

MIT Open Access Articles

Stable Encoding of Task Structure Coexists with Flexible Coding of Task Sensorimotor Striatum

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Kubota, Y. et al. "Stable Encoding of Task Structure Coexists With Flexible Coding of Task Events in Sensorimotor Striatum." *Journal of Neurophysiology* 102.4 (2009) : 2142-2160. Copyright © 2009 The American Physiological Society

As Published: <http://dx.doi.org/10.1152/jn.00522.2009>

Publisher: American Physiological Society

Persistent URL: <http://hdl.handle.net/1721.1/64674>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Stable Encoding of Task Structure Coexists with Flexible Coding of Task Events
in Sensorimotor Striatum

**Yasuo Kubota,¹ Jun Liu,¹ Dan Hu,¹ William.E. DeCoteau,² Uri T. Eden,³
Anne C. Smith,⁴ and Ann M. Graybiel¹**

¹Department of Brain and Cognitive Sciences and McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, Massachusetts; ²Department of Psychology, St. Lawrence University, Canton, New York; ³Department of Mathematics and Statistics, Boston University, Boston, Massachusetts; and ⁴Department of Anesthesiology and Pain Medicine, University of California Davis, Davis, California

Running Head: Striatal Activity during Successive Learning

Address for reprint requests and other correspondence: A.M. Graybiel, Massachusetts Institute of Technology, 46-6133, 43 Vassar Street, Cambridge, MA 02139 (E-mail: graybiel@mit.edu).

ABSTRACT

The sensorimotor striatum, as part of the brain's habit circuitry, has been suggested to store fixed action values as a result of stimulus-response learning, and has been contrasted with a more flexible system that conditionally assigns values to behaviors. The stability of neural activity in the sensorimotor striatum is thought to underlie not only normal habits, but also addiction and clinical syndromes characterized by behavioral fixity. By recording in the sensorimotor striatum of mice, we asked whether neuronal activity acquired during procedural learning would be stable even if the sensory stimuli triggering the habitual behavior were altered. Contrary to expectation, both fixed and flexible activity patterns appeared. One, representing the global structure of the acquired behavior, was stable across changes in task cuing. The second, a fine-grain representation of task events, adjusted rapidly. Such dual forms of representation may be critical to allow motor and cognitive flexibility despite habitual performance.

KEYWORDS

procedural learning, basal ganglia, plasticity, cognitive flexibility, habit

INTRODUCTION

Habits, once learned, become notoriously difficult to change. Understanding the neural correlates of this remarkable stability is a central issue in the field of learning and memory and is critical to addressing the neurobiology of addiction and a range of clinical disorders in which behavioral fixity is a concern. Much evidence suggests that the striatum and associated cortico-basal ganglia loops are important for laying down habits and procedural memories through a dynamic engagement of successive cortico-basal ganglia loops culminating in fixed, stimulus-response representations in the sensorimotor striatum (Belin et al. 2009; Graybiel 2008; Voorn et al. 2004; Yin and Knowlton 2006). Experimental studies show that performing procedural tasks activates neurons in the sensorimotor striatum, and that large-scale changes in neural firing patterns can develop in the sensorimotor striatum as a result of learning (Atallah et al. 2007; Barnes et al. 2005; Blazquez et al. 2002; Costa et al. 2004; Kimchi et al. 2009; Schultz et al. 2003; Tang et al. 2007; Yin et al. 2009). Findings in human imaging experiments also demonstrate that changes in the sub-regions of striatal activation occur with learning, leading to activity in the sensorimotor striatum after learning (Doyon and Benali 2005; Graybiel 2005; Haruno and Kawato 2006; Hazeltine et al. 1997; Jueptner and Weiller 1998; O'Doherty et al. 2004; Poldrack et al. 2005; Willingham et al. 2002).

These experimental findings are in accord with reinforcement-learning models of cortico-basal ganglia circuits that implicate the striatum in developing behavioral policies and in updating them under the influence of inputs from the dopamine-containing substantia nigra, the thalamus and the neocortex (Daw and Doya 2006; Graybiel 2008; Niv et al. 2006; Redish et al. 2008). Reinforcement sensitive, flexible circuits successively involving limbic and associative cortico-striatal loops are thought to be engaged as habits are acquired, and sensorimotor cortico-

striatal systems are thought to be essential to the performance of habits through having activity that is impervious to changes in reinforcement conditions, thus allowing repetitive, habitual performance (Balleine et al. 2007; Packard 2009; White and McDonald 2002; Yin et al. 2004). Accordingly, the responses of single striatal neurons have been linked to the representation of sensory and motor events, timing, reinforcement-related parameters and action value (Aldridge et al. 2004; Aosaki et al. 1994; Blazquez et al. 2002; Kimchi and Laubach 2009; Lau and Glimcher 2007; Matell et al. 2003; Morris et al. 2004; Ravel et al. 2003; Samejima et al. 2005; Schmitzer-Torbert and Redish 2008; Schultz et al. 2003; Tang et al. 2007; Watanabe and Hikosaka 2005; Williams and Eskandar 2006). Moreover, when there is a change in the behavioral procedure required in a task, single unit activity in the striatum changes rapidly (Eschenko and Mizumori 2007; Pasupathy and Miller 2005; Watanabe and Hikosaka 2005). But what happens when the behavioral procedure to be performed remains the same, yet the sensory cues instructing the procedure change? This situation is also common in daily life, as in suddenly having a new store at the corner where we always turn right, or, for addicts, in having a new cue that indicates the availability of a sought-for drug.

We focused on this issue in the experiments reported here. We hypothesized that the introduction of a new set of cues requiring a new stimulus-response (S-R) association, in the absence of a requirement to change the S-R procedure, would trigger new patterns of task-related activity in the dorsolateral striatum, the striatal region thought to underlie performance of learned procedures and habits. We trained mice on a simple T-maze task instructed by sensory cues, and then switched the cues after they reached a high performance level with the first set of sensory cues. We recorded single unit and ensemble activity in the sensorimotor striatum throughout training as the mice first learned to turn right or left in response to auditory cues, and then as they

had to perform the same maze task after the sensory cues in the task were switched from auditory tones to tactile patterns. During acquisition of the first task-version, as predicted from earlier work in the rat (Barnes et al. 2005; Jog et al. 1999), the ensemble activity of the task-related projection neurons gradually came to accentuate the start and end of the maze runs as the mice learned the initial tone-instructed task. Remarkably, after the switch to the tactile version of the task, these ensemble firing patterns were largely maintained, despite the fact that the mice needed to learn a new S-R association in order to succeed in the new task version. This stability suggests that task-boundary representations acquired in the sensorimotor striatum through learning can generalize across differences in sensory cuing. By contrast, changes did occur in the sensorimotor striatum in the local event-related neuronal activity of subsets of individual projection neurons and, most strikingly, in the firing of fast-firing striatal interneurons. Based on these findings, we suggest that the sensorimotor striatum can maintain dual patterns of task representation so as to maintain a chunked task representation of the procedure (a stable, rule-based “policy”) while details of the firing patterns change as needed for new sensory guidance of performance (a representation allowing on-line updating). Such dual representational capacity may be fundamental to the functions of the sensorimotor striatum as part of the habit system of the brain. Disruptions of these functions may contribute to behavioral syndromes associated with repetitive and compulsive behaviors.

METHODS

Experimental Animals and Surgical Protocols

Fifteen C57BL/6 mice (3-5 months old males, 22-32 g) were trained in these experiments. Experimental procedures were approved by the Committee on Animal Care of the Massachusetts Institute of Technology. Three mice were trained behaviorally without neuronal recording. For implantation of the recording headstages, each of the 12 remaining mice was anesthetized with ketamine hydrochloride (100 mg/kg) and xylazine (20 mg/kg), and, under stereotaxic guidance, a burr hole was made in the skull at coordinates corresponding to the left dorsolateral caudoputamen (AP = +0.2 mm, ML = +2.4 mm). The dura mater was opened to allow later penetration of the tetrodes, and the headstage for neuronal recording was placed over the opening and was fixed to the skull with dental acrylic anchored to 6-7 watchmaker's screws (S. LaRose, Inc., Greensboro, NC) inserted into the skull near the headstage. A tungsten wire attached to one of the screws served as the ground.

Tetrode Headstage and Recording Protocols

We designed a miniature, light-weight implantable headstage for recording with tetrodes in the mouse striatum (Jog et al. 2002) (Specialty Machining, Wayland, MA). Four tetrodes made of twisted 10- μ m Ni/Cr wire (Kanthal Palm Coast, Palm Coast, FL) were held in a series of nested polyimide tubes (23, 30 and 36 gauges, Phelps Dodge, Inman, SC). The tubes attached to microdrives were held parallel to the screws that were used to advance each tetrode independently. The four recording tips of each tetrode (15-50 μ m apart) were exposed by cutting each to a length appropriate for the traveling distance of the microdrive. They were plated with

gold to reduce the impedance to 200-250 k Ω . The free end of each wire was connected to a circuit board mounted on the headstage.

Each tetrode was lowered, day by day, during the first post-operative week until it reached the dorsolateral caudoputamen (1.5-2.5 mm ventral to the brain surface). Neuronal activity was monitored continuously during each daily session via an oscilloscope and audio monitor to track the movement of tetrodes through the neocortex and underlying white matter. Typically, the tetrodes entered the caudoputamen on the fifth or sixth post-operative day, as judged by the loss of typical cortical activity, passage through a quiet white matter layer, and entry into a zone in which neurons tended to fire at very low spontaneous rates but could exhibit bursty phasic firing (Wilson and Groves 1981). The recording sites were then adjusted to maximize the number of units recorded by each tetrode (usually ca. 150-300 μm into the caudoputamen), and the tetrodes then were left in place as behavioral training began. The recording quality of the tetrodes was tested before each training session, and the locations of tetrodes were changed as little as possible. If activity was lost, fine adjustments were made to maintain recordings with multiple units during the span of training.

Behavioral Procedures

All behavioral training and recording sessions were conducted in a T-shaped maze made from black Plexiglas and consisting of an alleyway (55.2 cm long, 3.2 cm wide) with two short alleys (39.4 cm long). The alleys were surrounded by outward-sloping walls (60° angle). A sliding gate prevented the mouse from leaving the start site until trial start. Circular wells (1 cm diameter) located 2 cm away from the end of right and left choice arms could be filled from outside the maze, out of sight from the mouse, with chocolate milk delivered through a metal

tube. An audio speaker was located behind the choice point of the maze. Photobeam units (Med Associates Inc., St. Albans, VT) were placed along the maze to determine the times of the opening of the start gate, the onset of auditory and tactile cues at 22 cm from the start gate, and goal reaching. A CCD camera for monitoring and tracking the mouse (Cohu, Inc., San Diego, CA) and two axially placed, dimly lit projection lamps that illuminated reflective material on the headstage were located above the center of the T-maze. The maze apparatus was situated at one corner of the experimental room, and a number of extra-maze cues (e.g., computers, computer monitors, amplifiers and humidifier) were present at constant locations.

Mice were maintained on a food restriction regimen to hold their body weight at 85% of normal, and they were given 5-7 habituation sessions in the training apparatus prior to surgery. During these sessions, mice were allowed to move freely in the T-maze and to consume chocolate milk placed randomly on the floor of the maze. No conditional stimuli were presented to the mouse during the habituation sessions.

Once the tetrodes were in place, mice were trained to run the T-maze with either auditory cues or tactile cues as conditional stimuli (Fig. 1A). Twelve of the 15 mice were first trained on the task with auditory cues and then with tactile cues. In the auditory version of the task, the mouse was placed behind the closed start gate in the T-maze. A click (70 dB) was presented as a warning cue to indicate the start of each trial, and the start gate was then opened to allow the mouse to move along the maze. One of two auditory cues (a 1 or 8 kHz pure tone, 85 dB) was turned on when the mouse broke a photobeam placed 22 cm from the start gate. These cues instructed the animal about the location of reward in each trial: one tone indicated the availability of reward (chocolate milk) at the right goal, and the other tone indicated that the left goal was baited. The cues were presented according to one of 15 pre-determined pseudo-random

sequences, each of which contains 10 trials with each cue in any 20-trial block. The tone conditions were assigned randomly to the different mice, and remained constant for each mouse throughout training. The tone remained on until the mouse reached either goal or reversed its direction of movement before reaching the goals. If the mouse reached the correct goal as instructed by the cue, it received reward. After a trial, the mouse either returned spontaneously or was guided back by the experimenter to the start site for the next trial, which started after an inter-trial interval typically lasting 30-180 sec. Each mouse received up to 40 trials (the goal for single sessions), up to 20 with each tone, during each daily session. In the tactile version of the T-maze task, a reversible plastic insert with rough and smooth surfaces was placed on the maze floor before each trial. As in the auditory task, one cue (rough or smooth) indicated reward at the right goal, and the other reward at the left goal. The mouse stepped onto the insert at the same site at which the tone was turned on in the auditory task, and the insert, like the auditory cues, extended to both goals. All other aspects of the task were identical to those of the auditory task, except that the floor insert was cleaned with 70% alcohol or replaced by a pre-cleaned insert after each trial.

For each cue-version of the task, the criterion for task acquisition was 29 (72.5 %) or more correct trials in a 40-trial session ($P < 0.01$, chi-square test) on two consecutive days. The mice were given overtraining sessions until the performance exceeded the criterion level in 10 consecutive sessions. The numbers of sessions with criterial performance given to each mouse varied from 10 to 18.

Three mice were trained only on the tactile version of the T-maze task. Performance of these 3 mice during tactile training was compared with auditory-task performance of 12 mice trained first on the auditory task in order to compare the difficulty of the tactile and auditory

versions of the task. In the 12 mice trained on both task versions, we compared the behavioral learning and neuronal activity during successive acquisition of the two tasks.

Six of the mice trained on both versions of the task then received a further series of test sessions in which sessions with auditory and tactile cues were alternated daily or given within a single day.

Data Collection

During neuronal recording, a 16-channel preamplifier (1.7 g) connected to lightweight wires (Neuralynx Inc., Tucson, AZ) was attached to the headstage on the mouse. A tetrode channel without spike activity served as reference for differential unit recording. Unit activity was recorded from the 3 remaining tetrodes that did not have the reference channel.

