

Improvements in the Signal-to-Noise Ratio of Motor Cortex Cells Distinguish Early versus Late Phases of Motor Skill Learning

William J. Kargo and Douglas A. Nitz

The Neurosciences Institute, San Diego, California 92121

There are numerous experience-driven changes in cortical circuitry that correlate with improved performance. Improved motor performance on a reach-to-grasp task in rodents is associated with changes in long-term potentiation (LTP), synaptogenesis, and movement representations in primary motor cortex (M1) by training days 3, 7, and 10, respectively. We recorded single-cell activity patterns in M1 during reach-to-grasp training to test how neural-spiking properties change with respect to LTP, synaptogenesis, and motor map changes. We also tested how neural-spiking changes relate directly to improved performance by monitoring muscle activity patterns. We found that signal-to-noise ratios (SNRs) of M1 spiking were significantly improved with practice but only after 7–12 d. Three sources of noise were assessed: signal-dependent noise exemplified by the slope of the relationship between mean spike count and count variance per burst, signal-independent noise exemplified by the offset of this relationship, and background firing rates before and after bursts. Signal-independent noise and pre-burst firing rates were reduced with practice. Early performance gains (days 1–6) were dissociated from SNR improvements, whereas later performance gains (day 7–12) were related directly to the magnitude of improvement in both muscle recruitment reliability and success rates. With training, an increased number of cells exhibited firing rates that were correlated with muscle recruitment patterns, with lags suggesting a primary direction of influence from M1 to muscles. These results suggest a functional linkage from local synaptogenesis in M1 to improved spiking reliability of M1 cells to more reliable recruitment of muscles and finally to improved behavioral performance.

Key words: motor learning; motor cortex; EMG; variability; plasticity; reliability

Introduction

The simplest type of motor learning entails practice-related performance gains (Sanes, 2003). Practice-related gains appear to depend on functional changes in primary motor cortex (M1) (Karni et al., 1998; Laubach et al., 2000; Sanes and Donoghue, 2000; Kargo and Nitz, 2003). In human subjects, repetitive transcranial magnetic stimulation of M1 immediately after skill practice disrupts the retention of behavioral improvements (Muellbacher et al., 2002). Extended skill training with a specific body part increases the area in M1 from which movements of that part can be evoked (Nudo et al., 1996; Kleim et al., 1998; Pearce et al., 2000) and produces changes in the topography and volume of M1 that are activated during performance (Karni et al., 1995, 1998). Manipulations that prevent such map changes diminish performance gains (Conner et al., 2003; Kleim et al., 2003). M1 also contains “memory cells” that encode novel (learned) mappings between limb motion and required muscle forces (Gandolfo et

al., 2000; Li et al., 2002; Paz et al., 2003). Functionally related M1 cells also tend to fire more synchronously with training (Schieber, 2002). A major challenge is to integrate these different levels of physiological analysis and to understand better the underlying neural mechanisms.

In a rodent model of motor skill learning (reach-to-grasp task) (Whishaw and Pellis, 1990), there is long-term potentiation (LTP) of layer II–III synapses, increased dendritic growth and synapse formation, and reorganized movement representations that develop in M1 over 1–2 weeks of daily training (Greenough et al., 1985; Kleim et al., 1998, 2002, 2003; Rioult-Pedotti et al., 1998; Bury and Jones, 2002; Conner et al., 2003). It is unclear whether and how these neural changes, which are local to M1, affect both spiking patterns in M1 cells and muscle activation patterns during skilled performance. Linking these levels of analysis might be addressed on the basis of findings that the observed neural changes exhibit different time courses: motor maps change significantly after 8–10 d (Kleim et al., 2004), synaptogenesis increases significantly after 4–7 d (Kleim et al., 2004), and LTP develops by days 3–5 (Rioult-Pedotti et al., 1998, 2000). In the present study, we first tested whether a particular feature of M1 spiking is changed with a time course similar to one of the above neural changes. We next tested how this change in M1 spiking affects muscle recruitment patterns during the task.

To address the above two questions, we recorded single-cell

Received Feb. 17, 2004; revised April 12, 2004; accepted May 3, 2004.

This work was supported by the Neurosciences Research Foundation. W.J.K. is the Clayton Fellow in Motor Control. We are grateful to M. Kuoppamäki for assistance with surgeries, training, and electrode fabrication and to G. Davis who performed histology.

Correspondence should be addressed to William J. Kargo, Neurosciences Institute, Neuroscience, 10640 John Jay Hopkins Drive, San Diego, CA 92121. E-mail: kargo@nsi.edu.

DOI:10.1523/JNEUROSCI.0562-04.2004

Copyright © 2004 Society for Neuroscience 0270-6474/04/245560-10\$15.00/0

activity patterns of M1 cells and the activity patterns of forelimb muscles during reach-to-grasp training in rodents. We found that signal-to-noise ratios (SNRs) of M1 cell firing were significantly improved with practice in the form of reduced spiking variability within task-related bursts and reduced pre-burst (and post-burst) firing rates. The time course of SNR improvement was correlated best temporally with that of increased synaptogenesis. Late performance gains were related directly to the magnitude of SNR improvement and associated functionally with reductions in muscle recruitment variability.

Materials and Methods

Surgery. All procedures were done in accordance with Institutional Animal Care and Use Committee protocol. A total of nine male Sprague Dawley rats (~1 year old; weighing between 225 and 260 gm) were used in this study. Rats were anesthetized with isoflurane. Incisions were made on the forelimb, and specific muscles were implanted with a pair of EMG wires [single-stranded, Teflon-coated, stainless-steel wire (A-M System, Carlsborg, WA); “ball” electrodes]. Muscles included trapezius, pectoralis major, deltoid, teres major, lateral and medial triceps, biceps brachii, brachialis, flexor carpi ulnaris and radialis, and extensor carpi ulnaris and radialis. EMG wires were tunneled under the skin to a scalp incision and soldered into pin connectors. Fascia on the skullcap was removed, and the skullcap was dried with etchant (3M; purchased through Henry Schein). A small window over M1 was drilled through the skull. Dura covering the exposed M1 was removed. An array of stereotrodes (25 μ m Teflon-coated tungsten wires twisted together; gold-plated tips with 0.1–1 M Ω resistance at the tip) were stereotactically placed over M1, 2.5 mm lateral and 1.0 mm rostral to bregma. The array was lowered 0.5 mm into M1, the exposed window was filled with agarose (0.9% in water), and the array was then lowered slowly an additional 1.1 mm to presumed layer 5. The array and EMG connectors were secured to the skull using bone screws and dental cement. The array was not a movable drive.

Correct placement of the array into forelimb M1 was assessed at the end of training. Animals were anesthetized with ketamine (50 mg/kg) and xylazine (10 mg/kg). A stimulus train (300 msec in duration; bipolar pulses, 0.5 msec duration; 20–60 μ A; 200 Hz) was delivered through each wire, and evoked forelimb responses were observed. Animals were deeply anesthetized after verification of forelimb responses and perfused with 4% paraformaldehyde. Fixed brain tissue showed that the depth and site of electrode tips were in layer V of forelimb M1 (forelimb M1 was ascertained by the observation of microstimulation-evoked forelimb movements).

