Published in final edited form as:

Exp Brain Res. 2006 February ; 169(1): 69–84.

Comparison of Population Activity in the Dorsal Premotor Cortex and Putamen During the Learning of Arbitrary Visuomotor Mappings

Ethan R. Buch, Peter J. Brasted, and Steven P. Wise

Section on Neurophysiology, Laboratory of Systems Neuroscience, National Institute of Mental Health, National Institute of Health, Bethesda, Maryland, USA

Abstract

A previous study found that as monkeys learned novel mappings between visual cues and responses, neuronal activity patterns evolved at approximately the same time in both the dorsal premotor cortex (PMd) and the putamen. Here we report that, in both regions, the population activity for novel mappings came to resemble that for familiar ones as learning progressed. Both regions showed activity differences on trials with correct responses versus those with incorrect ones. In addition to these common features, we observed two noteworthy differences between PMd and putamen activity during learning. After a response choice had been made, but prior to feedback about the correctness of that choice (reward or nonreward), the putamen showed a sustained activity increase in activity, whereas PMd did not. Also in the putamen, this prereward activity was highly selective for the specific visuomotor mapping that had just been performed, and this selectivity was maintained until the time of the reward. After performance reached an asymptote, the degree of this selectivity decreased markedly to the level typical for familiar visuomotor mappings. These findings support the hypothesis that neurons in the striatum play a pivotal role in associative learning.

Keywords

Conditional motor learning; Striatum; Basal ganglia; Frontal cortex; Associative learning

Introduction

Arbitrary visuomotor mapping, also known as conditional motor learning and conditional discrimination, requires the trial-and-error based acquisition of associations among antecedents and consequents. Antecedents can include visual and other sensory inputs, and consequents usually involve movements or the targets of movement. Changes in task-related neuronal activity during arbitrary visuomotor learning have been described for dorsal premotor cortex (PMd) (Mitz et al., 1991; Brasted and Wise, 2004), supplementary eye field (Chen and Wise, 1995a, b), frontal eye field (Chen and Wise, 1995b), prefrontal cortex (Asaad et al., 1998; Pasupathy and Miller, 2005), hippocampus (Cahusac et al., 1993; Wirth et al., 2003), striatum (Tremblay et al., 1998; Jog et al., 1999; Hadj-Bouziane and Boussaoud, 2003; Brasted and Wise, 2004; Pasupathy and Miller, 2005) and globus pallidus (Inase et al., 2001). Neuropsychological studies point to the same neural network (Petrides, 1985; Halsband and Passingham, 1985; Rupniak and Gaffan, 1987; Gaffan and Harrison, 1988; Canavan et al., 1989; Eacott and Gaffan, 1992; Murray and Wise, 1996; Ridley and Baker, 1997; Wang et al.,

2000; Bussey et al., 2001, 2002; Brasted et al., 2002, 2003; Barefoot et al., 2002, 2003; Nixon et al., 2004), as does the neuroimaging literature (Deiber et al., 1991, 1997; Paus et al., 1993; Toni and Passingham, 1999; Toni et al., 2001, 2002; Eliassen et al., 2003).

Although the network that participates in the learning and recall of arbitrary visuomotor mappings includes the regions mentioned above, little is known about how these structures differ in supporting such learning. To better understand the differential contributions of PMd and putamen to this form of associative learning, we have examined learning-related activity at the population level. Other results from this study have appeared previously (Brasted and Wise, 2004).

Materials and methods

Subjects and apparatus

The data in this report come from a male rhesus monkey (*Macaca mulatta*), which was referred to as Monkey 1 by Brasted and Wise (2004). That previous report details the methods and results concerning the animal's behavior, histology, electromyographic activity, and eye movements, as well as other methodological details not mentioned here. All procedures conformed with the *Guide for the Care and Use of Laboratory Animals* (rev. 1996, ISBN 0-309-05377-3) and were institutionally approved. The monkey sat, with its head fixed, in a primate chair. A video monitor was placed 30 cm in front of the monkey and it was covered by a transparent 25 cm by 17 cm touch screen (C.A.M. Graphics, Amityville, NY). We monitored eye position with an infrared oculometer (Bouis Instruments, Karlsruhe, Germany) at 500 or 1000 Hz.

Behavioral design

The monkey learned an arbitrary visuomotor mapping between an instruction stimulus (IS) and a reaching movement to one of four target locations (Fig. 1). The four response targets were located in the four corners of the screen. On each behavioral trial, an IS appeared, which consisted of two colored ASCII characters, ~5 cm and ~3 cm high, superimposed at the center of the video monitor. The monkey's task was to learn which one of the four targets on the touch screen would, if contacted, produce reinforcement on trials with a given IS.

On each trial, a stimulus was selected pseudorandomly from a set of eight stimuli. Each stimulus instructed one and only one correct response. Expressed differently, each IS was associated with one of the four response targets. All eight stimuli in each set of stimuli were highly familiar to the monkey. For four of those stimuli, the visuomotor mappings never changed, i.e., each of these four ISs always instructed the same response. We called these kinds of associations *familiar mappings*. For the other four stimuli, although stimuli *per se* were highly familiar to the monkey, the responses that they instructed changed frequently so that they had to be learned and relearned through trial and error. We called these kinds of associations *novel mappings*. Previous work has shown that monkeys learn novel mappings based on familiar stimuli in the same way that they learn such mappings based on novel stimuli. Fig. 9 of Brasted et al. (2003) shows this result for eight monkeys. All eight of those monkeys showed the same learning rate and the same number of trials to a learning criterion of 90% correct responses for novel and familiar stimuli. On this basis, we assume that use of familiar ISs had no impact on the interpretation of the present results.

Each trial began with the monkey touching and maintaining contact with a metal bar below the screen, which led to the appearance of the four response targets. After 200 ms, a white circle, 4 mm in diameter, appeared at the center of the screen. The monkey fixated this point for 600 ms, after which the IS appeared for 400 ms. We termed this interval of each trial the

IS-on period. In the rare event that the monkey broke fixation during the first 150 ms of the IS-on period, the trial aborted. After the IS-on period, a gray, 3-cm square replaced the IS for a duration of 0.75–1.75 s, which we called the *delay period*. The monkey had to maintain contact with the metal bar until the gray square disappeared, an event termed the trigger stimulus (TS or TS onset) because it cued the monkey to reach toward and touch one of the four targets within 1.9 s. After 100 ms of contact with one of the four targets, all four target outlines filled with white, which indicated that the computer had registered a response. Immediately afterwards, we imposed a variable *pre-reward period*, which lasted 0.75 s, 1.0 s, or 1.25 s (pseudorandomly selected). The monkey maintained contact with the chosen target for this entire period. After a correct choice, the monkey received ~0.1 ml of water at the end of the pre-reward period as reinforcement, and the targets disappeared at the same time. After an incorrect choice, the targets disappeared after the same variable period, but no reinforcement was delivered. The monkey learned four novel visuomotor mappings concurrently, which means that they did not train on one stimulus–response pair until achieving some criterion level of performance. However, a correction procedure was run so that, after errors, the IS was repeatedly presented for consecutive trials until the monkey made a correct response. After that correct response, the next IS was chosen pseudorandomly, and it could either be the same or different from the IS on the previous trial. The intertrial interval was 2.5 s.