Unit activity was recorded for each trial from 2 sec prior to the click warning cue to 1 sec after trial-end, defined by either goal reaching or reversal of running direction. Neuronal activity recorded on each recording channel (tetrode wire) was sent through the preamplifier with unity gain to two 8-channel programmable amplifiers (gain: 2000-10000, filter: 0.6-6 kHz) and then to a Cheetah Data Acquisition System (Neuralynx). The spike waveform of each spike exceeding a preset voltage threshold was digitized at 32 kHz and stored with a microsecond-precision time stamp. Spike waveforms on all recording channels and scatter dot plots were displayed in real time to estimate the number and quality of putative single units on the record. The unit activity on different selected channels was also monitored by an oscilloscope display and audio speaker.

Behavioral and stimulus event-markers sent from a separate computer controlling the behavioral training were recorded by the Cheetah Data Acquisition System with time stamps that were synchronized with time stamps of recorded spikes. The movements of the mice in the T-

maze were monitored continuously by a video tracker system that detected light from the headstage reflector and recorded its location at a sampling rate of 60 Hz.

Spike sorting and unit classification

Unit activity containing spikes of multiple neurons was sorted offline into putative single units (“clusters”) according to multiple spike parameters (e.g., peak height, valley depth, peak time) on the four channels of each tetrode (DataWave Technologies; Fig. 1B). The accuracy of spike-sorting and the quality of the single units were then evaluated by (1) t-tests for spike variability, (2) spike waveform overlays to confirm uniform waveforms for a given unit and different waveforms across units (Fig. 1C) and (3) autocorrelograms to detect the presence of an absolute refractory period (Fig. 1D). Based on these tests, clusters containing noise (artifacts and activity of other units) were excluded from further analyses. All accepted units were classified as putative medium-spiny projection neurons (MS), fast-firing interneurons (FF) or tonically-active interneurons (TAN) based on properties of discharge patterns using interspike intervals, autocorrelograms, and firing rates (Barnes et al. 2005) (Supplemental Fig. S1). We could not differentiate other known types of striatal interneurons (e.g., somatostatin-positive interneurons (Tepper and Bolam 2004)), and these neurons could have been included in one of the three categories we used here.

Task-related activity

The stability of spike waveforms during individual daily recording sessions permitted us to sum activity over individual trials to detect patterns of task-related activity even in the low-firing MS units. Peri-event time histograms (PETHs) were constructed for each unit by summing

the activity over all of the trials of a daily training session for the ± 1 sec period around each task event (the click warning cue, the opening of the start gate, the onset of locomotion, the onset of auditory or tactile cues, the beginning and the end of turning, and goal reaching). We then compared activities around each of the task-events monitored during the ca. 40-trial sessions to the activities recorded during a 500-ms pre-trial baseline period 1900 to 1400 ms before the warning cue delivery. Event-related neuronal responses were defined as increases or decreases in neuronal discharge greater than two standard deviations (SDs) around the mean baseline spike counts in at least 4 consecutive 20-ms intervals during ± 200 -ms period around each task event with spike counts of greater than 1. Units with such responses were classified as task-responsive units (Fig. 1E).

These task-responsive units were further categorized based on specific task events to which they responded (e.g., turn-related and goal-related units). These categories were named to reflect the peri-event period in which the unit activity deviates significantly from the mean, without making inferences about what unit activity represents. Units were classified in two ways: (1) units that responded to a given event regardless of presence or absence of responses to other events (Fig. 5C), and (2) units that responded only to a given event (Supplemental Fig. S6). The numbers of units with responses time-locked exclusively to either warning cue, gate opening, locomotion onset, out of start area or cue onset were too small for analysis ($n = 0-4$ per learning stage). There was a tendency for the proportion of units with single-event responses to increase with training ($P = 0.02-0.04$ with multiple chi-square tests), but we detected no significant differences in response properties between groups of units classified using the two methods, even for exclusively turn- and exclusively goal-related units for which sample sizes were relatively

large (Supplemental Fig. S6). We thus principally report results obtained with the first method in this study.

Units that did not have phasic activity reaching the threshold for being accepted as task-responsive were designated as non-task-responsive units. In previous work on the rat striatum, we encountered such units and found that despite their lack of significant phasic responses to specific task events, their activity did change, when they were considered as a group, over the course of learning (Barnes et al. 2005). Their average activity fell during acquisition from pre-training levels and tended to have highest levels before trials began. We found similar “non-task-responsive” units in the present study, and we treated them with analyses that we performed on task-responsive units. Units with fewer than 100 spikes during an entire training session were discarded from the analysis.

Behavioral analysis and learning stages

The behavioral performance of the mice was evaluated by three measures. The accuracy of behavioral responses (right or left turns as instructed by the sensory cues) was determined for each daily session. The running speed of the mice was calculated for the entire trial period (locomotion onset to goal-reaching) and for segments of the maze runs in between two successive task events (e.g., locomotion onset to cue onset, cue onset to turn onset, and the end of turn to goal-reaching). The reaction time from the opening of the start gate to locomotion onset was also measured.

For the analysis of changes in behavior and neuronal responses during the course of learning, data from the 12 mice given sequential auditory-tactile training were combined by employing learning stages based on response accuracy (% correct, Figs. 2, *C* and *D*, 3, 6, *A-B*, 7,

A and *B*, and 8, Supplemental Figs. S2-S4). Thirteen training stages were defined for the auditory task as follows: (A1) first acquisition session; (A2) second acquisition session; (A3) first acquisition session with response accuracy of $\geq 60\%$; (A4) first acquisition session with response accuracy of $\geq 70\%$; and (A5-A13) first to 18th sessions with above criterion ($\geq 72.5\%$) performance, combining every two such consecutive sessions (i.e., the first and second criterial sessions combined for stage 5, the third and fourth criterial sessions for stage 6, and the 17th and 18th criterial sessions for stage 13). For the tactile task, training stages were defined similarly, except that stages with $\geq 60\%$ and $\geq 70\%$ performance were skipped because of rapid learning. Thus, stages were: (T1) first acquisition session; (T2) second acquisition session; and (T3-T7) first to 10th sessions with at or above criterion performance by combining every two sessions. Learning-related changes in behavioral accuracy, running times, and locomotion-onset latency were analyzed with ANOVA. To analyze changes in neuronal activity during learning, ANOVA was performed to compare 6 phases of training (Figs. 3, *C* and *D*, 6*B* and 7*A* and Supplemental Fig. S3, *C* and *D*): early acquisition training on the auditory task (Stages A1-2), late acquisition training on the auditory task (Stages A3-5), early overtraining on the auditory task (Stages A6-9), late overtraining on the auditory task (Stages A10-13), acquisition training on the tactile task (Stages T1-2) and overtraining on the tactile task (Stages T3-7). Post-hoc tests were performed with Bonferroni correction.

In addition to these learning stages, changes in behavior and neuronal activity in response to the switch in modality of the conditional cues were tested by aligning data at the cue switch and comparing activity averaged for each of the last 3 or 5 sessions of auditory training and the first 3 or 5 sessions of tactile training (Figs. 2*E*, 4, 5, 6, *C* and *D*, and 7*G*, Supplemental Fig. S6). This analysis was necessary because the learning curves of each individual mouse were different,

resulting in different numbers of training sessions to reach the criterion of acquisition of the first auditory task-version. Thus, the last learning stage of the auditory version (A13) does not represent the last training session for all mice. By aligning daily sessions at the time of the cue switch and by detecting changes that occurred selectively on the first tactile session, we identified behavioral and neuronal responses to the introduction of a new set of cues in a different modality. Statistical differences in activity before and after the cue switch were tested with ANOVA performed to compare activity averaged over the last 5 auditory sessions and that averaged over the first 5 tactile sessions (Figs. 5A and 6E).

Stimulus and behavioral correlates of unit activity

To test whether differential unit activity occurs in relation to specific conditional cues or behavioral performance, trial-by-trial firing rates of each unit in 200-ms intervals before and after each task event were compared, with ANOVA, between (1) trials with stimuli indicating that right or left turns would lead to reward, (2) trials with right and left turns, and (3) trials with correct and incorrect behavioral responses (Fig. 7, A-C). Differences in the proportions of units with significant discrimination were tested with chi-square tests across peri-event windows and across learning stages.

To test whether activity of striatal units is modulated by the locomotor activity of the mouse, correlations were calculated between firing rates in each pre- and post-event window and (1) running speed in the same window, (2) total trial running time, and (3) inter-event durations, as well as between coefficients of variation calculated for each of these measures (Fig. 7, D and E). Chi-square tests were performed to test significance of differences in proportions of units with significant correlation across task time and over the course of learning.

Analysis of neuronal ensemble activity

Changes in striatal neuronal activity during the course of training were assessed by comparing, across training sessions, (1) average firing frequencies of ensembles of units during the task time, (2) proportions of single units with responses to task-events, and (3) properties of phasic task-related discharges of single units.

First, to determine whether there were changes in the overall patterns of neuronal activity recorded during task-time as training proceeded, we calculated average activity of ensembles of neurons (e.g., task-responsive MS units) during consecutive ± 200 -ms peri-event intervals (20-ms bins) for each learning stage in three ways. One, per-unit mean raw firing rates were averaged across all units of each neuronal sub-population (e.g., putative medium spiny neurons with task-related responses, fast-firing units with responses at goal reaching) recorded during each stage (Supplemental Figs. S3, A, C and D, S4B and S5B). This measure reflects the total number of raw spikes in the neuronal ensemble in response to task events. Two, the activity of each unit was first smoothed by using averages of 3 consecutive 20-ms bins running across the peri-event windows. These smoothed spike counts were normalized by setting the maximum spike count as 1 and the minimum as 0 and then were averaged across units (Supplemental Fig. S3B). This minimum-maximum normalization allowed us to detect relative amplitude of event-related responses during task time. Three, the baseline mean firing rates were subtracted from the smoothed spike counts, and these values were divided by the SDs for the entire trial time: $FR_N = (X - M_B) / SD_{B+T}$, where FR_N is the normalized firing rate, X is spike count per bin, M_B is the mean firing rate for the baseline, and SD_{B+T} is SD for the entire trial time including baseline. These normalized scores were averaged across units for each learning stage (Figs. 3, 4A, 5 and 6, Supplemental Figs. S2, S4A, S5A and S6). This normalization procedure was derived by

modifying Z-score conversion relative to the baseline mean and SD. This modification was necessary because some units did not have a single spike during the baseline period, thus yielding the SD of 0. These normalized scores reflect both magnitude of raw spiking activity and relative changes for each unit, and we accordingly present data primarily in this form. However, plots of ensemble activity made with the other averaging methods are also described in the main text or shown in figures. Learning stages with fewer than 6 recorded units were discarded from analysis because average activity can be dominated in a small sample by a few single units that have large responses and may not reflect activity of neuronal ensembles.

The mean scores were calculated for 500-ms pre-trial baseline periods and 200-ms pre- and post-event windows (Figs 3, *C* and *D*, 5*A* and 6, *B* and *E*, and Supplemental Fig. S3, *C* and *D*). These averages were submitted to ANOVA to detect differences in activity magnitude between different task events and across learning stages/sessions. Individual differences observed by these analyses were tested by post-hoc comparisons with Bonferroni correction.

Second, to determine whether there were changes in the proportions of units with phasic task responses, the numbers of units responding to any task events relative to that of all accepted units were compared across the performance-defined training stages by chi-square tests (Fig. 7*F*). Similarly, the proportions of units with activity related to a specific task event (e.g., opening of start gate, turn onset, goal-reaching) relative to all task-related units were compared across training stages, so that learning-related changes in neuronal activity specific to different task events could be detected (Supplemental Fig. S8).

Third, to test whether the properties of phasic responses by single neurons changed during learning and at the switch from auditory to tactile cues, peaks in peri-event histograms detected by an algorithm were compared across sessions and learning stages, as described

previously (Barnes et al. 2005) (Fig. 7G). Briefly, noisiness of activity of each unit was evaluated by calculating the SD of differences between the average raw peri-event histogram (± 1 sec) and the histogram smoothed with a Savitzky-Golay filter (window size = 300 ms). Units with the range of minimum and maximum spike counts greater than $5 \times \text{SD}$ were accepted, and for each of them, onset and offset of candidate peaks were defined as 10 consecutive, significant positive and negative slopes calculated with 21 neighboring time bin values, respectively. For each accepted peak, the shape (height, width, proportion of spikes that occurred inside the peak, etc.) and the timing (onset, offset and peak times) were measured. Proportions of spikes that occurred inside the peak (between onset and offset times of the peak) relative to all spikes recorded during trial times were calculated to quantify how spiking was devoted to a single event-related phasic response. Averages of these measures for neuronal sub-populations were compared across learning stages using ANOVA.

Analysis of putative single-unit activity tracked over sessions

To test whether the task-related activity of single units changed in response to the switch in task version, we first sought to identify putative single neurons (PSNs) recorded across daily recording sessions from tetrodes that remained stationary by calculating spike waveform correlations based on the methods described by Emondi et al. (2004). For each of 3281 individual units (including all 3 cell types, task-related and non-task-responsive units and units with fewer than 100 spikes) recorded on the 28 tetrodes that had at least a total of 30 units over the entire training, the average spike waveforms of each cluster recorded on the four channels of a tetrode were joined together consecutively (Fig. 1F), and a waveform correlation coefficient was calculated between the conjoined waveforms of all combinations of cluster pairs across two

consecutive sessions. Cluster pairs with correlation coefficients greater than 0.98 ($P < 0.01$ based on correlations for shuffled data, Supplemental Fig. S1D) were considered to be a match and were putatively identified as the identical neurons recorded on separate sessions (Fig. 1F). This method is based on the assumption that the relative size and shape of spikes recorded on four channels of a single tetrode are stable across different recording sessions, despite changes in the absolute values of spike parameters over time.

For each of 43 PSNs recorded on a minimum of 6 consecutive daily sessions, with at least 3 sessions of each cue type, differences in event-related responses in peri-event histograms between the two task-versions were tested in two ways. First, average firing rates during 200-ms windows before and after each task event were compared between auditory and tactile sessions by using Wilcoxon rank sum test (Figs. 9-12, Supplemental Figs. S9 and S10). This test was applied to both units recorded across the cue switch and those recorded during daily alternation of the two task versions.

To test in which pre-/post-event windows the activity differentiating the auditory and tactile sessions occurred at above chance level, the firing rates of each unit in all pre/post-event windows were normalized to the maximum firing rate across all sessions for the unit. These normalized firing rates were reassigned randomly to the same pre/post-event windows of all sessions, and a Wilcoxon test was performed to determine, for each window, how many of the 43 units showed significant discrimination. We repeated this procedure 1000 times to estimate p-values for results obtained for real data (i.e., whether there were significantly more units than chance that showed differential firing between the auditory and tactile sessions).

