

Direct Instrumental Conditioning of Neural Activity Using Functional Magnetic Resonance Imaging-Derived Reward Feedback

Signe Bray,¹ Shinsuke Shimojo,^{1,2} and John P. O'Doherty^{1,3}

¹Computation and Neural Systems Program, ²Division of Biology, and ³Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, California 91125

Successful learning is often contingent on feedback. In instrumental conditioning, an animal or human learns to perform specific responses to obtain reward. Instrumental conditioning is often used by behavioral psychologists to train an animal (or human) to produce a desired behavior. Shaping involves reinforcing those behaviors, which in a stepwise manner are successively closer to the desired behavior until the desired behavior is reached. Here, we aimed to extend this traditional approach to directly shape neural activity instead of overt behavior. To achieve this, we scanned 22 human subjects with functional magnetic resonance imaging and performed image processing in parallel with acquisition. We delineated regions of interest (ROIs) in finger and toe motor/somatosensory regions and used an instrumental shaping procedure to induce a regionally specific increase in activity by providing an explicit monetary reward to reinforce neural activity in the target areas. After training, we found a significant and regionally specific increase in activity in the ROI being rewarded (finger or toe) and a decrease in activity in the nonrewarded region. This demonstrates that instrumental conditioning procedures can be used to directly shape neural activity, even without the production of an overt behavioral response. This procedure offers an important alternative to traditional biofeedback-based approaches and may be useful in the development of future therapies for stroke and other brain disorders.

Key words: conditioned; conditioning; feedback; fMRI; motor cortex; operant; reward

Introduction

In instrumental conditioning, an animal learns to increase the probability of making a particular response to obtain reward or avoid punishments. Traditionally, the response consists of overt behavioral actions, such as pulling a lever, traversing a maze, or pressing a button (Small, 1901; Thorndike, 1911; Skinner, 1938; Ljungberg et al., 1992; Balleine and Dickinson, 1998). Because the ability to measure neural responses has improved, it has become possible to perform experiments in which an animal is rewarded merely for generating neural activity instead of actually performing an overt motor response (Fetz, 1969). Musallam et al. (2004) demonstrated that by recording from neurons in parietal cortex, monkeys could be trained to generate neural responses to obtain juice rewards, without emitting any behavior.

Parallel advances in human neuroimaging techniques have enabled neural activity measured by functional magnetic reso-

nance imaging (fMRI) to be processed and analyzed in parallel with image acquisition (real-time fMRI), making it possible to provide rapid feedback of activity in specific brain regions to the subject during an ongoing experiment (Cox et al., 1995; Gembris et al., 2000; Yoo and Jolesz, 2002). This technique has previously been used to assess human subjects' ability to modulate their own brain activity, by providing an on-line graphical representation of activity in a specific brain region (Weiskopf et al., 2003; deCharms et al., 2004, 2005). This approach has much in common with traditional biofeedback techniques that have provided on-line feedback of physiological responses such as heart rate or scalp EEG (Schwartz, 1995; Birbaumer et al., 2000).

In the present study, we explore an alternative approach for modulating neural activity to the standard biofeedback paradigm. Here, instead of providing an on-line representation of neural activity and requiring subjects to actively modulate that activity to reach a specified goal, we used procedures derived from instrumental conditioning, whereby an actual reward (monetary gain) is the only feedback subjects receive contingent on their performance. This instrumental training procedure allows one to use "shaping" (Skinner, 1953), in which the threshold for reward is gradually increased to induce incremental improvements in performance.

The aim of the present study was to determine whether it is possible to use instrumental conditioning techniques to modulate neural activity in the human brain. For this, we delineated

Received Jan. 15, 2007; revised June 4, 2007; accepted June 4, 2007.

This work was supported by a Dana Foundation grant to J.P.O., by the Japan Science and Technology Agency ERATO (Exploratory Research for Advanced Technology) Shimojo Implicit brain function project, and by a Gordon and Betty Moore grant to the Caltech Brain Imaging Center. We thank Mike Tzyska for invaluable assistance with technical setup as well as Steven Flaherty and Ralph Lee. We also thank Masamichi Sakagami and Shinsuke Kobayashi for their input.

Correspondence should be addressed to Dr. John P. O'Doherty, Division of Humanities and Social Sciences, California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125. E-mail: johdoherty@hss.caltech.edu.

DOI:10.1523/JNEUROSCI.2118-07.2007

Copyright © 2007 Society for Neuroscience 0270-6474/07/277498-10\$15.00/0

two regions of left sensorimotor cortex (activated by imagined flexion and extension of fingers and toes) and attempted to train subjects to activate one region in response to a visual cue, while suppressing the second region. An additional aim was to determine the extent to which learned modulation of motor cortex in the absence of movement might subsequently influence overt motor behavior as assessed by a speeded reaction-time task. Subjects performed a task in which the cues used during conditioning were alternately displayed on a screen, and intermittent cues instructed them to respond as quickly as possible with fingers or toes. Modulation of reaction times by exposure to instrumental cues offers a measure of how learning of cue contingencies affects concurrent processing of motor responses.

Materials and Methods

Experiment 1

Subjects. A total of 26 right-handed healthy normal human subjects participated in the experiment (14 males and 12 females; age, 18–39 years; mean age, 25.4 years). All subjects gave informed consent, which was approved by the local research ethics committee. The first seven subjects performed only the pretraining and conditioning components of the study. The remaining 19 subjects also performed a reaction-time task before and after conditioning.

Four subjects were removed from the imaging analysis, three of which were also removed from the reaction-time analysis. One subject was eliminated from the imaging analysis because of excessive head movements during the final run. Two other subjects were eliminated from all analyses because of inability to learn the task. An additional subject was removed from all analyses for failing to comply with task instructions. For one subject, the experiment terminated on the ninth trial of the last block because of equipment failure. This left a total of 16 subjects in the reaction-time analysis and 22 subjects in the imaging analysis.

Stimuli. During the conditioning task, subjects were presented with one of three brightly colored abstract fractal images (100×100 pixels) centered on a gray background (800×600 pixels). One of the fractals had the word “Rest” written across it in white letters, to clarify that this cue meant that the subject should be resting. A different set of stimuli were used during the pretraining and functional localizer, in which each task (real or imagined/hand or foot movements) was associated with a centrally presented colored circle with a radius of 100 pixels and lettering coding for each task as described below. During the reaction-time task, the fractal images were presented at an offset of 125 pixels above center, and responses were prompted using the brightly colored circles used in the localizer task. All stimuli were presented using Cogent 2000 (developed by the Cogent 2000 team at the Functional Imaging Laboratory and the Institute of Cognitive Neuroscience, University College London, London, UK) and Cogent Graphics (developed by J. Romaya, Laboratory of Neurobiology, Wellcome Department of Imaging Neuroscience, University College London, London, UK).

