




Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

The influence of feedback valence in associative learning

Amanda Bischoff-Grethe^a, Eliot Hazeltine^b, Lindsey Bergren^c, Richard B. Ivry^d, Scott T. Grafton^{c,*}

^a Psychology Service, VA San Diego Healthcare System, and Department of Psychiatry, University of California, San Diego, La Jolla, CA 92093, USA

^b Department of Psychology, University of Iowa, Iowa City, IA 52242-1407, USA

^c Center for Cognitive Neuroscience, Dartmouth College, Hanover, NH 03755, USA

^d Department of Psychology, University of California, Berkeley, Berkeley, CA 94720-1650, USA

ARTICLE INFO

Article history:

Received 17 July 2007

Revised 26 August 2008

Accepted 27 August 2008

Available online 12 September 2008

Keywords:

fMRI

Nucleus accumbens

Punishment

Reward

Stimulus–response mapping

Striatum

ABSTRACT

The neural systems engaged by intrinsic positive or negative feedback were defined in an associative learning task. Through trial and error, participants learned the arbitrary assignments of a set of stimuli to one of two response categories. Informative feedback was provided on less than 25% of the trials. During positive feedback blocks, half of the trials were eligible for informative feedback; of these, informative feedback was only provided when the response was correct. A similar procedure was used on negative feedback blocks, but here informative feedback was only provided when the response was incorrect. In this manner, we sought to identify regions that were differentially responsive to positive and negative feedback as well as areas that were responsive to both types of informative feedback. Several regions of interest, including the bilateral nucleus accumbens, caudate nucleus, anterior insula, right cerebellar lobule VI, and left putamen, were sensitive to informative feedback regardless of valence. In contrast, several regions were more selective to positive feedback compared to negative feedback. These included the insula, amygdala, putamen, and supplementary motor area. No regions were more strongly activated by negative feedback compared to positive feedback. These results indicate that the neural areas supporting associative learning vary as a function of how that information is learned. In addition, areas linked to intrinsic reinforcement showed considerable overlap with those identified in studies using extrinsic reinforcers.

© 2008 Elsevier Inc. All rights reserved.

Reinforcement is a fundamental mechanism for shaping behavior. A basic distinction can be made between positive and negative reinforcement. Positive reinforcers, or rewards, serve a number of basic functions. They promote selected behaviors, induce subjective feelings of pleasure and other positive emotions, and maintain stimulus–response associations (Thut et al., 1997). Negative reinforcement also plays an essential role in shaping behavior. Error signals generated during movement can be used to make rapid on-line adjustments (Ito, 2000). Negative reinforcers such as the reprimand of a parent or the loss of money are intended to help change behavior in the future by promoting more acceptable actions or wiser decisions.

An important question for investigations of learning is to establish similarities and differences between neural networks that process positive and negative reinforcement signals. Overlapping systems allow the control of behavior in a bidirectional manner: a positive reward promotes the reinforced behavior, whereas a negative reward attenuates that behavior. On the other hand, the computations required for using positive and negative feedback in terms of modifying synaptic efficiency can be quite different (Albus, 1971;

Schultz and Dickinson, 2000), and these might be performed by distinct neural systems. For example, computational models have suggested different learning methodologies for the basal ganglia, cerebellum, and cerebral cortex (Doya, 1999). The basal ganglia have long been considered to employ reinforcement learning. Neurophysiological recordings have demonstrated that midbrain dopamine neurons initially respond to rewards, but that eventually this response shifts to the conditioned stimulus, suggesting dopamine neurons encode both present and future rewards (Schultz, 1998). Dopaminergic projections terminate upon corticostriatal synaptic spines, thereby modulating synaptic plasticity and enabling learning to take place. In comparison, motor learning with the cerebellum depends upon the error signal generated by climbing fiber input to the Purkinje cell synapses (Thompson et al., 1997). This enables both the on-line correction of individual behaviors as well as long-term adjustments for improved performance. The cerebral cortex, however, may engage in unsupervised learning, based upon the concepts of Hebbian plasticity and the reciprocal connections both within and between regions of the cortex (Sanger, 1989; von der Malsburg, 1973). The integration of these different methodologies may account for a wide variety of motor and cognitive behaviors.

It has been well established through animal and human studies that the basal ganglia play a role in reinforcement and motivation. In humans, neuroimaging has shown, for example, that the ventral

* Corresponding author. Present address: Sage Center for the Study of Mind, Department of Psychology, 3837 Psychology East (Bldg 251), UC Santa Barbara, Santa Barbara, CA 93106, USA. Fax: +1 805 893 4303.

E-mail address: grifton@psych.ucsb.edu (S.T. Grafton).




Fig. 1. Examples of (A) positive, (B) unknown, and (C) negative feedback trials. A single trial lasted 3.8 s, and subjects had 1.8 s to respond once the stimulus (e.g., the letter 'A') was displayed. Subjects received one of three forms of feedback; positive feedback (occurring in the positive feedback blocks), negative feedback (occurring in the negative feedback blocks), or unknown (occurring in both positive and negative feedback blocks).

striatum responds to numerous rewarding stimuli, including cocaine (Breiter et al., 1997), money (Breiter et al., 2001; Delgado et al., 2000; Elliott et al., 2003; Thut et al., 1997), and pleasurable tastes (Berns et al., 2001; McClure et al., 2003a). It has also frequently been shown to be engaged during conditional motor learning (Toni and Passingham, 1999). Other regions involved in reward processing include the amygdala and medial frontal areas, which process both appetitive and aversive stimuli (Baxter and Murray, 2002; Everitt et al., 2003; Seymour et al., 2004), and the insula, which assesses sensory input (Mesulam and Mufson, 1982) and is engaged during uncertainty in decision-making (Casey et al., 2000; Huettel et al., 2005; Paulus et al., 2003, 2005).

