

Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology

D. Shohamy,¹ C. E. Myers,² S. Grossman,¹ J. Sage,³ M. A. Gluck¹ and R. A. Poldrack⁴

¹Center for Molecular and Behavioral Neuroscience and

²Department of Psychology, Rutgers University, Newark,

³Department of Neurology, UMDNJ/Robert Wood Johnson Medical School, New Brunswick, New Jersey, and

⁴Department of Psychology, UCLA, Los Angeles, California, USA

Correspondence to: Daphna Shohamy, Center for

Molecular and Behavioral Neuroscience, Rutgers

University, 197 University Avenue, Newark, NJ, USA

E-mail: shohamy@axon.rutgers.edu

Summary

The striatum has been widely implicated in cognition, but a precise understanding of its role remains elusive. Here we present converging evidence for the role of the striatum in feedback-based learning. In a prior functional imaging study, healthy controls showed striatal activity during a feedback-based learning task, which was decreased when the same task was learned without feedback. In the present study, we show that individuals with striatal dysfunction due to Parkinson's disease are impaired on the feedback-based task, but not on a

non-feedback version of the same task. Parkinson's patients and controls also used different learning strategies depending on feedback structure. This study provides direct behavioural evidence from humans that cortico-striatal systems are necessary for feedback-based learning on a cognitive task. These findings also link between learning impairments in Parkinson's disease and the physiological and computational evidence for the role of midbrain dopaminergic systems in feedback processing.

Keywords: cognition; learning and memory; basal ganglia; dopamine; Parkinson's disease

Abbreviations: MMSE = Mini-Mental State Examination; UPDRS = Unified Parkinson's Disease Rating Scale.

Received July 17, 2003. Revised December 11, 2003. Accepted December 12, 2003 Advance Access publication March 10, 2004

Introduction

Recent advances in understanding the neural bases of learning and memory have emphasized a critical role for cortico-striatal circuitry in supporting a 'habit' or 'procedural' learning system (Squire, 1994; Knowlton *et al.*, 1996; Robbins, 1996; Squire and Zola-Morgan, 1996; Gabrieli, 1998; Jog *et al.*, 1999; Eichenbaum and Cohen, 2001). Evidence for the role of cortico-striatal contributions to learning comes from findings that individuals with disrupted striatal function, such as occurs in Parkinson's disease, are impaired at a wide range of tasks that rely upon procedural learning (Knowlton *et al.*, 1996; Saint-Cyr *et al.*, 1988; Vriezen and Moscovitch, 1990; Swainson *et al.*, 2000; Myers *et al.*, 2003). In Parkinson's disease, patients suffer from a profound loss of nigro-striatal dopamine neurons, disrupting striatal function (Agid *et al.*, 1987). However, individuals with Parkinson's disease are not impaired on all procedural tasks (e.g. Heindel *et al.*, 1989; Harrington *et al.*,

1990; Bondi and Kaszniak, 1991; Koivisto *et al.*, 1996; Reber and Squire, 1999; Smith *et al.*, 2001; Witt *et al.*, 2002), and currently there is no clear understanding of why the striatum appears to be critical for learning under some conditions, but not others. Thus, a precise understanding of the role of the striatum in learning and memory remains elusive.

Significant advances have been made in recent years into functional neurophysiological, neurochemical and neurocomputational characteristics of the striatum (Schultz *et al.*, 1997; Schultz, 2002; Beiser and Houk, 1998; Horvitz, 2000). Collectively, these studies emphasize a role for dopaminergic projections to the striatum in modifying behavioural responses to environmentally salient stimuli based on response-contingent feedback (Ljunberg *et al.*, 1992; Schultz *et al.*, 1997; Hollerman *et al.*, 2000; Horvitz, 2000; Schultz, 2002). These findings suggest, therefore, that striatal disruption may lead to impaired learning when a task relies

upon trial-by-trial feedback, but learning may be spared if the same task is learned in an ‘observational’ manner, with no feedback.

Results from functional neuroimaging in humans support this idea. In a previous functional neuroimaging study of probabilistic classification learning, we found that the striatum was significantly more active during feedback-based learning than during observational learning with no feedback, despite the fact that performance levels were similar in both cases (Poldrack *et al.*, 2001). The same effect was also found in the midbrain dopaminergic regions. Because neuroimaging cannot establish the necessity of particular regions for task performance (see, for example, Poldrack, 2000), it is critical to establish that patients with damage to striatal function are specifically impaired at feedback-based learning. Preliminary evidence for this claim can be found in previous neuropsychological studies, since many of the tasks where Parkinson’s disease patients showed impaired learning did involve trial-by-trial feedback-based learning (e.g. Knowlton *et al.*, 1996; Myers *et al.*, 2003), while many of the tasks that were spared in Parkinson’s disease patients did not (for discussion, see Reber and Squire, 1999). However, those tasks differed in many ways besides the lack of feedback. Thus, although converging evidence points towards a role for the striatum in feedback-based learning, the specific necessity of the striatum for feedback-based learning remains to be established.

For example, in the study by Knowlton *et al.* (1996), Parkinson’s disease patients and amnesic patients were tested on a probabilistic classification learning task. The study found that while Parkinson’s disease patients were impaired compared with control subjects, the amnesic patients initially performed as well as controls. On a post-test questionnaire, the individuals with amnesia failed to remember facts about the testing episode or to recognize visual stimuli that were used in the task; conversely, the Parkinson’s disease patients could report facts related to the testing episode and the visual stimuli that appeared on the screen, despite having been unable to master the task. Thus, feedback *per se* was neither manipulated nor examined in that study. In fact, Knowlton *et al.* (1996) assumed that the critical feature of the task related to the Parkinson’s disease deficit was its probabilistic nature rather than the feedback-based nature of the learning. If so, then Parkinson’s disease patients should presumably be impaired at learning the task even if the training involves non-feedback (observational) learning—as long as the category structure is probabilistic.

