Published in final edited form as:
Neuroimage. 2010 April 1; 50(2): 644–656. doi:10.1016/j.neuroimage.2009.11.083.

Dissociating the Contributions of Independent Corticostriatal Systems to Visual Categorization Learning Through the Use of Reinforcement Learning Modeling and Granger Causality Modeling

Carol A. Seger[1,3], **Erik J. Peterson**[1,3], **Corinna M. Cincotta**[1], **Dan Lopez-Paniagua**[1], and **Charles W. Anderson**[2,3]

[1]Department of Psychology, Colorado State University

[2]Department of Computer Science, Colorado State University

[3]Program in Molecular, Cellular and Integrative Neurosciences, Colorado State University

Abstract

We dissociated the contributions to learning of four corticostriatal "loops" (interacting striatal and cortical regions): *motor* (putamen and motor cortex), *visual (*posterior caudate and visual cortex), *executive* (anterior caudate and prefrontal cortex), and *motivational* (ventral striatum and ventromedial frontal cortex). Subjects learned to categorize individual repeated images into one of two arbitrary categories via trial and error. We identified (1) Regions sensitive to correct categorization, categorization learning, and feedback valence (2) Regions sensitive to prediction error (violation of feedback expectancy) and reward prediction (expected feedback associated with category response), using reinforcement learning modeling and (3) Directed influences between regions using Granger causality modeling. Each loop showed a unique pattern of sensitivity to each of these factors. Both the motor and visual loops were involved in acquisition of categorization ability: activity during correct categorization increased across learning, and was sensitive to reward prediction. However, the posterior caudate received directed influence from visual cortex, whereas the putamen exerted directed influence on motor cortex. The motivational and executive loops were involved in feedback processing: both regions were sensitive to feedback valence, which interacted with learning across scans. However, the motivational loop activity reflected prediction error, whereas executive loop activity reflected reward prediction, consistent with the executive loop role in integrating reward and action. Granger causality modeling found directed influences between striatal and cortical regions within each loop. Across loops, the motor loop exerted directed influence on the executive loop which is consistent with the role of the executive loop in integrating feedback with stimulus-response history.

Visual categorization is known to recruit the striatum and interacting cortical regions. However, it is still unclear which striatal regions are associated with specific aspects of categorization tasks, or how cortex and striatum interact during categorization. Visual categorization tasks

Address for Correspondence: Carol A. Seger Department of Psychology 1876 Campus Delivery Colorado State University Fort Collins, CO 80523 Voice: 970 491 3540 FAX: 970 491 1032 Carol.Seger@colostate.edu.

typically involve trial and error learning in which subjects view a stimulus, perform a motor response indicating the category membership of the stimulus, then receive feedback indicating whether the response was correct or not. This basic design of stimulus - response - feedback is common to a wide variety of tasks from the human and animal learning literatures; in addition to categorization, human tasks that follow this design include reward learning, rule learning, and decision making tasks. Nonhuman primate tasks include arbitrary visuomotor learning and conditional discrimination learning, and rodent tasks include instrumental conditioning and habit learning (see Seger (2009) for a review). Across all these tasks and species, the basal ganglia and their interacting cortical regions have been shown to be active during learning and necessary for learning to take place.

Our hypotheses about how cortex and striatum subserve categorization are based on the known anatomy of corticostriatal systems. Different regions of cortex project to different regions of the basal ganglia, and from there to the thalamus and back to cortex, forming independent corticostriatal loops. Following, we differentiate between four loops, which we refer to as the visual, motor, executive, and motivational loops. Information processing takes place both within particular corticostriatal loops, and by interactions between loops[1]. A primary function of the basal ganglia within corticostriatal loops is to balance inhibition and excitation of cortical representations in order to select the most appropriate one. The type of representations that the striatum acts upon depends on the region of cortex; for example, motor cortex typically represents movements, whereas lateral frontal cortex may represent cognitive processes or strategies. The selection role of the basal ganglia in motor and frontal cortex is consistent with the motor initiation deficits and cognitive set-shifting deficits, respectively, seen in basal ganglia disorders such as Parkinson's disease and Huntington's disease.

The visual loop connects visual cortex with the posterior caudate, particularly the body and tail regions. Large regions of inferior temporal and middle temporal cortex project to the posterior caudate These cortical regions are part of the ventral “what” system, and underlie visual object identification and integration of visual information with other sensory modalities and semantic representations. The posterior caudate is active during visual categorization in humans. Studies in nonhuman animals also find that the posterior caudate is important for learning in similar tasks. We hypothesize that in categorization the visual loop assists in selection of appropriate visual representations via recurrent connections to visual cortex, and in selection of responses via feed forward connections to the motor corticostriatal loop

The motor loop consists of the putamen and interacting regions of the primary motor cortex, primary somatosensory cortex, and premotor and supplementary motor cortex. Within the motor loop there is a gradient such that more anterior cortical regions (e.g., pre-SMA) project to the more anterior regions of the putamen, whereas more posterior cortical regions (e.g., primary motor cortex) project to more posterior regions of putamen. The motor loop is often active during visual categorization. We hypothesize that its role is in selecting appropriate category-related motor responses (e.g., the button press indicating category A).

The executive loop connects the anterior dorsal striatum, in particular the head of the caudate, with dorsal and lateral portions of the prefrontal cortex. This region is commonly activated during fMRI studies of visual categorization that use a block design. Activity during categorization learning in the head of the caudate is affected by the presence of feedback, and by feedback valence, typically resulting in greater activity for positive than negative feedback.

[1]Interactions between loops occur through several anatomical routes. Most notably, projections from striatum to basal ganglia output nuclei can “split”, so that return projections from the thalamus extend both to the originating cortical region and to additional cortical regions. In addition, DA projections to and from the VTA and SNc form an “ascending spiral” that allows striatal regions to affect regions that are relatively more lateral and superior.

We hypothesize that the executive loop is important for selection of learning strategies and switching between strategies on the basis of feedback.

The motivational loop connects ventral striatum (consisting of the nucleus accumbens and ventral regions of the caudate and putamen) with ventromedial frontal cortex, hippocampus, and amygdala. Fewer studies have investigated the role played by ventral striatum in categorization, but previous studies that do report activity in this area have linked it to feedback or error processing, or to categorization uncertainty. Although both the executive (head of the caudate) and motivational (ventral striatum) loops are sensitive to feedback, they differ in that the ventral striatum is sensitive to unexpected reward regardless of whether it has anything to do with the organism's behavior, whereas the head of the caudate is sensitive to feedback only when it comes in response to an action the subject performed. Furthermore, the ventral striatum is associated with sensitivity to feedback regardless of whether the subject learns, but dorsal striatal activity is specific to learners.

In light of the hypothesized role of the visual loop in visual category selection we predicted that the posterior caudate would (1) be active during correctly categorized trials and increase in activity across the time course of learning in the general linear model analyses (2) be associated with reward prediction rather than prediction error in the reinforcement learning analyses, and (3) that Granger causality modeling would find directed influence from visual cortex to the posterior caudate. We hypothesized that the motor loop subserves selection of the motor response indicating category membership and accordingly predicted that the putamen would (1) be active during correctly categorized trials and would increase in activity across trials as subjects learned about appropriate responses (2) be associated with reward prediction rather than prediction error, and (3) exert directed influence onto motor cortical regions consistent with its role in motor selection. Note that our predictions for the visual and motor loops are similar for the general linear model and reinforcement learning analyses, but differ for the Granger causality modeling analysis.

We hypothesized that the executive loop and motivational loops would both contribute to feedback processing during categorization, but that the executive loop would additionally integrate feedback with learning. Our predictions for the current study were that the striatal regions in both loops would be sensitive to feedback valence in the general linear model analysis, as was found in our previous research (Seger and Cincotta, 2005). However, we predicted that the reinforcement learning analysis would dissociate these two loops: the ventral striatum reflects reward prediction regardless of the subjects behavior, and therefore should be related to prediction error. The head of the caudate is concerned with the integration of reward expectancy and action, and should be related to reward prediction. Finally we predict the executive loop, due to its role in integrating feedback with the stimulus and response made on the trial, should receive directed influence from the motor and visual corticostriatal loops.

Materials and Methods

Subjects

Eleven healthy adult members (6 male, 5 female) of the Stanford community participated as paid subjects. All subjects were right handed and spoke fluent English. One subject's Scan 3 behavioral data was lost, so behavioral and functional analyses of this scan include only 10 subjects.

Tasks and Procedure

Categorization task

During the categorization task subjects were presented with individual images of faces or houses. They were instructed to learn which arbitrary category, which were given the labels 1 or 2, each image belonged to via trial and error. Category membership was not a function of whether the image was a face or a house.

Sixteen visual images were used as stimuli, eight of faces and eight of houses. House pictures were selected from Northern Colorado region real estate websites; they included houses of various colors, stories, and architectural styles. Face pictures were obtained from the Face Database available through the Productive Aging Laboratory website (agingmind.cns.uiuc.edu/facedb/). The faces were chosen from the collection of images of male Caucasians between the ages of 18 and 29 with neutral facial expressions. For the houses and faces separately, the 8 stimuli were assigned to categories as follows. Six had a deterministic relationship with category: 3 were arbitrarily assigned to Category 1 and 3 to Category 2. Care was taken to ensure that stimuli within an category did not share any unintended similarities (e.g., we did not assign brick houses to one category, and wood frame houses to the other category). Four stimuli (two faces and two houses) were randomly assigned to category, such that each was in Category 1 on half of the trials and Category 2 on half of the trials.

Each categorization trial consisted of the following sequence of events. A stimulus (face or house) was displayed first (1700 ms) followed by a blank screen (100 ms) and the feedback screen (500 ms). Between trials subjects viewed a blank screen that was presented for a minimum of 200 ms, plus a variable “jittered” amount of time of 0 ms (half the trials), 1500 ms (one quarter of trials), or 3500 ms (one quarter of trials). Each block of 8 trials included 4 trials with 0 ms jitter, 2 trials with 1500 ms jitter, and 2 trials with 3000 ms jitter. The order of jitter in each block was randomized. Participants were to make a response during the stimulus display by pressing one of two stimulus buttons with the index or middle finger of the right hand. The feedback consisted of the words “Right!!!” in blue, “Wrong” in red, or “No response made” in red.