In much of this work, the direct manipulation of one type of reinforcer may also, albeit indirectly, provide information about the state of another reinforcer (Breiter et al., 2001; Delgado et al., 2000). Consider the parent who chooses to use only positive feedback to promote good behavior in her child at a restaurant. Whenever the child uses her utensils properly or makes a polite request, the parent offers praise. If the child reaches over the table or eats with her fingers, the parent becomes silent. The child may readily come to recognize the silence as a negative reinforcer; that is, the absence of positive reinforcement serves as a negative reinforcer. The presence of multiple reinforcers, whether explicit or not, makes it difficult to determine if a neural system is preferentially engaged by a particular class of reinforcement signals.

The preceding example also makes clear that much of human behavior is shaped and guided by intrinsic rewards and motivation as well as extrinsic rewards. It is difficult to tease the two apart, as often an intrinsically rewarding event, such as the subjective pleasure from working on a jigsaw puzzle, might be modulated by the outcome of the event, such as solving the puzzle. Both behavioral and neuroimaging reward paradigms have long focused upon external rewards; few studies have relied on intrinsic reward in studying the reward circuitry (Breiter et al., 2001; Elliott et al., 1997).

In this pilot study, we examined the neural systems engaged by positive or negative intrinsic reinforcement signals in an associative learning task. Through trial and error, participants learned the arbitrary assignments of a set of stimuli to one of two response categories. An important feature of the study was that the task was designed to assess how neural responses were influenced by feedback valence under conditions in which the absence of one type of reinforcement (e.g., positive) would not indirectly provide the opposite type of reinforcement (e.g., negative). To this end, informative reinforcement was provided on some trials, whereas on other trials neutral, uninformative feedback was given. The type of reinforcing (informative) feedback, positive or negative, was manipulated between scanning runs. Within a scan, the response to a stimulus might be followed by informative feedback or it might be followed by uninformative feedback. Thus, in the positive reinforcement condition, the absence of positive feedback did not provide

information that the selected response was incorrect (i.e., serve as negative reinforcement). Similarly, the absence of negative feedback in the negative reinforcement condition did not serve as an indirect positive reinforcer. This design allowed us to isolate neural regions associated with positive or negative reinforcement, uncontaminated by the indirect engagement of systems associated with the other form of reinforcement. The basic question in this study focused upon the degree of overlap and difference between neural regions involved in processing positive and negative intrinsic reinforcement signals.

We employed a region of interest (ROI) analysis to establish further evidence that the amygdala, nucleus accumbens, caudate nucleus, and putamen would play a role in intrinsic feedback and would be sensitive to feedback valence, whereas the insula, a region associated with punishment and negative emotionality, would respond to negative feedback. Because we used an association learning paradigm, we hypothesized that two areas linked to stimulus–response learning – the cerebellum and supplementary motor area (SMA) – would be differentially influenced by feedback valence, due to their established involvement in either error learning and stimulus–response remapping (cerebellum) (Bischoff-Grethe et al., 2002; Rushworth et al., 2002), response execution (SMA-proper) or response inhibition (pre-SMA) (Garavan et al., 1999).

Materials and methods

Experimental design

Twelve right-handed participants (five male, seven female) aged 18 to 27 years (average age 20 years) gave written informed consent in accordance with the Dartmouth College human subjects committee. Subjects were told that the study would examine their ability to learn simple response associations through trial and error.

On each trial, the subject saw a single letter stimulus and pressed one of two response keys with either the index or middle finger of the right hand. The stimulus set consisted of all 26 letters, with each letter arbitrarily assigned to one of two response categories (see below). The stimulus was presented for 1.0 s and then replaced by a fixation cross for 0.8 s (Fig. 1). Subjects were instructed that they had to respond during the time of the stimulus presentation. If the response was made within 1.8 s of stimulus onset, feedback was provided (see below). The feedback screen was presented for 1.0 s and then replaced by a fixation cross for 1.0 s until the start of the next trial. The entire trial duration was 3.8 s. In addition to these stimulus-present trials, we also included 23% fixation trials in which no stimulus was presented, and the “+” fixation point remained on the screen for the entire 3.8 s duration. The inclusion of these fixation trials allowed for the characterization of event-specific responses in a pseudorandomly designed rapid presentation of all trial types (Friston et al., 1999b).

Two types of runs were used that differed in terms of feedback. On positive feedback runs, the word “CORRECT” was displayed when a)

the trial was a candidate for informative feedback and b) the response was correct. These trials allowed the subject to learn the correct response association for that stimulus. On average, positive feedback was provided on 27% of the stimulus-present trials. Of the remaining stimulus-present trials, the word “UNKNOWN” was presented as feedback. Similarly, on negative feedback runs, the word “INCORRECT” was displayed when a) the trial was a candidate for informative feedback and b) the response was incorrect; otherwise the word “UNKNOWN” was presented after the response. On average, negative feedback was provided on 18% of the stimulus-present trials on negative feedback runs. The differential occurrence of negative and positive feedback was due to the fact that, as participants learned, the probability for both informative feedback conditions being met was less in the negative feedback condition than in the positive feedback condition. By making the presentation of informative feedback conditional in terms of overall probability as well as the specific response for that trial, subjects were not able to infer the correct response on most of the trials. For example, the absence of positive feedback following a response in a positive run did not imply that the response was incorrect. This design allowed us to compare associative learning on the basis of either positive feedback or negative feedback.

A new subset of sixteen stimuli was selected for each scanning run. During each run, a stimulus was considered “new” if it were being presented for the first time. Subsequent viewings of the same stimulus were either “paired” or “unpaired.” Paired stimuli were those in which a prior trial gave informative feedback (i.e., the stimulus was associated with a specific response); unpaired stimuli were those in which no prior trials had provided informative feedback (i.e., all prior trials were uninformative, meaning the stimulus had yet to be associated with a response). Because it was possible within a single run to view the same stimulus multiple times, an adaptive procedure (described below) was used so that new, paired, and unpaired stimuli were equally likely to receive reinforcement and were distributed evenly throughout the scan. This enabled us to look at behavioral differences based upon levels of knowledge with respect to the stimulus–response (S–R) pairs.