However, an alternate interpretation of the results of Knowlton *et al.* (1996) is that Parkinson’s disease patients could learn some details of the task through observation (and hence show good performance on the questionnaire, reporting what they had seen during the experiment), but could not learn the category structures based on feedback. This latter interpretation would be consistent with the prior imaging study showing basal ganglia activation during feedback-based learning of a probabilistic categorization task, but not

during non-feedback (observational) learning of the same task (Poldrack *et al.*, 2001). This would predict that Parkinson’s disease patients’ impairment on probabilistic categorization might be ameliorated if the training involves non-feedback (observational) learning.

The purpose of the present study was to directly assess the role of feedback in learning and memory impairments in patients with Parkinson’s disease. Patients and age-matched controls were required to learn a probabilistic classification learning task, similar to tasks previously found to be sensitive to striatal function in behavioural and imaging studies (Knowlton *et al.*, 1996; Poldrack *et al.*, 2001). Here, we compared performance on two versions of a probabilistic classification learning task: a ‘feedback-based’ version and an ‘observational’ version. In the ‘feedback-based’ version, subjects were provided with trial-by-trial feedback based on their response to each trial. In the ‘observational’ version, subjects were shown the stimuli together with the correct outcome, with no behavioural response and no feedback.

Consistent with evidence from electrophysiological and imaging studies, we predicted that Parkinson’s disease patients would be impaired at the feedback-based version, but would perform as well as controls on the observational version. The role of cortico-striatal circuitry in feedback-based versus observational learning was further examined by investigating learning strategies in Parkinson’s disease and control subjects in each version, using mathematical analyses of learning strategies described previously (see Gluck *et al.*, 2002). These strategy analyses allow a more detailed analysis of differences in how task performance is influenced by instructional conditions or brain disorders.

Methods

Participants

Participants included 28 individuals with a diagnosis of idiopathic Parkinson’s disease, randomly assigned to either the feedback ($n = 13$) or the observational ($n = 15$) learning conditions. Patients for this study were recruited from the motor disorders clinic at Robert Wood Johnson University Hospital (New Jersey, USA), having met diagnostic criteria for Parkinson’s disease as assessed by a neurologist (J.S.), and having given informed consent to participate.

All Parkinson’s disease patients were in the mild to moderate stages of the disease, with scores on the Hoehn–Yahr scale of motor function (Hoehn and Yahr, 1967) that ranged from 1 to 3. Patients’ motor function was also rated according to the Unified Parkinson’s Disease Rating Scale (UPDRS). All Parkinson’s disease patients were non-demented, as indicated by scores >24 on the Mini-Mental State State Examination (MMSE; Folstein *et al.*, 1975). Parkinson’s disease patients were also screened for clinical depression, as indicated by scores <15 on the Beck Depression Inventory (Beck *et al.*, 1996). All patients included in the study were treated with L-dopa, were stable on their medication doses for at least 3 months, and were responding well to the medication. Some patients were additionally being treated with dopamine receptor agonists (11 patients). No patients included in the study were treated with

Table 1 Clinical characteristics and demographic information for Parkinson's disease patients and controls

		Age	Education	MMSE	Hoehn–Yahr	UPDRS	Parkinson's disease duration (years)
Feedback task	Parkinson's disease ($n = 13$)	61.3 (8.4)	16.6 (2.4)	29.2 (0.8)	2.3 (0.8)	21.2 (6.3)	6.2 (3.2)
	Control ($n = 13$)	59.0 (6.4)	17.0 (2.4)	29.8 (0.4)	N/A	N/A	N/A
Observational task	Parkinson's disease ($n = 15$)	64.5 (6.0)	15.9 (3.3)	29.0 (1.2)	2.1 (0.7)	19.7 (6.8)	5.4 (4.7)
	Control ($n = 15$)	64.5 (6.0)	16.9 (2.3)	29.0 (0.9)	N/A	N/A	N/A

Age, education and Parkinson's disease duration in years. SD given in parentheses. N/A = not applicable

anti-cholinergics. Patients were tested within 3 h since their last dose of medication. Information about Parkinson's disease patients and controls is presented in Table 1.

Twenty-eight healthy control participants were recruited and randomly assigned to either the feedback ($n = 13$) or the observational ($n = 15$) learning conditions. Controls were screened for the presence of any neurological disorder or history of psychiatric illness including depression. Controls and Parkinson's disease patients did not differ in terms of age, education or MMSE [ANOVA (analysis of variance) on age, education and MMSE as dependent variables, with group (Parkinson's disease or control) and condition (feedback or observational) as independent variables; all $P > 0.05$].

All participants signed statements of informed consent before participating in behavioural testing. All studies conformed to research guidelines established by Rutgers University, Robert Wood Johnson and the Federal Government.

Stimuli

Cues were features on a Mr PotatoHead™ toy (hat, glasses, moustache or bow tie) and subjects were required to predict the preferred flavour of ice cream (vanilla or chocolate). Sample stimuli for each version are shown in Fig. 1.