Data acquisition during reach training. The experimental setup has been described thoroughly in Kargo and Nitz (2003) and similar to that developed by Whishaw and Pellis (1990). Reach-to-grasp training was initiated 5–7 d after surgery. Head connectors (one carrying EMG signals and a second carrying signals from stereotrode wires in M1) were secured to an amplifier deck via a tethered wire system. Each tether had a preamplifier array (Field Effect Transistor; NB Labs, Denison, TX) to increase the signal-to-noise ratio. EMG signals were filtered (0.1–3 kHz) and digitized at 3 kHz and stored on hard disk using custom-written software [Matt Wilson and Loren Frank, Massachusetts Institute of Technology (MIT)]. Signals from individual stereotrode wires were filtered (0.6–6 kHz), and amplitude thresholds for triggering data collection (± 5 msec) were individually set for each wire. An amplitude-crossing event on one wire triggered data collection from the stereotrode pair.

A digital video camera was placed 15 feet away and perpendicular to the long axis of the reach box. The camera recorded data at 60 fields per second. Videotaped reaches were analyzed off-line by transferring digital video to hard disk and importing video frames into Matlab. A piezoelectric strip (Measurement Specialties) was placed on the tray, under the pellet. Voltage levels were amplified, stored on hard disk at 1–3 kHz, and synchronized to EMG data. The piezoelectric strip was extremely sensitive to small forces and therefore monitored precisely the time at which the animal's hand contacted the pellet.

On training days, animals were placed into the reach box. A single pellet was placed on the outside tray. A successful trial was one in which

the rat grasped, retrieved, and brought the pellet to the mouth. Failures were reaches that missed, failed to grasp, or dropped the pellet. After each trial, rats were given 1 min before another pellet was placed on the tray. This prevented sequential reaching movements. Each rat was given ~70–80 pellets per day (one session per day) for 12 d. Because of this training pattern and learning over time, animals reached more on early days (e.g., 200 reaches on day 1 when 125 were failed attempts and 75 were successful attempts) compared with later days (e.g., 110 reaches on day 12 when 77 were successful attempts). The number of consumed pellets was nearly constant, whereas the number of attempts and the time spent in the reach box were reduced with training.

EMG analysis. EMGs were analyzed off-line using Matlab (Mathworks, Natick, MA). Data for each day (~30–60 min) were parsed to provide 2 sec of EMG data before and after tray contact for each reaching attempt. Parsed EMG records were then rectified and filtered with a 40-point, tapered boxcar filter (Basmajian and DeLuca, 1985).

We calculated pattern similarity between EMG patterns on individual trials during the day and the mean EMG pattern for the day. To calculate pattern similarity, the 12-dimensional muscle vector at each time point (± 2 sec relative to pellet contact) was normalized to a unit vector with magnitude = 1.0. We normalized the EMG vector because we were interested in how similar the balance (or ratio) of muscle activity was at a particular time point during the reach and how this value changed with practice. The dot product was calculated for time-matched unit vectors (e.g., EMG vector at 50 msec for trial 1 \cdot mean EMG vector at 50 msec, where \cdot refers to the dot product operator). A dot product of 1.0 represents perfect similarity between vectors, and -1.0 represents oppositely pointing, or most dissimilar, vectors. In addition to comparing muscle balance across time, we determined the magnitude of the (non-normalized) EMG vector across time and compared magnitudes between failures and successes within each day and across days. Because EMG signals may change differentially across days (e.g., the signal may improve in one muscle and degrade in another), we did not determine pattern similarity across days but only within a day. Also, we did not compare EMG vector magnitudes across days but only within a day.

We calculated a value that we termed muscle signal-to-noise ratio by the following. For individual muscles, we computed the following: $EMG^{mean} / \Sigma (EMG^i - EMG^{mean})$, where i is the trial number and $mean$ is the EMG mean for the day. The set of 10–12 values for individual muscles was then averaged for the day to produce one mean value. This mean value was determined on each training day for each animal.

Analysis of stereotrode recordings. M1 cells were distinguished off-line and based primarily on the relative amplitudes of spikes recorded on the stereotrode wires (McNaughton et al., 1983; Kargo and Nitz, 2003). As many as five to six cells per stereotrode were identified and isolated with this cluster-cutting technique. This method involves extraction of a set of spike waveform parameters (mainly spike height and width) for each spike from the two stereotrode channels and the separation of units on the basis of these parameters using interactive graphics software (Matt Wilson, MIT). Different combinations of parameter pairs were projected as two-dimensional scatter plots. When this was performed, points derived from single cells tended to form recognizable clusters (Fig. 1). The spikes within a cluster were enclosed in a polygon drawn using the computer mouse. The data points were then projected into new two-dimensional plots in which the earlier partitions of the data were preserved by color coding the points lying within the polygon boundaries. This process was performed until a multidimensional set of boundaries was established that provided the subjectively best separation of spike waveform clusters. The times of spikes in each cluster (or for each cell) were exported to Matlab for analysis.

In Matlab, peri-event histograms (PETHs) of cell spike times were computed relative to the instant of pellet contact (± 2 sec). All reach trials (e.g., 100 trials) were aligned at the instant of pellet contact. PETHs were computed in 20 msec bins and expressed as a rate (spikes per second). A three-point boxcar filter was used to smooth PETHs. Most M1 cells fired a single, primary burst of spikes during the task that was observed as a monophasic waveform in the PETH (Fig. 2A). We determined the mean onset and offset time (duration), peak amplitude, and integrated area of bursts for each reaching day. We also counted the number of spikes

within the burst window for each trial. The burst onset and offset times were determined as the times when the PETH initially rose to $>30\%$ of the baseline firing rate (baseline rates were determined in quietly resting, but not sleeping, animals that were in a holding pot before transfer to the reach chamber) and when the PETH subsequently fell below 30% of the baseline rate (Fig. 2A, dashed lines represent onset and offset times of the burst). We determined the count mean and count variance within the burst period for each day.

Linear cross-correlations between PETHs and EMG activations were performed using the Matlab system identification toolbox m-file *cra(Z,M)*, where $Z = [y\ u]$, u is the firing rate vector for one cell (e.g., firing rate for 100 trials, each ± 2 sec relative to pellet contact, are aligned head to tail), y is the rectified EMG vector, and M is the number of lags (± 500 msec) for which the cross-correlation was computed. Positive lags correspond to an influence from u to y , but do not necessarily imply causality. To remove spurious correlations caused simply by covariations in firing rate and EMG from trial to trial, we used a shift-predictor correction (Miller et al., 1993; Kargo and Nitz, 2003). The shift predictor is constructed by taking the firing rate and EMG vectors and shifting one of them a certain distance in time, so that the firing rate of one trial corresponds to a different trial of the EMG pattern. This shift was performed 30 times by (*rand**4.0) sec amounts, where *rand* is a random integer from 2 to the number of reach trials and 4.0 is the length of one trial. The cross-correlogram was recomputed using each shifted pair, averaged for the 30 randomly shifted pairs and then subtracted from the original cross-correlogram. Significant cross-correlations (at $p < 0.05$) were determined by testing whether the peaks or valleys of the original cross-correlation were ± 2 SD outside the mean shift predicted-correlation in the functionally important time window of $+150$ msec.

Results

During the reach-to-grasp task, many M1 cells exhibited discrete bursts of increased activity. The duration of these bursts differed between cells and was defined as the period between the onset and offset of firing $>30\%$ of the baseline rate (Fig. 2A) (see Materials and Methods). In the following sections, we examine (1) how spiking patterns within task-related bursts changed with practice, (2) how spiking patterns outside of these bursts changed with practice, and (3) how changes in spiking patterns were related to changes in muscle recruitment patterns.