A run consisted of 80 trials. Of these, 24 were fixation trials, randomly distributed among the behavioral trials. The 5 behavioral trial types were defined by the type of stimulus (new, unpaired, or paired) and whether it received informative or uninformative feedback. New trials (whether receiving informative or uninformative feedback) were those in which the stimulus had not been viewed previously. The number of “new” stimuli was distributed across the run by equating them to the number of stimuli previously viewed regardless of prior informative feedback. Unpaired, uninformative trials were those in which only uninformative feedback had been given previously to that stimulus, and no feedback was presented on the current trial. Unpaired, informative trials were those in which only uninformative feedback had been given previously, and informative feedback was presented on the current trial. Because either response could, in theory, be correct on these trials, an on-line algorithm was used to categorize the stimuli after the response to ensure that the desired distribution of trial types was maintained. Thus, the response category for each stimulus was not fixed a priori, but determined after the first response to that stimulus on a trial with informative feedback. Paired, uninformative trials were those in which informative feedback had been given previously, and no feedback was presented on the current trial. Paired, informative trials were those in which informative feedback had been given previously. On these trials, feedback was given only if the participant made the appropriate response (the category assignment for a stimulus was fixed once informative feedback had been presented). Each of these 5 trial types was evenly distributed, so that on any given trial, the 5 trial types were equally likely. The first 10 trials in the run consisted of 6 trials with new stimuli and 4 fixation trials. For the remaining 50 non-fixation trials, each trial type occurred 10 times.

The order of the two feedback conditions was counterbalanced: half of the participants started with four blocks of positive feedback followed by four blocks of negative feedback; the order was reversed for the other participants. Participants were given one practice block prior to their first positive scan and first negative feedback scan. At the end of each run, a learning score was presented to motivate the participants. Perfect performance led to a score of 100, chance was 0, and a score of –100 indicated that all of the responses were incorrect. The program only scored trials on which the subject should have known the correct stimulus–response association based upon feedback from a previous trial.

MRI

Eight fMRI runs of 152 scans each were obtained. Functional MRI was performed with gradient-recalled echoplanar imaging (reaction time, 2000 ms; echo time, 35 ms; flip angle, 90°; 64×64 matrix; 27 5.5 mm contiguous axial slices) on a GE 1.5 T scanner (Kwong et al., 1992; Ogawa et al., 1992). A coplanar T1-weighted structural and a high resolution MRI were obtained for each individual for subsequent spatial normalization.

Statistical analysis

The data were analyzed using Statistical Parametric Mapping (SPM2; Wellcome Department of Cognitive Neurology, London UK) (Friston et al., 1995). Motion correction to the first functional scan was performed within subject using a six-parametric rigid-body transformation. The mean of the motion-corrected images was first coregistered to the individual's high-resolution MRI using mutual information, followed by coregistration of the structural MRI. The images were then spatially normalized to the Montreal Neurologic Institute (MNI) template (Talairach and Tournoux, 1988) by applying a 12-parameter affine transformation followed by a nonlinear warping using basis functions (Ashburner and Friston, 1999). The spatially normalized scans were then smoothed with a 6 mm isotropic Gaussian kernel to accommodate anatomical differences across participants. Six subjects had one volume in which the slices were improperly reconstructed. These volumes were discarded and replaced with average volumes based upon the volumes directly preceding and following them. Due to technical difficulties, two subjects each lost one positive feedback run, and one subject lost two negative feedback runs. Their remaining data was included in the analysis.

The data were first analyzed with the general linear model on an individual subject basis with an event-related design and convolved with the SPM canonical hemodynamic response function with temporal and dispersion derivative terms. Next, a random effects model was performed to make group statistical inferences (Friston et al., 1999a), and contrasts (described below) were applied using a false detection rate correction (Genovese et al., 2002) at $p < 0.05$. We then employed a region of interest (ROI) analysis with small volume correction. Anatomically defined ROIs from an automated atlas (Tzourio-Mazoyer et al., 2002) were used to restrict analyses to the putamen, caudate nucleus, amygdala, insula, supplementary motor area, and cerebellar lobule VI. Because no anatomical definition of the nucleus accumbens was available, a sphere centered at $x, y, z = \pm 10, 8, -4$ with radius = 8 mm was defined. These coordinates have been used in other ROI analyses and are near the ventral striatal foci as defined in previous studies (Breiter et al., 2001; Cools et al., 2002; Delgado et al., 2000). This ROI partially overlapped the MNI ROI for the putamen (left: 21.2% overlap; right: 2.4% overlap); however, given the variability of the nucleus accumbens location we felt this partial overlap was acceptable.

To determine the extent the ROIs were engaged by informative feedback, we implemented a contrast that compared informative

feedback trials to uninformative feedback trials: the runs containing positive feedback were treated separately from those containing negative feedback. That is, we conducted two contrasts, one comparing positive (“CORRECT”) feedback trials to uninformative (“UNKNOWN”) feedback trials from the same blocks, and a second involving a similar contrast for negative (“INCORRECT”) feedback trials. These contrasts included all trials that were followed by feedback, regardless of the overall definition of the trial type (e.g., new, paired, or unpaired).

Results

Behavior

We first examined how well the participants learned the associations under the limited feedback conditions employed in the current study. To this end, we looked at responses on trials to stimuli that had been previously linked to informative feedback (paired trials). Participants responded correctly on $80.69\% \pm 9.89$ (average $\pm$ SD) of the paired trials in the positive feedback condition, and $73.62\% \pm 13.05$ in the negative feedback condition. Both values were significantly greater than a chance response of 50% (positive feedback: $t(11)=10.744$, $p<0.001$; negative feedback: $t(11)=6.271$, $p<0.001$). Although performance was numerically lower in the negative feedback conditions, this difference was not significant (two-tailed paired sample t -test, $t(11)=1.719$, $p=0.114$).

