

Dynamic Coding of Goal-Directed Paths by Orbital Prefrontal Cortex

James J. Young and Matthew L. Shapiro

Fishberg Department of Neuroscience and Alfred B. and Gudrun J. Kastor Neurobiology of Aging Laboratories, Mount Sinai School of Medicine, New York, New York 10029-6574

Adapting successfully to new situations relies on integrating memory of similar circumstances with the outcomes of past actions. Here, we tested how reward history and recent memory influenced coding by orbital prefrontal cortex (OFC) neurons. Rats were trained to find food in plus maze tasks that required both the OFC and the hippocampus, and unit activity was recorded during stable performance, reversal learning, and strategy switching. OFC firing distinguished different rewarded paths, journeys from a start arm to a goal arm. Activity of individual cells and the population correlated with performance as rats learned newly rewarded outcomes. Activity was similar during reversal, an OFC-dependent task, and strategy switching, an OFC-independent task, suggesting that OFC associates information about paths and outcomes both when it is required for performance and when it is not. Path-selective OFC cells fired differently during overlapping journeys that led to different goals or from different starts, resembling journey-dependent coding by hippocampal neurons. Local field potentials (LFPs) recorded simultaneously in the OFC and the hippocampus oscillated coherently in the theta band (5–12 Hz) during stable performance. LFP coherence diminished when rats adapted to altered reward contingencies and followed different paths. Thus, OFC neurons appear to participate in a distributed network including the hippocampus that associates spatial paths, recent memory, and integrated reward history.

Introduction

To respond appropriately in new circumstances, people remember similar situations and the consequences of past actions. The orbitofrontal cortex (OFC) may contribute to flexible responding by integrating information about outcome with contextual knowledge (Wallis, 2007). Rats with OFC lesions are impaired when learning requires withholding previously rewarded responses and producing previously unrewarded ones (e.g., in “reversal learning”) (Schoenbaum et al., 2002, 2003a; McAlonan and Brown, 2003; Kim and Ragozzino, 2005; Stalnaker et al., 2007; Young and Shapiro, 2009). In contrast, “strategy switching” requires attending to a new class of stimuli or executing a new category of responses and is impaired in rats with lesions to prelimbic–infralimbic cortex (PL–IL) (Ragozzino et al., 1999a,b; Birrell and Brown, 2000; Rich and Shapiro, 2007; Young and Shapiro, 2009).

OFC may support reversal learning by representing aspects of reward. OFC neurons in rats fire selectively with the presentation of rewarded odors (Schoenbaum and Eichenbaum, 1995a,b; Schoenbaum et al., 1998, 1999, 2000), selectivity emerges with the same time course as learning (Schoenbaum et al., 1998, 1999;

Ramus and Eichenbaum, 2000; Alvarez and Eichenbaum, 2002), and selectivity changes on reversal (Schoenbaum et al., 1999, 2003b; Roesch et al., 2007). OFC neuronal activity reflects the expected or relative reward associated with stimuli (Tremblay and Schultz, 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006). These results are consistent with the hypothesis that OFC encodes reward expectancy, the probability that a stimulus presented on a particular trial will be followed by reward delivery (Schoenbaum et al., 2009).

Most studies of OFC encoding have used tasks that associate explicit stimuli with rewards presented after relatively brief intervals. In contrast, we hypothesized that, beyond associating simple stimuli with expected reward, the OFC may help guide discriminative responses by associating temporally extended, goal-directed actions with reward history. Furthermore, if associating reward history with similar past experiences is required to solve a task, OFC coding may be coordinated with activity in brain regions required for memory, such as the hippocampus. To test these predictions, we simultaneously recorded OFC unit activity and hippocampal local field potentials (LFPs) as rats learned and performed spatial reversals or strategy switches in plus maze tasks that require the hippocampus and OFC (White and McDonald, 2002; Young and Shapiro, 2009) (see Fig. 1).

Materials and Methods

Experimental overview. Rats were assigned to reversal and switch groups (Fig. 1), implanted with recording electrodes in the OFC and the dorsal hippocampus, and allowed to recover from surgery. Each rat was trained on an initial task, and then tested on either repeated spatial reversals or strategy switches while neural activity was recorded. Each spatial reversal and strategy switch session began with a 20 trial retention test of the task

Received Oct. 17, 2010; revised Jan. 23, 2011; accepted Feb. 21, 2011.

This work was supported by National Institutes of Health Grants MH065658, MH073689, and MH084436, and the Mount Sinai School of Medicine. We thank Mark Baxter, Denis Pare, and Geoffrey Schoenbaum for comments on a previous draft. We also thank Maojuan Zhang for assistance building hyperdrives, and Prasad Shirvalkar and Janina Ferbinteanu for technical advice.

Correspondence should be addressed to Matthew L. Shapiro, Mount Sinai School of Medicine, One Gustave Levy Place, Box 1065, New York, NY 10029. E-mail: matthew.shapiro@mssm.edu.

DOI:10.1523/JNEUROSCI.5436-10.2011

Copyright © 2011 the authors 0270-6474/11/315989-12\$15.00/0

of the previous day, followed by training and testing on the new task. Single-unit and LFP recordings from OFC and LFP recordings from the hippocampus were analyzed and compared with behavioral variables such as path, region, and learning performance.

Animals. Long-Evans rats ($n = 9$; males; 2 months of age; ~ 300 g) were housed individually in a room on 12 h light/dark cycle. After acclimating to the colony room for 1 week, rats were food-restricted to 85% of their *ad libitum* body weight and maintained on a food-restricted diet for the duration of the experiments. All procedures with animals followed National Institutes of Health and Institutional Animal Care and Use Committee guidelines.

Maze. The plus maze had four wooden arms (65 × 8 cm) meeting at 90° angles, two start and two goal arms (Fig. 1A). Food wells drilled into the ends of each arm held cereal reward (one-half Froot Loops cereal; Kellogg's). Mesh screens covered the bottom of the wells, where an inaccessible food reward was placed to minimize the influence of odor cues in the task. A waiting platform was placed next to the maze. The maze and waiting platform were open to the testing room, the walls of which had several distal visual cues.

Task. On a given trial, either the east or the west arm was designated as the “start” arm, and either the north or south arm was designated the “goal” arm. In all tasks, the rat was placed on a start arm and trained to enter one of the goal arms for food reward. A wooden block prevented access to the unused start arm on each trial. On each trial, both goal arms were open, but only one held food. This study distinguishes paths, tasks, and strategies as hierarchically related, increasingly abstract descriptions of behavior in the plus maze (Rich and Shapiro, 2009) (Fig. 1B). A path describes a specific journey through the maze (e.g., east arm-to-north arm). A task is a rule describing the set of paths leading to reward [e.g., in the spatial task “go north,” both paths ending in the north goal arm (east-to-north and west-to-north paths) were rewarded]. A strategy is an abstract rule that does not dictate reward contingencies, but rather the relevant stimuli, cognitive processes, and responses required to solve the task (e.g., the spatial tasks “go north” and “go south” require approaching specific allocentric locations, whereas the response tasks “turn right” and “turn left” require executing specific egocentric body turns). Thus, the plus maze includes four potential tasks—each composed of two paths—divided into two strategies.

Maze acclimation. Rats were handled, acclimated to the testing environment, and allowed to forage for cereal spread over the maze until the rat consumed all the available food in <10 min during a single testing day.

Recording equipment. Assemblies of 12 independently driveable tetrodes (Neuro-Hyperdrive; David Kopf Instruments) were constructed using 12.7- μm -diameter Nichrome wire (RO-800; Kanthal Precision Technology). Each tetrode wire was gold plated to an impedance <0.5 M Ω . A 0.01 inch Nichrome Formvar-coated wire (California Fine Wire) attached to the hyperdrive recorded hippocampal LFPs. The drive assembly was attached via a head stage to a Digital 64 Cheetah data acquisition system (Neuralynx). The Cheetah system recorded simultaneously from the 12 tetrodes and the hippocampal electrode. The animal's location was recorded using light-emitting diodes mounted on the head stage that