There were a total of 720 trials: 480 categorization trials (240 houses and 240 faces) and 240 baseline trials. A mixed block-event related design was utilized. Face and House stimuli were separated on a blockwise basis. Each block consisted of 8 categorization trials (faces or houses) and 4 baseline trials. Each categorization block was 30 seconds in length and each baseline block was 10 seconds in length. Baseline trials consisted of a plus sign for 2000 ms and a blank screen for 500 ms. Blocks of faces alternated with blocks of houses with baseline blocks in between each categorization block. There were 120 blocks total; 60 categorization and 60 baseline. The 120 blocks were broken down into 4 scans of 30 blocks each. Note that even though subjects completed 4 scans, the results from only the first 3 scans are presented here due to loss of data from the 4th scan from several subjects.

Localizer task

During the localizer task subjects viewed images of faces, houses, animals, and tools. There were a total of 152 stimuli presented in the localizer task, 38 each of faces, houses, animals, and tools. House and face pictures were taken from the same websites as in the Categorization task. For the localizer task, faces images were selected on the basis of being young adults (between the ages of 18 and 29) with neutral expressions; however, for the localizer task we included equal numbers of male and female faces, and included all the ethnicities represented in the face database (which has separate collections of images of Caucasians, African-Americans, Asian-Americans, and Latinos). Tool pictures were of both hand and power tools.

Animal pictures were of various animals including reptiles, birds, and mammals. Most animal pictures were profiles, with no pictures emphasizing animal faces.

In order to ensure attention to the images, subjects performed a one-back task: on each trial they had to decide whether the current stimulus matched or did not match the immediately preceding stimulus. Of the 200 trials, 48 were matches, 12 for each of the stimulus types. Each trial consisted of a stimulus display for 1500 ms and a blank screen for 500 ms. Participants were to make a response during the stimulus display. There was no feedback given. There were 20 blocks consisting of 10 trials per block. Each block was 20 s. There were five blocks of each stimulus type, presented in the following repeating order: Faces, Houses, Animals, Tools.

MRI image acquisition

Imaging was performed with a custom-built whole head coil in a 3.0 Tesla MRI Signa LX Horizon Echospeed (General Electric Medical Systems, Milwaukee, WI). Head movement was minimized for participants using a “bite-bar” formed with the participant's dental impression. In addition to the functional scans, three anatomical scans were performed: a coronal T1-weighted localizer scan, a three-dimensional high-resolution T1-weighted spoiled gradient echo scan with 124 contiguous 1.5 mm slices [minimum full echo time (TE), 30 degrees flip angle, 24 cm field of view, 256 × 256 acquisition matrix], and a inplane anatomical T1-weighted spin-echo scan with 22 contiguous 5 mm axial slices [minimum full TE; 500 ms TR, 24 cm field of view, 256 × 256 acquisition matrix]. Functional scanning was performed using a T2* sensitive gradient echo spiral in-out pulse sequence (Glover & Law 2001; Preston et al., 2004) [30 ms TE; 1500 ms TR; 65 degree flip angle; 24 cm field of view; 64 × 64 acquisition matrix] of the same 22 contiguous 5 mm axial slices as the inplane images.

Stimuli were presented using a magnet-compatible projector (Resonance Technology, Inc., Van Nuys, CA) that back-projects visual images onto a screen mounted above the participant's head. E-prime software (Psychology Software Tools, Pittsburgh, PA) running on a personal computer was used to generate visual stimuli and control experimental parameters. Responses were obtained using a magnet-compatible response system.

Data analyses

Image preprocessing

Image analysis was performed using Brain Voyager QX 1.1.9 (Brain Innovation, Maastricht, The Netherlands). The functional data were first subjected to preprocessing, consisting of three dimensional motion correction, slice scan time correction, and temporal data smoothing with a high pass filter of 3 cycles in the time course and linear trend removal. Each participant's high resolution anatomical image was normalized to Tailarach space (Tailarach & Tournoux, 1988). The normalization process in Brain Voyager consists of two steps, an initial rigid body translation into the AC-PC plane, followed by an elastic deformation into the standard space performed on 12 individual subvolumes. The resulting transformations were applied to the participant's functional image volumes to form volume time course representations to be used in subsequent statistical analyses. Finally, the volume time course representations were spatially smoothed with a Gaussian kernel, full width at half maximum of 8.0 mm.

General Linear Model analyses

Brain Voyager was used to analyze contrasts between conditions, using the general linear model with separate subject predictors and random effects analysis. Only contrasts with potential theoretical interest were analyzed. The False Discovery Rate method with a threshold of $q < .05$ was used to control for proportion of false positives; $q < .05$ results in no more than 5% of the voxels being significantly active by chance. Our examination of correct

categorization compared activity during correctly categorized deterministic trials with an implicit baseline consisting in effect of all other time points, not just those occurring during the baseline blocks. In order to perform analyses of changes in activity in response to feedback valence across scans, we performed an exploratory interaction analyses; due to limitations in the Brain Voyager software package, this required treating both subject and scan as fixed effects. We used the false discovery rate with a threshold of $q < .01$ to control for the proportion of false positives, and examined only striatal regions for which we had a preexisting hypothesis that activity would be affected by feedback valence. The analyses examined 2 × 2 interactions with scan (scan 1 vs. scan 3) as one factor and feedback valence (random trials receiving positive feedback vs. random trials receiving negative feedback) as the other factor.

Reinforcement learning modeling analyses

The reward prediction error (δ; equation 1) and reward prediction (*Q*; equation 2), were calculated using SARSA (State-Action-Reward-State-Action) reinforcement learning modeling technique, with no discount factors or eligibility trace.

For modeling purposes, each trial was considered to contain two time points (*t*, *t+1*), with the first encompassing stimulus onset and category response, and the second reward delivery. Reward, represented as r_t, was `1' for positive feedback, `0' for negative. There were thus a total of 64 possible *Q* values that we calculated for each trial, corresponding to the 16 possible visual stimuli (that defined the state, *s*), the two categories (synonymous with action, *a*) and the two time points within a trial. *Q* was initially set to 0 and was updated on a trial by trial basis depending on the stimuli presented and the category the subject selected (equation 2). The optimal learning rate (α) was determined on a subject-by-subject basis. Each trial's Q value, for a given α (range: 0.01–1, increments of 0.01), was transformed to an action selection probability via the softmax distribution (a standard method in Reinforcement Learning;. The optimal values of α and β (temperature, which controls the probability bias of softmax; range: 0.01–5, increments of 0.01) maximized the negative log-likelihood (*L*; $L = -1*\text{sum}(\ln(P(Q,t)))$. Trials in which the subject made no response were excluded from the model. For regression analysis *Q* and δ, from time t, were combined by stimulus type (face or house) and condition (deterministic or random), convolved with the conical HRF and converted to Brain Voyager compatible text files using a set of in-house perl scripts (version 5.8.8; www.perl.org). Modeling and HRF convolution was done using the R programming language (version 2.7.2; www.R-project.org). The *Q* and δ regressors were decorrelated using the orthvec() routine in the BVQX tools MATLAB toolbox (version 0.8a; http://wiki.brainvoyager.com/BVQXtools). Note that for prediction error (δ) analyses we used predictors from both deterministic and random stimuli but report only results from Scan 1 and for House stimuli; deterministic learning was quite rapid, resulting in prediction error values of (near) zero for later scans for house stimuli (see Figure 4) and for all three scans for face stimuli. For reward prediction (*Q)* we also limited our analysis to Scan 1 but include both House and Face stimuli.

Granger Causality Modeling analyses

Granger causality modeling analyses were performed using the Granger Causality Mapping plug-in within Brain Voyager. The algorithms implemented in this program are described in more detail by Goebel, Roebroeck, and colleagues. Briefly, Granger causality modeling is used to examine directed influences between brain regions. A small reference region, referred to as a seed region, is defined as a set of voxels. The average time course of activation across these voxels is compared to the time course of activation in all other voxels in the brain via a vector autoregressive algorithm. This allows for the identification and mapping of two kinds of directed effect: regions that are influenced by the reference region (that is, regions in which later activation can be predicted by earlier activation in the reference region, abbreviated as

Ref2Vox), and regions that influence the reference region (that is, regions in which earlier activation can predict later activation in the reference region, abbreviated as Vox2Ref). Granger causality modeling does not require any assumptions about the underlying anatomical or functional connectivity between regions. argue that the most unbiased measure of directed influence is achieved by taking the difference between the two measures of directed influence, Ref2Vox-Vox2Ref, in which a positive difference is interpreted as directed influence onto the seed region, and a negative difference as directed influence from the seed region. The Ref2Vox-Vox2Ref measurement is calculated separately at each voxel outside of the reference region, and the voxelwise measures are used to form maps that are overlaid on the anatomical images. Maps were defined on an individual subject basis, then combined across subjects via a voxelwise t-test identifying voxels in which activity was significantly different from zero. The analyses were performed using preprocessed data, which included spatial smoothing.

Two sets of Granger causality modeling analyses were performed. The first set examined directed influence from visual cortical regions. Face and house category specific regions in the inferior temporal cortex were identified individually for each subject using the localizer scan. The fusiform face area (FFA) was identified with a contrast comparing faces with houses and tools; animals were excluded because some of the images had faces visible. The parahippocampal place area (PPA) was identified with a contrast comparing houses with the other three categories (faces, animals, and tools). Across subjects the average Tailarach coordinates for the center voxel of the FFA were x = 41.6 (range: 38–46), y = −40.6 (range: −33 to −53), z = −18.5 (range: −8 to −23). Overall, this is a good match with previous studies reporting a location of the FFA (Gobbini & Haxby 2006), although in some subjects the identified region was somewhat more anterior than is typical. However, recent research has found in humans that face selective patches often extend anterior to the FFA. For the PPA the average was x = 28 (range: 25 – 36), y = −38 (range: −31 to −47) and z = −7 (range 2 to −14). This location is consistent with previous studies of the PPA. The FFA and PPA regions of interest were defined using region growing algorithms implemented in Brain Voyager. The size of each ROI was constrained to be between 100 and 200 voxels.

The second set of GCM analyses examined directed influence onto or from striatal regions. Striatal ROIs were identified on a cross-subject basis, based on the analysis comparing deterministic correct trials with baseline (see Table 1 for coordinates). We examined regions of the right and left putamen and right and left anterior caudate. The size of each ROI was constrained to be between 100 and 200 voxels. With both the localizer based reference regions (FFA and PPA) and striatal reference regions we calculated positive and negative influence during baseline blocks, face categorization blocks, house categorization blocks, and both house and face blocks combined.