Practice-related changes in intraburst spiking reliability

M1 spiking within task-related bursts became more reliable with training. To quantify spiking reliability, we determined the mean spike count μ_N within a burst and the SD in the spike count σ_N for all cells on each training day (in nine animals). σ_N and μ_N were normalized by the burst duration of the cell (e.g., $\sigma_N^2 = \sigma_N^2/300$ msec). The mean burst duration was 320 msec ± 32 SD on day 1 and was significantly reduced to 258 msec ± 33 SD ($p < 0.05$) on day 10. The linear log-log relationship between μ_N and σ_N was shifted downward with training (Fig. 2B, black dots) [day 1 ($r^2 = 0.64$; significant linear fit at $p < 0.01$); gray squares, day 10 ($r^2 = 0.56$; significant linear fit at $p < 0.01$)]. There was a significant difference in the y -intercept values of the day 1 and day 10 relationships using Tukey's honestly significant difference criterion ($p < 0.02$; day 1 intercept, -0.32 ; day 10, -0.63), i.e., a twofold reduction. There was no significant difference, however,

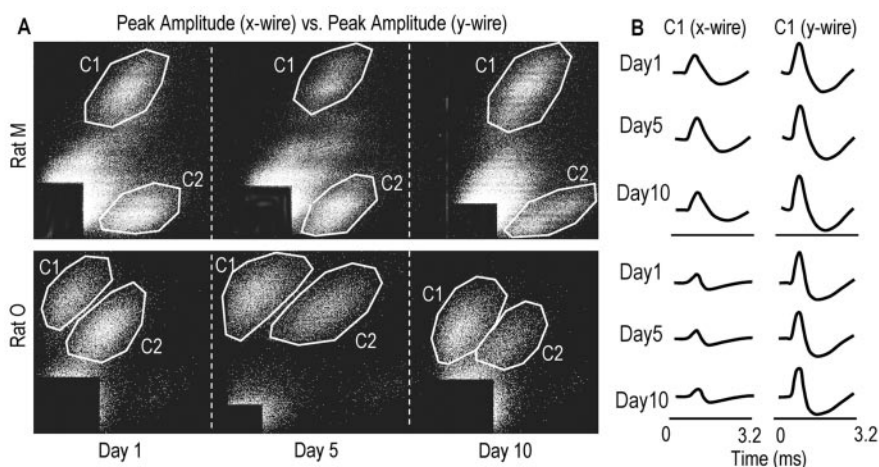


Figure 1. M1 cells were distinguished across days on the basis of relative amplitudes and waveforms of the spikes recorded on the stereotrode wires. *A*, Scatter plots of spike amplitudes on the stereotrode wires across days (top 3 panels, data for rat M on days 1, 5, and 10; bottom panels, data for rat O). Points formed recognizable clusters across days (C1, cluster 1; C2, cluster 2). *B*, Waveforms for spikes belonging to cluster C1 in rat M (top 3 rows, days 1, 5, and 10, respectively) and cluster C1 in rat O (bottom rows). On certain stereotrodes, spike waveforms and relative spike amplitudes were constant across the training period, as seen here.




Figure 2. M1 cells fired more reliably with training. *A*, Peri-event rasters are shown for one cell (each row is one trial). Burst width (between vertical lines) was the period between the onset and offset of firing $>30\%$ of the baseline rate; $p(F)$ is the probability of firing in 20 msec bins. *B*, The relationship between the mean spike rate and spike variance is shown [black dots, day 1 ($n_1 = 101$ cells); gray squares, day 10 ($n_{10} = 73$ cells)]. Y -intercepts were significantly different ($p < 0.02$). *C*, Twenty cells were tracked for the entire training period (in 9 animals). Black vectors represent the normalized values for mean spike rate and variance (origin, day 1 values; vector tips, day 10 values; values are log transformed). The solid gray line represents the fit for the linear relationship between mean rate and variance on day 1 [$\log(\text{variance}) = 0.94 \cdot \log(\text{mean rate})$], and dashed gray lines represent 95% confidence intervals. Fourteen of 20 cells followed a trajectory during learning that was not directed along this day 1 operating characteristic. *D*, The relationship between mean rate and variance within locomotor-related bursts of activity did not change during reach training [black dots, days 1–2 ($n_1 = 35$ cells); gray boxes, days 8–10 ($n_{10} = 38$ cells); 5 animals].




Figure 3. Increased spiking reliability was dissociated from early performance gains and associated with late performance gains. *A*, Mean success rates from days 1–12 are shown (black lines, mean for 6 animals with success rates >70% on day 12; gray lines, mean for 3 animals that had success rates <70%; vertical bars represent ± 1 SD). *S* and *S'* represent the first day in which the better and worse reaching groups, respectively, achieved significantly increased success rates relative to day 1 values ($p < 0.05$). Asterisks mark days in which success rates between groups were significantly different ($p < 0.05$). *B*, Mean SNR ratios of cells ($1/\text{Fano}$ factor; see Results) in the two groups are shown. Individual values for each animal were normalized to the animal's mean SNR on day 1. *M* and *M'* represent the first day in which the better and worse reaching groups achieved significantly improved M1 SNRs relative to day 1 values ($p < 0.05$; vertical bars are ± 1 SEM). Asterisks mark days in which M1 SNRs were different between groups ($p < 0.05$). *C*, Mean EMG signal-to-noise ratios (summed mean/summed variance; see Results) are shown. Values were normalized on the basis of the day 1 mean. *E* and *E'* represent the first day in which the better and worse reaching groups achieved significantly improved EMG SNRs relative to day 1 values ($p < 0.05$; vertical bars are ± 1 SEM). Asterisks mark days in which EMG SNRs were different between groups.

in the slopes (day 1, 0.94; day 10, 0.89) or σ_N with training (day 1, 7.08 ± 3.80 SD; day 10, 5.50 ± 3.98 SD). The mean spike count per burst or firing rate (μ_N) was nearly doubled and significantly increased with training ($\mu_N = 16.15 \pm 16.04$ Hz SD, day 1; $\mu_N = 31.03 \pm 28.17$ Hz SD, day 10; $p < 0.05$). Thus, after consecutive days of skill training, recruited M1 cells exhibited increased burst-related firing rates. Signal-dependent noise was not changed as inferred from the unchanged slope of the relationship between μ_N and σ_N ; however, a signal-independent noise source was reduced with training as inferred from the shift in the y -intercept of the relationship between μ_N and σ_N .

We recorded from a subset of neurons in the nine animals ($n = 20$ U), each of which appeared to have similar physiological identifying characteristics for 9–12 d (see Materials and Methods), which suggested that we potentially recorded from the same 20 cells across the entire examination period. Like the M1 population, these individual cells exhibited a shifted relationship between μ_N and σ_N with training. We normalized μ_N and σ_N for each cell and training session on the basis of its day 1 values (e.g., $\mu_N^i = \mu_N^i - \mu_N^1$, where i refers to the day number and 1 is day one). Figure 2C shows 20 vectors in which the origin represents the day 1 values for μ_N and σ_N and the vector tips represent the day 10 values (values are log transformed). All cells exhibited an increased μ_N with training (day 1, $\mu_N = 11.4 \pm 2.5$ Hz SD; day 10, $\mu_N = 33.5 \pm 2.1$ Hz SD), whereas 19 of 20 cells showed an increased σ_N with training (day 1, $\sigma_N = 6.3 \pm 4.2$ Hz SD; day 10, $\sigma_N = 7.0 \pm 3.8$ Hz SD). The magnitude of the σ_N increase, however, was significantly less than predicted in 14 of 20 cells on the basis of the initial relationship between μ_N and σ_N for the M1 population (Fig. 2B, dashed gray lines represent 95% confidence intervals for the day 1 linear fit). These data indicate that individual cells fired spikes at higher rates during the reach task but with less than predicted increases in trial-to-trial spiking variability.

