

Triple Dissociation of Information Processing in Dorsal Striatum, Ventral Striatum, and Hippocampus on a Learned Spatial Decision Task

Matthijs A.A. van der Meer,¹ Adam Johnson,^{1,2} Neil C. Schmitzer-Torbert,^{1,3} and A. David Redish^{1,*}

¹Department of Neuroscience, University of Minnesota, Minneapolis, MN 55455, USA

²Department of Psychology, Bethel University, St. Paul, MN 55112, USA

³Department of Psychology, Wabash College, Crawfordsville, IN 47933, USA

*Correspondence: redish@umn.edu

DOI 10.1016/j.neuron.2010.06.023

SUMMARY

Decision-making studies across different domains suggest that decisions can arise from multiple, parallel systems in the brain: a flexible system utilizing action-outcome expectancies and a more rigid system based on situation-action associations. The hippocampus, ventral striatum, and dorsal striatum make unique contributions to each system, but how information processing in each of these structures supports these systems is unknown. Recent work has shown covert representations of future paths in hippocampus and of future rewards in ventral striatum. We developed analyses in order to use a comparative methodology and apply the same analyses to all three structures. Covert representations of future paths and reward were both absent from the dorsal striatum. In contrast, dorsal striatum slowly developed situation representations that selectively represented action-rich parts of the task. This triple dissociation suggests that the different roles these structures play are due to differences in information-processing mechanisms.

INTRODUCTION

A key insight from decision-making studies across different domains is that decisions can arise from multiple, parallel systems in the brain (O'Keefe and Nadel, 1978; Schacter and Tulving, 1994; Poldrack and Packard, 2003; Daw et al., 2005; Redish et al., 2008). One system, broadly characterizable as "model-based," relies on internally generated expectations of action outcomes, while the other, "model-free" system uses learned (cached) values of situation-action associations. This distinction between different decision-making systems has been demonstrated behaviorally (e.g., stimulus-response versus response-outcome learning [Balleine and Dickinson, 1998], as well as response learning versus place learning [Packard and McGaugh, 1996]), has been articulated computationally (Daw et al., 2005; Niv et al., 2006), and maps onto dissociable brain

structures (Packard and McGaugh, 1996; Yin et al., 2004). In rodent navigation studies, lesion and inactivation studies have shown that the model-based system (as engaged by place navigation) depends on hippocampal integrity, while the model-free system (as engaged by response navigation) depends on dorsal striatal integrity (Packard and McGaugh, 1996); ventral striatum may play a role in both systems (Atallah et al., 2007).

Computationally, the model-based system is thought to rely on world knowledge in order to generate specific expectations about the outcomes of actions, which may range from anticipating the outcome of a simple lever press to mental simulation or planning over extended spatial maps or Tower of London puzzles (Shallice, 1982). While this process may be computationally expensive, it allows for adaptive behavior in novel situations and under changing goals. In contrast, a typical model-free system associates actions with values, reflecting how well each action has turned out in the past. This system is efficient but also inflexible because cached action values reflect past experience rather than current goals (Daw et al., 2005; Niv et al., 2006; Redish et al., 2008). Thus, computational theories of decision making have suggested potential information processing differences that capture the behavioral and anatomical distinctions between model-based and model-free decision-making systems. However, in order to reveal the mechanisms actually used in the brain to *specifically* support these different decision-making algorithms, it is necessary to compare neural activity between structures on a task that engages both systems.

The multiple-T task is a spatial decision task that engages different decision-making strategies (Schmitzer-Torbert and Redish, 2002). On this task, Johnson and Redish (2007) found that ensembles of hippocampal neurons transiently represented locations ahead of the animal, sweeping down one arm of the maze, then another, before the animal made its choice. Such "lookahead" operations are a critical element of model-based decision making. However, given that dorsal striatum can represent locations as well (Wiener, 1993; Yeshenko et al., 2004), an important question is whether this property is in fact unique to the hippocampus. Similarly, slow changes in dorsal striatal firing patterns (Barnes et al., 2005) demonstrate reorganization that could support gradual model-free learning. However, slow changes have also been observed in the hippocampus (Mehta et al., 1997; Lever et al., 2002), so in the absence of direct comparison it is not clear if such effects are specific to how

dorsal striatum operates. Finally, van der Meer and Redish (2009) found ventral striatal firing patterns relevant to roles in both model-free and model-based decision making, such as anticipatory “ramping” and covert activation of reward-responsive neurons at decision points. However, it is not known if dorsal striatal neurons show ramping or reward activation at decision points.

Thus, in order to determine which of these information processing mechanisms are unique to these areas—a requirement if we are to understand the neural basis of their distinct behavioral roles—we compared the firing properties of dorsal striatal, ventral striatal, and hippocampal neurons on the multiple-T task. Because several of these analyses require large neural ensembles, we used three different groups of animals, one for each structure. The data sets used here include data used in previously published work (dorsal striatum: Schmitzer-Torbert and Redish [2004], ventral striatum: van der Meer and Redish [2009], hippocampus: Johnson and Redish [2007]). However, here we use a comparative approach applying the same, new analyses to each structure, allowing direct comparisons and the identification of a triple dissociation in information processing mechanisms.