Second, to determine whether the peri-stimulus histogram of the spikes around each event changed significantly when the task-version was changed from auditory to tactile, we used

a Bayesian generalized linear model (GLM) approach (Nelder and Wedderburn 1972) (Fig. 10). Our approach relied on fitting the following three time series models: (1) a Bayesian Poisson GLM, (2) a Bayesian Poisson GLM with a step function at the transition from auditory to tactile versions, and (3) a Bayesian Poisson GLM with a step function at a location estimated by the model. For each model, we estimated goodness-of-fit by computing the deviance information criterion (Spiegelhalter et al. 2002), a Bayesian analogue of the Akaike information criterion. If the fit with step functions is better than without (i.e., DIC for models 2 and 3 lower than the DIC for model 1), and if the estimated step location in model 3 coincide with the actual auditory-to-tactile transition, then we concluded that the neural rate changed significantly in response to the cue switch.

Histology

Following the completion of training, each mouse was deeply anesthetized with Nembutal (150 mg/kg), and micro-lesions were made at the final recording sites by passing electrical current (200 μ A, 2 pulses/sec, 2 min) through all tetrodes. The mouse was then given a lethal dose of Nembutal and was perfused transcardially with 4% paraformaldehyde in 0.1 M NaKPO₄ buffer, in compliance with the recommendations on the Panel on Euthanasia of the American Veterinary Medical Association. Transverse 30 μ m frozen sections were stained with Cresylecht violet, and the sites of recording were found by analysis of lesions, track sites, and the extension of tetrodes measured post-mortem. These sites were in the dorsolateral striatum in all 12 mice (Fig. 1G).

RESULTS

Acquisition of the Auditory and Tactile T-maze Tasks

The mice learned the auditory and tactile versions of the maze-running task at nearly equivalent rates whether the conditional cues were auditory or tactile (Fig. 2A; $F = 1.34$, $P = 0.21$), and the average numbers of sessions required to reach the 72.5% correct criterion for behavioral acquisition were statistically indistinguishable (8.3 sessions for the tactile version and 9.0 sessions for the auditory version; $t = 0.175$, $P = 0.8638$, Fig. 2B). These results suggest that the two versions of the procedural task were nearly equivalent in task difficulty. The learning curves for the two task versions were different, however, when the tasks were presented serially (Fig. 2, A-C). The mice required significantly fewer training sessions to reach the behavioral criterion on the tactile version (average = 3.2) than during the preceding training on the auditory task version (average = 9.0, $t = 3.55$, $P < 0.005$; Fig. 2B). Running speeds were also sensitive to the cue switch (Fig. 2, C and D), especially in the period from cue onset to turn onset (Figs. 2E and 4B). After tactile training, the average trial durations and running speeds eventually reached about the same level as those in the late stages of auditory training.

These results suggest that there was behavioral transfer from the initial auditory task to the second tactile task. In order to ensure consistency across recording data collection and across mice, all 12 mice that received the two-version protocol were trained with the same task-version order (auditory, then tactile).

Recorded Units

We recorded the activity of 3683 well-isolated striatal units in the 12 mice trained successively on the two cue-versions of the T-maze task, over 23-89 daily sessions altogether (a

total of 500 sessions) lasting 34 to 77 calendar days (up to 217 days with alternation sessions). The average number of units recorded per session was 7.4 (maximum = 22). A total of 656 units were excluded from further analysis because of their low firing rates (< 100 spikes during a recording session). The daily recordings were made with minimal movement of the tetrodes (on average once in ca. 6 training sessions, ca. $90 \mu\text{m}$ per move). Due to these fine adjustments recordings during tactile training and alternation sessions were made at sites on average ca. 170 and $365 \mu\text{m}$ ventral to those made during initial auditory training, respectively. From our tracking analysis of units across days of training, described below, we estimate that we successfully recorded from some single neurons across multiple days. We included all well-isolated units in assessing the activity patterns of neuronal ensemble activity.

The 3027 units accepted for ensemble analysis were classified putatively into MS neurons, FF neurons or TANs (see METHODS and Supplemental Fig. S1). A majority of these (78%, 2347 of 3027) met the criteria for classification as MS neurons: they fired at low rates (1.3 Hz on average) during the pre-trial baseline period but could exhibit bursts of firing (up to >100 Hz) during the task (Wilson and Groves 1981). Units classified as FF neurons (18%, 539) had average baseline firing rates of over 17 Hz and short interspike intervals, typical of fast-firing striatal interneurons (Kawaguchi et al. 1995; Koós and Tepper 2002). We categorized 4% (141) of the units as TANs (Aosaki et al. 1994) but their small numbers prevented further analysis. Other types of interneurons have been identified in the striatum (Tepper and Bolam 2004), but we were unable to identify these; some of these may have been included in the categories defined here.

MS and FF units exhibiting phasic responses time-locked to one or more of the task events (click warning cue, gate opening, locomotion onset, conditional cue onset, turning, and

goal-reaching) were designated as “task-related” neurons (Fig. 1E, see also Figs. 6G and 9-12, and Supplemental Figs. S1, S9 and S10). By contrast, neurons lacking such phasic responses to task events were designated as “non-task-responsive” neurons.

Ensemble Activity of Striatal Projection Neurons during Acquisition of the Auditory Task

The ensemble activity of the task-related MS neurons developed an emphasis on the early and late parts of the maze runs during the course of training on the auditory task (Fig. 3A and Supplemental Fig. S2). At the start of training, ensemble activity occurred throughout the maze runs, but this pattern quickly gave way to one in which activity was highest during the gate-opening/locomotion onset period and then again in the late run period, especially from turn to goal reaching. There also was a more variable band of activity just before the time of turn onset. ANOVA performed to compare the average normalized activity in 200-ms time windows before and after each task event with activity during the pre-trial baseline period in each successive phase of auditory training confirmed these changes (Fig. 3C). This pattern was evident whether we normalized the data relative to trial-time standard deviations and baseline mean firing rates (Fig. 3A) or to per-neuron minimum-to-maximum firing rates (Supplemental Fig. S3B), whether we plotted average raw activity without any normalization (Supplemental Fig. S3, A and C), whether the mice made the correct behavioral responses or not (Supplemental Fig. S4), and whether the data were plotted by learning stages (Fig. 3A and Supplemental Figs. S3, A-B, and S4) or by daily sessions (Supplemental Fig. S5).

Heightened activity early and/or late in the runs was visible in the normalized and raw activity plots for the auditory task in 8 of the 12 individual mice (data not shown). Of the 4

remaining mice, 2 did not show this pattern, and the other 2 did not have sufficient unit recordings to permit ensemble analysis.

The MS neurons that did not show phasic task-related responses fired, throughout learning, at lower rates during the maze runs, relative to the pre-run baseline levels ($P < 0.05$, Fig. 3, *B* and *D*). The average raw firing rates showed similar reductions (Supplemental Fig. S3A). The raw activity during the baseline period itself tended to fluctuate, first rising then falling, during acquisition and overtraining ($P = 0.055$). It was relative to this fluctuating pre-run activity that the spiking of the non-task-responsive units was suppressed during the runs as the procedural behavior was acquired.

Ensemble Responses of Striatal Projection Neurons after the Switch of Cue Modality

Remarkably, the dominant pattern of early and late firing that developed during training on the auditory task persisted virtually unchanged after the modality of the conditional cues was changed from auditory to tactile (Fig. 3, *A* and *C*). The pre-switch beginning-and-end pattern was also visible in the raw firing rates after the switch; there was even stronger firing at gate opening/locomotion onset (Supplemental Fig. S3A).

Because individual mice reached the criterion for acquisition and overtraining on the first task-version after different numbers of sessions, it was important to realign the ensemble data with respect to the time of the version switch for each individual and to test directly for possible changes in ensemble firing at the cue switch. We did so, and then computed the day-by-day ensemble firing of the MS neurons for the last 5 days of auditory training and the first 5 days of tactile training (Figs. 4A and 5A) and for blocks of 5 trials for 6 training sessions around the cue switch (Fig. 5B). The results were striking: there was almost no difference in the ensemble firing

pattern before and after the switch, despite the large changes in trial duration and running speed that occurred at the cue switch, particularly during the cue to turn period (Figs. 2E and 4B). No significant differences occurred for any pre/post-event period between grouped data combining the last 5 auditory sessions and those combining the first 5 tactile sessions (Fig. 5A, $P > 0.05$). Raw spike plots did show that the activity at locomotion onset was high after the switch ($P < 0.05$), but such increased activity already appeared during the last sessions of auditory training and continued to increase across the switch, suggesting that this activity developed with extended training on the T-maze task, not in response to the cue switch.

We also plotted separately the ensemble activity aligned at the switch for subsets of units with responses to individual task events (e.g., to warning cue). Units in a given subset were included even if they also had responses to other events. These switch-aligned plots showed equally striking results (Fig. 5C): cell ensembles with responses at the earliest events in the task—warning cue, gate opening and locomotion onset—and units with ensemble responses at the latest events in the task—turn and goal-reaching—exhibited similar patterns of firing across the switch in cue modality. The activity of units with responses exclusively during turn and those with responses only around goal reaching was also unchanged at the cue switch (Supplemental Fig. S6). The out-of-start units, however, largely lost their responses to events after start and responded more robustly from gate opening to the post-start period (Fig. 5C). The activity of the cue-responsive units seemed, if anything, less stable after the switch, but the numbers of units were small.

These findings demonstrate that the overall ensemble pattern of early and late firing persisted, despite changes in details of the response patterns, especially in responses in the mid-run intervals. The MS units without phasic task-related responses also retained their suppressed

levels of spiking during the maze runs throughout the tactile training (Fig. 3B, $P < 0.05$ compared to baseline firing rates, ANOVA).

Ensemble Activity of Fast-firing Neurons during Training on the Successive Task-versions

The ensemble of FF units recorded during acquisition of the first task-version acquired a beginning-and-end firing pattern that was similar to that of the task-related MS neurons, although the late activity gradually weakened during overtraining (Fig. 6, A and B, and Supplemental Fig. S2). Also like the MS neurons, the FF neurons maintained the enhanced early and late activity across the cue switch. In striking contrast to the MS neurons, however, the FF ensembles abruptly exhibited at the onset of tactile training a sharp “new” response at cue onset (Fig. 6, C and E).

We tested whether this response was related to the reduction in running speed at the onset of new tactile cues by measuring average firing rates in 5-trial blocks across the first 3 tactile sessions. We found no correlation with the average running speeds during these blocks (Fig. 6F) and detected no shift in activity within single sessions (Supplemental Fig. S7). Furthermore, the response at cue onset was detectable in ensemble activity plotted after removing trials with long (mean + 2SD) cue-turn run durations that had occurred early in tactile training (Fig. 6D). Seven of 12 individual task-related FF units recorded during the first tactile session responded to the cue onset, and only one of them exhibited a significant correlation between firing rates and running speed in the post-cue period (e.g., Fig. 6G). This intense peri-cue response persisted for about three days but then became weaker and disappeared during overtraining, but this diminution did not parallel the changes in running speed of the animals ($P = 0.108$). These results suggest that the appearance of peri-cue responses by the FF neurons likely did not solely

reflect behavioral changes in response to the new cues. Finally, the goal-related activity of the FF neurons, which had decreased toward the end of the auditory training, was reinstated at the onset of the tactile training ($P = 0.059$ for post-goal period), a pattern also not observed in the MS ensembles. These results suggest a sensitivity of the task-related FF neurons to the newly presented tactile cue that appeared after the task-switch and to rewards newly associated with these new instruction cues.

The numbers of FF neurons that did not have phasic task-related activity were too small for a full analysis, but we saw no indication of suppression of in-trial activity even in stages with moderate numbers (9-14) of units in stages A1-A8 (data not shown).

Behavioral Correlates of Striatal Neurons Insensitive to Cue Switch

We found gradual but clear changes during training in the proportions of MS and FF units that fired differentially for right turns and left turns, but these changes did not exhibit sudden shifts at the onset of tactile training (Fig. 7A). The proportions of MS units with differential activity around the onset of turns (and at goal reaching) increased significantly during training (from ca. 0 to 15%, $P < 0.00-0.01$), but the proportions of units with responses around the completion of turns was stable at ca. 25% throughout training. FF units gradually developed turn-discriminative activity over the course of learning, from almost 0% early in auditory training to 20-25% by the end of tactile training in periods from pre-turn onset to post-goal reaching ($P < 0.00-0.05$), again without evidence of any relation to the cue switch. It was as though these changes were related to the general process of learning rather than to the specifics learned.

The task-switch also did not appear to influence the total percentage of task-related MS neurons (40-50%), and this proportion did not change significantly during training (Fig. 8, A and C). The proportions of FF units that exhibited task-related responses, by contrast, increased from 45% on the first day of training on the auditory task version to 92% late in overtraining and then stabilized at around 70% thereafter (data not shown).

Lack of Differential Responses to Cue Submodalities and Response Accuracy

Surprisingly few units showed activity differentiating the submodalities of the two conditional auditory cues (high and low pitch) or the two different tactile cues (rough or smooth), and very few showed activity that was correlated with behavioral accuracy during training on the two task versions. We found no evidence for changes in these proportions across training stages or across the cue switch. Fewer than 1% of task-related MS and FF units recorded over the entire phases of training fired differentially to the two auditory cues or to the two tactile cues that signaled the baited goal location (Fig. 7B). Similarly, only about 1% of units responded differently in correct and incorrect trials during each pre-/post-event period, except at goal reaching (Fig. 7C). About 3-11% of the units showed significant correlations between firing rates and trial duration (Fig. 7D), or between firing rates and running speed (during periods of movement; Fig. 7E). No consistent relationship between degree of per-session variability in spiking and running speed was found.

Resetting of Signal-to-Noise Ratio in Activity of Projection Neurons at Cue Switch

Despite the stability of the ensemble responses of the MS neurons across the task-version switch, the prominence of the phasic responses of individual MS neurons, relative to the spike

firing during other parts of the runs, was sensitive to the switch (Fig. 7G). The measure of how many of the spikes recorded for a single MS neuron during each maze run were concentrated in phasic peaks around task events (the spike proportion index) increased significantly during the course of auditory training ($F = 3.65$, $P < 0.0001$), and then dropped abruptly at the switch, as indicated by significant differences between the last 5 auditory sessions and the first 5 tactile sessions (Fig. 7G, right, $F = 3.91$, $P < 0.001$), and then rose again to the level that had been reached during overtraining on the auditory task. Thus the “noise” in the spike production of the task-related MS neurons, considered as responses outside their phasic response periods, first fell during initial task acquisition, then rose at the switch, and then fell again as the second task version was acquired.