Effects of trial type on reaction time (Fig. 2) were also of interest given that they would reflect differences in the processing strategies applied on new, unpaired, and paired stimuli trials. A 3×2 ANOVA with trial type (new, unpaired, paired) and feedback condition (positive or negative) as factors revealed a significant difference in the median reaction times for trial type ($F(2, 10)=16.405$, $p=0.001$) and for condition ($F(1, 11)=6.254$, $p=0.029$) but no interaction between these factors ($F<1$). Thus, it appeared that the effects of feedback on RT were the same for the two types of feedback.

Neural activation

We performed separate comparisons of each type of informative feedback to the uninformative feedback trials from the same runs using random-effects analysis with small volume correction in SPM. Informative feedback trials, whether positive or negative, were




Fig. 2. Median reaction times for the positive and negative feedback conditions by trial type. New trials were stimulus–response pairs that were presented for the first time within the block; unpaired trials were previously presented pairs that had only received uninformative feedback; paired trials were previously presented pairs which had received informative feedback. An ANOVA with trial type and feedback condition as factors revealed significant differences in the reaction times across trial types.

associated with significant changes within several of our regions of interest, including the bilateral caudate nucleus, right cerebellar lobule VI, bilateral insula, bilateral nucleus accumbens, left putamen, and bilateral SMA (Table 1, Fig. 3). In general, these ROIs tended to activate more strongly (in extent and/or peak T score) in the positive comparison than in the negative comparison. Other regions were specifically activated by one form of feedback but not the other. For runs containing positive feedback, ROI activation was greater on trials with informative feedback for the bilateral amygdala, left cerebellar lobule VI, and right putamen. All of the active ROIs for the negative feedback trials were either more strongly activated or equally activated for the positive feedback trials.

In a direct comparison of the two types of informative feedback trials, the activation in the amygdala, insula, putamen and SMA was reliably greater on positive feedback trials. Interestingly, the bilateral insula and right putamen showed a significant difference in the uninformative feedback > negative feedback contrast, with a larger BOLD response on the uninformative trials. The insula response was located within the posterior portion, whereas it was more anterior in the negative feedback > uninformative feedback and positive feedback > uninformative feedback comparisons.

In order to determine if there were any run effects or feedback $\times$ run interactions, the mean beta values for the predefined, hypothesis driven ROIs were entered into a 2 (feedback type) $\times 4$ (run) ANOVA. Because the two feedback conditions were not presented within the same functional runs, the feedback conditions were restricted to either positive feedback vs. uninformative feedback or negative feedback vs. uninformative feedback. None of the ROIs exhibited a main effect for run. Similar to the small volume correction results, most regions showed a main effect of feedback (all $p<0.05$) (Fig. 4). The bilateral amygdala showed a quadratic run effect for the negative feedback condition $\times$ run interaction (left: $F(3, 30)=3.257$, $p=0.035$; right: $F(3, 30)=4.154$, $p=0.014$).

Discussion

Learning requires the use of feedback. Behavior can be positively reinforced to promote a desired behavior; alternatively, behavior can be negatively reinforced in an effort to decrease the likelihood of an undesirable behavior. In most empirical studies of reinforcement, positive and negative reinforcements are intermixed, at least implicitly. In the current study, we introduced a task in which feedback was completely uninformative on a large proportion of trials. In this way, we sought to isolate the neural responses to positive and negative reinforcement as well as compare the neural systems involved when the learning was based on one type of reinforcement or the other.

Our results demonstrated two main findings: First, all our ROIs were significantly responsive to positive feedback, and a majority of them were likewise responsive to negative feedback. This suggests that the valence of feedback may be less important than the motivational significance of the feedback. In our association task, informative feedback was highly significant to learning associations, whereas noninformative feedback was not. However, our task limited potential motor responses to just two; if additional responses were possible we may have seen a greater discrimination between positive and negative feedback. Second, our findings demonstrated that regions commonly associated with extrinsic reinforcers also respond to symbolic cues. This supports the notion that the context in which reinforcement is presented influences its perceived significance.

Neural regions responding to informative feedback

Most of our predetermined ROIs were associated with informative feedback, regardless of valence, when compared to uninformative feedback. This observation held true for regions typically associated

Table 1

Activation peaks within the ROIs for the contrasts (positive feedback>uninformative feedback; negative feedback>uninformative feedback; uninformative feedback>negative feedback; and (positive feedback–uninformative feedback)>(negative feedback–uninformative feedback))