Recordings

Single-unit potentials were isolated with platinum-iridium electrodes (0.5–1.5 M Ω at 1 KHz) driven by a 16-electrode microdrive (Thomas Recording, Giessen, Germany). The signal from each electrode was filtered with a bandpass of 0.6–6.0 KHz, amplified, and discriminated using a Multispike Detector (Alpha-Omega Engineering, Nazareth, Israel) or a Multichannel Acquisition Processor (Plexon, Dallas, Texas). CORTEX software (<http://www.cortex.salk.edu/>) was used to control behavior and collect data. The recordings alternated between the striatum and PMd from day to day, and we avoided recording the activity of tonically active striatal neurons (TANs).

Data analysis

We used MatOFF (<http://dally.nimh.nih.gov/matoff/matoff.html>), SPSS (<http://www.spss.com/>), and custom programs for analysis. We focused the neuronal-activity analysis on nine task periods: a baseline period, reference period, and seven principal task periods (Fig. 1). The *baseline period* was a 100-ms interval beginning 500 ms before IS onset; the *reference period* also began 500 ms prior to the IS but lasted until IS onset. We defined the seven principal task periods as follows: (1) the *IS-on period* extended from 80–400 ms after the IS onset; (2) the *IS-off period* was 80–250 ms after the offset of the IS; (3) the *delay period* extended from 500–250 ms prior to the TS; (4) the *pre-movement period* was the 250 ms interval before breaking screen contact; (5) the *movement period* was the 250 ms after breaking screen contact; (6) the *pre-reward period* was the last 750 ms (the shortest of the three possible pre-reward delay intervals) before the reward was delivered, provided that the monkey had chosen the correct target; and (7) the *post-reward period* was measured for 1000 ms after reward. Task-related activity (one of the seven principal task periods vs. the reference period) and mapping selectivity (four IS–response mappings) were assessed for each task period by a two-factor analysis of variance (ANOVA, $\alpha = 0.01$). The change-point test for continuous variables ($\alpha = 0.01$) was used to detect significant time-trends in the neurophysiological and behavioral data (Siegel and Castellan, 1988).

The population perievent time histograms were computed in two ways. In one method, each cell contributed its mean activity and the averages were baseline corrected, thus showing the net modulation. In the other, each cell average was normalized to its maximum discharge rate for a given IS and response, and then the histograms were baseline corrected. The latter kind

of average could not be dominated by the activity of one or a few neurons, but in practice these two techniques yielded virtually identical results. Some of the population averages included all of the neurons recorded in either PMd or in the putamen. Others were restricted to those cells showing significant changes in activity according to the change-point test. These two kinds of population measures also yielded very similar results, and therefore the former will be emphasized here. Other, more selected population histograms were prepared for cells either increasing activity during learning or decreasing activity, both overall and for particular task periods. These averages did not generate any noteworthy findings.

Learning phases were defined in a manner similar to that of Chen and Wise (1995a), taking into account the method for detecting the onset of learning devised by Wirth et al. (2003). Their method is called a state-space filter, and it detects the trial on which the lower 95% confidence limit of averaged performance first exceeds chance levels, which in this case and in theirs, was 25% correct. We divided trials for novel mappings into three groups: (1), an early phase of learning that incorporated all trials up to (but not including) the learning-criterion trial of Wirth et al. (2003); (2) a middle phase of learning comprising trials from (and including) the same learning-criterion until the attainment of asymptotic performance, as identified by the first occurrence of three consecutive correct responses; and (3), all trials after reaching the asymptote. The latter group included the last trial of the series of three consecutive correct responses that defined the attainment of asymptotic performance.

The mapping selectivity index S_i , devised by Moody et al. (1998), was calculated as:

$$S_i = \frac{\left(k - \sum_{n=1}^k \left(\frac{i_n}{i_{\max}} \right) \right)}{k - 1},$$

where k is the total number of mappings, neuronal activity level i_n is observed for the n^{th} mapping, and $i_{\max}$ is the modulation rate for the mapping that elicited the highest activity level.

Results

Learning rates

Fig. 2 shows the learning rates for the data used in the present analysis, from the same sessions as the neuronal-activity data presented below. As documented in Brasted and Wise (2004), the monkey did not differ in its learning rate for the sessions in which we collected neuronal data from the putamen versus PMd. The plots in Fig. 2A–D show an increase in the percentage of correct responses as a function of trial number for each response choice, averaged over 178 visuomotor mappings for each target. As the monkey learned that a given IS instructed a particular response, its performance score increased to asymptotic levels. At the same time, the proportion of responses to the three incorrect response targets decreased. The monkey showed a slight bias toward the upper left target, but otherwise the performance curves for the four response targets resemble each other closely. Fig. 2E shows the average performance curve for all four targets, i.e., the mean of the learning curves shown in Figs. 2A–D.

Note, however, that the monkey did not learn the individual visuomotor mappings one at a time, but instead learned the four novel mappings concurrently. For example, the IS for the upper left target appeared in an unpredictable, intermixed sequence of trials that included the other seven ISs in a given stimulus set. In each stimulus set, four ISs involved familiar mappings and four involved novel mappings. (The latter occurred twice as frequently, on average, as the former.) Fig. 2F shows the overall learning curve, excluding the trials with familiar mappings. On average, the monkey reached a behavioural plateau in ~60 trials, after a rapid initial improvement in performance.

Comparison of PMd and putamen population activity levels

We recorded from 120 PMd cells and 75 cells in the putamen, most of which were located in the zone of putamen that receives inputs from PMd (see Brasted and Wise, 2004). Except where explicitly noted, all recorded neurons were included in the computations reported here. Some of these cells lacked task relationships (11 PMd and 13 putamen cells), and some had modulations of less than 5 impulses/s (31 PMd and 22 putamen cells). Accordingly, the absolute levels of modulation observed with this nonselective inclusion criterion underestimate the activity of task-related and learning-related neurons. However, these populations were free of any bias other than those involved in selecting the neurons for recording during the daily recording sessions.

Fig. 3 shows perievent time histograms for the activity of selected populations of neurons in PMd (blue) and in the putamen (red). It does not include the entire populations of putamen and PMd neurons, but instead conveys a more accurate impression of the magnitude of activity modulation for the cells that show significant learning-dependent increases in activity, as defined by Chen and Wise (1995a). Activity early in learning appears on the top row, that for the middle and late phases of learning appears on the next two rows, and the bottom row shows the population activity for familiar mappings. The Materials and Methods provides precise definitions, but for heuristic purposes it may be useful to consider these three phases of learning as follows: the early phase corresponds to the period prior to a significant improvement in performance, as judged by the proportion of correct response choices, the middle phase corresponds to the steepest improvement in performance, and the late phase includes all trials after performance reached an asymptote. Each column of perievent time histograms shows population activity aligned on a particular event, noted at the bottom of the figure. Neurons contributing to the averages in Fig. 3 all showed statistically significant increases in activity during learning in at least one of four task periods: the IS-on period, the delay period, the premovement period, and the prereward period, marked by asterisks at the top of the figure. Fig. 3 also shows averages for task periods “attached” to one selected for increasing activity, i.e., task periods immediately before or after the selected period. (The averages aligned on target acquisition are “attached,” in this sense, to the premovement period.) Note that, in this analysis, a given neuron could contribute to some columns of the figure, but not others, and that only correctly executed trials contribute activity to the averages. As a further selection criterion, activity during the delay period included only neurons with near-baseline activity levels during the early phase of learning. Because the cells contributing to this subpopulation histogram were selected for increases during learning, the magnitude rather than the occurrence of the increase is of interest for the present analysis.