Tracking Single-unit Activity across Training Sessions

The changes in phasic response properties of the MS neurons at the cue switch suggested that, despite the retention of the global ensemble firing patterns after the cue switch, there were nevertheless changes at a neuron-by-neuron level. We analyzed in detail the waveforms of individual units and identified 43 putative single neurons (PSNs) that could be tracked over the minimum of 6 consecutive recording sessions either across the initial switch from the auditory version to the tactile version or during the period of daily alternation between the two versions (at least 3 with each cue modality). These PSNs were recorded by tetrodes that did not move, and they exhibited spike waveforms on all 4 tetrode channels that were nearly identical across sessions, as indicated by high correlations ($r > 0.98$) of composite waveforms in consecutive sessions (Fig. 1F). Of the 43 PSNs, 32 (74%) were classified as MS type, 8 (19%) as FF type, and 3 (7 %) as TAN type.

Despite the small numbers of these PSNs, their ensemble activity showed the heightened early-and-late patterning seen for the total MS population, and notably the stability across the cue switch was comparable to that found for the total population (Fig. 8A). We found high stability also when we tracked the ensemble activity of task-related units ($n = 14$) just from the last auditory session to the first tactile session (Fig. 8B). Task-related responses of these PSNs were either present or absent in both sessions in 83 of 98 (84.7%) peri-event periods (7 task events x 14 units, $P = 0.001$ based on 1000 shuffling of the pre-/post-switch pairings).

Remarkably, despite this overall stability, in the sample of 43 PSNs followed across the switch or during daily alternation, we found activity that differentiated the auditory and tactile task versions for all peri-event periods analyzed, not just for activity around cue presentation. The results of 1000 bootstrap calculations (see METHODS) indicated that the numbers of PSNs exhibiting task-version-related differences in activity were significantly above chance not only for periods before and after cue onset ($n = 6$ and 7 and $P = 0.002$ and < 0.001 , respectively), but also for periods after the gate opening ($n = 4$, $P = 0.04$), after right and left turn ($n = 4$ and 7 , $P = 0.04$ and < 0.001 , respectively), before right and left goal reaching ($n = 6$ and 4 , $P = 0.002$ and 0.042), and after right and left goal reaching ($n = 8$ and 11 , $P = 0.001$ and < 0.001 , respectively).

Of the 32 MS-type PSNs, 5 responded differently to the auditory and tactile cues, suggesting the presence of activity responsive to a particular cue modality or to a novel cue. A neuron from this group, shown in Fig. 9 as an example, fired at higher rates during the pre- and post-cue periods with auditory cues than those with tactile cues over the course of 65 training sessions (Wilcoxon rank-sum test, $P < 0.0001$ for both pre- and post-cue periods). During this entire time, this neuron responded robustly during right turns, but this activity did not change

with the switch in cue modalities ($P = 0.13$). Supplemental Fig. S9 illustrates the activity of a second PSN that showed a similar pattern of activity.

In 13 of the 32 MS-type PSNs, we found changes in activity related to task events other than cue onset when the modality of instruction cues changed from auditory to tactile. Fig. 10 illustrates the activity of one such neuron. The goal-related responses of this neuron decreased during the last 3 auditory sessions, but suddenly rebounded on the first day of tactile training and then again gradually decreased as training on the tactile task progressed, without changes in responses to other task events. To test whether the rebound did represent a sudden shift in activity, we fitted these data with Bayesian generalized linear models with and without an assumed change-point (see METHODS). The models with a change-point performed significantly better (deviance information criterion [DIC] ≈ 46) than the model without a change-point (DIC ≈ 64), suggesting that the rebound of goal-related activity was related to the cue switch.

Five PSNs classified as FF-type units were tracked across the cue switch. Of these, 3 exhibited significantly greater responses to the tactile cues after the switch than to the auditory cues before the switch, in agreement with the appearance of robust peri-cue ensemble activity of FF units early in tactile training. One such neuron is shown in Fig. 11. It fired more during the period preceding the onset of tactile cues than during that preceding the onset of auditory cues ($P < 0.005$). This neuron also exhibited differential activity after left turn (auditory $>$ tactile, $P < 0.005$). Heightened discharges around tactile cue onset were found in two other units (e.g., Supplemental Fig. S10). Differential activity in auditory and tactile sessions was not detected in 3 FF-type PSNs recorded during the alternation phase.

None of the units with detectable activity changes, including those shown in Figs. 9-11, showed consistent relationships between trial-by-trial firing rates and either running speed or trial duration. Post-cue running speeds were especially reduced during the first 5-10 trials of the first few tactile sessions (Fig. 4B), but the trial-by-trial firing rates did not parallel such patterns (Supplemental Fig. S7). Thus the changes in unit responses at cue switch were not directly related to the reduction of running speed that occurred early in tactile training

The task-related activity of about half of the PSNs followed across the cue switch was stable. Fig. 12 gives an example of such stability for a neuron that exhibited nearly identical activity pattern throughout the 10-day period around the switch.

DISCUSSION

Two of the most characteristic features of procedural memories are that they are acquired slowly and that they can be maintained for long periods of time once they are acquired. The striatum has been shown to be engaged during this prolonged type of learning, but paradoxically, the striatum also has been implicated as critical for behavioral and cognitive flexibility and for updating of motor performance. How these apparently disparate functions of the striatum can coexist is not understood. The most frequent suggestion for their co-occurrence, based on lesion studies, is that activity in different parts of the striatum and their corresponding corticostriatal loops underlie these capacities. Evidence from lesion studies suggests that the sensorimotor striatum is mostly responsible for prolonged, stable procedural or habit memories, and the associative striatum is mostly responsible for behavioral flexibility (Cools et al. 2006; Daw et al. 2005; Dayan and Balleine 2002; Everitt and Robbins 2005; Graybiel 2008; Kimchi et al. 2009; Ragozzino 2007; Yin and Knowlton 2006). Our findings suggest that neuronal activity patterns consistent with both stability and flexibility coexist in the sensorimotor striatum during procedural learning. Thus both the generation of behavioral policies and their updating can be implemented by this large striatal territory.

Global and Local Task Representations in the Sensorimotor Striatum

We attempted to set up an experimental situation in which the motor performance requirements and reward outcomes of the auditory and tactile task-versions were similar, so that the cue change would be the single variation when the task-versions were switched. Moreover, we did not signal the task-version switch to the mice, so that we could observe the uninstructed behavioral responses of the mice and the corresponding responses of their striatal neurons when

the switch occurred. Given evidence that the dorsolateral striatum is necessary for executing S-R associations (Atallah et al. 2007; Balleine et al. 2007; White and McDonald 2002; Yin and Knowlton 2006), and evidence that the projection neurons in this striatal region acquire new patterns of responding in similar T-maze tasks in the rat (Barnes et al. 2005; Jog et al. 1999), a plausible outcome of this cue-switch design was that the projection neurons in the dorsolateral striatum would revert to their pre-acquisition activity patterns at the task-version switch, and that they then would acquire new patterns based on the new S-R associations learned in the second task-version. Standard reinforcement learning theories similarly would suggest that the activity patterns would change after the switch as new S-R associations had to be acquired (Daw and Doya 2006; Sutton and Barto 1998).

What we found, instead, was that the global pattern of strong, early and late firing of the projection neurons was largely maintained across the cue switch: the switch was scarcely discernable in the ensemble activity. Yet we did observe small but significant changes in the specific task-related responses of as many as 40% of the putative projection neurons tracked across days before and after the cue change, together with a decline in signal-to-noise ratios, reflected by the increase in spikes per trial outside of these phasic event-related responses. Moreover, the putative fast-firing interneurons recorded exhibited new ensemble responses after the switch.

These findings suggest that in the sensorimotor striatum, two types of response pattern may have been acquired as a result of the T-maze training and over-training. One, represented by the early and late task-bracketing firing pattern, may represent the boundary-structure of the task as a whole. This pattern did not need to be reconstructed as a result of the cue switch, and it was not. The second response pattern was a more detailed and flexible form of local representation of

task events, evident among some of the projection neurons and many of the fast firing interneurons, updating detailed firing patterns when the cues were changed. We propose that this combination of characteristics could be critical to the formation of striatum-based representations of learned procedures and to their maintenance and updating across time.

Global Task Structure Representations in the Sensorimotor Striatum Can Favor Behavioral Flexibility as Well as Behavioral Stability

Ensemble spike activity, relatively evenly distributed throughout the maze runs early in training, fell in mid-run and increased early and late in the runs with extended training to and beyond the acquisition criterion. These results for the mouse are similar to those observed previously in rats performing a T-maze task similar to the auditory version used here (Barnes et al. 2005). We have suggested that this pattern could represent a marking of the boundaries of the task acquired as behavioral exploration yields to behavioral exploitation (Barnes et al. 2005). We did not have recordings during other task versions that would allow testing of whether the task-related firing represented a spatial code (as in place-cell firing) as opposed to another representation of the procedure. For example, the heightened firing near the start of the maze runs could be related to prediction of eventual reward at the end of the runs, and the heightened late firing could be related to anticipation of reward receipt after the turns, the actions that determined reward outcome. Alternatively, the firing patterns might be related to a more abstract representation of the task for the purpose of facilitating release of the behavior after it was chunked. These different possibilities are not mutually exclusive and could all contribute to the firing patterns observed.

If this early-late firing pattern indeed provides a representation related to the global structure of the task, then it could contribute to the behavioral savings shown by the mice as they learned the second task-version. The low level of activity of the non-task-responsive projection neurons likely further promoted this stability by producing a context-specific increase in the signal-to-noise ratio of the task representation (Barnes et al. 2005; Tang et al. 2007). Moreover, by having a global representation of the task structure in place, the learning of the second task could have been accelerated.

The mid-run decrease in firing of the task-related neurons could also have contributed to the rapid and fluid acquisition of the second task-version. If the striatal representation of this part of the procedure was refined so as to be represented efficiently by a smaller number of “expert neurons” (Barnes et al. 2005; Jog et al. 1999), then most of the neuronal population should have been available to register the new cues during training on the second task-version and to engage in new S-R associations within the framework of the general representation given by the early-late firing pattern. Our finding of a significant increase in the spike proportion index for the projection neurons accords with this view: when these task-related neurons developed their phasic responses, they tended to decrease those spikes not in the phasic response profiles.

This interpretation is consonant with findings from human brain imaging experiments suggesting that activity in the striatum is critical for transfer of learning requiring updating (Dahlin et al. 2008) and for the filtering out of extraneous stimuli that is required for focused, efficient performance (McNab and Klingberg 2008). Our results demonstrate a potential neural correlate of this attentional filter adjustment during learning: we propose that during procedural learning, the sensorimotor striatum makes available neuronal space for representing novel events during a task and for switches in task demands. Thus the beginning-and-end accentuation that we

observed in the ensemble firing pattern may not only provide a general structure for the boundaries of the task and thus the performance strategy, but also enable new associations to form.

The stability of ensemble activity across the cue switch may help to account for another remarkable characteristic of procedural learning, that extended training on a task actually makes it easier for subjects to switch their behavior when the cues or rules or strategies of the task are reversed. Striatal inactivation greatly impairs this paradoxical overtraining reversal effect (Van Golf Racht-Delattour and Massiou 2000). If, as we suggest, this pattern both provides a representation of the general structure of the task (the rules) and represents a release of the attentional demands of performance after prolonged learning, then learning new S-R associations to the new cues in a reversal paradigm should be facilitated. This representation, combining task bracketing and local flexibility, could be critical also for the phenomenon of Pavlovian-instrumental transfer, which in rats depends on the dorsolateral striatum (Corbit and Janak 2007).

In the human, the putamen, which includes the sensorimotor striatum, is active under conditions requiring set-shifting (Monchi et al. 2006) and is required for reversal learning based on a switch in sensory cues, as opposed to a switch in performance strategy (Cools et al. 2006). Like the striatal region from which we recorded in the mice, the putamen is the part of the striatum most associated with the execution of slowly acquired procedural memories (Doyon et al. 2003; Floyer-Lea and Matthews 2004).

Interestingly, the relative stability that we observed here in the striatal recordings after a cue switch contrasts sharply with the lack of stability of the otherwise similar acquired firing patterns that we found in the dorsolateral striatum when rats encountered a switch to extinction training, in which the T-maze cues remained the same but reward was removed (Barnes et al.

2005). During extinction, the originally acquired beginning and end pattern nearly disappeared. This difference suggests that the cue switch was not adequate to induce a full change in firing patterns that are nevertheless vulnerable to reversal or suppression after major context changes including changes in reward contingencies and the need to modify the learned behavioral procedures. It is possible, and even likely, that such ensemble patterns would change in medial parts of the striatum, associated by experimental studies with modulating flexible behavior (Kimchi et al. 2009; Ragozzino et al. 2002; White 2009; Yin et al. 2009).

Changes in Fine-scale Task-related Activity in the Sensorimotor Striatum also Favor Behavioral Flexibility

Despite the relative stability of the MS neuron ensemble-level firing patterns recorded across the cue-switch, some putative single MS neurons tracked did change their response properties after the switch. This finding suggests that there were adjustments of the fine-scale selectivity of the neuronal firing patterns underlying the acquired ensemble patterns. Significant changes in peri-event spike counts occurred not only at cue presentation, but also at other time-windows in the task, including gate opening, turn offset and goal-reaching. The fact that these fine adjustments of task-firing occurred across multiple parts of the task after the switch suggests that the overall representation of the task may have been destabilized. Our criteria for classifying units as putatively being the same on successive days depended on comparing the waveforms recorded on all four channels of recording tetrodes unmoved for the successive days in question, and requiring that these waveforms be similar with a probability of <0.01 . As this is a probabilistic, not categorical, classification, we interpret our results in accord with this probabilistic bound. It seems highly likely, however, given the supportive evidence in a few

instances from Bayesian GLM analyses, that many and we suspect most of the putative single neurons were indeed single neurons.

It is difficult, if not impossible, to quantify all aspects of motor activity in a study with freely moving mice and to identify behavioral correlates of neuronal activity, but parameters of the maze running behavior that we could measure failed to account for these shifts in putative single unit activity. Most notably, the cue switch produced transient changes in running speed, but we found no consistent correlation between unit activity and running speed around the switch, or at other times during training.