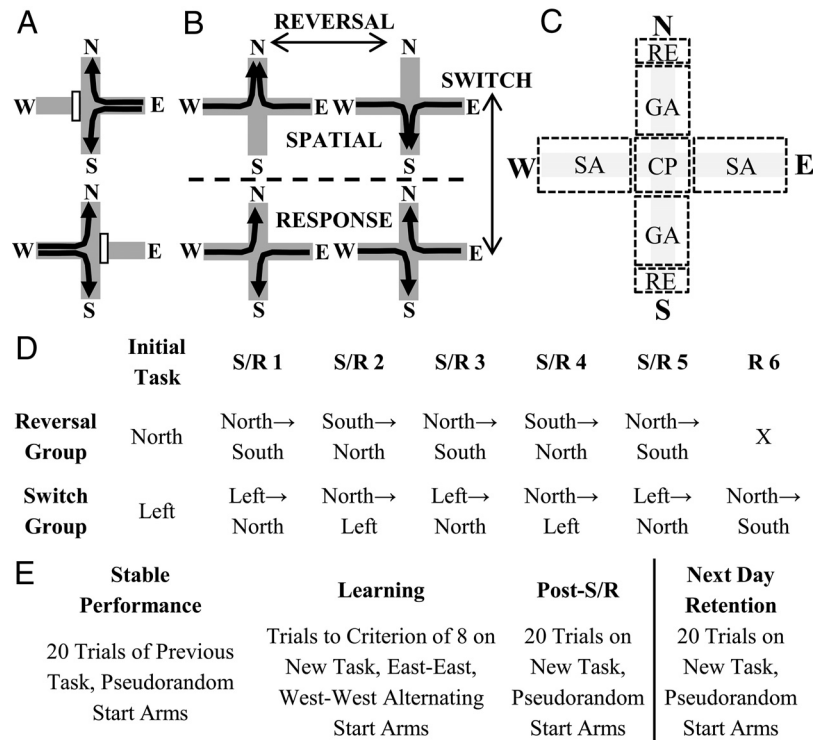


Figure 1. Experimental design. **A**, The plus maze had two start arms (east or west). On each trial, the unused start arm was blocked. The animal could enter one of two goal arms (north or south), only one of which was rewarded. **B**, The plus maze assessed four different tasks. Two spatial tasks were solved by entering the goal arm in one spatial location (north or south). Two response tasks were solved by executing a specific body turn (right or left) to enter a goal arm. A strategy switch occurred when the task was changed from spatial to response or vice versa. A spatial reversal occurred when the task was changed from one spatial task to the other. Each task included two paths [e.g., the north task was the combination of the west-to-north (WN) and east-to-north (EN) paths]. **C**, Along each path, the trial was divided into regions and marked by computer with behavioral flags: the start arm (SA), the choice point (CP), the goal arm (GA), and the reward cup (RE). **D**, Rats in the reversal group ($n = 5$) were trained initially to go north. They were then tested on five serial spatial reversals (north→south or south→north). Rats in the switch group ($n = 4$) were trained initially to go left and then tested on five serial strategy switches (left→north or north→left). A rat that completed five strategy switches was then tested on an additional north→south spatial reversal. **E**, Each reversal or strategy switch was divided into four trial stages. At the start of the testing day, the animal was tested for 20 trials in the previously rewarded task with pseudorandomly ordered start arms (never more than 3 in a row of 1 start arm) (SP, stable performance). Then the reward contingency was changed to the new task. The animal was tested on the new task with alternating pairs of start arms (east→west, west→west, etc.) until it performed eight in a row correct on the new task (learning). The animal was then tested for 20 trials with pseudorandomly ordered start arms on the new task (post-5/R). The following day, the animal was tested on the new task for another 20 trials with pseudorandomly ordered start arms (retention).

were detected by an overhead video camera (640 × 480 camera pixels; 16.7 ms sampling rate).

Surgery. Rats were anesthetized with isoflurane and placed in a stereotaxic frame. The scalp was shaved, scrubbed with Betadine, locally anesthetized with lidocaine, and incised. A hole was drilled at +4 mm anteroposterior (AP), +2.2 mm left mediolateral (ML) from bregma. The dura was removed, and the hyperdrive electrode bundle was lowered to the cortical surface. Six of the nine animals were implanted with a fixed electrode in the dorsal hippocampus [−4 mm AP, +2 mm left ML, −2.5 mm dorsoventral (DV)]. Two ground wires were attached to skull screws immediately behind and to either side of lambda. The drive assembly, electrodes, and skull screws were fixed to the skull with dental acrylic. All tetrodes were immediately lowered ~ 3 –3.5 mm DV into the cortex. Rats were allowed to recover for 7 d after surgery. Tetrodes were then lowered 100 $\mu\text{m}/\text{d}$ until each had identifiable cells as detected by stable waveform clusters. After lowering a tetrode to its final depth, no testing was performed for at least 48 h. Furthermore, once tetrodes had reached their correct depth, they were no longer moved until the conclusion of testing.

Pretraining. After recovery from surgery, the rats were reacclimated to the maze with another day of foraging for randomly distributed food rewards. The next day, cereal rewards were placed only in the food cups on both goal arms and one start arm was blocked. Rats were placed on

each start arm twice and given access to both goal arms for a total of four trials.

Initial task training. On each trial, the rat was placed at the distal end of a start arm facing the center of the maze and allowed to enter one of the goal arms. Entering one full body length into a goal arm defined a choice. The trial ended when the rat either proceeded to the end of the arm or attempted to turn around and enter the other arm. Intertrial intervals were 5–8 s. If the rat chose the correct arm, it was allowed to consume the food and was placed on the waiting platform until food had been replaced. If the rat entered the incorrect arm, it was returned to the waiting platform with no reward. During training, the start arms were alternated (two west, two east, etc.) until the animal reached a criterion of eight-in-a-row correct. The animal was limited to 40 trials on each training day. Twenty-four hours after reaching the eight-in-a-row correct criterion, the rats were tested for 20 trials with pseudorandomly ordered start arms (such that no more than three consecutive trials used the same start arm). Rats were required to perform at least 80% of trials correctly (less than four errors) each day until they met this performance criterion on 2 consecutive days.

Testing. Five rats were assigned to a reversal group, and four rats were assigned to a switch group (Fig. 1D). The reversal group was initially trained on the “go south” task; the switch group on the “go left” task. Each testing day, a rat was given 20 trials with pseudorandomly ordered start arms on the previously rewarded task (Fig. 1E). For the first reversal or strategy switch, this was the animal’s initial task. Then, the reward contingencies were changed so that the rewarded paths reflected a new task. In the reversal group, the task was changed from “go south” to “go north.” In the switch group, the task was changed from “go left” to “go north.” Testing on the new task proceeded until the animal performed eight consecutive correct trials. On reaching criterion, the animal was tested for 20 trials with pseudorandomly ordered start arms on the new task. The next day, the animal was tested again for 20 trials with pseudorandomly ordered start arms on the new task. If the animal failed to perform >80% of the trials correctly on the retention test, it was given the same task on subsequent days for 20 trials with pseudorandomly ordered start arms until it achieved >80% correct. During subsequent sessions, animals in the reversal group were reversed between “go north” and “go south,” and animals in the switch group were switched between “go left” and “go north.” In addition, animals in the switch group that successfully concluded five strategy switches were tested on a single reversal from “go north” to “go south.” Testing proceeded until the animal completed five spatial reversals or five strategy switches followed by one spatial reversal.

Data acquisition. OFC and hippocampal activity was recorded each trial from the time the animal was placed on the start arm until immediately before it received a reward or, during an error trial, reward was omitted. Recording was terminated before the rat consuming the reward to exclude motion artifacts from chewing. Recording was also paused while the animal was on the waiting platform. The gain (500–5000 \times) was adjusted for each tetrode to maximize waveform discrimination at the start of each testing day. Action potentials (threshold, $\geq 100 \mu\text{V}$) were sampled at 32,556 Hz (600–6000 Hz), and LFPs were sampled at 2012 Hz (0–475 Hz).

Unit discrimination. Single units were defined by homogenous waveforms quantified by sets of waveform parameters (peak–valley duration, peak–valley height, etc.) clustered in a multidimensional parameter space. Each waveform defined a point in this space and was assigned to clusters with a semiautomated, nonlinear elliptical cluster cutting software (Ferbinteanu and Shapiro, 2003). To confirm tetrode stability, the average waveform for the action potentials of the cell during the first 10 trials and last 10 trials of a testing day were compared and shown to be highly correlated (supplemental Fig. 1, available at www.jneurosci.org as supplemental material).

Identification of cell types. Pyramidal cells and interneurons can be distinguished by electrophysiological characteristics such as waveform shape and average firing rate (McCormick et al., 1985; Wilson et al., 1994). To distinguish between pyramidal cells and interneurons, the waveform from the tetrode channel with the largest spike height for each unit was normalized to its peak-to-trough voltage difference. The wave-

forms from all cells were then subjected to *k*-means clustering (two clusters) to distinguish two groups (Diester and Nieder, 2008). The group that showed delayed afterpolarization—as indicated by longer average peak-to-trough time—and a lower average firing rate was labeled putative pyramidal cells. The group that showed symmetrical waveforms and a higher average firing rate was labeled putative interneurons.

Histology. After testing was complete, rats were deeply anesthetized with isoflurane and pentobarbital (50 mg/ml, i.p.). One wire in each tetrode was pulsed with 300 μA DC to label the tetrode tips. The animal was then transcardially perfused with ice-cold PBS followed by 4% paraformaldehyde. Brains were removed, placed in 4% paraformaldehyde for at least 24 h, and then cryoprotected in 18% followed by 30% sucrose solution. Fifty micrometer coronal sections were cut on a cryostat, mounted on slides, and stained with formol-thionin.

Behavioral correlates and statistical analysis. To ensure that overt behavior was homogenous across testing conditions, all trials with duration or distance moved more than ± 2 SD from the session median were excluded from additional analysis. Trials in which the animal doubled back and went the wrong direction on the path were also excluded. In addition, spike data were speed-filtered to exclude times when the animal was moving < 2 cm/s for > 200 ms. A speed-filtered rate was defined as the number of spikes that occurred in a given recording interval when the rat was moving > 2 cm/s. Behavior and unit analyses were performed using Matlab (The MathWorks) and Systat (Cranes). Variation of firing rate because of path, region, etc., were compared by ANOVA. *Post hoc* comparison of different groups used Bonferroni-corrected Student’s *t* test ($\alpha = 0.05$).

General linear model. The relative influence of running speed, heading, path, performance, and region on OFC activity was quantified by creating a general linear model (GLM) to predict the firing rate of each OFC cell. The activity of each OFC cell was calculated for every 10×10 cm grid unit traversed by the animal during the testing day. The firing rate—the dependent variable—was then predicted by five independent variables: running speed, heading, region of that grid unit, the path traversed, and the performance during that trial. Running speeds were *Z*-scored to session averages, and the *sin* of the heading was used to normalize the data. Performance was quantified by generating a learning curve representing the probability of the animal receiving reward on a given trial (Smith et al., 2004). Grid units that were significant upper outliers in terms of running speed (running speed $>$ mean speed + 1SD, approximately > 150 cm/s) were excluded. Interaction terms were excluded because of missing data. The GLM for each cell generated a regression coefficient associated with each independent variable. The average absolute regression coefficient across all cells represented the relative influence of each factor on OFC activity.