The population histograms in Fig. 3 show several features. Beginning with PMd’s population activity, the phasic increase that followed the appearance of the instruction stimulus increased by a factor of approximately 2.5:1 (late phase of learning:early phase). The buildup of activity prior to the IS, called anticipatory activity (Mauritz and Wise, 1986), was absent in the earliest phase of learning, but increased to a modulation of approximately 10 impulses/s in the late phase (blue arrows). During the delay period, cells that began each learning block near reference or baseline activity levels increased as learning progressed, also to about 10 impulses/s. After movement onset, but before target acquisition, activity in PMd decreased to below baseline levels.

As for the putamen, the change in activity after the IS was very modest in comparison with that in PMd, reaching less than a fifth of its modulation level. Perimovement activity showed a similar pattern. For putamen cells with activity that increased during the premovement period, their greatest modulation never reached the level seen in PMd and, furthermore, this activity increase occurred later than in PMd. After movement onset, but before target acquisition, average activity in putamen remained above baseline levels as the time for reward approached

(unlike the PMd population). The period immediately after the reward showed a phasic decrease in putamen activity to baseline levels early in learning, but this decrease was not seen later in learning or for familiar mappings.

Fig. 3 presents population histograms based on neurons selected for increases in task-related activity during learning. To eliminate the effects of this selection criterion on the population averages, we computed averages for the entire populations of PMd and putamen cells, and Fig. 4 presents these results. Thus, the data presented in Fig. 4 come from all 120 PMd cells and all 75 putamen cells recorded. Only correctly performed trials contributed to these averages, and each cell contributes in proportion to its baseline-corrected activity modulation. To ensure that a few cells did not dominate the population, we also computed population histograms in which each neuron contributed in proportion to its maximal activity. As noted in the Material and Methods, these two analyses yielded virtually identical results, as did analyses restricted to task-related neurons (by ANOVA) or neurons with significant learning-related activity changes (by the change-point test).

Several features of the population activity changed during arbitrary visuomotor learning. Reward-related signals were observed in PMd that resembled those previously reported for midbrain dopamine neurons and striatal cells (see Discussion). For example, like dopaminergic neurons, there was a general tendency for postreward discharge modulations to decrease as learning progressed, along with an increase in activity shortly after the IS (blue arrows in Fig. 4). In both task periods, the overall pattern of activity for novel mappings progressively resembled that for familiar mappings as learning progressed. Fig. 5 quantifies this development by plotting the difference between population activity for novel-mapping trials and familiar-mapping trials, for each of the learning phases. It shows both the peak difference for 10-ms bins (dashed lines) and the mean difference for each point composing Fig. 4's population histograms (solid lines). This difference decreased in each task period as learning progressed. No such systematic changes were observed when familiar-mapping trials were divided by the learning phase in which they occurred (see Electronic Supplementary Materials, Figure 1).

As Fig. 4 shows, the population activity in the putamen differed from that in PMd. Perhaps the most significant difference between the PMd and putamen populations can be observed for the prereward period (red arrows in Fig. 4). The putamen population showed a persistent, increased activity relative to baseline activity during this period, but PMd did not. Fig. 4 shows that for the overall population of putamen neurons, this elevated activity level did not appear to change dramatically as a function of learning. (In the selected subpopulation in Fig. 3, the prereward activity increased, but this change reflects only the fact that this subpopulation was selected for increases in activity during learning.) The two regions also differed in their anticipatory activity, with a clear anticipatory buildup in PMd, but not in the putamen population (Fig. 4).

Fig. 6 shows population averages for IS-on period, but plotted over the entire learning period rather than divided into early, middle, and late phases of learning. In accord with previous results of PMd (Mitz et al., 1991; Brasted and Wise, 2004), the post-IS activity increased most dramatically at the time that the monkey's performance improved most rapidly, the middle phase of learning as defined here.

Comparison of PMd and putamen error signals

We also examined the difference in population activity for correctly executed trials and trials on which the monkey chose an incorrect response. Figs. 7 and 8 show these results for PMd (blue) and the putamen (red). The population activity for error trials appears as the heavier line and the difference (correct – error trial activity) is shown in the right half of both figures. The monkey made very few errors for familiar mappings, so those data are not presented.

These averages contain several noteworthy features. For PMd, Fig. 7 shows that the post-IS activity on error trials initially, i.e., in the early phase of learning, exceeded that on correct trials (upward green arrow in Fig. 7), as did delay-period activity. Indeed, not until just prior to movement onset did the correct-trial and error-trial activity reach the same level relative to baseline. In the middle and late phases of learning, however, this difference reversed (downward green arrows in Fig. 7), and average activity on correct trials came to exceed that on error trials. The putamen population showed smaller differences between correct and incorrect trials, but there was a slight elevation of delay-period activity on trials in which the monkey later chose an incorrect response, especially early in learning.

As illustrated in Fig. 8, PMd showed the most dramatic difference between correct and error trials in the postreward period (green arrows in Fig. 8), with a much greater level of activity for correct trials, immediately after the reward, compared to trials in which the reward was not delivered (error trials). Although postreward phasic bursts were observed for both correct trials (after reward delivery) and error trials (after reward delivery was omitted), the bursts observed in correct trials were of greater magnitude and occurred earlier, relative to the time of reward. After the initial phasic burst, correct-trial activity in the PMd population dropped significantly below baseline levels and remained so during the latter stages of the postreward period. Conversely, on error trials, late postreward activity remained at or above baseline levels. Unlike the PMd population, the putamen population showed a sustained postreward increase in activity in correct trials relative to error trials, especially during the late phase of learning (magenta arrow in Fig. 8). During that phase of learning, the putamen population decreased its activity less when a reward was delivered than when no reward occurred.

Comparison of PMd and putamen mapping selectivity

Fig. 9 explores the time course of the changes in mapping selectivity for a population of neurons in both PMd and the putamen. The selectivity of neural signals just prior to the reward is of considerable theoretical interest because the monkey had to retain information about the mapping (or, alternatively, the IS and response) until the “knowledge of results” was received, in this case reward or nonreward. The Discussion explains this idea in more detail.

We calculated the mapping selectivity index S_i (see Materials and Methods) for each neuron showing a minimal level of activity (5 impulses/s) for at least one of the novel visuomotor mappings in at least one phase of learning. Fig. 9 shows that, during the prereward period, mapping selectivity in PMd decreased as the time for reward approached (Fig. 9A, C). The PMd population also had a relatively low level of activity, relative to baseline, during this period (Fig. 4). In contrast to the PMd population, not only did the overall activity level become much higher in the putamen just prior to reward (Fig. 4), but the cells maintained a high level of mapping selectivity up to the time of reward, most notably during the early and middle phases of learning (Fig. 9B, D). In the later phase of learning, after performance had ceased improving, the putamen population lost much of its mapping selectivity during this task period and adopted a level similar to that seen in PMd and for the interleaved familiar mappings in both PMd and the putamen. Both the PMd (paired t-test, $t_{76}=4.68$, $p<0.001$) and putamen ($t_{38}=3.35$, $p<0.01$) populations displayed significantly greater mapping selectivity over the 300 ms preceding the reward in the early and middle learning phases, when compared with the late learning phase and familiar mapping trials.

Discussion

The mechanisms of arbitrary visuomotor mapping are of particular interest because this form of learning allows individuals to link any goal-directed action with any sensory input that can be detected. This level of behavioral flexibility characterizes much of what we consider

intelligent behavior and may underlie not only the symbolic guidance of actions, but also aspects of visual and vocal communication (Murray et al., 2002).

