learning. To do so, FF was computed after eliminating the two trials with the most extreme spike counts using two approaches (see Materials and Methods). In both approaches (Fig. 6*B*), we obtained FF values that were, by virtue, somewhat lower than the original ones but maintained the typical dynamics of modulation during learning (namely phases 7–9, and only these were found to increase significantly above chance; bootstrap, $p < 0.01$). Third, we computed the FF for different window lengths (e.g., different number of trials in a phase; tested for nine, eight, six, and five trials per phase). Note that as the number of trials in a phase decreases, the FF captures more local changes from trial to trial and cannot capture slower changes; however, as shown in Figure 6*C*, for all window lengths we observed similar dynamics of FF modulation during learning.

SMA analysis

The same analysis was applied to recordings of single-cell activity in SMA (monkey W only). Figure 7*A* shows examples of two single cells: spike counts of PA along trials (left) and the corresponding FFs of PA along phases in black; and FFs of MRA in gray (right). For each cell, the spike counts of the phase with the maximal FF are denoted in red circles and bounded by two red dash lines. Averaged spike counts before, during, and after this phase are denoted above the trace. Similarly to M1 cells (Fig. 4), FFs of PA increase in SMA cells occur because of different patterns of spike count modulations.

Figure 7*B* depicts the averaged FFs (top row) and spike counts (bottom row) across all SMA cells. No significant dynamics were found in the FFs and spike counts of MRA (Fig. 7*B*, blue traces). The FF of the PA is depicted in Figure 7*B* (top row; red). As in Figure 2*A*, the typical FF dynamics during learning were evident for SMA cells as well, without a significant modulation in the spike counts (Figure 7*B*, bottom row, red). Neuronal activity recorded ipsilaterally to the movement (Fig. 7*B*, top row, dashed red line) show weak and insignificant FF modulation (again, as in M1) (Fig. 2*C*). However, there was a marked difference between FF modulation in M1 and SMA; in SMA, the increase in FF occurred earlier in learning and became significant as early as phase 4. For easy comparison, both FF curves are aligned together by normalization in Figure 7*C*. In SMA (red), FFs started to diverge almost immediately at the beginning of the learning block (phase 3), reached a maximum by phase 5, and converged to baseline by phases 8–9. In M1 (black), it occurred later with a delay of about three phases. This difference was statistically significant as found by bootstrap analysis ($p < 0.05$; see Materials and Methods). Cosine fits for the SMA and M1 curves peaked at 4.4 and 7.9, respectively, showing a marked temporal difference of 3.5 phases. A generic peak-finding algorithm (see Materials and Methods) resulted in peaks at 5.2 and 8.1, again showing a significant difference of 2.9 phases ($p < 0.05$).

Inspection of the relationship between Maxphases of single cells in SMA and LCPs (see Materials and Methods) showed, as expected from Figure 7*B*, the opposite distribution of cells compared with M1. As shown in Table 1, for most SMA cells, the peak of FF did not exceed LCP ($p < 0.05$), whereas in M1 only few cells had a Maxphase that preceded LCP. Moreover, this group of SMA cells had a significantly higher VI than the group of SMA cells that peaked after LCP ($p < 0.01$).

The SMA has weaker directional tuning properties in general, and in our recordings specifically. Therefore, we did not study the relationship of FF modulation and directional tuning in the SMA.