LFP spectral power. All LFP analyses were performed using Chronux (<http://chronux.org/>) (Mitra and Bokil, 2008). LFP signals from OFC and hippocampus were first scrubbed for 60 Hz noise using a 1 s window with 0.5 s steps (*rmlinesmovingwinc*) and locally detrended (*locdetrend*) with the same window. Then spectral power analyses for either single trials or groups of trials were generated by averaging 1 s windows of the LFP data with 0.5 s steps. The bandwidth product was 3, and five tapers were used. The 95% confidence intervals (CIs) were calculated by jackknife. Sampling frequency for the LFP signals was 2012 Hz, and the bandpass was 0–100 Hz. The spectral power was calculated for the period of stable performance before task change (SP), the learning period before achieving criterion (learning), and trials recorded after reaching criterion (post).

LFP coherence. Data were preprocessed as in spectral power. In addition, a 10 trial moving window was used in which all the trials were averaged to generate coherence (i.e., the coherence value is an average of 1 s windows with 0.5 s steps through 10 trials). Mean coherence measures during stable performance were calculated for delta (0.5–4 Hz), theta (5–12 Hz), beta (15–30 Hz), and gamma (30–80 Hz) bands by averaging the observed *R* values within that band for each of the 10 trial moving window (after Fisher’s *Z* transform). Changes in coherence from stable performance baseline were calculated by dividing the mean *R* value for a given 10 trial moving window by the average *R* value during stable performance for that animal for that session. The data of each trial were then

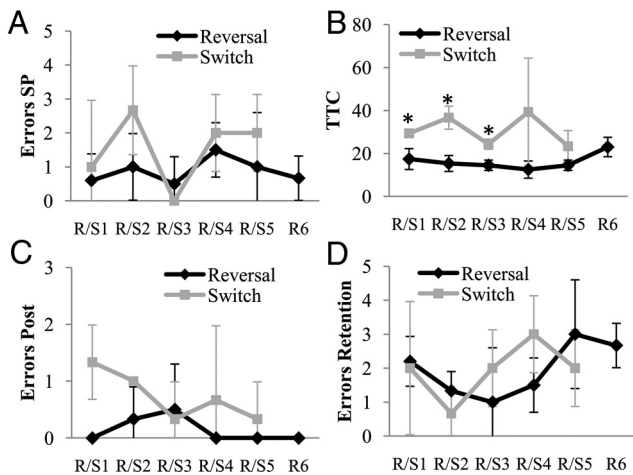


Figure 2. Performance in the reversal and switch groups across multiple spatial reversals and strategy switches (R/S) and during a spatial reversal after strategy switching (R6). **A, C, D**, Errors during stable performance of the previous days task (errors SP) (**A**), errors during stable performance of the new task (errors post) (**C**), and retention errors the following day (errors retention) (**D**) did not differ between groups. **B**, Trials to the eight-in-a-row criterion (TTC) were greater for the first three strategy switches than the first three spatial reversals ($*p < 0.05$, Bonferroni-corrected t test). Error bars represent 95% CIs.

averaged with all other trials within that stage [i.e., the learning period before achieving criterion (learning) and trials recorded after reaching criterion (post)].

Linearized trial diagrams. Linearized trial diagrams were created by generated firing rate maps (1×1 cm resolution) for each trial on the maze for a particular path. The firing rate map for each trial was then projected onto a single line at angle connecting the start and goal for that particular path using the *radon* transformation. The trialwise rate lines were averaged to create a mean rate at different points of the path. The 95% confidence intervals were generated by bootstrap (1000 iterations with replacement). The resulting mean rate line and CIs were smoothed using a 5 cm moving window. Because linearization required the beginning and end of the path to be the same for each trial, trials from only a single path are depicted. Note that linearized trial diagrams were generated to help visualize the data. Statistics were generated by measuring the filtered firing rate between behavioral flags using actual position data.

Results

Behavioral performance and electrode placement

Performance in the reversal and switch groups was comparable except that switches 1 through 3 required more trials to acquire than reversals 1 through 3 (TTC, $F_{(1,9)} = 24.6$, $p < 0.001$) (Fig. 2). Correct electrode placement was confirmed histologically using a standard brain atlas (Paxinos and Watson, 1998) (Fig. 3).

Most OFC cells fired throughout paths and maze regions

If the OFC contributes to learning the reward expectancy associated with goal-directed actions, then OFC activity should discriminate different paths on the plus maze. Indeed, many OFC neurons were selectively active as rats followed specific paths in the maze, defined as the journey from a particular start arm to a particular goal arm. To quantify this “path selectivity,” the first 20 and last 20 correct trials were divided into path categories, and the firing rate of each cell on each path was compared by ANOVA (e.g., east-to-north vs west-to-north). These trials were selected to compare OFC activity in the different paths before and after learning, when the rats were performing each task stably and obtaining reward consistently. The analysis thereby compared different paths associated with equivalent reward magnitude and




Figure 3. Electrode placement. **A**, Left side shows a formol-thionin-stained section from the OFC on an implanted animal (approximately bregma +3.5 mm AP). The arrowhead points to the lacunae near the end of a tetrode, created by running DC through the tetrode before perfusing the animal. The asterisk (*) indicates a visible tetrode track. Right side shows the placement of tetrodes in all nine animals (indicated by color) that varied between +3.2 and +4.0 mm on the AP axis. **B**, Left side shows a formol-thionin-stained section in the hippocampus of an animal implanted (approximately bregma -3.8 mm AP). The arrowhead indicates the hippocampal electrode. Right side indicates the placement of the hippocampal electrodes in the six implanted animals (indicated by color). Electrode placement varied between -4.0 and -3.5 mm AP, but placements are shown collapsed on a single section. The black line coronal section diagrams are modified from Paxinos and Watson (1998).

density. The path (or combination of paths) with trialwise firing rates significantly larger by t statistic compared with all other paths was defined as the path specificity of the cell. Of the 332 cells recorded during reversals and 318 during strategy switches, most were active on the plus maze (229 during reversals and 192 during strategy switches; firing rate > 1 Hz in any path condition). Of these active cells, 28% (73) were path selective in reversals and 31% (76) in strategy switches (ANOVA, $p < 0.05$). Most cells were specific to one path (51%; 76) or two paths (44%; 65), but some were specific to three paths (5%; 8) (Fig. 4).

Furthermore, most path-selective cells fired throughout journeys, so that activity was distributed evenly across paths (Fig. 5A). To quantify how OFC activity varied across different regions of the maze, each trial was divided into nonoverlapping regions—the start arm (SA), choice point (CP), goal arm (GA), and reward cup (RE)—and firing rates were compared by ANOVA. Although 55% (40) of cells recorded during reversals and 74% (56) during strategy switches fired differently across regions (ANOVA, $p < 0.05$), few were significantly more active in one region than the others [8% (6) during reversals, 26% (20) during strategy switches; Bonferroni-corrected t test, $p < 0.05$]. Of the 96 region-selective cells, 7% (7) fired specifically on the SA, 7% (7) at the CP, 4% (4) on the GA, and 8% (8) at the RE (Fig. 5B). Hence, most path-selective cells (74%, 70) fired equivalently in all regions, verifying that OFC activity was distributed throughout the paths. This tendency to fire evenly across regions of the maze is also reflected in the non-path-selective cells. Seventy-eight percent (212 of 272) also fired equivalently when activity in different maze regions was compared (Bonferroni-corrected t test, $p > 0.05$).

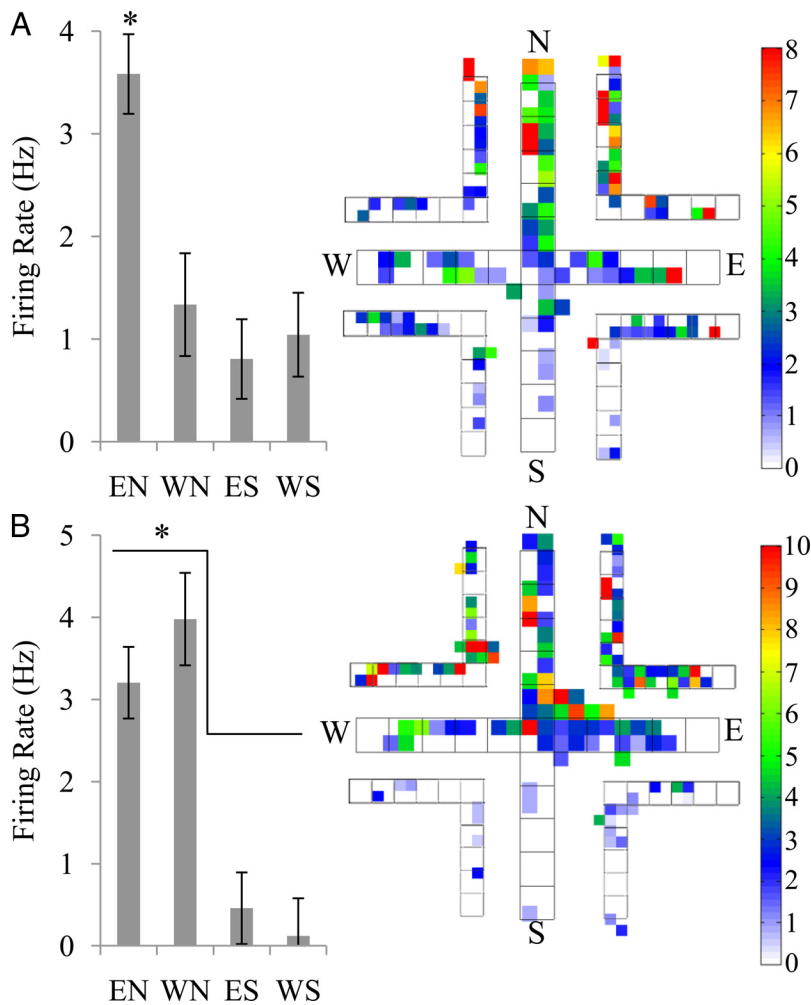


Figure 4. Path-selective cells. **A**, Example of a cell that showed path specificity for a single path (EN). Left shows mean firing rate on each path. The asterisk (*) indicates significantly greater firing rate in the EN path (Bonferroni-corrected t test, $p < 0.05$ between all groups). Right shows a heat plot of mean firing rates in 5×5 cm grid units. The black lines depict approximate location of the maze. The internal figure shows the mean firing rate across all trials, and the four external plots show mean firing rates on specific paths. The color bar shows firing rate in hertz. **B**, Example of a cell with path specificity for two paths (WN and EN). The asterisk (*) indicates that WN and EN paths had greater firing rates than other two paths (comparison by Student's t test, $p < 0.05$, between two groups).