Reaction-time task. Subjects performed a simple reaction-time task before and after conditioning. They were randomly presented with one of the three fractal cues used in the conditioning task (referred to here as the background cue), slightly above center on the screen for a time uniformly distributed between 1 and 2 s. After this time, the background cue remained on the screen, and a second cue appeared in the center of the screen (referred to here as the response cue), either a green circle containing the letters “HaT” or an orange circle containing the letters “FoT.” This second cue instructed subjects to respond by pressing a button on the keypad in their hand (HaT) or strapped to the bottom of their foot (FoT). Both cues remained on the screen for 1 s. Subjects responded 30 times to each of the six possible combinations of background cue and response cue.

Pretraining and functional localizer tasks. The functional localizer task consisted of blocks of real and imagined movement, alternating with periods of rest. Subjects ran through this task once outside the scanner as pretraining, so that they could familiarize themselves with the task. Movement tasks consisted of (1) bending fingers II–V at the metacarpophalangeal joint and (2) flexing and extending all five toes through their

full range of movement. During imagination blocks, subjects were instructed to imagine what it would feel like to produce these movements without actually moving. The functional localizer sequence of [resting, finger tapping, resting, imagined finger tapping, resting, toe tapping, resting, imagined toe tapping] blocks was repeated five times. During pretraining, blocks were 10 s in duration, and during the functional localizer performed in the scanner, blocks lasted 15 s. Subjects were cued as to which task to perform by brightly colored visual stimuli with letters coding for the task: a red circle with “R” for rest, a green circle with HaT for hand/finger tapping, a blue circle with “HaI” for imagined hand/finger tapping, an orange circle with FoT for foot/toe tapping, and a yellow circle with “FoI” for foot/toe imagined tapping. During both the pretraining and scanner sessions, subject motion was recorded, as described below. Additionally, during pretraining, we were able to observe subjects at close range to confirm that they were not moving during the imagined movement periods.

Region-of-interest selection. After completion of the functional localizer task in the scanner, the resulting images were sorted into resting and task periods, and *t* tests were applied to generate probabilistic activation maps. Two regions of interest (ROIs) were selected for each subject: one for hand–motor areas activated by imagined finger tapping and one for foot–motor areas activated by imagined toe tapping. In both cases, a mask was generated from the contrast of actual movement versus resting periods. The mask was used to spatially constrain the results of a second contrast comparing imagined movement of fingers to imagined movement of toes. This contrast was chosen to identify regions associated with imagining moving each body part specifically, rather than areas activated by motor imagery in general. From this second map, an ROI center was chosen among the most significant regions, using previous anatomical knowledge of where finger/toe motor cortical areas should be located. A rectangular area of 6×6 voxels in the *x–y* plane and 3 voxels in the *z*-direction was generated around the chosen center. The ROI for each subject comprised a maximum of 108 voxels; in some subjects, this number was smaller if the volume defined by the rectangle stretched beyond the spatial extent of the brain.

Neuroconditioning procedure task and instructions. Subjects were instructed that during this part of the experiment they should never perform any real movements but must only use their imagination or state of mind to increase activity in the specific brain regions defined during the localizer task, corresponding to imagined finger and toe tapping, respectively. A reinforced conditioning trial is illustrated in Figure 1*a*. Each trial began with a resting cue for a variable duration between 15 and 20 s. Next, the subjects saw one of two fractal cues for 15 s. Each “active” cue meant that if the subject sufficiently activated one of the ROIs, they could earn a reward. Data were analyzed on-line after 14 s, and after the 15th second, subjects received visual feedback indicating whether they had successfully earned a reward. Positive reward feedback consisted of a picture of a dollar and the phrase “You have won ONE dollar,” whereas negative feedback was represented by a picture of a scrambled dollar, along with the phrase “You have not won ONE dollar.” Dollars earned during the task corresponded to real money paid to the subject at the end of the experiment. At the start of the experiment, subjects did not know which cue corresponded to which brain region. They were told that they would have to proceed by trial and error to discover the meaning of each cue and that once they learned the meaning, it would stay the same for the duration of the experiment.

Subjects were told that the “resting” period preceding each active period would serve as a baseline against which the activity during the “active” periods would be compared. Therefore, they should try to relax as much as possible during rest periods and not practice mental imagery similar to during the active periods. They were also told that to earn a reward they would have to activate one region specifically and not both regions. Subjects were told that any kind of mental imagery could be appropriate as long as it specifically activated brain regions delineated by the imagined finger- and toe-tapping tasks but that strategies involving motor imagery might be more likely to succeed, given the known functional responses of these regions. Subjects were told that the threshold defining the minimum activity required to get rewarded would be slowly

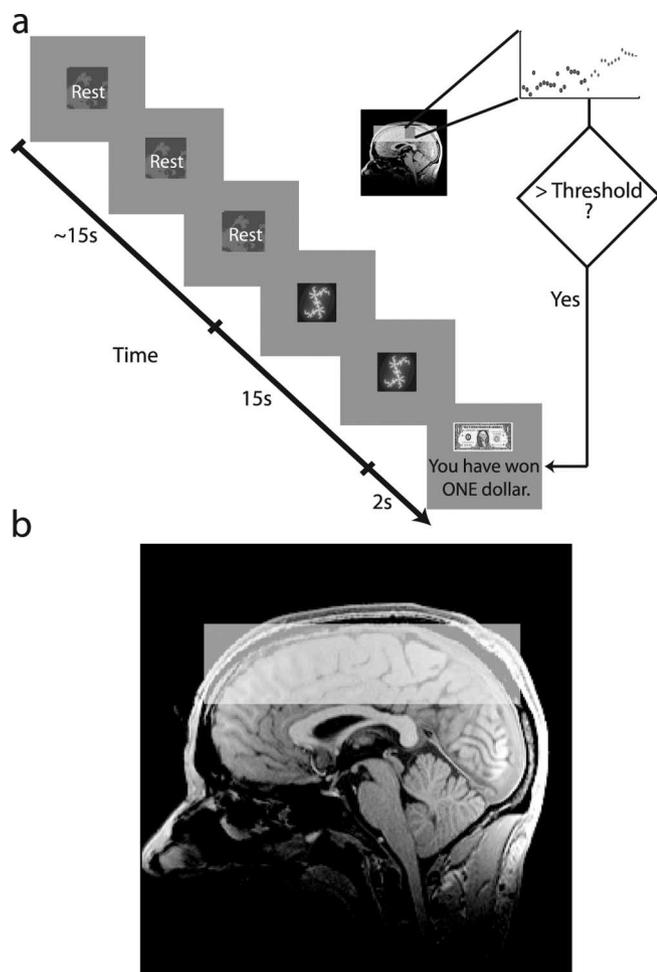


Figure 1. *a*, Example time course for a conditioning trial. Subjects were presented with a resting cue for a variable interval between 10 and 20 s, followed by a cue to activate a specific brain region for 15 s. A percentage signal change value from resting to active was computed on-line and compared with the current threshold. If the threshold was exceeded, subjects were shown a picture of a dollar bill, indicating that they had won one dollar, otherwise a scrambled picture of a dollar was shown, for 2 s. *b*, Diagram showing typical fMRI slice coverage, overlaid on a sagittal slice from a single subject's anatomical scan. We imaged 16 3 mm slices, straight across the top of cortex.

increasing; therefore, they would have to improve on their strategy to continue earning rewards.