RESULTS

We recorded 1646, 2323, and 1473 spike trains from 98, 96, and 31 recording sessions from dorsal striatum, ventral striatum, and hippocampus, respectively, as rats ($n = 5$ each for dorsal and ventral striatum, $n = 6$ for hippocampus) performed the multiple-T task (Figure 1A). On this task, three low-cost choice points (T1–T3) with dead ends on one side were followed by a high-cost choice (T4) between the left or right “return rail,” with only one side rewarded during any given session. Although rats were trained on the task prior to electrode implant surgery, both the rewarded side (left or right choice at T4) as well as the correct sequence of preceding turns (T1–T3) could be varied from day to day, such that the rats started out uncertain about the correct choices at the beginning of each session. Rats rapidly learned to choose the rewarded side, reaching asymptotic performance (>90%) within ten laps (Figure 1C) with each group improving at a comparable rate (Figure S1A). Coincident with this rapid performance increase, rats exhibited pausing behavior at the high-cost choice point (T4) during early laps, looking back and forth between left and right before making their choice (a hippocampus-dependent behavior known as *vicarious trial and error*, or VTE [Tolman, 1938; Hu and Amsel, 1995]). Pausing was absent at a control choice point (T2; Figures 1B and S1B). Following this initial VTE phase, choice performance reached asymptote, yet lap times continued to decrease (Figure 1C, inset), indicating a change in behavior beyond choice performance (Schmitzer-Torbert and Redish, 2002). These behavioral characteristics indicate the engagement of different decision-making strategies within single recording sessions.

Differential Coding of Task Structure in Dorsal Striatum, Ventral Striatum, and Hippocampus

As both striatum and hippocampus are known to contain different cell types (Ranck, 1973; Kawaguchi, 1993), we sepa-

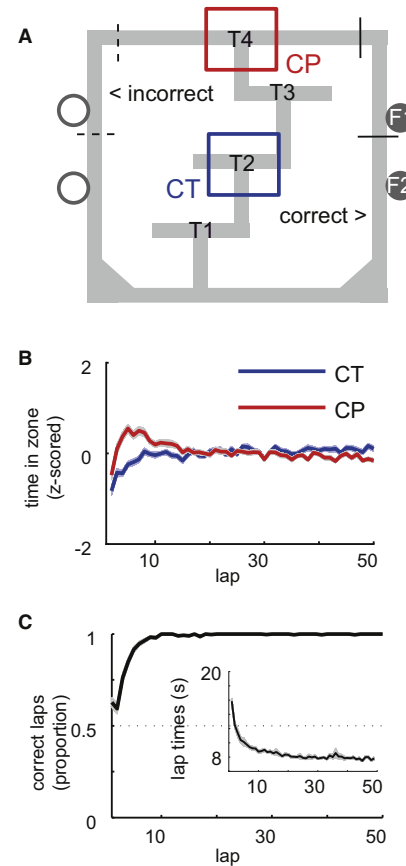


Figure 1. Behavior on the Multiple-T Task

(A) Diagram of a single Multiple-T configuration (“RRLR”). T1–T4 indicate turns, with T4 the final, high-cost choice point. Food reward is delivered at the reward sites (F1 and F2) when the rat crosses the active feeder trigger lines (on the right side for this session, indicated by black lines in the diagram but not present on the actual track used). Which side was rewarded (left or right choice at T4) as well as the correct sequence of preceding turns (T1–T3) could be varied from day to day. (B) Rats paused specifically at the final choice point (CP/T4) during early laps. Plot shows normalized time spent in the two zones indicated in (A); note the increase in time spent at the choice point (CP, red line) absent from the control point (CT, blue line). Time in zones was normalized by z-scoring the time spent in each zone on each lap relative to the distribution of times spent over laps within each session. (C) Behavioral choice performance. The proportion of correct choices over laps, averaged across sessions, increased to >90% within ten laps. (Inset) Lap times continued to decrease even when choice performance was stable. Times spent at the feeder sites were excluded in computing lap times, such that the decrease in lap times was not due to a change in immobility at the feeder sites. Shaded area is S.E.M. over sessions.

rated putative projection neurons and interneurons based on firing statistics (Barnes et al., 2005; Schmitzer-Torbert and Redish, 2008). Consistent with previous reports, putative hippocampal pyramidal neurons tended to show spatially focused firing fields (“place fields” [O’Keefe and Dostrovsky 1971]), while phasically firing neurons (PFNs; putative medium spiny neurons) in both striatal subregions exhibited a wider range of firing correlates, including maze-related activity and responsiveness to reward (Lavoie and Mizumori 1994; Schmitzer-Torbert and