Stimuli were created using a Mr Potato Head™ set (Toy Story 2™ and Silly Suitcase™ versions) (Playschool/Hasbro Inc., Pawtucket, RI, USA; items 2260/2289 and 2279). Each stimulus was based on a face to which eyes, ears and other features could be added. Stimuli were photographed directly into the computer using a digital camera; stimuli were then edited further using Adobe Photoshop to ensure consistent image size (2.95" high \times 2" wide), resolution (100 pixels/inch) and alignment of the figure within the pictures.

All stimuli consisted of the basic Mr Potato Head™ face with black eyes, red nose, white arms and green feet. The face had a visible hole and 'smiling' surface texture where the mouth would appear. This basic face was altered by addition of one or more features: cue 1 = black hat, cue 2 = black moustache, cue 3 = red eyeglasses, cue 4 = white bow tie. Fourteen stimuli were devised following the scheme shown in Table 2. In Table 2, each pattern is encoded via a numeric four-digit pattern corresponding to whether each of the four features (tie, glasses, moustache, hat) is present (1) or absent (0). Thus, pattern A = 0001 had the black hat, pattern B = 0010 had the black moustache, pattern C = 0011 had both hat and moustache, and so on. Background was a constant light brown with minimal visible shadows.

Once all stimuli were constructed, an additional two versions were made of each stimulus pattern: (i) one with the figure holding a vanilla ice cream cone in its right hand (appearing on the left of the picture); and (ii) one with the figure holding a chocolate ice cream

cone in its left hand (appearing on the right of the picture). The ice cream cones were taken from the Lil Chef's Bakery Ice Cream Party set (Toys 'R' Us, Paramus, NJ, USA; item 9326). The ice cream cones were photographed separately into the computer and then morphed onto the existing photographs using Adobe Photoshop, to ensure that the appearance of the Mr Potato Head™ figure was identical in each version of the stimulus.

Two hundred trials were constructed from the 14 patterns. The two outcomes ("vanilla" and "chocolate") were equally probable, but each feature was independently associated with each outcome with a fixed probability as shown in Table 2: the probabilities of "vanilla" given feature 1 (hat), 2 (moustache), 3 (glasses) and 4 (bow tie) were 0.8, 0.6, 0.4 and 0.2, respectively. The probability of "chocolate" given each feature was correspondingly 0.2, 0.4, 0.6 and 0.8. Trials were then constructed to adhere to these probabilities. Table 2 shows the number of times each pattern occurred with each outcome. Thus, for example, feature 1 (hat) appears in seven patterns (A–G), which together account for 100 trials; the outcome of 'vanilla' occurs on 80 of these trials, so $P(\text{"vanilla"}/\text{feature 1 present}) = 0.8$. The 200 trials defined in Table 2 were presented in a random, but fixed order, for all subjects.

Apparatus

The experiment was conducted on a Macintosh G3 or iBook computer with colour screen, programmed in the SuperCard language (Allegiant Technologies, San Diego, CA, USA). The keyboard was masked except for two keys labelled "vanilla" and "chocolate", which the subject used to enter responses. During the observation phase of the observational condition, the keyboard was masked except for one key, labelled 'NEXT', which the subject pressed to see the next stimulus.

Procedure

Feedback-based condition

The subject was seated in front of the computer at a comfortable viewing distance. Instructions appeared on the computer screen. The subject read these instructions aloud:

Welcome! In this game, you are working in an ice cream shop. Customers will come in and buy vanilla or chocolate ice cream cones. Each time a customer visits, try to guess whether he wants vanilla or chocolate. If you guess correctly, you will earn an extra \$1 tip. Try to collect as many tips as you can. Good luck!

Participants were told that at first they would have to guess, but that they would gradually improve their performance.

On each of 200 training trials, the screen showed a PotatoHead™ figure (without ice cream) along with the prompt: *Which flavor do*

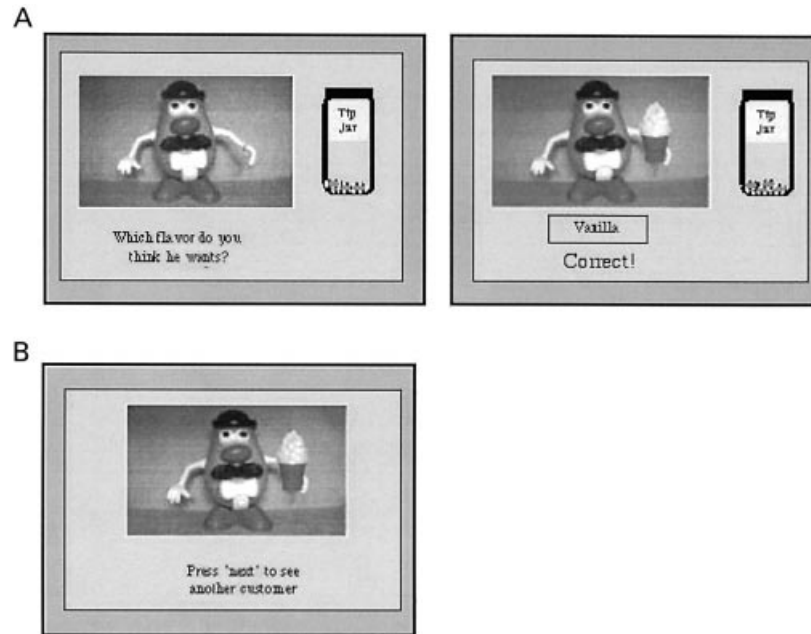


Fig. 1 Probabilistic classification learning task. Subjects learn to predict which flavour of ice cream Mr PotatoHead™ prefers (vanilla or chocolate) based on the presentation of four visual cues (hat, glasses, moustache, bow tie). **(A)** In the feedback-based version, subjects are required to guess the outcome on each trial and are provided with response-contingent feedback. **(B)** In the observational version, subjects are shown the stimulus together with the outcome on each trial, providing no response and receiving no feedback.