	Left hemisphere						Right hemisphere					
	Cluster size	P (FDR-corr)	T score	MNI coordinates (mm)			Cluster size	p (FDR-corr)	T score	MNI coordinates (mm)		
<i>Positive feedback > uninformative feedback</i>												
Amygdala	111	<0.001	7.68	-22	-4	-12	165	<0.001	8.28	20	-2	-12
Caudate nucleus	816	<0.001	12.56	-12	14	-8	792	<0.001	12.80	14	14	-12
Cerebellar lobule VI	701	0.021	6.03	-36	-50	-24	943	0.013	5.65	32	-46	-22
Insula	968	0.002	7.51	-28	16	-8	968	0.030	5.74	50	8	-6
	18	0.035	2.61	-32	-18	22						
Nucleus accumbens	122	<0.001	10.43	-12	12	-8	123	<0.001	9.21	12	12	-8
Putamen	972	<0.001	10.43	-12	12	-8	867	<0.001	9.82	16	14	-10
SMA	664	0.002	7.27	0	10	54	831	0.001	6.99	6	14	46
	46	0.032	2.67	-14	-14	64	122	0.009	3.52	14	-22	54
<i>Negative feedback > uninformative feedback</i>												
Caudate nucleus	258	0.044	5.05	-14	16	0	530	0.017	5.15	10	18	-8
Cerebellar lobule VI							744	0.011	6.07	32	-74	-24
Insula	628	0.004	6.74	-28	24	-8	600	0.011	5.83	48	14	-2
Nucleus accumbens	89	0.012	4.46	-14	12	-2	98	0.011	4.30	14	4	-2
Putamen	296	0.009	6.09	-20	20	-6						
SMA	821	0.003	6.23	-2	2	54	721	0.006	6.21	2	0	52
<i>Uninformative feedback > negative feedback</i>												
Insula	559	0.024	5.91	-38	-14	2	398	0.043	5.76	38	-10	4
Putamen							177	0.048	5.54	36	-10	2
<i>(Positive feedback – uninformative feedback) > (negative feedback – uninformative feedback)</i>												
Amygdala	96	0.022	4.74	-24	0	-12	134	0.003	6.00	22	-2	-14
Insula	680	<0.001	11.61	-38	-4	8	437	0.009	7.33	38	-8	8
Putamen	734	<0.001	10.64	-28	-10	8	637	0.001	9.95	34	-2	2
SMA	269	0.010	6.82	-6	-20	52	278	0.014	6.28	12	-22	54

Significance was set at $p < 0.05$ with small volume correction for the regions of interest. Coordinates are based on the Montreal Neurological Institute brain template.

with positive feedback (caudate nucleus, nucleus accumbens) as well as those commonly linked to negative feedback (insula). The nucleus accumbens, for example, has traditionally been associated with anticipation of impending reward, such that only positive or more favorable cues produce an increased response, whereas negative or

undesirable cues have no effect (Breiter et al., 2001; Schultz, 2000). However, growing evidence suggests that salience may also invoke a response (Cooper and Knutson, 2008; Tricomi et al., 2004; Zink et al., 2004), such that negative events, if motivationally salient, will also induce a response in the same region (Seymour et al., 2004).




Fig. 3. Statistical parametric maps of the functional localization within the coronal plane ($p < 0.05$, FDR-corrected). a) Regions responding to positive feedback>uninformative feedback included the bilateral cerebellum ($z = -50$), the bilateral amygdala ($z = -4$), and the nucleus accumbens, caudate nucleus, putamen, insula, and SMA bilaterally ($z = +12$). b) For the negative feedback>uninformative feedback comparison, activated regions included the right cerebellum ($z = -74$) and the bilateral nucleus accumbens, caudate nucleus, putamen, insula, and SMA ($z = +16$).




Fig. 4. Region of interest analysis comparing mean beta values defined at each location for the contrasts positive feedback–uninformative feedback and negative feedback–uninformative feedback. Beta values were determined by the general linear model in SPM for each individual. Beta weights were averaged across an anatomical region of interest, and then averaged across subjects. Anatomical regions were as defined in an automated atlas (Tzourio-Mazoyer et al., 2002), except for the nucleus accumbens, which was defined as a spherical region of interest ($x, y, z = \pm 10, 8, -4$; radius = 8 mm). Regions with statistically significant differences between mean beta value and baseline are indicated as follows: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.005$.

The caudate nucleus is known for its involvement in action contingency (Knutson and Cooper, 2005; Tricomi et al., 2004). The nonhuman primate literature has supported a role for the caudate nucleus in reward association learning (Kawagoe et al., 1998; Schultz, 1998; Schultz et al., 1998). It has strong connections with orbitofrontal and prefrontal cortices; these areas are associated with sensory rewards (Alexander et al., 1990; Kringelbach and Rolls, 2004; Rolls, 2000). As such, it may play a central role in goal-directed behavior. One functional hypothesis concerning the caudate nucleus is that it is involved in the detection of actual and predicted rewards (Delgado et al., 2004). Interestingly, the caudate nucleus responded to negative feedback as well as to positive feedback in the present study. This differs from prior studies which show increased activation within the caudate nucleus for positive feedback and little change in response to negative feedback (Delgado et al., 2004; Seger and Cincotta, 2005).

Although this is a small study and we should therefore interpret this finding with caution, one possible explanation for the significant response in caudate nucleus following negative reinforcement may be related to the context of our task compared to others studies providing informative feedback. In studies showing transient changes within the caudate nucleus (Delgado et al., 2000; Elliott et al., 2000), a positive reward of some kind was always possible within each trial. In our study, however, positive reinforcement was never provided in the negative feedback runs and, correspondingly, there should never have been an expectation of positive reinforcement. Moreover, the absence of feedback in this condition was ambiguous. Thus, these results suggest that the caudate nucleus response may have been modulated by the cognitive aspect of the task. There is some support for this in the literature, where a response has been associated with either salience or contingency detection (Tricomi et al., 2004; Zink et al.,

2004). An intriguing corollary here is that negative reinforcement is not equivalent to the absence of positive reinforcement, at least when the context contains considerable uncertainty (e.g., the high percentage of uninformative trials).

In contrast to the activation in the caudate, the activation within the putamen varied as a function of the type of feedback. While the left ventral putamen responded to all kinds of informative feedback, the right ventral putamen responded only to positive feedback when compared to uninformative feedback. Unlike the caudate nucleus, which receives prefrontal inputs and is associated with high level cognition, the putamen receives a mixture of prefrontal, premotor and motor inputs and is more closely associated with the motoric aspect of S-R learning (Middleton and Strick, 2000, 2001). Overall, these results are consistent with response reinforcement.