The total duration of the experiment was ~ 1.5 h in a single session. In this time, subjects performed reaction-time tasks, pretraining, a functional localizer, and four conditioning blocks consecutively with 14 trials in each; trials were ordered pseudorandomly so that each trial type appeared seven times within a block without three consecutive trials being of the same type. Each block was ~ 8 min long for a total of 32 min of training.

Postexperimental debriefing. After the experiment was completed, subjects were asked to complete a short questionnaire. This form asked them to briefly describe what they were thinking about when they saw each of the three cues (one rest, two active) on the display and to indicate how their strategy might have changed across runs.

Motion recordings. To control for subject motion during periods of imagined movement, we recorded an EMG from the forearm (flexor digitorum superficialis muscle) to measure muscle activity related to finger flexion and extension. We also used a finger-twitch sensor (Biopac Systems, Goleta, CA), placed lengthwise along the bottom of the foot and attached by Velcro around the big toe and below the ball of the foot. The sensor is essentially a variable resistor sensitive to bending and compression and therefore generated a potential difference when subjects bent their toes downward. Both movement recording devices that we used are

MRI compatible, but fMRI scanning introduced noise into the recordings. These data were analyzed by comparing the root mean square (RMS) signal value during resting, active and imagined periods. Recordings obtained while the scanner was running were smoothed using a 15-point (twitch sensor) or 25-point (EMG) median filter to reduce the impact of scanner noise on signal detection.

fMRI scanning procedure. fMRI data were acquired on a Siemens (Erlangen, Germany) 3T TRIO MRI scanner; blood oxygenation level-dependent (BOLD) contrast was measured with gradient echo T2*-weighted echo-planar images (EPIs). We used an eight-channel phased array coil. The first five volumes were discarded to permit T1 equilibration. To keep the repetition time at 1 s, we imaged only 16 3 mm slices across the top of cortex. Typical slice coverage is illustrated on a single subject's anatomical scan in Figure 1*b*. Scan coverage was therefore limited to superior and middle frontal gyri, precentral and postcentral gyri, and superior parietal lobule. Other scan parameters were as follows: in-plane resolution, 3×3 mm; echo time, 30 ms; field of view, 192×192 mm. After the conditioning procedure, a T1-weighted structural image was acquired for each subject, as well as a set of approximately six 32-slice EPIs (to improve coregistration and normalization of images to a template).

Concurrent fMRI analysis and processing. As soon as images were reconstructed, they were transferred in real time via a TCP/IP socket to an external Intel Xeon workstation (3.8 MHz, 64-bit processor running Redhat Linux); data processing was performed using Matlab 7.0 (Mathworks, Natick, MA).

Preprocessing

Image preprocessing consisted of motion correction using AFNI (Cox and Jesmanowicz, 1999) and linear detrending to correct for low-frequency scanner drift. During functional localizer scans, spatial smoothing using a two-dimensional Gaussian of 5 mm width was completed before performing statistical tests. During the conditioning task, no temporal or spatial smoothing was performed.

Reward criterion

Two thresholds shared equal priority in the decision rule for determining whether a subject had earned a reward on a particular trial: one threshold on the minimum percentage signal change within a region and a second threshold on the difference between the percentage signal change in the rewarded region and the nonrewarded region. Both thresholds had to be exceeded for a subject to earn reward on a given trial, and both were adapted according to a modified percentile reinforcement schedule (Galbicka, 1994). They both started at 0 and increased only after the current threshold had been exceeded four times. At this time, both thresholds were set to be the lowest of the four values that had beaten the previous value. If a reward was not obtained on one of the next four trials, one or both thresholds was reset to its previous value, depending if one or both conditions was not met. In this way, the thresholds for the signal level and the difference increased together but were reset separately.

As images arrived on the external workstation, they were preprocessed, and the signal was averaged over all voxels in the previously defined ROIs. After one variable-length baseline period (10–20 s) and one 14 s active period had elapsed, the ROI signal was averaged over each time period (excluding the first 2 s to allow for some lag in the hemodynamic response), and a percentage signal change from baseline to active was computed for each ROI. For the ROI being rewarded, the percentage signal change was compared with the current threshold, and the difference between the percentage signal change in the two ROIs was compared with the difference threshold. If both conditions were met, the current trial was “rewarded” after the 15th second, when the subject would see the “reward” feedback for 2 s. If the reward conditions were not met, they would see the “no reward” feedback for 2 s.

Performance-based grouping of subjects. For some analyses, we divided subjects into groups depending on their performance during the last experimental run: subjects who earned fewer than five rewards on the final run were classified as poor learners, relative to those who earned more than five rewards. Performance during the last conditioning run was especially relevant to analysis of the reaction-time measures taken

immediately afterward, because that should give the most current estimate of the subjects' level of learning. Some subjects reported tiring toward the end of the experiment, which could corrupt learning-related effects in the reaction-time analysis. This criterion put 17 subjects in the "good-learner" category and 5 subjects in the "poor-learner" category.

Group fMRI percentage signal change analysis. We performed a group analysis on the trial-by-trial percentage change values measured during conditioning. Trials in which twitching movements were visible in the EMG traces were eliminated, as were trials in which large head movements caused sharp deflections in the BOLD signal time course. We performed a repeated-measures ANOVA on the averaged percentage signal change values during each run, with within-subject factors of ROI (three levels: hand ROI, foot ROI, and whole-brain background ROI), rewarded ROI (two levels: hand rewarded and foot rewarded) and run (four levels) and a single between-subjects factor, a binary value indicating whether or not the subject was a good learner.

To look for trial-by-trial increases in signal difference between the two ROIs, we averaged the percentage signal change value across subjects on each trial and performed a linear regression on the difference between the signals in each ROI.

Statistical parametric mapping analysis. Data were preprocessed using the SPM5 software package (statistical parametric mapping; <http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>). Images were corrected for slice timing and spatially realigned to the first image from the functional localizer. One of the 32-slice EPIs collected at the end of the experiment was used to improve coregistration and spatial normalization. The 16-slice EPIs were coregistered to a 32-slice EPI, which in turn was coregistered to the T1-weighted anatomical scan. The T1 image was segmented into white and gray matter, and the gray matter was coregistered and normalized to the template gray matter image distributed with SPM5 [in Montreal Neurological Institute (MNI) space]. These parameters were subsequently applied to the T1 image itself as well as the set of 16-slice EPIs. Spatial smoothing was then applied to the 16-slice EPIs using a Gaussian kernel with a full-width at half-maximum of 8 mm.

The four conditioning sessions for each subject were modeled in SPM using a finite impulse response model, with separate regressors for hand and foot rewarded trials, for each run. The six ongoing motion parameters estimated during realignment were included as regressors of no interest.