Individual path-selective cell activity correlated with performance

If the OFC contributes to learning goal directed actions by integrating reward history, then OFC firing rates should correlate with performance, and path-selective cells did so. During each testing day, learning curves were estimated by calculating the probability of success on a given trial (Smith et al., 2004). The probability of a correct response during each trial was then correlated with the firing rates of path-selective cells during trials as rats followed the preferred path of each cell. Activity in 20% (30) of the path-selective cells recorded during strategy switches and reversals correlated significantly with performance (Spearman's ρ , $p < 0.05$), consistent with the idea that OFC codes the integrated history of reward associated with paths.

Rather than reflecting the integrated history of reward associated with a particular path, however, OFC firing may respond to recent error history, such as whether the immediately preceding trial was correct (CT) or an error (ER). From this view, path-selective activity should decline immediately after errors. To distinguish between these possibilities, path-selective firing was

compared during trials that immediately followed either by a CT versus an ER trial. Activity differed significantly in only 11% (17) of path-selective cells depending on whether the previous trial was CT or ER (Student's t test, $p < 0.05$). Thirty percent of these (5 of 17) fired significantly less when the previous trial had been an ER, whereas 70% (12 of 17) fired significantly less when the previous trial was CT. This result suggests that OFC activity is better correlated with the reward history associated with a path than recent error history. However, because the recent error history and performance are highly correlated measures, the present experiments cannot entirely distinguish between these two models.

Running speed, heading, and cell type do not significantly modify OFC activity

Behavioral factors such as running speed and heading could affect OFC activity and confound other behavioral measures. To assess these potential influences, a GLM was generated for each path-selective cell to measure the relative influence of running speed and heading versus path, performance, and region of the maze (for model specifications, see Materials and Methods). Running speed and heading influenced unit activity far less than other factors. In path-selective cells, 26% (39) were significantly influenced by running speed and in 5% (8) by heading (t test, $p < 0.05$). In comparison, path, region, and performance significantly affected 53% (79), 38% (57), and 60% (90) of cells, respectively. Thus, fewer cells were influenced by running speed and heading than by path, region, or performance. Furthermore, the average regression coefficient

(β)—a measure of effect size—was lower for running speed and heading than for other factors. β indicates how much a marginal change in a predictor variable (1 for a category or 1 SD for a continuous variable) will change the output—in this case, firing rate—with all other variables held constant. For running speed and heading, the average absolute regression coefficients were 0.33 ± 0.06 and 0.27 ± 0.05 , respectively—meaning that a change in running speed by 1 SD (~ 30 cm/s) would result in a 0.33 Hz change in firing rate—as opposed to 0.75 ± 0.09 for path, 1.21 ± 0.19 for region, and 2.49 ± 0.46 for performance. The average absolute regression coefficients for path, region, and performance were all significantly greater than those for running speed and heading (Student's t test, $p < 0.001$). These analyses suggest that, although running speed and heading may affect the firing of OFC neurons, they did not account for the effect of other behavioral correlates. In addition, the firing characteristics of putative pyramidal cells and interneurons were compared. Although both types of cells could be identified, their firing correlates did not differ significantly in these experiments

(supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

The population of OFC path-selective neurons code reward expectancy during reversal and strategy switching when spatial trajectory is constant

If OFC neurons integrate reward history for particular stimuli or responses (Schoenbaum et al., 2009), then path-selective OFC neurons should fire differently as rats learned new reward expectations for a path. To test this hypothesis, we measured the activity of path-selective neurons in OFC during reversal and strategy switching when the animal was on the start arms (i.e., where behavior and peripheral stimuli were consistent even as the animal learned to approach a different goal). Start arm activity was measured from the time the animal was placed on the maze until it entered a goal arm. The data were divided into three sets based on the dynamics of cell firing during stable performance and learning. One set included all cells with specificities for the paths rewarded during stable performance at the beginning of the testing day (before paths). A second set included cells with specificities for paths rewarded during stable performance of the new task (end paths). Finally, during strategy switching, a third set included cells with specificities for the path rewarded throughout the testing day, because it was common between the old task and the new task (common path). If the activity of OFC neurons reflects the integrated history of reward associated with particular paths, then as learning proceeds, cells coding before paths should fire less and cells coding end paths should fire more. Cells specific to the common path during strategy switching should fire stably across learning.

The firing rates of path-selective cells on the start arm correlated with performance dynamics during both reversals and strategy switching (Fig. 6). Learning curves were calculated to reflect the probability of success on a given trial (Smith et al., 2004), and adjacent trials were blocked to compare different phases of learning and performance. Cells specific to before paths showed initially high activity that declined with learning. During all reversals and all switches, activity declined immediately and significantly when the task was changed [SP2 (last 10 trials of stable performance) vs trial epoch 1 (mean performance, 53% in switches and 8% in reversals); Bonferroni-corrected t test, $p < 0.05$]. Changes in activity in the reversal 1 or switch 1 were less dramatic than subsequent learning sessions, but still declined significantly by the end of the testing day. Cells specific to end paths fired with increasing rates as learning proceeded. In all reversals, activity was higher than during stable performance (SP1) as the animals reached criterion (trial epoch 4, mean performance, 92%; Bonferroni-corrected t test, $p < 0.05$). In all switches, firing was significantly greater during the last 10 trials after achieving criterion (post2). Finally, cells specific to the consistently rewarded

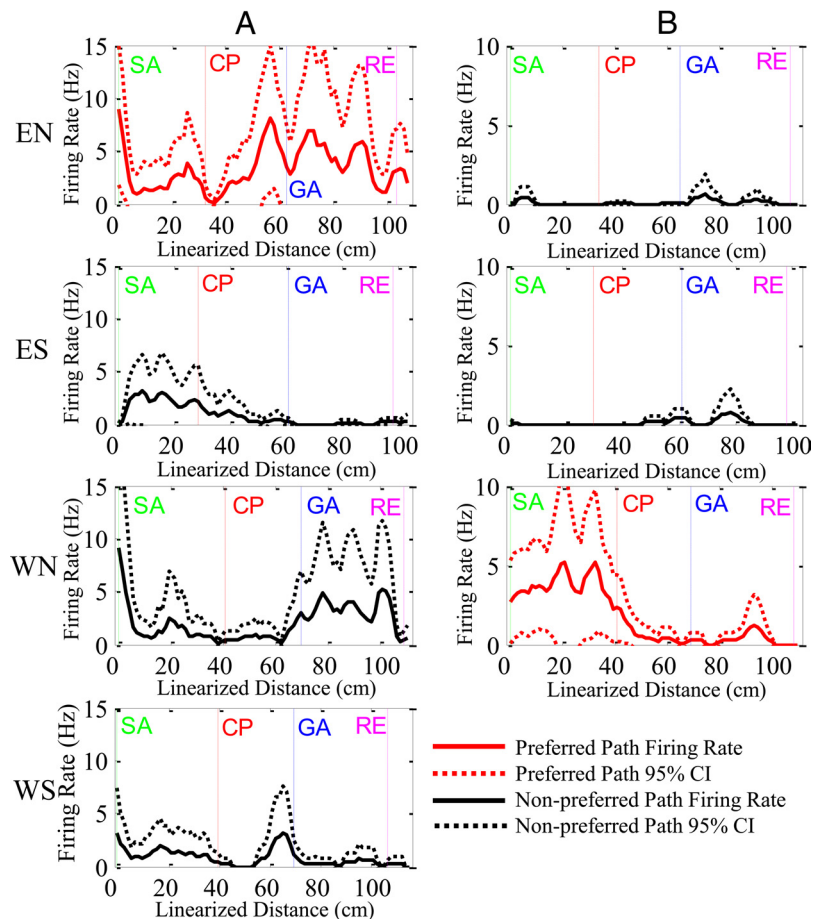


Figure 5. Region-selective and -specific cells. Region-selective cells fired differently across paths (e.g., EN vs ES), whereas region-specific cells fire differently within paths (e.g., start arm vs goal arm). **A**, An example of an EN path-selective cell with region-selective (ANOVA, $p < 0.05$ by region) but not region-specific firing (no region higher than others by Bonferroni-corrected t test). Linearized trial firing diagrams for trials on the EN, ES, WN, and WS paths (top to bottom). The line graph in each shows firing rate across regions with mean rate and 95% confidence intervals plotted as a function of distance along the path. The red line indicates the path selectivity of the cell, whereas the black lines are for nonselective paths. The vertical lines show the approximate positions of behavioral flags [start arm (SA), choice point (CP), goal arm (GA), and reward cup (RE)]. **B**, An example of WN path-selective cell that is also fired specifically on the SA (Bonferroni-corrected t test, $p < 0.05$, comparison with other groups). WS is not shown because the cell was recorded during a strategy switch that did not include this path.

path during strategy switching (common path) fired relatively consistently throughout the testing day. In switch 1, firing declined during trial epoch 3 (mean performance, 78%) but returned to baseline during the last 10 trials after reaching criterion (post2). Thus, even as overt behavior on the start arm was unchanged, firing rate dynamics correlated with learning new goals, consistent with the hypothesis that OFC neurons integrate reward history.