The anterior insula responded to informative feedback regardless of valence. While activation in the insula has been associated with negative events including pain (Craig, 2003), rejection within a social context (Eisenberger et al., 2003), and punishment in reward paradigms (Abler et al., 2005; O'Doherty et al., 2003; Ullsperger and von Cramon, 2003), it has also been implicated in decision-making and behavioral monitoring (Casey et al., 2000; Huettel et al., 2005; Paulus et al., 2003, 2005). While these processes would be engaged following informative reinforcement, the current results would indicate that the engagement of anterior insula here may be greater when the reinforcement is negative. It is possible that subjects experienced emotional arousal in relation to informative feedback, and that this drove the anterior insula activation.

The right cerebellar lobule VI was also activated following both positive and negative feedback. Both the cerebellum and the putamen have long been associated with motor learning. However, computational models of these two systems typically emphasize different feedback methodologies. Basal ganglia learning models focus upon the reward signaling properties of dopamine (e.g., McClure et al., 2003b), whereas cerebellar learning models focus upon error signals generated by climbing fiber input (e.g., Albus, 1971; Marr, 1969). Neuropsychological and neuroimaging studies have suggested that the cerebellum plays a critical role in higher cognition rather than being strictly related to motor learning and motor control. One functional account of these effects is that the cerebellum may be involved in the planning and rehearsal of actions as part of a process that anticipates potential outcomes (Bischoff-Grethe et al., 2002). It may be that informative feedback of either valence engages these operations as the individual attempts to associate a response to a stimulus that just received informative feedback.

Learning with intrinsic cues

An important feature of this task is that even symbolic cues, associated with intrinsic reinforcement, activated regions commonly associated with more extrinsic feedback. Much of the reward literature has used some form of extrinsic reinforcement. Neuroimaging studies in humans have demonstrated responses within the limbic circuit for both primary (e.g., appetitive stimuli) (Berns et al., 2001; McClure et al., 2007) and secondary (e.g., monetary gain or loss) reinforcers (Breiter et al., 2001; Delgado et al., 2004). However, the valence of these reinforcers may also influence neural response due to the context in which they are presented. Reward processing systems determine whether an outcome is favorable based upon the range of possible outcomes (Breiter et al., 2001; Nieuwenhuis et al., 2005); an outcome of \$0 is undesirable when other outcomes provide a monetary reward, but is desirable when only negative outcomes are possible. This implies that, within the right framework, symbolic cues could indeed drive the reward system. Our results support this hypothesis by demonstrating a strong limbic response to informative feedback trials.

The receipt of informative feedback was based upon action-contingencies: in the positive feedback condition, a correct stimulus–

response association was necessary in order to allow the possibility of positive feedback. In the negative condition, the inverse was true. Action-contingencies may also have played a role in the involvement of the reward circuitry. Without the association learning aspect, it is likely that our symbolic cues are meaningless, and therefore would fail to induce reward-related responses. This is in contrast to primary and secondary reinforcement paradigms, in which simply attending to the stimuli activated the reward system (Breiter et al., 2001).

Successful performance of an associative task requires forming and maintaining the correct visuomotor association. Overall, participants responded more slowly for unpaired and paired trials than they did for new trials. The effect of trial type likely reflects response retrieval: unpaired and paired trials may have been slower because they involved retrieval processes that determined whether the stimulus had been associated with a particular response on a previous presentation. Recognition that a stimulus is novel would have precluded or aborted such retrieval processes. For all trial types, participants were slower to respond in the negative feedback condition, consistent with the finding that learning tended to be more difficult with negative feedback.

The right dorsoposterior putamen and bilateral posterior insula both demonstrated an increased response for uninformative feedback when compared to negative feedback. This peak activation was distinct from that seen for informative feedback when compared to uninformative feedback. The posterior insula is associated with motor sensation and other homeostatic information (Craig, 2003). This region responds to predictable aversive events, in contrast to the anterior insula, which responds to unpredictable aversive events (Carlsson et al., 2006). The insula has bidirectional connections to both the amygdala and to the nucleus accumbens (Reynolds and Zahm, 2005) as well as the orbitofrontal cortex (Ongur and Price, 2000). In the amygdala, neurons can rapidly adjust their activity to reflect both positive and negative value of an external stimulus, which is predictive of how fast monkeys learn to respond to a stimulus (Paton et al., 2006). Therefore, the insula is centrally placed to receive information about the salience (both appetitive and aversive) and relative value of the stimulus environment and integrate this information with the effect that these stimuli may have on the body state. As anxiety can induce homeostatic changes in an individual, we suggest that the receipt of uninformative feedback in the negative condition induced more anxiety due to its context. Unless an S-R pairing had previously received negative feedback, subjects did not know the correct response in this condition. This uncertainty may have increased anxiety levels regarding task performance.

It is possible that, despite efforts for an unbiased, neutral stimulus, the uninformative feedback trials might be viewed negatively and be a source of frustration to the participant. If uninformative feedback trials had negative connotations, we would expect a reduced likelihood of activation in comparison of negative informative feedback to uninformative feedback. However, both positive and negative feedback, when compared to uninformative feedback, produced responses that 1) support the proposal that positive informative feedback is rewarding, and 2) show salience is also critical, regardless of feedback valence. It is therefore likely that uninformative feedback produced little negative interference that would confound our results.

Conclusion

An important form of associative learning requires the arbitrary linkage of an action to a stimulus. Most studies of associative learning use a trial and error approach (Toni and Passingham, 1999; Toni et al., 2001) in which feedback is provided on every trial. This procedure makes it difficult to dissociate the effects of feedback unless a large temporal gap separates the stimulus, response, and feedback signals.

By separating blocks that used positive and negative feedback in combination with a partial reinforcement schedule, we were able to identify areas associated with specific types of feedback. While some neural regions were associated with these processes independent of whether the feedback was positive or negative, we also observed regions that were sensitive to the valence of feedback. This experiment also demonstrates that regions commonly associated with extrinsic rewards and punishments can also respond to more intrinsic forms of feedback.