Results

Behavioral results

As shown in Figure 1, across scans subjects learned to categorize the deterministic face and house stimuli, whereas performance on the random stimuli remained at chance levels as would be expected. This pattern was confirmed using a 3 × 3 within subjects ANOVA with factors of stimulus type (Faces, Houses, Random) and Scan (1–3), which found an interaction between scan and stimulus type, $F(4, 36) = 6.7$, $p < .005$, a main effect of stimulus type, $F(2, 18) = 43.8$, $p < .001$, and a main effect of scan, $F(2, 18) = 8.9$, $p < .005$.

Focusing on the deterministic stimuli, it appears that subjects were faster initially at learning the face stimuli, but were ultimately able to learn both stimulus types. This observation is confirmed by paired t-tests, which found that in scan 1, subjects were more accurate categorizing Faces (M =77.3) than Houses (M= 64.9), $t(10) = 2.9$, $p < .05$, but by scan 3

performance was equivalent on both stimulus types, $t(9) < 1.0$ (Faces: $M = 84.6$; Houses: $M = 84.1$).

Categorization related activation

Categorization across scans—The primary contrast for examining correct categorization is activity during deterministic correct trials in comparison with baseline. This contrast revealed a wide network of activity in cortical and striatal regions, as shown in Table 1. Within the striatum, activity extended throughout the caudate nucleus, including the head, body, and tail regions, and to a region of the anterior putamen. Within the frontal cortex there were regions of activation in bilateral inferior frontal gyri that extended to the anterior insula; activity in this region has been associated with a wide variety of visual categorization tasks in fMRI studies, and is particularly associated with early stages of learning, and for novel in comparison with familiar associations (Boettiger and D'Esposito, 2005). Lesions to the homologous area in monkeys eliminates the ability to learn new abstract rules. There was also activation in right premotor cortex; previous fMRI and monkey research has linked this region to both acquisition and performance of learned associations. Broadly, this region has been associated with cognitive aspects of response selection: it is generally more active for response selection than for simple response execution, and overlaps with regions containing "mirror neurons" (Chouinard & Paus 2006). Electrophysiological studies of monkeys have found that individual lateral premotor neurons are tuned to perceptual decision, action, and monitoring outcomes.

In addition to frontal regions, there was widespread activity within extrastriate visual areas, the thalamus, the midbrain, and the cerebellum. A region of medial prefrontal cortex exhibited reduced activation during correctly categorized deterministic trials. This region is commonly activated during "rest" or simple baseline tasks and has been associated with internally focused attention and attention to ones mental states.

An additional contrast compared activation between correctly and incorrectly categorized deterministic trials. This contrast revealed two regions of activation in the striatum: one in the right body of the caudate at x = 19 y = 3, z = 22, which overlapped with a region of activation in the deterministic correct vs. baseline contrast, and one in the left lateral putamen, at x = −28 y = 3, z = 2, which was a novel region of activity not present in the deterministic correct vs. baseline contrast at the statistical threshold.

Changes in striatum across the time course of categorization learning—We examined how activity changed across the time course of learning within striatal regions of interest (ROIs). We created ROIs for each of the regions within the striatum that were active in the deterministic correct vs. baseline comparison (left head of the caudate, bilateral body of the caudate, and bilateral tail of the caudate; see Table 1 for center coordinates of these ROIs), plus the putamen region active in the deterministic correct vs. incorrect contrast described above. We then calculated beta values for the deterministic correct predictor within each ROI separately for each scan. As shown in Figure 2, the activity in the anterior caudate (head and body ROIs) was stable across scans. In contrast, activity within the posterior caudate (tail ROIs) and putamen increased across scans. Examination of activity in scans 1 through 3 separately confirmed this pattern of results: the tail of the caudate and putamen regions were significantly activated in scans 2 and 3 but not scan 1, whereas the head and body regions of the caudate were significantly activated in all three scans.

Feedback related activity

We examined the effects of feedback through a contrast comparing random trials on which positive feedback was received with random trials on which negative feedback was received.

Previous research in our lab using similar tasks found greater activity in the head of the caudate for positive than negative feedback (Seger and Cincotta, 2005). However, for this study we found no feedback valence differences when all three scans were included in the analysis. We considered the possibility that an interaction between learning and feedback processing might preclude finding a main effect of feedback valence across scans; in our earlier study the differences between positive and negative feedback decreased substantially as learning progressed and were not present in the final block. Therefore, we examined the interaction between feedback valence and scan, as described in Methods, above, and identified four regions of the striatum affected by the interaction: the right and left ventral striatum, the right head of the caudate and the right tail of the caudate. As shown in Figure 3, in the ventral striatum there was significantly more activity for positive than negative feedback in scan 1, which was followed by a much smaller activation difference in scans 2 and 3. In the head and tail of the caudate nucleus, a cross over interaction was found: in scan 1, there was more activity in response to positive feedback than negative, but by scan 3 there was more activity in response to negative feedback than positive.

Previous research has found a general tendency for greater activity to positive than negative feedback during early learning or in randomly structured "gambling" tasks. However, some studies have found equivalent activity, or even greater activity for negative feedback than positive feedback. The effect of feedback valence may depend on factors such as the degree of learning and the utility of the feedback, with positive feedback being more useful early in learning, and negative feedback more useful as learning progresses and errors are rarer.

Reinforcement Learning model based analyses

The prediction error (δ) and reward prediction (*Q*) functions across scans for a typical subject are shown in Figure 4. Within the striatum, prediction error predicted activity in the ventral striatum bilaterally, as shown in Figure 5. Within cortical regions, prediction error predicted activity most notably in several prefrontal cortical regions, and in the fusiform and parahippocampal gyri (see Table 2). These latter locations are consistent with regions involved in processing the visual stimuli.

Reward prediction (*Q*) predicted activity in the putamen and anterior caudate nucleus. These regions are clearly dorsal to the regions of ventral striatum predicted by the error function, as shown in Figure 5. Within cortex, reward prediction predicted activity in three notable groups of cortical regions (see Table 3): high level visual cortical regions of the fusiform and inferior temporal gyri, the bilateral supramarginal gyrus, and right premotor cortex.

Granger Causality Modeling

Directed influence from fusiform face area and parahippocampal place area— For this analysis, we examined positive directed influence from the reference regions to other brain regions. As described in Methods, the fusiform face area (FFA) and parahippocampal place area (PPA) were defined on an individual subject level based on activity during the independent locator scan. Activity in the FFA during face categorization predicted activity in the posterior caudate, as shown in Figure 6. The pattern using the PPA as the reference region was similar: activity in the PPA predicted activity in the posterior caudate nucleus. However, the directed influences between FFA and caudate, and PPA and caudate were not specific to face and house categorization blocks, respectively: activity was similar across both types categorization blocks, and t-tests comparing face and house blocks found no significant differences across subjects. Face and house blocks were combined for follow-up analyses. Both the PPA and FFA reference regions also had directed influence on other visual regions during categorization. In addition, the PPA but not the FFA showed directed influence to bilateral inferior frontal cortex / anterior insula and to the cerebellum.

In addition to examining directed influence during the categorization blocks, we also examined influences during the baseline blocks. There were no directed influences between FFA or PPA and other regions during the baseline blocks. This implies that the directed influences identified during categorization blocks are not due solely to intrinsic connectivity independent of cognitive task.

Directed influences to and from striatal regions—We also examined directed influence with striatal regions (anterior caudate and putamen) as reference regions. As with the PPA and FFA, no differences were found between house and face categorization blocks, so we collapsed across stimulus type for these analyses. These patterns of directed influence were specific to categorization blocks and did not occur during baseline blocks.

One salient pattern, illustrated in Figure 7 and with coordinates given in Table 4, was that striatal regions exerted directed influence on other regions of the striatum. Both right and left putamen showed positive directed influence on other regions of the putamen bilaterally and on all regions of the caudate nucleus. Both left and right anterior caudate nucleus received directed influence from the putamen, and exerted directed influence on other regions of the caudate bilaterally.

The putamen and anterior caudate had differing patterns of directed influence with cortical regions. The putamen showed positive directed influence onto medial parietal and medial frontal (anterior cingulate and SMA), and motor regions in the left precentral gyrus. The directed influence to the SMA is consistent with who found that coherence between putamen and supplementary motor area increased as subjects learned to perform different responses in different cued contexts. The putamen did not receive directed influence from any cortical regions. The anterior caudate received directed influence from many of the same regions that the putamen showed directed influence to, including right premotor, medial prefrontal regions, and right parietal cortex; these are illustrated in Figure 8. Anterior caudate exerted directed influence on primary visual cortex and the cerebellum.

These analyses show an overall pattern of directed influence across time in which the putamen influences motor and executive cortical regions, which in turn influence the anterior caudate. This pattern is consistent with the hypothesis that the putamen is involved in earlier categorization response selection and execution, whereas the anterior the caudate is involved in later feedback related functions.

Discussion

The results support our hypotheses about the roles played in visual categorization by each corticostriatal loop. Consistent with our hypothesis that the visual loop plays a role in visual stimulus processing necessary for categorization, we found (a) posterior caudate (bilateral tail of the caudate) recruitment during correct categorization that increased across scans as subjects learned, (b) directed influence from visual cortical regions onto the posterior caudate during categorization, and (c) visual cortex activity during correct categorization that was also predicted by both prediction error and reward prediction. The motor loop was hypothesized to contribute to learning to select appropriate motor responses to indicate category membership; consistent with this hypothesis, (a) the putamen was more active during correct categorization trials than incorrect trials, with this activity increasing across scans as learning progressed, (b) activity in the putamen was predicted by reward prediction, and (c) the putamen exerted a directed influence on ventral lateral premotor cortex and the supplementary motor area.

We predicted that both the motivational loop and executive loop would be involved in feedback processing, but would differ in what aspects of feedback they are sensitive to. Both ventral

striatum (motivational loop) and the head of the caudate (executive loop) showed an interaction between positive and negative valence across scans, but the shape of the interaction differed in the two regions; in the head of the caudate there was a cross-over interaction such that by the third scan activity was higher for negative feedback than positive. The ventral striatum and head of the caudate also differed in sensitivity to learning parameters from the reinforcement learning analysis. The ventral striatum was sensitive to prediction error, whereas the head of the caudate was sensitive to reward prediction. This pattern is consistent with previous research finding that ventral striatum is sensitive to reward in a classical conditioning task, whereas the dorsal striatum is important for integrating reward with behavioral choice.