Table 2 Probability structure of the task

Pattern	Cue				P (cue combination)		
	1	2	3	4	P (pattern)	Frequency (No. per 200 trials)	P (outcome)
A	0	0	0	1	0.095	19	0.89
B	0	0	1	0	0.045	9	0.78
C	0	0	1	1	0.130	26	0.92
D	0	1	0	0	0.045	9	0.22
E	0	1	0	1	0.060	12	0.83
F	0	1	1	0	0.030	6	0.50
G	0	1	1	1	0.095	19	0.89
H	1	0	0	0	0.095	19	0.11
I	1	0	0	1	0.030	6	0.50
J	1	0	1	0	0.060	12	0.17
K	1	0	1	1	0.045	9	0.55
L	1	1	0	0	0.130	26	0.08
M	1	1	0	1	0.045	9	0.44
N	1	1	1	0	0.095	19	0.11
Total					1.00	200	

On any trial, one of 14 possible combinations of four cues could appear with the probability indicated [P (pattern)]. Each combination of cues predicted one outcome with the probability P (outcome) shown above and predicted the other outcome with a probability of $1-P$ (outcome).

you think he wants? The subject responded by pressing one of the labelled keys; the word "vanilla" or "chocolate", corresponding to the subject's response, appeared below the prompt. At this point, the stimulus pattern was replaced by a picture of the same figure holding either a vanilla or chocolate ice cream cone. If the subject's guess was correct, the word "Correct" appeared at the bottom of the screen, a few coins were added to the image of the tip jar, and a sound of

dropping coins was played through the computer speaker. If the guess was incorrect, the word "Incorrect" appeared at the bottom of the screen and the tip jar was unchanged. The figure with ice cream and the feedback remained on the screen during a one-second intertrial interval. If the subject did not respond within 2 s, a prompt appeared: *Answer Now!* If the subject did not respond within the next 3 s, the trial was terminated and the correct answer was shown.

Observational condition

The procedure was generally similar to that described for the feedback condition above. In the observational condition, training was broken into two phases: an observation phase and a test phase. Before starting the observation phase, the following instructions appeared on the screen:

Welcome to the ice cream parlor. You will see pictures of customers, along with their favorite flavor of ice cream – either vanilla or chocolate. Pay attention closely. Later on, you will be asked to remember which flavor of ice cream each customer wants. When you are finished looking at each customer, press the 'NEXT' key to see the next one. Good Luck!

On each trial, a Mr PotatoHead™ figure appeared with his favourite ice cream (vanilla or chocolate) in hand. When the subject was ready to move to the next customer, they pressed the "NEXT" button on the keyboard. There was no time constraint; however, subjects spent on average 1–2 s for each observational trial, which was similar to the amount of time allotted for each trial in the feedback-based condition. The observation phase consisted of 100 trials. After the last observation trial, instructions appeared on the screen stating that a new phase was beginning, that subjects would now be shown customers without their ice cream, and that their job was to try and predict the correct flavour of ice cream for each customer. On each trial of the test phase, the screen showed a PotatoHead™ figure (without ice cream) along with the prompt: *Which flavor do you think he wants?* The subject responded by pressing "vanilla" or "chocolate" on the keyboard. After selecting a response, subjects were shown the next trial, with no feedback. After the task was completed, subjects were told how many correct responses they had made overall.

Trial order for the 100 observation trials was identical to that used for the first 100 trials in the feedback condition. During the test phase, subjects were tested three times for each of the 14 patterns, for a total of 42 test trials; these trials were presented in a random, but fixed order, for all subjects.

Data collection

On each trial, the computer recorded the pattern, the subject's response and the actual outcome. The subject's response was defined as optimal if it matched the outcome that was most often associated with that pattern across the course of the experiment. For example, since pattern A = 0001 is most often associated with "vanilla" (see Table 2), a "vanilla" response is optimal for that pattern—even though on a few trials the actual outcome is "chocolate". Following earlier studies by Knowlton *et al.* (1994, 1996) and others, we accordingly defined a 'correct' response as one that obeyed this 'optimal response' rule, regardless of the actual outcome (i.e. whether the participant accurately predicted the outcome). Note that optimal response is undefined for patterns F = 0110 and I = 1001, which are equally often associated with each outcome.

For the feedback condition, percent correct scores were analysed in blocks of 50 trials. For the observational condition, percent correct scores were analysed for performance on the test phase, as described above.

Following prior studies (e.g. Poldrack *et al.*, 2001; Gluck *et al.*, 2002), subjects failing to reach a set criterion of 60% correct by the last block (feedback) or test phase (observational) were excluded from further analyses. Based on this criterion, three controls and five

Parkinson's disease subjects failed the feedback-based task; four Parkinson's disease and six controls failed the observational task [these did not differ significantly; $\chi^2(1) = 0.002$, $P > 0.5$].

