J Neurophysiol 91: 2259–2272, 2004.
First published January 21, 2004; 10.1152/jn.00687.2003.

Neuronal Activity in the Rodent Dorsal Striatum in Sequential Navigation: Separation of Spatial and Reward Responses on the Multiple T Task

Neil Schmitzer-Torbert and A. David Redish
Department of Neuroscience, University of Minnesota, Minneapolis, Minnesota 55455

Submitted 17 July 2003; accepted in final form 15 January 2004

Schmitzer-Torbert, Neil and A. David Redish. Neuronal activity in the rodent dorsal striatum in sequential navigation: separation of spatial and reward responses on the multiple T task. *J Neurophysiol* 91: 2259–2272, 2004. First published January 21, 2004; 10.1152/jn.00687.2003. The striatum plays an important role in "habitual" learning and memory and has been hypothesized to implement a reinforcement-learning algorithm to select actions to perform given the current sensory input. Many experimental approaches to striatal activity have made use of temporally structured tasks, which imply that the striatal representation is temporal. To test this assumption, we recorded neurons in the dorsal striatum of rats running a sequential navigation task: the multiple T maze. Rats navigated a sequence of four T maze turns to receive food rewards delivered in two locations. The responses of neurons that fired phasically were examined. Task-responsive phasic neurons were active as rats ran on the maze (maze-responsive) or during reward receipt (reward-responsive). Neither maze- nor reward-responsive neurons encoded simple motor commands: maze-responses were not well correlated with the shape of the rat's path and most reward-responsive neurons did not fire at similar rates at both food-delivery sites. Maze-responsive neurons were active at one or more locations on the maze, but these responses did not cluster at spatial landmarks such as turns. Across sessions the activity of maze-responsive neurons was highly correlated when rats ran the same maze. Maze-responses encoded the location of the rat on the maze and imply a spatial representation in the striatum in a task with prominent spatial demands. Maze-responsive and reward-responsive neurons were two separate populations, suggesting a divergence in striatal information processing of navigation and reward.

INTRODUCTION

Much recent work on the striatum, the major input structure of the basal ganglia, has focused on the involvement of the dorsal striatum in "habitual," stimulus-response (S-R) learning and memory (Cohen and Squire 1980; Graybiel 1998; Knowlton et al. 1996; Mishkin and Appenzeller 1987; Packard and Knowlton 2002). Striatal damage or inactivation impairs performance in S-R tasks while not affecting performance on tasks requiring "explicit" or cognitive learning and memory (Doyon et al. 1998; Kesner et al. 1993; McDonald and White 1993; Packard 1999; Packard and McGaugh 1992; Packard et al. 1989). The dorsal striatum in primates is also critical for the learning and performance of visuomotor sequences (Matsumoto et al. 1999; Miyachi et al. 1997). In the rodent, the dorsal striatum is critical for the performance of natural sequences of grooming movements (Berridge and Whishaw 1992; Cromwell and Berridge 1996). Neurons in the dorsal striatum respond to a variety of parameters, including grooming movements, auditory cues, head direction, head movements, and turning (Aldridge and Berridge 1998; Carelli et al. 1997; Gardiner and Kitai 1992; Jog et al. 1999; Kermadi and Joseph 1995; Kermadi et al. 1993; Ragozzino et al. 2001; Tremblay et al. 1998; Wiener 1993). Many studies have indicated the specificity of these responses, showing that neurons that are responsive during a task may not be responsive outside of the task and vice versa (Aldridge and Berridge 1998; Carelli et al. 1997; Gardiner and Kitai 1992) and that in sequential tasks, neurons can have preferences for specific sequence locations or specific sequences (Aldridge and Berridge 1998; Kermadi and Joseph 1995; Kermadi et al. 1993; Kimura 1986, 1990). The striatum is thus involved in the learning and performance of simple S-R tasks as well as longer chains of S-R relationships in sequential tasks.

While striatal neurons have been examined in primates during the performance of long chains of sequential movements, studies in the rodent have tended to use tasks in which rats make a single response (such as turning on a T maze, lever pressing, or head movements). Little is known about how neurons in the rodent dorsal striatum respond as rats perform more complex tasks, and an understanding of such responses may shed light on how striatal activity could be used to perform long chains of automated behavior. Many tasks that have been used in striatal recordings have strong temporal components: subjects pay attention to instruction cues and make reactions after fixed or variable delays to receive reward at some later time point. These tasks have shown that striatal neurons have responses that cover the temporal interval between the start of a trial and the receipt of reward (summarized by Schultz et al. 1995). However, less is known of what type of striatal representation will be obtained in a task that emphasizes spatial information.

In this experiment, we examined neurons in the dorsal striatum of rats during the performance of arbitrary navigational sequences. Rats ran through a complex route formed by four T choices arranged sequentially and were rewarded for correctly navigating the maze with food delivered at two different locations. The use of multiple T choices and multiple food-delivery sites allowed us to examine neural responses when rats made the same action (turning left, turning right, food consumption) observed in different parts of the task sequence. We have previously reported that rats running on a multiple T maze quickly eliminated errors from their path, demonstrating that rats could perform well in a single session

Address for reprint requests and other correspondence: A. D. Redish, Dept. of Neuroscience, 6–145 Jackson Hall, 321 Church St. SE, Minnneapolis MN 55455 (E-mail: redish@ahc.umn.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

even when presented with a sequence of T choices that they had never experienced (Schmitzer-Torbert and Redish 2002). Here, we present behavioral data from rats running novel and familiar multiple T mazes and data from neurons recorded in the dorsal striatum as rats ran the multiple T maze. Some of this work has been previously presented in abstract form (Redish et al. 2002a; Schmitzer-Torbert et al. 2002).

METHODS

Animals

Five male Brown-Norway-Cross rats obtained from the National Institute on Aging were used (aged 13–15 mo at the beginning of recordings). During behavioral training, rats were food-restricted, and in most cases, rats received all of their food during the experimental task. Additional food was provided following the experimental task when required to maintain a rat's weight >80% of its free-feed weight. Two of the rats had prior experience running for food on an operant conditioning task, whereas the other three rats did not. Before and during experiments, rats were handled for 15 min each day. All procedures were in accordance with National Institutes of Health guidelines for animal care and approved by the IACUC at the University of Minnesota.

Task

Rats were trained to run an elevated linear multiple T task which consisted of four T choices arranged sequentially to form a *turn sequence* (see Fig. 1). On either side of the turn sequence, *return rails* led from the end of the maze back to the beginning, so that rats ran the maze as a continuous loop. On each return rail, two automatic food dispensers (Med-Associates, St. Albans, VT) delivered pellets to locations on the track separated by ~45 cm. On completion of each trial, the rat received two 45-mg pellets (Research Diets, New Brunswick, NJ) at each food-delivery location, for a total of four pellets per trial. If the rat made an incorrect turn on the final T and ran back along the wrong return rail (thus passing the incorrect pair of pellet dispensers), no pellets were delivered, and the rat had to repeat the turn sequence to finish the trial and receive food. Throughout the task, rats were blocked from moving backward on the maze but were allowed to make incorrect choices. In practice, rats tended not to turn around and were rarely blocked.

The maze was constructed of plywood boards measuring 10 cm wide and covered with carpet. Each T consisted of a stem (30 cm long), and two choice arms (each 18.5 cm long) oriented at 90° to the stem. The return rails were 212.5 cm long, and were separated by two rails (142.5 cm long), located at either end of the return rails, which led from the turn sequence to the return rails (see Fig. 1).

A *trial* was defined as the interval between successive arrivals at the second food-delivery site on the rewarded return rail. A *lap* was defined as each time the rat passed through the turn sequence; thus multiple laps could occur during a single trial if the rat made an incorrect choice on the final T. Rats were not removed from the track between trials or laps; rats ran the task as a continuous loop. Trials and laps were defined for analysis purposes only.

Initial training was conducted using a shortened three T maze with only three T choices. When rats were able to run the task, they began training on a five T maze, which used five T choices. They were trained on the five T maze for ≥1 wk. For both three T and five T maze training, the sequence of turns used for each rat was changed daily. Rats were taken off the task ≥2 days before surgery and given ad lib access to food. Beginning 2 days after surgery, rats ran three T mazes until they were able to run the task while connected to the recording system and tetrodes had been advanced into the striatum. Rats were then moved to a novel-1/familiar/novel-2 protocol in which they ran one four-T maze per day for 3 wk (7 mazes per condition). In the first and third week (novel maze conditions), rats were presented with a different sequence of turns each day, whereas in the second week (familiar maze condition), rats were presented with the same sequence of turns each day. The mazes in the familiar condition used the last sequence of turns presented in the novel-1 condition. To control for odor cues in the familiar maze condition, specific T choices were swapped daily, but the position of the turn sequence relative to the

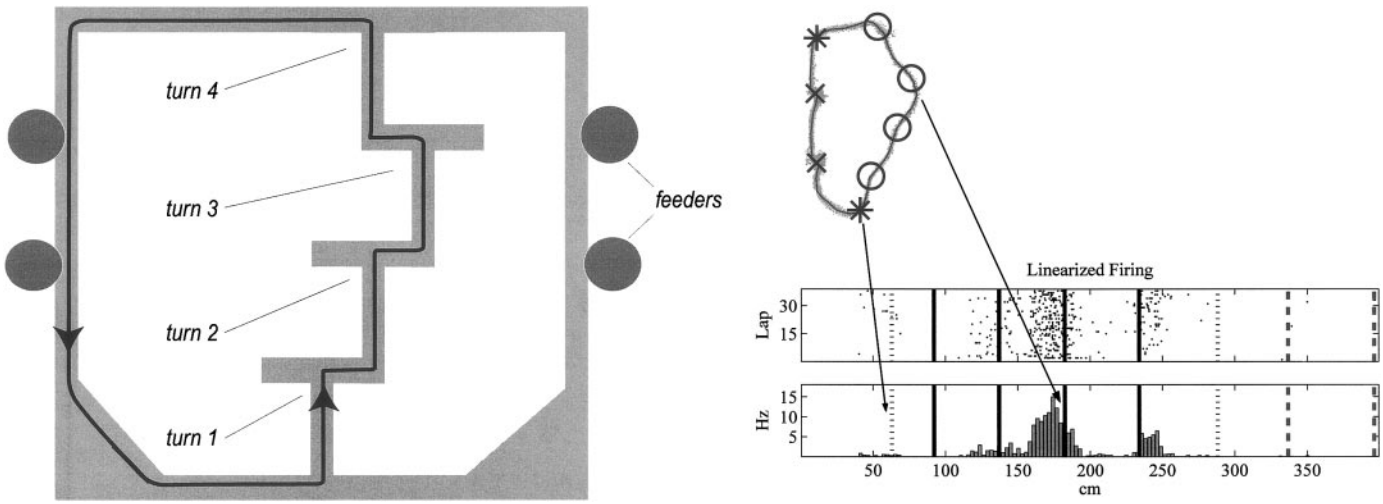

FIG. 1. *Left*: schematic of the multiple T maze. The path of the animal is indicated by the dark line. Filled circles: the locations of the feeders on the return rails. Each day, the turn sequence (which in this case was right-right-left-left) remained constant, but between days the turn sequence could be changed. On each day, only 1 pair of feeders was active (either the left or the right pair of feeders), providing a 4th choice to the turn sequence. The animal ran a continuous one-way loop, receiving food at the correct feeders on each trial. *Right*: linearized spatial plot. *Top*: position of the rat during a single session (gray points) and idealized path (solid line). The turn sequence (RRLL) is the same as in the schematic (*left*). Symbols indicate the location of the 8 spatial landmarks (circles, the 4 turns; asterisks, turns preceding and postceding the turn sequence; and ×, the 2 food-delivery sites). *Bottom*: spatial rastergram and histogram for a dorsal striatal phasic-firing neuron. Dotted lines, the start and end of the turn sequence; solid lines, turns; and dashed lines, the 2 food-delivery sites. Arrows, the correspondence between spatial landmarks (*top*) and the linearized plots (*bottom*). (R010-2001-12-19-TT03-01 maze = RRLL 38 trials)

experimental room remained constant from day to day. T choices were also swapped daily in Novel maze conditions. Sessions lasted for 40 min.