The results provide clear evidence that OFC neuronal activity correlates with learning the paths associated with a new task, even when traversing the same sequence of locations. However, this correlation could be interpreted as showing that OFC neurons are important for coding paths per se, rather than the integrated reward history associated with those paths. In other words, the firing dynamics that accompanied learning may have reflected the changing composition of the paths that were performed at the beginning versus the end of the testing day. To control for the effect of path composition on changes in firing rate associated with learning a new task, we analyzed path-selective cell activity on preferred paths before and immediately after the task change (i.e., during the first error trials). Cell activity on preferred paths

Start Arms Activity Across Learning

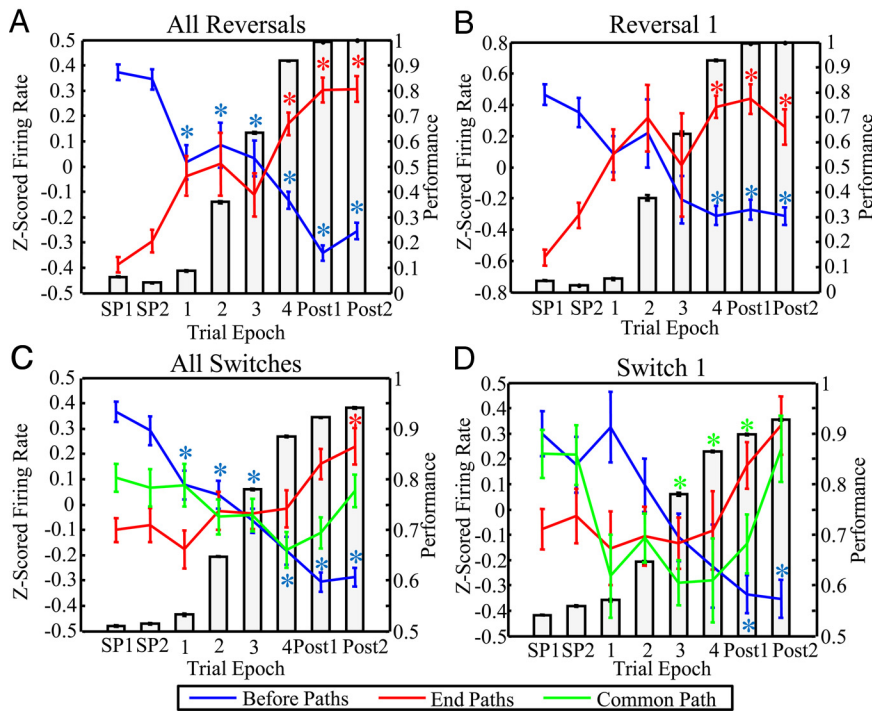


Figure 6. Dynamics of OFC activity across learning on the start arm. The solid lines indicate the Z-scored firing rate (\pm SEM) for path-selective cells recorded before, during, and after learning. Only activity in the start arm is included. The blue lines indicate the activity of cells specific to paths rewarded at the beginning of the testing day (before paths); the red lines indicate activity of cells specific to paths rewarded at the end of the day (end paths). The gray bars indicate the probability of correct responses (Smith et al., 2004). SP1 and SP2 indicate the first and second blocks of 10 trials of stable performance in the previously rewarded task; post1 and post2 indicate the first and second blocks of 10 trials after the rat reached criterion in the new task. In all reversals (A) and in reversal 1 (B), numerals 1–4 indicate trials with the following ranges of correct response probabilities: 1, 0–0.25; 2, 0.25–0.5; 3, 0.5–0.75; 4, 0.75–1.0. Similar diagrams show firing rate changes during switch learning in all switches (C) and for switch 1 (D). The green lines indicate the Z-scored activity in the start arm for cells that were specific to the common path during strategy switching. For C and D, the following ranges of correct response probabilities are indicated: 1, 0–0.60; 2, 0.6–0.7; 3, 0.7–0.85; 4, 0.85–1.0. The asterisk (*) indicates a significant difference between that epoch and SP2 (Bonferroni corrected *t* test, $p < 0.05$).

was Z-scored to the mean and SD during stable performance before the task change. During strategy switches, cells specific to the path rewarded throughout the testing day on both tasks—the common path—were excluded from this analysis. Two measures were calculated. First, we compared the trialwise average firing rate to the mean firing rate of all cells on each trial before and after the task changed (one-way Student’s *t* test vs mean before task change). Second, we compared the mean firing rate for all trials before the task change to the mean firing rate for all trials during the learning period before reaching criterion (Student’s *t* test).

OFC activity specific for rewarded paths declined immediately after the task was changed even when rats traversed identical journeys (Fig. 7). During all reversals and reversal 1, activity recorded during stable performance on rewarded paths declined immediately after the task changed, even as the path executed by the animal was constant. The decline in activity associated with the task change during all switches was smaller than during all reversals. This difference may be an artifact because of limited data; excluding the common path dramatically reduced the number of assessable trials. Alternatively, the size of the change in reward expectation during reversals and strategy switches also differs. During reversals, both paths rewarded at the beginning of the day are incorrect when the task is changed. By contrast, during strategy switching, only one of the paths rewarded at the beginning of the day is incorrect when the task is changed. This

difference in the number of errors the animal will experience if it continues to perform the old task may also explain the smaller decline in activity observed during all switches. This analysis shows that changes in path composition cannot explain changes in activity. OFC activity reflected the integrated reward history for paths rather than paths per se. The observed OFC firing dynamics during learning are consistent with the hypothesis that OFC codes the expectancy for particular paths by integrating reward history (Schoenbaum et al., 2009).

Path-selective cells show journey-dependent activity

Spatial reversal learning in the plus maze requires both the OFC and the hippocampus, so comparing the activity in the two structures might help illustrate how each structure contributes to guiding behavior. Hippocampal cells tend to fire in restricted regions of space. As a rat performs spatial memory tasks in the plus maze, individual hippocampal neurons are virtually silent until the rat enters the place field of the cell—a restricted spatial region where the firing rate increases dramatically. Beyond signaling location, however, place fields are modulated by the start and goal of spatial journeys (Ferbin-teanu and Shapiro, 2003). For example, place fields in the start arm are modulated by the pending goal arm (prospective coding), and place fields in the goal arm are modulated by the starting location of the journey (retrospective coding). Such

journey-dependent coding suggests that hippocampal activity may represent temporally extended behavioral episodes (Ferbin-teanu and Shapiro, 2003).

Although the behavioral correlates of OFC firing are not well described in terms of place fields, path-selective OFC activity can also be described operationally in terms of both prospective and retrospective coding. To quantify such coding by OFC neurons, we compared firing on either the start arm or goal arm in path-selective OFC cells during stable performance before and after learning. For example, if a path-selective cell was specific for the east-to-south path but not the east-to-north path, then the activity on the start arm was compared using a *t* test. Goal arm activity included all firing from the time from the first entry into the goal arm until the end of the trial. During reversals, 42% (53 of 125) of path-selective OFC cells fired differently in the start arm depending on the current goal, and 28% (18 of 63) of goal arm paths fired differently depending on where the path started. During strategy switches, 48% (27 of 56) of start arm paths and 56% (24 of 43) of goal arm paths showed similar journey-dependent coding (Fig. 8). The presence of journey-dependent coding in the OFC and the hippocampus suggests that the two regions may interact to guide behavior. Path coding by OFC and journey coding by hippocampal activity could reflect a bidirectional interaction between the two regions that integrate information about memory and reward.

Although both show journey-dependent coding, the behavioral correlates of OFC and hippocampal activity differ substantially along several important dimensions. For example, hippocampal neurons fire in small, spatially restricted place fields, whereas OFC neurons fire throughout the maze. To compare the spatial firing properties of OFC and hippocampal neurons, we compared the spatial distribution of activity of 650 OFC and 132 hippocampal neurons recorded in a similar plus maze task using identical methods (Ferbin-teanu and Shapiro, 2003). OFC firing fields were approximately fourfold larger ($61 \pm 3\%$ of visited grid units) than hippocampal fields ($16 \pm 2\%$ visited grid units; $t_{(476)} = 12.5$; $p < 0.001$). Therefore, if OFC neurons convey spatial signals, they are less informative than hippocampal neurons. Furthermore, because hippocampal place fields are usually limited to a subregion of one maze arm, they typically have either prospective or retrospective correlates, but seldom both. In contrast, because path-selective OFC neurons typically fired throughout the maze, a single cell could appear to have both prospective and retrospective correlates. Finally, the temporal dynamics of hippocampal journey-dependent coding and OFC path coding differ substantially. Hippocampal journey coding changes immediately during spatial reversals, whereas OFC path coding emerges gradually with learning. Thus, the presence of journey-dependent activity in the hippocampus and the OFC does not indicate that the two regions performed similar computations.

LFP oscillations in the OFC and hippocampus

To investigate potential interactions between the OFC and hippocampus during learning and spatial memory performance, we analyzed LFPs (0–100 Hz) recorded simultaneously in both structures. Spectrograms showed stable power within each structure during reversal and strategy switching. After normalization to the spectral power within a frequency band—delta (0.5–4 Hz), theta (5–12 Hz), beta (15–30 Hz), and gamma (30–80 Hz)—for that animal during SP, groups were compared with Bonferroni-corrected t test ($p < 0.05$ comparison between trial stages). Although statistically significant changes were observed in some groups—all reversals, reversal 1, all switches, and switch 1—the overall effect size was small, and the spectral power of no band changed more than $\pm 5\%$ of SP baseline during the testing day (data not shown).

