THE BASAL GANGLIA AND TRAINING OF ARITHMETIC FLUENCY

by

Andrea Ponting

B.S., University of Pittsburgh, 2007

Submitted to the Graduate Faculty of

Arts and Sciences in partial fulfillment

of the requirements for the degree of

Master of Science

University of Pittsburgh

2010

UNIVERSITY OF PITTSBURGH

SCHOOL OF ARTS AND SCIENCES

This thesis was presented

by

Andrea Ponting

It was defended on

April 14, 2010

and approved by

Beatriz Luna, Associate Professor, Departments of Psychiatry and Psychology

Christian Schunn, Associate Professor, Department of Psychology

Thesis Director: Julie Fiez, Professor, Departments of Psychology and Neuroscience

THE BASAL GANGLIA AND TRAINING OF ARITHMETIC FLUENCY

Andrea Ponting, M.S.

University of Pittsburgh, 2010

The role of dopamine neurons in reward processing is well-established, as is the observation of reward-related responses in the striatum, a region to which these midbrain dopamine neurons project. The reward-prediction error signals generated in the midbrain may play a role in the striatum in learning, as they help to shape expectations about future events based on prior experiences. The goal of the current experiment was to use principles of striatal function in order to optimize learning in an arithmetic domain. We created a training program that we believed would lead to increased arithmetic fluency by maximally engaging the striatum, through the use of contingent feedback, uncertainty regarding performance, and incentives for correct responses. Both experimental and control participants, who completed training focusing on arithmetic calculation and digit-entry respectively, showed improvement on a task involving the addition of a double-digit and a single-digit number following training, as successful performance on the task required accurate computations and entry of the solution within a narrow response window. We conducted functional magnetic resonance imaging before and after training while participants performed this task, in order to examine the effect of feedback on activity in the caudate nucleus and to determine if learning signals generated by the striatum during arithmetic training are able to modify quantity representations in parietal cortex. Results indicated activation of both the caudate nucleus and the horizontal regions of the intraparietal sulcus (hIPS). Activation of the caudate nucleus replicated previous work, as it showed the prototypical pattern of activity that

distinguished between positive and negative feedback. Activation of the hIPS region was not surprising, given the focus on arithmetic calculation, but this region also exhibited feedback-sensitive activation that differed between sessions and groups, possibly indicating the common influence of a reinforcement learning system.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1.0 INTRODUCTION

The role of dopamine neurons in reward processing is well-established, as is the observation of reward-related responses in the striatum, a region to which these midbrain dopamine neurons project. The reward-prediction error signals generated in the midbrain may play a role in the striatum in learning, as they help to shape expectations about future events based on prior experiences. The goal of the experiment presented here is to use principles of striatal function in order to optimize learning in an arithmetic domain. We created a training program that we believed would lead to increased arithmetic fluency by maximally engaging the striatum. We conducted functional magnetic resonance imaging (fMRI) before and after training to examine the effect of feedback on activity in the caudate and to determine if learning signals generated by the striatum during arithmetic training are able to modify quantity representations in parietal cortex.

1.1 DOPAMINE AND REWARD PROCESSING

Dopaminergic neurons, located in the pars compacta of the substantia nigra and the ventral tegmental area within the midbrain, have been implicated in the processing of reward-related information. The earliest, and most powerful, evidence for the importance of dopamine neurons in reward processing came from electrophysiological studies in primates. These midbrain dopamine neurons exhibit a short, phasic activation following the presentation of a reward, as

well as visual or auditory stimuli that predict reward. Dopamine neurons, as well as some neurons in orbitofrontal cortex, have been shown to be best activated by unpredicted rewards (Schultz, 1998). When a reward is preceded by a well-learned, predictive stimulus, dopamine neurons are not activated following the reward, but instead respond following the predictive stimulus. When a predicted reward fails to occur, the activity of these neurons is depressed at the time at which the reward would have occurred (Mirenowicz & Schultz, 1994; Tremblay & Schultz, 2000).

Reward-directed learning depends on the predictability of the reward, as the occurrence of the reward has to be surprising or unpredicted for a stimulus or action to be learned. The degree to which a reward cannot be predicted is indicated by the discrepancy between the reward obtained for a given behavioral action and the reward that was predicted to occur as a result of that action. This is known as the prediction error and underlies a class of error-driven learning mechanisms (Schultz, 2000). A closer examination of the properties of the phasic dopamine response suggests that it might encode a reward prediction error rather than reward per se. As described above, the phasic dopamine activity is enhanced by surprising rewards, but not by those that are fully predicted, and it is depressed by the omission of predicted rewards (Schultz, 2000).

1.2 THE BASAL GANGLIA AND REWARD PROCESSING

Reward-prediction error signals reported by a phasic dopamine response could play an important role during learning. The striatum, the main input structure of the basal ganglia, receives widespread projections from cortex, as well as from dopamine neurons in the pars compacta of the substantia nigra and the ventral tegmental area. The reward-prediction error signals generated by dopamine neurons are thought to serve as teaching signals that modify

cortico-basal ganglia-thalamic processing loops. With learning, these loops become able to accurately predict the value of a current stimulus so that it can drive the appropriate behavioral response (Doya, 2000). Given the tight interconnection between midbrain dopamine neurons and the striatum, it is not surprising that reward-related activity has also been observed in the striatum.

1.2.1 Primates

Previous work in nonhuman primates has shown that striatal neurons exhibit two different forms of reward-related activity. These neurons show expectation-related activations that begin shortly after a reward-predicting stimulus and terminate after a reward is delivered, and detection-related responses that follow the delivery of a reward (Hikosaka et al. 1989).

Some striatal neurons have been found to fire more vigorously for preferred rewards, and also modulate their activity based on the magnitude of the reward. Some of these neurons respond preferentially to certain types of rewards over others (e.g., apple juice vs. water), differentiating them from the dopaminergic neurons which do not discriminate between different rewards (Schultz et al., 1998). One study involving a visually-guided saccade task in monkeys found that the preferred direction of neurons in the caudate nucleus, a structure in the dorsal striatum, changed as a function of which direction was rewarded, showing that the activity of caudate neurons can be modulated over short time periods through expectation of reward (Kawagoe et al., 1998). Another study, involving a delayed go/no-go task in monkeys, found task-related activation in the caudate nucleus when a liquid reward was delivered at the end of the trial, but not when a sound reinforcement was given (Hollerman et al., 1998).

1.2.2 Humans

The work in nonhuman primates provided a springboard for fMRI work, which has worked to

further elucidate the reward circuitry and how the basal ganglia is involved in reward-

processing. Early work in humans showed that reward-related responses could be detected with

fMRI, and that the locus of these responses were consistent with the circuitry revealed by

primate work. These fMRI studies have shown that the caudate nucleus is activated by primary

rewards (Berns et al., 2001) as well as extrinsic monetary rewards (Delgado et al., 2000; Elliott

et al., 2000; Knutson et al., 2000).

In a study conducted by Berns, McClure, Pagnoni and Montague (2001), subjects

received small amounts of orally delivered fruit juice or water, in either a predictable or

unpredictable manner, while in an fMRI scanner. The medial orbitofrontal cortex and the

nucleus accumbens, a ventral striatal structure, showed greater activation during the

unpredictable run compared to the predictable run. These two brain regions both receive

projections from the dopamine neurons of the midbrain. The results from this experiment

converge with those from the monkey literature to illustrate that neurons in the striatum are

more responsive to unpredictable errors, though this enhancement may be due to errors in

reward prediction.

Previous studies have also investigated striatal learning mechanism during goal-directed

actions (Delgado et al., 2000; Tricomi et al., 2004; Tricomi et al., 2006). During a simple

guessing task, subjects attempted to accurately guess whether the value of a "card" was higher

or lower than five on each trial, earning $1 for each correct guess and losing $0.50 for each

incorrect guess. Results from Delgado, et al. indicate that the striatum is engaged during the

choice period, in which participants selected a response. In the subsequent period during which

feedback was delivered, activity in the caudate nucleus differentiates between reward and

punishment trials. Sustained activation is seen following feedback indicating a monetary gain. However, following feedback indicating a monetary loss, activity in the caudate tends to increase rapidly, and then decrease below baseline (Delgado et al., 2000).

1.3 THE BASAL GANGLIA AND REWARD-PREDICTION ERRORS

Dopamine projections into the striatum can provide a source of reward-prediction error signals that can be used to modify expectations about future rewards based upon unexpected outcomes, in order to impact future expectations and actions. As discussed previously, fMRI experiments in human subjects have suggested that the striatum is involved in processing information about the predictability of rewarding stimuli, though striatal neurons could be responding to the fact that the stimuli are unpredictable in character, time, or amount. Using a fMRI passive conditioning task, McClure, Berns, and Montague (2003) showed that positive and negative prediction errors in reward delivery time correlate with BOLD changes in the human striatum, with the strongest activation lateralized to the left putamen.

Activity in the striatum is enhanced when a primary rewarding stimulus is delivered in an unpredictable fashion compared to a predictable one (Berns et al., 2001). In a follow-up study, Pagnoni and colleagues used an fMRI operant conditioning task, in which small amounts of juice were administered to participants following a button press response, with a delay of either four seconds (Regular) or eight seconds (Delayed). The goal of their study was to determine whether the enhancement in striatal activity was due to errors in reward prediction. A significant effect of trial type (Regular vs Delayed) was observed in the BOLD response of the nucleus accumbens, with activity in the nucleus accumbens diverging at the timepoint corresponding to the missing delivery of the expected reward (Pagnoni et al., 2002).

5

1.4 LEARNING-RELATED CHANGES IN THE BASAL GANGLIA

The striatum has been previously implicated in habit learning, in which stimulus-response (S-R) associations are incrementally acquired (Packard & Knowlton, 2002). Acquisition and performance of instrumental actions are assumed to require both action-outcome and S-R habit processes. Over the course of extended training, control over instrumental performance shifts from goal-directed action-outcome associations to S-R associations that progressively dominate behavior. Lesions of the lateral part of the dorsal striatum disrupt this process, with rats having lesions to the lateral striatum showing selective sensitivity to devaluation of the instrumental outcome, indicating that this area is necessary for habit formation (Yin et al., 2004).

More recently, the role of the striatum has been extended beyond stimulus-response associations into the domain of more complex, cognitive tasks. A separate literature has implicated the basal ganglia in gradual, feedback-based learning of stimulus-response associations that involves integrating information across multiple experiences (Shohamy et al., 2008). Studies of patients with damage to the basal ganglia and patients with damage to the medial temporal lobe (MTL), a region known to be involved in declarative memory, have provided evidence that the basal ganglia are important for learning. Patients with basal ganglia damage have been found to have impaired performance on incremental, stimulus-response learning tasks, but intact performance on tasks involving declarative memory, suggesting a dissociation between the declarative and non-declarative memory systems (Knowlton et al., 1996).