Acknowledgments

Supported by PHS grants NS33504 and NS40813. Corresponding author: Scott T. Grafton, M.D., Sage Center for the Study of Mind, Department of Psychology, Psychology East, UC Santa Barbara, Santa Barbara CA 93106.

References

- Abler, B., Walter, H., Erk, S., 2005. Neural correlates of frustration. *Neuroreport* 16, 669–672.
- Albus, J.S., 1971. The theory of cerebellar function. *Math. Biosci.* 10, 25–61.
- Alexander, G.E., Crutcher, M.D., DeLong, M.R., 1990. Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor, “prefrontal” and “limbic” functions. In: Uylings, H.B.M., Eden, C.G.V., Bruin, J.P.C.D., Corner, M.A., Feenstra, M.G.P. (Eds.), *Progress in Brain Research*. Elsevier Science Publishers B. V., New York, pp. 119–146.
- Ashburner, J., Friston, K.J., 1999. Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* 7, 254–266.
- Baxter, M.G., Murray, E.A., 2002. The amygdala and reward. *Nat. Rev., Neurosci.* 3, 563–573.
- Berns, G.S., McClure, S.M., Pagnoni, G., Montague, P.R., 2001. Predictability modulates human brain response to reward. *J. Neurosci.* 21, 2793–2798.
- Bischoff-Grethe, A., Ivry, R.B., Grafton, S.T., 2002. Cerebellar involvement in response reassignment rather than attention. *J. Neurosci.* 22, 536–553.
- Breiter, H.C., Gollub, R.L., Weisskoff, R.M., Kennedy, D.N., Makris, N., Berke, J.D., Goodman, J.M., Kantor, H.L., Gastfriend, D.R., Riorden, J.P., Mathew, R.T., Rosen, B.R., Hyman, S.E., 1997. Acute effects of cocaine on human brain activity and emotion. *Neuron* 19, 591–611.
- Breiter, H.C., Aharon, I., Kahneman, D., Dale, A., Shizgal, P., 2001. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30, 619–639.
- Carlsson, K., Andersson, J., Petrovic, P., Petersson, K.M., Ohman, A., Ingvar, M., 2006. Predictability modulates the affective and sensory-discriminative neural processing of pain. *Neuroimage* 32, 1804–1814.
- Casey, B.J., Thomas, K.M., Welsh, T.F., Badgaiyan, R.D., Eccard, C.H., Jennings, J.R., Crone, E.A., 2000. Dissociation of response conflict, attentional selection, and expectancy with functional magnetic resonance imaging. *Proc. Natl. Acad. Sci. U. S. A.* 97, 8728–8733.
- Cools, R., Clark, L., Owen, A.M., Robbins, T.W., 2002. Defining the neural mechanisms of probabilistic reversal learning using event-related functional magnetic resonance imaging. *J. Neurosci.* 22, 4563–4567.
- Cooper, J.C., Knutson, B., 2008. Valence and salience contribute to nucleus accumbens activation. *Neuroimage* 39, 538–547.
- Craig, A.D., 2003. Interoception: the sense of the physiological condition of the body. *Curr. Opin. Neurobiol.* 13, 500–505.
- Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C., Fiez, J.A., 2000. Tracking the hemodynamic responses to reward and punishment in the striatum. *J. Neurophysiol.* 84, 3072–3077.
- Delgado, M.R., Stenger, V.A., Fiez, J.A., 2004. Motivation-dependent responses in the human caudate nucleus. *Cereb. Cortex* 14, 1022–1030.
- Doya, K., 1999. What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Netw.* 12, 961–974.
- Eisenberger, N.I., Lieberman, M.D., Williams, K.D., 2003. Does rejection hurt? An fMRI study of social exclusion. *Science* 302, 290–292.
- Elliott, R., Frith, C.D., Dolan, R.J., 1997. Differential neural response to positive and negative feedback in planning and guessing tasks. *Neuropsychologia* 35, 1395–1404.
- Elliott, R., Friston, K.J., Dolan, R.J., 2000. Dissociable neural responses in human reward systems. *J. Neurosci.* 20, 6159–6165.
- Elliott, R., Newman, J.L., Longe, O.A., Deakin, J.F., 2003. Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study. *J. Neurosci.* 23, 303–307.
- Everitt, B.J., Cardinal, R.N., Parkinson, J.A., Robbins, T.W., 2003. Appetitive behavior: impact of amygdala-dependent mechanisms of emotional learning. *Ann. N. Y. Acad. Sci.* 985, 233–250.
- Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.-B., Frith, C.D., Frackowiak, R.S.J., 1995. Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210.
- Friston, K.J., Holmes, A.P., Worsley, K.J., 1999a. How many subjects constitute a study? *Neuroimage* 10, 1–5.
- Friston, K.J., Zarahn, E., Josephs, O., Henson, R.N., Dale, A.M., 1999b. Stochastic designs in event-related fMRI. *Neuroimage* 10, 607–619.
- Garavan, H., Ross, T.J., Stein, E.A., 1999. Right hemispheric dominance of inhibitory control: an event-related functional MRI study. *Proc. Natl. Acad. Sci. U. S. A.* 96, 8301–8306.
- Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15, 870–878.
- Huettel, S.A., Song, A.W., McCarthy, G., 2005. Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *J. Neurosci.* 25, 3304–3311.
- Ito, M., 2000. Mechanisms of motor learning in the cerebellum. *Brain Res.* 886, 237–245.
- Kawagoe, R., Takikawa, Y., Hikosaka, O., 1998. Expectation of reward modulates cognitive signals in the basal ganglia. *Nat. Neurosci.* 1, 411–416.
- Knutson, B., Cooper, J.C., 2005. Functional magnetic resonance imaging of reward prediction. *Curr. Opin. Neurol.* 18, 411–417.
- Kringelbach, M.L., Rolls, E.T., 2004. The functional neuroanatomy of the human orbitofrontal cortex: evidence from neuroimaging and neuropsychology. *Prog. Neurobiol.* 72, 341–372.