Parameter estimates were modeled with a full factorial model with two factors: rewarded region (two levels) and session (four levels). This created an eight-column design matrix for each subject, with each column corresponding to a session \times rewarded-region interaction term. Linear contrast images from these design matrices were taken to the random-effects level by applying *t* tests between them to produce group statistical parametric maps.

Reaction time analysis. During the conditioning task, subjects learned to associate the fractal cues with either a hand-imagine or foot-imagine response, so that after the experiment, the background cues can be considered either compatible (e.g., hand-imagine cue and hand-response cue) or incompatible (e.g., hand-imagine cue and foot-response cue). Trial-by-trial reaction times measured before and after the conditioning task were divided into three blocks (early, middle, and late) and averaged within each trial type. The block-averaged reaction times were analyzed with a repeated-measures ANOVA, with within-subject factors of block (three levels), time (two levels: before and after conditioning), cue-response relationship (three levels: compatible cue, incompatible cue, rest cue), response type (two levels: hand and foot), and a single between-subjects factor, a binary value indicating whether or not a subject performed well on the task.

Experiment 2

In experiment 1, the behavioral reaction-time measure was taken outside the scanner, before and after the experiment. This meant that subjects were exposed to the cues in a different context and that any response evoked by the fractal cues could diminish because the test was performed in extinction (responses were not rewarded). We conducted a follow-up experiment to test the effect of performing the reaction-time measure in a similar context as the conditioning and interleaved reaction-time mea-

surements with conditioning trials to minimize the effects of extinction. In experiment 2, we also included a condition to control for the effects of repeated practice alone, without contingent feedback.

Subjects. We scanned an additional nine subjects in an alternate version of the conditioning task that included a reaction-time measure taken during the conditioning trials (age, 21–38 years; mean age, 24.9 years; three males). One subject performed only three of four sessions because of discomfort, and the imaging data from another were not analyzed because of excessive head movements. We also scanned nine subjects (age, 21–34 years; mean age, 23.3 years; four males) in a control condition.

Conditioning with interleaved reaction-time task. A separate group of nine subjects underwent a conditioning procedure nearly identical to experiment 1, but with additional sets of reaction-time trials randomly inserted among the regular conditioning trials in each block. For this task, subjects held a button pad in their right hand, and a second button pad was held against the bottom of their foot in a sandal so that they could push a button with their toe. The reaction-time trials began with a central fixation cross presented for 250 ms, followed by one of the two fractal cues from the conditioning trials for 1.75 s. Either HaT or FoT then appeared on the fractal for 250 ms, instructing subjects to respond by pressing the hand button or foot button, respectively. The fixation cross appeared for 1.75 s, during which time subjects made their response. Each session consisted of 14 conditioning trials with two sets of 30 consecutive reaction-time trials inserted at pseudorandom intervals. In the first session, they always appeared after the 12th and 14th trial, to give subjects the opportunity to learn the response associated with each fractal. In subsequent sessions, the blocks of reaction-time trials appeared at random intervals, with the condition that two blocks could not be presented consecutively. In each session, 15 of each type of trial (hand-cue/hand-response, hand-cue/foot-response, foot-cue/hand-response, foot-cue/foot-response) were presented in random order in two sets of 30 trials. Across the four conditioning blocks, a total of 60 reaction times from each type of trial was collected.

Control task. We ran a second version of the task designed to control for the effects of repeated practice of motor imagery. During the conditioning trials, these subjects were instructed to imagine either hand or foot movements when they saw the corresponding cue and to ignore the reward feedback. Unlike in the feedback task, the rewards delivered to subjects were not linked to neural activity, but instead each subject in the control group experienced the rewards obtained by a randomly assigned "yoked" subject from the feedback group. The control task also included reaction-time trials identical to those in the feedback task.

Reaction-time analysis. During the conditioning task, subjects learned to associate the fractal cues with either a hand-imagine or foot-imagine response, so that the background cues can be considered either compatible (e.g., hand-imagine background cue and hand-response cue) or incompatible (e.g., hand-imagine background cue and foot-response cue) with the response. We hypothesized that there would be a facilitation for compatible stimuli relative to incompatible stimuli (i.e., faster reaction times). We log transformed these data and entered them into a three-way repeated-measures ANOVA, with within-subject factors of cue (hand/foot), response (hand/foot), and session (1–4), separately for the feedback and control groups.

Results

Experiment 1

Behavioral results

Postexperimental debriefing. According to the questionnaire responses, all subjects, except two, correctly discriminated between the two cues and were aware which cue instructed them to activate hand or foot areas. The two subjects who did not learn correctly performed very poorly at the task and were eliminated from both the behavioral and imaging analyses. The self-reported strategies for activating these areas all involved motor imagery of some kind. During the resting period, subjects either relaxed and let their mind wander or distracted themselves by repeating a

Table 1. Movement recording comparisons

	Condition compared with rest	Number of subjects showing significant difference
Pretraining (Wilcoxon rank sum; $n_1 = 5, n_2 = 5; p < 0.05$)	Finger tapping	16/22
	Imagined finger tapping	0/22
	Toe tapping	21/22
Functional localizer (Wilcoxon rank sum; $n_1 = 5, n_2 = 5; p < 0.05$)	Imagined toe tapping	0/22
	Finger tapping	19/22
	Imagined finger tapping	0/22
	Toe tapping	20/22
Conditioning task (Wilcoxon rank sum; $n_1 = 7, n_2 = 7; p < 0.05$)	Imagined toe tapping	0/22
	Imagined hand movement	0/22
	Imagined foot movement	0/22

song or numbers in their head. Some subjects reported making eye movements during these periods.

Subject performance on conditioning task. Most subjects were able to successfully obtain rewards in the task. The mean number of rewards obtained per run were [$3.27 \pm 0.40, 3.41 \pm 0.31, 2.82 \pm 0.30, 3.27 \pm 0.35$] for hand rewarded trials and [$3.05 \pm 0.32, 2.82 \pm 0.37, 3.09 \pm 0.40, 2.82 \pm 0.39$] for foot rewarded trials. The number of rewards remained relatively constant across runs; a repeated-measures ANOVA with between-subjects factors of rewarded region (two levels) and run (four levels) yielded no significant main effects or interactions. However, the threshold for the activation level that subjects had to achieve to obtain reward increased across trials. A linear regression on the trial-by-trial mean threshold across subjects shows a significant increase, both for the hand rewarded threshold ($\beta = 9.05 \times 10^{-5}; R^2 = 0.8717; p < 0.001$) and foot rewarded threshold ($\beta = 1.51 \times 10^{-4}; R^2 = 0.9634; p < 0.001$). Because subjects were able to maintain a constant rate of reward despite the increasing difficulty of the task, we consider this a measure of overall success of the conditioning procedure.