We found surprisingly little evidence that individual MS neurons there code for the sub-modalities of the auditory and tactile cues that instructed the mice where to go to receive reward. Consistent with previous studies (Barnes et al. 2005; Berke et al. 2009; Jog et al. 1999), stimulus-driven activity was rare, except for that of the FF units at the onset of tactile training. A majority of “cue-period” units responded around, not only after, the cue onset, suggesting that their activity may have been anticipatory of presentation of critical stimuli and that it did not reflect stimulus properties such as rough or smooth. By contrast, we found that many units became differentially active in relation to turning direction as a result of training, and these acquired responses were insensitive to cue switch. These action-related responses may have reflected decision-making, behavioral selection or evaluation of completed behaviors based on incoming pre-processed cue-related information—or combinations of these functions. We may have missed neurons with more specific sensory coding properties for technical reasons, but the combination of a lack of detailed differentiation of sensory cue properties with a significant buildup of turn-direction selectivity supports the view that neurons in the dorsolateral striatum

participate especially in action selection as a new S-R association is acquired rather than in the encoding of conditional cues.

Detection of New Task-Version Cues by Fast-Firing Striatal Neuron Ensembles

In contrast to the MS neuron ensembles, the FF neuron ensembles abruptly gained a brief, time-locked response to the new cue after the cue switch. Parallel changes were found in putative FF-type single units tracked across the cue switch. These sharp responses were not directly related to changes in running speed, and they remained prominent for several days of training on the second task-version, suggesting that they did not simply reflect a novelty response or locomotor behavior, but a more protracted response. This response could have been a somatosensory response to the tactile cue, as suggested for the rat (Thorn et al. 2008). There is a strong projection to the rodent dorsolateral striatum from somatosensory and motor cortical areas (Alloway et al. 2006; Brown 1992; McGeorge and Faull 1989). The response could also have been related to the evaluation of the novel cuing conditions. However, the phasic FF cue response declined over succeeding days of training, and differential activity was not found in the few putative single FF units tracked during the period of daily alternation between auditory and tactile versions. These results suggest that the responses of FF neurons to tactile cues, induced by a shift in task cuing or a novel encounter specifically with tactile stimuli, are modulated as the behavioral procedures are adjusted with the new cues. The FF neurons also showed an augmented ensemble response at goal-reaching when the new task-version began, and this response was similar to the FF goal-reaching response acquired during training on the first task. Both of these characteristics of the FF ensemble activity, difficult to detect in the MS ensemble activity, would be predicted by reinforcement learning models.

The FF interneurons of the striatum have been singled out as receiving fast input from the neocortex and as producing powerful feed-forward inhibition of other striatal neurons, including projection neurons (Kawaguchi et al. 1995; Koós and Tepper 2002). It has been suggested that the FF neurons may be part of a filtering mechanism for cortico-striatal inputs, serving as an intrastriatal selection mechanism (Berretta et al. 1997; Parthasarathy and Graybiel 1997). This function fits well with the suggestion, made here, that the decreased mid-run firing that occurred in the striatum as learning occurred may represent a form of attentional filtering favoring behavioral and cognitive flexibility upon task switch.

Building and Updating Cognitive Policies

At the core of reinforcement learning theory is the notion that behavioral policies are built-up through trial and error search, but that these are sensitive to reinforcement-based updating. Our findings suggest that the sensorimotor striatum at once may develop a policy that constitutes a structural boundary-representation of procedural tasks and habits, and at the same time develop fine-scale forms of representation that favor updating after the initial learning has occurred. The global structural representation of the task may help to stabilize cognitive strategy for the habitual performance of the procedure, while both the reduction in neuronal firing during each time-step of behavioral execution and the local response changes across the cue-switches provide cognitive flexibility. Such cognitive functions are thought to be impaired in disorders affecting cortico-basal ganglia circuits, but they have been difficult to analyze electrophysiologically because of the need for long-term chronic recording. Our demonstration that such recordings are feasible in the mouse opens the way to analysis of these circuit functions

in genetically engineered mouse models of neurologic and neuropsychiatric disorders and addiction.

ACKNOWLEDGMENTS

We thank Patricia Harlan, Henry F. Hall and Daniel J. Gibson for their help.

GRANTS

This work was supported by National Institute of Neurological Disorders and Stroke Grant P50-NS038372 and by the Office of Naval Research Grant N00014-04-1-0208.

REFERENCES

- Aldridge JW, Berridge KC, and Rosen AR.** Basal ganglia neural mechanisms of natural movement sequences. *Can J Physiol Pharmacol* 82: 732-739, 2004.
- Alloway KD, Lou L, Nwabueze-Ogbo F, and Chakrabarti S.** Topography of cortical projections to the dorsolateral neostriatum in rats: multiple overlapping sensorimotor pathways. *J Comp Neurol* 499: 33-48, 2006.
- Aosaki T, Tsubokawa H, Ishida A, Watanabe K, Graybiel AM, and Kimura M.** Responses of tonically active neurons in the primate's striatum undergo systematic changes during behavioral sensorimotor conditioning. *J Neurosci* 14: 3969-3984, 1994.
- Atallah HE, Lopez-Paniagua D, Rudy JW, and O'Reilly R C.** Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. *Nat Neurosci* 10: 126-131, 2007.
- Balleine BW, Delgado MR, and Hikosaka O.** The role of the dorsal striatum in reward and decision-making. *J Neurosci* 27: 8161-8165, 2007.
- Barnes T, Kubota Y, Hu D, Jin DZ, and Graybiel AM.** Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* 437: 1158-1161, 2005.
- Belin D, Jonkman S, Dickinson A, Robbins TW, and Everitt BJ.** Parallel and interactive learning processes within the basal ganglia: Relevance for the understanding of addiction. *Behav Brain Res* 199: 89-102, 2009.
- Berke JD, Breck JT, and Eichenbaum H.** Striatal versus hippocampal representations during win-stay maze performance. *J Neurophysiol* 101: 1575-1587, 2009.

Berretta S, Parthasarathy HB, and Graybiel AM. Local release of GABAergic inhibition in the motor cortex induces immediate-early gene expression in indirect pathway neurons of the striatum. *J Neurosci* 17: 4752-4763, 1997.

Blazquez P, Fujii N, Kojima J, and Graybiel AM. A network representation of response probability in the striatum. *Neuron* 33: 973-982, 2002.

Brown LL. Somatotopic organization in rat striatum: evidence for a combinatorial map. *Proc Natl Acad Sci U S A* 89: 7403-7407, 1992.

Cools R, Ivry RB, and D'Esposito M. The human striatum is necessary for responding to changes in stimulus relevance. *J Cogn Neurosci* 18: 1973-1983, 2006.

Corbit LH, and Janak PH. Inactivation of the lateral but not medial dorsal striatum eliminates the excitatory impact of Pavlovian stimuli on instrumental responding. *J Neurosci* 27: 13977-13981, 2007.

Costa RM, Cohen D, and Nicolelis MA. Differential corticostriatal plasticity during fast and slow motor skill learning in mice. *Curr Biol* 14: 1124-1134, 2004.

Dahlin E, Neely AS, Larsson A, Backman L, and Nyberg L. Transfer of learning after updating training mediated by the striatum. *Science* 320: 1510-1512, 2008.

Daw ND, and Doya K. The computational neurobiology of learning and reward. *Curr Opin Neurobiol* 16: 199-204, 2006.

Daw ND, Niv Y, and Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8: 1704-1711, 2005.

Dayan P, and Balleine BW. Reward, motivation, and reinforcement learning. *Neuron* 36: 285-298, 2002.

Doyon J, and Benali H. Reorganization and plasticity in the adult brain during learning of motor skills. *Curr Opin Neurobiol* 15: 161-167, 2005.

Doyon J, Penhune V, and Ungerleider LG. Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning. *Neuropsychologia* 41: 252-262, 2003.

Emondi AA, Rebrik SP, Kurgansky AV, and Miller KD. Tracking neurons recorded from tetrodes across time. *J Neurosci Methods* 135: 95-105, 2004.

Eschenko O, and Mizumori SJ. Memory influences on hippocampal and striatal neural codes: effects of a shift between task rules. *Neurobiol Learn Mem* 87: 495-509, 2007.

Everitt BJ, and Robbins TW. Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nat Neurosci* 8: 1481-1489, 2005.

Floyer-Lea A, and Matthews PM. Changing brain networks for visuomotor control with increased movement automaticity. *J Neurophysiol* 92: 2405-2412, 2004.

Graybiel AM. The basal ganglia: learning new tricks and loving it. *Curr Opin Neurobiol* 15: 638-644, 2005.

Graybiel AM. Habits, rituals and the evaluative brain. *Annu Rev Neurosci* 31: 359-387, 2008.

Haruno M, and Kawato M. Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. *J Neurophysiol* 95: 948-959, 2006.

Hazeltine E, Grafton ST, and Ivry R. Attention and stimulus characteristics determine the locus of motor-sequence encoding. A PET study. *Brain* 120 (Pt 1): 123-140, 1997.

Jog M, Kubota Y, Connolly CI, Hillegaart V, and Graybiel AM. Building neural representations of habits. *Science* 286: 1745-1749, 1999.

Jog MS, Connolly CI, Kubota Y, Iyengar DR, Garrido L, Harlan R, and Graybiel AM.

Tetrode technology: advances in implantable hardware, neuroimaging, and data analysis techniques. *J Neurosci Methods* 117: 141-152, 2002.

Jueptner M, and Weiller C. A review of differences between basal ganglia and cerebellar control of movements as revealed by functional imaging studies. *Brain* 121: 1437-1449, 1998.

Kawaguchi Y, Wilson CJ, Augood SJ, and Emson PC. Striatal interneurons: chemical, physiological and morphological characterization. *Trends Neurosci* 18: 527-535, 1995.

Kimchi EY, and Laubach M. Dynamic encoding of action selection by the medial striatum. *J Neurosci* 29: 3148-3159, 2009.

Kimchi EY, Torregrossa MM, Taylor JR, and Laubach M. Neuronal correlates of instrumental learning in the dorsal striatum. *J Neurophysiol* 102: 475-489, 2009.

Koós T, and Tepper JM. Dual cholinergic control of fast-spiking interneurons in the neostriatum. *J Neurosci* 22: 529-535, 2002.

Lau B, and Glimcher PW. Action and outcome encoding in the primate caudate nucleus. *J Neurosci* 27: 14502-14514, 2007.

Matell MS, Meck WH, and Nicolelis MA. Interval timing and the encoding of signal duration by ensembles of cortical and striatal neurons. *Behav Neurosci* 117: 760-773, 2003.

McGeorge AJ, and Faull RLM. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* 29: 503-537, 1989.

McNab F, and Klingberg T. Prefrontal cortex and basal ganglia control access to working memory. *Nat Neurosci* 11: 103-107, 2008.

Monchi O, Petrides M, Strafella AP, Worsley KJ, and Doyon J. Functional role of the basal ganglia in the planning and execution of actions. *Ann Neurol* 59: 257-264, 2006.

Morris G, Arkadir D, Nevet A, Vaadia E, and Bergman H. Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron* 43: 133-143, 2004.

Nelder JA, and Wedderburn RW. Generalized linear models. *J R Stat Soc Ser A* 135: 370-384, 1972.

Niv Y, Joel D, and Dayan P. A normative perspective on motivation. *Trends Cogn Sci* 10: 375-381, 2006.

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, and Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304: 452-454, 2004.

Packard MG. Exhumed from thought: Basal ganglia and response learning in the plus-maze. *Behav Brain Res* 199: 24-31, 2009.

Parthasarathy HB, and Graybiel AM. Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. *J Neurosci* 17: 2477-2491, 1997.

Pasupathy A, and Miller EK. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433: 873-876, 2005.

Poldrack RA, Sabb FW, Foerde K, Tom SM, Asarnow RF, Bookheimer SY, and Knowlton BJ. The neural correlates of motor skill automaticity. *J Neurosci* 25: 5356-5364, 2005.

Ragozzino ME. The contribution of the medial prefrontal cortex, orbitofrontal cortex, and dorsomedial striatum to behavioral flexibility. *Ann N Y Acad Sci* 1121: 355-375, 2007.

Ragozzino ME, Ragozzino KE, Mizumori SJ, and Kesner RP. Role of the dorsomedial striatum in behavioral flexibility for response and visual cue discrimination learning. *Behav Neurosci* 116: 105-115, 2002.

Ravel S, Legallet E, and Apicella P. Responses of tonically active neurons in the monkey striatum discriminate between motivationally opposing stimuli. *J Neurosci* 23: 8489-8497, 2003.

Redish AD, Jensen S, and Johnson A. A unified framework for addiction: vulnerabilities in the decision process. *Behav Brain Sci* 31: 415-437; discussion 437-487, 2008.

Samejima K, Ueda Y, Doya K, and Kimura M. Representation of action-specific reward values in the striatum. *Science* 310: 1337-1340, 2005.

Schmitzer-Torbert NC, and Redish AD. Task-dependent encoding of space and events by striatal neurons is dependent on neural subtype. *Neuroscience* 153: 349-360, 2008.

Schultz W, Tremblay L, and Hollerman JR. Changes in behavior-related neuronal activity in the striatum during learning. *Trends Neurosci* 26: 321-328, 2003.

Spiegelhalter DJ, Best NG, Carlin BR, and van der Linde A. Bayesian measures of model complexity and fit. *J Roy Stat Soc B* 64: 583-616, 2002.

Sutton RS, and Barto AG. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.

Tang C, Pawlak AP, Prokopenko V, and West MO. Changes in activity of the striatum during formation of a motor habit. *Eur J Neurosci* 25: 1212-1227, 2007.

Tepper JM, and Bolam JP. Functional diversity and specificity of neostriatal interneurons. *Curr Opin Neurobiol* 14: 685-692, 2004.

Thorn CA, Atallah H, Howe M, and Graybiel A. Ensemble unit activity differs between lateral and medial striatal regions during T-maze task learning Program No. 578.7. 2008 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2008. Online.

Van Golf Racht-Delatour B, and Massioui NE. Alleviation of overtraining reversal effect by transient inactivation of the dorsal striatum. *Eur J Neurosci* 12: 3343-3350, 2000.

Voorn P, Vanderschuren LJ, Groenewegen HJ, Robbins TW, and Pennartz CM. Putting a spin on the dorsal-ventral divide of the striatum. *Trends Neurosci* 27: 468-474, 2004.