Granger Causality Modeling

Granger Causality modeling proved useful in elucidating which neural regions are interacting with each other, and the directionality of these influences in time. We were able to confirm several important directed influences that were predicted by our theory. First, within the visual loop we have evidence that visual cortex exerts directed influence on the posterior caudate nucleus. Within the motor loop, we have evidence that the putamen interacts with ventrolateral premotor and supplementary motor cortices, exerting directed influence on each of them. Within the executive loop we found that the anterior caudate exerts directed influence on other regions of the caudate, and received directed influence from the putamen and frontal lobe regions. Thus, the executive loop can be seen as bridging between the motivational loop (ventral striatum) and the visual and motor loops (posterior caudate / putamen) through its integration of feedback processing with action plan updating.

Across loops, we found a pattern in which the putamen exerts directed influence on the caudate, which is consistent with the gradient of interactions within the striatum that progress from ventral-anterior-medial to dorsal-posterior-lateral regions (Haber, 2003). This pattern of directed influences also fits our predictions for how regions would interact across the time course of categorization. The typical sequence in a categorization task trial is first to view the stimulus and decide on the category, then select and execute a response, and finally, receive feedback and use it to updating the category representation. These three phases of a categorization trial map to the directed influences we found from visual cortex to posterior caudate (stimulus categorization), putamen to premotor cortex and SMA (response selection), and the SMA and premotor to caudate (feedback processing).

Functional MRI is limited in its temporal resolution, but electrophysiological studies have been used to examine the behavior of individual neurons in corticostriatal regions, and to associated their activity with stimulus, response, and feedback phases of a learning trial. Studies overall typically find more activity for putamen neurons around the time of the response, and caudate neurons at the time of reward receipt (Lau and Glimcher, 2007; Williams and Eskandar, 2006). However, neuron types are typically mixed within any particular region. For example, Lau and Glimcher (2008) found separate populations of cells in the body of the caudate tuned for action values (active before a response was made), and tuned for choice value (active after movement execution). Even among cells active after a movement was executed, some coded for the reward value of the movement, and some for specific features (e.g. direction) of the movement. Individual neurons also differ in terms of their changes across the time course of learning, with some cells increasing in firing rate, and others in the same region decreasing in firing rate.

Granger causality modeling does have some limitations that should be kept in mind. First, directed influence is a measure of whether previous time course values in one region can predict values in another region. This leaves open the possibility that some directed influences are not straight from the seed region to the target region, but instead pass through intermediate regions. Since Granger causality modeling is typically performed over blocks of many seconds

containing several complete acquisitions of brain slices, it is possible for events at the very beginning of the block to significantly affect events seconds later at the end of the block. Second, Granger causality modeling does not incorporate a model of the hemodynamic response function, and as a result it is possible for interregional variation in hemodynamic properties to result in misleading temporal precedence information.

One puzzling result from the Granger causality modeling analysis is that the directed influence from putative face and house selective regions (FFA and PPA) did not differ between face blocks, in which all the stimuli were faces, and house blocks, in which all the stimuli were houses. One possibility is that visual cortex projects to the posterior caudate in a widespread manner, and that all visual cortical activity will affect striatal activity. This would imply that the striatum does not have a separate representation of the same visual object classes that are found in the inferior temporal cortex. Alternatively, visual object class information may be kept separate in the striatum, but the relatively coarse spatial resolution of fMRI precludes us from detecting this pattern. Anatomically, the projections of cortex to striatum are relatively diffuse (each projection neuron from cortex makes few synapses onto a few striatal modules organized longitudinally down the caudate or putamen) and projections from individual cortical columns often project to common striatal units. This pattern of projections implies that visual cortical information is kept separate at least at the level of individual striatal spiny cells. Future progress on this question will require improving the resolution of the fMRI scanning, using alternative techniques for identifying distributed voxels that participate in a representation (such as multi voxel pattern analysis techniques), and convergent research from single unit recording in nonhuman animals.

Reinforcement Learning Modeling

We used reinforcement learning methods to identify regions that were sensitive to prediction error and reward prediction. There are many different ways of calculating prediction error; the most popular are temporal difference and Rescorla-Wagner rules. In the present study we examined signed prediction error and found that this predicted activity in the ventral striatum, consistent with previous studies.

In addition to prediction error, reinforcement learning modeling approaches include methods for estimating how prediction error is integrated with action selection. We used SARSA to estimate reward prediction, the total predicted reward associated with choosing the categorical response for the stimulus. We found that reward prediction predicted activity across the dorsal striatum, including the anterior caudate and putamen. Previous fMRI research has not used SARSA, or applied reinforcement learning to categorization learning tasks, but similar methods such as actor-critic and Q-learning (Haruno and Kawato, 2006) have been used in the context of studies of instrumental learning[2]. Our results are in accord with those of O'Doherty et al. (2004), who found that the actor function was associated with activity in the head of the caudate and putamen. Haruno and Kawato (2006) found that the putamen was most strongly predicted by reward prediction, as in our study, but they found that activity in the head of the caudate was more strongly predicted by prediction error than reward prediction. As we argued above, the anterior caudate can be seen as a transitional region between the ventral striatum and the other regions of the dorsal striatum: functionally, it integrates feedback with category and action selection, and anatomically it receives an intermediate level of dopaminergic innervation (Haber, 2003).

[2]SARSA and Q learning differ only in the prediction error term. In Q learning, the prediction error requires the maximum Q value for next state, maximized over all actions. SARSA simply uses the Q value for the next state and actual next action taken. In Q learning and SARSA, the reward prediction function is used to determine the best action by selecting the action that results in the highest predicted Q value. Actor-critic differs from both SARSA and Q learning in that it required the addition of another function to predict the best action for each state, rather than basing the choice on the value of Q.

Reinforcement learning models proved to be a useful addition to the standard fMRI analysis approach of examining contrasts between conditions using the general linear model. Activity patterns identified in the contrast comparing deterministic correct trials versus baseline can be compared with activity patterns identified by the reward prediction regressor. Both measures should to some degree reflect successful categorization; however, the former measure weights all correct trials equally, including those on in which a subject may have only a slight bias towards the correct category, or even is guessing. In contrast, reward prediction weights trials more highly the stronger the knowledge is about the expected reward linked to the action to be performed. Thus, reward prediction may more precisely isolate neural activity associated with learning strength. In addition, deterministic correct trials may reflect all phases of learning (early to late), whereas reward prediction will emphasize late learning. In our study, we found that the deterministic correct trials were associated with activity across the striatum broadly, whereas reward prediction was associated with more restricted regions of activity in the head of the caudate and anterior putamen.

Prediction error provides an alternative to examining feedback valence. Feedback valence compares all trial on which subjects were told they were wrong with all trials in which they were told they were right; this contrast may reflect a multiplicity of factors, including knowledge of the correct answer, expectations about feedback, and emotional connotations of positive and negative feedback. In contrast, the prediction error measure specifically reflects the degree to which the feedback was expected or unexpected. Prediction error also allows us to link activity to the body of research that has shown that prediction error provides a good account for the activity of midbrain dopamine neurons that provide a major input to the basal ganglia (Schultz and Dickinson, 2000), particularly the ventral striatum. In our study, feedback valence led to similar patterns of activity in both dorsal (anterior caudate) and ventral striatum, but prediction error accounted for activity only in the ventral striatum, indicating that prediction error may be a more precise measure of feedback sensitivity.

In conclusion, we successfully dissociated the contributions to visual categorization made by individual corticostriatal loops. We found patterns of activity in the visual loop consistent with playing a role in learning related visual processing, patterns of activity in the motor loop consistent with playing role in category related response learning, and patterns in the motivational loop and executive loops consistent with both playing a role in feedback processing, with the executive loop integrating feedback and category related learning. Additionally, we found interaction between the loops, with strong involvement of visual loop early in learning associated with visual processing, involvement of motor loop in response selection, and finally an influence of both loops on the executive loop as it integrates behavior and feedback.

Acknowledgments

This research was supported by a grant from the National Institutes of Health (R01MH079182) to Carol A. Seger. Marc V. Richard assisted with developing the behavioral task. This study was performed at Stanford University during CAS's sabbatical leave; she would like to thank Anthony Wagner for sponsoring her and all the members of the Wagner lab for their feedback and assistance.

References

Abler B, Roebroeck A, Goebel R, Höse A, Schönfeldt-Lecuona C, Hole G, Walter H. Investigating directed influences between activated brain areas in a motor-response task using fMRI. Magn Reson Imaging 2006;24:181–185. [PubMed: 16455407]

Alexander GE, DeLong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. Annu Rev Neurosci 1986;9:357–381. [PubMed: 3085570]

Ashby FG, Ennis JM, Spiering BJ. A neurobiological theory of automaticity in perceptual categorization. Psychol Rev 2007;114:632–656. [PubMed: 17638499]

Atallah HE, Lopez-Paniagua D, Rudy JW, O'Reilly RC. Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. Nat Neurosci 2007;10:126–131. [PubMed: 17187065]

Badre D, Wagner AD. Left ventrolateral prefrontal cortex and the cognitive control of memory. Neuropsychologia 2007;45:2883–2901. [PubMed: 17675110]

Balleine BW, Delgado MR, Hikosaka O. The role of the dorsal striatum in reward and decision-making. J Neurosci 2007;27:8161–8165. [PubMed: 17670959]

Bar-Gad I, Morris G, Bergman H. Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. Prog Neurobiol 2003;71:439–473. [PubMed: 15013228]

Bischoff-Grethe A, Hazeltine E, Bergren L, Ivry RB, Grafton ST. The influence of feedback valence in associative learning. Neuroimage. 2008

Boettiger CA, D'Esposito M. Frontal networks for learning and executing arbitrary stimulus-response associations. J Neurosci 2005;25:2723–2732. [PubMed: 15758182]

Brasted PJ, Wise SP. Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum. Eur J Neurosci 2004;19:721–740. [PubMed: 14984423]

Brovelli A, Laksiri N, Nazarian B, Meunier M, Boussaoud D. Understanding the neural computations of arbitrary visuomotor learning through fMRI and associative learning theory. Cereb Cortex 2008;18:1485–1495. [PubMed: 18033767]

Brown VJ, Desimone R, Mishkin M. Responses of cells in the tail of the caudate nucleus during visual discrimination learning. J Neurophysiol 1995;74:1083–1094. [PubMed: 7500134]

Buch ER, Brasted PJ, Wise SP. Comparison of population activity in the dorsal premotor cortex and putamen during the learning of arbitrary visuomotor mappings. Exp Brain Res 2006;169:69–84. [PubMed: 16284756]

Chouinard PA, Paus T. The primary motor and premotor areas of the human cerebral cortex. Neuroscientist 2006;12:143–152. [PubMed: 16514011]

Cincotta CM, Seger CA. Dissociation between striatal regions while learning to categorize via feedback and via observation. J Cogn Neurosci 2007;19:249–265. [PubMed: 17280514]

Cohen MX. Neurocomputational mechanisms of reinforcement-guided learning in humans: A review. Cognitive, Affective, & Behavioral Neuroscience 2008;8:113–125.