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., et al., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. U. S. A.* 89, 5675–5679.
- Marr, D., 1969. A theory of cerebellar cortex. *J. Physiol.* 202, 437–470.
- McClure, S.M., Berns, G.S., Montague, P.R., 2003a. Temporal prediction errors in a passive learning task activate human striatum. *Neuron* 38, 339–346.
- McClure, S.M., Daw, N.D., Montague, P.R., 2003b. A computational substrate for incentive salience. *Trends Neurosci.* 26, 423–428.
- McClure, S.M., Ericson, K.M., Laibson, D.I., Loewenstein, G., Cohen, J.D., 2007. Time discounting for primary rewards. *J. Neurosci.* 27, 5796–5804.
- Mesulam, M.M., Mufson, E.J., 1982. Insula of the old world monkey. III: efferent cortical output and comments on function. *J. Comp. Neurol.* 212, 38–52.
- Middleton, F.A., Strick, P.L., 2000. Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brains Res. Rev.* 31, 236–250.
- Middleton, F.A., Strick, P.L., 2001. A revised neuroanatomy of frontal-subcortical circuits. In: Lichter, D.G., Cummings, J.L. (Eds.), *Frontal-Subcortical Circuits in Psychiatric and Neurological Disorders*. Guilford Press, New York, pp. 44–58.
- Nieuwenhuis, S., Heslenfeld, D.J., Alting von Geusau, N.J., Mars, R.B., Holroyd, C.B., Yeung, N., 2005. Activity in human reward-sensitive brain areas is strongly context dependent. *Neuroimage* 25, 1302–1309.
- O’Doherty, J., Critchley, H., Deichmann, R., Dolan, R.J., 2003. Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J. Neurosci.* 23, 7931–7939.
- Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc. Natl. Acad. Sci. U. S. A.* 89, 5951–5955.
- Ongur, D., Price, J.L., 2000. The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cereb. Cortex* 10, 206–219.
- Paton, J.J., Belova, M.A., Morrison, S.E., Salzman, C.D., 2006. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* 439, 865–870.
- Paulus, M.P., Hozack, N., Frank, L., Brown, G.G., Schuckit, M.A., 2003. Decision making by methamphetamine-dependent subjects is associated with error-rate-independent decrease in prefrontal and parietal activation. *Biol. Psychiatry* 53, 65–74.
- Paulus, M.P., Feinstein, J.S., Leland, D., Simmons, A.N., 2005. Superior temporal gyrus and insula provide response and outcome-dependent information during assessment and action selection in a decision-making situation. *Neuroimage* 25, 607–615.
- Reynolds, S.M., Zahm, D.S., 2005. Specificity in the projections of prefrontal and insular cortex to ventral striatopallidum and the extended amygdala. *J. Neurosci.* 25, 11757–11767.
- Rolls, E.T., 2000. The orbitofrontal cortex and reward. *Cereb. Cortex* 10, 284–294.
- Rushworth, M.F., Hadland, K.A., Paus, T., Sipila, P.K., 2002. Role of the human medial frontal cortex in task switching: a combined fMRI and TMS study. *J. Neurophysiol.* 87, 2577–2592.
- Sanger, T., 1989. Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Netw.* 2, 459–473.
- Schultz, W., 1998. Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27.
- Schultz, W., 2000. Multiple reward signals in the brain. *Nat. Rev., Neurosci.* 1, 199–207.
- Schultz, W., Dickinson, A., 2000. Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500.
- Schultz, W., Tremblay, L., Hollerman, J.R., 1998. Reward prediction in primate basal ganglia and frontal cortex. *Neuropharmacology* 37, 421–429.
- Seger, C.A., Cincotta, C.M., 2005. The roles of the caudate nucleus in human classification learning. *J. Neurosci.* 25, 2941–2951.
- Seymour, B., O’Doherty, J.P., Dayan, P., Koltzenburg, M., Jones, A.K., Dolan, R.J., Friston, K.J., Frackowiak, R.S., 2004. Temporal difference models describe higher-order learning in humans. *Nature* 429, 664–667.
- Talairach, J., Tournoux, P., 1988. *Co-planar Stereotaxic Atlas of the Brain*. Thieme Medical Publishers, New York.
- Thompson, R.F., Bao, S., Chen, L., Cipriano, B.D., Grethe, J.S., Kim, J.J., Thompson, J.K., Tracy, J.A., Wengler, M.S., Krupa, D.J., 1997. Associative learning. *Int. Rev. Neurobiol.* 41, 151–189.

- Thut, G., Schultz, W., Roelcke, U., Nienhusmeier, M., Missimer, J., Maguire, R.P., Leenders, K.L., 1997. Activation of the human brain by monetary reward. *NeuroReport* 8, 1225–1228.
- Toni, I., Passingham, R.E., 1999. Prefrontal-basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: a PET study. *Exp. Brain Res.* 127, 19–32.
- Toni, I., Ramnani, N., Josephs, O., Ashburner, J., Passingham, R.E., 2001. Learning arbitrary visuomotor associations: temporal dynamic of brain activity. *NeuroImage* 14, 1048–1057.
- Tricomi, E.M., Delgado, M.R., Fiez, J.A., 2004. Modulation of caudate activity by action contingency. *Neuron* 41, 281–292.
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15, 273–289.
- Ullsperger, M., von Cramon, D.Y., 2003. Error monitoring using external feedback: specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by functional magnetic resonance imaging. *J. Neurosci.* 23, 4308–4314.
- von der Malsburg, C., 1973. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85–100.
- Zink, C.F., Pagnoni, G., Martin-Skurski, M.E., Chappelow, J.C., Berns, G.S., 2004. Human striatal responses to monetary reward depend on saliency. *Neuron* 42, 509–517.