Two rats were implanted with hyperdrives before learning the multiple T task. Training for these animals began with the three T version of the task and then moved directly to the 3-wk novel-1/familiar/novel-2 protocol using four T mazes.

Surgery

Rats were implanted with 14-tetrode hyperdrives (David Kopf Instruments, Tujunga, CA) targeting the striatum. Twelve tetrodes were used to record neural activity, and two electrodes were used as references for common noise rejection. Tetrodes were constructed from four lengths of 0.0127-mm wire insulated with polyamide (Kanthal Precision Wire, Palm Coast, FL). Rats were anesthetized with Nembutal (pentobarbital sodium, 40–50 mg/kg, Abbott Laboratories, North Chicago, IL), and the area of the implantation was shaved. Rats were then placed on a stereotaxic apparatus (Kopf) and 0.1 ml Penicillin G benzathine and penicillin G procaine (dual-cillin, Phoenix Pharmaceutical, St. Joseph, MO) was injected intramuscularly into each hindlimb. During surgery, anesthesia was maintained using isoflurane (0.5–2% isoflurane vaporized in medical grade O_2). The scalp was then disinfected with alcohol and swabbed with Betadine (Purdue Frederick, Norwalk, CT). The skin overlying the skull was incised and retracted, and the underlying fascia was cleared from the surface of the skull. Excess bleeding was stopped by application of hydrogen peroxide followed by cautery of the retracted fascia. Anchor screws and one ground screw were placed in the skull, and a 1.8-mm diam craniotomy was opened using a surgical trephine (Fine Science Tools, Foster City, CA). The hyperdrive was positioned over the striatum (Bregma +0.5 mm AP, 3.0 mm ML) (Paxinos and Watson 1998) and lowered to 1 mm below the surface of the skull. The craniotomy was protected using silastic (Dow Corning 3140), and the hyperdrive was secured in place with dental acrylic (Perm Reline and Repair Resin, The Hygenic Corp, Akron, OH). After surgery, 10 ml sterile saline (0.9%) was administered subcutaneously, and all tetrodes were advanced ~1 mm. Animals were allowed to recover in an incubator until they were ambulatory, which was usually 1–2 h after surgery. Once animals were ambulatory, 0.8 ml acetaminophen (children's Tylenol) was administered orally. For 2 days after surgery, rats received water containing children's Tylenol (25 ml in 275 ml of water). Rats were allowed 2 days to recover from surgery before resuming experiments. Three animals received right-side implants, and two animals received left-side implants.

Histology

After the completion of all experiments, the location of each tetrode was marked by passing a small amount of anodal current through each tetrode (5 μA for 5 s), which causes a small lesion to form in the region of the tetrode tip. At least 48 h after gliosis, rats were killed with 1.0 ml Nembutal and perfused intracardially with saline followed by 10% formalin. Brains were removed and placed in 10% formalin overnight, then transferred to a 30% sucrose/formalin mixture for several days. Brains were sliced on a freezing microtome into 40-μm coronal sections (3 animals) or horizontal sections (2 animals) through the region of the hyperdrive implantation. Slices were stored in formalin at 4°C until staining. Sections were then mounted on gelatin-coated slides, dipped in ethidium bromide (Sigma Aldrich, St. Louis, MO) for 15 s, rinsed in dH_20, dehydrated, and coverslipped with DPX (Fluka Chemical, Ronkonkoma, NY).

Neurophysiology

RECORDING. After surgery, the 12 recording tetrodes were advanced ~160–640 μm per day until reaching the striatum. Arrival in the striatum was determined on the basis of the estimated depth of each tetrode and the passing of corpus callosum (which is quiet relative to the overlying cortex and underlying striatum). The striatum was further identified by the presence of slow firing cells, consistent with known properties of medium-spiny GABAergic neurons. After reaching the striatum, each tetrode was moved no more than 40 μm per day. The two reference electrodes were advanced in a similar manner to the corpus callosum.

Neural activity was recorded using a 64-channel Cheetah recording system (Neuralynx, Tucson, AZ). During recording sessions, a 72-channel motorized commutator (AirFlyte, Bayonne, NJ; Dragonfly, Ridgeley, WV; Neuralynx) allowed the rats to run the task without twisting the tether cables that connected the hyperdrive to the recording system. Tetrode channels were sampled at 32 kHz and filtered between 600 Hz and 6 kHz. When the voltage on any of the four channels of a single tetrode exceeded a threshold set by the experimenter, a 1-ms window of the spike waveform on each of the four channels on the tetrode was recorded to disk and timestamped with microsecond resolution. Spikes were clustered off-line into putative cells on the basis of their waveform properties using MClust 3.0 (Redish et al. 2002b), with automatic preclustering using KlustaKwik 1.0 (Harris 2002).

CLUSTER QUALITY. In each session, spike waveforms recorded on a single tetrode represented a mixture of spikes obtained from multiple sources including both neurons and noise events (chewing artifact, mechanical artifacts). After clustering spikes into putative cells, it was important to ensure that spikes assigned to one cluster were well separated from other spikes recorded simultaneously. A quantitative measure of cluster quality, L_{ratio}, was applied to measure how well each cluster (i.e., each putative cell) was separated from other clusters and noise events recorded on the same tetrode (Jackson et al. 2003).

For each spike, four features of the waveform recorded on each tetrode channel were calculated (energy and the 1st 3 principal components of the energy normalized waveform). Energy was calculated using the square root of the sum of squares of each sample in the waveform. For the principal-components analysis, a set of principal components based on a sample of striatal neurons were applied to the waveforms after normalization by energy.

The 16 feature quantities (4 tetrode channels $\times$ 4 features) defined each spike as a point in 16 dimensional space. For each cluster, the squared Mahalanobis distance (D^2) from the center of the cluster was calculated to every spike in the data set using the covariance matrix based on spikes assigned to the cluster. The Mahalanobis distance had the effect of scaling the spikes in the cluster to unit variance in all 16 dimensions. Under the assumption that the spikes in the cluster distribute normally in each dimension, D^2 for spikes in a cluster will distribute as χ^2 with 16 degrees of freedom (D'Agostino and Stephens 1986).

The amount of contamination of a given cluster is denoted by L and is calculated as the sum of the probabilities that each spike that is not a cluster member should actually be a part of the cluster. The probability of cluster membership for each spike is taken to be the inverse of the cumulative distribution function for the χ^2 distribution with 16 degrees of freedom. Spikes that are close to the center of the cluster will have probabilities approaching 1, whereas spikes far from the center of the cluster will have probabilities approaching 0. For each cluster, L was calculated by

$$L(C) = \sum_{i \notin \mathbf{C}} 1 - \mathrm{CDF}_{\chi^2_{\mathrm{df}}}(D^2_{i,C}) \qquad (1)$$

where $i \notin \mathbf{C}$ is the set of spikes that are not members of the cluster, $\mathrm{CDF}_{\chi^2_{\mathrm{df}}}$ is the cumulative distribution function of the χ^2 distribution with df = 16, and $D^2_{i,C}$ is the squared Mahalanobis distance of spike i from the center of cluster C. Spikes from other clusters or noise events that are close to the center of cluster C will have high probabilities and contribute strongly to this sum, whereas spikes far from the center of

cluster C will contribute little. A low value of L indicates that the cluster has a good "moat" and is well separated from other spikes recorded on the same tetrode. In contrast, a high value of L indicates that the cluster is not well separated and is likely to include both spikes that are not part of the cluster and exclude spikes that are part of the cluster. The cluster quality measure, L_{ratio} was defined as L divided by the total number of spikes in the cluster. Using a criterion based on L_{ratio} rather than L allows clusters with larger numbers of spikes to tolerate more contamination. Examples of representative clusters and their L_{ratio} values are shown in Fig. 4. For the analyses described in the following text, clusters with values of $L_{ratio} < 0.05$ were used. The results described in the following text were consistent with more strict thresholds of L_{ratio}.

UNIQUE SPIKE TRAINS. After reaching the striatum, tetrodes were often not advanced if cells were observed. In many cases, spike trains recorded in successive sessions represented multiple observations of the same cells. This repeated sampling of the same neurons in multiple sessions could create a bias in analyses of the proportion of cells which fired phasically or were task responsive. To correct for repeated sampling, a set of unique spike trains was defined. Spike trains obtained from each tetrode were matched across successive sessions on the basis of the correlation of their extracellular waveforms on all four tetrode channels in both sessions. Spike trains with very similar waveforms in successive sessions were considered likely to represent multiple observations of the same cell (for an example, see Fig. 4). After matching cells across sessions, a set of unique cells was created by selecting the spike train from each set of matched spike trains with the smallest L_{ratio}. This set of matched spike trains was also used to examine how the responses of some neurons changed between sessions.

CELL TYPE CATEGORIZATION. On the basis of the extracellularly recorded spike trains, striatal cells were grouped into two categories: phasic-firing neurons (PFNs) and tonic-firing neurons (TFNs). PFNs were distinguished from TFNs on the basis of the proportion of time spent in interspike intervals (ISIs) >5 s. This long-ISI proportion was calculated by finding all ISIs >5 s, summing these ISIs and dividing by the total session time. PFNs were then defined as cells with values of the long-ISI proportion >0.4, which indicated that these cells spent >40% of the session in ISIs >5 s. TFNs were defined as cells with values of the long-ISI proportion <0.4. The choice of 5 s as the definition of a long-ISI was based on a preliminary consideration of the data and definitions using 2- or 10-s ISIs yielded similar separations of PFNs and TFNs. Analyses described in the following text were applied to the set of phasic-firing neurons.

Behavioral data collection

Position of the rats during the task was determined using a battery-operated LED backpack constructed in the lab. The LED was secured in an elastic wrap and fastened together with Velcro, which allowed for snug fitting to the rat. Wearing the backpacks, rats were able to move without obvious difficulty, and the LEDs appeared to maintain stable positions on the rat over the course of a session. After implantation, additional LEDs were also present on the headstage, which plugged into the hyperdrive.

LED position was monitored by video tracking input to the Cheetah recording system sampled at 60 Hz. Food delivery was controlled with signals generated by in-house software running in Matlab (MathWorks, Natick, MA). Delivery of food pellets was recorded and timestamped by the Cheetah recording system.

Data analysis

IDEALIZED PATH. For the purpose of identifying errors and constructing linearized spatial rastergrams and histograms (described in the following text), an *idealized path* was created for each session by selecting a set of points that followed the path that the rat traveled through on a typical lap (one without errors). This set of points was interpolated linearly so that the distance between points was 1 pixel ($\sim 0.4 \times 0.4$ cm). A set of eight spatial landmarks on the idealized path was also selected: one for each food-delivery site, one for each turn, and two that marked where the rat turned to enter and exit the turn sequence (see Fig. 1).