Population activity for evaluation of feedback

The present experimental design incorporated a variable and relatively long interval from the time that the computer registered the monkey's response until the delivery of fluid reinforcement, provided that the monkey chose the correct response. The inclusion of this variable period arose from theoretical considerations and a specific prediction about a possible distinction between PMd and putamen activity.

To understand this prediction, consider the nature of error feedback in arbitrary visuomotor learning and how it differs from feedback in what might be called "traditional" forms of motor learning, such as the learning of smooth reaching movements to a visible target. In traditional motor learning, sensory transducers can provide online error signals based on proprioceptive or visual feedback, in relation to acquiring the target of reach. These error signals can arrive during the course of a movement, can be based on a comparison of predicted and actual feedback, and can reflect changes in either the location of the target, the position of the hand, or the state of the motor system (Kawato, 1989; Thoroughman and Shadmehr, 1999). These error signals can be graded, i.e., they can convey an amount of error and not simply success or failure in reaching a goal.

Error feedback in arbitrary visuomotor learning differs dramatically from that in traditional motor learning. Because of the arbitrary linkage between the IS and the correct target, no assessment of response correctness is possible during the course of the movement or immediately afterward. Until either the reward occurs or sufficient time elapses to indicate that reward will not occur—i.e., until the monkey receives "knowledge of results"—there can be no assessment of whether it has made the correct choice. Thus, arbitrary visuomotor learning cannot occur if, at the time of the reward, the monkey has forgotten the IS–response mapping or, alternatively, either of those elements separately. In the present experiment, the monkey could view its hand during the prereward period, so there was no requirement that the monkey remember its motor response or its spatial goal. Accordingly, learning probably depended mostly on the memory of the IS or, more likely, the IS–response mapping. Any neural system that contributes directly to arbitrary visuomotor mapping must retain that information in some form, and it must do so until the knowledge of results is obtained.

On this basis, we predicted that a sustained, mapping-selective neural signal should occur in any neuronal population involved directly in this form of learning. The present result showed a strong indication of such a memory signal in the putamen, especially during the early and middle phases of learning. This memory signal was characterized, in the putamen population, by relatively high levels of discharge above baseline (Fig. 4) and a relatively high degree of selectivity for the mapping that had just been chosen (Fig. 9). During the later phases of learning, when the monkey performed at asymptotic levels and occasionally below, the putamen population continued to show relatively high levels of activity in terms of discharge rate (Fig. 4), but that activity showed much less mapping selectivity than during the early and middle phases of learning. It came to resemble the modest levels of mapping selectivity observed for familiar mappings (Fig. 9).

In contrast to the properties of the putamen population, in PMd the mean discharge rate remained relatively near baseline levels during the prereward period (Figs. 4), and the mapping selectivity decreased as the reward time approached (Fig. 9).

Examination of the time course of these signals (not illustrated) showed that, in the putamen, mapping selectivity became especially prominent at approximately the time of movement

onset, especially during the early and middle phases of learning. In general, at approximately the time that the reaching movement began, the population activity in PMd and putamen changed in different ways. PMd activity returned to baseline levels or below, with little deviation until after the reward (Fig. 4) or nonreward (Fig. 8) and it became less selective for the mapping (Fig. 9). By contrast, putamen activity increased just prior to movement onset and maintained this heightened level of activity (Fig. 4) and mapping selectivity (Fig. 9) until after the knowledge of results had been obtained.

It might seem surprising in this context that the putamen's prereward activity increased as robustly, if not more so, for familiar visuomotor mappings as for novel ones, at least in terms of discharge rates (Fig. 4). Thus, it would be misleading to relate this sustained prereward activity, *per se*, to learning, if that term is construed to exclude the maintenance of already learned associations. Rather, the data suggest that the mapping selectivity of the neurons in the putamen population, rather than their overall discharge rate, contributed most to evaluating the correctness of a choice during learning (Fig. 9). However, mapping selectivity persisted to a lesser degree in the late phase of learning and for familiar mappings, and this selectivity could have played a role in maintaining learned mappings in memory. It seems likely that the arbitrary nature of the associations, and the monkey's experience with the fact that the instructional significance of visual cues changed often, made continual evaluation of mappings an important cognitive operation. Evaluation thus probably occurred for already learned mappings as well as new ones, reminiscent of a role has been proposed for the supplementary eye field (Schall et al., 2002) and anterior cingulate cortex (Ito et al., 2003).

Comparison with other brain areas

The strength of mapping selectivity in the putamen and its weakening in PMd as the reward time approached (Fig. 9) suggest that the putamen plays a more direct role in arbitrary visuomotor learning than does PMd. This conclusion accords, in many respects, with that from a recent study of mapping selectivity in the caudate nucleus during arbitrary visuomotor learning (Pasupathy and Miller, 2005), and it will be of interest to explore the relative contributions of these two parts of the striatum to this form of associative learning.

Similar changes occur in the hippocampus during the learning of arbitrary visuomotor mappings (Wirth et al., 2003), also called object-in-scene learning. A comparison of the precise relative timing of neuronal activity changes in these various regions remains limited by small, yet significant differences in the methodology used from study to study. Nevertheless, relative to the learning criterion trial as defined by Wirth et al. (2003), hippocampal cells change their activity several trials earlier than do cells in the putamen and PMd (S. Wirth, E. Buch, S. Wise, and W. Suzuki, unpublished data). It seems likely that, together, the caudate nucleus (Pasupathy and Miller, 2005) and the hippocampus (Wirth et al., 2003) contribute most to the earliest phases of learning, and that PMd and the putamen play a critical role in more permanent storage for later recall. The results of several brain-lesion studies point to the same conclusion (Murray and Wise, 1996; Bussey et al., 2002; Brasted et al., 2002, 2003; Nixon et al., 2004, see also Wise and Murray, 1999; Hadj-Bouziane et al., 2003), as do the neuroimaging studies cited in the Introduction. We propose, in large part along the lines already suggested by Schultz et al. (2003) and by Pasupathy and Miller (2005), that information processing in the striatum contributes most to linking an arbitrary visuomotor mapping to a current value or predicted outcome, by virtue of its inputs from dopaminergic neurons of the midbrain (see below), and that the hippocampus contributes most to the establishment of a visuomotor mapping *per se* (see Brasted et al., 2003). Compared to the caudate nucleus and the hippocampal system, PMd and the putamen participate more in recall and retention over the long term. According to this idea, the prefrontal cortex plays a pivotal role in both rapid learning and long-term retention, through its participation in networks that link it both to the caudate nucleus and inferotemporal

cortex, on the one hand, and to the remainder of the frontal lobe, including the premotor cortex, on the other. Hadj-Bouziane et al. (2003) have expressed a similar view.

Interpretational limitations

Taken on its own, the sustained prereward signal seen in the putamen population could have reflected neural operations other than, or in addition to, short-term maintenance of visuomotor mappings. We have already mentioned the possibility that this activity could have been due to the maintenance of information about the visual stimulus that instructed a given response. In the present context, this distinction has little relevance because an IS always instructed a single response during the recording of any neuron's activity. Therefore, we cannot distinguish between signals that reflected stimulus memory from those involved with maintaining the memory of a IS–response mapping.