A probabilistic classification task commonly used in the literature is the "weather prediction" task. In this category learning task, subjects view one or more cards with different geometric shapes on each trial, are asked to predict a category outcome (''rain'' or ''sunshine''), and receive feedback on their decision. The actual weather outcome is differentially associated with each card with a particular probability. (Gluck & Bower, 1988; Shohamy et al., 2008).

Knowlton and colleagues demonstrated that this sort of probabilistic classification depends on the basal ganglia, and not the MTL (Knowlton et al., 1996). Since the association between a card and a weather outcome is probabilistic, declarative memory of a single trial cannot support learning. Instead, learning in this task is supported by non-declarative, associative processes that integrate information across multiple trials. By implicating the basal ganglia in category learning in humans, this study has added to a growing literature suggesting that the same learning processes involving the basal ganglia may underlie both simple stimulus-response learning, as well as higher cognitive processes, such as categorization.

1.5 DIRECTION OF LEARNING SIGNALS TO TASK-RELEVANT AREAS

Although the bulk of human fMRI work looking at feedback processing has focused on the striatum, additional feedback-sensitive regions can be detected in a task-dependent way. However, it remains an open question whether these additional regions reflect an interaction with the striatum or some other feedback-based system.

An event-related fMRI study of native Japanese speakers investigated how performance feedback is processed when participants were asked to distinguish words beginning with /r/ and /l/ phonemes (Tricomi et al., 2006). Successful learning of the task and activation in the caudate nucleus was seen when participants received feedback, with the right caudate nucleus showing a differential response to positive and negative feedback. Greater activation was also seen in the superior temporal gyrus, an area involved in auditory processing, for incorrect compared to correct trials. In another study, involving a paired associate word-learning task, participants were asked to correctly choose which of two response words went with the target words.

Greater activation was observed in the left fusiform gyrus, the location of the visual word form area, on trials associated with subsequent correct performance than on trials associated with subsequent incorrect performance (Tricomi & Fiez, 2008).

These two experiments illustrate that learning signals generated by the basal ganglia may be directed toward the representations that are relevant for task performance, as additional regions in both studies showed differential activation for positive versus negative feedback. In the current study, we chose to examine these potential learning signals in the caudate nucleus in the context of numerical computation. This provides us with the opportunity to test the generality of the hypothesis that feedback sensitivity will be observed in both the caudate, and also in task-relevant areas.

1.6 MODULATION OF ACTIVATION IN THE CAUDATE NUCLEUS

If the caudate is important for learning, it suggests that if it can be maximally engaged, learning will likewise be greatest. Activation of the caudate nucleus seems to be modulated by multiple factors during feedback-based learning. Feedback is most informative when it differs from expectation, a situation that occurs when uncertainty in task performance is high. Once a task is well learned, feedback becomes completely expected, and therefore ceases to provide new information. Reward-related activity is greatest in both the striatum and midbrain regions when the reward is unpredictable. Activation of the caudate nucleus has been shown to vary as learning occurs, with reward becoming more predictable as performance on a task, and thus certainty, increases. As contingencies between stimuli and correct responses become well learned, caudate activity decreases (Delgado et al., 2005). The striatum appears to be sensitive to the informational value of feedback, with uncertainty in task performance resulting in more meaningful feedback and driving caudate activation.

An additional factor that appears to modulate activity in the caudate nucleus is the value of feedback to the individual. Since the feedback display itself is not intrinsically rewarding, it is the meaning in relation to the task that gives the feedback value and motivates future actions. This incentive value of an outcome can drive activity in reward-related brain regions, including the striatum, which varies with the motivational state of an individual. These regions respond to the value of an outcome relative to alternatives, rather than to its absolute value (Nieuwenhuis et al., 2005). For example, when the alternative is a monetary gain, winning no money results in a punishment response in reward-sensitive regions, including the caudate nucleus, while the same outcome produces a reward response when the alternative is a monetary loss.

A third factor that has been found to modulate caudate activation is not only the presence of feedback during a task, but the perceived relationship between one's response and the feedback received. For example, activation in the caudate nucleus during the "weather prediction" task, a probabilistic category learning task, was found to occur only when the task was feedback-based (Poldrack et al., 2001). Similarly, a phoneme learning task resulted in caudate activation and successful performance only when feedback was present (Tricomi et al., 2006). The caudate is not strongly activated simply in response to randomly delivered or cued monetary rewards or punishments. Instead, it shows strong activation only when there is a perceived contingency between a response and the outcome that follows the response (Tricomi et al., 2004). Therefore, feedback might be expected to more strongly activate the caudate when participants feel that the outcome they receive is determined by their response.

1.7 THE HORIZONTAL REGION OF THE INTRAPARIETAL SULCUS AND BASIC NUMBER SENSE

One general hypothesis is that the learning signals that are processed within the caudate are also reflected in task-relevant cortical areas, either through direct or indirect connections involving basal ganglia regions. Alternatively, it is possible that feedback can influence multiple learning systems. Previous work did not provide clear *a priori* regions in which task-relevant sensitivity might be expected. However, the current study uses numerical computation as the context in which to examine learning and its effects on activation in the caudate nucleus and task-relevant areas.

Our experimental design is based upon the idea that mathematical proficiency involves various types of knowledge that are in turn supported by different brain regions. Analogic representations, located in the horizontal intraparietal sulcus (hIPS) brain region in parietal cortex (Dehaene et al., 2003), provide a basic sense of number in both humans and non-human primates. This sense of quantity becomes more broadly tuned as the quantity magnitude increases (Nieder & Miller, 2004). Humans can also engage in symbolic representations, using Arabic numerals and number names to represent abstract information. These representations of number and rule-based procedures, which are supported by the inferior frontal gyrus and other regions (Goldberg et al., 2007), provide a precise and flexible way to manipulate numerical knowledge. Finally, learned mathematical facts, which may be supported by an angular gyrus region involved in verbal memory, may function to reduce the effort involved in symbolic computation (Dehaene et al., 2004; Delazer et al., 2005).

We hypothesize that the learning signals from the basal ganglia can serve to modify the strength and precision of these quantity representations in parietal cortex. Through error-

prediction signals generated by feedback, the representations and patterns of functional connectivity that are associated with rewarding outcomes could be strengthened, so that a stimulus can more automatically induce the response associated with positive outcomes.

1.7.1 Developmental changes

Neuroimaging studies of numerical development support the idea that changes in the hIPS regions can underlie changes in arithmetic ability. During a task in which participants, aged 8 to 19 years old, viewed arithmetic equations and were asked to judge whether they were correct or incorrect, older subjects showed greater activation in the left parietal cortex. Younger subjects showed greater activation in the prefrontal cortical areas, the hippocampus, and dorsal basal ganglia, indicating that they require comparatively more working memory and attentional resources, placing greater demands on the declarative and procedural memory systems in order to achieve similar levels of mental arithmetic performance (Rivera et al., 2005).

Ansari and Dhital investigated developmental changes in the functional neuroanatomy underlying magnitude processing by asking adults and children to perform numerical magnitude comparisons using arrays of squares. Although behavioral performance was similar across groups, adult participants exhibited greater effects of numerical distance on the left intraparietal sulcus compared to children (Ansari & Dhital, 2006). The findings of these studies suggest that the left intraparietal cortex exhibits an increase in activation and in the precision of its number code as arithmetic abilities emerge from childhood to adulthood.

1.7.2 Learning of arithmetic

Previous fMRI studies have investigated the learning of arithmetic by comparing trained problems to untrained problems. As problems become well-learned, a shift in activation is observed from frontal and parietal areas to the angular gyrus, indicating a switch from direct calculation to retrieval of a learned fact from memory. In a study investigating the effect of training on brain activation patterns, participants were trained for five sessions to answer specific multiplication and subtraction problems (Ischebeck et al., 2006). Trained problems yielded significantly higher activation in the left angular gyrus, while untrained problems showed stronger activation of inferior frontal and parietal regions, areas involved in working memory and quantity-based processing.

A similar study investigating the learning process (Ischebeck et al., 2007), showed a similar pattern of results, even when no training was provided prior to the scanning session. In this case, training consisted of a high frequency of repetition for one set of problems and a lower frequency for a second set. Following training, a shift of activation from the intraparietal sulci to the angular gyrus was observed, again indicating a shift from calculation to result retrieval from long-term memory.

1.8 FEATURES OF THE CURRENT TRAINING PROGRAM

The past studies of arithmetical learning do provide important evidence that activation changes within the hIPS regions can occur as a consequence of training, but they were not designed with a particular neural learning system in mind. In contrast, the central aim of the current work is to test whether the existing literature focusing on the involvement of the caudate in feedback processing can be used to create optimal learning via engagement of the caudate. Based upon

12

the results of previous work, we chose to design our training regime in a way that would maximally engage the striatum (in particular, the caudate nucleus), with the idea that learning can be optimized when the caudate is most active and can use reward-prediction error signals to modify expectations about future events.

This resulted in a training program involving uncertainty about response outcomes, high incentives for correct responses, and contingent feedback. Feedback was presented to participants following their response on each trial, and uncertainty was achieved by imposing short response windows. Participants were motivated to succeed by a bonus given for each correct response made. In applying this training procedure to particular problem types, we also took a very different approach than used in past studies. Since the goal was to promote broad adaptive change within hIPS, and not simply to support the memorization of particular problems, training focused on randomly-selected multi-digit addition and subtraction problems. If successful, we anticipated that participants would exhibit broad gains in computation speed and accuracy, with continued engagement of the hIPS region.

2.0 METHODS

2.1 PARTICIPANTS

Forty healthy, right-handed college students or recent graduates were recruited through newspaper advertisements and posted flyers and were paid an average of $317 for their participation in the experiment. All participants passed a screening process to ensure that they were native English speakers, not experts in the field of mathematics, did not have any problems performing basic math, and had a math SAT score between 600 and 700. Participants also completed an additional MR screening prior to enrollment in the study. This included excluding those who had metal or a history of metal in any part of their bodies, had a history of mental illness or illicit drug use, weighed over 300 pounds, or were uncomfortable in confined spaces.

Twenty participants (10 males, 10 females) were assigned to the experimental group and another twenty participants, each matched in gender and math SAT score to an experimental participant, were assigned to the control group. The math SAT score of a control participant was considered to be a match if it was within 20 points in either direction of the experimental participant's score.

One participant from the control group was removed from the analysis due to a lack of compliance during the fMRI session, leaving a total of 19 participants (9 males and 10 females; age range = 18-23, mean age = 20.8, average math SAT score = 645 ± 31). One participant from the experimental group was removed due to a cyst located in the cerebellar region, leaving

a total of 19 participants (9 males and 10 females; age range = 19-24, mean age = 20.9, average math SAT score = 643 $\pm$ 33). All participants gave written informed consent according to the Institutional Review Board at the University of Pittsburgh.