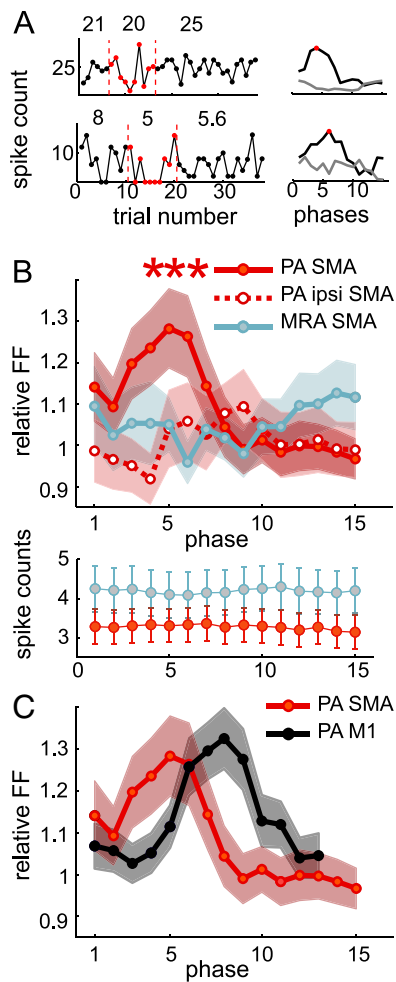


Figure 7. FF analysis in the SMA ($N_{\text{cells}} = 78$) shows the typical dynamics of PA that temporally precede that in M1. **A**, Activity modulations in two SMA cells. Spike counts of PA along learning trials (left) and FFs along phases for PA (black) and MRA (gray). For each cell, spike counts of the phase with maximal FF of PA are marked by red circles and bounded by two dashed red lines. FFs of PA and MRA in this phase are also marked by red circles. Asterisks denote 1% significant difference from chance. **B**, Top row, Relative FFs averaged across cells for PA in SMA cells recorded contralateral to the moving hand display the typical dynamics (solid red), whereas ipsilateral cells did not (dashed red line). FFs for MRA did not show the dynamics either (blue). Bottom row, Averaged spike counts for PA (red) and MRA (blue) for contralateral SMA cells showed no modulation. **C**, FF for PA in SMA (red) and M1 (black) showing that the typical dynamics in SMA preceded M1.

Discussion

This study shows dynamic modulations of trial-to-trial variability of firing rate of single cells in the primary and premotor cortices when monkeys adapt to a novel visuomotor rotational task. We found that variability increased during learning and returned to baseline as learning progressed. These changes occurred early in learning in the SMA and only later in M1 after the learning curve had almost reached its plateau. The modulations were related to the learning curves in each monkey and consistently reflected the different learning rate of the two monkeys (the slow learner showed the modulation later). The changes were much more pronounced in cells with PDs close to the one direction experienced during the adaptation. Importantly, the changes occurred primarily in the premovement delay period, after the TO but before the go signal. Below, we discuss these findings in detail and propose possible roles for increased variability in learning and memory.

Premovement variability and motor states

The increased variability occurred preferentially in the premovement delay period, when the animal could know the required movement but was not allowed to initiate it. Numerous studies have suggested that the activity in motor areas during this period reflects preparation of movement (Kurata and Wise, 1988; Mitz et al., 1991; Tanji and Mushiaki, 1996; Wise et al., 1998; Paz et al., 2003; Padoa-Schioppa et al., 2004).

It was recently shown that trial-to-trial variability during this period reflected properties of the upcoming movement (Churchland et al., 2006a) and is correlated with movement reaction time (Nawrot et al., 2003; Churchland et al., 2006b). We suggest that the increased variability during the delay period could reflect several aspects of adjustments and consolidation of the internal models that control sensorimotor associations (Wolpert et al., 1995; Kawato and Wolpert, 1998; Shadmehr and Krakauer, 2008). For example, modulation of variability might represent adjustments of the forward model that the brain can use to test different movement plans to minimize on-line corrections. Alternatively, the variability may represent the process selecting the association between an instruction (TO) and a new specific action (movement direction) (Wise et al., 1996; Zach et al., 2008). The association requires few (if any) changes in the kinematics to dynamics transformations since there was no alternation in the relationship between movement and muscle activity. We therefore hypothesize that in force-field adaptation tasks (Shadmehr and Thoroughman, 1998; Li et al., 2001) that require changes of kinematics to dynamics transformations, trial-to-trial variability should be observed during movement as well.