Alternatively, the putamen's prereward activity might have reflected the anticipation or prediction of reward. If so, then this signal should have increased as the learning of novel visuomotor mappings progressed and reward therefore became more likely. Certain population activity in PMd, discussed below, did indeed display this property in the present data and in that of others (Roesch and Olson, 2004), but the prereward activity level in putamen changed little over the course of learning, contrary to what would be expected for reward anticipation. Furthermore, the mapping selectivity went down, not up during learning (Fig. 9). Thus, reward prediction or anticipation probably contributed little to this aspect of the putamen's neural activity. The sustained prereward activity might also have reflected attentional factors, such as those related to the events that occurred simultaneously with reward, but this interpretation would not account for the changes in mapping selectivity observed during learning.

Many of the other, more general problems in interpreting the present data were discussed in detail by Brasted and Wise (2004). For example, in analyses that consider only correctly executed responses, we know that both the stimulus and response were the same from the first correctly executed trial to the last one, and thus sensory and motor factors, *per se*, cannot account for the learning-related changes in activity. The present study involved forelimb reaching movements, which raises a myriad of interpretational issues involving kinematics, dynamics, reaction time, movement time, and the variation in those variables from reach to reach. Brasted and Wise (2004) cite the studies and findings that eliminate all of those variables as general accounts for learning-related changes in activity.

Beyond the general issues discussed previously, some problems are specific to the population analyses presented here. For example, we assume that neurons collected serially, for the most part, represent a valid sample of what would be observed if populations of neurons were monitored simultaneously. Although new methods will make the latter approach feasible in future studies (Laubach et al., 2000; Nicolelis et al., 2003), we did not attempt to use those methods in this one. Our conclusions thus rely on the validity of the assumption stated above. A more serious issue, perhaps, is the fact that we examined these neuronal signals in one monkey rather than the traditional two. This approach, although it has ample precedent, remains controversial. We acknowledged that another monkey might have generated different results, although we have no reason to suppose that it would. To allay such concerns, the Electronic Supplementary Materials, Figure 2, presents data from a second monkey, called Monkey 2 by Brasted and Wise (2004). These data come only from the putamen, but confirm that aspect of Fig. 9, at least in broad terms.

Error and reward-prediction signals in population activity

At the single-cell level, generalized correlates of neuronal activity such as reward expectancy or motivation can be ruled out for the majority of neurons, because those factors would not be

selective for a given mapping, and individual neurons typically showed such selectivity. At the population level, however, and for neurons lacking mapping selectivity, it is likely that generalized computations such as reward prediction were reflected in the population activity shown in Figs. 4, 7, and 8. Early in learning, when reward expectation must have been relatively low, PMd showed greater postreward modulation than later in the learning of novel mappings and for familiar mappings. The small decrease in population activity after the reward in the putamen, which occurred only early in learning, showed the same properties in reverse. Thus, a deviation from the expected reward appeared to contribute to neural activity in both PMd and in the putamen. This activity probably reflected the reward-prediction error signal described for dopamine neurons in the brainstem and in striatal cells by Schultz and his colleagues (Waelti et al., 2001; Fiorillo et al., 2003) and by Hikosaka and his (Itoh et al., 2003; Kawagoe et al., 2004; Takikawa et al., 2004). Chen and Wise (1995a,b) showed similar population averages for the supplementary eye field and frontal eye field. In that study, postreward activity failed to show the learning-related increase in activity that was seen in the illustrated subpopulation as a whole (Fig. 13 in Chen and Wise, 1995a; Fig. 7 in Chen and Wise, 1995b).

The population activity reported here for PMd and the putamen had other similarities to that reported previously for the striatum and its dopaminergic inputs, both of which have, in several neurophysiological studies, displayed activity that reflected reward prediction in addition to reward-prediction error (Hollerman and Schultz, 1998; Hollerman et al., 1998). In the present study, we found that the population activity after the IS increased at approximately the same times during learning that postreward activity decreased. In past studies, this finding has been interpreted in terms of the assignment of reinforcement value or valence to the IS, as it came to be predictive of and correlated with a positive outcome. This signal was found in PMd in the present study (Figs. 3, 4 and 6), and its properties supported the possibility that it reflects inputs from dopaminergic cells, perhaps those projecting directly to PMd (Lewis et al., 1987; Berger et al., 1988; Gaspar et al., 1992; Williams and Goldman-Rakic, 1993). If so, then blockade of dopamine receptors at the cortical level should eliminate this signal.

The population activity also differed on error trials versus trials with correct responses. As one difference between PMd's activity after correct and incorrect responses, population activity quickly returned to baseline levels and below after a correct response, whereas on trials with incorrect responses the activity increased much less but remained at or above baseline levels (Fig. 8). The putamen activity after incorrect responses showed no such tendency. At the times marked by the green arrows in Fig. 8 (left), the population activity in the putamen changed little after correct responses occurred, but after incorrect responses a phasic decrease in putamen activity occurred at about the same time as the postreward discharge in PMd, especially during the early and middle phases of learning. Recall that after incorrect responses, the temporal alignment point labeled "reward time" in Fig. 8 was not associated with any concrete event. Rather, at the time that reinforcement would have been delivered had the response been a correct one, nothing happened. These findings are consistent with the idea that one aspect of the PMd activity signaled relatively unexpected reward, whereas putamen activity signaled the absence of a relatively expected reward (on error trials). However, the activity of the PMd population after incorrect responses varied from this pattern: at a delayed latency from the "reward time", i.e., after the "event" of nonreward, the population activity increased, although not to the extent that followed a reward. As a consequence of both of these properties, there was greater population activity after the "reward time" in correct versus incorrect trials, but this occurred earlier in PMd both in terms of the latency from the event (green versus purple arrows in the right part of Fig. 8) and earlier during the course of learning (i.e., during the early and middle phases of learning). Of course, we acknowledge the speculative nature of these possible interpretations, and offer these views as hypotheses for future testing rather than as conclusions that follow from the present data.

Prior to the reward, population activity also differed on error trials versus correct trials. PMd activity during the IS-on, IS-off, and delay periods was higher when the upcoming response was going to be wrong in the early phase of learning, but it was lower during the middle and late phases (Fig. 7, green arrows). At first glance, this results seems paradoxical. But perhaps this difference reflected a variable amount of neural noise in decision networks. According to this idea, in a winner-take-all network, the lower the neural noise, the more reliably that the network can make a decision. In the presence of higher noise levels, the overall activity modulation might be higher, but the relative activity for the competing decisions might be smaller, thus leading to less accurate output.

Timing differences in movement-related activity

These population histograms confirm the well-known timing difference between the PMd and putamen populations for movement-related activity (Crutcher and DeLong, 1984; Turner and DeLong, 2000). Putamen cells show phasic movement-related modulations later than cells in cortical motor areas, including PMd (Fig. 4).

Conclusion

Striatal input connectivity seems well suited to detecting context and using that context as a basis for predicting the outcome of actions (Houk et al., 1995; Schultz, 1998; Suri and Schultz, 1999; Bar-Gad et al., 2000; Suri et al., 2001; Zheng and Wilson, 2002). If the putamen detects the context for a given action, it can provide a signal that leads to the programming and execution of cognitive or motor output through its recurrent influence over the cortex (Groves, 1983; Houk et al., 1993; Houk and Wise, 1995; Mink, 1996; Bar-Gad et al., 2000; Gurney et al., 2001; Bar-Gad and Bergman, 2001; Djurfledt et al., 2001). The present results add support to ideas from many studies that point to a central role for the striatum in the learning of arbitrary mappings (Tremblay et al., 1998; Hadj-Bouziane and Boussaoud, 2003; Pasupathy and Miller, 2005) and other kinds of associations, perhaps by assigning value to particular mappings (for recent reviews, see Shohamy et al., 2004; Schultz, 2004).