Movement recordings. Movement recordings during the localizer task, both during pretraining and in the scanner, confirmed that subjects were able to perform the imagination task without actually moving. We compared RMS values during resting periods to real and imagined movement periods, during the pretraining, functional localizer, and conditioning task. The results are summarized in Table 1, and example recordings for real and imagined movements with fingers and toes are shown in Figure 2. For some subjects, the difference between rest and movement did not reach significance; inspection of the movement time courses showed that these subjects had probably adjusted their position during the resting period of one of the blocks, and because of the small number of blocks ($n = 5$), the comparison did not reach significance.

Although subjects were instructed that they should keep their hands and feet still, some subjects showed evidence of hand twitches during certain trials. Trials in which sharp spikes in the EMG indicated a small twitch in the hand or arm, during rest, hand-imagined, or foot-imagined periods, were removed from further analysis. The mean number of trials eliminated per subject was 5 of 56, with a SD of 4.33.

Reaction times. The repeated-measures ANOVA performed on the reaction times yielded a significant main effect of response type ($p < 0.001; F_{(1,14)} = 48.374$) and significant interactions of time \times cue–response relationship \times learner type ($p < 0.01; F_{(2,28)} = 6.011$) and time \times response type ($p < 0.01; F_{(1,14)} = 13.292$). Planned *t*-contrasts showed that in good learners there was a significant difference between both compatible (paired *t* test, $p < 0.05; n = 11; |t| = 2.1871$) and incompatible (paired *t*

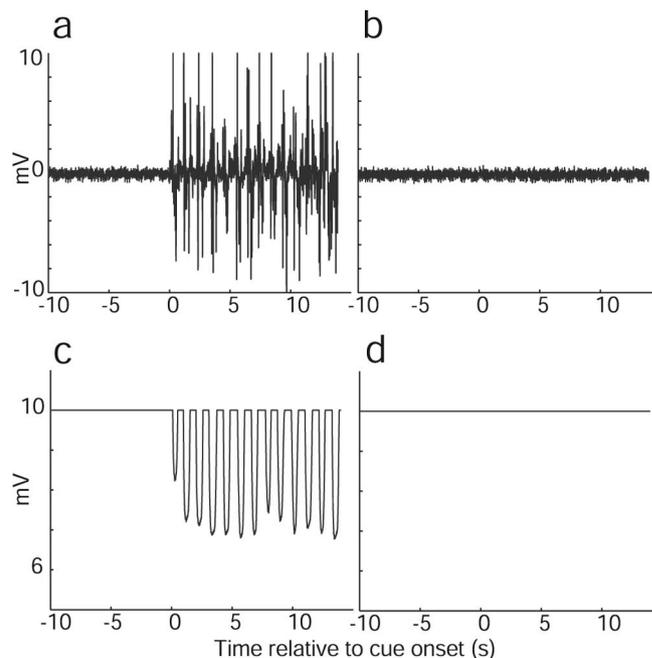


Figure 2. Sample movement recordings from a single subject during experiment 1. *a*, EMG during real hand movement. *b*, EMG during imagined hand movement. *c*, Variable resistor recording during real foot movement. *d*, Variable resistor recording during imagined foot movement.

test, $p < 0.05; n = 11; |t| = 1.8013$) cue types after conditioning compared with before, with both becoming slower after conditioning. However, contrary to our hypothesis, we did not observe a significant difference between responses that were compatible or incompatible with the background cue. To address the possibility that the absence of this effect was a consequence of the reaction time measure being performed outside the scanner and therefore in extinction, in experiment 2, reaction times were tested in the scanner interleaved with conditioning trials to reduce extinction of the response (for more details, see Materials and Methods, Experiment 2).

fMRI results

ROI location. The mean ROI center for the hand region in MNI space was [$-35 \pm 1.24, -26 \pm 1.9, 65 \pm 0.941$] located on the left precentral gyrus (Brodmann areas 4a, 6, and 1) (Eickhoff et al., 2005); individual-subject ROI centers were located near the hand knob (Yousry et al., 1997) on the precentral and postcentral gyri. The mean ROI center for the foot region was [$-6 \pm 0.729, -25 \pm 1.5, 69 \pm 1.1$] located on the left paracentral lobule (Brodmann areas 4a and 6); individual-subject ROI centers were distributed from the posterior part of the superior frontal gyrus

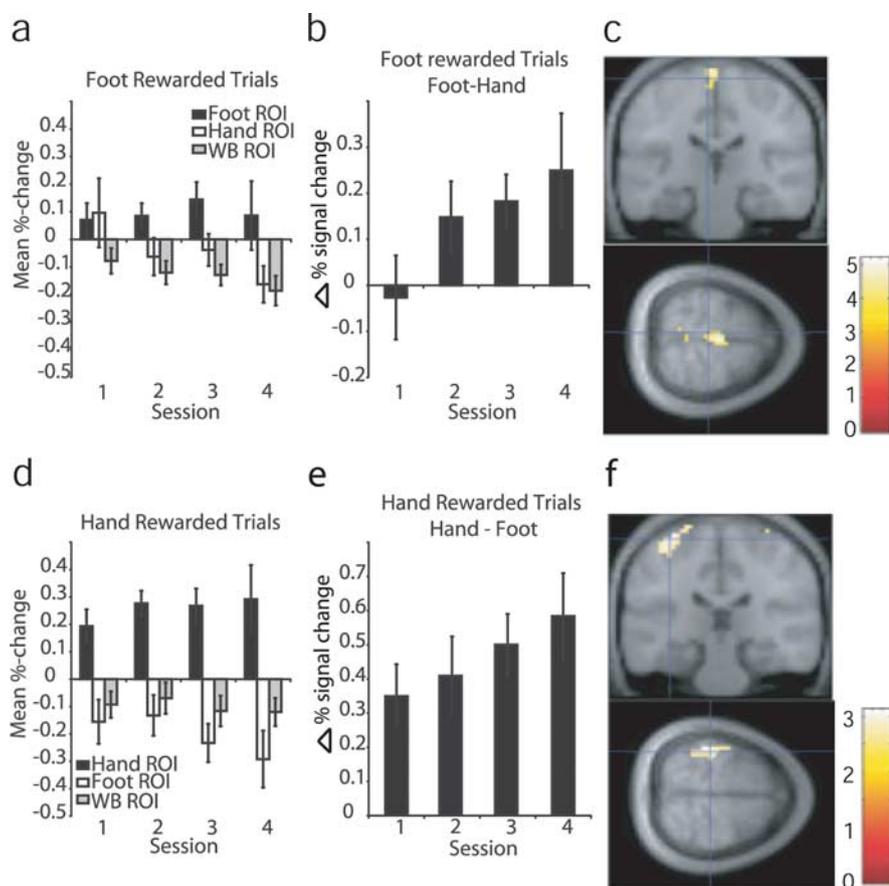


Figure 3. fMRI results from experiment 1. *a*, Mean percentage signal change data averaged over subjects within runs in each ROI during trials in which the foot ROI was rewarded. *b*, Difference in mean percentage signal change data, averaged over subjects within runs, between foot and hand ROIs during foot rewarded trials. *c*, Results of random-effects analysis in SPM from experiment 1; *t* test on contrast increasing during foot rewarded trials and decreasing during hand rewarded trials, thresholded at $p < 0.01$; cross-hairs indicate mean of subjects' ROI centers for the foot ROI [−6, −25, 69]. *d*, Averaged responses in each ROI during trials in which the hand ROI was rewarded. *e*, Difference between hand and foot ROIs during hand rewarded trials. *f*, Results of random-effects analysis in SPM from experiment 1; *t* test on contrast increasing during hand rewarded trials and decreasing during foot rewarded trials, thresholded at $p < 0.001$; cross-hairs indicate mean of subjects' ROI centers for the hand ROI [−35, −26, 65]. Error bars indicate SE. WB, Whole brain background.

along the length of the paracentral lobule. These areas are highly consistent with finger and toe imagery-specific locations described by Ehrsson et al. (2003).