Learning in motor cortical hierarchy

We report that changes in variability occur in the SMA in early trials of learning and only later in M1. Several studies have shown that the late period, but not the early period, is accompanied by changes of activity in M1 (Karni et al., 1995; Sanes and Donoghue, 2000; Ungerleider et al., 2002; Kleim et al., 2004). These changes may evolve even within one single session (Laubach et al., 2000; Li et al., 2001; Cohen and Nicolelis, 2004). Specifically, in the same task used here, changes in the neuronal activity in M1 arose in the late learning trials (Paz et al., 2003). It was suggested that initial learning is mediated by other areas such as the cerebellum, basal ganglia, and parietal cortex (Doyon et al., 1997; Hikosaka et al., 2002; Ungerleider et al., 2002), as well as the premotor cortices and the SMA (Nakamura et al., 1998; Hikosaka et al., 2002; Padoa-Schioppa et al., 2002; Lee and Quessy, 2003; Krakauer et al., 2004). These findings suggest that the fast and slow stages of learning (Karni et al., 1998; Krakauer et al., 2000) are represented differentially in the two sampled cortical areas. The neuronal changes during the early and fast improvement of behavior are reflected in the SMA but not in M1, whereas the slower stage with moderate improvement of behavior is reflected more in M1. This finding supports previous results showing similar differences for cross-neuron variability (but within a single trial) (Paz et al., 2005). The trial-to-trial variability measured here also supports the notion of hierarchical processing of motor learning. The first stage is reflected in premotor areas (the SMA, in this case), and only later in M1. We speculate that the SMA is involved directly in finding a neuronal configuration that results in a minimal error of trajectory, whereas M1 involvement is not as straightforward. It might be involved in modifying the forward model (Wolpert et al., 1995; Karni et al., 1998) or in consolidation, since new representations were shown to exist when returning to standard conditions after learning (Paz et al., 2003; Zach et al., 2008).

The temporary increase in variability in each area may reflect increases in variability in the inputs or in the synapses to the sampled cells (Riultz-Pedotti et al., 1998, 2000) because of a gradual recruitment and reorganization of brain areas (Imamizu et al., 2007) (e.g., from the parietal cortex to SMA and other premotor cortices) (Andersen and Buneo, 2002), and from SMA to M1.

Specificity for movement directions

Visuomotor tasks have been used extensively in various studies (Abeele and Bock, 2001; Krakauer et al., 2005; Bock et al., 2009). Here, we used only one movement direction during learning (out of the eight possible ones). It is well known that this mode of learning induces “local learning” with limited generalizations. Namely, the effect on movement directions decays as a function of the distance from the LD (the hand movement direction that was experienced during learning). This was shown for humans (Pine et al., 1996; Krakauer et al., 2000) and monkeys (Paz et al., 2005).

We found that cells with a PD (Georgopoulos et al., 1982) close to the LD ($<90^\circ$) showed increased trial-to-trial variability, whereas the population of cells with a PD far from the LD did not present this increase. This is in line with our previous findings that these cells underlie the adaptation process (Paz et al., 2003; Paz and Vaadia, 2004). Moreover, it provides an important inherent control and support for the hypothesis that this variability reflects searching within a restricted range of putative states, rather than a general increase in noise in the system that would be exhibited in all neurons regardless of their tuning.

Suggested framework for motor learning

In a simple stationary deterministic network, under noiseless conditions, there is a one-to-one correspondence between neuronal configurations and actions; namely, a single configuration always results in the same action. If the motor cortex is seen as the output stage of such a control system, its activity should be fixed and stable to generate precise and replicable movements. However, if the motor cortex is also part of an adaptive system, such fixed activity would not allow flexibility and plasticity that can cope with noise and changing conditions.

Using the concepts of attractor states (Hopfield, 1982; Hopfield and Herz, 1995; Goldberg et al., 2004) it may be the case that under constant conditions (control of highly stereotypical precise movements) this network stays in a fixed point, which represents a single neuronal configuration in the neuronal space. It was recently suggested that redundancy in the motor cortex provides robustness to noise in the system by allowing changes in neuronal activity that do not significantly affect behavior. Thus, the initial low trial-to-trial variability (which we observed) has been defined as “background” variability, which reflects a stable optimal manifold (rather than a single point) in the neuronal space in which all points produce a similar end result, but with different motor commands (Rokni et al., 2007).

This approach is attractive since trial-to-trial variability, while slightly impairing accuracy (Chestek et al., 2007), can facilitate plasticity. When confronted with a novel sensorimotor environment (like visuomotor rotation), the neuronal configuration is no longer optimal. Variability facilitates plasticity by allowing the network to leave the current configuration and search for a new appropriate configuration (Sutton and Barto, 1998). We suggest that the increase in variability during learning reflects this search, and when a new appropriate configuration is found, variability decreases once again. Furthermore, the increase to relatively large

variability during learning could serve to prevent convergence to local minima and thus may make it possible to achieve optimal neuronal configurations (Kirkpatrick et al., 1983).

In summary, this notion implies that the observed variability facilitates learning since the required transitions can end when one of many optimal solutions spanned by the stable optimal manifold (network configurations) is reached, rather than a single stable point.