ERROR QUANTIFICATION. Errors were defined as deviations from the idealized path at any of the choices, which occurred when rats made incorrect turns. At each turn, a 28-cm segment of the path centered on the location of the turn was examined. An error was scored on a lap if the rat deviated by >5 cm from the idealized path in this region for a total of ≥10 position samples (166 ms).

LINEARIZED SPATIAL PLOT CONSTRUCTION. To examine the responses of phasic-firing neurons as rats ran on the maze, a linearized spatial rastergram and histogram were developed. The location of the rat at the time of each spike was mapped to the nearest point on the idealized path, and the location of the spike relative to the idealized path was used to construct linearized rastergrams, and average firing across laps was used to construct linearized histograms (see Fig. 1).

WARPED LINEARIZED SPATIAL PLOTS. When rats ran a new maze each day, the spatial layout of the turn sequence varied from session to session. For presenting data from a set of PFNs recorded in different sessions, the linearized spatial histograms were warped to 20 bins between each of the eight spatial landmarks. The 20 bins surrounding each spatial landmark (the 10 bins before and the 10 bins after) were taken to represent activity of the PFN near the landmark.

FIRING RATE. To better estimate the instantaneous firing rate of each cell, a measure of continuous firing rate was calculated for each spike train by dividing the session into 10-ms bins and assigning each spike to one bin. The binned spike train was convolved with a Gaussian ($\sigma = 100$ ms) to create a continuous function of estimated firing rate sampled at discrete intervals of 10 ms. This firing rate estimate was used in calculations of task responsiveness, in the creation of phasic firing fields, and in the temporal versus spatial encoding analysis (described in the following text).

TASK RESPONSIVENESS. The responses of PFNs to task parameters were classified using firing rate relative to the time of arrival at each food-delivery location and position on the maze. Two-sample Kolmogorov-Smirnov tests were applied to determine if the population of PFNs was modulated by these task parameters. The null hypothesis was that PFN firing rates were not modulated by either location or food delivery. Then the distribution of firing rates with respect to either task parameter would not differ from the distribution of average PFN firing rates (calculated across the entire recording session). A significant Kolmogorov-Smirnov test would indicate that these task parameters (food delivery and location on the maze) did modulate PFN firing rates.

After determining if PFNs as a population were responsive to task parameters, individual PFNs were tested for task responsiveness. PFNs were classified as *reward-responsive* if they showed a significant increase in firing rate during the 5 s after arrival at either food-delivery site. PFNs were classified as *maze-responsive* if they showed a significant increase in firing rate when the rat was running on the maze. To determine if a firing rate was significantly elevated, the mean firing rate of each PFN relative to task events was compared with a distribution of expected mean firing rates created from the same spike train using shuffled event times (i.e., a bootstrap) (see Efron 1982).

With a large number of expected mean firing rates, the mean ± SD of the distributions of expected mean firing rates can be used as estimates of what the cell's firing rate should be if the cell is not responsive to task parameters. The distributions of expected mean firing rates that were obtained for PFNs frequently exhibited a skew toward positive values, and a square-root transform was applied to normalize the distribution (Sokal and Rohlf 1995). Under the assump-

tion of normality, estimates of μ and σ from the expected mean firing rate distributions were used to calculate the probability of observing the cell's actual mean firing rate using the inverse of the normal cumulative distribution function. An α of 0.05 was adopted, and a Bonferroni correction for multiple comparisons was applied.

PFNs were classified as reward-responsive if the mean firing rate of the PFN in the 5 s after arrival at either food-delivery location was significantly larger than the distribution of expected mean firing rates created from 5-s time segments selected randomly from the session. PFNs were classified as maze-responsive if the mean firing rate at any location on the maze was significantly larger than the distribution of expected mean firing rates created from similar length time segments selected randomly from the session. For determining maze-responsiveness, the idealized path in each session was divided into eight regions using the spatial landmarks described in the preceding text. If any of these regions was >1.5 times the average distance between successive turns on the maze, these regions were divided in half.

PFNs that fired very infrequently during the session tended to produce quantized distributions of expected mean firing rates. Such quantized distributions were not normal following the square-root transform. Therefore all PFNs which fired <100 spikes were not considered any further in our analyses because not enough spikes were observed in the session to accurately estimate the cell's responsiveness to task parameters.

PHASIC FIRING FIELDS. To examine the size and distribution of maze responses in PFNs, a quantification of each maze-responsive PFN′s activity on the maze was defined. For each maze-responsive PFN, phasic firing fields (PFFs) were defined as each set of continuous 5-cm bins on the linearized spatial histogram that exceeded 50% of the PFN′s maximum firing rate in any bin.

SPATIAL VERSUS TEMPORAL ENCODING. A correlation analysis was performed to determine if maze responses were better related to the location of the rat or to temporal events. For each maze-responsive PFN, the average firing rate at each position along the maze was determined for each lap using the continuous firing rate measure. The correlation of the firing rate as a function of position was calculated between every pair of laps, and the average correlation was taken to represent how well related the PFN′s activity was to the location of the animal on the maze.

For temporal measures, the firing rate of each maze-responsive PFN was calculated for each lap over two temporal windows: the 20 s preceding the arrival at the first food-delivery site and the 20 s after departure from the second food-delivery site. For each measure, the correlation of the PFN′s firing rate relative to either arrival or departure was calculated for every pair of laps, and the average correlation served to describe how well related a cell's activity was to the time of arrival at the first food-delivery site and the time of departure from the second food-delivery site.

RESULTS

Error elimination

Over 105 sessions, rats ran an average of 52.4 ± 8.3 laps per session. Figure 2 shows the average probability of making an error on each lap over the 76 sessions (23 from novel-1, 26 from familiar, 27 from novel-2) in which rats ran ≥40 laps. Sessions in which rats ran novel and familiar mazes are plotted separately. For convenience, weeks 1 and 3 (novel-1 and novel-2, in which rats ran novel mazes) have been combined. Rats performed at chance on the first lap when running novel mazes and eliminated errors over the first five laps of the session.

When rats ran familiar mazes, performance on the first lap was better than chance, and they made fewer errors compared with novel mazes during the first five laps. To examine the differences in error-elimination across the session and between novel and familiar mazes, a repeated-measures ANOVA was conducted on the average number of errors obtained in the first 40 laps (divided into 4 10-lap blocks) in each week (novel-1, familiar, novel-2). To control for the departures from sphericity that are possible when using a repeated measures design, all significant ANOVAs were checked using the degrees of freedom correction of Greenhouse and Geisser (1959). There was a significant effect of week [novel-1, familiar, novel-2, $F(2,8) = 11.64$, $P < 0.05$], and block [$F(3,12) = 81.24$, $P < 0.05$], and a significant interaction [$F(6,24) = 9.92$, $P < 0.05$]. Post hoc tests (Tukey-Kramer HSD, $\alpha = 0.05$) revealed that in each week, more errors were made in the first 10 laps than in the rest of the session. In the first 10 laps, the number of errors made in the two novel maze conditions did not differ, and significantly fewer errors were made in the familiar maze condition than either novel maze condition.

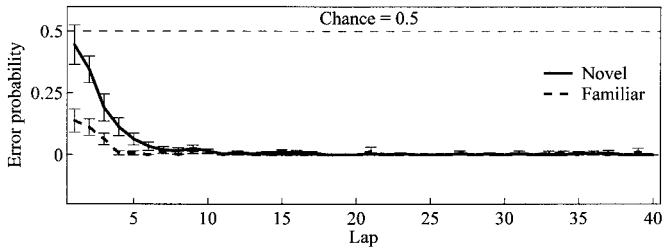

FIG. 2. Errors across laps. Shown is the average error-probability (number of errors made divided by 4) over the 1st 40 laps, using 76 sessions in which rats completed ≥40 trials (23 sessions from novel-1, 26 sessions from familiar, 27 sessions from novel-2). Bars represent SE across 5 rats. The solid line represents the probability of making an error on novel mazes, whereas the dotted line represents the probability of making an error on familiar mazes. On novel mazes, errors were within chance levels on the 1st trial and declined prominently over the 1st 4 laps. On familiar mazes, errors were below chance on the 1st lap and also declined prominently over the first four laps. These data show that rats quickly eliminated errors and that fewer errors were made in the error-elimination phase when rats ran familiar mazes.

Cell-type categorization

All final tetrode locations were histologically verified to be in the dorsal striatum. Tetrode tracks were observed in a region extending approximately −0.5 to 1.5 mm anterior/posterior and 1.6–3.6 mm medial/lateral with respect to bregma (see Fig. 3).

From 104 sessions, 2,125 spike trains were obtained (21 sessions from 4 rats, 20 sessions from one rat. 425 ± 64 spike trains per rat). See Table 1 for a description of the spike train data collected from each rat. Between successive sessions, spike trains recorded on the same tetrode had an average waveform correlation of 0.729 ± 0.009 (SD). Spike trains which were matched across sessions (i.e., considered to be the same cell in both sessions) had an average waveform correlation of 0.977 ± 0.002 (SD). As an example of the waveform correlations and stability of cells across sessions, two successive sessions from a representative tetrode are shown in Fig. 4. As can be seen in the figure, the waveform correlations successfully identified clusters that were stable across the two sessions. 1,144 spike trains were judged to be unique on the basis of the correlation of the spike waveforms between sessions, as described in METHODS (229 ± 28 unique spike trains per rat).

Two hundred seventy-five spike trains with values of

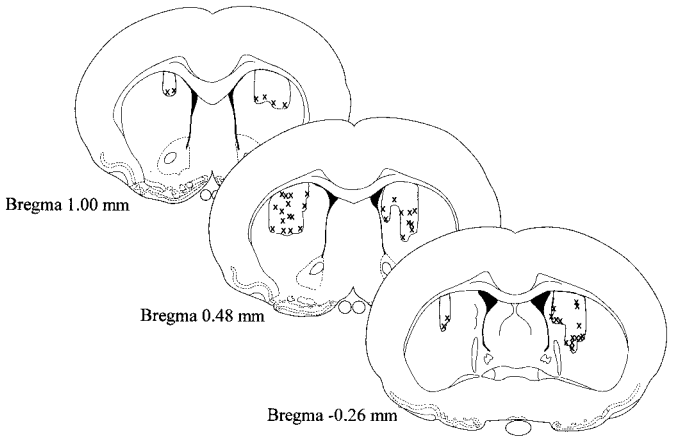

FIG. 3. Recording locations verified histologically. ×, final tetrode positions; ▨, the region of striatum that was sampled as tetrodes were advanced through the striatum. All tetrode locations have been mapped to the nearest of the 3 coronal sections shown. Tetrodes were observed in a region extending appoximately −0.5–1.5 mm anterior/posterior relative to bregma. Diagrams adapted from Paxinos and Watson (1998).

$L_{ratio} > 0.05$ were removed from this set of unique spike trains, leaving a total of 867 spike trains (173 ± 22 spike trains per rat). Analyses were restricted to this set of unique spike trains, but the results reported here were consistent with the entire data set of 2,125 spike trains.

From the set of unique spike trains, 589 (68%) were classified as phasic-firing neurons (PFNs). The remaining 278 (32%) were classified as tonic-firing neurons (TFNs). As shown in Fig. 5, PFNs were well separated from TFNs.