Strategy analysis

Strategy analysis followed procedures described in Gluck *et al.* (2002). To investigate response strategies, for the entire training session (200 training trials for the feedback-based and 42 trials for the observational condition), we generated response profiles based on how an 'ideal' participant would respond on each trial if they had been following each strategy: multi-cue, one-cue or singleton (see details below). For each participant, we then calculated the degree to which each model fit the participant's data, using a least-mean-square measure, with 0.0 indicating a perfect fit. Comparing all strategies examined, the model that most closely approximated a participant's individual response profile was defined as the best-fit model for that participant. Because some participants may not be well fit by any pre-defined model, we excluded strategy analysis data from any participant who was not fit by any model within a tolerance of 0.1. Prior studies found that at least 95% of young control subjects were fit within this defined tolerance by one of the strategies described (Gluck *et al.*, 2002).

As described previously (Gluck *et al.*, 2002), we considered the following three classes of learning strategies:

(i) Multi-cue strategy: this is the optimal strategy for learning this task. Under this strategy, a participant should respond to each pattern of cues with the outcome most often associated with that pattern. This involves attending to the entire pattern (i.e. all four cues) present on each trial. A participant reliably following this strategy would be scored as making 100% correct optimal responses over the course of the experiment. In addition, we considered two sub-optimal strategies, in which participants focus on single cues or single patterns, rather than on all four cues:

(ii) One-cue strategy: using this strategy, a participant should respond to each pattern based on the presence or absence of a single cue, disregarding the other cues. For example, a participant might respond "vanilla" whenever cue 1 is present and "chocolate" otherwise, regardless of what other cues are present. A participant reliably following this strategy should generate 90% correct optimal responses. (Cue 4, which predicts chocolate with high accuracy, could also be used to generate 90% correct responses. Cues 2 or 3, which are associated less reliably with the two outcomes, could each be used to generate 67% correct responses.)

(iii) Singleton strategy: in this strategy, a participant should learn the outcomes associated with those patterns in which only a single cue appears. For example, a participant would learn that cues 1 and 2 each reliably predict "vanilla", while cues 3 and 4 reliably predict "chocolate". Thus, whenever cues 1 or 2 (alone or together) were present, participants responded "vanilla"; whenever cues 3 or 4 (alone or together) were present, participants responded "chocolate". However, whenever a combination of cues appeared which differed in association (e.g. pattern E with cues 2 and 4 present), responding was random. Note that in this strategy, the response to a pattern cannot be different than the sum of the responses to individual cues. Nevertheless, since patterns A, B, C, D, H and L occur with such high frequency during the experiment (accounting for 54% of all trials), a participant responding correctly to these patterns and randomly to the remaining patterns could achieve up to 77% correct over the course of the experiment.




Fig. 2 Percent correct for Parkinson's disease and controls on the observational task (*right*) and the test phase of the feedback-based task (*left*).

Results

The results include analyses for those subjects who reached criterion performance of at least 60% correct by the last block of training, including 10 controls and eight Parkinson's disease patients for the feedback-based task, and nine controls and 11 Parkinson's disease on the observational task.

Overall learning

Figure 2 shows performance for Parkinson's disease patients and controls on the test phase of the observational task, compared with the corresponding block of the feedback-based task (trials 100–150). Consistent with our prediction, Parkinson's disease patients were impaired at learning the feedback-based version, but were not impaired at learning the observational task. An ANOVA on performance by group (dependent variable) and condition (independent variables) revealed a significant main effect of group [$F(1,33) = 12.9$, $P < 0.01$] and a significant group \times condition interaction [$F(1,33) = 5.2$, $P < 0.05$]. *Post hoc* Tukey analyses confirmed that this was due to a significant difference between Parkinson's disease and controls on the feedback-based task ($P < 0.01$), but not on the observational task ($P = 0.8$). *Post hoc* analyses of performance across conditions showed that Parkinson's disease patients were significantly worse on the feedback-based task compared with the observational task ($P < 0.05$), whereas there was no difference between the tasks for control subjects ($P = 0.9$). For comparison with previous studies, learning curves for the feedback-based condition are shown in Fig. 3.

Learning strategies

We found that subjects engaged in different types of learning strategies in the feedback-based version compared with the observational version. In the feedback-based version, all control subjects and all but two Parkinson's disease patients




Fig. 3 Performance of Parkinson's disease patients and controls on the feedback-based task across the first 50 trials of training.

showed performance consistent with one of three models of probabilistic classification learning ('multi-cue', 'one-cue' or 'singleton' strategies; for details see Gluck *et al.*, 2002). These three learning models each assume different patterns of acquisition of cue-outcome associations during learning of the task: for all four of the cues (a 'multi-cue' model); for just one of the cues, such as the hat (a 'one-cue' model); or for those patterns where only a single cue is presented (a 'singleton' model; patterns A, B, D, H in Table 2; more details are described in Methods). By contrast, in the observational version, only three control subjects and none of the Parkinson's disease subjects performed in a manner consistent with any of the strategies described previously.

Additionally, within the feedback-based task, there were differences between Parkinson's disease and controls in the kind of strategy used during learning as shown in Fig. 4. Specifically, while 50% of control subjects were fit by an optimal 'multi-cue' strategy, only 16% of Parkinson's disease were fit by this strategy. Instead, most Parkinson's disease subjects (>60%) appeared to engage in a useful but sub-optimal 'singleton' strategy. This was a significant difference [Yates corrected Chi-square, $\chi^2(1) = 3.88$, $P < 0.05$].