Curtis CE, D'Esposito M. Persistent activity in the prefrontal cortex during working memory. Trends Cogn Sci 2003;7:415–423. [PubMed: 12963473]

Delgado MR. Reward-related responses in the human striatum. Ann N Y Acad Sci 2007;1104:70–88. [PubMed: 17344522]

Delgado MR, Miller MM, Inati S, Phelps EA. An fMRI study of reward-related probability learning. Neuroimage 2005;24:862–873. [PubMed: 15652321]

Downing PE, Chan AW, Peelen MV, Dodds CM, Kanwisher N. Domain specificity in visual cortex. Cereb Cortex 2006;16:1453–1461. [PubMed: 16339084]

Fernandez-Ruiz J, Wang J, Aigner TG, Mishkin M. Visual habit formation in monkeys with neurotoxic lesions of the ventrocaudal neostriatum. Proceedings of the National Academy of Sciences 2001;98:4196.

Filoteo JV, Maddox WT, Simmons AN, Ing AD, Cagigas XE, Matthews S, Paulus MP. Cortical and subcortical brain regions involved in rule-based category learning. Neuroreport 2005;16:111–115. [PubMed: 15671857]

Foerde K, Knowlton BJ, Poldrack RA. Modulation of competing memory systems by distraction. Proc Natl Acad Sci U S A 2006;103:11778–11783. [PubMed: 16868087]

Frank MJ. Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. J Cogn Neurosci 2005;17:51–72. [PubMed: 15701239]

Gaffan D, Eacott MJ. Visual learning for an auditory secondary reinforcer by macaques is intact after uncinate fascicle section: indirect evidence for the involvement of the corpus striatum. Eur J Neurosci 1995;7:1866–1871. [PubMed: 8528460]

Genovese CR, Lazar NA, Nichols T. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage 2002;15:870–878. [PubMed: 11906227]

Goebel R, Roebroeck A, Kim DS, Formisano E. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magn Reson Imaging 2003;21:1251–1261. [PubMed: 14725933]

Grinband J, Hirsch J, Ferrera VP. A neural representation of categorization uncertainty in the human brain. Neuron 2006;49:757–763. [PubMed: 16504950]

Gusnard DA, Akbudak E, Shulman GL, Raichle ME. Medial prefrontal cortex and self-referential mental activity: Relation to a default mode of brain function. Proceedings of the National Academy of Sciences 2001:71043098.

Haber SN. The primate basal ganglia: parallel and integrative networks. J Chem Neuroanat 2003;26:317–330. [PubMed: 14729134]

Haber SN, Fudge JL, McFarland NR. Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. J Neurosci 2000;20:2369–2382. [PubMed: 10704511]

Hadj-Bouziane F, Boussaoud D. Neuronal activity in the monkey striatum during conditional visuomotor learning. Exp Brain Res 2003;153:190–196. [PubMed: 14610634]

Haruno M, Kawato M. Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. J Neurophysiol 2006;95:948–959. [PubMed: 16192338]

Hoshi E, Tanji J. Distinctions between dorsal and ventral premotor areas: anatomical connectivity and functional properties. Curr Opin Neurobiol 2007;17:234–242. [PubMed: 17317152]

Houk JC, Bastianen C, Fansler D, Fishbach A, Fraser D, Reber PJ, Roy SA, Simo LS. Action selection and refinement in subcortical loops through basal ganglia and cerebellum. Philos Trans R Soc Lond B Biol Sci 2007;362:1573–1583. [PubMed: 17428771]

Houk JC, Wise SP. Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. Cereb Cortex 1995;5:95–110. [PubMed: 7620294]

Humphries MD, Stewart RD, Gurney KN. A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. J Neurosci 2006;26:12921–12942. [PubMed: 17167083]

Joel D, Weiner I. The organization of the basal ganglia-thalamocortical circuits: open interconnected rather than closed segregated. Neuroscience 1994;63:363–379. [PubMed: 7891852]

Lau B, Glimcher PW. Action and outcome encoding in the primate caudate nucleus. J Neurosci 2007;27:14502–14514. [PubMed: 18160658]

Lawrence AD, Sahakian BJ, Robbins TW. Cognitive functions and corticostriatal circuits: insights from Huntington's disease. Trends Cogn Sci 1998;2:379–388.

Leh SE, Ptito A, Chakravarty MM, Strafella AP. Fronto-striatal connections in the human brain: a probabilistic diffusion tractography study. Neurosci Lett 2007;419:113–118. [PubMed: 17485168]

Mayka MA, Corcos DM, Leurgans SE, Vaillancourt DE. Three-dimensional locations and boundaries of motor and premotor cortices as defined by functional brain imaging: a meta-analysis. Neuroimage 2006;31:1453–1474. [PubMed: 16571375]

Monchi O, Petrides M, Petre V, Worsley K, Dagher A. Wisconsin Card Sorting revisited: distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. J Neurosci 2001;21:7733–7741. [PubMed: 11567063]

Murray EA, Bussey TJ, Wise SP. Role of prefrontal cortex in a network for arbitrary visuomotor mapping. Exp Brain Res 2000;133:114–129. [PubMed: 10933216]

Nomura EM, Maddox WT, Filoteo JV, Ing AD, Gitelman DR, Parrish TB, Mesulam MM, Reber PJ. Neural correlates of rule-based and information-integration visual category learning. Cereb Cortex 2007;17:37–43. [PubMed: 16436685]

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science 2004;304:452–454. [PubMed: 15087550]

Packard MG, Knowlton BJ. Learning and memory functions of the Basal Ganglia. Annu Rev Neurosci 2002;25:563–593. [PubMed: 12052921]

Pardo-Vazquez JL, Leboran V, Acuña C. Neural Correlates of Decisions and Their Outcomes in the Ventral Premotor Cortex. J Neurosci 2008;28:12396–12408. [PubMed: 19020032]

Poldrack RA, Gabrieli JD. Characterizing the neural mechanisms of skill learning and repetition priming: evidence from mirror reading. Brain 2001;124:67–82. [PubMed: 11133788]

Poldrack RA, Prabhakaran V, Seger CA, Gabrieli JDE. Striatal Activation During Acquisition of a Cognitive Skill. NEUROPSYCHOLOGY-NEW YORK- 1999;13:564–574.

Poldrack RA, Rodriguez P. How do memory systems interact? Evidence from human classification learning. Neurobiol Learn Mem 2004;82:324–332. [PubMed: 15464413]

Rodriguez PF, Aron AR, Poldrack RA. Ventral-striatal/nucleus-accumbens sensitivity to prediction errors during classification learning. Hum Brain Mapp 2006;27:306–313. [PubMed: 16092133]

Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using Granger causality and fMRI. Neuroimage 2005;25:230–242. [PubMed: 15734358]

Saint-Cyr JA, Ungerleider LG, Desimone R. Organization of visual cortical inputs to the striatum and subsequent outputs to the pallido-nigral complex in the monkey. J Comp Neurol 1990;298:129–156. [PubMed: 1698830]

Schönberg T, Daw ND, Joel D, O'Doherty JP. Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. J Neurosci 2007;27:12860–12867. [PubMed: 18032658]

Seger CA. How do the basal ganglia contribute to categorization? Their roles in generalization, response selection, and learning via feedback. Neurosci Biobehav Rev 2008;32:265–278. [PubMed: 17919725]

Seger CA, Cincotta CM. The roles of the caudate nucleus in human classification learning. J Neurosci 2005;25:2941–2951. [PubMed: 15772354]

Seger CA, Cincotta CM. Dynamics of frontal, striatal, and hippocampal systems during rule learning. Cereb Cortex 2006;16:1546–1555. [PubMed: 16373455]

Seger CA, Prabahkaran V, Poldrack RA, Gabrieli JDE. Neural activity differs between explicit and implicit learning of artificial grammar strings: An fMRI study. Psychobiology 2000;28:283–292.

Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS. Temporal difference models describe higher-order learning in humans. Nature 2004;429:664–667. [PubMed: 15190354]

Sutton, RS.; Barto, AG. Reinforcement Learning: An Introduction. MIT Press; 1998.

Teng E, Stefanacci L, Squire LR, Zola SM. Contrasting effects on discrimination learning after hippocampal lesions and conjoint hippocampal-caudate lesions in monkeys. J Neurosci 2000;20:3853–3863. [PubMed: 10804225]

Tricomi E, Delgado MR, McCandliss BD, McClelland JL, Fiez JA. Performance feedback drives caudate activation in a phonological learning task. J Cogn Neurosci 2006;18:1029–1043. [PubMed: 16839308]

Tricomi EM, Delgado MR, Fiez JA. Modulation of Caudate Activity by Action Contingency. Neuron 2004;41:281–292. [PubMed: 14741108]

Tsao DY, Moeller S, Freiwald WA. Comparing face patch systems in macaques and humans. Proc Natl Acad Sci U S A 2008;105:19514–19519. [PubMed: 19033466]

Tunik E, Schmitt PJ, Grafton ST. BOLD coherence reveals segregated functional neural interactions when adapting to distinct torque perturbations. J Neurophysiol 2007;97:2107–2120. [PubMed: 17202232]

van Duijvenvoorde AC, Zanolie K, Rombouts SA, Raijmakers ME, Crone EA. Evaluating the negative or valuing the positive? Neural mechanisms supporting feedback-based learning across development. J Neurosci 2008;28:9495–9503. [PubMed: 18799681]

Webster MJ, Bachevalier J, Ungerleider LG. Transient subcortical connections of inferior temporal areas TE and TEO in infant macaque monkeys. J Comp Neurol 1995;352:213–226. [PubMed: 7536756]

Williams ZM, Eskandar EN. Selective enhancement of associative learning by microstimulation of the anterior caudate. Nat Neurosci 2006;9:562–568. [PubMed: 16501567]

Yamada H, Matsumoto N, Kimura M. Tonically active neurons in the primate caudate nucleus and putamen differentially encode instructed motivational outcomes of action. J Neurosci 2004;24:3500–3510. [PubMed: 15071097]

Yamada H, Matsumoto N, Kimura M. History- and current instruction-based coding of forthcoming behavioral outcomes in the striatum. J Neurophysiol 2007;98:3557–3567. [PubMed: 17928555]