PFNs tended to have at least two firing modes (an active mode during which the cells fired at ~10 Hz and a quiescent mode), whereas TFNs tended to fire in one tonic mode. Figure 5 (*A* and *D*) shows the average ISI histogram for PFNs and TFNs, constructed using normalized ISI histograms. The average TFN ISI histogram was unimodal (Fig. 5*D*) as was the ISI histogram of a typical TFN (Fig. 5*E*). This indicates that TFNs tended to fire regularly at a single rate (see also Fig. 5*F* for an example). The average PFN ISI histogram had a peak at short ISIs of ~100 ms, corresponding to a firing rate of ~10 Hz and a shoulder at longer ISIs. This pattern was also observed in individual PFNs (Fig. 5*B*). PFNs tended to have their spikes organized into bouts of activity, separated by periods of quiescence (see Fig. 5*C* for an example).

As described in METHODS, PFNs with <100 spikes had too few spikes to estimate responsiveness to task parameters. Of the 589 PFNs in the unique spike train sample, 194 contained <100 spikes and were eliminated from further analysis.

Task responses

To determine if PFN firing rates were modulated by two task parameters (the location of the rat on the maze and the delivery of food) Kolmogorov-Smirnov goodness-of-fit tests for two samples were applied. The null hypothesis was that PFN firing rates were not modulated by either location or food delivery and that these distributions of firing rates would not differ from the distribution of average PFN firing rates (calculated across the entire recording session). The distribution of average firing rates of PFNs on the maze was significantly different from the distribution of average PFN firing rates ($P < 0.0001$). Also the distribution of PFN firing rates in the 5 s after arrival at either food-delivery site was significantly different from the distribution of average PFN firing rates ($P < 0.0001$). These results indicate that the firing rates of at least some PFNs were modulated by the location of the animal and by the delivery of food. Posthoc tests (described in TASK RESPONSIVENESS) were conducted to identify which PFNs were task-modulated. In general, PFNs had low average firing rates and thus only increases in activity were tested. Of 395 PFNs, 108 (22 ± 9 cells per rat) PFNs were responsive in one or more regions on the maze and 81 (16 ± 7 cells per rat) PFNs were responsive during the 5 s after arrival at one or both food-delivery sites.

Maze-responsive PFNs

Of the PFNs which fired ≥100 spikes a session, 27% were classified as maze-responsive. Figure 6 shows an example of a maze-responsive PFN that was active at one location on the maze (as the rat ran between the last turn on the maze and the 1st food-delivery location). Maze-responsive PFNs had between one and six phasic firing fields (PFFs, median number of PFFs per cell = 1), with a median PFF width of 3 bins (~15 cm).

MAZE RESPONSES WERE RELATED TO SPATIAL CUES. Maze responses were better related to the position of the rat on the maze than to the time at which the rat arrived at the first food-delivery site or the time at which the rat departed the second food-delivery site. Across laps, the correlation of each maze-responsive PFN′s activity was examined with respect to: the position of the rat on the maze, the 20-s preceding arrival at the first food-delivery site, and the 20 s after departure from the second food-delivery site. Higher correlations were obtained for the spatial reference frame than either temporal reference frame. Shown in Fig. 7, *left*, is a PFN with maze responses that were well related to the location of the rat on the maze and poorly related to both temporal reference frames. The jitter in spike timing observed in the temporal plots is indicative of behavioral variability of the animal across laps.

TABLE 1. *Spike trains for each rat*

Rat	Sessions/Rat	No. of Spike Trains	No. of Unique Spike Trains	# L_{ratio} <= 0.05	No. of PFNs	No. of TFNs
R010	20	465	278	187	150	37
R011	21	350	210	150	99	51
R016	21	322	200	157	132	25
R018	21	355	160	131	94	37
R023	21	633	296	242	114	128
Mean		425.0 ± 64.2	228.8 ± 28.3	173.4 ± 21.7	117.8 ± 11.7	55.6 ± 20.8

Values are means ± SE. Number of unique spike trains identifies those spike trains from the third column that were classified as unique. The next column list the number of unique spike trains that had values of the cluster quality measure L_{ratio} <= 0.05. The last 2 columns refer to spike trains from the 5th column.

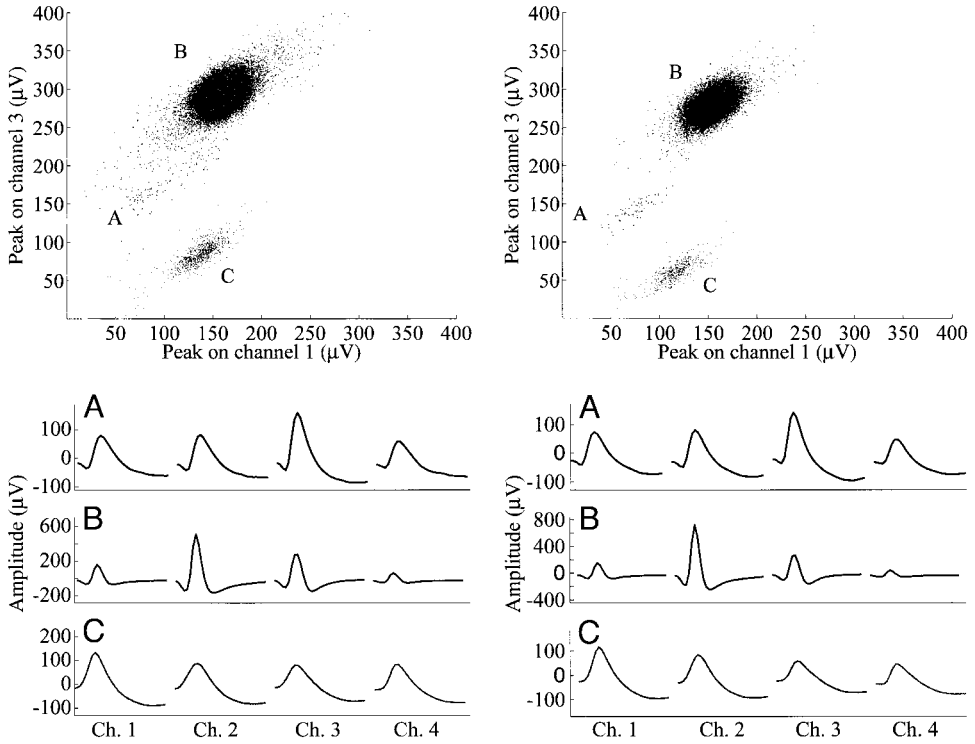

FIG. 4. Example of cells matched across sessions. Data from the same tetrode recorded in successive sessions shows recording stability across days. Three clusters were isolated from each tetrode recording and were matched across sessions on the basis of the correlation of the average waveforms in each session. *Top*: peak waveform amplitudes on channel 3 vs. channel 1 for both sessions. Points identified as members of 1 of the 3 clusters identified in these recordings are plotted black, all other points are shown in gray. *Bottom*: average waveforms of the 3 clusters shown above. Waveform correlations of the matched spike trains between sessions (i.e., correlation of the average waveform of cluster A in session 1 with the average waveform of cluster A in session 2, etc.) were >0.98. Values of L_{ratio} for these clusters were: *left*: cell A = 8.9×10^{-6}, cell B = 0.00, cell C = 0.0028; *right*: cell A = 1.3×10^{-4}, cell B = 0.00, cell C = 0.0004. Data from R018-2002-09-22-TT01 and R018-2002-09-23-TT01.

Across rats, the median amount of time rats took from leaving the second food-delivery site to arriving at the first food-delivery site was 13.5 ± 1.0 s. The within-session SD of this departure-to-arrival time was 9.3 ± 2.7 s. This within-session variability allowed for the dissociation of responses to spatial and temporal events.

Across the set of 108 maze-responsive PFNs, there were significantly larger correlations in the spatial reference frame than either temporal reference frame (see Fig. 7 *right*). Paired *t*-test across five animals: space versus arrival, $t(4) = 3.593$, $P = 0.023$, space versus departure, $t(4) = 3.560$, $P = 0.024$).

MAZE RESPONSES DISTRIBUTE EVENLY ON THE TURN SEQUENCE. To determine if the PFN maze responses favored specific locations on the maze, the distribution of phasic firing fields (PFFs) over the maze was examined. Figure 8 shows the PFFs obtained from maze-responsive PFNs sorted according to the center of each PFF. Maze-responsive PFNs responded at every location along the length of the turn sequence. The R^2 from a linear regression on the sorted PFF centers was 0.98, indicating that the centers of the PFFs were well described by a linear fit. This implies that maze-responses were uniformly distributed on the maze and did not concentrate at specific landmarks.

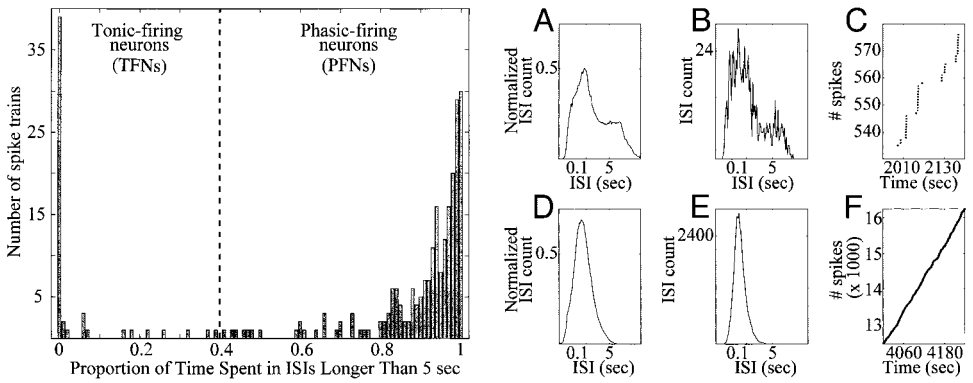

FIG. 5. Categorization of striatal cells. *Left*: histogram of the proportion of time spent in interspike intervals (ISIs) >5 s. The dashed line indicates the threshold separating PFNs from other striatal cells (40% of time spent in ISIs >5 s). All cells to the right of the dotted line were classified as PFNs. *Right*: average and typical striatal ISI histograms and spike count functions. Average ISI histograms are shown for all PFNs (*A*) and TFNs (*D*). *B*: a typical ISI histogram from a PFN. *E*: a typical TFN ISI histogram. All ISI histograms in *A, B, D,* and *E* are plotted on the log scale. PFNs typically showed a bimodal distribution of ISIs on the log scale with a peak at ~100 ms, corresponding to a firing rate of 10 Hz, whereas TFNs typically showed a unimodal distribution of ISIs. *C*: a representative 4-min window showing the number of spikes obtained from the same PFN as in *B*. This PFN had bouts of activity separated by periods of quiescence. *F*: a representative 4-min window showing the number of spikes obtained from the same TFN as in *E*. This TFN did not pause for long periods of time and did not show activity organized into bouts.