2.2 PROCEDURE

Each participant attended 9 sessions over a 10 day time period. During the first and ninth sessions, participants completed a 1.5 hour behavioral testing sessions (one before training and one following training), during which they completed multiple computerized tasks. Stimulus presentation and recording of behavioral data were controlled by E-Prime software (Psychology Software Tools, Inc.). The goal of these sessions was to assess potential benefits that arose from training. To minimize repetition effects, the stimuli for each measure were randomly selected, or in the case of a complex math assessment, drawn from one of two sets, the order of which was counterbalanced across participants. During the second and eighth sessions, participants completed a one hour-long pre- and post-training fMRI session, consisting of three tasks (one of which will be discussed in the current thesis). During the third through seventh sessions, participants completed two versions of a training task for 30 minutes each, resulting in an hour of training each day (Figure 1). Results from these pre/post behavioral tasks are not included in the current paper and will be reported in-depth elsewhere; instead, the behavioral data from the pre/post imaging sessions will be used as the primary measures to assess changes in performance.

Session	1	2	3	4	5	6	7	8	9
Description	Pre-test	Pre-scan	Train1	Train2	Train3	Train4	Train5	Post-scan	Post-test
Length	1.5 hours	1 hour	1 hour	1 hour	1 hour	1 hour	1 hour	1 hour	1.5 hours

Figure 1. Outline and timing of each session of the experiment.

2.2.1 Arithmetic training task

For each day of the five-day training period, participants in the experimental group spent approximately 30 minutes solving addition problems and 30 minutes solving subtraction problems. The specific features of our training program were designed with principles of basal ganglia function in mind. The training program incorporated contingent feedback on each trial, with the outcome of a trial dependent upon the response of the participant. On each trial, an addition or subtraction problem was briefly displayed on the screen, and participants had a brief amount of time to type in the correct answer after the problem disappeared from the screen (2.2 – 3.2 seconds, depending upon the type of problem). Following this response period, three green ✔'s or three red ✘'s were presented, indicating whether the response was correct or incorrect. If the participant did not enter an answer during the response window, three white hyphens were presented. After a delay of 500 ms, the next trial commenced, with a short rest period between blocks of 50 trials (Figure 2).

Arithmetic training task
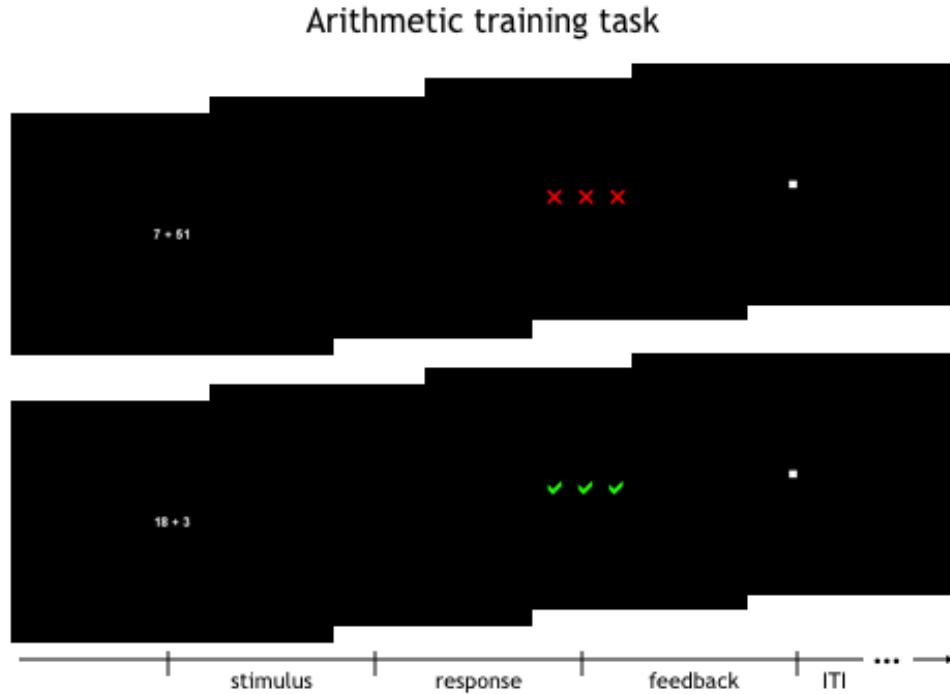
Figure 2. Trial outline for the arithmetic training task.

To maximize uncertainty about the outcomes, a short response window was imposed. For each trial, participants had a response deadline of 2.9 – 4.3 seconds from stimulus onset, depending on problem difficulty. High uncertainty was maintained by adjusting problem difficulty when participants attained 90% accuracy on a block of 50 trials. Monetary rewards were used to provide high incentive value for correct responses. Participants earned $0.02-$0.06 for each correct response, with larger rewards earned for more difficult problems.

On each trial, two numbers were randomly selected to form the presented problem, with some exclusionary constraints. Single-digit operands were randomly chosen from the numbers 2 through 9, while double-digit operands were randomly chosen from the numbers 11 through 99, excluding rounded decades (such as 20, 30, 40, etc.) For subtraction, the ones digits of the two operands were not allowed to be identical, in order to prevent answers that ended in 0. In problems involving two double-digit operands, identical operands were not allowed.

There were three levels of difficulty for training, each created through successive changes in the operand sizes. For Level 1, one operand was a single-digit number and the other a double-digit number, for Level 2, both were double-digit numbers, and for Level 3, one was a double-digit number and the other was a triple-digit number. With each increase in difficulty, the response deadline was increased slightly, in part to provide additional time for perception of the stimuli and entry of the response (Figure 3).

	Digits	Stimulus	Response (blank screen)	feedback (✔ or ✖)	ITI	Total time	Earnings
Level 1	double/single	700ms	2200ms	500ms	500ms	3900ms	$0.02
Level 2	double/double	900ms	2600ms	500ms	500ms	4500ms	$0.04
Level 3	triple/double	1100ms	3200ms	500ms	500ms	5300ms	$0.06

Figure 3. Stimulus timing for each level of arithmetic training.

2.2.2 Control training task

Rather than performing the arithmetic training task, control participants spent the five-day training period being exposed to and typing numbers. The control task had levels similar to those in the arithmetic training task. On each trial, two numbers containing the correct number of digits for the corresponding level were randomly selected, given the same exclusionary constraints described above. These two numbers were presented on either side of the screen with a colored shape appearing in the center. The size and type of shape were chosen randomly, and indicated to the participant which number they were required to enter in order to respond correctly on a given trial. An adaptive algorithm was used to adjust the length of the response window, so that the average accuracy on each block was yoked to that achieved by the experimental participant to whom the control participant was matched. Like the experimental

participants, participants in the control group received feedback immediately following each trial and monetary earnings associated with correct responses during the response window. Thus, each pair of matched experimental and control participants received approximately the same distribution of feedback and accumulated approximately the same pattern of earnings across the five days of training.

2.2.3 Functional magnetic resonance imaging task

2.2.3.1 Task description

An addition task, similar to what was seen during the training period, was used to understand which neural systems are engaged by the training tasks, and to investigate how and why the brain may change following training. The task involved the presentation of single-digit/double-digit addition problems (700 ms), and participants entered the solution to each problem using an MR-compatible numeric keypad. Participants were given an additional 2200 ms following the disappearance of the problem in order to enter their responses. Contingent feedback was presented for 500 ms after the response window and was followed by a fixation square, denoting the end of the trial (Figure 4). Three green ✔'s were presented following a correct response, three red ✗'s were presented following an incorrect response, and if the participant made no response, three white hyphens were shown and the trial was excluded from analysis. There was a break of twelve seconds between trials, resulting in a slow event-related task design that allowed time for the hemodynamic response being measured to return to baseline between trials. Participants solved a total of 40 addition problems, divided in 2 runs of 20 problems each.
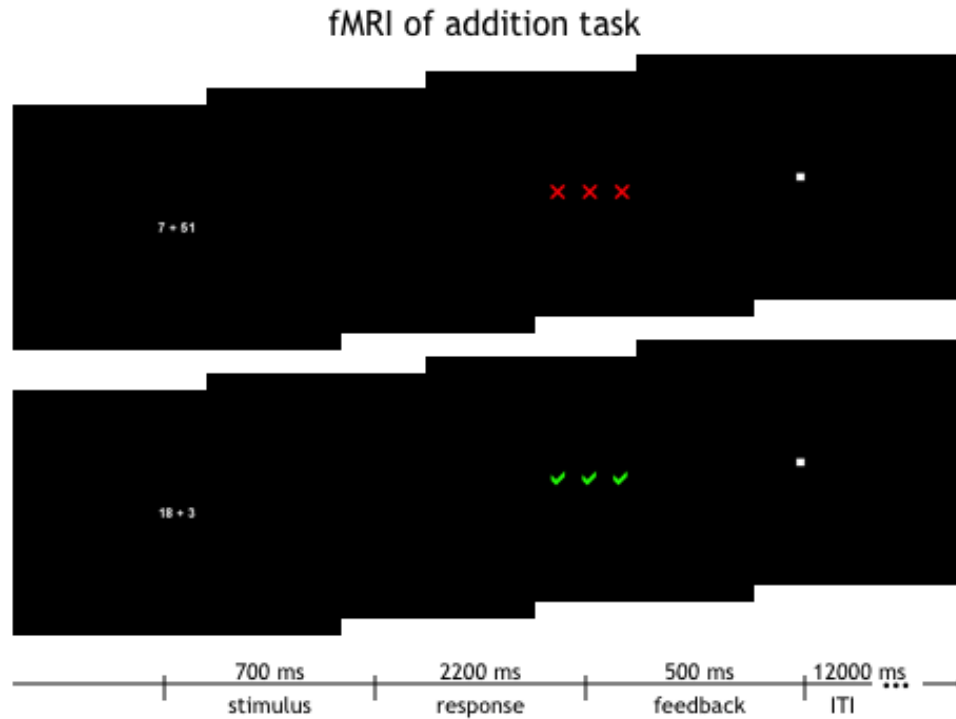
Figure 4. Trial outline and timing for the fMRI addition task.

2.2.3.2 Technical details

All structural and functional scans were performed on a 3.0 Tesla Siemens Magnetom Allegra head-only research scanner at the Brain Imaging Research Center (BIRC), a joint University of Pittsburgh and Carnegie Mellon University facility. An FDA approved research head coil was used for all of the scans and a certified MR technician was on hand at all times to monitor the participants while they were in the scanner.