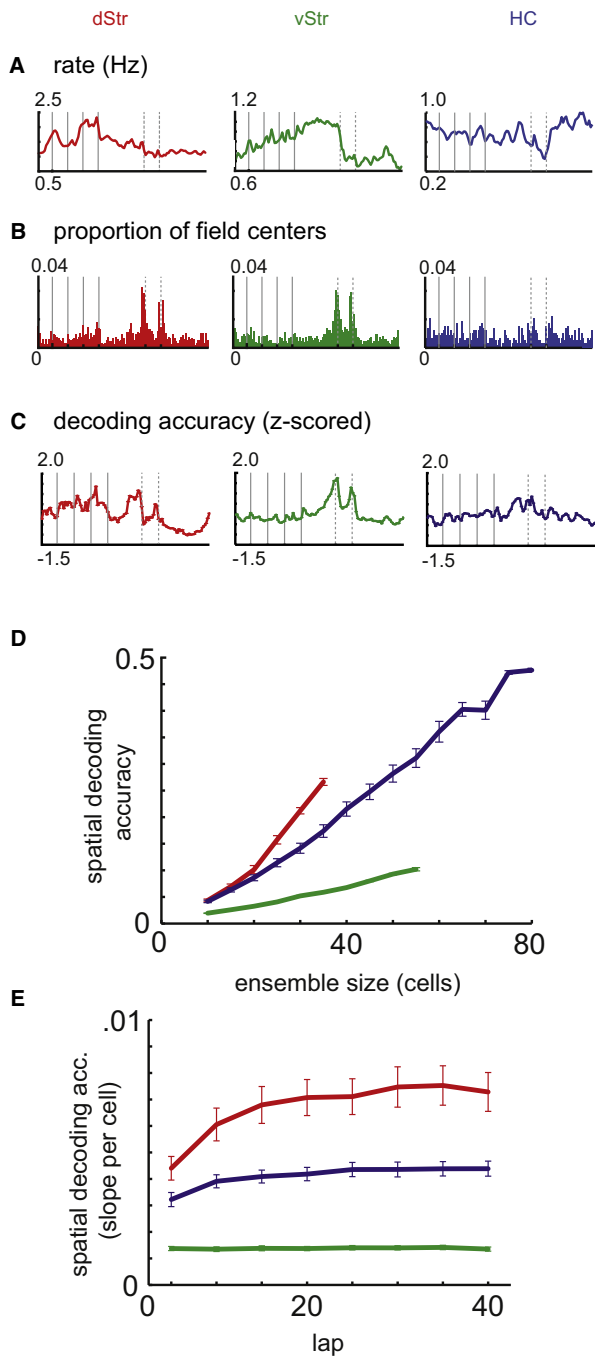


Figure 2. Differences in Neural Coding between Dorsal Striatum (Left Column, Red), Ventral Striatum (Middle Column, Green), and Hippocampus (Right Column, Blue)

(A) Average firing rate of putative projection neurons (MSNs in striatum, pyramidal neurons in hippocampus) by location on the track (solid lines, T1–T4 choice points; dashed lines, reward sites). Note the difference in spatial firing distributions, with dorsal striatum most active on the sequence of turns (S–T4), ventral striatum showing a gradual ramping up, and hippocampus relatively uniform apart from a sharp dip at the reward locations. (B) Frequency histograms of the spatial location of peak firing rates. Both striatal subregions show a clear peak at the reward sites absent in hippocampus, reflecting the presence of reward-responsive cell populations. (C) Spatial decoding accu-

redish 2004; Barnes et al., 2005; Berke et al., 2009). In order to examine differences between the three structures at the population level, we plotted the average firing rates for putative pyramidal neurons or PFNs in each of the three structures over the track (Figure 2A; interneurons, Figure S2B). Dorsal striatal PFNs were most active on the sequence of turns (S–T4), especially between T3 and T4, and least active on the bottom return rail (F2–S). Ventral striatal PFNs showed a ramping up of activity through the turn sequence, dropping off sharply at the first reward site. Hippocampal pyramidal neuron firing rates were relatively uniform over the track (Levene’s test for uniformity, see Table S1), although a decrease between the two reward sites was visible, which may reflect effects of low running speed (pyramidal cells in hippocampus are sensitive to running speed [McNaughton et al., 1983]). Because increases in population firing rate at a given location can result from (1) an increase in the number of cells that have fields there and (2) increased firing rates at that location, we also plotted the distribution of peak firing locations (Figure 2B). Both striatal subregions showed a clear increase in the number of active cells at the reward sites, reflecting a population of reward-responsive neurons absent from hippocampus. In dorsal striatum, a decline in the number of firing fields was visible after the navigation sequence, while hippocampal firing fields were more uniformly distributed (χ^2 test, see Table S1). These characteristics support a distinction in information processing in which dorsal striatum emphasized the turn sequence (consistent with situation-action encoding), ventral striatum showed a ramp (consistent with representation of motivationally relevant information), and hippocampus represented the track relatively uniformly (consistent with a spatial, map-like representation).

The above comparison suggests underlying differences in neural coding but does not reveal how informative these codes are. To address this, we measured the extent to which neural ensembles in the three structures contained information about location on the track. A Bayesian ensemble decoding algorithm was applied, which computes a probability distribution over the track given the numbers of spikes fired by each neuron within a given time window (Zhang et al., 1998; Figure S3). The average probability (over all time windows) at the rat’s actual location was used as a measure of decoding accuracy: it indicates how good the ensemble is at representing the rat’s actual location. Given the spatial modulation of firing rates on the track, we

racy, z-scored within each session to account for overall accuracy differences (see panel D), plotted over position on the maze. Dorsal striatal decoding accuracy was best on the section of the navigation sequence that contained choice points (T1–T4), as well as the reward sites, but poor on the (horizontal) top and bottom segments. In contrast, ventral striatal decoding accuracy was relatively good during approach of the reward sites, and hippocampus showed more spatially uniform decoding accuracy. (D) Between structures, spatial decoding accuracy changed differentially as a function of the number of cells in the ensemble. Dorsal striatal ensembles were the most efficient (steepest slope), while in ventral striatum, spatial decoding accuracy increased the least with ensemble size (shallowest slope). (E) With learning, spatial decoding efficiency (the slope of accuracy as a function of ensemble size, panel D) changed differentially between the structures. Ventral striatum did not show a change in efficiency, while dorsal striatum showed the largest increase. Error bars show SEM over sessions.

asked whether, at the ensemble level, certain parts of the track could be decoded more accurately than others. To account for differences in decoding accuracy between recording sessions and structures (explored in the next section), we normalized the decoding accuracy distribution over the track (Figure 2C). Dorsal striatal decoding accuracy was best on the sequence of choice points (T1–T4), as well as the reward sites, but poor on the return segment (F2–T1). For hippocampus, the relatively uniform decoding accuracy was in agreement with its spatial firing rate distribution (Levene's test for uniformity, see Table S3). In contrast, ventral striatal decoding accuracy did not change over the sequence of turns as the firing rate distribution did. Thus, on the ensemble level, dorsal striatal decoding accuracy focused on the turn sequence on the track as well as the reward sites, while hippocampal decoding accuracy was most uniform, and ventral striatum highlighted the reward sites only.