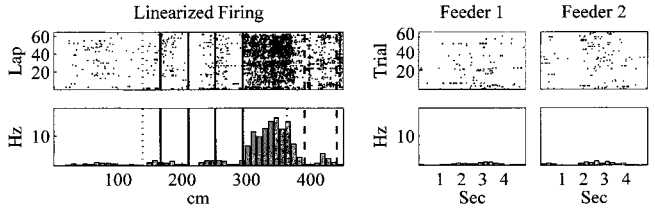

FIG. 6. A maze-responsive PFN. *Left*: rastergram and histogram of linearized firing on the maze. Key as shown in Fig. 1. *Right*: peri-event time histograms (PETHs) of the firing rate of the cell as the rat arrived at the 1st and 2nd food-delivery sites. This cell was active as the rat ran between the 4th turn and the end of the turn sequence. (R023-2002-08-27-TT07-02 maze = LRLL trials = 61)

MAZE RESPONSES DID NOT ENCODE GENERAL ACTIONS. Given that maze responses did not favor specific locations on the maze, did maze responses depend on the specific actions that the rats were making? From the example shown in Fig. 1, we might not expect this to be the case. This PFN had responses at two left turns on the maze: a strong response at the third turn (~40 Hz) and a weak response at the fourth turn (~10 Hz). This PFN was not equally responsive to all left turns, however. It was silent at two other left turns, where the rat was entering and exiting the turn sequence. To examine how the activity of maze-responsive PFNs as a group depended on the actions that the rats were making, the firing rate in the phasic-firing fields of maze-responsive PFNs was compared with the firing rate of the same PFN at other regions of the turn sequence where the shape of the rat's path was highly similar or dissimilar. Similar and dissimilar paths were defined as locations where the rat's path was well and poorly correlated, respectively. In cases where the shape of the paths were very similar, the rat was likely to be making similar motor actions, such as turning in the same direction. In cases where the shapes of the paths were dissimilar, the rat was likely to be making different motor actions, such as turning in opposite directions.

Of the 108 maze-responsive PFNs, 58 cells (from 5 animals, 11.6 ± 6.0 cells per rat) had at least one PFF on the turn sequence. The firing rate and path of the rat in each PFF was compared with other locations on the turn sequence using a sliding comparison window the size of which was equal to that of the PFF. For each comparison window, the firing rate in the PFF as a function of position along the idealized path was correlated with the firing rate as a function of position in the

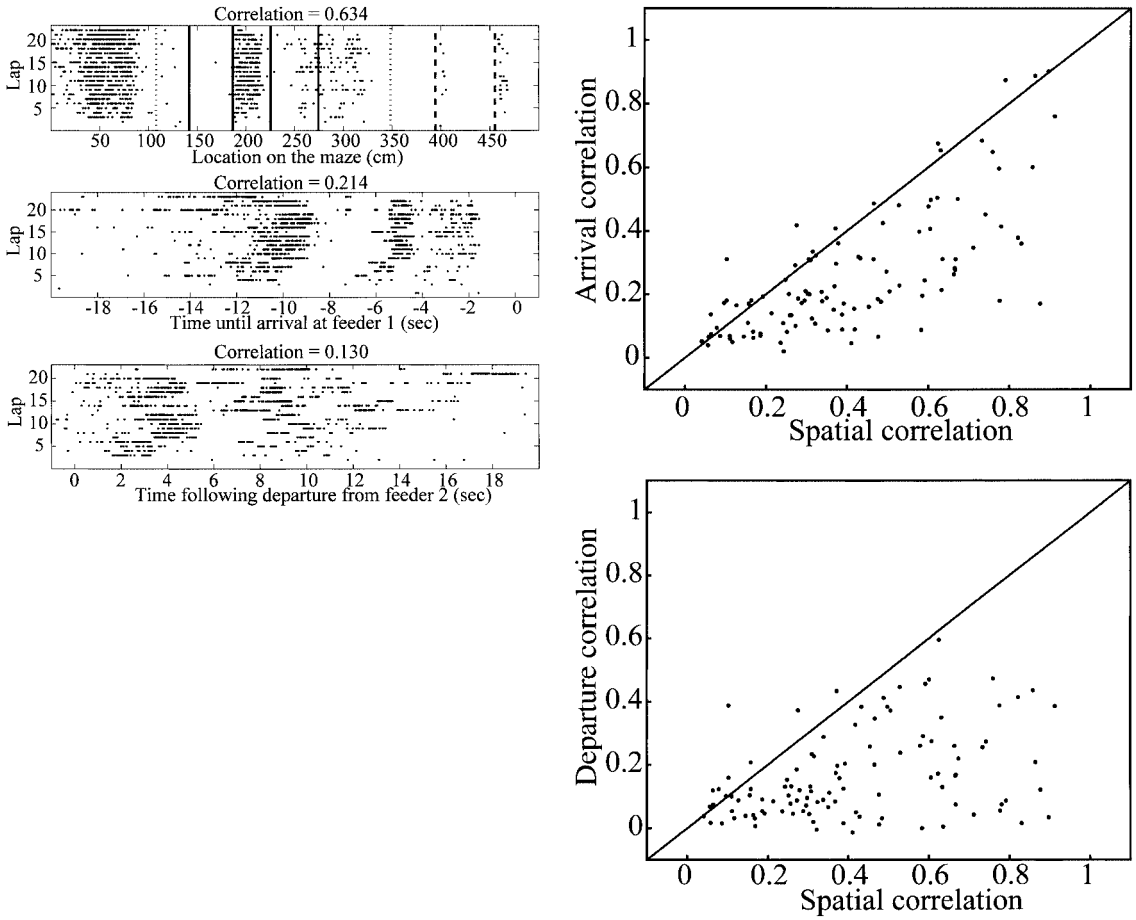

FIG. 7. Maze responses were better related to spatial than temporal parameters. *Left*: example of the activity of a maze-responsive PFN with respect to location on the maze (*top*), in the 20 s preceding arrival at the 1st food-delivery site (*middle*), and in the 20 s after departure from the 2nd food-delivery site (*bottom*). The same set of spikes is shown in each plot. The activity of this PFN was stable with respect to space and variable with respect to the two temporal events. Across laps, this maze-responsive PFN′s activity was better correlated in the spatial reference frame (correlation = 0.63) than in either temporal reference frame (correlations of 0.21 and 0.13). *Right*: across all maze-responsive PFNs, there was a significant bias toward higher correlations in the spatial reference frame than in the 20 s preceding arrival at the 1st food-delivery site (*top*) or in the 20 s after departure from the 2nd food-delivery site (*bottom*).

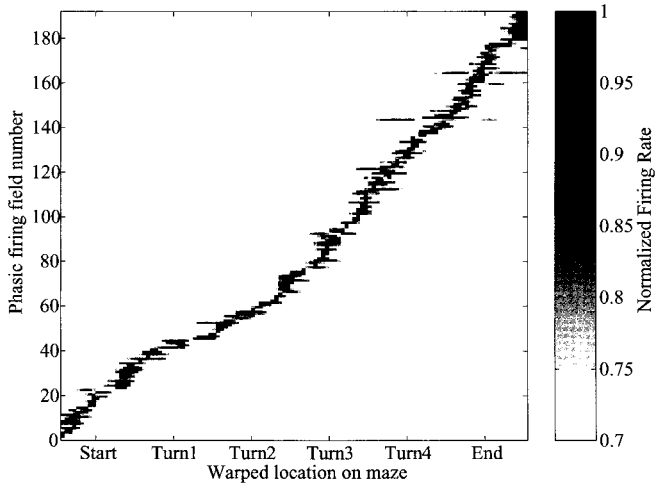

FIG. 8. Distribution of spatial activations. To make comparisons across sessions (i.e., across different maze configurations), the response of each PFN on the turn sequence was warped to a fixed number of bins between landmarks (see METHODS for a description of the warping process). For each responsive PFN, phasic firing fields (PFFs) were defined as each set of continuous bins on the warped spatial histogram which exceeded 50% of the PFN's maximum firing rate in any bin considered. A PFN could have multiple phasic firing fields. In the plot above, each PFF has been aligned to its center, and normalized to a maximum of 1 for display purposes. The field centers are well fit with a linear regression ($R^2 = 0.98$), implying that PFF centers are distributed evenly across the turn sequence.

comparison window. The rat's path in the PFF was also correlated with the rat's path in the comparison window to quantify how similar the route taken by the rat in the comparison window was to the route taken in the PFF. If maze-responsive PFNs encoded general actions, then the firing pattern in the PFF should be highly similar to the firing pattern observed at other locations where the rat's route was similar to the route taken through the PFF. Also we would expect that the firing pattern in the PFF should be poorly correlated with the firing pattern observed at other locations where the rat's route was dissimilar to the route taken through the PFF. In these analyses, a similar route was defined as having a path correlation >0.85, and a dissimilar route was defined as having a path correlation less than -0.85. In Fig. 9, we can see that there was a significant bias toward similar firing patterns in other regions on the maze where the rats' paths were similar compared with other regions on the maze where the rats' paths were dissimilar [Kruskal-Wallis $\chi^2(1, n = 116) = 9.98$, $P = 0.0016$]. However, neither group of correlations was biased toward positive correlations, which indicates that even in the same route group, there was no tendency for a maze-responsive PFN to fire similarly at other locations on the maze in which the rat's path was similar to that taken through the PFF. To the extent that the shape of the rat's path is an indication of the motor activity of the rat, these results indicate that maze-responsive PFNs did not purely encode movements.

SEQUENCE SPECIFICITY. Maze responses of PFNs were not well described by the actions rats were taking, but they were well described by a combination of action, sensory-context, and the specific sequence of turns rats were presented with. Twenty-two maze-responsive PFNs were recorded in at least two successive sessions and had at least one PFF on the turn sequence in at least one of these sessions. For each maze-responsive PFN, the correlations between firing pattern and path of the rat were tested. Pairs of sessions were considered in which the center of the PFF and center of the same region in the matching session were in the same location in the environment. Three groups of correlations were examined: same maze: cases where the rat ran the same sequence of turns in both sessions; same route: cases where the rat ran a different sequence of turns in each session, but the path taken through the PFF and the path taken through the same region of the matching session were correlated by ≥ 0.85; and different route: cases where the rat ran a different sequence of turns in each session, but the path taken through the phasic-firing field and the path taken through the same region of the matching session were correlated by less than -0.85. To correct for multiple observations of the same cell, the set of correlations in each group obtained from one maze-responsive PFN and its matching spike trains were averaged.

In the group of same-maze pairs, correlations from 14 cells were included (from 4 rats, 3.5 ± 1.5 cells per rat). In the group of same-route pairs, correlations from four cells were included (from 3 rats, 1.3 ± 0.4 cells per rat). In the group of different-route pairs, correlations from 10 cells were included (from 5 rats, 2.0 ± 0.5 cells per rat). Shown in Fig. 10 are the firing rate correlations for each group. There was a significant difference between groups [Kruskal-Wallis $\chi^2(2, n = 28) = 12.85$, $P = 0.0016$], and pairwise comparisons using Wilcoxon rank-sum tests ($\alpha = 0.05/3$) revealed that firing rate correlations in the same-maze group were significantly higher than the different-route group ($P = 0.0006$) but did not differ from the same-route group ($P = 0.574$). The firing rate correlations in the same-route group were higher than those of the different-route group, but these differences were not significant ($P = 0.036$). These results indicate that maze responses were highly similar

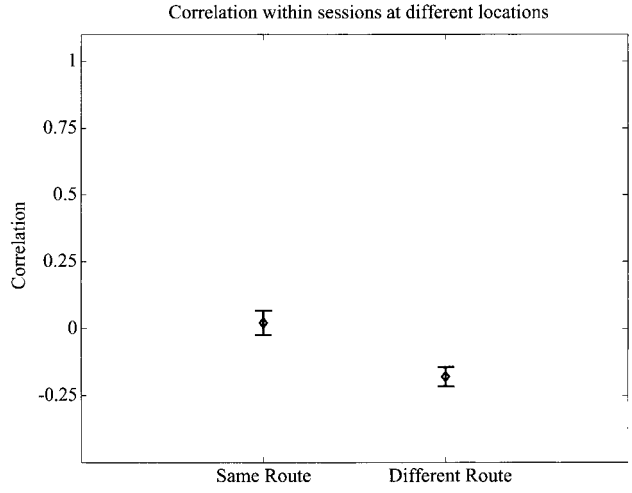

FIG. 9. Distribution of firing rate correlations for similar and dissimilar routes. For 56 maze-responsive PFNs that had phasic-firing fields (PFFs) on the turn sequence, the firing pattern and path of the rat in the PFF was correlated with the firing pattern and path of the rat in windows shifted over the length of the turn sequence. The same-route group includes firing rate correlations from cases where the rat's path was correlated by ≥ 0.85, indicating that the rat ran through a very similar route in both windows. The different-route group includes firing rate correlations from cases where the rat's path was correlated by less than -0.85, indicating that the rat ran through a very dissimilar route in both windows. There was no bias in either group toward positive correlations, indicating that these maze-responsive PFNs did not respond entirely on the basis of the shape of the rat's route. Bars represent means $\pm$ SE across cells.