T2-weighted in-plane and high resolution structural scans were obtained using a standard EPI pulse sequence. Thirty-eight 3.2 mm thick oblique slices, aligned to the AC/PC line, were obtained during the functional scans. Acquisition parameters used in this study were

designed to maximize the BOLD response: TR=2000 ms, TE= 25 ms, flip angle= 70°, with a field of view of 205 mm. These parameters gave us coverage of temporal lobe structures, as well as nearly full coverage of the parietal lobe.

We performed the analysis of the neuroimaging data using the NeuroImaging Software package (NIS 3.6) developed at the University of Pittsburgh and Princeton University. After the raw data were reconstructed and checked for quality, a 6-parameter rigid-body automated image registration procedure was used to correct for motion (AIR 3.08; Woods et al., 1993). No runs exceeded 4 mm of subject movement, so all runs were included in the analysis. The images were then corrected for linear trends. A reference brain was chosen from among the participants, and after stripping the skull off each participant's brain, the structural scans were transformed into this reference space. The functional data were scaled to a global mean and a three-dimensional Gaussian filter (4 mm FWHM) was used to smooth the data in order to account for between subject anatomical differences. All statistical analyses were conducted within the space of the reference brain, followed by a conversion of the reference brain into Talairach space (Talairach & Tournoux, 1988) so that statistical results could be reported within the context of a standard atlas space.

3.0 RESULTS

3.1 BEHAVIORAL RESULTS

3.1.1 Accuracy

In order to investigate performance on the addition task during the scanning sessions, the accuracy for each participant was calculated by dividing the number of correctly answered problems by 40 (the total number of problems presented during a session). A 2-way repeated measures ANOVA was then conducted to contrast the effects of session (pre- or post-training) and group (experimental or control). There was a significant main effect of session ($F(1, 37) = 30.891$, MSE = 0.022, $p < 0.001$, $\eta_p^2 = 0.455$), with the average accuracy during the post-training session (mean accuracy = 0.782) being significantly higher than during the pre-training session (mean accuracy = 0.649) (Figure 5).
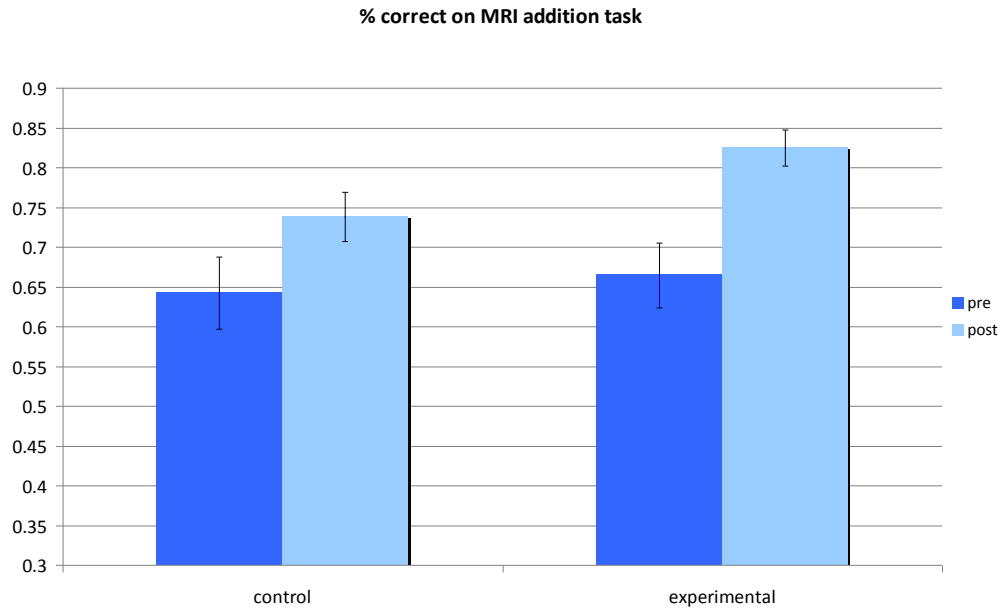
% correct on MRI addition task

Figure 5. Average accuracy (illustrated as % correct) on the addition task shown for each group and session.

Performing a simple-effect contrast of session within each group revealed that both groups improved significantly following training. The experimental group improved from a mean accuracy rate of 0.655 to 0.825 ($F(1, 37) = 25.889$, $p < 0.001$, $\eta_p^2 = 0.412$), while the control group showed a slightly smaller, though still significant, improvement, from 0.642 to 0.738 ($F(1, 37) = 7.851$, $p = 0.008$, $\eta_p^2 = 0.175$). There was no significant main effect of group and no significant difference between the accuracy of each group before training, indicating that there was not a pre-training difference in accuracy between the two groups. Although both groups demonstrated learning, there was a significant difference in post-training accuracy between groups. The experimental group (mean accuracy = 0.825) performed with a significantly higher accuracy during the post-training session than the control group (mean accuracy = 0.738) ($F(1, 37) = 5.149$, MSE = 0.014, $p = 0.025$, $\eta_p^2 = 0.128$).

3.1.2 Reaction time

A 2-way repeated measures ANOVA was also conducted on the mean reaction time (RT) data for each subject in order to contrast the effects of session (pre- or post-training) and group (experimental or control). For the current analysis, RT is defined as the time in milliseconds at which the last digit of the response was entered during the response window on correctly answered trials. There was a significant main effect of session ($F(1, 37) = 28.815$, MSE = 15578.294, $p < 0.001$, $\eta_p^2 = 0.438$), with RTs during the post-training session (mean RT = 1998.494 ms) being significantly faster than during the pre-training session (mean RT = 2105.814 ms). The interaction between session and group was also significant ($F(1, 37) = 9.008$, MSE = 15578.294, $p = 0.005$, $\eta_p^2 = 0.196$) (Figure 6).

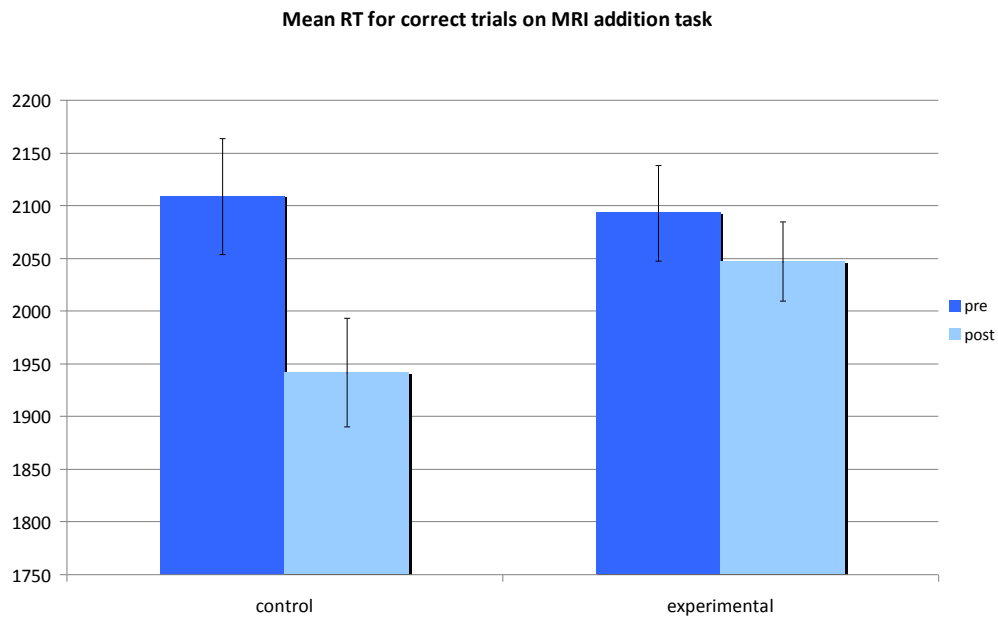**Mean RT for correct trials on MRI addition task**

Figure 6. Mean RT (shown in milliseconds) for correct trials on the addition task shown for each group and session.

In order to investigate this interaction in more detail, a simple-effect contrast of session within each group was performed. This contrast revealed that the control group, but not the experimental group, became significantly faster between the pre- (mean RT = 2108.687 ms) and post-training (mean RT = 1941.364 ms) sessions (control: $F(1, 37) = 34.146$, $p < 0.001$, $\eta_p^2 = 0.480$), presumably as a result of the control training task, which focused on rapid digit entry. There was no significant main effect of group ($F(1, 37) < 1$) and no significant difference between the RT of each group before training; however, the difference in RT between groups following training approached significance (experimental group: mean RT = 2055.625 ms, control group: mean RT = 1941.364 ms, $F(1, 37) = 3.312$, MSE = 38408.084, $p = 0.077$, $\eta_p^2 = 0.082$).

3.1.3 Accuracy on completed trials

The primary analysis of accuracy described above indicates that both groups learned following training. Both the experimental and control groups improved in accuracy from before to after training, raising the question of what each group learned, as the training programs completed by each group were focused on different aspects of the task. Although both groups demonstrated learning, an increase in performance could be accomplished in two different ways. First, an increase in the speed of digit entry could allow for the entry of more digits during the response window, which may be a likely explanation for the increase in performance observed for control participants. Second, an increase in computational ability can decrease the time needed to calculate the solution and increase the likelihood of arriving at a correct solution, which seems to be more likely for the experimental participants. Both of these changes could lead to a greater likelihood of entering a complete and correct solution for a given problem, and could account for the improvements in accuracy seen in the current study.

In order to examine these two forms of learning in more detail, and minimize the effect of response speed, a *post-hoc* secondary analysis of accuracy was conducted using only the trials in which a complete response was made. A complete response is defined here as a response containing the same number of digits as the solution of the presented problem. A new accuracy measurement was calculated for each participant, in which the number of accurate responses was divided by the number of complete responses that were produced. A 2-way repeated measures ANOVA was then conducted to determine the effects of session (pre- or post-training) and group (experimental or control). There was a significant main effect of session ($F(1, 37) = 7.041$, MSE = 0.008, $p = 0.012$, $\eta_p^2 = 0.160$), with accuracy during the post-training session (mean accuracy = 0.877) being significantly higher than during the pre-training session (mean accuracy = 0.824) (Figure 7).