Although in the observational task subjects were not fit by these same strategies, nonetheless, many subjects showed responses that appeared to be consistent with particular patterns of learning. Many subjects in each group made consistent responses (zero errors on certain patterns and 100% errors on other patterns), indicating that they may have formulated a specific rule-based strategy they followed in order to predict the outcomes for each pattern. Other subjects showed mixed responses with little or no consistency in responding to particular patterns. Figure 5 presents sample data from two subjects with similar overall levels of performance following each of these response strategies. To quantify the frequency of these strategies among the groups, we defined 'consistent responders' as any subject who made uniform responses to at least nine patterns, and 'inconsistent responders' as any subject who provided non-uniform responses to at least nine patterns. Using these measures,

we found that approximately half the subjects in each group were consistent responders; there was no evidence for differences between Parkinson's disease and controls in terms of number of subjects showing consistent versus inconsistent response patterns [Chi-square comparison of number of subjects in each group who responded consistently with number of subjects who did not: $\chi^2(1) = 1.9, P > 0.1$].

For comparison, we applied the same approach to the data from the feedback-based task, looking at consistent versus mixed responding to the last three encounters with each pattern. This analysis revealed that two Parkinson's disease patients and seven control subjects were defined as consistent




Fig. 4 Percentage subjects in the feedback-based task in each group best fit by the multi-cue or the singleton strategies. The multi-cue strategy refers to subjects learning to respond correctly based on the association between all four cues and each outcome. The singleton strategy refers to a useful, but sub-optimal strategy in which subjects respond correctly to those stimulus patterns that consist of only a single cue.

responders. In line with what would be expected based on the strategy analyses, eight of those subjects defined as consistent responders were also best-fit by the optimal multi-cue strategy (one Parkinson's disease patient was not fit by any of the strategies).

Discussion

The results presented here provide behavioural evidence from humans that cortico-striatal systems are necessary for feedback-based learning on a cognitive task. The results present a direct confirmation of a prediction inspired by previous neuroimaging results from normal humans (Poldrack *et al.*, 2001), which had demonstrated differences in engagement of striatal and midbrain dopaminergic regions between feedback-based and observational learning. The present findings also provide a link between the learning impairments in Parkinson's disease and the substantial physiological and computational evidence for the role of midbrain dopaminergic systems in reward processing. In doing so, our findings propose a means by which to understand the varied pattern of spared and impaired learning in Parkinson's disease across different tasks.

The present findings are consistent with converging evidence suggesting that the basal ganglia process feedback-related information to modify learning. Electrophysiological and anatomical studies show that cortico striatal synapses are modified based on dopaminergic signals from the SNc/VTA (substantia nigra pars compacta/ventral tegmental area), which carry stimulus-specific reward-related information (Calabresi *et al.*, 1992; Cepeda *et al.*, 1993; Wickens *et al.*, 1996; Reynolds *et al.*, 2000). These feedback-related signals are thought to be critical in modifying the organism's response to future encounters with the same

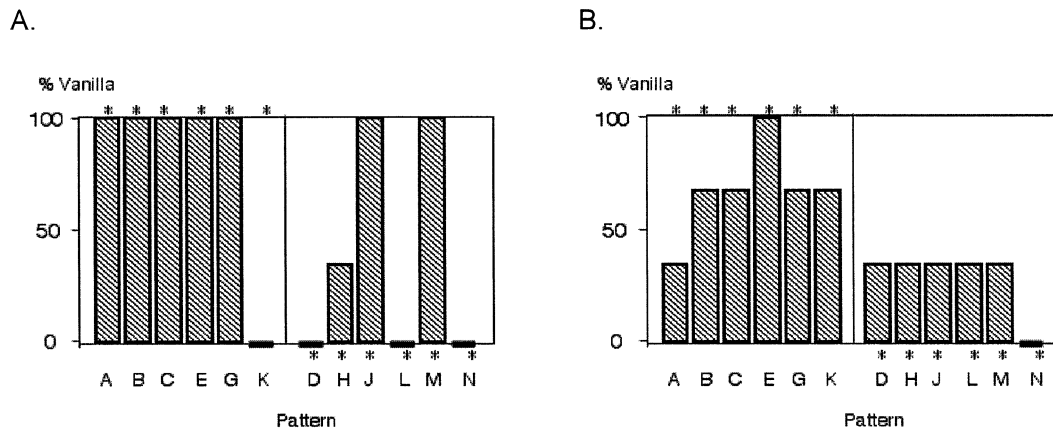


Fig. 5 Data from two representative subjects showing the proportion of 'vanilla' responses for each trial type (asterisks indicate the optimal response for each pattern). Optimally, a subject should respond "vanilla" to patterns A, B, C, E, G, K (left of graph) and "chocolate" to patterns D, H, J, L, M, N. (A) One subject consistently mapped most patterns to either always chocolate or always vanilla. This subject consistently responded "vanilla" to patterns A, B, C, E, G (correctly), but also consistently responded "vanilla" to patterns M and J (incorrectly) (B) Another subject showed mixed responses, responding to almost every pattern as sometimes chocolate, sometimes vanilla, but rarely responded always with either chocolate or vanilla to any single pattern.

stimulus (Wickens *et al.*, 1996; Reynolds *et al.*, 2000). In Parkinson's disease, patients suffer from profound loss of nigro-striatal dopamine neurons, disrupting striatal function (Agid *et al.*, 1987). Thus, the present findings support the idea that input from midbrain dopaminergic systems is critical in modifying behavioural responses based on trial-by-trial feedback.