Trial-by-trial percentage signal change in ROIs. Averaging the trial-by-trial percentage signal change data across trials within each session and over subjects, we see a general increase in signal in the rewarded ROI and a decrease in the nonrewarded ROI, corresponding to an overall increase in the signal difference between the rewarded and nonrewarded regions; these data are plotted in Figure 3. In addition to the two predefined ROIs, we also looked at the signal in a large background ROI that included all brain voxels outside of the two task-related ROIs. The background ROI did not show the same increase as the rewarded ROI, confirming that the activation in response to the cue was specific to the rewarded ROI rather than reflecting a nonspecific increase in brain activity.

To test for a learning effect, we performed a repeated-measures ANOVA on the trial-averaged percentage signal change measures within each session from each ROI. Across all 22 subjects, we found a significant main effect of ROI ($p < 0.005$; $F_{(6,120)} = 6.246$), and significant interactions of ROI × rewarded ROI ($p < 0.001$; $F_{(2,40)} = 14.308$), as well as an interaction be-

tween ROI × rewarded ROI × session that approached significance ($p = 0.064$), suggesting a learning effect. By restricting our analysis to a subgroup that successfully met a learning criterion of five or more rewards during the last session ($n = 17$), this interaction became significant ($p < 0.05$; $F_{(6,96)} = 3.907$). Taking the trial-by-trial average across all subjects and regressing the mean difference between ROIs onto trial number, we found a significant positive increase, both for [hand ROI − foot ROI] in trials when an increase in the hand ROI was rewarded ($\beta = 0.0001$; $R^2 = 0.3840$; $p < 0.05$) and [foot ROI − hand ROI] in trials when an increase in the foot ROI was rewarded ($\beta = 0.0001$; $R^2 = 0.2557$; $p < 0.05$).

Random-effects analysis with SPM. We generated a contrast to detect regions in which signal increased during hand rewarded trials and decreased during foot rewarded trials and, likewise, a second contrast to detect regions with signal increase during foot rewarded trials and decreases during hand rewarded trials. Taken to the random-effects level, the contrast to detect activity during foot rewarded trials showed a significant cluster with peaks surviving small-volume correction around the mean foot ROI center [−6, −24, 69] at [−3, −24, 75] [$t = 5.06$; $p < 0.01$, false discovery rate (FDR) corrected], [0, −21, 72] ($t = 4.95$; $p < 0.01$, FDR corrected), and [−3, −18, 69] ($t = 4.71$; $p < 0.01$, FDR corrected). The results of this contrast are shown in Figure 3*c*.

The contrast to detect activity during hand rewarded trials shows a large cluster with a peak at [−39, −33, 66], which survives small-volume correction in an 8 mm sphere around the mean of subjects' hand ROI centers [−36, −27, 66] ($k = 38$; $t = 3.13$; $p < 0.05$, FDR corrected). The results of this contrast are shown in Figure 3*f*.

The hand-region and foot-region activation tasks engaged a network of brain regions in addition to the ROIs, although activations in these regions remained relatively constant across the study (Table 2). As would be expected, there was substantial overlap between regions activated by imagined hand and foot movements, in the dorsal premotor (PMd) region extending into the supplementary motor area (SMA) and pre-SMA as well as bilateral regions of the parietal cortex and precentral gyri. In Figure 4, we have plotted the parameter estimates for the hand and foot rewarded trials in each of the four sessions; Figure 4*a* shows the significant regions in the foot-region activation task, and Figure 4*b* shows the hand-region activation task. Despite the fact that these regions were generally activated by subjects performing the task, our protocol caused selective enhancement and depression of activity only in the delineated ROIs. This can be seen from the slopes and divergence of the curves in the topmost plots (from the peak voxel near the ROI centers described above) compared with the other regions significantly activated by a general task–baseline contrast.

Table 2. Imagine to activate tasks: Z scores and MNI coordinates of peak activation foci

Region	Contrast			
	Hand rewarded		Foot rewarded	
	Number of voxels	Z	Number of voxels	Z
PMd	1001	5.79 (−12, −3, 69)	997	5.94 (−12, −6, 72)
Left precentral gyrus			6	3.29 (−48, 0, 54)
Right precentral gyrus	68	4.79 (57, 0, 48)	44	4.22 (57, 0, 48)
Right parietal (supramarginal gyrus)	64	5.40 (60, −27, 51)	47	5.47 (60, −27, 54)
Left parietal (supramarginal gyrus)	79	4.49 (−36, −48, 60)		
Left superior parietal gyrus			71	4.35 (−18, −60, 69)
Right middle frontal gyrus	5	3.51 (30, −3, 72)		

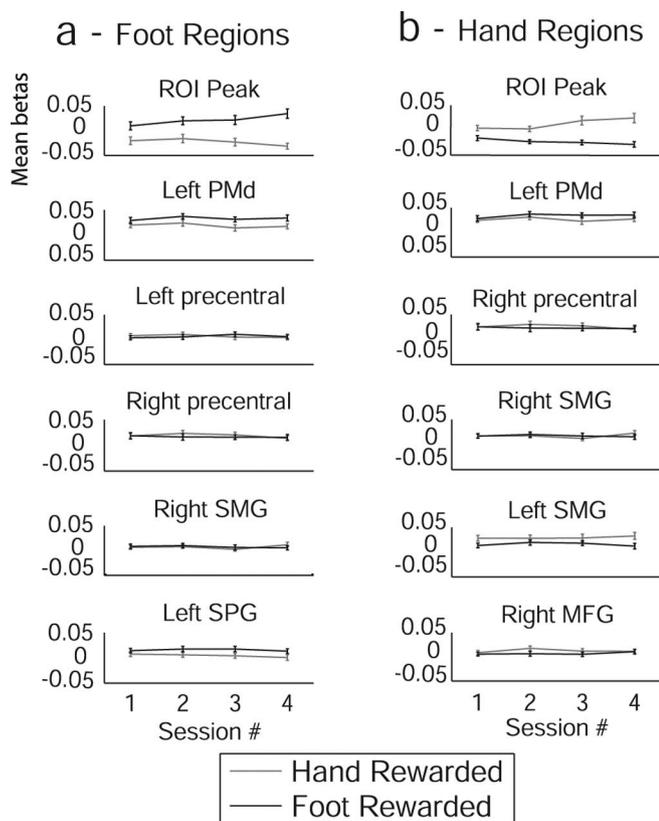


Figure 4. Subject-averaged parameter estimates across sessions from experiment 1. Hand and foot rewarded trials are plotted separately. Error bars indicate SEs. Regional coordinates are as in Table 2. **a**, Regions identified as significant during trials when subjects were rewarded for activating the foot region. **b**, Regions identified as significant during trials when subjects were rewarded for activating the hand region. MFG, Middle frontal gyrus; SMG, supramarginal gyrus; SPG, superior parietal gyrus.