% correct on MRI addition given complete response
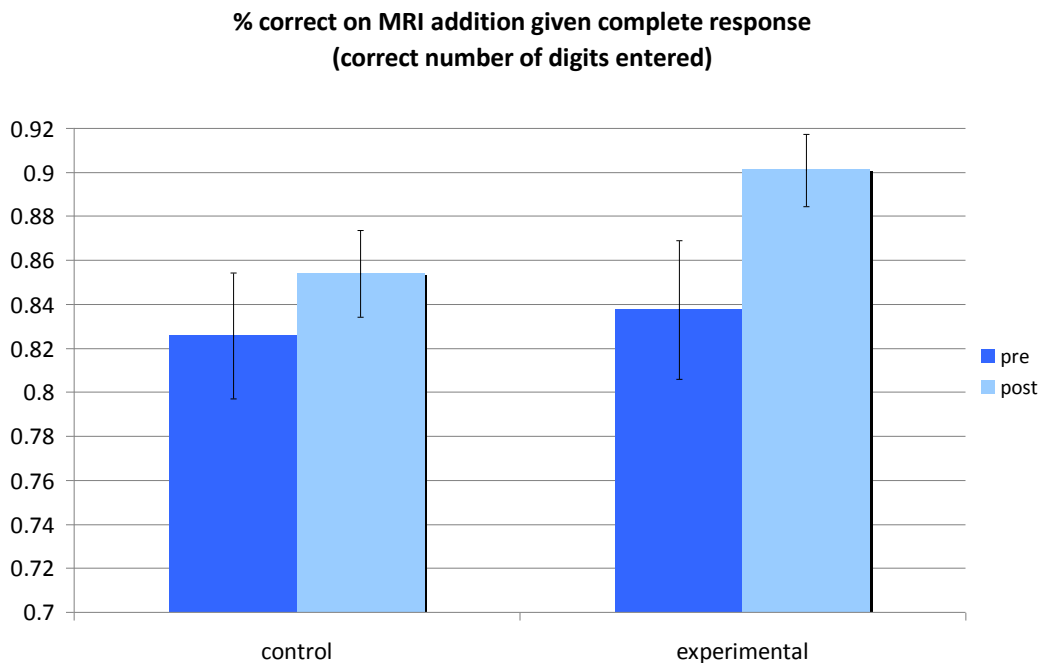**(correct number of digits entered)**

Figure 7. Average accuracy (illustrated as % correct) on addition task given a complete response, in which the correct number of digits were entered, shown for each group and session.

Performing a simple-effect contrast of session within each group revealed that only the experimental group improved significantly from before to after training. The experimental group improved significantly from a mean accuracy rate of 0.831 to 0.903 ($F(1, 37) = 6.618$, $p = 0.014$, $\eta_p^2 = 0.152$). The control group showed a slightly smaller improvement, from 0.816 to 0.850, though this difference was not significant ($F(1, 37) = 1.435$, $p = 0.239$, $\eta_p^2 = 0.037$). There was no significant main effect of group and no significant difference between the accuracy of each group before training. There was, however, a significant difference in accuracy between groups following training: the experimental group performed with a significantly higher accuracy during the post-training session than the control group ($F(1, 37) = 4.254$, $p = 0.046$, $\eta_p^2 = 0.103$).

3.2 NEUROIMAGING RESULTS

We had two main predictions for the fMRI data. First, we believed that the addition task performed during the scanning session would engage both the caudate nucleus and the hIPS region, with both regions responding differentially to positive versus negative feedback. Second, we believed that both regions would show an effect of training, with the experimental group showing a decrease in caudate activity following training, and the control group showing a smaller decrease in activity or no decrease at all. The current state of the literature precluded strong predictions about the effect of training on activity in the hIPS, though we speculated that the representation of quantity would become more refined (more precisely tuned) following training.

In order to test the predictions of this study, we began by first localizing brain regions that responded differentially to positive versus negative feedback. This was done using a voxel-wise ANOVA with subject as a random factor and TR (TR1-TR8) and feedback (correct vs.

incorrect) as within-subject factors was conducted. The resulting images were thresholded using an alpha criterion of $p < 0.000000000005$ and a contiguity requirement of 5 voxels. These values were chosen *post-hoc* because they served to identify focal and distinct activation clusters within the *a priori* regions of interest (bilateral caudate and hIPS). In order to reach more reasonable cluster sizes and better focus on the centroid of activation for each region, the stringency of the alpha threshold for the left and right hIPS, as well as the left caudate, was further increased (to $p < 0.0000000000000005$). The location and extent of each of the four regions of interest used in the current study are shown below (Figures 8-9).
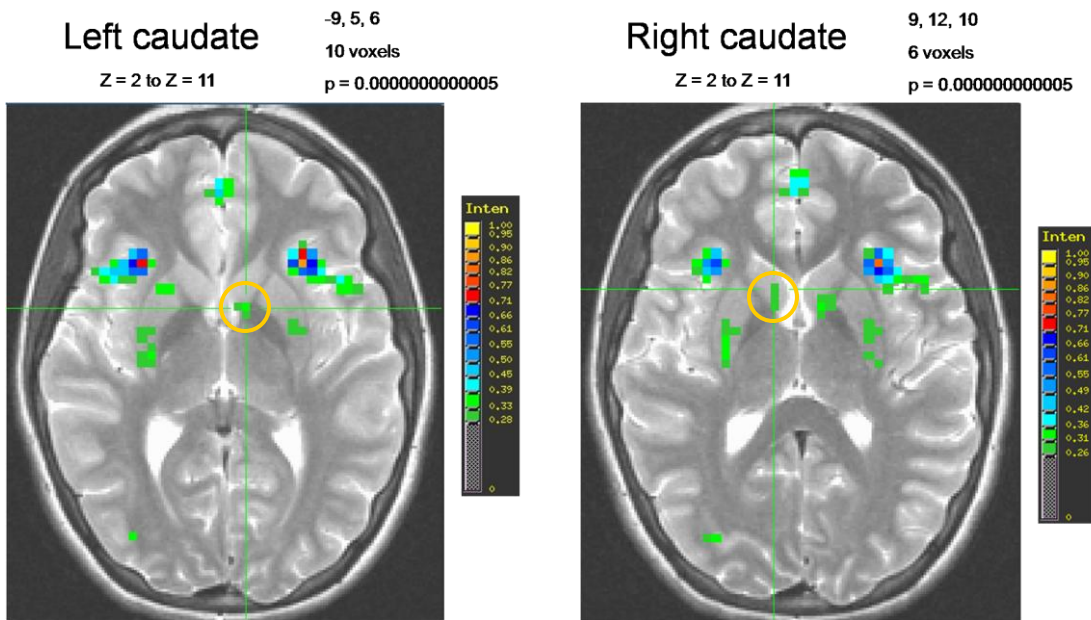
Figure 8. Axial view of activation in the left caudate and right caudate ROIs (circled in yellow).
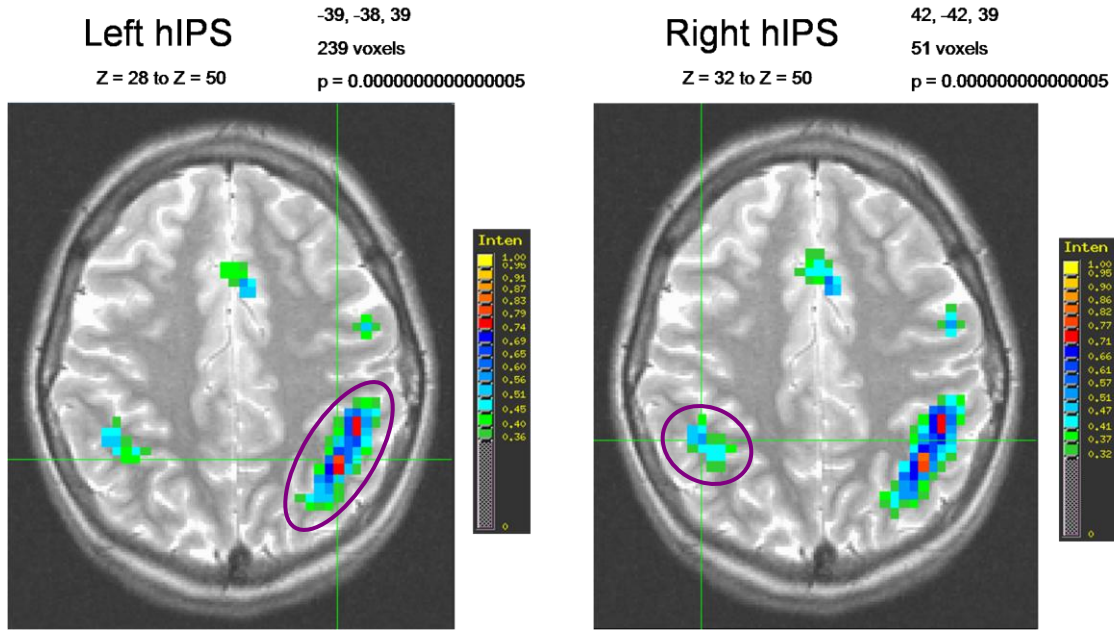
Figure 9. Axial view of activation in the left hIPS and right hIPS ROIs (circled in purple).

In addition to the four *a priori* regions of interest, 23 other clusters of activation surpassed the statistical and contiguity criteria (see Table 1). The activation patterns within these additional regions will not be discussed further, though it should be noted that all regions exhibit a highly significant TR x feedback interaction (the corrected *p*-value for all regions was well less than $p < 0.001$ given a contiguity threshold of 5, based on AFNI AlphaSim) and thus future research on the role that these regions may play in learning the training task is warranted.

Table 1. Regions emerging from a Feedback x TR ANOVA at a threshold of p=0.000000000005 with a cluster contiguity of 5 voxels

ROI	AREA	TALAIRACH COORDINATES (x, y, z)	SIZE (in voxels)
ROI-001	L medial frontal gyrus	-5, 3, 52	197
ROI-002	* L inferior parietal lobule	-39, -48, 39	538 (239)
ROI-003	R middle frontal gyrus	27, -7, 56	21
ROI-004	* R inferior parietal lobule	42, -42, 39	110 (51)
ROI-005	R postcentral gyrus	41, -19, 43	13
ROI-006	R paracentral lobule	1, -38, 43	6
ROI-007	L inferior frontal gyrus	-46, 3, 29	182
ROI-008	R inferior frontal gyrus	45, 3, 33	36
ROI-009	R middle frontal gyrus	45, 25, 33	16
ROI-010	R cuneus	12, -86, 29	42
ROI-011	L cingulate gyrus	-2, -28, 26	7
ROI-012	L cingulate gyrus	-12, -46, 29	5
ROI-013	R cingulate gyrus	5, -52, 26	5
ROI-014	L middle occipital gyrus	-22, -86, 20	22
ROI-015	R anterior cingulate	2, 48, 6	56
ROI-016	L insula	-29, 21, 3	102
ROI-017	R insula	31, 21, 6	73
ROI-018	L lentiform nucleus & L putamen	26, -3, -3	46
ROI-019	R lentiform nucleus & L putamen	23, 6, 0	61
ROI-020	R middle occipital gyrus	34, -79, 6	12
ROI-021	R caudate (body)	9, 12, 10	6
ROI-022	* L caudate (body)	-9, 5, 6	16 (10)
ROI-023	R parahippocampal gyrus	23, -15, -11	12
ROI-024	L parahippocampal gyrus	-22, -12, -11	8
ROI-025	R culmen (cerebellum)	20, -45, -21	15
ROI-026	R culmen (cerebellum)	27, -56, -28	24
ROI-027	within 1 mm of L pyramis (cerebellum)	-26, -59, -28	16

* The actual ROIs for these regions were taken at a more stringent threshold in order to better focus on the centroid of activation. The number of voxels in the ROIs at these higher thresholds are indicated in parentheses.