The shift in cell firing properties was task specific and did affect firing patterns during other already learned behaviors. We

recorded the firing patterns of M1 cells during locomotion in five animals before and after each reaching session. We measured μ_N and σ_N for locomotor-related bursts of activity (during over-ground running). We examined bursts that occurred when animals locomoted at approximately similar, constant speeds across days (0.02–0.03 m/sec; position was tracked by an infrared light-emitting diode on the animals' heads and an infrared tracker system, and speed was calculated accordingly). The linear fit between μ_N and σ_N on early days (days 1 or 2; $n = 35$ cells) was not significantly different from the fit on later days during the training period (days 8–10; $n = 38$ cells), i.e., there was no significant difference in y -intercepts, slopes, and mean μ_N and σ_N across days (day 1–2, y -intercept = -0.88 ; slope = 1.08 , $\mu_N = 22.11 \pm 18.37$ Hz SD; $\sigma_N = 0.5 \pm 0.3$ Hz SD; day 8–10, y -intercept = -1.01 , slope = 1.14 , $\mu_N = 24.25 \pm 17.20$ Hz SD; $\sigma_N = 0.6 \pm 0.3$ Hz SD). Thus, the same population of M1 cells that exhibited changed firing properties across the reach training period exhibited stable firing properties during a well learned

behavior.

There was no significant difference in the slopes of the locomotor and reaching relationships between μ_N and σ_N (for either day 1 or day 10); however, there was a large significant difference in the y -intercepts between locomotor and day 1 reach data ($p < 0.001$). This difference became reduced with training (from 0.56 on day 1 to 0.25 on day 10) but was still significantly different at $p < 0.05$. With training, the day 10 relationship between μ_N and σ_N approached, but did not quite achieve, the relationship observed for these same cells during locomotion. In summary, the recruited M1 population did not exhibit a shift in the relationship between μ_N and σ_N during locomotor bursts across the training period, but these same cells exhibited a shift in the relationship between μ_N and σ_N during learning of the reach-to-grasp task. We show below (Training-related changes in muscle recruitment) that this task-specific shift in M1 firing properties was causally related to a shift in muscle recruitment patterns and thus to a fundamental change in motor behavior.

We examined more closely the time course over which the relationship between μ_N and σ_N was shifted. Figure 3B shows the mean values for μ_N/σ_N [i.e., the inverse Fano factor or signal (μ_N) to noise (σ_N) ratio] over days 1–12 for the recorded M1 population. The gray lines represent combined data for six animals that achieved success rates well above 70% by day 12 (Fig. 3A shows changes in success rate; animals were 72, 73, 75, 80, 81, and 85% on day 12); the black lines represent combined data for three animals that had success rates well below 70% on day 12 (animals were 60, 63, and 66%). Our target performance criterion was 70% because a recent study showed that a large group of trained rats ($n = 17$) obtained a mean success rate of 70% after 12 d of training (Conner et al., 2003). Thus, in the present study and for illustrative purposes, animals with success rates >70% were defined as belonging to a “better reaching” subgroup (i.e., exhibited above-mean performance levels), and animals with success rates <70% were defined as belonging to a “worse reaching” subgroup

(i.e., exhibited below-mean performance levels). There was a significant increase in success rates by day 5 relative to day 1 in the better subgroup and by day 6 relative to day 1 in the worse reaching subgroup ($p < 0.05$). There was no significant difference between group success rates, however, up until day 10. In both groups, the early improvements in success rate occurred without any improvement in M1 SNRs. M1 SNRs were not significantly improved relative to day 1 values until day 7 ($p < 0.05$) in the better reaching group and not until day 9 in the less successful group (M and M' mark the first days of significant improvement in M1 SNRs relative to day 1). M1 SNRs were gradually improved over days 7–12 in parallel with an additional 15–30% improvement in success rate. The change in M1 SNRs was directly associated with later performance gains (on days 7–12) because the better reaching group had significantly increased M1 SNRs relative to the less successful group (significant differences between groups on days 8–12; $p < 0.05$) and achieved higher success rates (significant differences between groups on days 10–12; $p < 0.05$).

Practice-related changes in extra-burst spiking activity

M1 cells showed spiking pattern changes outside of the individual bursts of activity. In particular, pre-burst firing rates were significantly reduced with training. The pre-burst duration was defined as the period from 300 msec before burst onset up to burst onset and was kept constant across days. For the population of all recorded M1 cells, there was a significant reduction in pre-burst firing rates by more than one-half between days 1 and 12 (day 1, 9.2 ± 4.2 Hz SD; day 12, 4.2 ± 3.8 Hz SD; $p < 0.01$). Figure 4A–D shows examples of four cells that were tracked over the 12 day training period. It is important to note that cells showing reductions in pre-burst firing rates discharged bursts at different time points during the reach, i.e., the reduction in pre-burst firing did not occur solely during the anticipatory-preparatory phase before reach onset but occurred throughout the reach (Fig. 4, the first vertical dashed line represents the approximate time of movement onset, whereas the second vertical line represents the time of pellet contact in which the trials were aligned). The first day in which pre-burst firing rates became significantly reduced relative to day 1 values was day 8 in the better reaching group (4.8 ± 4.0 Hz; $p < 0.05$) and day 9 in the worse reaching group (5.1 ± 3.1 Hz; $p < 0.05$). Thus, the reduction in pre-burst rates had a time course similar to that found for improved spiking reliability within bursts.

Firing rates after task-related bursts (for a period 300 msec after) tended to be reduced with practice, but this reduction was not significant by day 12 for the M1 population (day 1, 7.32 ± 6.01 SD; day 12, 5.11 ± 3.82 SD). Only 5 of 20 cells that were tracked across the training period showed significant reductions in post-burst firing rates (significant difference in means at $p < 0.05$ for these individual cells). The effect of significantly reducing pre-burst firing rates in the majority of cells and reducing post-




Figure 4. Baseline firing rates before and after task-dependent bursts were reduced with practice. A–D, Examples of four cells that were tracked over a 10 d training period. The left column of each panel shows data on day 1, and the right column shows data on day 10. The top row of each panel shows surface plots of firing rate across time (each row is one trial; the highest rate, white, = 40 Hz; black = 0 Hz). The bottom row of each panel shows averaged peri-event histograms of firing rate (bin size = 20 msec; vertical bars, ± 1 SEM).

burst firing rates in some cells was to increase the differentiation between the task-related burst and background firing before and after the burst, even in cells that exhibited similar peak firing rates across the training period (Fig. 4A,D). Thus, in one respect, the signal (the burst) to noise (background firing) was enhanced with training. Also, as shown in the previous section, the signal within a burst (mean spike count) was enhanced relative to the noise within a burst (spike count variance).