Experiment 2

Behavioral results

Reaction times. The results of the ANOVA on the feedback group showed that subjects were significantly faster to make a response when the background cue was compatible with the type of response, as demonstrated by a significant interaction between cue and response ($p < 0.05$; $F_{(1,7)} = 7.23$). We also found significant main effects of session ($p < 0.05$; $F_{(3,21)} = 4.134$) and response ($p < 0.01$; $F_{(1,7)} = 17.7$); subjects responded more quickly with fingers than toes. In the control group, only the main effect of response was significant ($p < 0.05$; $F_{(1,7)} = 11.29$), in that subjects were faster responding during hand than foot movements,

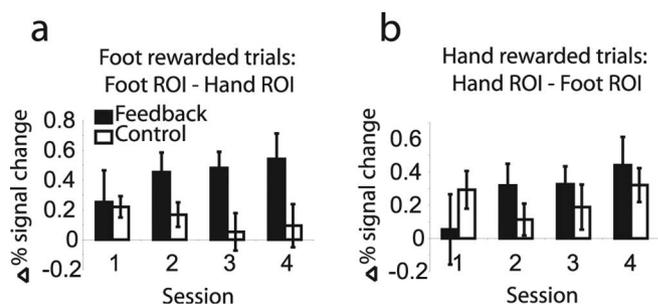


Figure 5. Percentage signal change plots across sessions for feedback and control groups from experiment 2. **a**, Difference in percentage signal change between foot and hand ROIs when foot responses were rewarded. **b**, Difference in percentage signal change between hand and foot ROIs when hand responses were rewarded. Error bars indicate SE.

but no significant cue or cue \times response effects were found in this group.

fMRI results

ROI location. The ROIs identified in experiment 2 were similar to experiment 1. For the feedback group, the mean ROI center in MNI space was $[-39 \pm 2.2, -25 \pm 2.1, 58 \pm 1.8]$ for the hand region and $[-6 \pm 0.7, -25 \pm 1.5, 69 \pm 1.1]$ for the foot region. The ROI centers for the control group were statistically indistinguishable from the feedback group, with the mean hand ROI center at $[-37 \pm 2.2, -23 \pm 1.1, 56 \pm 1.3]$ and the mean foot ROI center at $[-7.6 \pm 0.6, -29 \pm 2.5, 69 \pm 1.0]$.

Trial-by-trial percentage signal change in ROIs. We averaged the trial-by-trial percentage signal change data across trials within each session and over subjects. In the feedback group, we see a general increase in signal in the rewarded ROI and a decrease in the nonrewarded ROI, corresponding to an overall increase in the signal difference between the rewarded and nonrewarded regions. In the control group, the difference between the two regions is stable or decreasing. These data are plotted in Figure 5; Figure 5a shows the percentage signal change difference for the foot-region activation task, and Figure 5b shows the hand-region activation task. Taking the trial-by-trial average across all subjects and regressing the mean difference between ROIs onto trial number, in the feedback group we found a significant positive increase, both for [hand ROI – foot ROI] in trials when an increase in the hand ROI was rewarded ($R^2 = 0.23$; $p < 0.05$) and [foot ROI – hand ROI] in trials when an increase in the foot ROI was rewarded ($R^2 = 0.22$; $p < 0.05$). In contrast, no significant linear increase was seen in the control group, either in the hand-imagine or in the foot-imagine conditions, suggesting that repeated practice of motor imagery is not sufficient to explain the shaping of neural responses demonstrated here and in experiment 1.

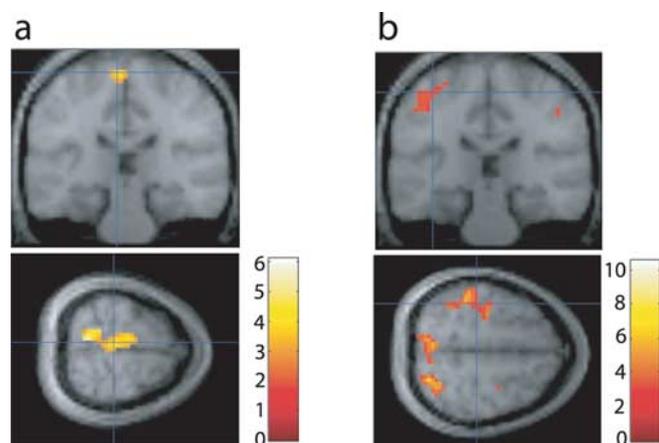


Figure 6. Random-effects and ROI analyses from experiment 2. *a*, Results of a contrast of foot-cue versus hand-cue conditions across all four sessions. Cross-hairs are centered on mean of subjects' ROI centers for the foot ROI [$-6, -30, 69$]. Results are shown at $p < 0.01$ for visualization but survive correction for small volume at $p < 0.05$. *b*, Results of a contrast of hand-cue rewarded versus foot-cue rewarded trials across all four sessions from the feedback group. Cross-hairs are centered on mean of subjects' ROI centers for the hand ROI [$-39, -27, 57$]. Results are shown at $p < 0.01$ for visualization but survive correction for small volume at $p < 0.05$.

Random-effects analysis with SPM. We generated contrasts comparing activity during hand-*imagine* periods and foot-*imagine* periods and took them to the random-effects level. Consistent with the results from experiment 1, significant activity was found in the foot region in the contrast of foot-cue trials $>$ hand-cue trials (Fig. 6*a*), within an 8 mm sphere corrected for small volume around the mean center of the foot ROIs for the feedback group at [$-6, -27, 69$] ($t = 4.04$; $p < 0.05$, FDR corrected). Significant activity was also found in the hand region in the contrast of hand-cue trials $>$ foot-cue trials (Fig. 6*b*), which survived correction for small volume within an 8 mm sphere centered around the mean of the hand ROIs for the feedback group at [$-42, -33, 54$] ($t = 5.66$; $p < 0.05$, FDR corrected).

ROI-based comparison of effects in feedback and control groups. We next compared the mean parameter estimates from each ROI between the feedback and control groups. During the hand-cue condition, neural activity in the hand ROI was significantly greater in the feedback than the control group during the last two sessions once learning was consolidated in the feedback group ($t_{(15)} = 1.9$; $p < 0.05$, one-tailed). During the last two sessions of the foot-cue condition, neural activity in the foot ROI was also significantly greater in the feedback group than the control group ($t_{(15)} = 3.2$; $p < 0.005$).