In the next step of the analyses, the localized activation clusters were used as regions of interest for secondary analyses that were designed to assess the impact of the other experimental factors (training group, imaging session) on the feedback-sensitive BOLD responses in the caudate and hIPS. For these secondary analyses, one participant from the control group and one participant from the experimental group were removed due to an incomplete experimental design (no incorrect responses following training). First, for each region, the mean signal intensity was computed for each individual for each type of feedback trial, separately for the pre vs. post imaging session and for each of the time point (TR1-TR8) within each trial. The average signal values for each time point of correct and incorrect trials were then corrected for baseline differences, by subtracting the activation in the given region during the first TR from each of the eight trial time points. Since graphical inspections of the results did not point towards significant laterality effects, and since there were no strong *a priori* reasons to expect laterality differences, the baselined data were then averaged across the left and right regions of each structure of interest. This final step was done with the intention of decreasing variance within the dataset, thereby increasing our statistical power to observe the predicted four-way interactions (i.e., the prediction that the shape of the BOLD responses in our region will vary as a function of the trial outcomes, with the size of the outcome differences specifically modulated in the experimental group as a result of training). These final values were used for all further analyses.

In order to determine if there were significant differences in the pattern of activation between any of the conditions, a 4-way repeated measures ANOVA with session (pre vs. post), feedback (correct vs. incorrect), and TR (TR1-TR8) as within-subject factors, and group (experimental vs. control) as a between-subject factor was performed for the bilateral hIPS (with activity averaged across the left and right hIPS ROIs) and bilateral caudate (also averaged across left and right). The results of these ANOVAs are described below for each of the regions.

31

3.2.1 Pattern of activation in hIPS

In bilateral hIPS, the 4-way repeated measures ANOVA with session, feedback, and TR as within-subject factors and group as a between-subjects factor, revealed a main effect of feedback ($F(1, 34) = 67.870$, MSE = 39.461, $p < 0.001$, $\eta_p^2 = 0.666$), with activity on incorrect trials (mean = 6.799) being significantly greater than on correct trials (mean = 3.749), as well as a main effect of TR ($F(7, 238) = 134.345$, MSE = 17.023, $p < 0.001$, $\eta_p^2 = 0.798$). The 4-way interaction (group x session x feedback x TR) was also significant ($F(7, 238) = 2.411$, MSE = 3.451, $p = 0.021$, $\eta_p^2 = 0.066$), indicating a different pattern of activation for each feedback type and session between groups, which seems to be driven by the greater activation on incorrect trials during the pre-training session compared to the post-training session in the control group, while the experimental group showed greater activation on incorrect trials for the post-training session compared to the pre-training session (Figure 10). A 3-way repeated measures ANOVA, conducted independently for the two groups, revealed that the experimental group showed a significant session x feedback x TR interaction ($F(7, 119) = 2.794$, MSE = 2.569, $p = 0.010$, $\eta_p^2 = 0.141$), while the control group did not. Similar results were observed within each hemisphere, with both hemispheres showing the same general pattern of activation, though the full 4-way interaction failed to reach significance when the analysis was restricted to a single hemisphere.
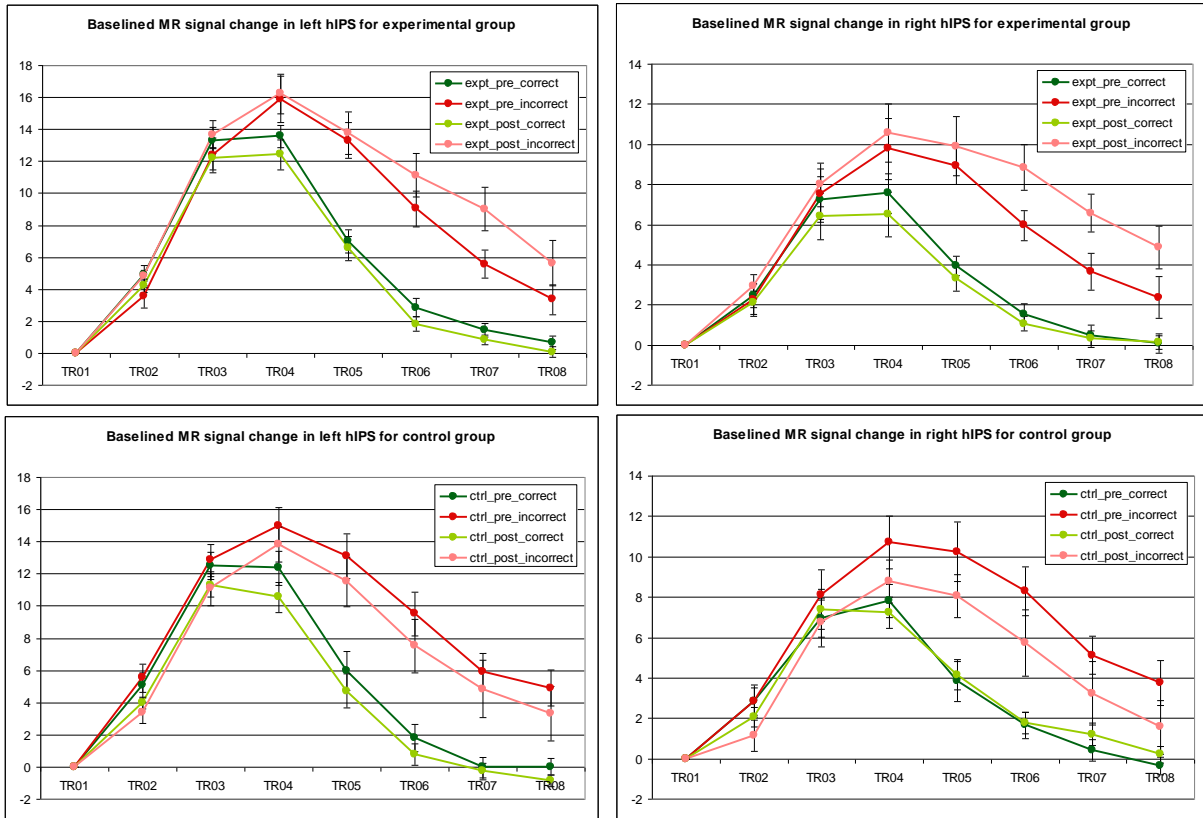
Figure 10. Activation in the left and right hIPS (illustrated as a change from baseline activation) during correct and incorrect trials shown for the experimental group (top) and the control group (bottom).

3.2.2 Pattern of activation in the caudate nucleus

Like the hIPS, the caudate did not show a main effect of group or session, though a main effect of TR ($F(7, 238) = 49.087$, MSE = 32.082, $p < 0.001$, $\eta_p^2 = 0.591$) was present. A significant feedback x TR interaction ($F(7, 238) = 13.715$, MSE = 7.200, $p < 0.001$, $\eta_p^2 = 0.287$) (Figure 11) was observed, but contrary to the initial hypotheses, the higher-order interactions involving feedback (e.g., group x session x feedback x TR, or session x feedback x TR) did not reach significance. However, a significant session x TR interaction was found ($F(7, 238) = 9.863$, MSE = 10.848, $p < 0.001$, $\eta_p^2 = 0.225$) (Figure 12). This interaction indicates a more muted BOLD response on each trial, during the post-training session for both the experimental and the control groups, and for both correct and incorrect trials.