Dorsal Striatum, but Not Ventral Striatum, Shows a Strong Increase in Coding Efficiency within Sessions

The preceding analysis normalized differences in decoding accuracy between sessions and structures. However, such differences can be informative when comparing neural coding properties between structures: we would like to ask “how well” each structure represents location on the track. In distributed representations, decoding accuracy depends on ensemble size (Zhang et al., 1998; Stanley et al., 1999; Wessberg et al., 2000). Thus, comparing decoding accuracy between sessions or structures with different ensemble sizes centers on the extent to which accuracy increases as a function of ensemble size (“coding efficiency”). We therefore used a neuron dropping procedure (Wessberg et al., 2000; Narayanan et al., 2005) to sample random subsets of ensembles (see Experimental Procedures) in order to plot overall decoding accuracy as a function of number of cells for the three structures (Figure 2D). Between structures, spatial decoding accuracy changed differentially as a function of the number of cells in the ensemble (two-factor ANOVA, structure \times ensemble size interaction, $F_{(2,1)} = 200.05$; $p < 10^{-10}$). Dorsal striatal ensembles were the most efficient (steepest slope; two-factor ANOVA for dorsal striatum and hippocampus, structure \times ensemble size interaction, $F_{(1,1)} = 7.72$, $p = 0.0058$) while in ventral striatum, spatial decoding accuracy increased the least with ensemble size, and hippocampus fell in between.

Previous studies have found slow changes (across days) in the distribution of dorsal striatal firing rates on a different T-maze task (Barnes et al., 2005). While within single sessions, we did not find evidence for systematic changes in dorsal striatal firing rates (Figure S2A), spatial decoding accuracy can vary independently of firing rate (compare Figures 2A and 2C) raising the possibility of reorganization with experience at the ensemble level. Thus, we asked how spatial coding efficiency changed over laps for the three structures. Decoding efficiency changed differentially between the three structures (Figure 2E; overall structure \times lap interaction, $F_{(2,1)} = 11.84$, $p < 10^{-10}$): in ventral striatum, there was no evidence of a change over laps, while dorsal striatum showed the strongest increase. This analysis relies on accurate estimation of the slope of decoding accuracy as a function of the number of cells (verified in Figure S2C). To

control for the possibility that behavioral differences between the groups of animals influenced these results, we used a multiple regression analysis to identify behavioral variables that explained a significant amount of variance in decoding efficiency and subtracted the best fits based on these variables from the data (running speed, distance from an idealized path through the maze, and proportion of correct choices; see Figure S2B); this did not affect the pattern of results. Thus, with experience, dorsal striatum showed the strongest increase in coding efficiency, while hippocampus showed a modest amount, and ventral striatum showed none. These results suggest the presence of a dynamic reorganization process in dorsal striatum that comes to reflect the structure of the task with experience (Nakahara et al., 2002; Barnes et al., 2005; Schmitzer-Torbert and Redish, 2008; Berke et al., 2009).

Hippocampus and Ventral Striatum, but Not Dorsal Striatum, Show Forward Representations at the Choice Point

Decoding provides access to representational content, allowing analysis of not just *how much* information a given ensemble contains but also of what that information actually is (Schneidman et al., 2003). Johnson and Redish (2007) found that as rats paused at the final choice point, hippocampal representations of space swept ahead of the animal, down one arm of the maze and then the other, before the rat made its choice. It is presently not known if other areas in which spatial information is present, such as dorsal striatum, exhibit a similar effect. To investigate this, we applied the decoding algorithm to data from all three structures, using a 20 ms time window. Note that unlike the analysis in Johnson and Redish (2007), this method is “memoryless,” treating each time window as independent. For all passes through the final choice point, the proportion of the decoded probability distribution that fell either ahead or behind the choice point (Figure 3A) was plotted as a function of lap. While for all three structures the decoding probability to the choice point itself increased over the first ten laps (data not shown), for dorsal striatum (top panel) decoding to both the behind and ahead zones was marginally increased in this same period, indicating a nonspecific improvement in decoding accuracy. In contrast, ventral striatum and hippocampus showed a different pattern, where decoding ahead of, but not behind, the animal was high during early laps (two-factor ANOVA, lap-decoding zone interaction, smallest $F_{(1,1)} = 4.69$, $p = 0.031$). This is not compatible with a nonspecific decoding improvement: instead, it suggests that during early laps there is increased representation of locations ahead of the animal. We did not find evidence for such events in dorsal striatum, neither when averaged across sessions (two-factor ANOVA, lap-decoding zone interaction, $F_{(1,1)} = 1.78$, $p = 0.18$) nor upon visual inspection of decoding during individual passes through the choice point. Thus, even though dorsal striatal position encoding is at least as good as that in hippocampus on this task (Figure 2D), it did not selectively represent locations ahead of the animal at the choice point.