The fact that Parkinson's disease patients showed no learning impairments on the observational version suggests that this type of learning is supported by brain systems which are spared in Parkinson's disease, and most likely by the medial temporal lobe. Our prior imaging studies found that while feedback-based learning primarily activated the striatum, learning this task with no feedback was associated with increased activation in the medial temporal lobe compared with feedback-based learning (Poldrack *et al.*, 2001). The finding in the present study that approximately half of the Parkinson's disease patients and controls learn this task by responding correctly to some patterns, but not others, is also reminiscent of the type of 'declarative' learning strategies typically associated with the medial temporal lobe (Squire, 1994).

The present findings are thus broadly consistent with prior studies demonstrating that individuals with Parkinson's disease are impaired at procedural learning, but are not impaired on an assessment of declarative memory for training events (such as the visual cues presented during training) (Knowlton *et al.*, 1996). Adopting a similar reasoning, one might suggest that in the present study, Parkinson's disease patients are impaired on the feedback-based version because it relies upon striatal-dependent procedural learning, while Parkinson's disease patients are spared on the observational version because this task relies upon striatal-independent declarative learning. Although this may be the case, the present study does not provide direct evidence to determine whether subjects were using procedural versus declarative learning. In fact, if the observational task were indeed learned by declarative rule-based mechanisms, all subjects might have been expected to show consistent responding to each stimulus during the test phase. However, we found that approximately half of the subjects in each group did not show such consistent responding, despite achieving overall similar levels of performance (see Fig. 4). Furthermore, recent studies have emphasized the difficulty in *a priori* defining a task as either procedural or declarative, since multiple strategies can be used to learn a probabilistic categorization task—some of which may be more easily verbalized than others (Gluck *et al.*, 2002). Furthermore, although subjects typically can verbalize a learning strategy when asked, modelling response patterns revealed that the actual strategy used during learning is not typically related to subjects' verbalized rules (Gluck *et al.*, 2002). One possibility, of course, is that healthy subjects normally make use of multiple parallel learning systems; in Parkinson's disease patients, where a feedback-based learning system may be damaged, learning must rely on alternate systems.

The present findings indicate that the Parkinson's disease patients perform significantly better when learning relies on observation rather than when the same information is trained in a feedback-based manner. Thus, one question is why the Parkinson's disease patients do not abandon the feedback-based learning strategies (which are impaired) and instead use the observational information that is presented during training to improve on the feedback-based task. The fact that the Parkinson's disease patients do not adopt observational strategies to learn the feedback-based task suggests that not only are they impaired at processing the feedback-related information, but that they can not modulate the use of feedback-based strategies when feedback-related information is presented during training. This may be related to a deficit in shifting or switching strategies, which is often attributed to cortico-striatal circuitry (Downes *et al.*, 1989; Owen *et al.*, 1993; Cools *et al.*, 2001).

The primary deficit in Parkinson's disease is a loss of dopaminergic neurons in the nigro-striatal pathways. However, this loss of dopamine in the striatum is also associated with disrupted frontal lobe function, as well as with disruption of other non-dopaminergic neurotransmitter systems, which are likely to contribute to the cognitive deficits found in Parkinson's disease patients. In addition, Parkinson's disease is commonly treated with L-dopa, which elevates dopamine levels in the brain and alleviates the motor symptoms of the disease. Studies examining the effect of L-dopa on cognitive deficits in Parkinson's disease have led to inconsistent results, with L-dopa either enhancing, having no effect, or impairing cognitive function, depending on task demands (e.g. Swainson *et al.*, 2000; Cools *et al.*, 2001). Cools *et al.* (2001) have suggested that L-dopa may impair cognitive function by 'overdosing' non-depleted fronto-striatal circuits. Since the patients in the present study were all tested while on-medication, future studies are necessary to determine the impact of L-dopa medication on feedback-based learning.

The present study manipulated the feedback-based structure of the learning task by eliminating all aspects of the feedback, including the reward that was associated with the outcome when it was correct. Therefore, in addition to feedback per se, the two versions of the task also differ in other respects. For example, in the observational condition, there is no requirement for explicit guessing of the predicted outcome, no decision-related motor response, and no reward associated with the correct outcome. Although collectively these aspects of the task are related to its feedback structure, future studies would be instrumental in dissociating the contribution of these feedback-related variables to the present results.

Conclusions

The present study suggests that the striatum plays an important role in feedback-based learning of cue-outcome associations. The findings presented here further suggest that

prior inconsistencies, i.e. with Parkinson's disease patients impaired on some learning tasks but not others, can be explained by the specific task demands. In addition, the present study provides a physiological context for understanding the relationship between striatal disruption and the procedural learning deficits found in patients with Parkinson's disease.