Training-related changes in muscle recruitment

Time-dependent improvements in M1 SNRs were similar to those observed in forelimb EMG signals. Muscle SNRs were computed as described in Materials and Methods and essentially represent the summed EMG mean divided by the summed EMG variance (averaged for the set of 10–12 recorded muscles). The mean muscle SNRs on each day are shown in Figure 3C. The gray lines represent data for six animals that achieved success rates $>70\%$ by day 12; the black lines represent data for three animals that had success rates less than our target performance criterion of 70%. Similar to M1 SNRs, muscle SNRs were not significantly improved relative to day 1 values until days 8 and 10 in the better and worse reaching groups, respectively (marked with E and E'), i.e., there was a dissociation between early performance gains and changes in EMG (and M1) SNRs. There was a significant difference between EMG SNRs in the two groups on days 10–12 (asterisks represent between-group differences at $p < 0.05$). This difference was associated with a significant difference in success rates between the two groups.

We examined EMG patterns across days more thoroughly to address what distinct mechanisms for skill improvement occurred early and separately from later reductions in motor pattern variability. We computed the similarity between motor pat-

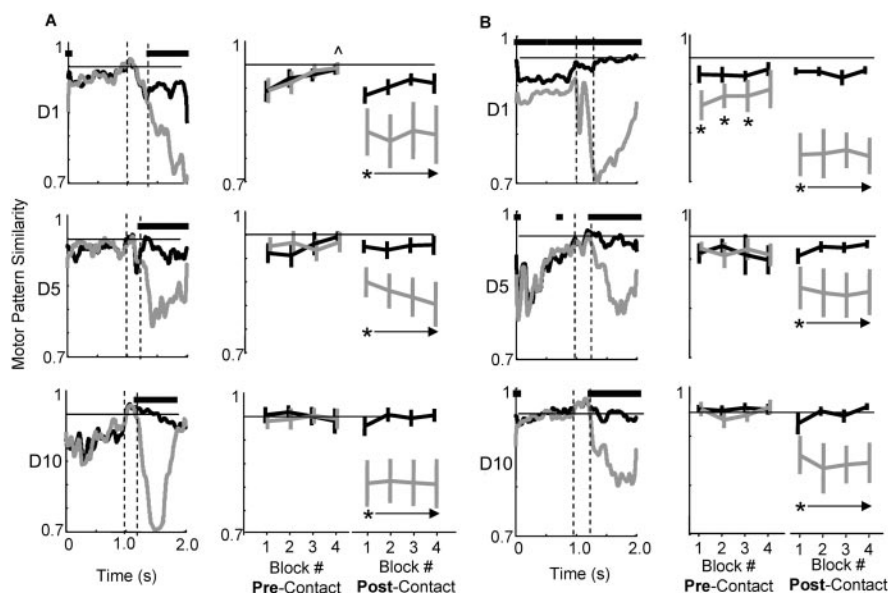


Figure 5. Practice-related changes in motor pattern similarity. *A* and *B* show data for two different animals. The first columns show how similar successful (black) and failed (gray) motor patterns were to the mean successful motor pattern for the last block of trials on that day (see Materials and Methods for dot product calculation). The first, second, and third rows are values for days 1, 5, and 10. Horizontal bars above each plot represent time points in which successful and failed patterns were significantly different (at $p < 0.05$). The two dashed vertical lines represent the approximate times of reach onset and pellet contact (pre-contact period = reach onset to pellet contact; post-contact period = pellet contact + 1.5 sec). The second two columns in each panel show how similar pre-contact and post-contact motor patterns were in each block of trials to the mean successful pattern of the last block (black, successes; gray, failures). In *A* (top row), last block values on day 1 were significantly different from first block values; the symbol * marks blocks that were significantly different from block 1 values. This gradual increase in similarity values within a day was not observed on days 5 and 10. In *B* (top row), pre-contact motor patterns for successful trials were significantly different from pre-contact patterns for failure trials; asterisks mark significant differences between success and failure blocks. This difference between pre-contact patterns was not observed on days 5 and 10. Similarity values in all blocks were increased on day 10 relative to days 1 and 5 (the horizontal black line = 0.95; this value was reached on day 10 but not on earlier days).

terns on individual trials and a template motor pattern. The training session for each day was divided into four blocks, e.g., if 120 attempts were made, there were four blocks of 30 trials. The mean (successful) motor pattern of block 4 was the template pattern to which we compared the motor patterns of individual trials. We used the block-4 pattern because mean success rates in block 4 were significantly increased on days 1 ($p < 0.01$) and 2 ($p < 0.05$) relative to block 1, i.e., the block-four pattern on early days was associated with more successful reaching. This was not the case on later days; the template pattern was not different from the mean pattern of earlier blocks. Figure 5, *A* and *B*, shows pattern similarity values for two animals on days 1, 5, and 10 (first, second, and third rows, respectively). The first columns show similarity values across time (black, successful trials; gray, failed trials). The second two columns in each panel show how similar each block of trials was to the template pattern during the pre-contact period (reach onset to pellet contact) and post-contact period (pellet contact plus 1.5 sec), respectively.

Similarity values exhibited three main changes across the training period. First, in all animals, motor patterns during the pre-contact period gradually became more similar (from trial to trial) to the template pattern across days. Motor patterns during the post-contact period also became more similar to the template with training, but primarily for successful trials (see black values in the last columns of each panel). Second, similarity values increased more abruptly across blocks on early training days. The animal in Figure 5*A* exhibits this clearly. On day 1, pre-contact motor patterns for the last block were significantly more similar

to the template pattern compared with how similar the first block was to the template pattern (significant difference in first and fourth block values at $p < 0.01$). This indicates that the animal adjusted its motor pattern over the day such that later motor patterns (or ratios of muscle activation) were different from the starting motor patterns. Third, successful (pre-contact) motor patterns tended to be significantly different from failure motor patterns on early training days. The animal in Figure 5*B* exhibited this clearly. Horizontal bars above the similarity values (in the first column on day 1) represent significant differences between motor patterns. Significant differences were observed only on day 1. This indicates that the animal switched between activating mainly one motor pattern for successful reaches and distinctly different patterns for failed reaches (Kargo and Nitz, 2003).

The three motor pattern changes exhibited different time courses. Time courses were discerned by quantifying the following parameters. First, the slope of the relationship between trial number and pattern similarity to the template was determined each day (just for successful trials). Positive slopes indicate that the starting pattern of muscle activation was changing over the day and becoming more similar to the template (final) pattern. Second, we quantified how many time points at which successful motor patterns were significantly different from failed motor patterns during the pre-contact phase, i.e., from 300 msec before deltoid EMG onset up to pellet contact. A large number of time points indicate that the animal activated different motor patterns during the pre-contact period, one associated mainly with successes and others associated with failed attempts. Third, we quantified the mean pattern similarity for the entire day. A large value would indicate that motor patterns were more similar to the template and thus less variable from trial to trial. Figure 6 summarizes these values across the training period for the nine animals (black, number of significantly different time points between failure and success patterns; dark gray, slope of trial versus similarity value within a day; light gray, mean similarity for the day). The values were normalized to compare their time courses (normalized to a maximum of 1.0 and a minimum of 0.0). The (colored) asterisks mark the first day in which the data were significantly different (at $p < 0.01$) from day 1 values. The number of different time points goes down significantly from 193 ± 181 SD on day 1 to 18 ± 17 SD on day 2. The slope of the relationship between trial number and pattern similarity within a day goes down significantly from 0.011 ± 0.007 SD to 0.007 ± 0.006 SD by day 4. Finally, the mean pattern similarity goes up significantly from 0.87 ± 0.04 SD on day 1 to 0.93 ± 0.4 SD on day 8. Importantly, improved M1 SNRs followed a time course most similar to increases in motor pattern similarity. Switching between the activation of different motor patterns occurred on day 1 only and (within-day) adaptation of motor patterns occurred across days 1–5. Thus, these early mechanisms were dissociated from improved M1 SNRs.