Figure 11. Activation in the left and right caudate nucleus (illustrated as a change from baseline activation) during correct and incorrect trials shown for the experimental group (top) and the control group (bottom).
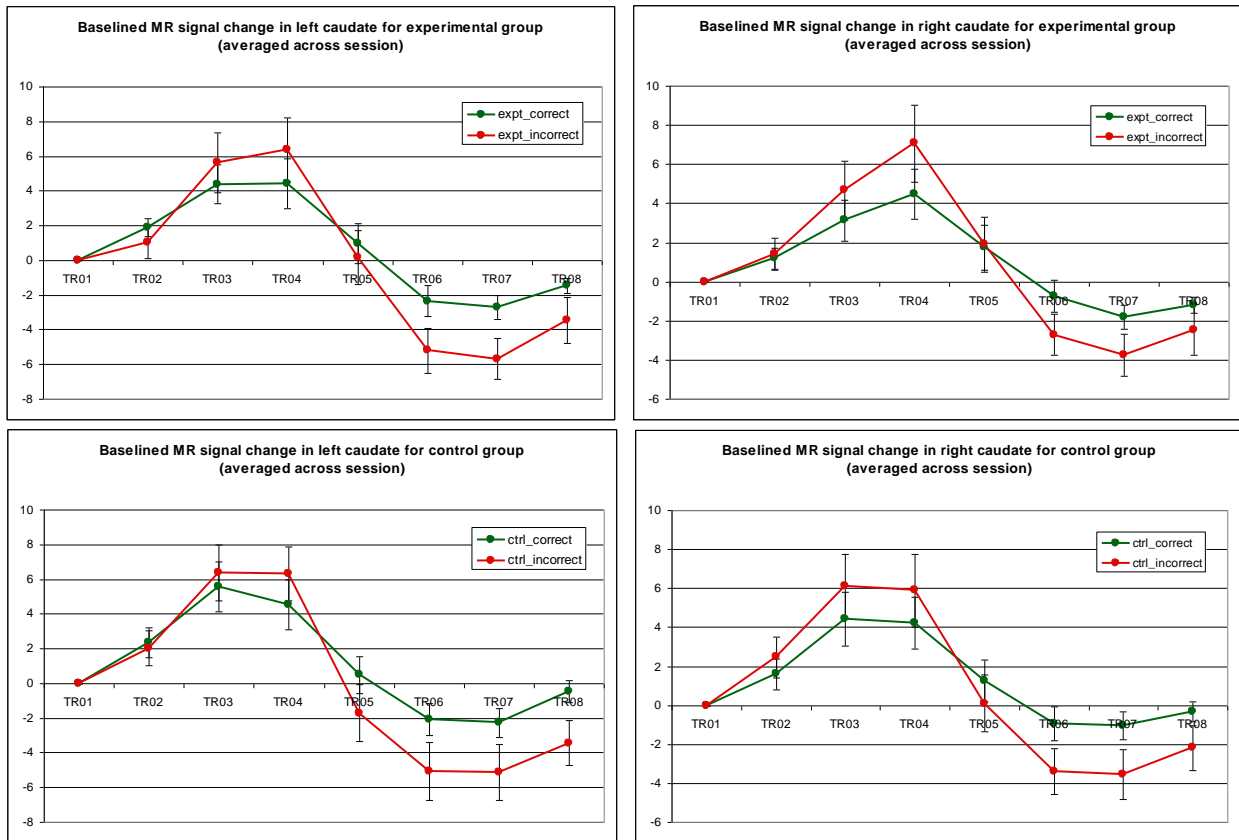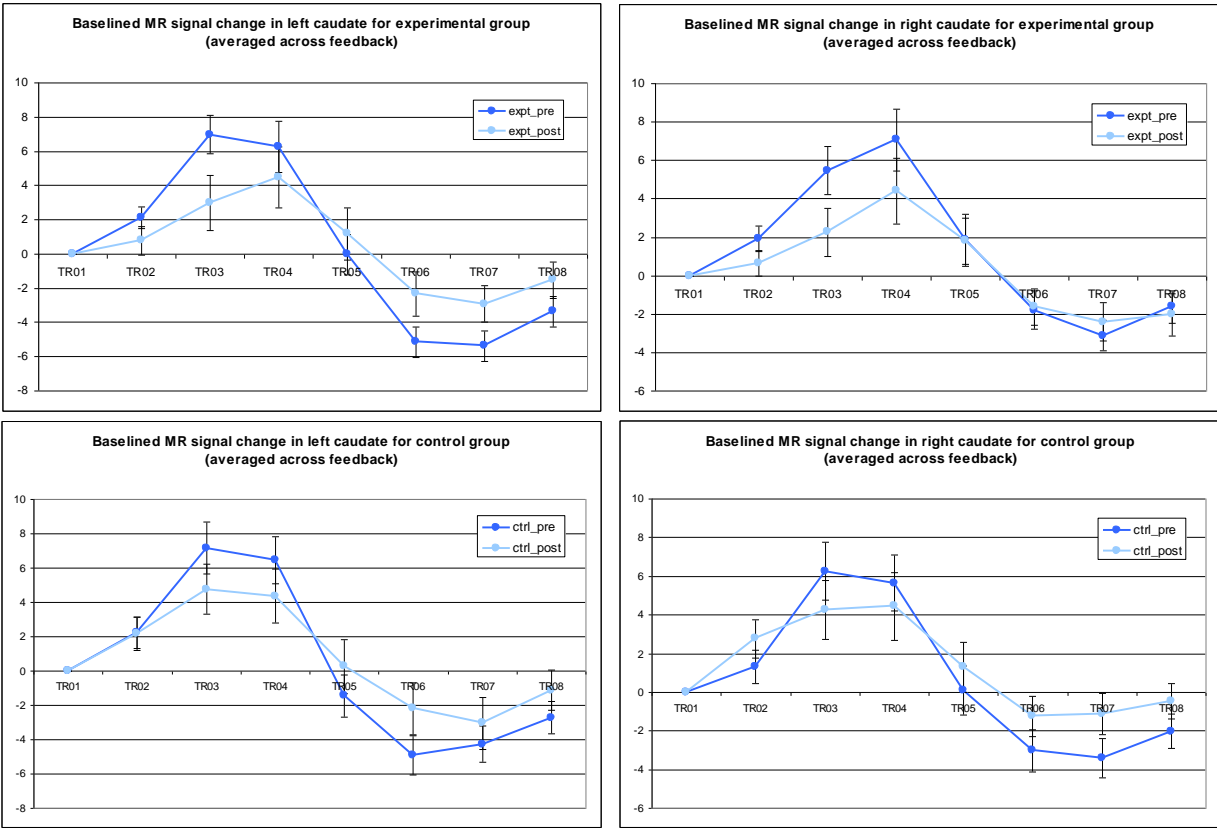
Figure 12. Activation in the left and right caudate nucleus (illustrated as a change from baseline activation) during the pre- and post-training session shown for the experimental group (top) and the control group (bottom).

A 3-way repeated measures ANOVA with session, feedback, and TR as within-subjects factors was conducted independently for each group. Both groups showed only a significant main effect of TR (experimental: $F(7, 119) = 30.320$, MSE = 24.148, $p < 0.001$, $\eta_p^2 = 0.641$; control: $F(7, 119) = 21.626$, MSE = 40.016, $p < 0.001$, $\eta_p^2 = 0.560$). A session x TR interaction and a feedback x TR interaction were also found to be highly significant for both groups, indicating that both groups showed a similar pattern of activation that differed across session and across feedback type.

4.0 DISCUSSION

We set out to demonstrate the effectiveness of a training regime inspired by principles of basal ganglia function and determine whether the training has a broad impact on mathematical proficiency. We accomplished our goal by creating an arithmetic training program that resulted both in the engagement of the basal ganglia and learning. Improvements were observed in accuracy on the addition task as well as other measures of numerical and mathematical abilities that are reported elsewhere (Kallai, et al., 2010; Ponting, et al., 2010).

4.1 IMPROVEMENT FOLLOWING TRAINING

The improvement in accuracy on the addition task that we observed for the experimental participants could have been due to mere exposure to symbolic representations of number, greater facility of digit entry using a keypad, and overall familiarity effects with our specific assessment measure. To rule these out as possible explanations, a second group of participants completed a digit-entry training program. Like the experimental participants, these control participants were exposed to numbers each day of training and had a great deal of experience typing digits using the number keypad. Both groups of participants performed identical tasks during the scanning sessions as well as the behavioral pre-/post-test sessions.

We expected that the experimental group would show greater improvement on the addition task performed during the scanning session since they had the advantage of 5 days of

training on addition and subtraction problems. Instead, results indicated that both groups showed a significant improvement in accuracy between sessions, though the experimental group did show a greater improvement in accuracy than the control group. Further investigating this improvement in accuracy, we examined the accuracy on trials in which participants entered a response consisting of the correct number of digits. Since the response window of the task was only 2200 ms, errors on the task were possible due to being slow to compute a solution and not being able to enter a full response in time, as well as computing the solution incorrectly and entering a complete, but incorrect, response within the response window.

We were surprised that the control group showed an improvement on the task following training, but a closer look at the accuracy on trials in which a complete response was made indicated that the experimental group showed a significant increase in accuracy, while the control group showed a slight, non-significant increase. This is also reflected in the RT data, which showed a significant decrease in RT for the control group compared to the experimental group. The increased accuracy for the control group during the post-training session seems to reflect the fact that they were able to enter their responses faster (their decreased RT), while the increase in accuracy of the experimental group reflects their increased computational ability. Both of these learning processes fit well with the emphasis of the training task completed by each group, with the experimental training program focusing on correct computation, and the control training program focusing on fast digit-entry.

4.2 FEEDBACK-SENSITIVE REGIONS

The first goal of the study was to determine the involvement of the caudate, as the task was designed with principles of striatal function in mind, having contingent feedback, uncertainty, and incentives for correct responses. We predicted that the caudate nucleus would be sensitive

to feedback, but this is not a trivial finding given that this study involved a higher-level cognitive task compared to previous studies investigating feedback processing. Previous work using a gambling task, phoneme learning, and paired associate learning, have also observed differential caudate activation for positive versus negative feedback, though in these studies, learning was achieved through multiple encounters with specific stimuli. In contrast, participants in the current study were not specifically learning each item, as there were thousands of possible operand combinations, yet the caudate still shows the prototypical response profile.

Activity in the caudate nucleus in the current study was sensitive to feedback, as indicated by the ability to identify significant clusters of activation in both the left and right caudate nuclei that exhibited a significant interaction in a voxel-wise ANOVA with TR (TR1-8) and feedback (positive, negative) as factors. Activation in both left and right caudate shows the prototypical response seen in previous work (Delgado et al., 2000; Tricomi et al., 2004; Tricomi et al., 2006), with an increase in activation on incorrect trials followed by a decrease below baseline, and a more sustained and lower magnitude activation on correct trials.

Since this study involved an arithmetic task, we believed that the hIPS region would be activated, given its well-established involvement in quantity representation. The observed response in hIPS in this study that differentiates between positive and negative feedback is an interesting finding, as feedback effects in hIPS have not been well examined. Like the caudate nucleus, activity in the hIPS region was sensitive to feedback (both left and right hIPS were two of the ROIs showing a feedback x TR interaction when a voxel-wise ANOVA with feedback and TR as factors was performed), though the overall pattern of the BOLD response was different across the two regions. In the caudate, incorrect trials are associated with an initially heightened positive response followed by a large dip below baseline. In the hIPS, incorrect trials appear to be best characterized as a delay in the return of the positive BOLD response to baseline, as compared to the time course for correct trials. The difference in activity between trials in which

positive feedback was given and trials in which negative feedback was given was not seen until TR4, which was 6 seconds after the onset of the trial and 3 seconds after the onset of feedback although feedback is displayed to the participant at TR2.

This feedback sensitivity in hIPS provides strong support for the hypothesis that task-relevant areas will show a feedback-sensitive response, as the hIPS plays a clear role in math and quantity representation. In previous work, feedback-sensitivity was seen during a phoneme learning study in the superior temporal gyrus, and during a paired associate learning task in the visual word form area. In both of these cases, it is not clear that the participants' relationship to the phonemes in an auditory manner or to the visual forms of the word pairs is the most relevant aspect of processing for task performance, as there is no well-established "phoneme region" in the brain and other representations of the word pairs may be more critical for task performance.

We did find other regions that were active in a TR x feedback voxel-wise ANOVA (see Table 1). These additional regions were not analyzed because they were not *a priori* regions of interest, though they are still interesting and warrant further study. Previous work in the lab using a paired associate task observed activation in the putamen and parahippocampal gyrus, both of which were seen in the original TR x feedback voxel-wise analysis. The inferior frontal gyrus, another region that showed a TR x feedback interaction in the current study, has been previously implicated in arithmetic tasks, with activation being related to phonological output (Dehaene et al., 2004).

4.3 TRAINING-EFFECTS ON CAUDATE ACTIVATION

As learning occurred in both groups, all participants should show a change somewhere within the brain that underlies this learning. We examined activity in the brain while participants performed the addition task in order to determine how the brain has changed following training,

with a focus on four areas of prior interest: the left and right caudate nucleus, and the left and right hIPS. Each of these regions was engaged by the addition task performed by participants within the scanner, and showed sensitivity to feedback, as evidence by the emergence of these regions in a TR x feedback voxel-wise ANOVA.

We expected a decrease in caudate activation following training, though we expected a larger decrease for the experimental group, who had experience solving addition problems and should therefore be more certain of their responses following training. The control group, on the other hand, had experience only in digit-entry, so the certainty of their computations should not have changed. Since we did see a decrease in caudate activation on all trials following training for both groups, it may be the case that the control group is also more certain of their answers following training, because they are able to enter their answers much faster and are more likely to respond fully during the response window, and not necessarily because their computations are more accurate. Although we did not expect the certainty of participants who completed the digit-entry training to change, both groups had reason to be more confident and certain of their responses, as both showed a significant increase in accuracy on the addition task.