Conclusion

An adaptive system like the brain is required to perform as accurately as possible and, at the same time, still adapt to ongoing changes. This induces a natural tension between the need for reliable (stable) neuronal configurations and (less stable) neuronal plasticity that enables adaptation and learning. Here, we suggest that variability represents a motor learning mechanism in which the system can balance these opposing requirements using limited and controlled variability. It allows for searching, reaching, and remaining in an optimal manifold with appropriate neuronal configurations that preserve reasonably accurate performance. This may serve as the basis for acquisition and consolidation of desired neuronal states, a mechanism that serves learning and memory of reaching movements.

References

- Abeele S, Bock O (2001) Mechanisms for sensorimotor adaptation to rotated visual input. *Exp Brain Res* 139:248–253.
- Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25:189–220.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996a) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996b) Predicting cortical evoked responses in spite of their large variability—the impact of ongoing activity. Rehovot: The Grodetsky Center for Research of Higher Brain Functions.
- Bach M, Kruger J (1986) Correlated neuronal variability in monkey visual cortex revealed by a multi-microelectrode. *Exp Brain Res* 61:451–456.
- Bialek W, Setayeshgar S (2008) Cooperativity, sensitivity, and noise in biochemical signaling. *Phys Rev Lett* 100:258101.
- Bialek W, Rieke F, De Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Bock O, Worringham C, Dawson S (2009) On the use of different spatial reference frames for movement control. *Motor Control* 13:43–53.
- Chestek CA, Batista AP, Santhanam G, Yu BM, Afshar A, Cunningham JP, Gilja V, Ryu SI, Churchland MM, Shenoy KV (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27:10742–10750.
- Churchland MM, Afshar A, Shenoy KV (2006a) A central source of movement variability. *Neuron* 52:1085–1096.
- Churchland MM, Yu BM, Ryu SI, Santhanam G, Shenoy KV (2006b) Neural variability in premotor cortex provides a signature of motor preparation. *J Neurosci* 26:3697–3712.
- Cohen D, Nicolelis MA (2004) Reduction of single-neuron firing uncertainty by cortical ensembles during motor skill learning. *J Neurosci* 24:3574–3582.
- Doyon J, Gaudreau D, Laforce R Jr, Castonguay M, Bedard PJ, Bedard F, Bouchard JP (1997) Role of the striatum, cerebellum, and frontal lobes in the learning of a visuomotor sequence. *Brain Cogn* 34:218–245.
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303.
- Fano U (1947) Ionization yield of radiations. II. The fluctuations of the number of ions. *Phys Rev* 72:26–29.
- Fiete IR, Fee MS, Seung HS (2007) Model of birdsong learning based on gradient estimation by dynamic perturbation of neural conductances. *J Neurophysiol* 98:2038–2057.
- Fusi S (2002) Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol Cybern* 87:459–470.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the rela-