References

- Agid Y, Javoy-Agid F, Ruberg M. Biochemistry of neurotransmitters in Parkinson's disease. *Mov Disord* 1987; 2: 166–230.
- Beck AT, Steer RA, Brown GK. *Manual for the Beck Depression Inventory*. San Antonio (TX): Psychological Corporation; 1996.
- Beiser DG, Houk JC. Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *J Neurophysiol* 1998; 79: 3168–88.
- Bondi MW, Kaszniak AW. Implicit and explicit memory in Alzheimer's disease and Parkinson's disease. *J Clin Exp Neuropsychol* 1991; 13: 339–58.
- Calabresi P, Maj R, Pisani A, Mercuri NB, Bernardi G. Long term synaptic depression in the striatum: physiological and pharmacological characterization. *J Neurosci* 1992; 12: 4224–33.
- Cepeda C, Buchwald NA, Levine MS. Neuromodulatory actions of dopamine in the neostriatum are dependent upon the excitatory amino acid receptor subtypes activated. *Proc Natl Acad Sci USA* 1993; 90: 9576–80.
- Cools R, Barker RA, Sahakian BJ, Robbins TW. Enhanced or impaired cognitive function in Parkinson's disease as a function of dopaminergic medication and task demands. *Cereb Cortex* 2001; 11: 1136–43.
- Downes JJ, Roberts AC, Sahakian BJ, Evenden JL, Morris RG, Robbins TW. Impaired extra-dimensional shift performance in medicated and unmedicated Parkinson's disease: evidence for a specific attentional dysfunction. *Neuropsychologia* 1989; 27: 1329–43.
- Eichenbaum HE, Cohen NJ. *From conditioning to conscious recollection: memory systems of the brain*. Oxford: Oxford University Press; 2001.
- Folstein MF, Folstein SE, McHugh PR. 'Mini-mental state': a practical method for grading the cognitive state of patients for the clinician. *J Psychiatr Res* 1975; 12: 189–98.
- Gabrieli JD. Cognitive neuroscience of human memory. *Annu Rev Psychol* 1998; 49: 87–115.
- Gluck MA, Shohamy D, Myers C. How do people solve the 'weather prediction' task?: individual variability in strategies for probabilistic category learning. *Learn Mem* 2002; 9: 408–18.
- Heindel WC, Salmon DP, Shults CW, Walicke PA, Butters N. Neuropsychological evidence for multiple implicit memory systems: a comparison of Alzheimer's, Huntington's and Parkinson's disease patients. *J Neurosci* 1989; 9: 582–7.
- Hoehn MM, Yahr MD. Parkinsonism: onset, progression and mortality. *Neurology* 1967; 17: 427–42.
- Hollerman JR, Tremblay L, Schultz W. Involvement of basal ganglia and orbitofrontal cortex in goal-directed behavior. *Prog Brain Res* 2000; 126: 193–215.
- Horvitz JC. Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 2000; 96: 651–6.
- Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM. Building neural representations of habits. *Science* 1999; 286: 1745–9.
- Knowlton BJ, Squire LR, Gluck MA. Probabilistic classification learning in amnesia. *Learn Mem* 1994; 1: 106–20.
- Knowlton BJ, Mangels JA, Squire LR. A neostriatal habit learning system in humans. *Science* 1996; 273: 1399–402.
- Koivisto M, Portin R, Rinne JO. Perceptual priming in Alzheimer's and Parkinson's diseases. *Neuropsychologia* 1996; 34: 449–57.
- Ljungberg T, Apicella P, Schultz W. Responses of monkey dopamine neurons during learning of behavioral reactions. *J Neurophysiol* 1992; 67: 145–63.
- Myers CE, Shohamy D, Gluck M, Grossman S, Kluger A, Ferris S, et al. Dissociating hippocampal versus basal ganglia contributions to learning and transfer. *J Cogn Neurosci* 2003; 15: 185–93.
- Owen AM, Roberts AC, Hodges JR, Summers BA, Polkey CE, Robbins TW. Contrasting mechanisms of impaired attentional set-shifting in patients with frontal lobe damage or Parkinson's disease. *Brain* 1993; 116: 1159–75.
- Poldrack RA. Imaging brain plasticity: conceptual and methodological issues. *Neuroimage* 2000; 12: 1–13.
- Poldrack RA, Clark J, Pare-Blagoev J, Shohamy D, Creso Moyano J, Myers C, et al. Interactive memory systems in the human brain. *Nature* 2001; 414: 546–50.
- Reber PF, Squire LR. Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behav Neurosci* 1999; 113: 235–42.
- Reynolds JN, Wickens JR. Substantia nigra dopamine regulates synaptic plasticity and membrane potential fluctuations in the rat neostriatum, *in vivo*. *Neuroscience* 2000; 99: 199–203.
- Robbins TW. Refining the taxonomy of memory. *Science* 1996; 273: 1353–4.
- Saint-Cyr JA, Taylor AE, Lang AE. Procedural learning and neostriatal dysfunction in man. *Brain* 1988; 111: 941–59.
- Schultz W. Getting formal with dopamine and reward. *Neuron* 2002; 36: 241–63.
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science* 1997; 275: 1593–9.
- Smith J, Siegart RJ, McDowall J, Abernethy, D. Preserved implicit learning on both the serial reaction time task and artificial grammar in patients with Parkinson's disease. *Brain Cogn* 2001; 45: 378–91.
- Squire LR. Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. In: Schacter DL, Tulving E, editors. *Memory systems*. Cambridge (MA): MIT Press; 1994. p. 203–31.
- Squire LR, Zola-Morgan S. Structure and function of declarative and non-declarative memory systems. *Proc Natl Acad Sci USA* 1996; 93: 13515–22.
- Swainson R, Rogers RD, Sahakian BJ, Summers BA, Polkey CE, Robbins TW. Probabilistic learning and reversal deficits in patients with Parkinson's disease or frontal or temporal lobe lesions: possible adverse effects of dopaminergic medication. *Neuropsychologia* 2000; 38: 596–612.
- Vriezen ER, Moscovitch M. Memory for temporal order and conditional associative-learning in patients with Parkinson's disease. *Neuropsychologia* 1990; 28: 1283–93.
- Wickens JR, Begg AJ, Arbutnot GW. Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex *in vitro*. *Neuroscience* 1996; 70: 1–5.
- Witt K, Nuhman A, Deuschl G. Intact artificial grammar learning in patients with cerebellar degeneration and advanced Parkinson's disease. *Neuropsychologia* 2002; 40: 1534–40.