Yeterian EH, Pandya DN. Corticostriatal connections of extrastriate visual areas in rhesus monkeys. J Comp Neurol 1995;352:436–457. [PubMed: 7706560]

Yin HH, Knowlton BJ. The role of the basal ganglia in habit formation. Nat Rev Neurosci 2006;7:464–476. [PubMed: 16715055]

Zeithamova D, Maddox WT, Schnyer DM. Dissociable prototype learning systems: evidence from brain imaging and behavior. J Neurosci 2008;28:13194–13201. [PubMed: 19052210]

Zheng T, Wilson CJ. Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. J Neurophysiol 2002;87:1007–1017. [PubMed: 11826064]

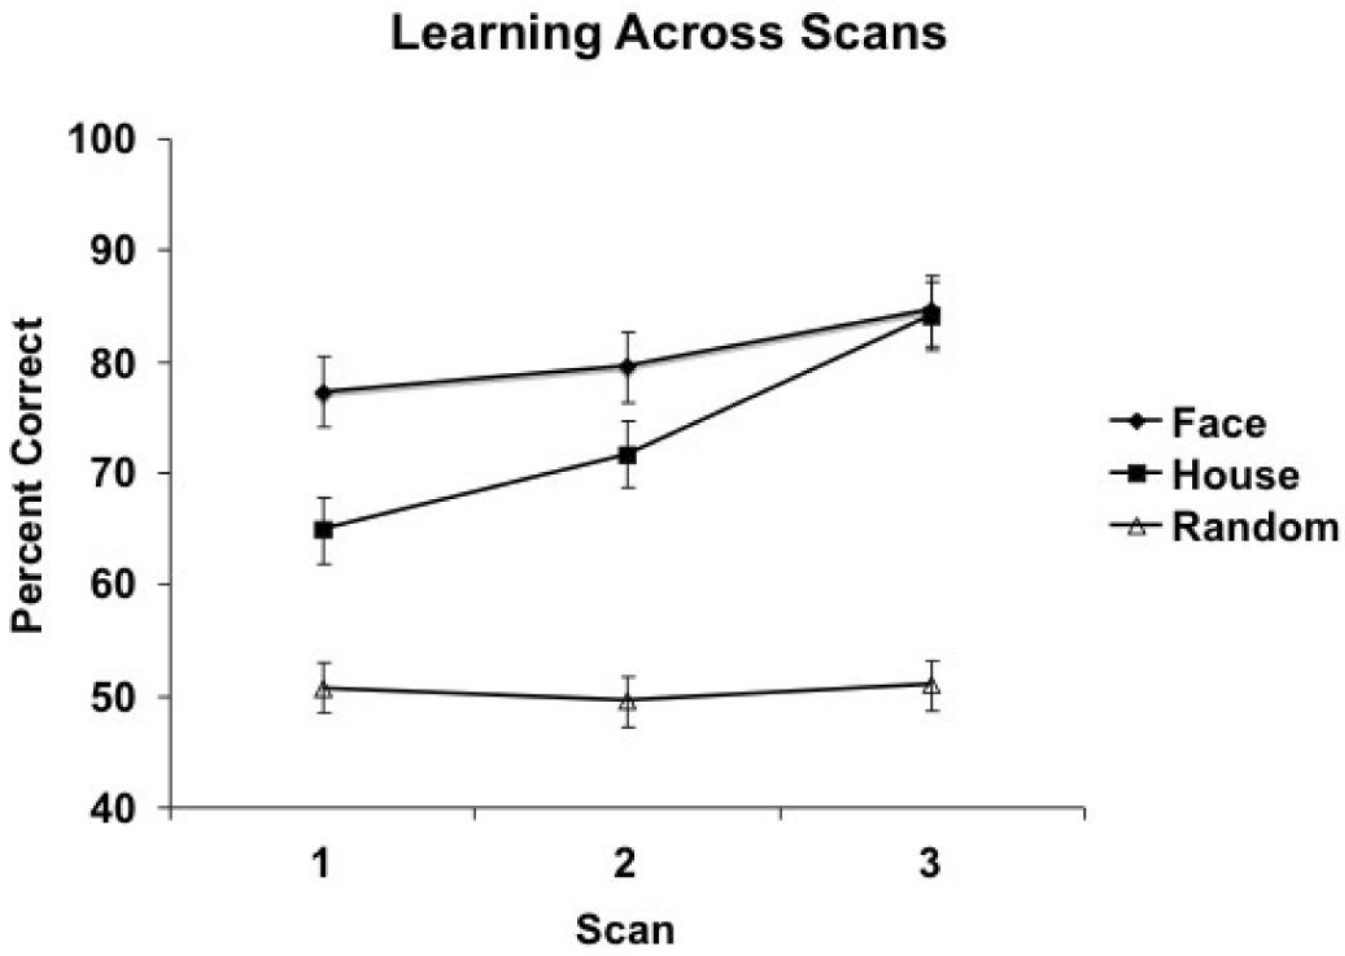

1.
Learning across scans by stimulus type: Faces (deterministic), Houses (deterministic) and Random (Faces and Houses combined). Error bars indicate standard error.

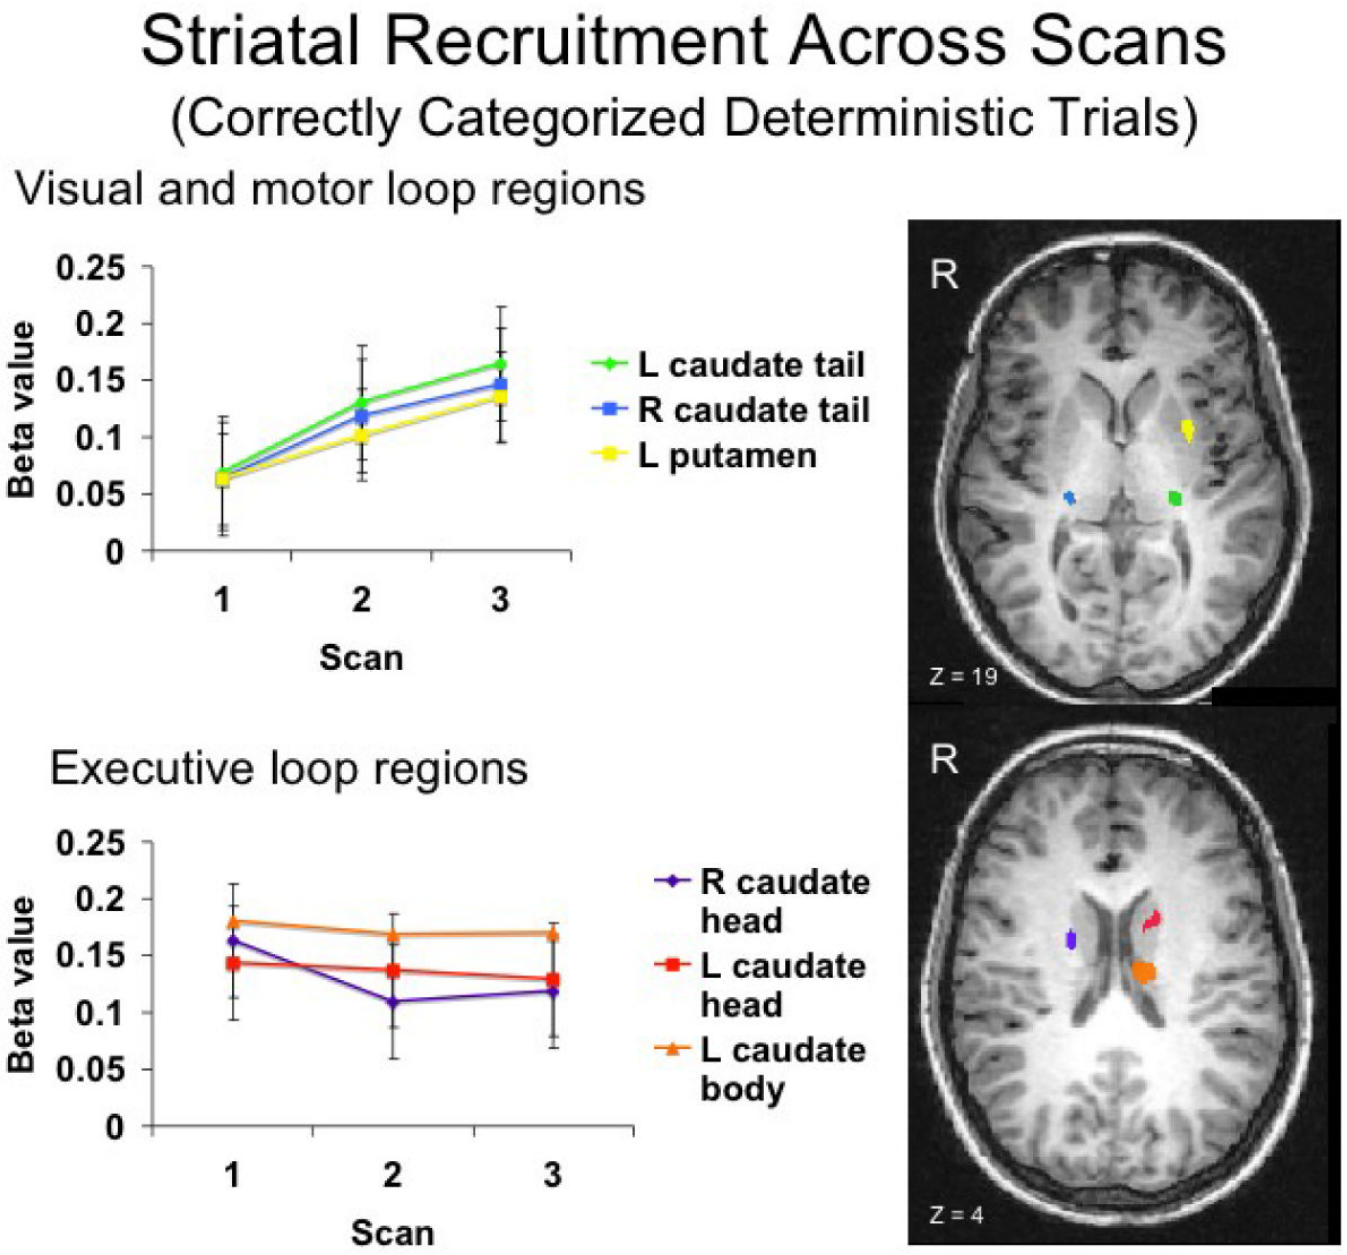

2.