- tions between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537.
- Goldberg JA, Rokni U, Boraud T, Vaadia E, Bergman H (2004) Spike synchronization in the cortex–basal ganglia networks of parkinsonian primates reflects global dynamics of the local field potentials. *J Neurosci* 24:6003–6010.
- Hamilton AF, Jones KE, Wolpert DM (2004) The scaling of motor noise with muscle strength and motor unit number in humans. *Exp Brain Res* 157:417–430.
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784.
- Hikosaka O, Nakamura K, Sakai K, Nakahara H (2002) Central mechanisms of motor skill learning. *Curr Opin Neurobiol* 12:217–222.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A* 79:2554–2558.
- Hopfield JJ, Herz AV (1995) Rapid local synchronization of action potentials: toward computation with coupled integrate-and-fire neurons. *Proc Natl Acad Sci U S A* 92:6655–6662.
- Imamizu H, Higuchi S, Toda A, Kawato M (2007) Reorganization of brain activity for multiple internal models after short but intensive training. *Cortex* 43:338–349.
- Jones KE, Hamilton AF, Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88:1533–1544.
- Kao MH, Wright BD, Doupe AJ (2008) Neurons in a forebrain nucleus required for vocal plasticity rapidly switch between precise firing and variable bursting depending on social context. *J Neurosci* 28:13232–13247.
- Karni A, Meyer G, Jezzard P, Adams MM, Turner R, Ungerleider LG (1995) Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377:155–158.
- Karni A, Meyer G, Rey HC, 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, Wolpert D (1998) Internal models for motor control. *Novartis Found Symp* 218:291–304.
- Kirkpatrick S, Gelatt C, Vecchi M (1983) Optimization by simulated annealing. *Science* 220:671–680.
- Kleim JA, Hogg TM, Vanderberg PM, Cooper NR, Bruneau R, Rempel M (2004) Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. *J Neurosci* 24:628–633.
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20:8916–8924.
- Krakauer JW, Ghilardi MF, Mentis M, Barnes A, Veytsman M, Eidelberg D, Ghez C (2004) Differential cortical and subcortical activations in learning rotations and gains for reaching: a PET study. *J Neurophysiol* 91:924–933.
- Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *J Neurosci* 25:473–478.
- Kurata K, Wise SP (1988) Premotor cortex of rhesus monkeys: set-related activity during two conditional motor tasks. *Exp Brain Res* 69:327–343.
- Laubach M, Wessberg J, Nicolelis MA (2000) Cortical ensemble activity increasingly predicts behaviour outcomes during learning of a motor task. *Nature* 405:567–571.
- Lee D, Quessy S (2003) Activity in the supplementary motor area related to learning and performance during a sequential visuomotor task. *J Neurophysiol* 89:1039–1056.
- Lee D, Port NP, 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.
- Maimon G, Assad JA (2009) Beyond Poisson: increased spike-time regularity across primate parietal cortex. *Neuron* 62:426–440.
- 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.
- Nakamura K, Sakai K, Hikosaka O (1998) Neuronal activity in medial frontal cortex during learning of sequential procedures. *80:2671–2687*.
- Nawrot MP, Aertsen A, Rotter S (2003) Elimination of response latency variability in neuronal spike trains. *Biol Cybern* 88:321–334.
- Nawrot MP, Boussein C, Rodriguez Molina V, Riehle A, Aertsen A, Rotter S (2008) Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods* 169:374–390.
- Padoa-Schioppa C, Li CS-R, 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, Vaadia E (2004) Learning-induced improvement in encoding and decoding of specific movement directions by neurons in the primary motor cortex. *PLoS Biol* 2:E45.
- 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, Natan C, Boraud T, Bergman H, Vaadia E (2005) Emerging patterns of neuronal responses in supplementary and primary motor areas during sensorimotor adaptation. *J Neurosci* 25:10941–10951.
- Pine ZM, Krakauer JW, Gordon J, Ghez C (1996) Learning of scaling factors and reference axes for reaching movements. *Neuroreport* 7:2357–2361.
- Riout-Pedotti MS, Friedman D, Hess G, Donoghue JP (1998) Strengthening of horizontal cortical connections following skill learning. *Nat Neurosci* 1:230–234.
- Riout-Pedotti MS, Friedman D, Donoghue JP (2000) Learning-induced LTP in neocortex. *Science* 290:533–536.
- Rokni U, Richardson AP, Bizzi E, Seung HS (2007) Motor learning with unstable neural representation. *Neuron* 54:653–666.
- Sanes JN, Donoghue JP (2000) Plasticity and primary motor cortex. *Annu Rev Neurosci* 23:393–415.
- Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. *Exp Brain Res* 185:359–381.
- Shadmehr R, Thoroughman K (1998) Learning and memory formation of arm movements. In: *Biomechanics and neural control of movement* (Winters JM, Crago PE, eds), pp 347–353. Berlin: Springer.
- Shalden MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Stein RB, Gossen ER, Jones KE (2005) Neuronal variability: noise or part of the signal? *Nat Rev Neurosci* 6:389–397.
- Sutton RL, Barto G (1998) Reinforcement learning: an introduction. Cambridge, MA: MIT.
- Tanji J, Mushiaki H (1996) Comparison of neuronal activity in the supplementary motor area and primary motor cortex. *Brain Res Cogn Brain Res* 3:143–150.
- Todorov E (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17:1084–1108.
- Ungerleider LG, Doyon J, Karni A (2002) Imaging brain plasticity during motor skill learning. *Neurobiol Mem Learn* 78:553–564.
- van Beers RJ, Haggard P, Wolpert DM (2004) The role of execution noise in movement variability. *J Neurophysiol* 91:1050–1063.
- Vogels R, Spileers W, Orban GA (1989) The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 77:432–436.
- 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.
- Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269:1880–1882.
- 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.
- Zacksenhouse M, Lebedev MA, Carmena JM, O'Doherty JE, Henriquez CS, Nicolelis MA (2007) Cortical modulations increase in early sessions with brain-machine interface. *PLoS One* 2:e619.