Figure 6. Practice-related motor pattern changes exhibited different time courses. Mean values ($n = 9$ animals) were normalized to compare their time courses (normalized to a maximum of 1.0 and a minimum of 0.0). The number of time points in which successful motor patterns were significantly different from failed motor patterns during the pre-contact period is shown in black. Successful and failed reaches were associated with the activation of different motor patterns mainly on the first day. The slope of the relationship between trial number and pattern similarity within each day is shown in dark gray. This measure represents how different motor patterns later in a given training day were from motor patterns earlier that day. Within-day adjustments occurred mainly on days 1–5. The mean similarity of pre-contact motor patterns to the block 4 mean is shown in light gray. The mean pattern similarity was significantly increased by day 8 and continued to increase. For each relationship, appropriately shaded asterisks mark the first day in which the data were significantly different from the day 1 values (at $p < 0.01$). Importantly, the values exhibited very different time courses.

Because improvements in M1 SNRs and increases in motor pattern similarity had similar time courses, it was difficult to ascertain what the primary direction of influence was from this data alone. To address this, we examined cross-correlograms between M1 firing rates and muscle EMGs. Twelve of 20 cells that had similar physiological identifying characteristics over the training period and thus were potentially the same cell from day to day had firing rates that were significantly, positively correlated with EMG waveforms of one to several muscles on day 1 [mean correlation strength = 0.21 ± 0.08 SD ($p < 0.05$); mean lag time of peak correlation = 75 msec ± 24 SD; mean number of muscles to which M1 cells had significant positive correlations = 3.2 msec ± 1.5 SD (10–12 muscles were recorded)]. On day 10, 14 of 20 cells had firing rates that had significant, positive cross-correlations with EMG waveforms (at $p < 0.05$; mean = 0.32 ± 0.09 SD; mean lag time of peak correlation = 84 msec ± 22 SD; number of muscles to which M1 cells had significant positive correlations = 3.8 msec ± 1.4 SD). Importantly, the positive lags suggest that the primary direction of influence was from M1 to muscles. These results are summarized in Figure 7. The top panels of Figure 7, *A* and *D*, show the firing rates of two different M1 cells, which were tracked over the entire training period, in two different animals [day 1, left panels; day 10, right panels; red represents highest activity (60 Hz) and dark blue represents lowest activity (0 Hz)]. The bottom panels show muscle activation patterns [Fig. 7*A*, flexor carpi ulnaris; Fig. 7*D*, trapezius; red represents highest activity (normalized to maximum of 1.0) and blue is lowest activity (0.0)]. Both M1 cells showed increased (mean) peak firing rates from day 1 to day 10 (30–62 and 18–45 Hz). Both cells also exhibited less than predicted increases in spike

count variability within task-related bursts (cell 1: day 1, mean rate = 21.77 ± 6.31 ; day 10, 39.81 ± 7.12 ; cell 2: day 1, 12.05 ± 5.31 ; day 10, 25.75 ± 5.80). The new relationships were associated with increased cross-correlation strengths between the M1 firing rates and the muscle activation patterns [correlations are shown in the insets to the right of firing rates; significant cross-correlations are greater than the dashed lines, which represent the 95% confidence interval of the shift-predicted correlations (see Materials and Methods)]. The cross-correlations had positive lags (89 msec in left panel and 50 msec in right panel) that correspond to a forward influence from the M1 to muscles but do not necessarily imply causality.

The percentage of all recorded cells that were recruited during the reach task and that had activity correlated with EMG patterns increased with training. The number of cells with significant positive cross-correlations with at least one EMG waveform during the reach task and within the functionally important time window of +150 msec increased from 37% (37 of 101 cells) on day 1 to 52% (38 of 73 cells) on day 10 (significant increase in the expected number of correlated cells on day 10 based on the day 1 percentages; $\chi^2 = 7.12$; $p < 0.01$). In addition to a significant increase in the number of correlated cells, the peak correlation strength of the population of positively correlated cells was significantly increased from 0.22 ± 0.11 SD (day 1; mean lag = 113 ± 41 msec SD) to 0.35 ± 0.12 SD at $p < 0.05$ (day 10; mean lag = 101 ± 34 msec SD). Thus, a larger percentage of M1 cells represented time-varying patterns of muscle activation during the task, and these cells represented EMG activations better with training. Most importantly, the cross-correlations were such that the lag times suggested a primary excitatory influence from M1 to muscles. Thus, increased firing reliability in M1 would translate into increased muscle activation reliability.

Discussion

Animals had improved SNRs of M1 cell firing during motor skill training in the form of increased spiking reliability within task-related bursts and reduced pre-burst firing rates. The time course of improvement was correlated best temporally with that of increased synaptogenesis in M1 as shown by Kleim et al. (2004). The late improvement in M1 SNRs was associated with more reliable recruitment of muscles. Animals with larger improvements exhibited more reliable recruitment of muscles and obtained higher success rates. Early performance gains were dissociated from M1 SNR changes and related to different learning mechanisms (Kargo and Nitz, 2003).

Animals reduced signal-independent noise in M1 firing with practice. The slope of the relationship between mean spike count and spike count variance (within task bursts) for the recruited population remained unchanged with practice. The slope indicates the level of signal-dependent noise in neural firing and was similar to values observed in other cortical areas (Lee et al., 1998; Shadlen and Newsome, 1998; Stevens and Zador, 1998). Unlike the slope, the y -intercept was modified with training. The y -intercept indicates the level of signal-independent noise in firing or noise that was reduced equally in task-related cells regardless of the mean firing rate during bursts. Signal-independent noise during the reach-to-grasp task was reduced to levels observed during locomotion, a stereotyped and well learned behavior. The reduction of signal-independent noise in M1 firing may be critical for skill improvement on the basis of findings that both signal-dependent and signal-independent noise in motor commands accounts for a large proportion of movement variability (Van Beers et al., 2004), and expert skill levels are associated with