Striatal regions across scans: motor and visual loop regions putamen and tail caudate (increases) and executive loop regions -- head and body caudate (stays about the same). All ROIs are from the contrast comparing activation on correctly categorized deterministic trials with the implicit baseline (see Table 1 for Tailarach coordinates), except the putamen ROI which is from the contrast comparing correctly and incorrectly categorized deterministic stimuli (see text; center at x = −28, y = 3, z = 2)

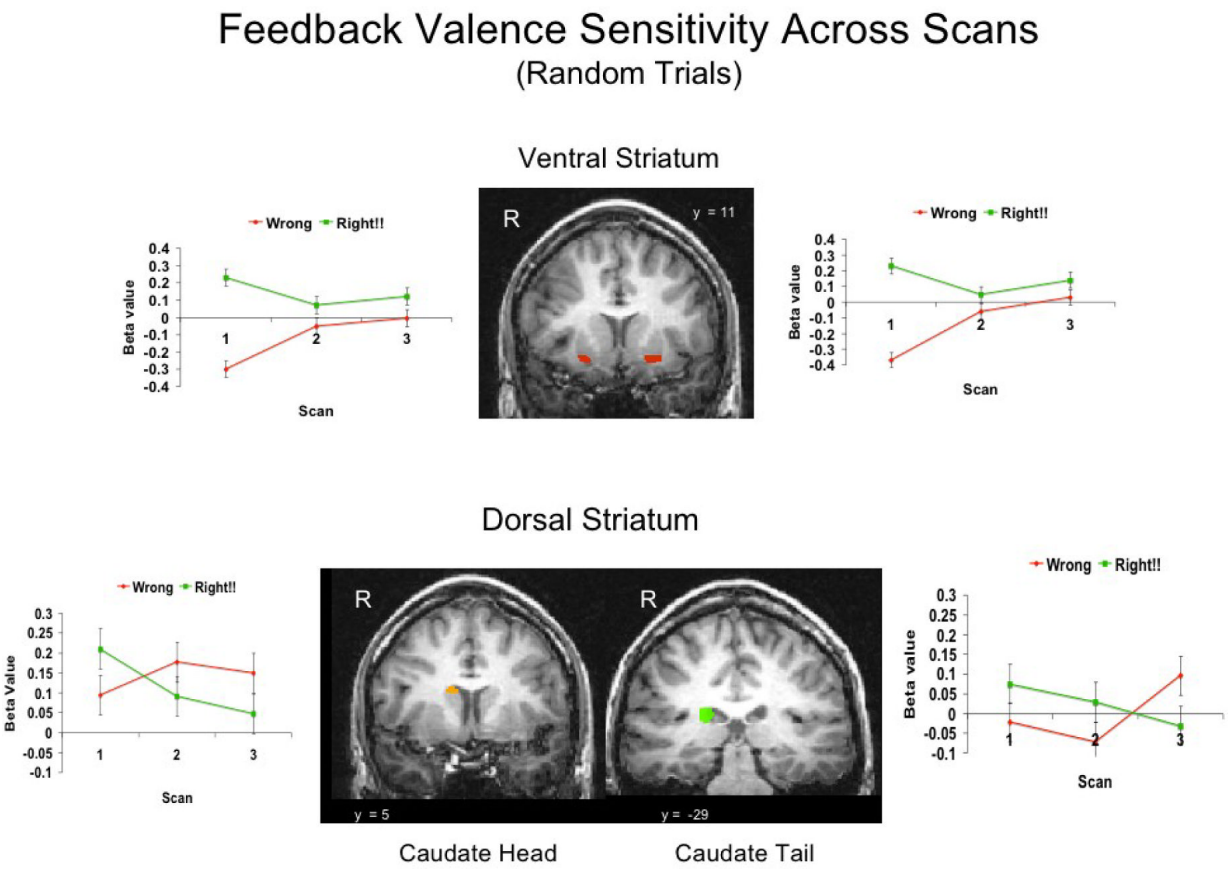

3.

Interaction of negative and positive feedback across scans in ventral and dorsal striatal regions.

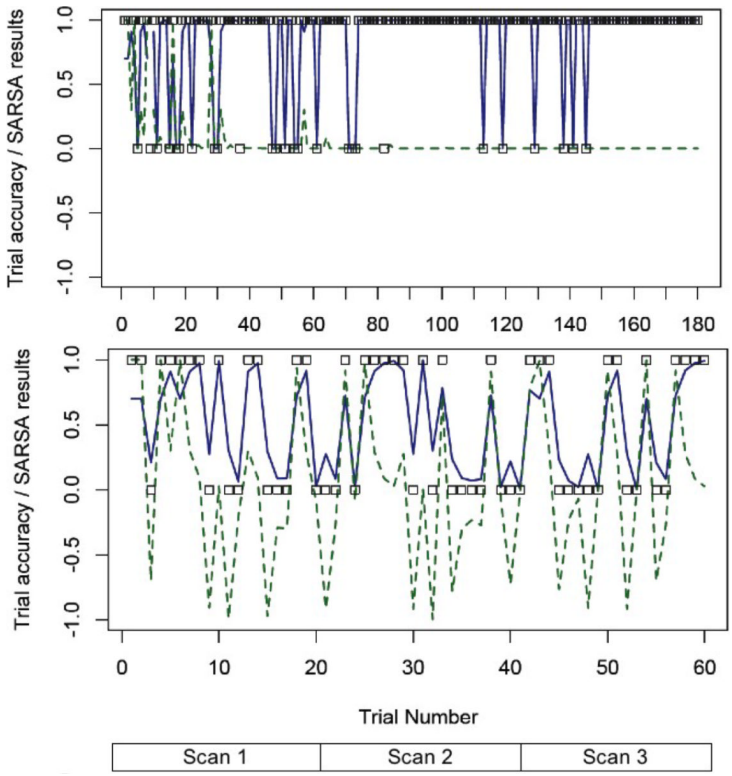

4.

Behavioral trial by trial accuracy and estimates of prediction error (dotted green line) and reward prediction (solid blue line) for a representative subject. Accuracy on each trial is indicated by the position of the black open square; the value is set to 1.0 if the categorization was correct, and to 0 if incorrect. For random trials, performance was considered correct if it received positive feedback. The top graph is for all deterministic house trials, and the bottom for random house trials; deterministic and random were plotted separately to better illustrate the convergence of the measures associated with learning of deterministic stimuli. Note that reward prediction varies between 0 and 1 across trials regardless of trial type, whereas prediction error for deterministic trials varies between 0 and 1 and on random trials varies between −1 and 1. Negative values can be thought of as unexpectedly getting positive feedback when negative feedback is expected; on deterministic trials this state is not possible. Also note that as the subject reaches a high level of proficiency on the deterministic trials reward prediction converges on the maximum value of 1, whereas prediction error is near zero except immediately after an error. On the random trials, net learning is not possible and both reward prediction and prediction error vary in response to the most recent trials.

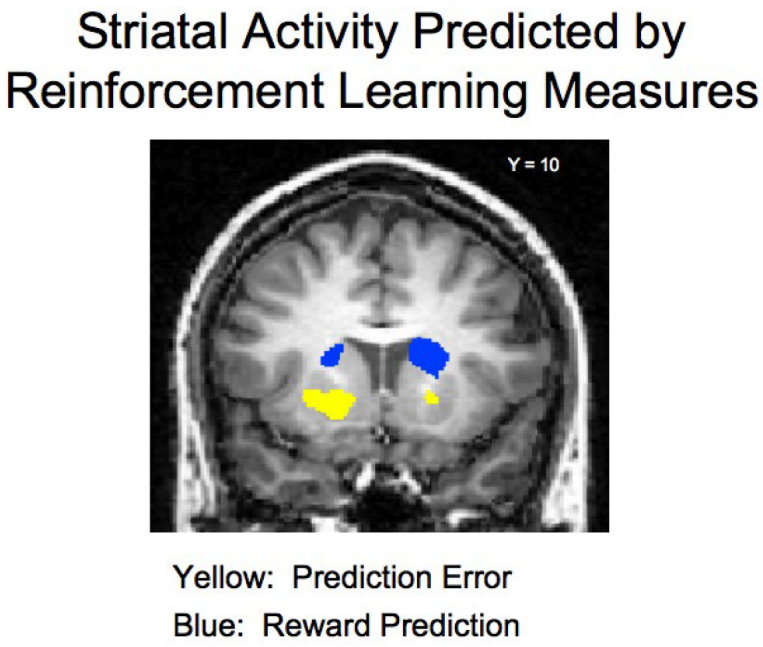

5.

Striatal regions in which activity was significantly predicted by prediction error (yellow), or by reward prediction (blue). Note that prediction error accounts for activity in the ventral striatum, and reward prediction accounts for activity in dorsal striatal regions, including the putamen and head of the caudate.

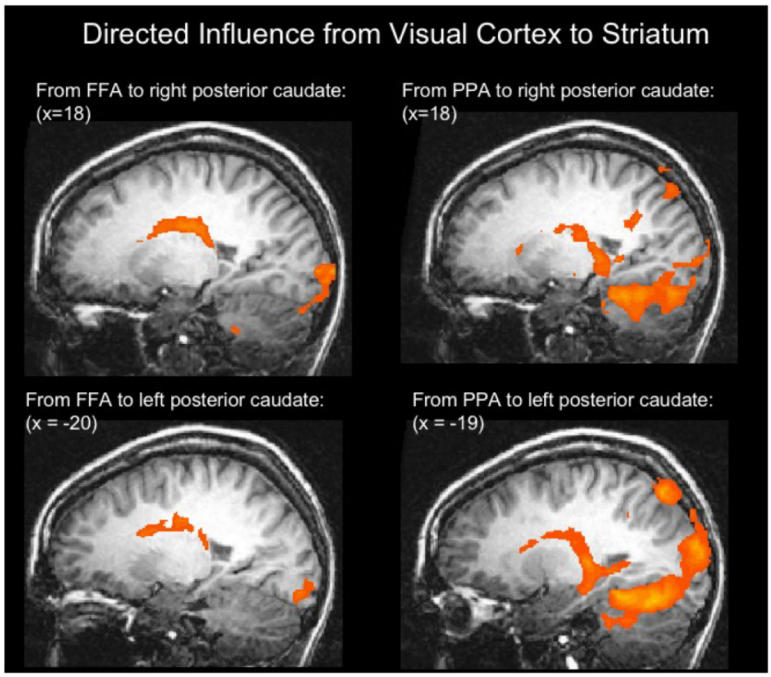

6.

Directed influence from the fusiform face area (left) and parahippocampal place area (right) on the caudate nucleus in the right (top) and left (bottom) hemispheres.