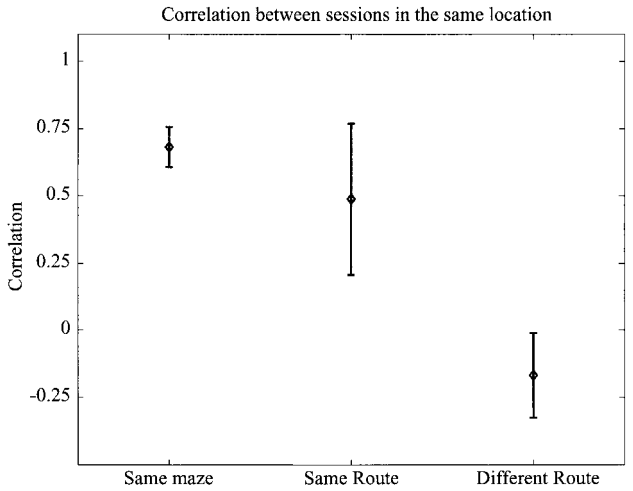

FIG. 10. Correlation of maze responses between sessions. For 22 maze-responsive PFNs that were observed in ≥2 sessions and had a PFF on the turn sequence, the correlation of the firing rate observed in the PFF on each session was correlated with the firing rate in the other session at the same location on the maze. These correlations were divided into 3 groups: same maze: correlations obtained from pairs of sessions in which rats ran the same sequence of turns; same route: correlations obtained from pairs of sessions in which rats ran different sequences of turns but had similar paths in the region of the PFF; and different route: correlations obtained from pairs of sessions in which rats ran different sequences of turns and had dissimilar paths in the region of the PFF. There was no significant difference between the same maze and same-route groups. The same maze condition was significantly higher than the different-route group, and there was a nonsignificant trend for the same-route group to be more highly correlated than the different-route group. These results indicate that maze-responsive PFNs maintained highly similar responses day to day when the same turn sequence was presented. When the turn sequence was changed, but the location of the PFF was the same and the route taken by the rat through the phasic firing field remained the same, there was a bias toward higher correlations, but not significantly so. Bars represent means ± SE across cells.

when the same turn sequence was presented and were not similar when rats took a different route through the same physical location in the environment. Based on the results of the same-route group, our data also indicate that maze responses did not purely encode sensory-context/action relationships: when rats ran through similar paths in the same physical locations in the environment, but ran a different turn sequence, correlations were intermediate between both the same-maze and different-route groups. The same-route group contained data from a small number of cells, and it could be the case that some PFNs were responding to purely sensory-context/action relationships while other PFNs were further modulated by the specific sequence of turns presented. Maze responses may thus have reflected a combination of information related to the specific actions performed, the sensory environment those actions were performed in, and in some cases the specific turn sequence presented.

Reward-responsive PFNs

Of the PFNs which fired ≥100 spikes a session, 21% were classified as reward responsive. Figure 11 shows an example of a reward-responsive PFN that was active after arrival at the first food-delivery site but not at the second food-delivery site. Of reward-responsive PFNs, 31 (38%) had a significant response only at the first food-delivery site, 33 (41%) had a significant response only at the second food-delivery site, and

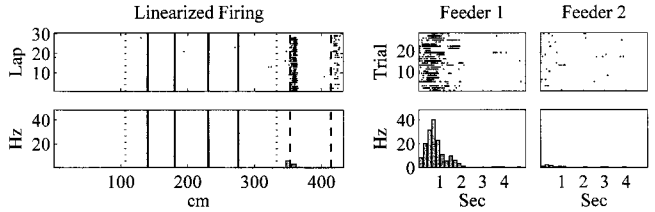

FIG. 11. A reward-responsive PFN. *Left*: rastergram and histogram of linearized firing on the maze. Key as shown in Fig. 1. *Right*: PETHs of the firing rate of the cell as the rat arrived at the 1st and 2nd food-delivery sites. This PFN had a large response at the 1st food-delivery site but not at the 2nd food-delivery site (R011-2002-02-16-TT09-03 Maze = LRRR trials = 28).

17 (21%) had a significant response at both food-delivery sites. Because 79% of the reward-responsive PFNs were responsive at only one of the food-delivery sites, these cells did not encode general aspects of food retrieval or consumption (e.g., chewing), which occurred at both food-delivery sites.

Maze- and reward-responsive PFNs are separate populations

In Fig. 6, the maze-responsive PFN was not active after arrival at either food-delivery location, and in Fig. 11, the reward-responsive PFN was not active while the rat was running on the maze. This segregation of maze and reward responses was a characteristic of the entire task-responsive PFN population. Based on the proportions of PFNs that were maze responsive (27.3%) or reward responsive (20.5%), we would expect that 5.6% of the PFNs (~22 cells) would have been responsive to both maze and reward if the probability of being a maze-responsive PFN and a reward-responsive PFN was independent. This was not the case: only 1/395 PFNs (0.3%)

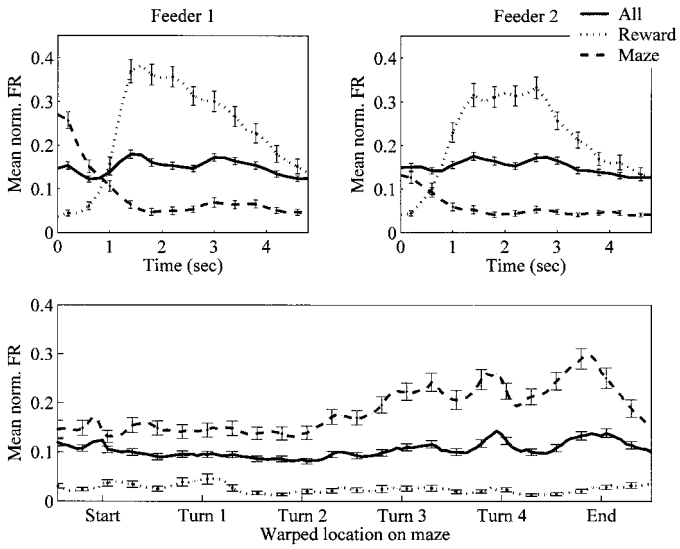

FIG. 12. Segregation of task responsive PFNs. Shown are the average normalized response of PFNs to the delivery of food (*top*) and on the turn sequence (*bottom*). To make comparisons across sessions (i.e., across different maze configurations), the response of each PFN on the turn sequence was warped to a fixed number of bins between landmarks (see METHODS for a description of the warping process). For all plots, —, the entire PFN population; - - -, maze-responsive PFNs; · · · , reward-responsive PFNs. After the delivery of food, maze-responsive PFNs were inhibited relative to reward-responsive PFNs and the entire PFN population. On the turn sequence, reward-responsive PFNs were inhibited relative to maze-responsive PFNs and the entire PFN population.

was classified as both reward and maze responsive. This is significantly less than we would expect by chance [$\chi^2(3) = 35.0$, $P < 0.001$].

Figure 12 shows the average normalized firing rate of PFNs after the arrival of the rat at the food-delivery sites (*top*) and on the turn sequence (*bottom*). Reward-responsive PFNs were more active after food delivery than the entire population of PFNs, whereas maze-responsive PFNs were more active on the maze than the entire population of PFNs. These results follow from the definitions of reward and maze responsiveness. However, reward-responsive PFNs were also less active on the maze than either maze-responsive PFNs or the entire PFN population. Likewise, maze-responsive PFNs were less active after food delivery than either reward-responsive PFNs or the entire PFN population. Our analyses allowed PFNs to be classified as both reward and maze responsive, but cells predominantly responded to one or the other parameter, not both. As such, these results further indicate that reward- and maze-responsive PFNs were separate populations of cells.

The differences between groups of PFNs were significant. In the 5 s after food delivery, there was a significant effect across animals of group [all PFNs, reward-responsive, maze-responsive, $F(2,12) = 31.2$, $P < 0.001$]. Post hoc comparisons (Tukey-Kramer HSD, $\alpha = 0.05$) revealed that the activity of maze-responsive PFNs was significantly less than the entire population of PFNs. Similarly, on the maze, there was a significant effect of group [all PFNs, reward-responsive, maze-responsive, $F(2,12) = 48.6$, $P < 0.001$]. Post hoc comparisons (Tukey-Kramer HSD, $\alpha = 0.05$) revealed that the activity of reward-responsive PFNs was significantly less than the entire population of PFNs.

These results indicate a separation of information processing such that PFNs that responded while the rats ran on the maze did not respond during food receipt and PFNs that responded during food receipt did not respond while the rats ran on the maze.

DISCUSSION

The present study examined the activity of neurons in the dorsal striatum as rats ran a multiple T maze. Analyses were conducted on neurons that fired phasically (PFNs), which are likely to have been medium spiny projection neurons (Kimura et al. 1990). Two types of task-responsive PFN were found: neurons that were active when rats were running on the maze (maze responsive) and neurons that were active when rats received reward (reward responsive). Maze responses were well related to the rats' position on the maze and not to temporal parameters. The locations of maze responses were uniformly distributed along the rat's path, while reward responses consisted of a phasic activation during reward receipt. Neither maze nor reward responses encoded general motor parameters such as turning or chewing. Over multiple sessions, maze responses tended to be similar if rats ran through a similar path in the same location in the environment and were most similar when rats ran the same sequence of turns. Two separate populations of PFNs encoded maze and reward responses, which suggests a segregation of navigation- and reward-related information processing in the dorsal striatum.