Furthermore, the present study shows that activity in a cortical area, PMd, which is known to be critical for arbitrary visuomotor mapping, differs at the population level from the activity of its striatal target, the dorsomedial putamen. Of particular note is the finding of sustained, mapping selective activity in the putamen during prereward period, after the response had been selected and executed (Fig. 9). Both the level of activity and the degree and time course of its mapping selectivity differ in the putamen and PMd. Although the recurrent nature of the network architecture imposes a degree of similarity on cortical activity and that of its targets in the striatum (for review, see Schultz et al., 2003), the present findings accord with those of Pasupathy and Miller (2005) and the idea that interconnected areas of cortex and the striatum play importantly different roles in the learning of arbitrary visuomotor mappings.

Acknowledgements

We thank Mr. Jim Fellows for assistance with behavioral training and Dr. Andrew R. Mitz for technical and engineering support.

References

- Asaad WF, Rainer G, Miller EK. Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 1998;21:1399–1407. [PubMed: 9883732]
- Bar-Gad I, Bergman H. Stepping out of the box: information processing in the neural networks of the basal ganglia. *Curr Opin Neurobiol* 2001;11:689–695. [PubMed: 11741019]

- Bar-Gad I, Havazelet-Heimer G, Goldberg JA, Ruppin E, Bergman H. Reinforcement-driven dimensionality reduction--a model for information processing in the basal ganglia. *J Basic Clin Physiol Pharmacol* 2000;11:305–320. [PubMed: 11248944]
- Barefoot HC, Baker HF, Ridley RM. Crossed unilateral lesions of temporal lobe structures and cholinergic cell bodies impair visual conditional and object discrimination learning in monkeys. *Eur J Neurosci* 2002;15:507–516. [PubMed: 11876778]
- Barefoot HC, Maclean CJ, Baker HF, Ridley RM. Unilateral hippocampal and inferotemporal cortex lesions in opposite hemispheres impair learning of single-pair visual discriminations as well as visuovisual conditional tasks in monkeys. *Behav Brain Res* 2003;141:51–62. [PubMed: 12672559]
- Berger B, Trottier S, Verney C, Gaspar P, Alvarez C. Regional and laminar distribution of the dopamine and serotonin innervation in the macaque cerebral cortex: a radioautographic study. *J Comp Neurol* 1988;273:99–119. [PubMed: 3209731]
- Brasted PJ, Bussey TJ, Murray EA, Wise SP. Fornix transection impairs conditional visuomotor learning in tasks involving nonspatially differentiated responses. *J Neurophysiol* 2002;87:631–633. [PubMed: 11784778]
- Brasted PJ, Bussey TJ, Murray EA, Wise SP. Role of the hippocampal system in associative learning beyond the spatial domain. *Brain* 2003;126:1202–1223. [PubMed: 12690059]
- Brasted PJ, Wise SP. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. *Eur J Neurosci* 2004;19:721–740. [PubMed: 14984423]
- Bussey TJ, Wise SP, Murray EA. The role of ventral and orbital prefrontal cortex in conditional visuomotor learning and strategy use in rhesus monkeys (*Macaca mulatta*). *Behav Neurosci* 2001;115:971–982. [PubMed: 11584930]
- Bussey TJ, Wise SP, Murray EA. Interaction of ventral and orbital prefrontal cortex with inferotemporal cortex in conditional visuomotor learning. *Behav Neurosci* 2002;116:703–715. [PubMed: 12148937]
- Cahusac PM, Rolls ET, Miyashita Y, Niki H. Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus* 1993;3:29–42. [PubMed: 8364681]
- Canavan AGM, Nixon PD, Passingham RE. Motor learning in monkeys (*Macaca fascicularis*) with lesions in motor thalamus. *Exp Brain Res* 1989;77:113–126. [PubMed: 2792254]
- Chen LL, Wise SP. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 1995a;73:1101–1121. [PubMed: 7608758]
- Chen LL, Wise SP. Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 1995b;73:1122–1134. [PubMed: 7608759]
- Crutcher MD, DeLong MR. Single cell studies of the primate putamen. II Relations to direction of movement and pattern of muscular activity. *Exp Brain Res* 1984;53:244–258. [PubMed: 6705862]
- Deiber MP, Passingham RE, Colebatch JG, Friston KJ, Nixon PD, Frackowiak RS. Cortical areas and the selection of movement: a study with positron emission tomography. *Exp Brain Res* 1991;84:393–402. [PubMed: 2065746]
- Deiber MP, Wise SP, Honda M, Catalan MJ, Grafman J, Hallett M. Frontal and parietal networks for conditional motor learning: a positron emission tomography study. *J Neurophysiol* 1997;78:977–991. [PubMed: 9307128]
- Djurfeldt M, Ekeberg O, Graybiel A. Cortex–basal ganglia interaction and attractor states. *Neurocomputing* 2001;38–40:573–579.
- Eacott MJ, Gaffan D. Inferotemporal-frontal disconnection: The uncinate fascicle and visual associative learning in monkeys. *Eur J Neurosci* 1992;4:1320–1332. [PubMed: 12106395]
- Eliassen JC, Souza T, Sanes JN. Experience-dependent activation patterns in human brain during visual-motor associative learning. *J Neurosci* 2003;23:10540–10547. [PubMed: 14627638]
- Fiorillo CD, Tobler PN, Schultz W. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 2003;299:1898–1902. [PubMed: 12649484]
- Gaffan D, Harrison S. Inferotemporal-frontal disconnection and fornix transection in visuomotor conditional learning by monkeys. *Behav Brain Res* 1988;31:149–163. [PubMed: 3202946]
- Gaspar P, Stepniewska I, Kaas JH. Topography and collateralization of the dopaminergic projections to motor and lateral prefrontal cortex in owl monkeys. *J Comp Neurol* 1992;325:1–21. [PubMed: 1362430]