Watanabe K, and Hikosaka O. Immediate changes in anticipatory activity of caudate neurons associated with reversal of position-reward contingency. *J Neurophysiol* 94: 1879-1887, 2005.

White NM. Some highlights of research on the effects of caudate nucleus lesions over the past 200 years. *Behav Brain Res* 199: 3-23, 2009.

White NM, and McDonald RJ. Multiple parallel memory systems in the brain of the rat. *Neurobiol Learn Mem* 77: 125-184, 2002.

Williams ZM, and Eskandar EN. Selective enhancement of associative learning by microstimulation of the anterior caudate. *Nat Neurosci* 9: 562-568, 2006.

Willingham DB, Salidis J, and Gabrieli JD. Direct comparison of neural systems mediating conscious and unconscious skill learning. *J Neurophysiol* 88: 1451-1460, 2002.

Wilson CJ, and Groves PM. Spontaneous firing patterns of identified spiny neurons in the rat neostriatum. *Brain Res* 220: 67-80, 1981.

Yin HH, and Knowlton BJ. The role of the basal ganglia in habit formation. *Nat Rev Neurosci* 7: 464-476, 2006.

Yin HH, Knowlton BJ, and Balleine BW. Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *Eur J Neurosci* 19: 181-189, 2004.

Yin HH, Mulcare SP, Hilario MR, Clouse E, Holloway T, Davis MI, Hansson AC, Lovinger DM, and Costa RM. Dynamic reorganization of striatal circuits during the acquisition and consolidation of a skill. *Nat Neurosci* 12: 333-341, 2009.

FIGURE LEGENDS

FIG. 1. Experimental methods and behavioral performance. *A*: Behavioral task. *B-D*: Sorting of recorded units. Clusters forming on spike dot plots defined putative single units (*B*). These clusters were then verified with shapes of spike waveforms (*C*) and a pause in spiking detected in autocorrelograms (*D*) and in interspike interval plots. *E*: Task-related activity of MS units with responses at locomotion onset (*top*) and with responses at goal reaching (*bottom*). Each peri-event time histogram (PETH) shows session average firing rates during ± 400 -ms intervals around task events. Red and blue horizontal lines indicate, respectively, mean baseline firing rates and levels 2 SD above and below the baseline mean. Raster plots show occurrence of spikes (black) and surrounding events (red) during trials of the training session. *F*: Methods for matching average spike waveforms across sessions to track single units over multiple sessions. The pairs of units across consecutive daily sessions that had correlation values of greater than 0.98 were judged putatively to be the same unit. See METHODS. *G*: Recording sites in the dorsolateral caudoputamen in the 12 mice with tetrode recording.

FIG. 2. Behavioral performance. *A*: Average percentage of correct responses during the first 12 training sessions of mice trained only on the tactile version of the T-maze task (green) and mice successively trained on the auditory version (red) and then on the tactile version (blue). Error bars indicate standard errors of the mean. *B*: The number of sessions required to reach the acquisition criterion. T: Tactile training only. A: Initial auditory training. T_A: Tactile training following auditory training. *C*: Average percentages of correct responses (*left*) and average maze-running times (*right*) during successive training on the auditory (red) and tactile (blue) task versions. Learning stages are defined in METHODS. *D*: Average running times for inter-event

intervals, from left to right, locomotion onset to cue onset, cue onset to turn onset (arrow marking sharp increase at cue switch), turn onset to turn offset, and turn offset to goal reaching. *E*: Average trial-by-trial running times (top) and 4 successive inter-event durations (rows 2-5) during the last 3 auditory and the first 3 tactile sessions.

FIG. 3. Activity of medium-spiny neurons of the striatum during successive training on the auditory and tactile task versions. *A* and *B*: Average population activity of task-related (*A*) and “non-task-responsive” (*B*) units averaged for each learning stage (auditory: A1-13; tactile: T1-7), plotted using the pseudocolor scale shown at bottom right. For each unit, spike counts in each 20-ms bin were normalized according to $FR_N = (X - M_B) / SD_{B+T}$ (FR_N : normalized firing rate; X : bin spike count; M_B : baseline mean firing rate; SD_{B+T} : baseline and trial time SD). FR_N values were then averaged across units recorded in each stage. Plots were made by abutting successive ± 200 ms peri-event windows (as labeled, separated by white vertical lines) according to their order of occurrence within a trial. Inset at bottom left depicts alignment of training sessions. Numbers of units and numbers of mice (in parentheses) available for each learning stage are listed at right. *C* and *D*: Average normalized activity of task-related (*C*) and non-task-responsive (*D*) units calculated for 500-ms pre-trial baseline period (BL) and for 200-ms intervals before and after the warning cue, gate opening (GO), locomotion onset (L), out of the start area, cue onset (C), turn (T) beginning and end and goal reaching (GR). Data are averaged across learning stages for 6 phases of training. ANOVA performed on these data indicated that activity of task-related units in the windows around locomotion onset and from beginning of turn to goal reaching was higher than that during baseline throughout training ($P < 0.05$). Activity before and after cue onset declined during acquisition, and by early overtraining (stages A6-9), activity was

significantly lower in peri-cue periods than in peri-turn offset and pre-goal periods. These patterns continued to be significant during acquisition and overtraining on the tactile task (stages T1-2 and T3-7; $P < 0.05$). Activity of non-task-responsive units during trial runs significantly declined from stages A1-2 to A6-9 ($P < 0.05$). Error bars represent standard errors of the mean.

FIG. 4. Ensemble activity of task-related MS neurons and running speed around cue switch. *A*: Average normalized ensemble activity, calculated as for Fig. 3, of task-related MS units during 10 sessions around cue switch (arrow). Session alignment is shown as in Fig. 3A. *B*: Trial-by-trial running speed in 200-ms pre- and post-event intervals (as labeled) during 6 sessions around cue switch.

FIG. 5. Stable activity of MS neurons across cue switch. *A*: Normalized activity of task-related MS units averaged over the last 5 auditory sessions (*left*) and the first 5 tactile sessions (*right*), for the pre-trial baseline period and for 200-ms intervals before and after each task event. *B*: Pseudo-color plots of normalized ensemble activity during 6 sessions around cue switch (arrow). For each session, average activity during consecutive 5-trial blocks is plotted in each row. Session alignment is shown as in Fig. 3A. *C*: Average ensemble activity of units with responses to individual task events, including those showing responses to more than one event. Gray entries represent days with no accepted units.

FIG. 6. Ensemble activity of fast-firing interneurons of the striatum. *A*: Average normalized activity of task-related FF units during successive training on the auditory and tactile task versions, calculated as for Fig. 3. Gray entries indicate stages with fewer than 6 recorded units,

excluded from analysis. *B*: Average normalized activity of task-related FF units over 200-ms pre- and post-event windows during, from left to right, stages A1-2, A3-5, A6-9, T1-2 and T3-7. These plots show that firing rates around locomotion onset and goal reaching were higher only than the rates before trial runs during auditory task acquisition ($P < 0.05$). Their activity during mid-run then decreased and became significantly lower than that in peri-locomotion onset (A6-9), post-turn onset (A6-9) and post-goal (A3-5) periods. Activity during periods from after turn to goal reaching did not differ significantly from pre-trial activity in stages A6-9. The beginning-and-end ensemble activity pattern was weaker during tactile acquisition (T1-2, $P = 0.02$) but returned to the pre-switch level in overtraining (T3-7, $P < 0.001$). *C* and *D*: Activity of task-related FF units during 10 sessions around cue switch (arrow), with all trials (*C*) and after discarding trials in which cue-turn durations were greater than mean + 2SD of last auditory session (*D*). Symbols as in Fig. 4. *E*: Firing rates in 200-ms pre- and post-event intervals averaged over the last 5 auditory sessions (*left*) and over the first 5 tactile sessions (*right*). The post-cue activity was greater in first 5 tactile sessions than in last 5 auditory sessions ($P = 0.027$). *F*: Scatter plot illustrating a lack of correlation between post-cue running speed and average activity calculated for 5-trial blocks in the first 3 tactile sessions. *G*: Peri-cue response histograms (*left*) and scatter plots of trial-by-trial running speed and firing rates during 200-ms post-cue window (*right*) for 4 FF task-related units that exhibited cue-related discharges on the first tactile session. Correlations coefficient (r) and p-value (P) are shown for each unit.

FIG. 7. Task-related activity profiles of MS and FF units. *A*: Proportions of MS (*top*) and FF (*bottom*) units that fired differentially for trials with right and left turns during successive 200-ms pre/post-event periods across 6 phases of training. Note increased values for FF units. *B*:

Proportions of MS (*top*) and FF (*bottom*) units with differential responses to conditional cues during post-cue period across learning stages. *C*: Proportions of MS (*top*) and FF (*bottom*) units with differential responses in correct and incorrect trials averaged over auditory (red) and tactile (blue) sessions. *D* and *E*: Proportions of task-related MS (*top*) and FF (*bottom*) units with significant correlations ($P < 0.01$) between firing rates during 200-msec pre- and post-event windows and trial running times (*D*) and between firing rates and running speed within each peri-event window (*E*). *F*: Percentages of MS units with responses to any task event relative to all accepted projection neurons. *G*: Changes in proportion of spikes that occurred within detected peaks of phasic neuronal responses during auditory and tactile training (*left*) and during 10 sessions around switch (*right*).

FIG. 8. Stability of putative single MS neuron ensembles across the cue switch. *A*: Average normalized activity of 24 MS units, including those with and without task-related responses, tracked over a minimum of 6 sessions across the switch in task version, plotted as described in Figs. 3A and 4A. *B*: Line plots illustrating average activity of 14 task-related MS units in the last auditory session (red) and in the first tactile session (blue). These units were identified as PSNs over these two sessions. Note that the two lines virtually overlap.

FIG. 9. Differential responses to auditory and tactile cues of putative single MS neuron tracked across the switch in task version. *A*: Session average PETHs, as described in Fig. 1E, of spike activity, aligned at gate opening, 8 kHz tone/smooth cue, and right turn-end during sequential training on the auditory (days 17-29) and tactile (days 30-42) task versions, followed by alternation between the task versions (days 43-49). Pre- and post-cue activity was significantly

different between auditory and tactile task versions ($P < 0.0001$), but activity at right turn was not ($P = 0.13$). *B*: Correlation matrix of peri-event histograms (± 400 -ms windows, 20-ms bins) among 10 daily sessions around the cue switch, calculated separately for each task event, as labeled. Scale is shown at right.

FIG. 10. Activity of putative single MS neuron with rebound of goal-related activity at cue switch. *A*: PETHs, constructed as in Fig. 1*E*, aligned at gate opening, left goal reaching, and right goal reaching during last 3 auditory sessions and first 7 tactile sessions. *B*: Correlation matrix of event-related discharges during these 10 sessions, as described in Fig. 9*B*. *C*: Poisson generalized linear models without a change-point (black), with a change-point at cue switch (dashed green), and with a change-point at undefined time (orange), fitted to average firing rates during 200-ms intervals after right goal reaching (solid circles). Note that green and orange lines almost completely overlap.

FIG. 11. Putative single FF neuron with changes in pre-cue activity at cue switch. *A*: PETHs aligned at gate opening, 1 kHz tone/smooth cue onset and left turn end during 10 consecutive sessions around cue switch. *B*: Correlation matrix of task-related responses, as described in Fig. 9*B*.

FIG. 12. Putative single neuron with stable activity across cue switch. *A*: PETHs aligned at warning cue, locomotion onset, and right goal reaching. *B*: Correlation matrix of event-related responses during 10 sessions around cue switch, as described in Fig. 9*B*.

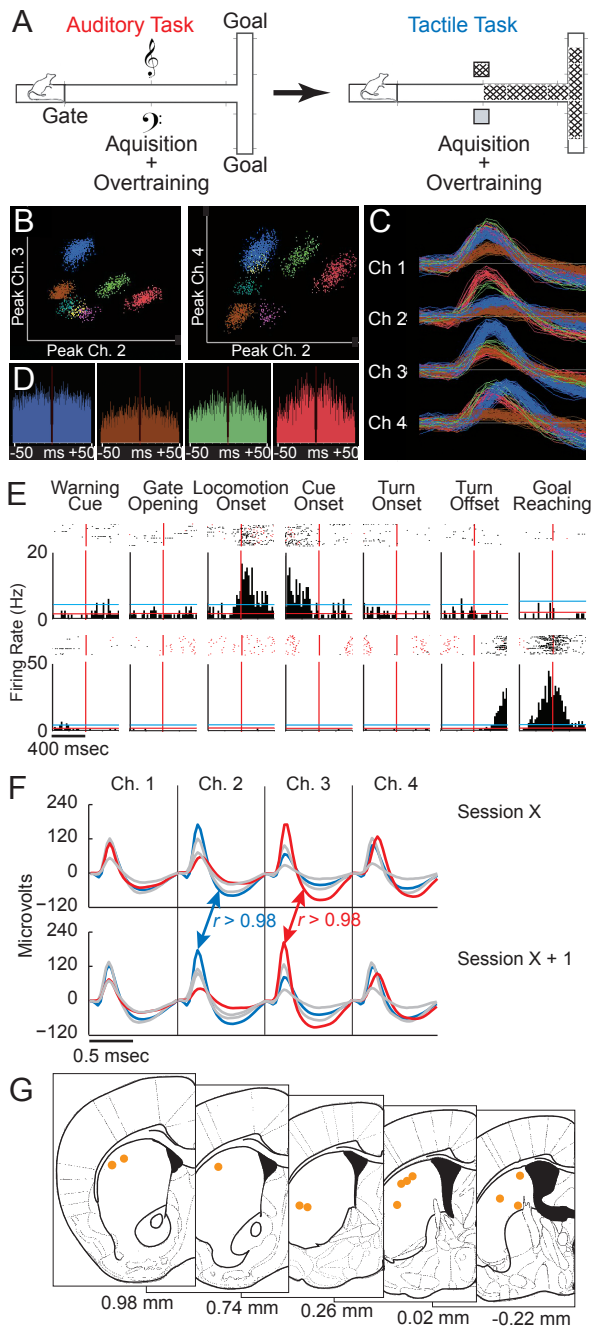


Figure 1 (Kubota et al.)