To assess whether look-ahead in hippocampus (and ventral striatum) was specific to the final choice point, we repeated the same analysis for passes through turn 2, a low-cost choice point (Figure 3B). At this point, decoding ahead of the animal




Figure 3. Decoding Ahead of the Animal Is Specific to Hippocampus, Ventral Striatum, and the High-Cost Choice Point

(A) For all passes through the final choice point (T4, black rectangle), the probability of decoding ahead of the animal (“forward,” pink rectangle) and behind the animal (“backward,” yellow rectangle) was computed. Dorsal striatum (dStr) showed a small and nonspecific increase in decoding accuracy: backward and forward both decreased equally. In contrast, ventral striatum and hippocampus showed an asymmetric pattern where, initially, decoding ahead of the animal was high compared to decoding behind the animal. (B) This effect was specific to the final choice point; compare the same analysis applied to a control point (T2). Note also that while ventral striatum showed a lookahead effect, its overall decoding accuracy was much lower than that in hippocampus or dorsal striatum (see Figure 2A). Error bars show SEM over sessions.

during early laps was much diminished, no longer reaching significance for either hippocampus or ventral striatum (two-factor ANOVA, lap-decoding zone interaction, largest $F_{(1,1)} = 0.25$, $p = 0.62$). Thus, lookahead occurred specifically at the final choice point, further supporting the notion that such processes are not permanently-on epiphenomena but can be dynamically engaged depending on task demands.

Covert Representation of Reward in Ventral Striatum, but Not Dorsal Striatum

van der Meer and Redish (2009) showed that ventral striatal reward-responsive neurons tended to be active at the choice point during early laps, suggesting covert expectation of reward congruent with model-based decision making. It is not known if dorsal striatal representations of reward also show this effect. To address this, we applied the same analysis to dorsal striatum, plotting the average (z-scored) firing rate of reward-responsive




Figure 4. Ventral Striatal, but Not Dorsal Striatal, Reward-Responsive Cells Show an Increase in Firing Rate at the Final Choice Point during Early Laps

Each panel shows average z-scored firing rates (over the spatial extent of the track) for reward-responsive (blue) and non-reward-responsive (red) cells. As reported in van der Meer and Redish (2009), during early laps (1–10), ventral striatal reward cells showed elevated activity at the final choice point (T4) compared to non-reward-responsive cells (A). No such difference was apparent in dorsal striatum (B) or in either structure during later laps (C and D). Plots were obtained by z-scoring each cell’s tuning curve using the mean and standard deviation of firing rates over the full tuning curve bins spanning the start of maze (S) to one-third between T4 and F1.

cells and non-reward-responsive cells over the maze (Figure 4). Note that, because this analysis is designed to address firing rates of reward and non-reward cells on the track in the absence of reward, the normalization and analysis was restricted to firing rates on the part of the track between the turn sequence start (S) to past the final choice point (T4). For completeness, we have included this analysis for hippocampal regions in Figure S4B, even though these did not show the characteristic reward response of the striatal regions (Figure 2B).

As reported in van der Meer and Redish (2009) for ventral striatum, a two-way ANOVA with location on the maze (nine bins, from the start of the first T to one-third of the way between T4 and F1) and cell type (reward or non-reward) as factors showed a significant interaction for early laps (1–10, $F_{(8,1)} = 4.0$, $p < 0.001$) with the reward cells having significantly higher activity in the T4 bin ($F_{(1,1)} = 12.56$, $p < 0.001$; see Figure S4A for full firing rate distributions and additional statistics). During late laps (31–40) there was no such difference ($F_{(1,1)} = 0.22$, $p = 0.64$). Thus, ventral striatal reward neurons showed a relative increase in activity specifically at the final choice point during early laps. In contrast, dorsal striatal reward neurons showed no difference in activity at T4 during early laps ($F_{(1,1)} = 2.47$, $p = 0.116$); for late laps, there was a difference ($F_{(1,1)} = 5.77$, $p = 0.0163$), but the non-reward cells were the more active group. Thus, even though we found similarly prominent reward-responsive activity in dorsal striatum compared to ventral striatum, only ventral striatal reward cells showed covert representation of reward at the choice point during early laps.

DISCUSSION

Hippocampus, dorsal striatum, and ventral striatum processed information differently on this task, consistent with their different roles in decision making.

Hippocampus as a Cognitive Map with Online Search

Hippocampal “place cells” are classically thought of as providing a cognitive map that supports flexible route planning (O’Keefe and Nadel, 1978). Such a map is an example of a world model that could be used for internal generation of potential outcomes during decision making. In support of this idea, Johnson and Redish (2007) found that hippocampal ensembles transiently represented possible outcomes at the final choice point of the multiple-T task. We extend this result in several important ways. First, we found that this “lookahead” is not a permanently enabled property of hippocampus, as would be expected from effects like theta phase precession (Maurer and McNaughton, 2007). Instead, lookahead was specific to the final choice point (T4) and absent from a control choice point (T2). This supports the idea that hippocampal lookahead can be dynamically engaged during decision making. Critically, using the same analysis on the same task, we found no evidence for lookahead in dorsal striatum, even though dorsal ensembles could represent location as well as or better than hippocampal ensembles on this task (Figure 2D). This demonstrates that lookahead is not a general, brain-wide phenomenon shared by all task-relevant representations but in the current dataset is restricted to brain systems known to play a role in “model-based” decision making.