By contrast, LFP coherence between the OFC and hippocampus varied significantly with learning. LFP coherence (coupled oscillation) between the OFC and hippocampus was calculated in the delta, theta, beta, or gamma bands during the period of stable performance before task change (SP), the learning period before achieving criterion (learning), and trials recorded after reaching criterion (post) for reversals and strategy switches (Fig. 9A). Al-

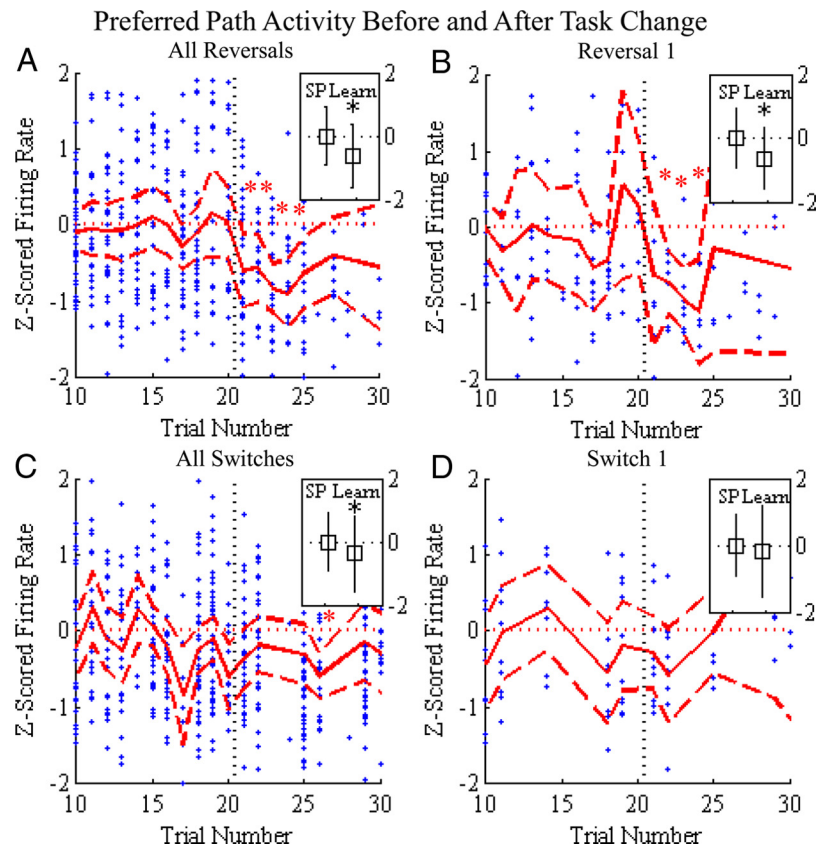


Figure 7. Firing rate dynamics before and after the task change. The figure shows normalized firing rates for individual cells (blue dots) recorded during identical behaviors. Firing rates for each cell was measured in its preferred path during the 10 trials before and after a task change (indicated by the vertical dotted lines). Only trials with more than five firing rate data points are included. The solid red line shows the moving average of the mean Z score firing rate; the dashed red lines show 95% confidence intervals. The inset shows the mean Z score firing rates (\pm SD) for all cells before the task change (SP) and during learning before criterion [learn; inset asterisk (*) shows Student's t test, $p < 0.05$]. **A**, Firing rate on preferred paths declined immediately after the task changed across all reversals [the asterisk (*) indicates significant decline compared with before the task change, one-way Student's t test, $p < 0.05$]. **B**, A similar decline in firing rates occurred in reversal 1. **C**, A mean firing rate reduction occurred during all switches (inset), but the trialwise firing rate change was smaller. **D**, The firing rate did not change during switch 1. The analysis shows that even when the rat followed the same path, OFC neuronal activity declined in correlation with falling reward probability when the task was changed.

though significant coherence was observed each analyzed band, only theta LFP coherence varied predictably during learning. Theta coherence was high during stable performance, declined significantly during reversals and the first switch, and remained lower than stable performance baseline even after rats reached criterion for both strategy switching and reversal (Fig. 9B).

LFP coherence may bind activity in different brain regions into distributed representations that guide behavior (Buzsáki and Chrobak, 1995; Buzsáki and Draguhn, 2004). Indeed, LFP theta coherence between the hippocampus and PL-IL correlates with behavior (Hyman et al., 2005; Siapas et al., 2005), and the theta phase of PL-IL spiking predicts spatial memory performance (Hyman et al., 2010). Thus, the high coherence between OFC and hippocampal LFPs during stable performance may indicate their integrated activity guides learning. Furthermore, the decline in theta LFP coherence may also have functional significance. The decline in theta coherence may reflect a necessary desynchronization for the brain regions to encode new task representations. Alternatively, it may reflect discordant predictions computed by the two structures (i.e., when recent memory contradicts an extensive reward history and new associations must be encoded). The lag between reaching criterion performance and the return of

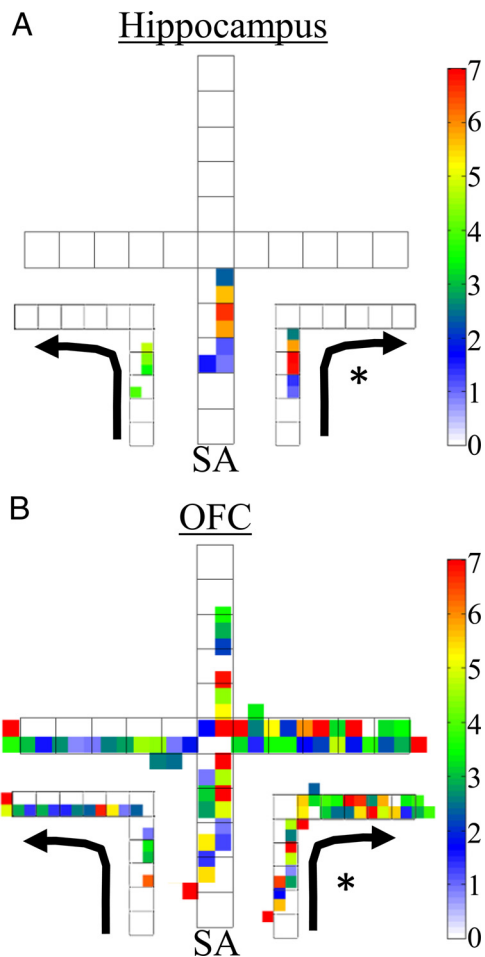


Figure 8. Prospective and retrospective coding in the OFC and hippocampus. *A*, An example of prospective coding for a cell with a place field on the start arm (SA) in the hippocampus (adapted from Ferbinteanu and Shapiro, 2003). Internal figure depicts a heat map (5×5 cm pixels) of the firing rate across the entire maze. External figures depict the firing rates on the two paths originating from that start arm. Units of the color bar are firing rate in hertz. *B*, An example of prospective coding for a cell in the OFC. The asterisk (*) indicates significantly greater activity on only the start arm between the two paths (comparison by Student’s *t* test, $p < 0.05$, between two groups).

high LFP coherence may thereby reflect the period of consolidation needed to align discordant task representations in multiple brain regions. This perspective is consistent with the relatively modest decline in theta coherence during strategy switching, which requires changing only one path and, thus, may produce less discordance between integrated reward probabilities and recent memory.

Discussion

A learning and memory task that required the OFC and hippocampus revealed that most individual OFC cells fired throughout the plus maze, and many fired selectively as rats traversed specific spatial paths. During learning, path-selective cell activity correlated closely with the probability that the animal would be rewarded on a particular path, consistent with the suggestion that the OFC integrates reward history and represents reward expectancy (Schoenbaum et al., 2009). Because spatial reversals require both the hippocampus and the OFC, successful behavior likely depends on their integrated activity. Although the mechanisms of such integration remain unknown, path-selective OFC cells showed activity that resembled journey-dependent coding in the

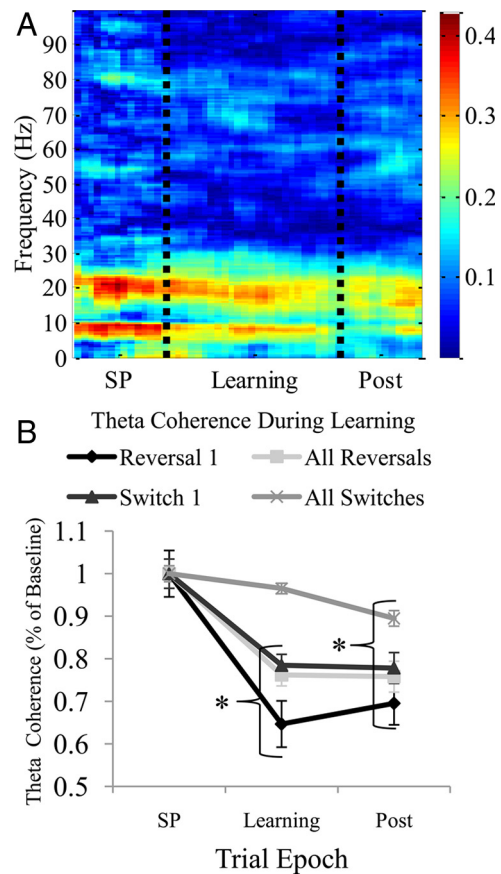


Figure 9. OFC–hippocampal theta LFP coherence. *A*, An example of a coherencegram for OFC and hippocampus LFPs during a first switch. Color map is the *R* value. The vertical lines indicate the division between the trial stages stable performance (SP), during learning (learning), and after achieving criterion (post). Note the high theta (5–12 Hz) and beta (15–30 Hz) coherence. *B*, Theta coherence declined significantly from stable performance (SP) to the learning period (learning) during the all reversals, reversal 1, and switch 1, but not the all switches. In all groups, theta coherence was lower in the postcriterion period (post). The asterisk (*) indicates significant reduction from SP baseline (Bonferroni-corrected *t* test, $p < 0.05$).

hippocampus. Moreover, LFPs in the two structures oscillated coherently in the theta band during stable memory performance, and this coherence declined during learning. Together, the results suggest that reward expectancy coded by the OFC is integrated with memory for spatial episodes coded by the hippocampus to support temporally extended, goal-directed behavior.