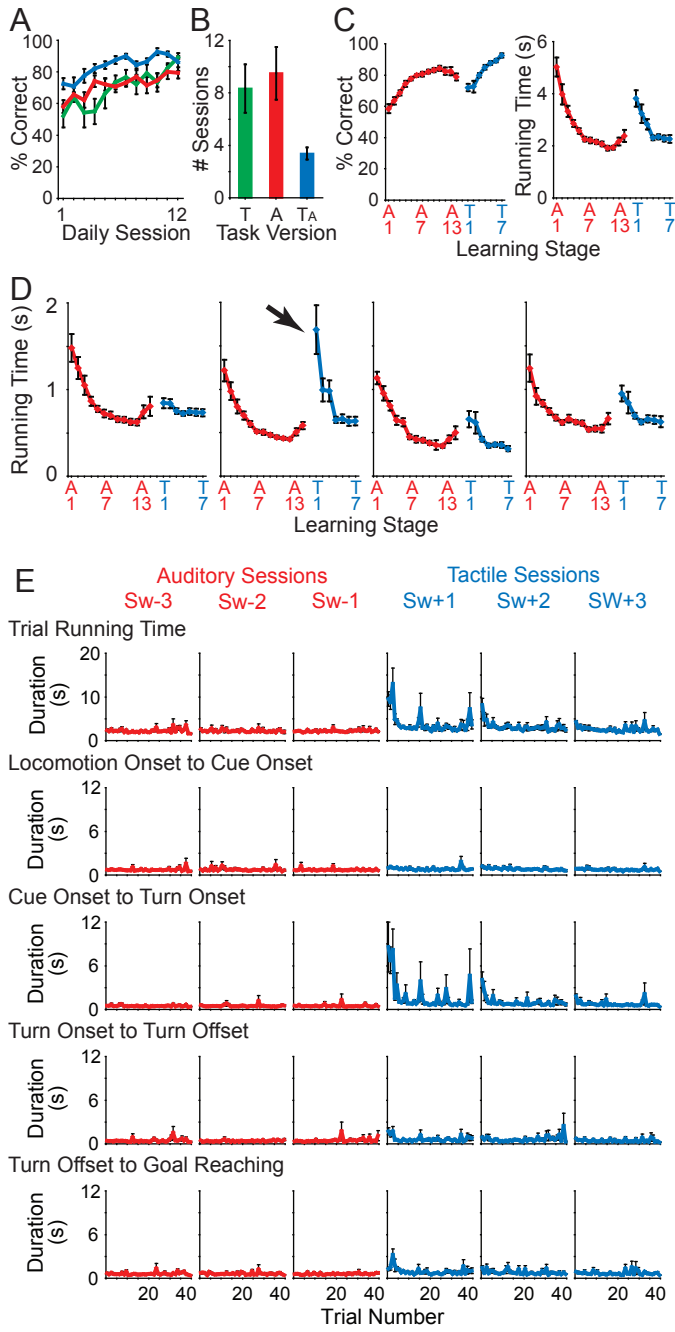


Figure 2 (Kubota et al.)

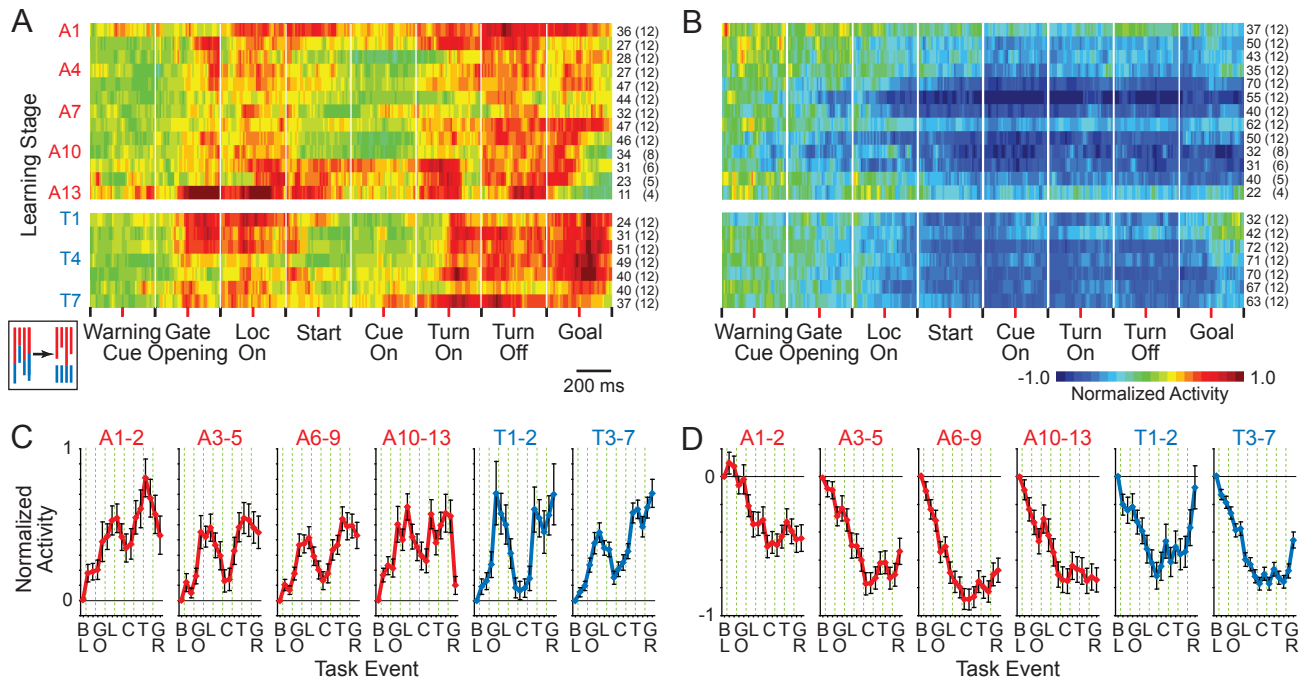


Figure 3 (Kubota et al.)

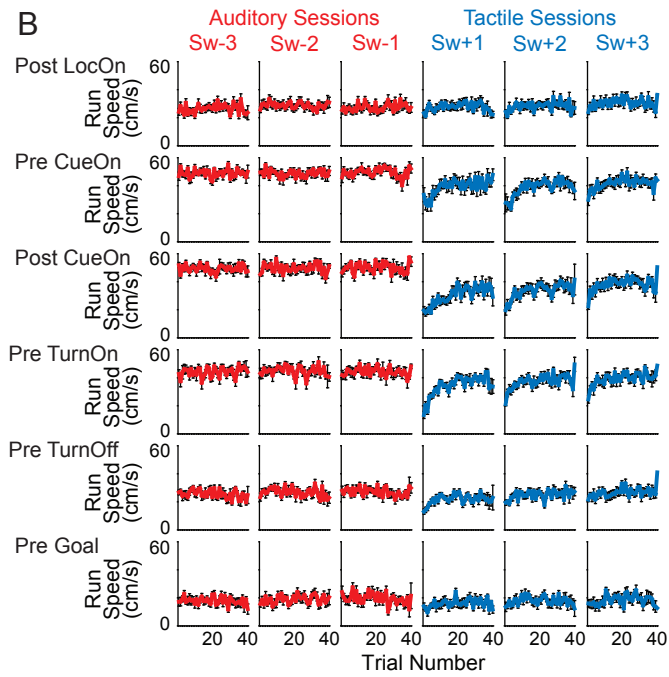
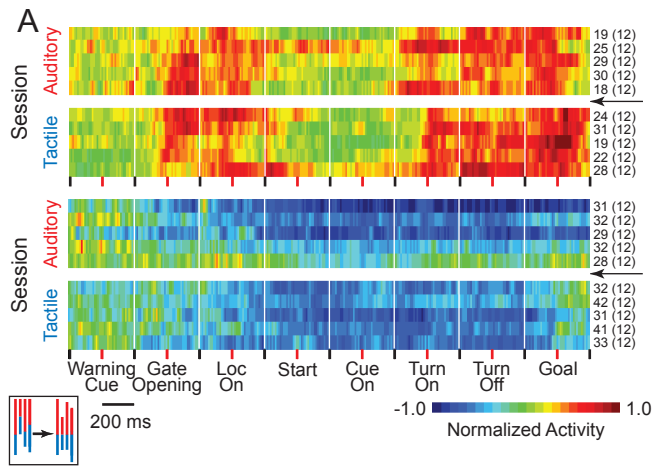


Figure 4 (Kubota et al.)

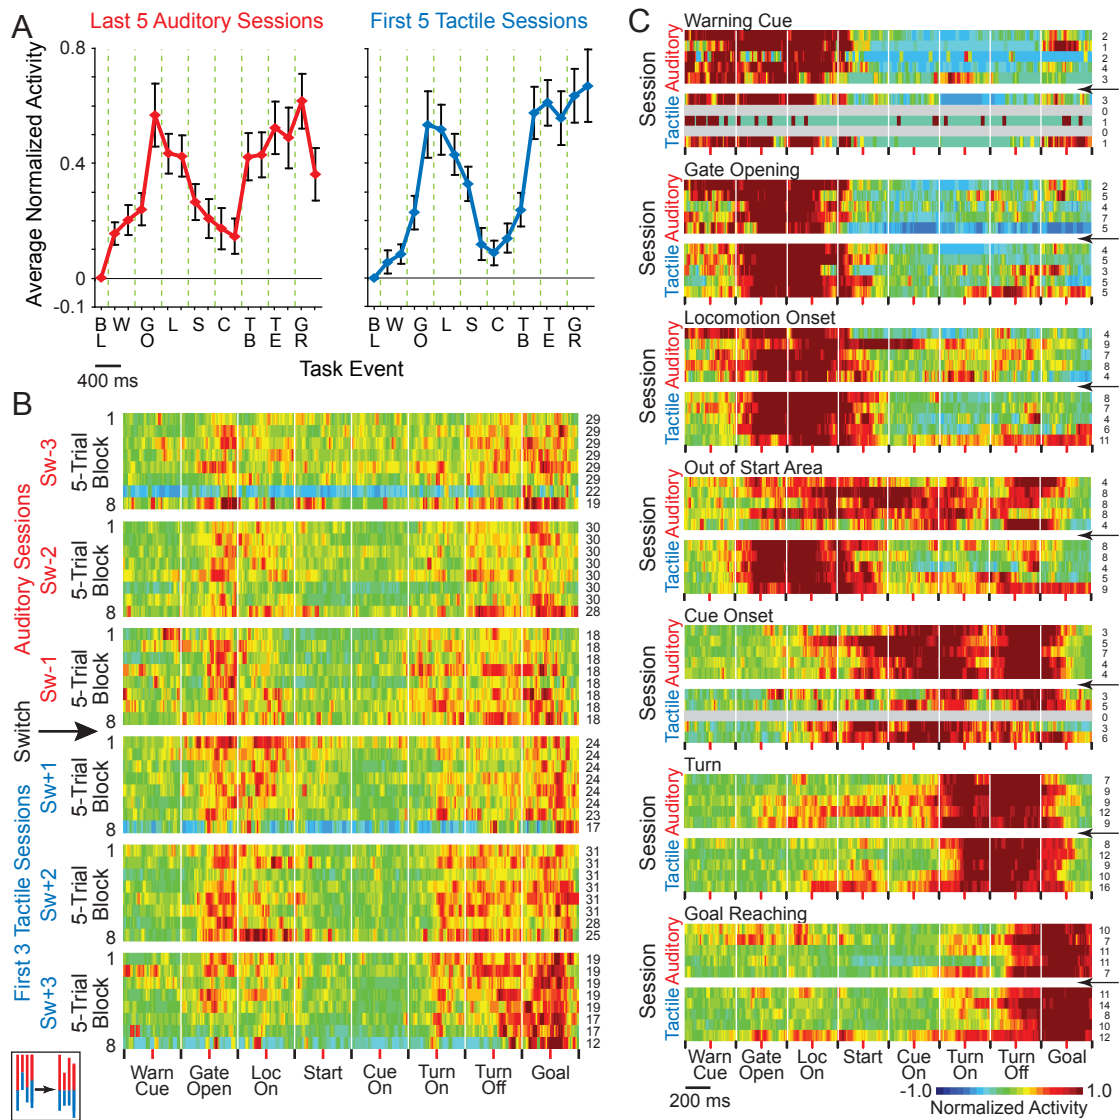


Figure 5 (Kubota et al.)

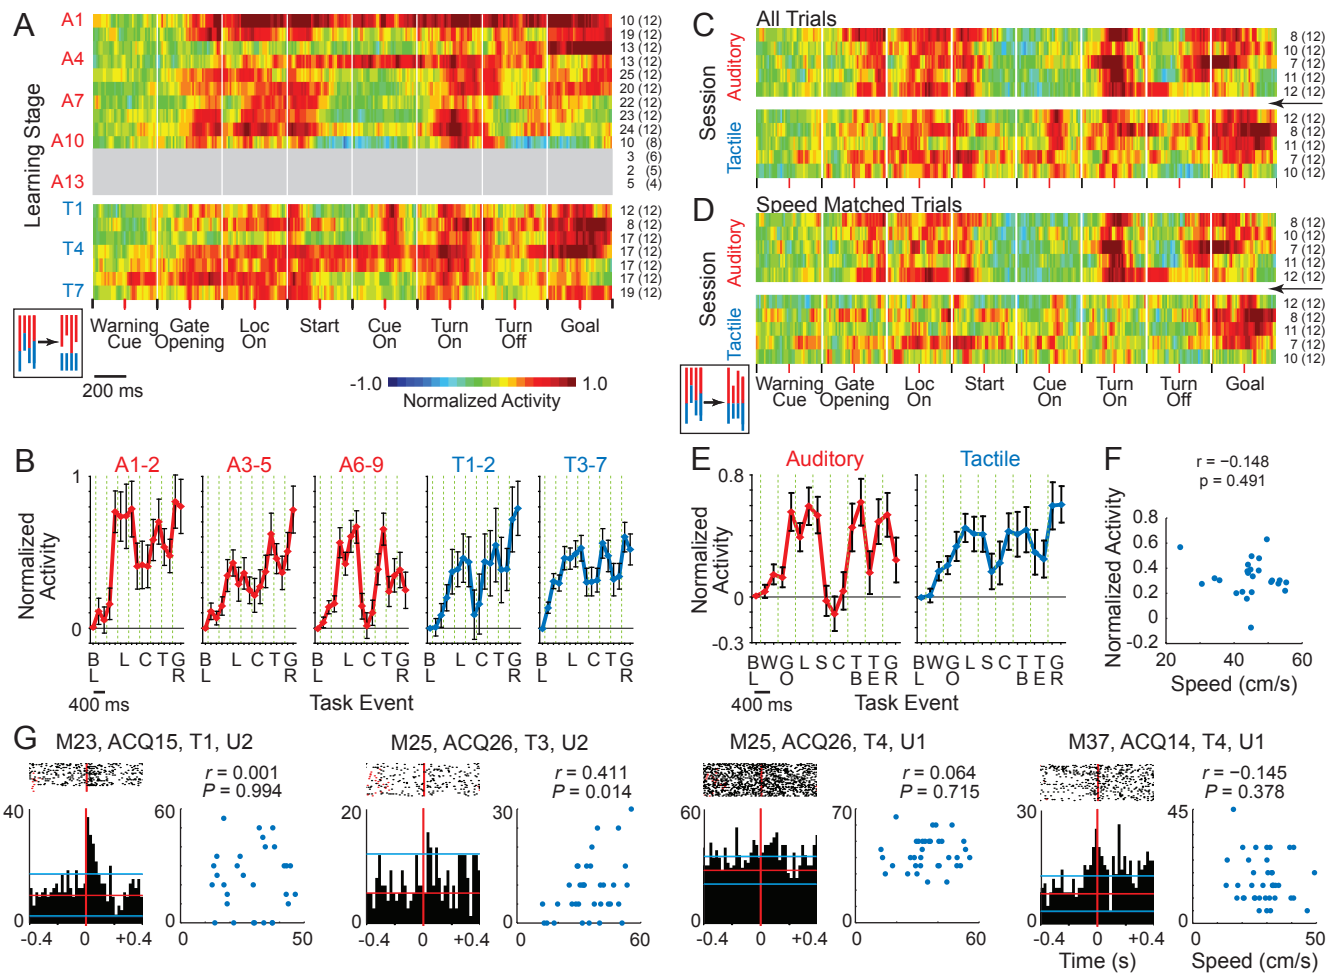


Figure 6 (Kubota et al.)