Dorsal Striatum as a Situation-Action Associator

We found that dorsal striatal firing, field, and decoding distributions were skewed toward the turn sequence of the task, as well as reward locations and cues predicting reward delivery (Figure 2). The turn sequence and reward cues together determine the structure of the task, i.e., how the actions the rat takes map onto motivationally relevant outcomes. Models that learn what action to take in what situation in order to maximize reward (such as temporal-difference reinforcement learning) need to represent this information (Sutton and Barto, 1998). Thus, dorsal striatum selectively represents those task aspects which computational accounts suggest are important for gradual, model-free learning. This is congruent with previous suggestions about the role of dorsal striatum as indicated by inactivation, recording, and imaging studies (Poldrack and Packard, 2003; Balleine et al., 2007; Redish et al., 2008). However, in our comparative approach, we can additionally show what dorsal striatum *does not* represent. It does not represent locations ahead of the animal at decision points, as hippocampus does; neither does it covertly represent rewards, as ventral striatum does. Although a population of dorsal striatal neurons responded to reward-predictive cues, these were not neurons that were activated by the rewards themselves (Figure 3B), consistent with a developing representation of cue-action value associations. Thus, dorsal striatum does not appear to employ model-based internal generation of possible outcomes.

Dorsal striatum did not represent all locations equally, even though animals executed similar actions at those locations, sug-

gesting that that dorsal striatum learns to disregard task-irrelevant aspects with experience (such as the maze segment from the reward sites to the start of the turn sequence, which is constant and does not contain decision points). The gradual increase in coding efficiency further supports such reorganization with experience, consistent with reports from Graybiel and colleagues (e.g., Barnes et al., 2005), although we show this effect within session (instead of across days) and using an ensemble measure (which addresses spatial information content rather than firing rates alone). Taken together, these results support the notion that dorsal striatum learns to represent situation-action associations as proposed by computational accounts of model-free, habitual, or response-driven decision making. It explicitly does not share properties important for model-based decision making, even though the same analysis reveals such properties in other structures on the same task.

It may be surprising that dorsal striatum appears to contain more information about location on the track than hippocampus, whose relatively uniform distribution of place cells is well suited to spatial representation. However, on this task, spatial location is an important element of task structure (whether to turn left or right depends on location; reward locations are fixed). As such, dorsal striatum would be expected to represent spatial information on this task. Others have also observed that dorsal striatal firing patterns contain information about spatial location (Wiener, 1993; Yeshenko et al., 2004). However, this does not mean that dorsal striatal representations are intrinsically spatial. In fact, studies on tasks where reward locations were explicitly dissociated from space (Schmitzer-Torbert and Redish, 2008) or where multiple locations were equivalently associated with rewards (Berke et al., 2009) found that dorsal striatum did not represent space well. More generally, these considerations serve as a reminder that it is important to consider both task structure and ensemble-level properties when making inferences about representational content.

Ventral Striatum as an Evaluator of Actual and Expected Situations

As shown by van der Meer and Redish (2009), ventral striatal reward neurons tended to reactivate at the final choice point during early laps. We show here that even though dorsal striatum has a similarly sized population of reward-responsive neurons, by the same analysis on the same task dorsal striatal neurons do not show this effect. This serves, first, as a particularly informative control that strengthens the original finding by illustrating that it is not due to nonspecific behavioral features such as simply pausing at the choice point. More importantly, this difference in information processing mechanisms in dorsal and ventral striatum maps onto the conceptual difference between situation-action representations and action-outcome representations: while dorsal striatal neurons learned to respond to reward-predictive cues, these neurons did not respond to actual rewards. This suggests a potential role for ventral striatum in model-based decision making.

Ventral striatum is generally acknowledged as an important structure in mediating motivated or goal-directed behaviors. A popular suggestion is that it acts as the “critic” component of a reinforcement learning system (Atallah et al., 2007). Anticipatory

ramp cells seen in primate and rodent studies could be interpreted as an instantiation of a critic-like value signal (Schultz et al., 1992; Lavoie and Mizumori, 1994; Miyazaki et al., 1998). The ramp nature of this signal suggests a certain motivational relevance. On our task, it is clear that ventral and dorsal striatum have very different information processing properties. Dorsal striatum lacks the population firing rate ramp of ventral striatum, while ventral striatal decoding accuracy was poor compared to dorsal striatum and hippocampus. This suggests that ventral striatum represents global quantities related to value or motivation, which may fluctuate throughout a session, resulting in poor decoding accuracy. Thus, our results imply that ventral striatum may carry multiple motivationally relevant signals: a global ramp that may serve as the value signal in model-free learning systems, but also the covert representation of reward important for model-based systems.

Synthesis

In conclusion, we observed multiple dissociations in information processing between dorsal striatum, ventral striatum, and hippocampus. While hippocampal neural ensembles encoded future paths during pauses at the choice point, dorsal striatal ensembles did not. While ventral striatal reward-related cells showed activity during pauses at the choice point, dorsal striatal reward-related cells did not. In contrast, dorsal striatal ensembles slowly developed a more accurate spatial representation than hippocampal ensembles on the action-rich navigation sequence of the task, and dorsal striatal non-reward-related cells slowly developed responses to high-value cues.

These differences reveal the different computations these structures are performing to accommodate their roles in model-based and model-free decision making: Hippocampus provides a cognitive map of the environment, which can dynamically represent potential future paths. During pauses at choice points, ventral striatal reward representations are reactivated, as an expectation of future reward outcome. Ventral striatum also develops an activity ramp through the task, which may provide a motivationally relevant signal. Dorsal striatum does not represent expectancies or show a firing rate ramp, but instead develops stimulus-action representations with experience. These data bridge the behavioral and lesion data with computational/theoretical models of decision making, directly linking distinct behavioral roles with unique information processing mechanisms at the neural level.