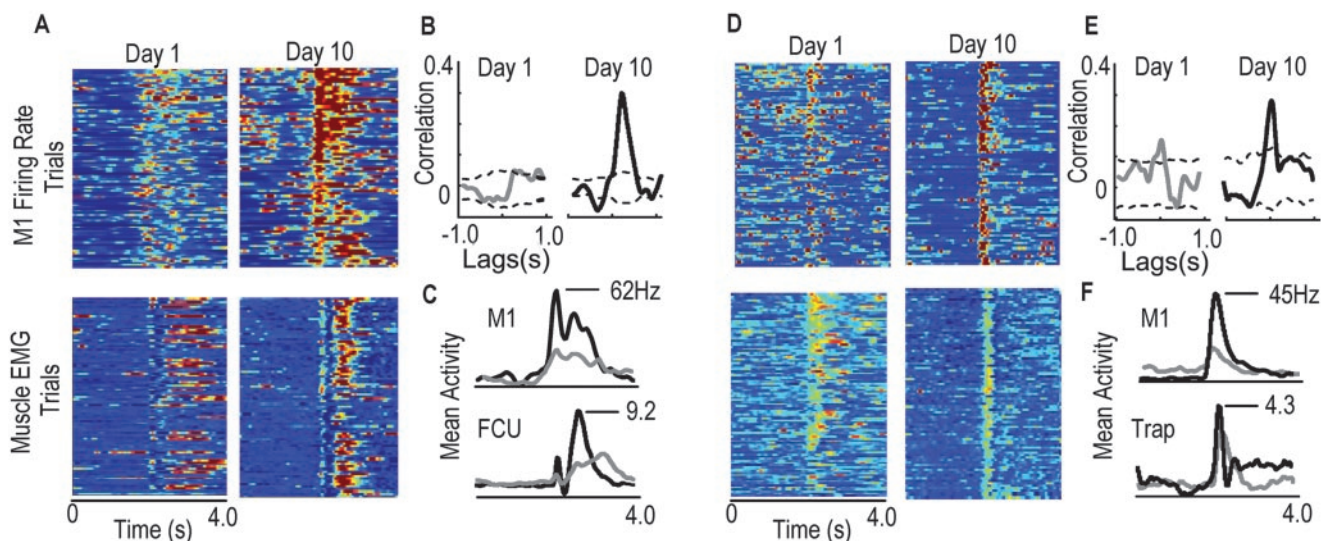


Figure 7. Evidence for a forward direction of influence and for transmission of enhanced reliability from M1 firing to muscle activation. *A*, Top panels show firing rates for an M1 cell (left, day 1 rates; right, day 10 rates; dark red = 60 Hz; dark blue = 0 Hz; each row corresponds to one trial). Activation patterns of the flexor carpi ulnaris muscle (FCU) are shown in the bottom two panels (left, day 1 pattern; right, day 10 pattern). *B*, The cross-correlation between the firing rate and FCU activation pattern of the cell was negligible on day 1 (gray line). The cross-correlation increased to significant levels on day 10 (solid black line); cross-correlations of 1.0 represent the highest degree of similarity at the appropriate time lag, and the dashed lines represent 2 SD from the mean shift-predicted correlation (see Materials and Methods). *C*, Both the mean firing rate (gray, day 1; black, day 10) and amplitudes of muscle activation (gray, day 1; black, day 10) were increased with training. *D–F*, Same as in *A–C* but for a different M1 cell and a proximal limb muscle (trapezius). Cross-correlations increased with training in parallel with increases in the firing rate and decreases in firing variability.

low levels of trial-to-trial variability (Young and Marteniuk, 1998; Gabriel, 2002).

Training-related improvements in M1 SNRs did not occur until days 7–9. Thus, the time course of M1 SNR changes was correlated best temporally with the delayed time course of increased synaptogenesis in M1 as shown in Kleim et al. (2004). Significant changes in forelimb motor maps are observed later, by day 10 (Kleim et al., 1998, 2004). Significant changes in intracortical (layer II–III evoked) LTP and long-term depression are observed earlier, by days 3–5 (Rioullet-Pedotti et al., 1998, 2000).

Despite the similar time courses, it is unclear how synaptogenesis is related to improved SNRs. New synapses might be GABAergic, glutamatergic, or neuromodulatory. This will constrain how spiking reliability is improved. One way to improve spiking reliability in response to repeated presentations of a stimulus is to reduce variability in the firing rate or membrane potential before the stimulus (Azouz and Gray, 1999). Many cells developed an enhanced pre-burst inhibition with training, which may be related to GABAergic synapse growth. The pre-burst inhibition developed with a similar time course as improved spiking reliability and acted to reduce M1 firing rates and firing variability before the task-related burst; however, we observed only a weak relationship between the magnitude of the developed pre-burst inhibition and the magnitude of improvement in spiking reliability. Thus, other synapse types like glutamate synapses might also be developing, because spiking reliability clearly depends on the number of glutamatergic synapses that drive spiking (Brenowitz and Trussell, 2001). Also, acetylcholine synapses may play an important role because selective lesions of acetylcholine cells in the nucleus basalis disrupt practice-related improvements on this task (Conner et al., 2003).

Increased spiking reliability in M1 cells was intimately linked to a reduction in muscle recruitment variability and enhancements in M1–EMG correlations. In the present study, we assessed task-related connectivity between M1 and muscles using cross-

correlograms. The finding that cross-correlations had positive lags from M1 firing to EMG (~70 msec) indicates a primary direction of influence from M1 to muscles during the task and an influence mainly through polysynaptic corticofugal pathways (Alstermark et al., 2004). With training, a significantly increased percentage of cells exhibited activity that was correlated with EMGs, and this neural activity was correlated more strongly. Pearson's correlation coefficient is designed to eliminate the problems of scale and unit selection in measuring correlation, so the scale of the variance (and its reduction with training) should not affect correlation strengths, provided that the distributions of firing rate and EMG vary normally and that spurious correlations attributable to trial-to-trial covariations are removed, e.g., with the shift-predictor correction (Miller et al., 1993); however, there are always interpretative challenges in using cross-correlograms. Cross-correlations do not necessarily reflect causality. It is unclear whether the M1 signal determines the EMG signal completely, in part, or not at all. Conceivably, M1 activity may reflect an efference copy of the upcoming motor pattern. Nevertheless, it is clear that the EMG signal did not determine M1 firing patterns because of the positive lag. Assuming that M1 determines to some extent the EMG signal, the mechanistic basis for the correlation increase between the firing of one cell and an EMG pattern is likely to reflect a population-level effect, e.g., increased signal strength, increased signal reliability in functionally related M1 cells, enhanced synchronous spiking (Schieber, 2002), or a combination of the three. Although not measured here, synchronous spiking is likely to increase with training because firing rates doubled on average and spiking reliability nearly tripled. Enhanced correlations might also be caused by plasticity downstream of M1 cells, e.g., in the strength of corticospinal synapses.

The present study cannot address whether learning-related changes in M1 activity occurred independently of performance-related changes because of the intimate linkage between performance change and learning in this task. Motor adaptation stud-

ies, which artificially dissociate movement kinematics from dynamics (Gandolfo et al., 2000; Li et al., 2002) or movement kinematics from visual feedback of the arm (Wise et al., 1998; Paz et al., 2003; Schwartz et al., 2004), have shown that M1 cells exhibit persistent, memory-like activity for the perturbation even after the perturbation was removed and performance returned to baseline. Because we examined a natural motor skill in which changes in motor performance define learning and changes are maintained with practice, the present study instead better addresses another issue: do cells that encode task parameters do so better with training? Ideally, one would use an information-theory analysis to show that with training M1 cells provide more information to decode some task parameter (Helms Tillery et al., 2003; Paz and Vaadia, 2004). Because we sampled a small subspace of movements and muscle activation patterns (i.e., learning of one behavior), it was difficult to compute mutual information between either cell activity and muscle activation or cell tuning curves. Nonetheless, two features of spiking were changed with training that could have contributed to improved encoding of muscle activation: response variability was reduced by three times in cells recruited during the task, and response rates were doubled. As noted above, these changes might have led to increased cross-correlation strengths between M1 firing and muscle activation. Importantly, the finding that increases in spiking reliability and M1–EMG correlations exhibited a similar time course as increased synaptogenesis in M1 (Kleim et al., 2004) suggests that learning-related changes within M1 potentially underlie the practice-related improvements in performance.