4.4 PATTERN OF ACTIVATION IN HIPS

Activity in the hIPS region, the site of analogic number representation in the brain, showed larger changes from baseline than the caudate nucleus. Activity during the first 3 TRs of the trial were very similar for correct vs. incorrect trials, with activity diverging for the two different outcomes on TR4, about 3 seconds following the onset of the feedback display. Since the divergence of activity occurs after feedback is given, activity in the hIPS region does not seem to reflect neural differences during the computation period that predict subsequent performance on a given trial (e.g., greater activity in hIPS could indicate better mathematical processing,

42

which in turn could increase the likelihood of a correct response). Instead, the timing of the feedback effects suggest that the hIPS activity is reactive to the delivery of feedback. Although the observed difference in activity seems to be in response to the performance of the participants, the idea that the differential activation may be predictive of performance cannot be ruled out. The current study did not use a jittered design, so the activation corresponding to different trial components cannot be separated, but the fact that activation for correct and incorrect trials are so close early in the trial before feedback is displayed provides support for the reactive hypothesis.

Activation of hIPS shows a prototypical hemodynamic response for all trials, with larger and more sustained activation for incorrect compared to correct trials, possibly indicating that participants are thinking more about the quantities involved in the addition problem after entering an incorrect response, though the learning signals of the two groups may be directed to different regions. For example, participants in the experimental group may be treating the negative feedback as evidence that they need to speed up their motor response, while control participants may treat it as evidence for more accurate computations.

The session during which the task was performed (pre- or post-training) did not seem to largely affect activity during correct trials, but activity on incorrect trials showed an interesting pattern. For the experimental group, activity on incorrect trials following training was higher than before training, while the control group showed the opposite pattern, with activity on incorrect trials following training being lower than before training. This increase in activation following training for the experimental group may indicate that the experimental participants are more focused on quantity following an incorrect response after training, thus more deeply engaging the hIPS region in order to better prepare for the next trial.

There are several alternative interpretations for what the difference in the punishment response may indicate, with the common theme that experimental participants are using the negative feedback to enhance processing in the hIPS region following training. The negative

43

feedback may indicate to the participant a need to recalculate the solution, provide an inhibitory signal to neurons representing the incorrectly-calculated response, or result in an increase in activation as participants prepare to use hIPS to a greater degree on the following trials. Further analysis is needed to test these ideas, including probing activation on trials following an error and looking for links to behaviors and performance on other tasks examining mathematical ability.

Previous work has indicated that following training, a shift of activation is seen in arithmetic tasks from the intraparietal sulci to the angular gyrus (Ischebeck et al., 2006; Ischebeck et al., 2007). This shift in activation has been interpreted as representing a shift from calculation to retrieval from long-term memory. These studies involved training on a specific subset of individual problems, though in the current study, participants were not trained on specific problems, and were instead trained focusing on general arithmetic competence. For this reason, we did not expect to see the same shift from hIPS activation to activation of the angular gyrus, observed in the previous work though the increased hIPS activation during incorrect trials observed for experimental subjects following training may indicate an increased use of estimation and general calculation, rather than relaying on the retrieval of math-facts from long-term memory.

4.5 CONNECTIVITY BETWEEN THE CAUDATE AND HIPS

One of the original goals of the study was to determine if learning signals generated by the striatum during arithmetic training are able to modify quantity representations in the hIPS region. Though feedback sensitivity in hIPS was observed in this study, it is not clear whether these changes could reflect some type of common influence of a reinforcement learning system (e.g., whether both regions could reflect an influence of dopaminergic cell firing) or whether the

44

differential activation for positive and negative feedback is mediated by another region. This study does not provide direct evidence that the caudate and hIPS are interconnected. The role of caudate in the feedback sensitivity of hIPS remains an open issue, as the changes in activation of the caudate do not track directly with changes seen in hIPS. Further study is needed, using functional connectivity, fiber tracking, or other measures, to determine the relationship between activation in the caudate and activation in hIPS.

4.6 CONCLUSIONS

Overall, this study provides evidence that learning can be achieved via optimal engagement of the caudate nucleus. Both experimental and control participants, who completed training focusing on arithmetic calculation and digit-entry respectively, showed improvement on a task involving the addition of a double-digit and a single-digit number following training, as successful performance on the task required accurate computations and entry of the solution within a narrow response window. This task utilized contingent feedback, uncertainty regarding performance, and incentives for correct responses, and resulted in activation of both the caudate nucleus and the hIPS region. Activation of the caudate nucleus replicated previous work, as it showed the prototypical pattern of activity that distinguished between positive and negative feedback. Activation of the hIPS region was not surprising, given the focus on arithmetic calculation, but this region also exhibited feedback-sensitive activation that differed between sessions and groups, indicating a possible common influence of a reinforcement learning system on the hIPS and the caudate nucleus.

BIBLIOGRAPHY

Ansari, D. & Dhital, B. (2006). Age-related changes in the activation of the intraparietal sulcus during nonsymbolic magnitude processing: An event-related functional magnetic resonance imaging study. *Journal of Cognitive Neuroscience, 18*(11), 1820-1828.

Berns, G. S., McClure, S. M., Pagnoni, G., & Montague, P. R. (2001). Predictability modulates human brain response to reward. *Journal of Neuroscience, 21*(8), 2793-2798.

Dehaene, S., Molko, N., Cohen, L., & Wilson, A. J. (2004). Arithmetic and the brain. *Current Opinion in Neurobiology, 14*, 218-224.

Dehaene, S., Piazza, M., Pinel, P., & Cohen, L. (2003) Three parietal circuits for number processing. *Cognitive Neuropsychology, 20*(3-6), 487-506.

Delazer, M., Ischebeck, A., Domahs, F., Zamarian, L., Koppelstaetter, F., Siedentopf, C. M., Kaufmann, L., Benke, T., & Felber, S. (2005). Learning by strategies and learning by drill – Evidence from an fMRI study. *NeuroImage, 25*, 838-849.

Delgado, M. R., Miller, M. M., Inati, S., & Phelps, E. A. (2005). An fMRI study of reward- related probability learning. *NeuroImage, 24*(3), 862-873.

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. *Journal of Neurophysiology, 84*(6), 3072-3077.

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology, 10*, 732-739.

Elliott, R., Friston, K. J., & Dolan, R. J. (2000). Dissociable neural responses in human reward systems. *Journal of Neuroscience, 20*, 6159-6165.

Gluck, M.A. & Bower, G.H. (1988). From conditioning to category learning: an adaptive network model. *Journal of Experimental Psychology- General, 117*(3), 227–247.

Goldberg, R. F., Perfetti, C. A., Fiez, J. A., Schneider, W. (2007). Selective retrieval of abstract semantic knowledge in left prefrontal cortex. *Journal of Neuroscience, 27*(14), 3790-3798.

Hikosaka, O., Sakamoto, M., & Usui, S. (1989). Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *Journal of Neurophysiology, 61*, 814–832.

Hollerman, J. R., Tremblay, L., & Schultz, W. (1998). Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of Neurophysiology*, *80*, 947-963.

Ischebeck, A., Zamarian, L., Egger, K., Schocke, M., & Delazer, M. (2007). Imaging early practice effects in arithmetic. *NeuroImage*, *36*, 993-1003.

Ischebeck, A., Zamarian, L., Siedentopf, C., Koppelstätter, F., Benke, T., Felber, S., & Delazer, M. (2006). How specifically do we learn? Imaging the learning of multiplication and subtraction. *NeuroImage*, *30*, 1365-1375.

Kallai, A., Ponting, A., Schunn, C., & Fiez, J. (2010). An arithmetical training regime motivated by principles of basal ganglia function. Poster pr*esented at the 17th Annual Meeting of the Cognitive Neuroscience Society (CNS)*, Montréal, Canada.

Kawagoe, R., Takikawa, Y., & Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience, 1*(5), 411-416.

Knowlton, B. J., Mangles, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science, 273*, 1399-1402.

Knutson, B., Westdorp, A., Kaiser, E., & Hommer, D. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *NeuroImage, 12*, 20-27.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron, 38*, 339-346.

Mirenowicz, J., & Schultz, W. (1994). Importance of unpredictability for reward responses in primate dopamine neurons. *Journal of Neurophysiol*ogy, *72*,1024-27.

Nieder, A., & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *PNAS*, *101*(19), 7457-7462.

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*(4), 1302-1309.

Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annual Review of Neuroscience*, *25*, 563-593.

Pagnoni, G., Zink, C. F., Montague, P. R., & Berns, G. S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nature Neuroscience*, *5*, 97-98.

Poldrack, R. A., Clark, J., Pare-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., et al. (2001). Interactive memory systems in the human brain. *Nature, 414*(6863), 546-550.

Ponting, A., Kallai, A., Schunn, C., & Fiez, J. (2010). General improvements in mathematical ability via a basal ganglia learning mechanism. *Poster presented at the 17th Annual Meeting of the Cognitive Neuroscience Society (CNS)*, Montréal, Canada.

Rivera, S. M., Reiss, A. L., Eckert, M. A., & Menon, V. (2005). Developmental change in mental arithmetic: Evidence for increased functional specialization in left inferior parietal cortex. *Cerebral Cortex, 15*, 1779-1790.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology, 80*, 1-27.

Schultx, W. (2000). Multiple reward signals in the brain. *Nature Reviews Neuroscience, 1*, 199-207.

Schultz, W., Tremblay, L., & Hollerman, J. R. (1998). Reward prediction in primate basal ganglia and frontal cortex. *Neuropharmacology, 37*, 421-429.

Shohamy, D., Myers, C. E., Kalanithi, J. & Gluck, M. A. (2008) Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews, 32*, 219–236.

Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Thieme.

Tremblay, L., & Schultz, W. (2000). Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. *Journal of Neurophysiology, 83*,1864-76.

Tricomi, E. M., Delgado, M. R., & Fiez, J. A. (2004). Modulation of caudate activity by action contingency. *Neuron, 41*, 281-292.

Tricomi, E., Delgado, M.R., McCandliss, B.D., McClelland, J.L., & Fiez, J.A. (2006). Performance feedback drives caudate activation in a phonological learning task. *Journal of Cognitive Neuroscience, 18*(6), 1029-1043.

Tricomi E., & Fiez, J.A. (2008). Feedback signals in the caudate reflect goal achievement on a declarative memory task. *NeuroImage, 41*, 1154-1167.

Woods, Cherry, & Mazziotta. (1993). Mri-pet registration with automated algorithm. *Journal of Computer Assisted Tomography, 17*, 536-546.

Yin H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience, 19*, 181-189.