Striatal representation

Lesions and inactivations of the dorsal striatum in rodents impairs performance of habitual, stimulus-response (S-R) tasks (Kesner et al. 1993; McDonald and White 1993; Packard 1999; Packard and McGaugh 1992; Packard et al. 1989) as well as longer chains of sequential behavior (Berridge and Whishaw 1992; Cromwell and Berridge 1996; Matsumoto et al. 1999; Miyachi et al. 1997). One theory of how the dorsal striatum learns and produces S-R behavior is that striatal projection neurons with connections to motor centers encode S-R relationships by responding specifically to complex cortical inputs (Graybiel 1998; Graybiel et al. 1994). A number of studies have shown that striatal neurons have highly specific responses to task parameters that could encode stimulus-response relationships. Studies in the rat (Carelli et al. 1997; Gardiner and Kitai 1992) and primate (Kermadi and Joseph 1995; Kermadi et al. 1993; Kimura 1986, 1990; Tremblay et al. 1998) have found that the responses of striatal cells often depend on behavioral context. In the rat, for example, Gardiner and Kitai (1992) report that cells in the dorsal striatum that responded to an auditory cue during a movement task usually did not respond to the same cue presented outside of the task, and some cells that responded during head movements during the task did not respond when rats made similar movements outside of the task. Carelli et al. (1997) report that in rats who have learned to barpress for food, dorsolateral striatal cells that responded to movement of the forelimb outside of the instrumental task were not active during lever pressing.

On the multiple T maze, neurons in the dorsal striatum responded as rats navigated the maze and during the delivery of food. Neither maze nor reward responses were described by general motor behavior. Less than one-third of reward-responsive PFNs responded at similar levels at both food-delivery sites, indicating that reward responses did not simply encode the action of chewing. Maze-responsive PFNs that responded at one location on the maze were not biased to respond similarly at other regions where rats took similar paths, indicating that maze responses did not simply encode motor activity during navigation. These results are consistent with the studies cited in the preceding text, which indicated that dorsal striatal neurons correlated with a movement or stimulus during a task are often not active during the same movement or stimulus presentation in a different behavioral context. Within a session, maze responses were well related to the spatial location of the animal. However, maze responses did not encode the spatial position of the animal independent of the animal's actions. First, maze responses were poorly correlated across sessions when rats took a different path through the same two-dimensional location in the environment. Second, maze responses were highly correlated across sessions when animals ran the same sequence of turns. Finally, maze responses were biased toward positive correlations across sessions when animals ran a different sequence of turns but took a similar path through the same two-dimensional location in the environment. These data indicate that maze-responsive cells were modulated by the location of the animal, what the animal was doing at that location, and, to some extent, by the specific sequence of actions the rat was performing.

This type of striatal sequence-specificity is consistent with the work done in primates and rats. In primates, Kermadi and

colleagues (Kermadi and Joseph 1995; Kermadi et al. 1993) have shown that striatal neurons in primates preferred specific visuomotor sequences. In rats, Aldridge and Berridge (1998) have shown that dorsal striatal neurons that were active during sequenced grooming were often not active during similar movements occurring outside of grooming sequences. To our knowledge, the data presented here from the multiple T task are the first evidence for sequence-specific striatal activity in rodents performing an arbitrary sequencing task.

Maze responses were also uniformly distributed over the turn sequence on the maze. If maze responses encode what actions need to be performed at a particular location/sensory context, then a uniform distribution of maze responses indicates that the striatal representation is rich enough to specify an action to perform at any point of the task. Such a result has important implications for theories of striatal function. Recent proposals have suggested that the striatum may implement a temporal-difference reinforcement-learning algorithm (Barto 1995; Sutton and Barto 1998) in which striatal neurons select an action to perform based on a policy that is modified to maximize the receipt of reward over time (Daw 2003; Doya 2000; Houk et al. 1995; Schultz et al. 1997; Suri and Schultz 1999; Suri et al. 2001).

Segregation of maze and reward responses

Maze-responsive PFNs often responded at multiple locations on the maze- and reward-responsive PFNs often responded after arrival at both food-delivery sites. However, maze-responsive PFNs did not respond after arrival at either food-delivery site, and reward-responsive PFNs did not respond while rats were running on the maze. A segregation of maze- and reward-responsive PFNs implies a segregation of information processing in the striatum and brings up two questions: what is the functional consequence of segregation and what properties of the striatum produce segregation?

FUNCTIONAL SIGNIFICANCE OF SEGREGATION. A segregation of maze and reward responses may shed light on the computational functions of the striatum. Recent proposals of basal ganglia function suggest that the striatum is involved in selecting appropriate actions in a task by implementing a reinforcement-learning algorithm (Barto 1995; Brown and Sharp 1995; Daw 2003; Daw and Touretzky 2000; Doya 1999, 2000; Foster et al. 2000; Houk et al. 1995; Montague et al. 1996; Schultz et al. 1997; Sutton and Barto 1998). In reinforcement-learning models of the striatum, the nigrostriatal dopaminergic system provides a reward-prediction error signal and the striatum implements an actor-critic architecture. The actor is responsible for selecting which action would be appropriate given the current sensory input, while the critic uses the reward-prediction error signal to change the value of sensory inputs so that the most advantageous action will be chosen. Segregation of maze- and reward-responsive PFNs may then reflect the separation of actor and critic components in the striatum.

POSSIBLE MECHANISM OF SEGREGATION. One possible mechanism of segregation of maze- and reward-responsive PFNs is related to striatal subcompartments. On the basis of markers such as μ-opiate receptor binding, the striatum can be divided into μ-opiate rich *striosomes* (also termed *patch*) (Gerfen 1985) and μ-opiate poor *matrix*, which is rich in acetycholinesterase (Gerfen 1985; Graybiel and Ragsdale 1978; Herkenham and Pert 1981). Matrix receives inputs from sensorimotor cortex and projects to the substantia nigra pars reticulata and palladial output nuclei (Gerfen 1984, 1989; Kawaguchi et al. 1990; Ragsdale and Graybiel 1984). Striosomes receive input from "limbic" cortex (including infralimbic, prelimbic, and anterior cingulate cortex) and project to dopaminergic cells in the substantia nigra pars compacta (Gerfen 1984, 1989; Ragsdale and Graybiel 1984). With its inputs to the substantia nigra pars compacta, striatal patches are well suited to be involved in reward-related processing, while matrix is well suited to be involved in action (Graybiel 1998; Houk et al. 1995; Kimura 1995; White 1989). White and Hiroi (1998) have shown that electrodes placed in striosomes, but not matrix, will support self-stimulation in rats, supporting a relationship between striosomes and reward. Trytek et al. (1996) have shown that motor related neurons tended to be located in the matrix, supporting a relationship between matrix and action. It could be that segregation of maze- and reward-responsive PFNs on the multiple T maze reflects an anatomical segregation of maze-responsive PFNs to the striatal matrix and reward-responsive PFNs to striosomes. The anatomical distributions of maze- and reward-responsive PFNs in the dorsal striatum is an important question to be addressed in future experiments.

Phasic/tonic neuron separation

On the basis of extracellularly recorded spike trains, we categorized neurons in the rodent dorsal striatum as PFNs or TFNs by the use of a new measure: the proportion of time spent in long ISIs. Our results are consistent with work in the primate showing a separation between phasic and tonic neurons (Alexander and DeLong 1985; Kimura et al. 1990). Kimura et al. (1990) have also found that the phasic neurons are the striatal projection neurons, which leads us to believe that our PFNs are likely to be striatal projection neurons. Histological studies indicate that between 90% and 95% of striatal neurons are projection neurons (Bennett and Wilson 2000), whereas in this study, only 68% of extracellularly recorded cells were identified as PFNs. This discrepancy is consistent with other experiments using extracellular recordings, which report proportions of phasic striatal neurons ranging from 59% to 92% in primates (Alexander and DeLong 1985; Kimura et al. 1990; Ueda and Kimura 2003). The variability in the proportions of phasic and tonic neurons observed likely results from sampling bias in extracellular recordings, which could be produced by a very low firing rate in some projection neurons and the large size of some interneurons.

Wilson and colleagues (Bennett and Wilson 2000; Stern et al. 1998; Wickens and Wilson 1998; Wilson 1993) have demonstrated in anesthetized rats that striatal medium spiny GABAergic neurons (MSPs) have a bistable membrane potential. MSPs show nonlinear shifts between hyperpolarized "down" states and depolarized "up" states, which are separated by 15–30 mV (Bennett and Wilson 2000; Stern et al. 1998; Wickens and Wilson 1998; Wilson 1993). MSP activity is organized into bouts of firing at ~10 Hz when in up states separated by quiescent periods when in down states (Bennett and Wilson 2000; Stern et al. 1998; Wickens and Wilson 1998; Wilson 1993). The firing pattern observed for MSPs is markedly different from the firing patterns observed for major

striatal interneurons, such as the cholingergic interneurons, which fire tonically at 3–10 Hz (Aosaki et al. 1995), and parvalbumin-immunoreactive GABAergic interneurons, which are thought to be fast-spiking striatal neurons (Kawaguchi et al. 1995).

The observed firing patterns of the major categories of striatal neurons suggest that our PFNs, which spent a high proportion of their time in long ISIs (i.e., quiescent periods) and fired in bouts at ~10 Hz, are likely to correspond to MSPs and imply that MSPs continue to exhibit bistable membrane potentials in the awake state. Kitano et al. (2002) also report evidence that MSPs continue to show bistable membrane potentials in the awake primate. The bistablity shown for PFNs presented here suggests that striatal MSPs in awake, behaving rats may show bistable membrane potentials.

Conclusions

A better understanding of how neurons in the dorsal striatum respond during complex tasks, such as the sequential navigation task we have used, may allow for a better understanding of the role of the dorsal striatum in the learning and expression of "habitual" types of learning and memory. An understanding of neural activity on these tasks will also allow the development of more realistic models of how the basal ganglia participate in learning and memory and may assist the implementation of striatal reinforcement-learning models. Many experimental paradigms that have been applied to the striatum have assumed that the basis for states that represent the world is temporal. The difference between our task and those that use a temporal task-structure may reveal how the striatal representation encoded by projection neurons depends on the demands of the task.

In our task, rats had to find the correct goal locations to obtain a food reward. In contrast, temporally structured tasks require animals to represent the temporal intervals between stimuli, such as instruction cues and stimuli predicting the availability of reward. Neural responses in these tasks are often well related to when a cue is delivered, when a reward is expected, or during the execution of a movement. Our results indicate that in a navigation task, striatal responses are well related to the location of an animal on the maze and to the delivery of food. Responses on the maze did not encode the animal's location nor the shape of the animal's path. Rather, the neural responses on the multiple T maze were best suited to encode *what/where* relationships: what action is required in the current location or sensory context. These maze responses covered space uniformly, indicating that at any point on the maze striatal activity could be used to guide the selection of actions. Such a finding will have an important impact in the application of reinforcement-learning models to the striatum.

Task-responsive cells that responded on the maze and during food receipt were separate populations, indicating a divergence of information processing in the dorsal striatum. This segregation is likely to play an important role in the computational functions accomplished by the basal ganglia.

ACKNOWLEDGMENTS

We thank P. Aia, M. Arudi, D. Bang, D. Bernal, C. Boldt, J. Jackson, A. Johnson, K. Seeland, S. Rao, T. Singewald, and J. Waataja for help running experiments as well as for helpful discussions. We are particularly indebted to D. Bang for technical work on the histology. We also thank M. Reidl for help with histology.

GRANTS

N. C. Schmitzer-Torbert was supported by National Science Foundation-IGERT Training Grant 9870633 and an NSF Graduate Research Fellowship. This work was supported by the University of Minnesota.

REFERENCES

Aldridge JW and Berridge KC. Coding of serial order by neostriatal neurons: a "natural action" approach to movement sequence. *J Neurosci* 18: 2777–2787, 1998.

Alexander GE and DeLong MR. Microstimulation of the primate neostriatum. II. Somatotopic organization of striatal microexcitable zones and their relation to neuronal response properties. *J Neurophysiol* 53: 1417–1430, 1985.