An examination of EMG patterns during the reach-to-grasp task indicates at least three mechanisms for skill improvement: (1) improved pattern selection, (2) motor pattern adaptation, and (3) reduction of pattern variability. Some animals switched between forelimb motor patterns at the very onset of training (or between forelimb and orally executed attempts, e.g., licking) and quickly learned to select the motor pattern with the higher probability of success more often (Kargo and Nitz, 2003). Switching was predicted by different patterns of neural activity in M1 (Kargo and Nitz, 2003) and was seldom observed after the first day. Despite the absence of switching on later days, M1 SNRs worsened in most animals on day 2. These high levels of spiking variability are presumably one reason why selected motor patterns still had low probabilities of early success (30–40%). To improve success further, animals used two additional strategies. Animals gradually adapted the ratios of muscle activation within training sessions over days 1–6. The finding that success rates usually increased within training sessions over days 1–6 indicates that the adaptations had functional importance. Finally, with extended training, animals reduced trial-to-trial variability in muscle recruitment patterns. Success rates were improved by 10–25% over the first 6 d during which primarily switching and adaptation were observed and another 10–25% over the subsequent 6 d during which mainly variability was reduced. The improvements of M1 SNRs paralleled and potentially caused the late reductions in motor pattern variability.

In summary, M1 spiking reliability was improved relatively late during skill training, was dissociated from early performance gains, and was associated directly with late performance gains. The time course of SNR improvement was most similar to that observed for increased synaptogenesis in M1. The functional effect of increased spiking reliability was a reduction in muscle recruitment variability, an enhancement of M1–EMG cross-correlations, and an increase in success rates.

References

- Altermark B, Ogawa J, Isa T (2004) Lack of monosynaptic cortico-motoneuronal EPSPs in rats: disynaptic EPSPs mediated via reticulospinal neurons and polysynaptic EPSPs mediated via segmental interneurons. *J Neurophysiol* 91:1832–1839.
- Azouz R, Gray CM (1999) Cellular mechanisms contributing to response variability of cortical neurons *in vivo*. *J Neurosci* 19:2209–2223.
- Basmajian JV, DeLuca CJ (1985) *Muscles alive: their functions revealed by electromyography*. Baltimore: Williams & Wilkins.
- Brenowitz S, Trussell LO (2001) Maturation of synaptic transmission at end-bulb synapses of the cochlear nucleus. *J Neurosci* 21:9487–9498.
- Bury SD, Jones TA (2002) Unilateral sensorimotor cortex lesions in adult rats facilitate motor skill learning with the “unaffected” forelimb and training-induced dendritic structural plasticity in the motor cortex. *J Neurosci* 22:8597–8606.
- Conner JM, Culbertson A, Packowski C, Chiba AA, Tuszynski MH (2003) Lesions of the basal forebrain cholinergic system impair task acquisition and abolish cortical plasticity associated with motor skill learning. *Neuron* 38:819–829.
- Gabriel DA (2002) Changes in kinematic and EMG variability while practicing a maximal performance task. *J Electromyogr Kinesiol* 12:407–412.
- 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 USA* 97:2259–2263.
- Greenough WT, Larson JR, Withers GS (1985) Effects of unilateral and bilateral training in a reaching task on dendritic branching of neurons in the rat motor-sensory forelimb cortex. *Behav Neural Biol* 44:301–314.
- Helms Tillery SI, Taylor DM, Schwartz AB (2003) Training in cortical control of neuroprosthetic devices improves signal extraction from small neuronal ensembles. *Rev Neurosci* 14:107–119.
- Kargo WJ, Nitz DA (2003) Early skill learning is expressed through selection and tuning of cortically represented muscle synergies. *J Neurosci* 23:11255–11269.
- Karni A, Meyer P, 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-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 USA* 95:861–868.
- Kleim JA, Barbay S, Nudo RJ (1998) Functional reorganization of the rat motor cortex following motor skill learning. *J Neurophysiol* 80:3321–3325.
- Kleim JA, Barbay S, Cooper NR, Hogg TM, Reidel CN, Remple MS, Nudo RJ (2002) Motor learning-dependent synaptogenesis is localized to functionally reorganized motor cortex. *Neurobiol Learn Mem* 77:63–77.
- Kleim JA, Bruneau R, Calder K, Pocock D, VandenBerg PM, MacDonald E, Monfils MH, Sutherland RJ, Nader K (2003) Functional organization of adult motor cortex is dependent on continued protein synthesis. *Neuron* 40:167–176.
- Kleim JA, Hogg TM, VandenBerg PM, Cooper NR, Bruneau R, Remple M (2004) Cortical synaptogenesis and motor map reorganization occur during late, but not early, phase of motor skill learning. *J Neurosci* 24:628–633.
- 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, 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 (2002) 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.
- McNaughton BL, O’Keefe J, Barnes CA (1983) The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci Methods* 8:391–397.
- Miller LE, van Kan PL, Sinkjaer T, Andersen T, Harris GD, Houk JC (1993) Correlation of primate red nucleus discharge with muscle activity during free-form arm movements. *J Physiol (Lond)* 469:213–243.
- Muellbacher W, Ziemann U, Wissel J, Dang N, Kofler M, Facchini S,

- Borojerdi B, Poewe W, Hallett M (2002) Early consolidation in human primary motor cortex. *Nature* 415:640–644.
- Nudo RJ, Milliken GW, Jenkins WM, Merzenich MM (1996) Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *J Neurosci* 16:785–807.
- 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* 2:264–274.
- 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.
- Pearce AJ, Thickbroom GW, Byrnes ML, Mastaglia FL (2000) Functional reorganisation of the corticomotor projection to the hand in skilled racquet players. *Exp Brain Res* 130:238–243.
- Rioul-Pedotti MS, Friedman D, Hess G, Donoghue JP (1998) Strengthening of horizontal cortical connections following skill learning. *Nat Neurosci* 1:230–234.
- Rioul-Pedotti MS, Friedman D, Donoghue JP (2000) Learning-induced LTP in neocortex. *Science* 290:533–536.
- Sanes JN (2003) Neocortical mechanisms in motor learning. *Curr Opin Neurobiol* 13:225–231.
- Sanes JN, Donoghue JP (2000) Plasticity and primary motor cortex. *Annu Rev Neurosci* 23:393–415.
- Schieber MH (2002) Training and synchrony in the motor system. *J Neurosci* 22:5277–5281.
- Schwartz AB, Moran DW, Reina GA (2004) Differential representation of perception and action in the frontal cortex. *Science* 303:380–383.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Stevens CF, Zador AM (1998) Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1:210–217.
- Van Beers RJ, Haggard P, Wolpert DM (2004) The role of execution noise in movement variability. *J Neurophysiol* 91:1050–1063.
- Whishaw IQ, Pellis SM (1990) The structure of skilled forelimb reaching in the rat: a proximally driven movement with a single distal rotatory component. *Behav Brain Res* 41:49–59.
- Wise SP, Moody SL, Blomstrom KJ, Mitz AR (1998) Changes in motor cortical activity during visuomotor adaptation. *Exp Brain Res* 121:285–299.
- Young RP, Marteniuk RG (1998) Stereotypic muscle-torque patterns are systematically adopted during acquisition of a multi-articular kicking task. *J Biomech* 31:809–816.