OFC and goal-directed action

The OFC activity described here includes more complex behavioral correlates than reported previously. Many previous studies report that OFC neurons respond selectively to complex cues (Schoenbaum and Eichenbaum, 1995a; Yonemori et al., 2000; Kadohisa et al., 2004; Rolls et al., 2006). However, selective response coding is reported by some studies (Feierstein et al., 2006; Furuyashiki et al., 2008) but not others (Tremblay and Schultz, 1999; Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006; Sul et al., 2010). Here, we observed complex selectivity for paths: extended sequences of actions leading to a goal. Three factors may account for these differing results. First, most previous studies used explicit cues such as odors or visual images as the predictors of reward. In the current experiment, the distal room cues were the discriminative stimuli, and these remained constant throughout testing even as response contingencies changed. Because distal stimuli predicted reward less reliably than memory

for rewarded paths, OFC activity may have become more closely related to responses than previously described. Second, the plus maze was larger than typical operant chambers. Previous studies tested responses limited to visual saccades or movements in small operant chambers, and the short duration and simplicity of the response might have limited the opportunity to observe response-correlated activity. Finally, most OFC neurons in these experiments fired a low mean rate (2–3 Hz). If the interval for statistical sampling of OFC activity was brief (e.g., 0.5 s maximum) (Sul et al., 2010), extended response coding may have been overlooked.

OFC encodes reward history during both reversals and strategy switching

This study compared the OFC activity in an OFC-dependent task, the first reversal, with later reversals and strategy switching tasks that are OFC independent. OFC neuronal activity was similar in both cases, firing in path-selective patterns that were continuous throughout single trials and correlated with performance. The similar neuronal activity during OFC-dependent and -independent learning suggests that OFC codes features common in both situations. The effects of previous training on the deficits observed during OFC inactivation also suggest that OFC coding is similar during strategy switching and reversal. During reversal, the OFC is required only transiently (McAlonan and Brown, 2003; Young and Shapiro, 2009). After initial training, the OFC is not required for subsequent reversals. Similarly, rats trained exclusively to perform repeated strategy switches learned reversals normally and were unimpaired by OFC inactivation (Young and Shapiro, 2009). Thus, some aspect of strategy switching may recapitulate important aspects of reversal, perhaps by engaging the similar activity in OFC.

The path-selective activity observed in these experiments during both reversal and strategy switching may explain this effect of training. Strategy switching and reversal share a common path in these experiments. An animal trained to switch strategies learned to perform a path tested subsequently during reversal. For example, an animal switched from “go left” to “go north” had been rewarded for performing the east-to-south, east-to-north, and west-to-north paths. When that animal was subsequently given reversal training to “go south,” it had to learn the west-to-south path for the first time, but it had already learned to associate reward with entering the south arm during “go left” (east-to-south) trials. If the OFC encoding associates paths with a reward history, OFC activity may have established reward expectancy for the east-to-south path in the previous context of “go left.” This representation may have improved later performance when the animal learned to “go south.” This model implies that the OFC

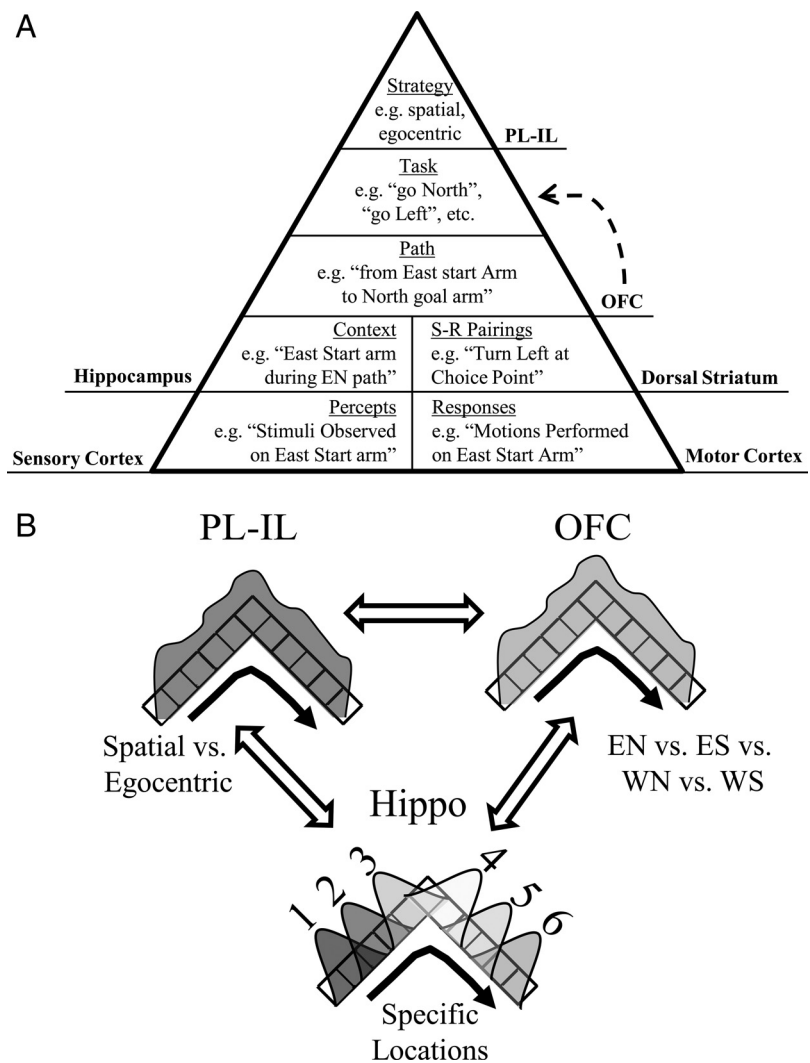


Figure 10. Hierarchy of action representation in the rat. *A*, Strategic representations such as spatial and egocentric are encoded in PL-IL. OFC encodes paths and possibly—in the combined activity of path-selective cells—tasks. Memory structures such as the hippocampus encode contextual information, whereas the dorsal striatum encodes stimulus–response pairings. Sensory cortices encode stimuli, whereas the motor cortex encodes responses. Higher levels of the hierarchy are created by combining the representations of lower levels. As the position in the hierarchy increases, the representations become more abstract (adapted from Rich and Shapiro, 2009). *B*, PL-IL, OFC, and hippocampus represent different levels of abstraction along rewarded paths. PL-IL represents a path only in the context of a particular strategy such as spatial. Ensembles in OFC are active continuously along rewarded paths. Individual hippocampal cells represent the component locations along the path. Interactions between the brain regions integrate these representations to guide behavior.

integrates reward history whether or not such coding is required for a given task. It also suggests that previously learned reward relationships can be combined to generate novel responses. From this view, OFC can facilitate learning by encoding reward contingencies for future use. Animals placed in a novel context can “draw on a library” of previously encoded reward representations to solve new tasks.

Interactions between the hippocampus and prefrontal networks

The tasks in the present experiment required both the OFC and hippocampus, allowing us to investigate functional interactions between the structures. First, by comparing their activity on the plus maze, we found that OFC showed activity that resembles prospective and retrospective coding in the hippocampus. Although OFC neurons did not fire in typical place fields, their firing rates discriminated different goal-directed journeys

through the same maze arms. Thus, coding in both regions was modulated by the significance of a location or action in the larger context of behavior. Second, simultaneous LFP recordings in the hippocampus and OFC showed high theta band coherence during stable performance that declined during learning. LFP coherence suggests that the interaction between the two regions may be crucial for adaptive behavior. LFP alignment may signal interaction between brain regions by indicating correlated firing and the potential for mutual information exchange (Buzsáki and Chrobak, 1995; Buzsáki and Draguhn, 2004). LFP theta coherence between the hippocampus and rat PL-IL has also been observed to correlate with spatial behavior (Hyman et al., 2005; Siapas et al., 2005). Together, these results suggest that additional experiments investigating prefrontal cortex (PFC)–hippocampal interaction through LFP coherence may reveal important mechanisms by which distributed networks interact to guide adaptive behavior.

Journey-dependent coding in both the OFC and hippocampus suggests several potential mechanisms by which the structures interact to combine reward expectancy coding by OFC with episodic-like coding by the hippocampus. First, OFC reward-related signals could modify hippocampal codes. From this view, OFC inputs to hippocampal neurons could modulate place field activity along rewarded paths and thereby facilitate the formation or expression of prospective and retrospective hippocampal coding. Conversely, hippocampal inputs could convey contextual information to OFC neurons, so that journeys become integrated with reward expectancy to produce path-selective firing. Thus, OFC path-selective activity could be the combined activity of many hippocampal place fields integrated with reward expectancy information. Finally, the interaction between the OFC and hippocampus may be bidirectional, so that OFC activity modulates hippocampal place fields along rewarded paths, while simultaneously integrating contextual input from the hippocampus with reward. Each of these potential mechanisms makes different predictions about the effects of OFC lesions on prospective and retrospective coding in the hippocampus, and vice versa. If OFC modulates hippocampal activity, then OFC lesions should alter prospective and retrospective coding in the hippocampus. In contrast, if the hippocampus influences OFC, then hippocampal lesions should alter path coding by OFC neurons. Finally, if the OFC and hippocampus interact bidirectionally, then lesions in each area should alter coding in the other. Ongoing experiments in our laboratory will test these predictions.