EXPERIMENTAL PROCEDURES

Subjects

Sixteen male Brown Norway-Fisher 344 hybrid rats (Harlan, IA), aged 8–12 months, were food deprived to no less than 85% of their free-feeding body weight during behavioral training; water was available ad libitum in the home cage at all times. All procedures were conducted in accordance with National Institutes of Health guidelines for animal care and approved by the IACUC at the University of Minnesota. Care was taken to minimize the number of animals used in these experiments and to minimize suffering.

Multiple-T Task

The Multiple-T task apparatus, a carpet-lined track elevated 15 cm above the floor, consisted of a turn sequence of three to five T-choices, a top and a bottom rail, and two return rails leading back to the start of the turn sequence (Figure 1A). The specific configuration of the turn sequence could be varied from day to day. Two feeder sites at each of the return rails could deliver

two 45 mg food pellets each (Research Diets, New Brunswick, NJ) through computer-controlled pellet dispensers (Med-Associates, St. Albans, VT), released when a ceiling-mounted camera and a position tracking system (Cheetah, Neuralynx, Bozeman, MT, and custom software written in MATLAB, Natick, MA) detected the rat crossing the active feeder trigger lines. Only one set of feeder sites (either on the left or the right return rail) was active in any given session. For presurgery training, rats ran 3-T and 5-T mazes with the turn sequence changed every day; once rats were running proficiently after surgery, recording sessions were run on 4-T mazes in a sequence of seven novel/seven unchanged/seven novel configurations, for a total of 21 sessions per rat. Novel sequences consisted of session-unique sequences (Figure 1A shows the “RRLR” sequence). 98, 96, and 31 recording sessions from dorsal striatum, ventral striatum, and hippocampus, respectively, were accepted for analysis: for the hippocampal recordings we obtained good ensemble sizes only for a few days out of the 21 day protocol. However, the proportions of Novel/Familiar sessions were comparable to those in the other data sets (20 novel/11 familiar, compared to 68/30 for dorsal and 68/28 for ventral striatum). Furthermore, all analyses reported are within-session only, so the number of sessions should not affect the results. Rats were allowed to run as many laps as desired in each 40 min recording session. Data collection was as described previously (Schmitzer-Torbert and Redish, 2004; Johnson and Redish, 2007; van der Meer and Redish, 2009).

Surgery

Following pretraining, rats were chronically implanted with an electrode array consisting of 12 tetrodes and two reference electrodes that could be moved in the dorsal-ventral plane (“hyperdrive,” Kopf, Tujunga, CA). Structures were targeted by centering the hyperdrive on stereotaxic coordinates relative to bregma: AP +1.2, ML ± 2.3–2.5 mm for ventral striatum, AP +0.5, ML ± 3.0 mm for dorsal striatum, and AP –3.8, ML ± 4.0 mm for hippocampus (subfield CA3) as described previously (Schmitzer-Torbert and Redish, 2004; Johnson and Redish, 2007; van der Meer and Redish, 2009).

Spatial Decoding

We used a Bayesian spatial decoding algorithm, designed to provide an estimate of the animal’s location given ensemble spiking activity at any given time in the session (Zhang et al., 1998). For each time window, this method takes the spike counts from each cell i and computes the posterior probability of the rat being at location x given spike counts s_i , $p(x|s)$. We used a time window of 200 ms (for the analysis in Figure 2) or 20 ms (for the analysis in Figure 3) and a uniform spatial prior. To obtain the decoding accuracy measure, for each time window, the probability of decoding to the animal’s actual location was taken from the decoded probability distribution for that time window; the pattern of results did not change when slightly wider windows of three and five bins around the animal’s actual location were used. This “local probability” was then averaged for each actual location on the track; this was done to minimize the contribution of long periods of inactivity at the reward sites when averaging over time windows alone. (This method is identical to that used in Schmitzer-Torbert and Redish [2008] and in van der Meer and Redish [2009].) Only recording sessions with at least 20 simultaneously recorded cells were used and only cells that fired at least 25 spikes in the session were included. For the slow timescale analysis, using a time window of 50, 100, or 150 ms did not change the pattern observed (data not shown).

Further experimental procedures, including cell classification, track linearization, and spatial distribution analyses, can be found in the Supplemental Information.

SUPPLEMENTAL INFORMATION

Supplemental Information includes full Experimental Procedures, four figures, and three tables and can be found with this article online at doi:10.1016/j.neuron.2010.06.023.

ACKNOWLEDGMENTS

We thank Anoopam Gupta, Jadin Jackson, and Paul Schrater for discussion, and David Crowe and Geoffrey Schoenbaum for comments on a previous

version of the manuscript. Supported by NIH grants MH068029, MH080318, T32HD007151 (Center for Cognitive Sciences, University of Minnesota), IGERT training grant 9870633, graduate student NSF fellowships, Dissertation Fellowships, and the Land Grant Professorship program at the University of Minnesota. This work has been published in abstract form (Frontiers in Systems Neuroscience, conference abstract: CoSyNe 2009); we are grateful to the conference participants for their suggestions.