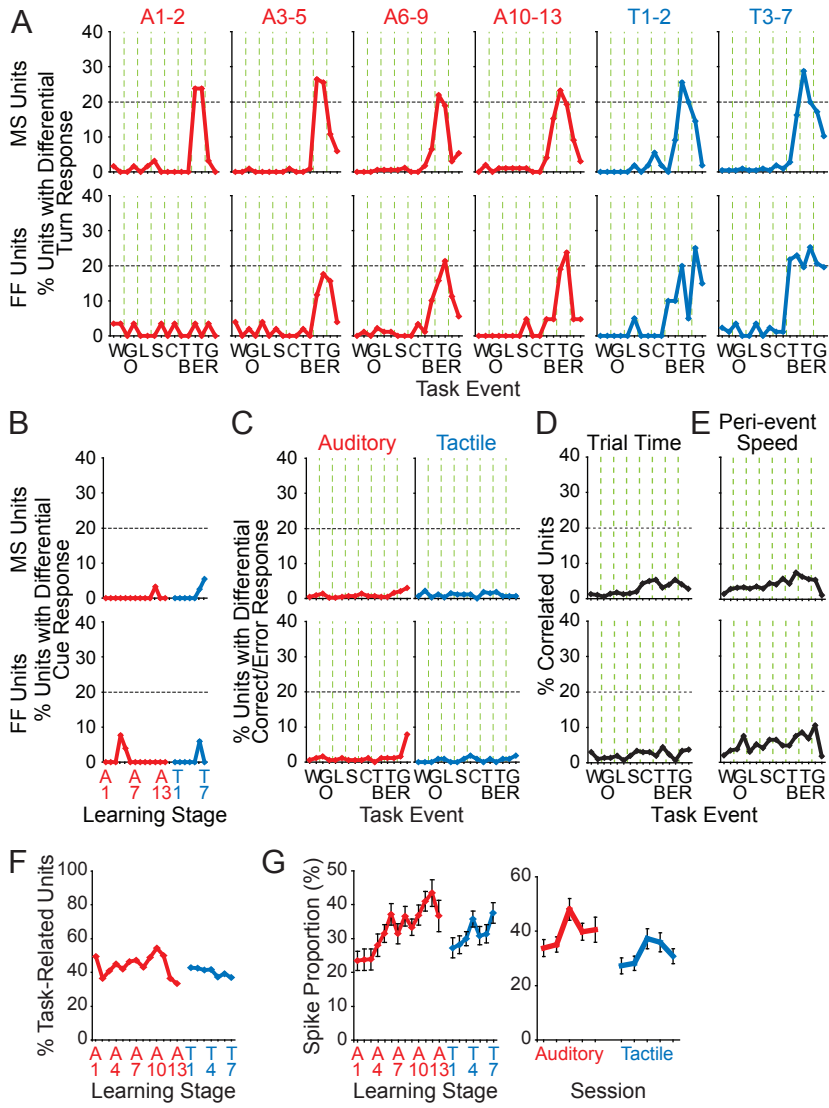


Figure 7 (Kubota et al.)

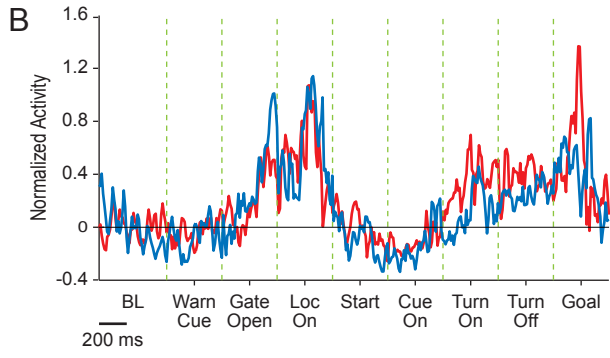



Figure 8 (Kubota et al.)

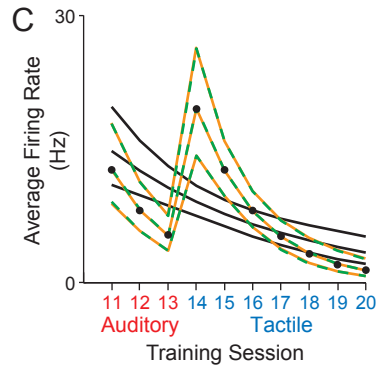
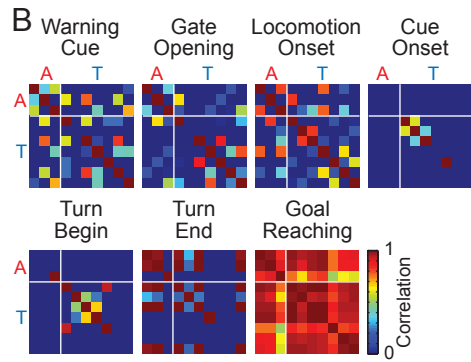
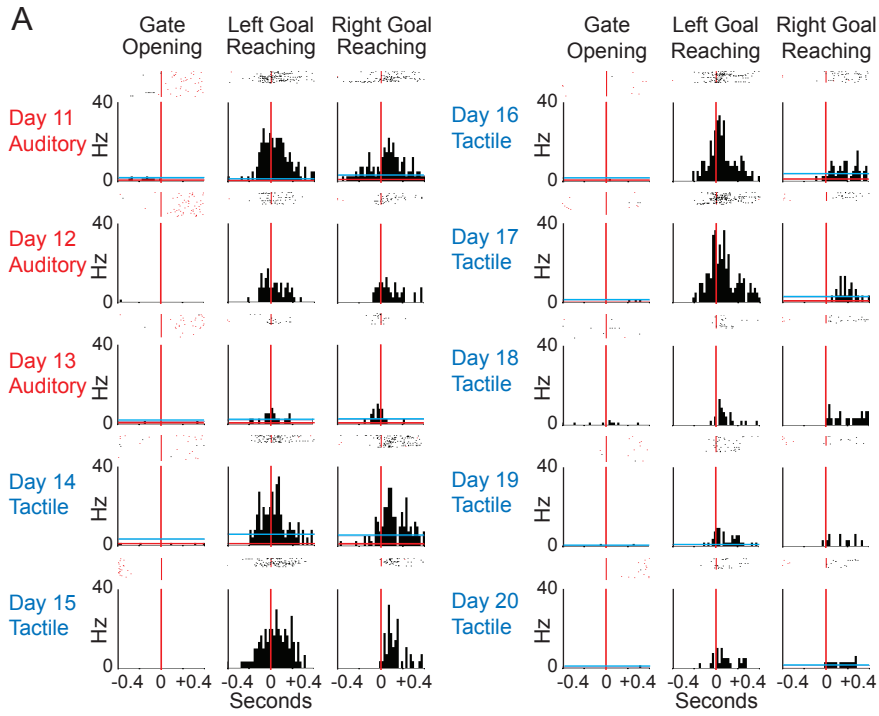


Figure 10 (Kubota et al.)


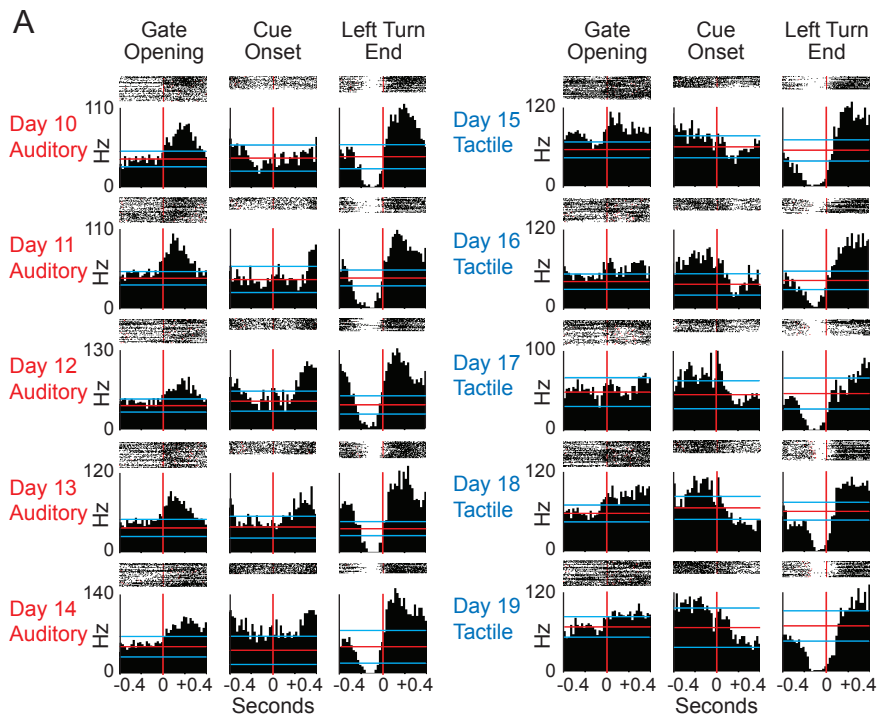


Figure 11 (Kubota et al.)


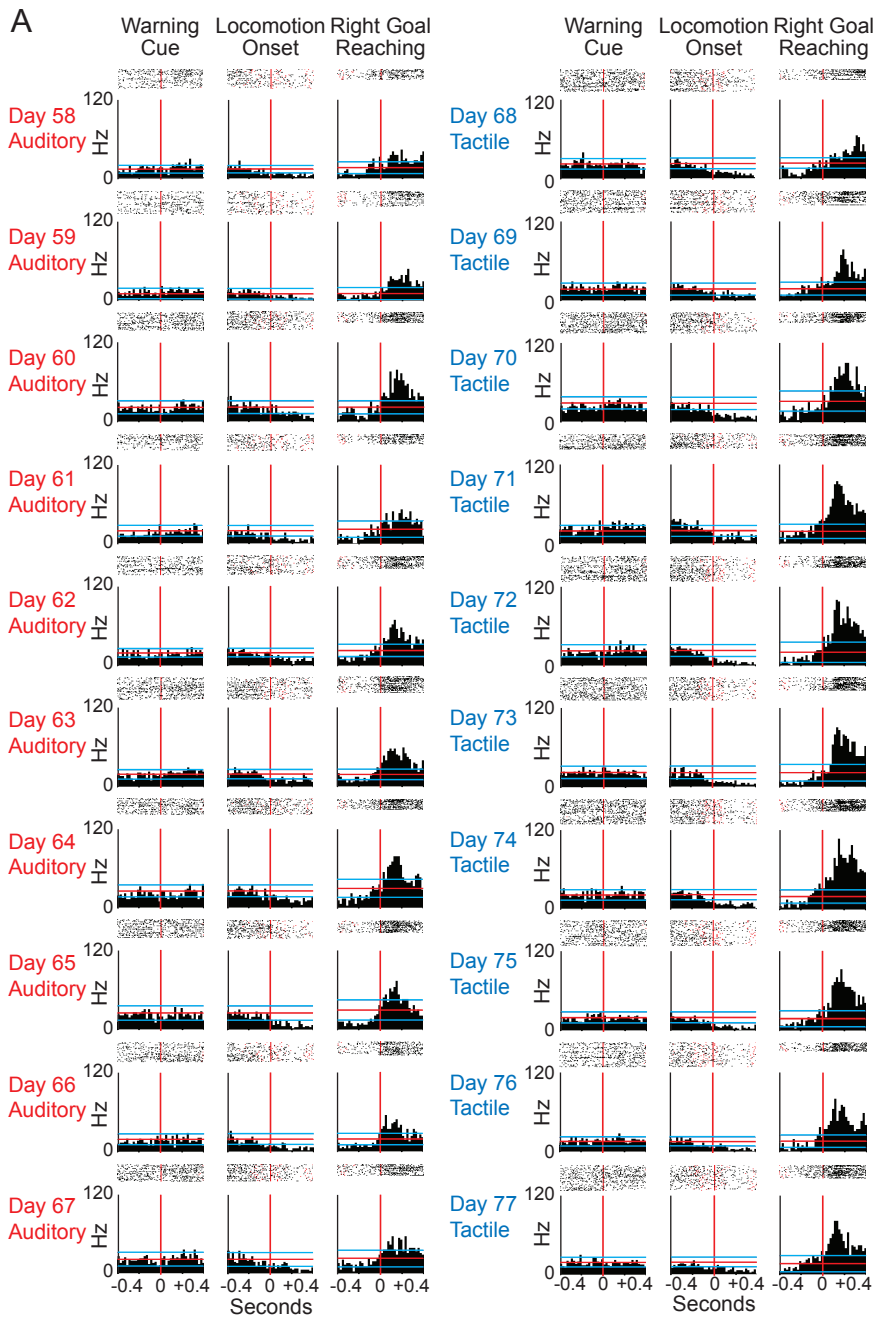


Figure 12 (Kubota et al.)