- Groves PM. A theory of the functional organization of the neostriatum and the neostriatal control of voluntary movement. *Brain Res Rev* 1983;5:109–132.
- Gurney K, Prescott TJ, Redgrave P. A computational model of action selection in the basal ganglia. I A new functional anatomy. *Biol Cybern* 2001;84:401–410. [PubMed: 11417052]
- Hadj-Bouziane F, Boussaoud D. Neuronal activity in the monkey striatum during conditional visuomotor learning. *Exp Brain Res* 2003;153:190–196. [PubMed: 14610634]
- Hadj-Bouziane F, Meunier M, Boussaoud D. Conditional visuo-motor learning in primates: a key role for the basal ganglia. *J Physiol (Paris)* 2003;97:567–579. [PubMed: 15242666]
- Halsband U, Passingham RE. Premotor cortex and the conditions for a movement in monkeys. *Behav Brain Res* 1985;18:269–277. [PubMed: 4091963]
- Hollerman JR, Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat Neurosci* 1998;1:304–309. [PubMed: 10195164]
- Hollerman JR, Tremblay L, Schultz W. Influence of reward expectation on behavior-related neuronal activity in primate striatum. *J Neurophysiol* 1998;80:947–963. [PubMed: 9705481]
- Houk JC, Adams JL, Barto AG (1995) A model of how the basal ganglia generate and use neural signals that predict reinforcement. In: Houk JC, Davis JL, Beiser DG (eds) *Models of Information Processing in the Basal Ganglia*. MIT Press, Cambridge, pp 249–270
- Houk JC, Keifer J, Barto AG. Distributed motor commands in the limb premotor network. *Trends Neurosci* 1993;16:27–33. [PubMed: 7679234]
- Houk JC, Wise SP. Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cereb Cortex* 1995;5:95–110. [PubMed: 7620294]
- Inase M, Li BM, Takashima I, Iijima T. Pallidal activity is involved in visuomotor association learning in monkeys. *Eur J Neurosci* 2001;14:897–901. [PubMed: 11576196]
- Ito S, Stuphorn V, Brown JW, Schall JD. Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 2003;302:120–122. [PubMed: 14526085]
- Itoh H, Nakahara H, Hikosaka O, Kawagoe R, Takikawa Y, Aihara K. Correlation of primate caudate neural activity and saccade parameters in reward-oriented behavior. *J Neurophysiol* 2003;89:1774–1783. [PubMed: 12686566]
- Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM. Building neural representations of habits. *Science* 1999;286:1745–1749. [PubMed: 10576743]
- Kawagoe R, Takikawa Y, Hikosaka O. Reward-predicting activity of dopamine and caudate neurons. A possible mechanism of motivational control of saccadic eye movement. *J Neurophysiol* 2004;91:1013–1024. [PubMed: 14523067]
- Kawato M. Adaptation and learning in control of voluntary movement by the central nervous system. *Adv Robotics* 1989;3:229–249.
- Laubach M, Wessberg J, Nicolelis MA. Cortical ensemble activity increasingly predicts behaviour outcomes during learning of a motor task. *Nature* 2000;405:567–571. [PubMed: 10850715]
- Lewis DA, Campbell MJ, Foote SL, Goldstein M, Morrison JH. The distribution of tyrosine hydroxylase-immunoreactive fibers in primate neocortex is widespread but regionally specific. *J Neurosci* 1987;7:279–290. [PubMed: 2879896]
- Mauritz KH, Wise SP. Premotor cortex of the rhesus monkey: neuronal activity in anticipation of predictable environmental events. *Exp Brain Res* 1986;61:229–244. [PubMed: 3948938]
- Mink JW. The basal ganglia: focused selection and inhibition of competing motor programs. *Prog Neurobiol* 1996;50:381–425. [PubMed: 9004351]
- Mitz AR, Godschalk M, Wise SP. Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. *J Neurosci* 1991;11:1855–1872. [PubMed: 2045890]
- Moody SL, Wise SP, di Pellegrino G, Zipser D. A model that accounts for activity in primate frontal cortex during a delayed matching-to-sample task. *J Neurosci* 1998;18:399–410. [PubMed: 9412516]
- Murray EA, Brasted PJ, Wise SP (2002) Arbitrary sensorimotor mapping and the life of primates. In: Squire LR, Schacter DL (eds) *Neuropsychology of Memory*. Guilford, New York, pp 339–348

- Murray EA, Wise SP. Role of the hippocampus plus subjacent cortex but not amygdala in visuomotor conditional learning in rhesus monkeys. *Behav Neurosci* 1996;110:1261–1270. [PubMed: 8986330]
- Nicolelis MAL, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, Wise SP. Chronic, multisite, multielectrode recordings in macaque monkeys. *Proc Nat Acad Sci USA* 2003;100:11041–11046. [PubMed: 12960378]
- Nixon PD, McDonald KR, Gough PM, Alexander IH, Passingham RE. Cortico-basal ganglia pathways are essential for the recall of well-established visuomotor associations. *Eur J Neurosci* 2004;20:3165–3178. [PubMed: 15579171]
- Pasupathy A, Miller EK (2005) Different time courses for learning-related activity in the prefrontal cortex and striatum. *Nature*: in press
- Paus T, Petrides M, Evans AC, Meyer E. Role of the human anterior cingulate cortex in the control of oculomotor, manual, and speech responses: A positron emission tomography study. *J Neurophysiol* 1993;70:453–469. [PubMed: 8410148]
- Petrides M. Deficits in non-spatial conditional associative learning after periarculate lesions in the monkey. *Behav Brain Res* 1985;16:95–101. [PubMed: 4041217]
- Ridley RM, ad Baker HF. Evidence for a specific information processing deficit in monkeys with lesions of the septo-hippocampal system. *Cortex* 1997;33:167–176. [PubMed: 9088729]
- Roesch MR, Olson CR. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* 2004;304:307–310. [PubMed: 15073380]
- Rupniak NMJ, Gaffan D. Monkey hippocampus and learning about spatially directed movements. *J Neurosci* 1987;7:2331–2337. [PubMed: 3612244]
- Schall JD, Stuphorn V, Brown JW. Monitoring and control of action by the frontal lobes. *Neuron* 2002;36:309–322. [PubMed: 12383784]
- Schultz W. Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Curr Opin Neurobiol* 2004;14:139–147. [PubMed: 15082317]
- Schultz W. Predictive reward signal of dopamine neurons. *J Neurophysiol* 1998;80:1–27. [PubMed: 9658025]
- Schultz W, Tremblay L, Hollerman JR. Changes in behavior-related neuronal activity in the striatum during learning. *Trends Neurosci* 2003;26:321–328. [PubMed: 12798602]
- Shohamy D, Myers CE, Grossman S, Sage J, Gluck MA, Poldrack RA. Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain* 2004;127:851–859. [PubMed: 15013954]
- Siegel S, Castellan NJ (1988) *Nonparametric statistics for the behavioral sciences*. McGraw-Hill: New York
- Suri RE, Bargas J, Arbib MA. Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience* 2001;103:65–85. [PubMed: 11311788]
- Suri RE, Schultz W. A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience* 1999;91:871–890. [PubMed: 10391468]
- Takikawa Y, Kawagoe R, Hikosaka O. A possible role of midbrain dopamine neurons in short- and long-term adaptation of saccades to position-reward mapping. *J Neurophysiol* 2004;92:2520–2529. [PubMed: 15163669]
- Thoroughman KA, Shadmehr R. Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 1999;19:8573–8588. [PubMed: 10493757]
- Toni I, Passingham RE. Prefrontal-basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: a PET study. *Exp Brain Res* 1999;127:19–32. [PubMed: 10424411]
- Toni I, Ramnani N, Josephs O, Ashburner J, Passingham RE. Learning arbitrary visuomotor associations: temporal dynamic of brain activity. *Neuroimage* 2001;14:1048–1057. [PubMed: 11697936]
- Toni I, Rowe J, Stephan KE, Passingham RE. Changes of cortico-striatal effective connectivity during visuomotor learning. *Cereb Cortex* 2002;12:1040–1047. [PubMed: 12217967]
- Tremblay L, Hollerman JR, Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate striatum. *J Neurophysiol* 1998;80:964–977. [PubMed: 9705482]
- Turner RS, DeLong MR. Corticostriatal activity in primary motor cortex of the macaque. *J Neurosci* 2000;20:7096–7108. [PubMed: 10995857]

- Waelti P, Dickinson A, Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 2001;412:43–48. [PubMed: 11452299]
- Wang M, Zhang H, Li BM. Deficit in conditional visuomotor learning by local infusion of bicuculline into the ventral prefrontal cortex in monkeys. *Eur J Neurosci* 2000;12:3787–3796. [PubMed: 11029649]
- Williams SM, Goldman-Rakic PS. Characterization of the dopaminergic innervation of the primate frontal cortex using a dopamine-specific antibody. *Cerebral Cortex* 1993;3:199–222. [PubMed: 8100725]
- Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA. Single neurons in the monkey hippocampus and learning of new associations. *Science* 2003;300:1578–1581. [PubMed: 12791995]
- Wise SP, Murray EA. Role of the hippocampal system in conditional motor learning: mapping antecedents to action. *Hippocampus* 1999;9:101–117. [PubMed: 10226772]
- Zheng T, Wilson CJ. Corticostriatal combinatorics: The implications of corticostriatal axonal arborizations. *J Neurophysiol* 2002;87:1007–1017. [PubMed: 11826064]

Fig. 1.