Accepted: June 3, 2010

Published: July 14, 2010

REFERENCES

- Atallah, H.E., Lopez-Paniagua, D., Rudy, J.W., and O'Reilly, R.C. (2007). Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. *Nat. Neurosci.* *10*, 126–131.
- Balleine, B.W., and Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* *37*, 407–419.
- Balleine, B.W., Delgado, M.R., and Hikosaka, O. (2007). The role of the dorsal striatum in reward and decision-making. *J. Neurosci.* *27*, 8161–8165.
- Barnes, T.D., Kubota, Y., Hu, D., Jin, D.Z., and Graybiel, A.M. (2005). Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* *437*, 1158–1161.
- Berke, J.D., Breck, J.T., and Eichenbaum, H. (2009). Striatal versus hippocampal representations during win-stay maze performance. *J. Neurophysiol.* *101*, 1575–1587.
- Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* *8*, 1704–1711.
- Hu, D., and Amsel, A. (1995). A simple test of the vicarious trial-and-error hypothesis of hippocampal function. *Proc. Natl. Acad. Sci. USA* *92*, 5506–5509.
- Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* *27*, 12176–12189.
- Kawaguchi, Y. (1993). Physiological, morphological, and histochemical characterization of three classes of interneurons in rat neostriatum. *J. Neurosci.* *13*, 4908–4923.
- Lavoie, A.M., and Mizumori, S.J. (1994). Spatial, movement- and reward-sensitive discharge by medial ventral striatum neurons of rats. *Brain Res.* *638*, 157–168.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* *416*, 90–94.
- Maurer, A.P., and McNaughton, B.L. (2007). Network and intrinsic cellular mechanisms underlying theta phase precession of hippocampal neurons. *Trends Neurosci.* *30*, 325–333.
- McNaughton, B.L., Barnes, C.A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res.* *52*, 41–49.
- Mehta, M.R., Barnes, C.A., and McNaughton, B.L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* *94*, 8918–8921.
- Miyazaki, K., Mogi, E., Araki, N., and Matsumoto, G. (1998). Reward-quality dependent anticipation in rat nucleus accumbens. *Neuroreport* *9*, 3943–3948.
- Nakahara, H., Amari, S., and Hikosaka, O. (2002). Self-organization in the basal ganglia with modulation of reinforcement signals. *Neural Comput.* *14*, 819–844.
- Narayanan, N.S., Kimchi, E.Y., and Laubach, M. (2005). Redundancy and synergy of neuronal ensembles in motor cortex. *J. Neurosci.* *25*, 4207–4216.
- Niv, Y., Joel, D., and Dayan, P. (2006). A normative perspective on motivation. *Trends Cogn. Sci.* *10*, 375–381.
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* *34*, 171–175.
- O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford: Clarendon Press).
- Packard, M.G., and McGaugh, J.L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol. Learn. Mem.* *65*, 65–72.
- Poldrack, R.A., and Packard, M.G. (2003). Competition among multiple memory systems: converging evidence from animal and human brain studies. *Neuropsychologia* *41*, 245–251.
- Ranck, J.B., Jr. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp. Neurol.* *41*, 461–531.
- Redish, A.D., Jensen, S., and Johnson, A. (2008). A unified framework for addiction: vulnerabilities in the decision process. *Behav. Brain Sci.* *31*, 415–437.
- Schacter, D.L., and Tulving, E., eds. (1994). *Memory Systems* (Cambridge, MA: MIT Press).
- Schmitzer-Torbert, N., and Redish, A.D. (2002). Development of path stereotypy in a single day in rats on a multiple-T maze. *Arch. Ital. Biol.* *140*, 295–301.
- Schmitzer-Torbert, N., and Redish, A.D. (2004). 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.
- Schmitzer-Torbert, N.C., and Redish, A.D. (2008). Task-dependent encoding of space and events by striatal neurons is dependent on neural subtype. *Neuroscience* *153*, 349–360.
- Schneidman, E., Bialek, W., and Berry, M.J., 2nd. (2003). Synergy, redundancy, and independence in population codes. *J. Neurosci.* *23*, 11539–11553.
- Schultz, W., Apicella, P., Scarnati, E., and Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *J. Neurosci.* *12*, 4595–4610.
- Shallice, T. (1982). Specific impairments of planning. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *298*, 199–209.
- Stanley, G.B., Li, F.F., and Dan, Y. (1999). Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J. Neurosci.* *19*, 8036–8042.
- Sutton, R., and Barto, A. (1998). *Reinforcement Learning: An Introduction* (Cambridge, MA: MIT Press).
- Tolman, E. (1938). The determiners of behavior at a choice point. *Psychol. Rev.* *45*, 1–35.
- van der Meer, M.A.A., and Redish, A.D. (2009). Covert expectation-of-reward in rat ventral striatum at decision points. *Front. Integr. Neurosci.* *3*, 1–15.
- Wessberg, J., Stambaugh, C.R., Kralik, J.D., Beck, P.D., Laubach, M., Chapin, J.K., Kim, J., Biggs, S.J., Srinivasan, M.A., and Nicolelis, M.A. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* *408*, 361–365.
- Wiener, S.I. (1993). Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *J. Neurosci.* *13*, 3802–3817.
- Yeshenko, O., Guazzelli, A., and Mizumori, S.J.Y. (2004). Context-dependent reorganization of spatial and movement representations by simultaneously recorded hippocampal and striatal neurons during performance of allocentric and egocentric tasks. *Behav. Neurosci.* *118*, 751–769.
- Yin, H.H., Knowlton, B.J., and Balleine, B.W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *Eur. J. Neurosci.* *19*, 181–189.
- Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* *79*, 1017–1044.