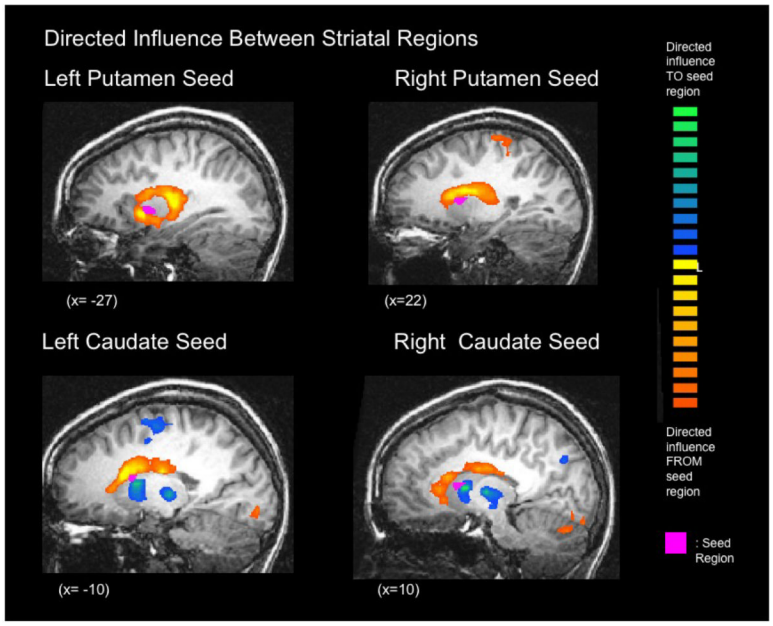

7.

Directed influence of ROIs in the left and right putamen and caudate nucleus on other regions of the striatum and thalamus. Magenta: seed region. Orange-yellow spectrum: Directed influence from the seed region. Blue-green spectrum: directed influence onto the seed region.

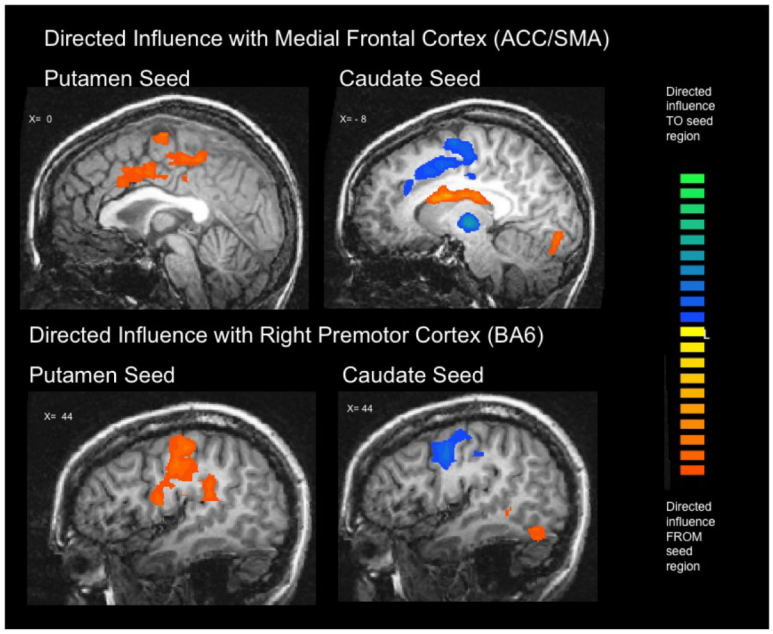

8.

Directed influence to and from cortical regions from the left and right putamen and caudate nucleus. Top: Medial saggital sections showing directed influence of the putamen on medial cortical regions and directed influence to the caudate. In addition, note the directed influence from the caudate onto the fusiform gyrus. Bottom: Lateral saggital sections showing directed influence on the precentral gyrus from the putamen, and from the precentral gyrus to the caudate. Orange-yellow spectrum: Directed influed from the seed region. Blue-green spectrum: directed influence onto the seed region. Note that the seed regions are the same as in Figure 7; they are not shown here because they do not fall within the selected slices.

Table 1

Regions of Activation: Correctly Categorized Deterministic Trials, All Scans

L/R/B	Region	Gyrus/subregion	Center Voxel (x y z)
Frontal Cortex:			
R	Ventral Lateral Premotor	Precentral Sulcus	39 0 31
R	Dorsolateral PFC	Middle Frontal	39 23 22
L	Posterior Ventrolateral PFC	Inferior Frontal / Insula	−28 24 4
R	Posterior Ventrolateral PFC	Inferior Frontal / Insula	27 22 4
Other Cortex:			
R	Medial Parietal	Posterior Cingulate	11 −29 33
L	Occipital	Lingual	−11 −91 −4
R	Occipital	Cuneus	22 −90 −3
R	Inferior Temporal	Fusiform	34 −42 −17
L	Inferior Temporal	Fusiform	−29 −52 −17
Striatum:			
L	Caudate	Head	−11 2 18
L	Caudate	Body	−10 −12 22
L	Caudate	Tail	−21 −23 2
R	Caudate	Tail	22 −23 4
R	Putamen	Anterior putamen	17 0 16
Other Subcortical:			
L	Thalamus		−9 −11 20
R	Thalamus		6 −7 8
B	Midbrain	Ventral Tegmental Area	−1 −15 −7
B	Cerebellum	Anterior lobe	6 −60 −26
Regions more active in baseline than deterministic correct trials:			
B	Medial Frontal	Medial frontal	1 57 9

Note: Regions of the lateral prefrontal lobes are identified in line with the distinctions laid out in (Badre and Wagner, 2007; Curtis and D'Esposito, 2003). Premotor regions are identified in line with the guidelines set out in. L: left hemisphere. R: Right hemisphere. B: bilateral. PFC: Prefrontal Cortex. The coordinates of the central voxel are given in Talairach space.

Table 2

Regions of Activation Predicted by Prediction Error (Scan 1, House stimuli).

L/R	Region	Gyrus/subregion	BA	Center Voxel (x y z)
Frontal:				
L	Posterior Ventrolateral PFC	Inferior Frontal / Insula		−24 22 1
R	Dorsolateral PFC	Middle Frontal		40 28 15
L	Dorsolateral PFC	Middle Frontal		−40 26 17
R	Premotor	Precentral gyrus		34 3 34
L	Anterior Prefrontal	Middle frontal gyrus		−36 51 −2
Other Cortex:				
R	Parietal	Inferior parietal lobule		33 43 43
B	Medial parietal	Posterior Cingulate		2 −25 36
R	Occipital	Middle occipital		30 −82 12
L	Occipital	Cuneus		−28 −82 22
R	Medial temporal	Parahippocampal		24 −33 −11
L	Medial temporal	Parahippocampal		−24 −30 −12
R	Medial temporal	Fusiform		24 −62 −16
L	Medial temporal	Fusiform		−26 −47 −12
L	Medial temporal	Fusiform		−22 −65 −11
Striatum:				
R	Putamen	Ventral	putamen	17 8 −1
L	Putamen	Ventral	putamen	−17 6 0

Note: Regions of the lateral prefrontal lobes are identified in line with the distinctions laid out in (Badre and Wagner, 2007; Curtis and D'Esposito, 2003). The coordinates of the central voxel are given in Talairach space. L: left hemisphere. R: Right hemisphere. B: bilateral. PFC: Prefrontal Cortex. SMA: Supplementary motor area.

Table 3

Regions of Activation Predicted by Reward Prediction (Scan 1)

L/R	Region	Gyrus/subregion	Center Voxel (x y z)
Frontal			
R	Premotor	Inferior frontal/Precentral sulcus	42 4 29
Other Cortex:			
L	Inferior Temporal	Fusiform	−42 −44 −18
R	Inferior Temporal	Fusiform	31 −60 −14
R	Inferior Temporal	Fusiform	30 −52 −22
R	Occipital	Middle occipital	34 −87 −13
L	Occipital	Middle occipital	−29 −79 −7
R	Parietal	Supramarginal	33 −54 33
L	Parietal	Supramarginal	−25 −57 39
Striatum:			
R	Putamen/ Globus Pallidus	Anterior putamen	16 3 4
R	Caudate, Putamen	Head	16 14 19
L	Caudate	Head	−16 9 17

Note: L: left hemisphere. R: Right hemisphere. B: bilateral. The coordinates of the central voxel are given in Talairach space.

Table 4

Directed Influences Measured via Granger Causality Modeling

Head of the caudate seed regions		**Left seed (−15 9 17)**	**Right (7 11 12)**
Directed influences to caudate:			
From Subcortical Regions:	center voxel		
L Putamen	−18 4 11	*	*
R Putamen	12 4 8	*	*
L Thalamus	−16 -19 3	*	*
R Thalamus	5 −14 −1	-	*
From Cortical Regions:			
Anterior Cingulate	6 23 32	*	*
Anterior Cingulate	−12 18 35	*	*
Medial Frontal / SMA	−15 −10 58	-	*
Right Premotor	41 −2 31	*	*
R Precuneus	17 −69 32	*	*
Directed influence from caudate to other regions:			
To Subcortical Regions:	center voxel		
R Caudate, body	13 −6 24	*	*
L Caudate, body	−15 −12 22	-	*
R Caudate, head	12 18 15	*	*
L Caudate, head	−15 18 14	-	*
To Cortical Regions:			
R Fusiform Gyrus	38 −68 −18	*	*
L Fusiform Gyrus	−36 −71 −18	*	*

Putamen seed regions		**Left (−27 4 4)**	**Right (20 6 11)**
Directed influences from putamen to:			
To Subcortical Regions:	center voxel		
R Caudate, body	15 −6 20	*	*
L Caudate, body	−17 −6 20	*	*
R Posterior putamen	28 −14 −3	*	*
L Posterior putamen	24 −18 0	*	*
R Putamen / head caudate	21 9 17	*	*
L Putamen / head caudate	−21 12 14	*	*
To Cortical Regions:			
Anterior cingulate	6 0 38	*	*
Medial Frontal / SMA	1 −10 64	*	*
R Premotor	54 −1 14	*	-
R Premotor	37 −13 41	*	*
L Premotor	−52 1 10	*	-
L Inferior Parietal	−48 −30 23	*	-
L Medial Parietal	−12 −32 53	*	*
Directed influences to putamen from other regions:			

Note:

L: left hemisphere. R: Right hemisphere. SMA: Supplementary Motor Area.

* directed influence was present to/from this region to the seed region.

- no directed influence was detected that met the statistical threshold.