Behavioral paradigm and task periods. The top row shows schematics of the video screen as each trial progressed. Each large square represents the video monitor and the small squares depict the four targets (not to scale). At the center of the screen, the fixation point (unfilled circle), the instruction stimulus, and a gray square appeared sequentially. The disappearance of the gray square served as the trigger stimulus. The rectangle beneath the screen represents the start bar and the position of the hand indicates the monkey's starting position and response. Abbreviations: Del, delay; IS, instruction stimulus; Mov, movement; Prm, premovement period; Ref, reference period.

Fig. 2. Performance curves. **A–D.** Proportion of correct responses with 95% confidence limits (shaded background) for novel, arbitrary visuomotor mappings to the four targets. Each curve is color coded for response target, according to the inset between C and D. The number of trials indicates the accumulated number of presentations of a given stimulus within a block of trials, regardless of the number of intervening stimuli that map to other targets. Note that, at the beginning of a session, the correct response and each of the three incorrect responses begin near the chance level of 25% (dashed horizontal line). **E.** Average and 95% confidence limit for the four curves in A–D. **F.** Average performance curve and 95% confidence limit for four concurrently learned mappings. The number of trials indicates the average over the entire

session, with trials instructing familiar mappings eliminated. Abbreviations: LL, lower left; LR, lower right; UL, upper left; UR, upper right.

Fig. 3. Composite population activity for PMd (blue) and putamen (red) for selected subpopulations of neurons. The activity rates represent modulation above baseline activity levels, and the black bar labeled “baseline” at the top left of the figure illustrates the interval used to calculate baseline activity. The arrows at the bottom indicate the alignment event for each population histogram. The early, middle, and late phases of learning (top three rows) correspond to trials prior to the learning criterion trial, during the fastest increase in performance, and during asymptotic performance, respectively. Population histograms for familiar visuomotor mappings appear in the bottom row. The asterisks at the top mark the task periods used to select cells with activity that increased significantly during learning. These four selected subpopulations contributed to all four rows for each column of population histograms, with the number of cells above each asterisk (blue numbers for PMd; red numbers for the putamen). The histograms aligned on target acquisition show activity for the same subpopulation of neurons as the histograms aligned on movement onset. The early, middle, and late learning phases were based on a mean of 4.9–5.2, 2.6–2.7, and 5.8–5.9 trials per cell, respectively, for PMd, depending on the column, and 4.8–5.1, 2.2–2.5, and 5.0–5.7 trials, respectively, for the

putamen, also depending on the column. Familiar mappings were based on 7.0–7.3 trials per cell for PMd and 6.6–6.8 trials per cell for the putamen. Abbreviations: Del, delay; imp, impulses (spikes); IS, instruction stimulus; Mov, movement period; Mvt, movement; Prm, premovement period; Ref, reference period; TS, trigger stimulus.

Fig. 4. Population histograms for the entire recorded population in PMd (blue, $n = 120$ cells) and putamen (red, $n = 75$ cells). All plots show modulations relative to baseline level, for correctly executed trials only, with the standard error indicated by the shaded background. The top, right plot shows baseline activity levels for the early (E), middle (M), and late (L) phases of learning, and for familiar mappings (F). The histograms are smoothed with a 7-point moving average. Arrows mark prominent features of the population histograms. The early, middle, and late learning phases were based on a mean of 4.7, 2.5, and 5.5 trials per cell, respectively, for PMd, and 4.4, 2.0, and 4.7 trials, respectively, for the putamen. Format as in Fig. 3.

Familiar Mapping Comparison

Fig. 5. Changes in the difference between activity for novel mappings during each phase of learning and activity for familiar mappings. Note that the activity during learning comes to differ less from that for familiar mappings as learning progressed. Data from Fig. 4, showing the average differences per 10-ms bin across the displays centered on IS onset, movement onset, and reward delivery (solid lines) and the peak difference (dashed lines).

Fig. 6.

Change in population activity in relation to performance for the IS-on period, PMd. This figure complements Fig. 4, which divides the data into discrete learning phases. Cells could have contributed to this average for more than one mapping, provided that its spatial tuning was sufficiently broad, but averages limited to the best mapping for each cell looked nearly identical. The background shading shows 95% confidence limits for both curves. Trial 0 is the learning criterion trial of Wirth et al. (2003) for each individual mapping (analogous to Fig. 2E), which often occurred one or two trials after the monkey reached that criterion for the stimulus set as a whole (analogous to Fig. 2F). The blue curve and background show the mean and 95% confidence limits for activity during the IS-on period as learning progressed. The gray curve and background show the performance rate and 95% confidence limit. The gray axis label is for the bar histogram, which shows the number of observations contributing to each point of the average. This number changes because of variation in learning rates and the duration that recording continued after the monkey had reached asymptotic performance.

Fig. 7.

Error effects for early task periods. Left: In the format of Fig. 4, population histograms for PMd (blue) and putamen (red) for correct (thin lines) and incorrect (thick lines) trials. Shaded background shows the standard error for those parts of the plot of interest (green arrows). The data points in the top, left show baseline activity during error trials for the early (E), middle (M), and late (L) phases of learning, and for familiar mappings (F). (For baseline activity during correct trials, see Fig. 4). Right: The differences between correct and incorrect trials for PMd (blue) and putamen (red). Green arrows mark features of interest. Errors were so rare for familiar mappings that the resulting population averages revealed little beyond the normal range of trial-to-trial variation. Abbreviations as in Fig. 3.

Fig. 8. Error effects for late task periods. In the format of Fig. 7, except that the magenta arrow also marks a feature of interest. The inset above the curves for the late phase of learning (bottom row, left half) shows the standard error for the data marked by the magenta arrow, and the magenta lines indicate its time boundaries. Abbreviations as in Fig. 3.

Fig. 9. Mapping selectivity (S_i) during the pre-reward period. **A, C.** Data from PMd cells ($n = 77$) that had >5 impulses per second during the prereward period for any direction in any learning phase (early, middle, or late). Correctly executed trials only. **B, D.** Data from putamen cells ($n = 39$) selected according to the same criteria. Background shading shows the standard error. Activity was averaged within a 250 ms window, moved along the time axis in 10 ms steps. The computed S_i was smoothed with a 31 ms gaussian in A and B, but appears without such smoothing in C and D. Abbreviation: rew, reward.

Electronic Supplementary Figure 1.

Familiar-mapping data. Population histograms for the entire recorded population in PMd (blue, $n = 120$ cells) and putamen (red, $n = 75$ cells) in the format of Fig. 4, but for familiar-mapping trials divided according to the phase of learning the novel mappings in which they occurred. The early, middle, and late learning phases were based on a mean of 2.5, 0.7, and 3.8 trials, respectively, for PMd, and 2.0, 0.5, and 3.8 trials, respectively, for the putamen.

Electronic Supplementary Figure 2.

Data from a second monkey, in the format of Fig. 9, putamen only. In this monkey, the delay between the touch of the screen and the reward was fixed period of 1.22 s.