OFC as part of a network hierarchy organizing goal-directed action

LFP and unit activity suggests that OFC coding contributes to the organization of complex, temporally extended behavior as part of a larger network hierarchy, forming a subset of the “perception–action cycle” (Fuster, 1997). In this model, perceptions and actions are coded hierarchically in increasing levels of abstraction as information processing proceeds from primary to highest order association cortices (Fig. 10A). Sensory and motor cortex represent component stimuli and responses. Further from the periphery, structures such as the hippocampus and striatum combine these components into contextual information or stimulus–response pairings. The PFC interacts with all of these systems and may be viewed as the apex of a coding hierarchy. From this view, path coding by OFC is intermediate between strategy coding reported in the PL-IL of rats (Rich and Shapiro, 2009) and the IPFC (lateral prefrontal cortex) of primates (Asaad et al., 2000; Mansouri et al., 2006), and action sequences coded by the striatum (Berke et al., 2009) or goal-directed journeys coded by the hip-

pocampus (Ferbinteanu and Shapiro, 2003). Specifically, OFC may associate paths, sequences of goal-directed actions, with the likelihood that the goal will be rewarded. Although we did not observe prominent task-specific coding, OFC may also contribute to task coding through the combined activity of path-selective neurons. The path selectivity observed here, together with the cue selectivity observed elsewhere, suggests that the OFC may associate collections of percepts and actions with the likelihood of reward. In doing so, OFC could help form abstract expectancies, sets of stimulus–action–reward associations that allow the animal to solve new tasks rapidly.

We propose that the OFC and PL-IL bind task correlates at different levels of abstraction (Fig. 10B). PL-IL represents the highest level of abstraction by showing activity along rewarded paths only when particular strategies (i.e., spatial or egocentric) guide responding. OFC represents the integrated reward history associated with a continuous sequence of locations as paths. Multiple path representations can be combined into tasks. The hippocampus represents spatiotemporal sequences of locations along paths. Simultaneous activation of the ensembles of cells in all three structures thereby represents strategies, rewarded paths, tasks, and memory for spatial episodes. The mechanisms, however, that integrate these widely distributed networks remain unknown. The structures may operate in parallel and perform independent computations that ultimately converge on downstream motor control systems to guide behavior. Conversely, the structures may act on one another directly. For instance, path-selective OFC codes may selectively activate the cells in the hippocampus that represent the spatial sequences of the path, or hippocampal representations of recent events may activate OFC representations of reward expectancies. This model predicts that coactive ensembles recorded simultaneously in these brain regions should discriminate the strategy, reward expectancy, and episodic history that guide identical spatial behaviors. Future experiments using simultaneous recording will attempt to determine whether these structures interact directly or function in parallel and, more generally, how distributed networks integrate different cognitive features.

References

- Alvarez P, Eichenbaum H (2002) Representations of odors in the rat orbitofrontal cortex change during and after learning. *Behav Neurosci* 116:421–433.
- Asaad WF, Rainer G, Miller EK (2000) Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol* 84:451–459.
- Berke JD, Breck JT, Eichenbaum H (2009) Striatal versus hippocampal representations during win-stay maze performance. *J Neurophysiol* 101:1575–1587.
- Birrell JM, Brown VJ (2000) Medial frontal cortex mediates perceptual attentional set shifting in the rat. *J Neurosci* 20:4320–4324.
- Buzsáki G, Chrobak JJ (1995) Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr Opin Neurobiol* 5:504–510.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Diester I, Nieder A (2008) Complementary contributions of prefrontal neuron classes in abstract numerical categorization. *J Neurosci* 28:7737–7747.
- Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF (2006) Representation of spatial goals in rat orbitofrontal cortex. *Neuron* 51:495–507.
- Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. *Neuron* 40:1227–1239.
- Furuyashiki T, Holland PC, Gallagher M (2008) Rat orbitofrontal cortex separately encodes response and outcome information during performance of goal-directed behavior. *J Neurosci* 28:5127–5138.

- Fuster JM (1997) The prefrontal cortex: anatomy, physiology, and neuropsychology of the frontal lobe, Ed 3. Philadelphia: Lippincott-Raven.
- Hyman JM, Zilli EA, Paley AM, Hasselmo ME (2005) Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus* 15:739–749.
- Hyman JM, Zilli EA, Paley AM, Hasselmo ME (2010) Working memory performance correlates with prefrontal-hippocampal theta interactions but not with prefrontal neuron firing rates. *Front Integr Neurosci* 4:2.
- Kadohisa M, Rolls ET, Verhagen JV (2004) Orbitofrontal cortex: neuronal representation of oral temperature and capsaicin in addition to taste and texture. *Neuroscience* 127:207–221.
- Kim J, Ragozzino ME (2005) The involvement of the orbitofrontal cortex in learning under changing task contingencies. *Neurobiol Learn Mem* 83:125–133.
- Mansouri FA, Matsumoto K, Tanaka K (2006) Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin Card Sorting Test analog. *J Neurosci* 26:2745–2756.
- McAlonan K, Brown VJ (2003) Orbital prefrontal cortex mediates reversal learning and not attentional set shifting in the rat. *Behav Brain Res* 146:97–103.
- McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* 54:782–806.
- Mitra, P., and Bokil, H. (2008) Observed brain dynamics. Oxford; New York: Oxford UP.
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226.
- Paxinos G, Watson C (1998) The rat brain in stereotaxic coordinates, Ed 4. San Diego: Academic.
- Ragozzino ME, Wilcox C, Raso M, Kesner RP (1999a) Involvement of rodent prefrontal cortex subregions in strategy switching. *Behav Neurosci* 113:32–41.
- Ragozzino ME, Detrick S, Kesner RP (1999b) Involvement of the prelimbic-infralimbic areas of the rodent prefrontal cortex in behavioral flexibility for place and response learning. *J Neurosci* 19:4585–4594.
- Ramus SJ, Eichenbaum H (2000) Neural correlates of olfactory recognition memory in the rat orbitofrontal cortex. *J Neurosci* 20:8199–8208.
- Rich EL, Shapiro ML (2007) Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. *J Neurosci* 27:4747–4755.
- Rich EL, Shapiro M (2009) Rat prefrontal cortical neurons selectively code strategy switches. *J Neurosci* 29:7208–7219.
- Roesch MR, Stalnaker TA, Schoenbaum G (2007) Associative encoding in anterior piriform cortex versus orbitofrontal cortex during odor discrimination and reversal learning. *Cereb Cortex* 17:643–652.
- Rolls ET, Critchley HD, Browning AS, Inoue K (2006) Face-selective and auditory neurons in the primate orbitofrontal cortex. *Exp Brain Res* 170:74–87.
- Schoenbaum G, Eichenbaum H (1995a) Information coding in the rodent prefrontal cortex. I. Single-neuron activity in orbitofrontal cortex compared with that in pyriform cortex. *J Neurophysiol* 74:733–750.
- Schoenbaum G, Eichenbaum H (1995b) Information coding in the rodent prefrontal cortex. II. Ensemble activity in orbitofrontal cortex. *J Neurophysiol* 74:751–762.
- Schoenbaum G, Chiba AA, Gallagher M (1998) Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci* 1:155–159.
- Schoenbaum G, Chiba AA, Gallagher M (1999) Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J Neurosci* 19:1876–1884.
- Schoenbaum G, Chiba AA, Gallagher M (2000) Changes in functional connectivity in orbitofrontal cortex and basolateral amygdala during learning and reversal training. *J Neurosci* 20:5179–5189.
- Schoenbaum G, Nugent SL, Saddoris MP, Setlow B (2002) Orbitofrontal lesions in rats impair reversal but not acquisition of go, no-go odor discriminations. *Neuroreport* 13:885–890.
- Schoenbaum G, Setlow B, Nugent SL, Saddoris MP, Gallagher M (2003a) Lesions of orbitofrontal cortex and basolateral amygdala complex disrupt acquisition of odor-guided discriminations and reversals. *Learn Mem* 10:129–140.
- Schoenbaum G, Setlow B, Saddoris MP, Gallagher M (2003b) Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala. *Neuron* 39:855–867.
- Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK (2009) A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nat Rev Neurosci* 10:885–892.
- Siapas AG, Lubenov EV, Wilson MA (2005) Prefrontal phase locking to hippocampal theta oscillations. *Neuron* 46:141–151.
- Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, Brown EN (2004) Dynamic analysis of learning in behavioral experiments. *J Neurosci* 24:447–461.
- Stalnaker TA, Franz TM, Singh T, Schoenbaum G (2007) Basolateral amygdala lesions abolish orbitofrontal-dependent reversal impairments. *Neuron* 54:51–58.
- Sul JH, Kim H, Huh N, Lee D, Jung MW (2010) Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making. *Neuron* 66:449–460.
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. *Nature* 398:704–708.
- Wallis JD (2007) Orbitofrontal cortex and its contribution to decision-making. *Annu Rev Neurosci* 30:31–56.
- Wallis JD, Miller EK (2003) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci* 18:2069–2081.
- White NM, McDonald RJ (2002) Multiple parallel memory systems in the brain of the rat. *Neurobiol Learn Mem* 77:125–184.
- Wilson FA, O'Scalaidhe SP, Goldman-Rakic PS (1994) Functional synergism between putative γ -aminobutyrate-containing neurons and pyramidal neurons in prefrontal cortex. *Proc Natl Acad Sci U S A* 91:4009–4013.
- Yonemori M, Nishijo H, Uwano T, Tamura R, Furuta I, Kawasaki M, Takashima Y, Ono T (2000) Orbital cortex neuronal responses during an odor-based conditioned associative task in rats. *Neuroscience* 95:691–703.
- Young JJ, Shapiro ML (2009) Double dissociation and hierarchical organization of strategy switches and reversals in the rat PFC. *Behav Neurosci* 123:1028–1035.