Aosaki T, Kimura M, and Graybiel AM. Temporal and spatial characteristics of tonically active neurons of the primate's striatum. *J Neurophysiol* 73: 1234–1252, 1995.

Barto AG. Adaptive critics and the basal ganglia. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG. Cambridge MA: MIT Press, 1995, p. 215–232.

Bennett BD and Wilson CJ. Synaptology and physiology of neostriatal neurons. In: *Brain Dynamics and the Striatal Complex*, edited by Miller R and Wickens J. Amsterdam: Harwood Academic, 2000, p. 111–140.

Berridge KC and Whishaw IQ. Cortex, striatum and cerebellum: control of serial order in a grooming sequence. *Exp Brain Res* 90: 275–290, 1992.

Brown MA and Sharp PE. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* 5: 171–188, 1995.

Carelli RM, Wolske M, and West MO. Loss of lever press-related firing of rat striatal forelimb neurons after repeated sessions in a lever pressing task. *J Neurosci* 17: 1804–1814, 1997.

Cohen NJ and Squire LR. Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science* 210: 207–210, 1980.

Cromwell HC and Berridge KC. Implementation of action sequences by a neostriatal site: a lesion mapping study of grooming syntax. *J Neurosci* 16: 3444–3458, 1996.

D'Agostino RB and Stephens MA (editors). *Goodness-of-Fit Techniques.* New York: Marcel Dekker, 1986.

Daw ND. *Reinforcement Learning Models of the Dopamine System and Their Behavioral Implications* (PhD thesis). Pittsburgh, PA: Carnegie Mellon University, 2003.

Daw ND and Touretzky DS. Behavioral considerations suggest an average reward TD model of the dopamine system. *Neurocomputing* 32–33: 679–684, 2000.

Doya K. What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? *Neural Networks* 12: 961–974, 1999.

Doya K. Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol* 10: 732–739, 2000.

Doyon J, Laforce R, Bouchard G, Gaudreau D, Roy J, Poirer M, Bedard PJ, Bedard F, and Bouchard JP. Role of the striatum, cerebellum and frontal lobes in the automatization of a repeated visuomotor sequence of movements. *Neuropsychologia* 36: 625–641, 1998.

Efron B. *The Jackknife, the Bootstrap, and Other Resampling Plans. Regional Conference Series in Applied Mathematics.* Philadelphia PA: Society for Industrial and Applied Mathematics, 1982, vol 38.

Foster DJ, Morris RGM, and Dayan P. A model of hippocampally dependent navigation using the temporal difference learning rule. *Hippocampus* 10: 1–6, 2000.

Gardiner TW and Kitai ST. Single-unit activity in the globus pallidus and neostriatum of the rat during performance of a trained head-movement. *Exp Brain Res* 88: 517–530, 1992.

Gerfen CR. The neostriatal mosaic: compartmentalization of corticostriatal input and striatonigral output systems. *Nature* 311: 461–464, 1984.

Gerfen CR. The neostriatal mosaic. I. Compartmental organization of projections from the striatum to the substantia nigra in the rat. *J Comp Neurol* 236: 454–476, 1985.

Gerfen CR. The neostriatal mosaic: striatal patch-matrix organization is related to cortical lamination. *Science* 246: 385–388, 1989.

Graybiel AM. The basal ganglia and chunking of action repertoires. *Neurobiol Learn Mem* 70: 119–136, 1998.

Graybiel AM, Aosaki T, Flaherty AW, and Kimura M. The basal ganglia and adaptive motor control. *Science* 265: 1826–1831, 1994.

Graybiel AM and Ragsdale CW Jr. Histochemically distinct compartments in the striatum of human, monkey, and cat demonstrated by acetylcholinesterase staining. *Proc Natl Acad Sci USA* 75: 5723–5726, 1978.

Greenhouse SW and Geisser S. On methods in the analysis of profile data. *Psychometrika* 24: 95–112, 1959.

Harris KD. KlustaKwik.Automatic Cluster Analysis, version 1.0. *http://osiris.rutgers.edu/Buzsaki/software/,* 2002.

Herkenham M and Pert CB. Mosaic distribution of opiate receptors, parafascicular projections and acetylcholinesterase in rat striatum. *Nature* 291: 415–418, 1981.

Houk JC, Adams JL, and Barto AG. A model of how the basal ganglia generate and use neural signals that predict reinforcement. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG. Cambridge MA: MIT Press, 1995, p. 249–270.

Jackson JC, Schmitzer-Torbert NC, Harris KD, and Redish AD. Quantitative assessment of extracellular multichannel recording quality using measures of cluster separation. *Soc Neurosci Abstr* 518.18, 2003.

Jog MS, Kubota Y, Connolly CI, Hillegaart V, and Graybiel AM. Building neural representations of habits. *Science* 286: 1746–1749, 1999.

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

Kawaguchi Y, Wilson CJ, and Emson PC. Projection subtypes of rat neostriatal matrix cells revealed by intracellular injection of biocytin. *J Neurosci* 10: 3421–3438, 1990.

Kermadi I and Joseph JP. Activity in the caudate nucleus of monkey during spatial sequencing. *J Neurophysiol* 74: 911–933, 1995.

Kermadi I, Jurquet Y, Arzi M, and Joseph J. Neural activity in the caudate nucleus of monkeys during spatial sequencing. *Exp Brain Res* 94: 352–356, 1993.

Kesner RP, Bolland BL, and Dakis M. Memory for spatial locations, motor responses, and objects: triple dissociation among the hippocampus, caudate nucleus, and extrastriate visual cortex. *Exp Brain Res* 93: 462–470, 1993.

Kimura M. The role of primate putamen neurons in the association of sensory stimuli with movement. *Neurosci Res* 3: 436–443, 1986.

Kimura M. Behaviorally contingent property of movement-related activity of the primate putamen. *J Neurophysiol* 63: 1277–1296, 1990.

Kimura M. Role of basal ganglia in behavioral learning. *Neurosci Res* 22: 353–358, 1995.

Kimura M, Kato M, and Shimazaki H. Physiological properties of projection neurons in the monkey striatum to globus pallidus. *Exp Brain Res* 82: 672–676, 1990.

Kitano K, Câteau H, Kaneda K, Nambu A, Takada M, and Fukai T. Two-state membrane potential transitions of striatal spiny neurons as evidenced by numerical simulations and electrophysiological recordings in awake monkeys. *J Neurosci* 22: 1–6, 2002.

Knowlton BJ, Mangels JA, and Squire LR. A neostriatal habit learning system in humans. *Science* 273: 1399–1402, 1996.

Matsumoto N, Hanakawa T, Maki S, Graybiel AM, and Kimura M. Nigrostriatal dopamine system in learning to perform sequential motor tasks in a predictive manner. *J Neurophysiol* 82: 978–998, 1999.

McDonald RJ and White N. Parallel information processing by hippocampal and dorsal striatal memory systems in the Morris water maze. *Soc Neurosci Abstr* 19:362, 1993.

Mishkin M and Appenzeller T. The anatomy of memory. *Sci Am* 256: 80–89, 1987.

Miyachi S, Hikosaka O, Miyashita K, Kárádi Z, and Rand MK. Differential roles of monkey striatum in learning of sequential hand movement. *Exp Brain Res* 115: 1–5, 1997.

Montague PR, Dayan P, and Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *J Neurosci* 16: 1936–1947, 1996.

Packard MG. Glutamate infused posttraining into the hippocampus or caudate-putamen differentially strengthens place and response learning. *Proc Natl Acad Sci USA* 96: 12881–12886, 1999.

Packard MG, Hirsh R, and White NM. Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. *J Neurosci* 9: 1465–1472, 1989.

Packard MG and Knowlton BJ. Learning and memory functions of the basal ganglia. *Annu Rev Neurosci* 25: 563–593, 2002.

Packard MG and McGaugh JL. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behav Neurosci* 106: 439–446, 1992.

Paxinos G and Watson C. *The Rat Brain in Stereotaxic Coordinates* (4th ed.). New York: Academic, 1998.

Ragozzino KE, Leutgeb S, and Mizumori SJ. Dorsal striatal head direction and hippocampal place representations during spatial navigation. *Exp Brain Res* 139: 372–376, 2001.

Ragsdale CW and Graybiel AM. A simple ordering of neocortical areas established by the compartmental organization of their striatal projections. *Proc Natl Acad Sci USA* 87: 6196–6199, 1984.

Redish AD, Schmitzer-Torbert NC, and Jackson JC. Classification of dorsal striatal neurons from extracellular recordings in awake behaving rats. *Soc Neurosci Abstr* 676.3, 2002a.

Redish AD, Schmitzer-Torbert NC. MClust Spike Sorting Toolbox, version 3.0. *http://www.cbc.umn.edu/~redish/mclust,* 2002b.

Schmitzer-Torbert NC, Jackson JC, and Redish AD. Behavioral correlates of neuronal activity in the rodent dorsal striatum: the Multiple-T task. *Soc Neurosci Abstr* 676.4, 2002.

Schmitzer-Torbert NC and Redish AD. Development of path stereotypy in a single day in rats on a multiple-T maze. *Arch Ital Biol* 140: 295–301, 2002.

Schultz W, Apicella P, Romo R, and Scarnati E. Context-dependent activity in primate striatum reflecting past and future behavioral events. In: *Models of Information Processing in the Basal Ganglia*, edited by Houk JC, Davis JL, and Beiser DG. Cambridge MA: MIT Press, 1995, p. 11–27.

Schultz W, Dayan P, and Montague R. A neural substrate of prediction and reward. *Science* 275: 1593–1599, 1997.

Sokal RR and Rohlf FJ. *Biometry: The Principles and Practice of Statistics in Biological Research* (3rd ed.). New York: Freeman, 1995.

Stern EA, Jaeger D, and Wilson CJ. Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. *Nature* 394: 475–478, 1998.

Suri RE, Bargas J, and Arbib MA. Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience* 103: 65–85, 2001.

Suri RE and Schultz W. A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. *Neuroscience* 91: 871–890, 1999.

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

Tremblay L, Hollerman JR, and Schultz W. Modifications of reward expectation-related neuronal activity during learning in primate striatum. *J Neurophysiol* 80: 964–977, 1998.

Trytek ES, White IM, Schroeder DM, Heidenreich BA, and Rebec GV. Localization of motor- and nonmotor-related neurons within the matrix striosome organization of rat striatum. *Brain Res* 707: 221–227, 1996.

Ueda Y and Kimura M. Encoding of direction and combination of movements by primate putamen neurons. *Eur J Neurosci* 18: 980–994, 2003.

White NM. A functional hypothesis concerning the striatal matrix and patches: mediation of S-R memory and reward. *Life Sci* 45: 1943–1957, 1989.

White NM and Hiroi N. Preferential localization of self-stimulation sites in striosomes/patches in the rat striatum. *Proc Natl Acad Sci USA* 95: 6486–6491, 1998.

Wickens JR and Wilson CJ. Regulation of action-potential firing in spiny neurons of the rat neostriatum in vivo. *J Neurophysiol* 79: 2358–2364, 1998.

Wiener SI. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *J Neurosci* 13: 3802–3817, 1993.

Wilson CJ. The generation of natural firing patterns in neostriatal neurons. In: *Progress in Brain Research. Chemical Signalling in the Basal Ganglia*, edited by Arbuthnott GW and Emson PC. New York: Elsevier, 1993, vol. 99, chapt. 18, p. 277–297.