Discussion

In this study, we have shown that it is possible to directly condition neural activity using reward feedback derived from fMRI. Subjects were able to discriminate between two cues and respond to each by activating the appropriate region of their left sensorimotor cortex, while suppressing activity in a second region. *Post hoc* analysis showed that the brain regions significantly increasing in response to rewarded cues and decreasing in response to non-rewarded cues were spatially limited to the specific brain regions where activity was reinforced in our procedure. We also demonstrated in a control group that repeated practice of motor imagery alone is not sufficient to account for this effect. A behavioral reaction-time measure showed that in the context in which the association was learned, a neural response to a cue can have a

facilitatory effect on reaction times, when the physical response engages regions similar to those activated by the learned neural response. Together, these findings could lead to development of therapies for patients who have suffered stroke damage to the motor system.

Behavioral shaping has long been known to be a powerful method for behavioral modification in both humans and animals (Thorndike, 1911; Skinner, 1953). Here, we have used the methods derived from behavioral shaping to directly shape neural activity. Our goal in this study was to show that by using a reward schedule based on behavioral shaping we could train subjects to increase the level of their neural responses in a specific brain region over time. Shaping schedules constantly adjust the threshold required to earn reward, based on subjects' previous performance, thus ensuring that subjects are in a state of constant learning (Keil et al., 2001). Our procedure succeeded not only in increasing activity over time but also in selectively increasing and decreasing activities in the specific ROIs, whereas activities in other regions recruited by this task remained stable.

The approach used here offers an important alternative to that used in previous fMRI neurofeedback training studies (Weiskopf et al., 2003; deCharms et al., 2004) (for review, see Weiskopf et al., 2004; deCharms et al., 2005). In these previous studies, explicit visual feedback was provided to subjects signaling the level of activity in a particular area. Subjects were then instructed to modulate their activity to attain a specific target level of activation. However, in the present study, no visual feedback was presented. Subjects were instructed to activate a specific brain region and received an actual tangible reward (here winning one U.S. dollar) if they succeeded in reaching a criterion on a given trial. One potential advantage of the present technique over the classical biofeedback approach is that provision of tangible rewards may be much more motivating for subjects than the instruction to reach a target activation level in the absence of extrinsic reward. Another possible advantage of the present technique is that the use of instrumental conditioning instead of a visual biofeedback procedure may render the task much less "cognitive" and thus less likely to require high-level or effortful cognitive processing. Thus, the present technique may be efficacious even under situations when subjects are either incapable or unwilling to engage in effortful cognitive processing, or when a cognitively demanding task is concurrently imposed. Furthermore, the present technique may not even require subjective conscious awareness of task progress to be effective, given that instrumental conditioning procedures are known to work in a wide variety of animal species including rats, pigeons, and even *Aplysia* (Chen and Wolpaw, 1995; Green et al., 1999; Hawkins et al., 2006), which one might speculate are unlikely to have developed conscious subjective awareness to the same degree as in humans. This raises the intriguing possibility that human brain regions may differ in the degree to which successful neural conditioning is associated with a subjective conscious correlate. Here, subjects reported using, and were instructed to use, a conscious strategy of imagining movement during task performance. Future studies could probe the subjective correlates of conditioning in different brain regions to examine whether, for example, subjective correlates of neural conditioning in higher cortical areas are qualitatively different than those associated with subcortical structures. Finally, the use of an approach based on instrumental conditioning means that we can benefit from the extensive work done in this area to inform our understanding of the neural and behavioral processes mediating this learning (Balleine and Dickinson, 1998;

Gottfried et al., 2002; O'Doherty et al., 2003; Samejima et al., 2005; Yacubian et al., 2006).

The task of differentially activating two motor cortical regions seemed to engage parallel learning processes: as the signal in the ROI being rewarded increased over time, we saw a corresponding decrease in the ROI not being rewarded. Subjects reported activating the rewarded ROI using kinesthetic motor imagery; however, the signal decrease observed in the nonrewarded ROI may not be attributable to the same deliberate control. To continue earning rewards throughout the task, subjects had to increase the difference in signal between the two ROIs. Such differential neural sensitivity to the reward conditions may tap into covert associative learning mechanisms over and above the explicit imagery strategy the subjects reported using, as demonstrated in previous instrumental conditioning experiments (Hefferline et al., 1959; Svartdal, 1995).

Although not all functional imaging studies of motor imagery have reported activations in primary motor cortex (M1) (Stephan et al., 1995; Gerardin et al., 2000), several fMRI studies have shown evidence for somatotopically organized activations in primary motor cortex during motor imagery (Leonardo et al., 1995; Ehrsson et al., 2003). We report here that activation in somatotopically specific regions of primary motor and sensory cortices increased over the course of conditioning. This enhancement could arguably be a side effect of repeated practice of mental imagery and not dependent on the reward feedback. However, the control task performed in experiment 2 suggests that feedback does facilitate learning. Furthermore, it is difficult to explain the suppression in the nonrewarded ROI without the requirement that we imposed for differential activity to earn reward, suggesting that in our study, provision of reward based on neural activity led to specific shaping of the neural response. Nyberg et al. (2005) compared the effects of mental practice to physical practice in a recent fMRI study. They found that practice in general led to a more regionally specific activation in motor cortex. They also found a differential increase in visual cortex activity in the mental practice group. Studies comparing kinesthetic and visual imagery have found that they evoke different patterns of neural activity (Neuper et al., 2005; Stinear et al., 2006). Because we found an increase in activity specific to sensorimotor cortex, perhaps the feedback from this area caused subjects to refine their imagery strategy to favor kinesthetic rather than visual. A similar effect was described by Yoo et al. (2006), in which verbal feedback of auditory cortex activation was found to influence subjects' strategies during selective attention to auditory stimuli. Similarly, Posse et al. (2003) gave subjects feedback of amygdala activation during sad mood induction, resulting in amygdala activations that correlated with sad mood. Generally speaking, training subjects to activate a particular part of their brain while performing a task could be a way of enhancing task performance or correcting deficits. Training subjects to make more efficient use of neural resources could potentially lead to long-term alterations in neural plasticity related to performance of specific tasks.

In summary, we have presented an instrumental conditioning technique that succeeds in shaping an increase in sensorimotor cortical responses over time, as measured with fMRI. We also used a behavioral measure to explore the effects of training on behavior. The method presented here extends previous work (Posse et al., 2003; Weiskopf et al., 2003; deCharms et al., 2004, 2005; Yoo et al., 2006) by incorporating a well studied operant conditioning paradigm with fMRI-derived neurofeedback training. This method was successful in conditioning a differential response between two regions with a very high neuroanatomical

precision, a finding that could have clear benefit in future clinical applications.

References

- Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37:407–419.
- Birbaumer N, Kubler A, Ghanayim N, Hinterberger T, Perelmouter J, Kaiser J, Iversen I, Kotchoubey B, Neumann N, Flor H (2000) The thought translation device (TTD) for completely paralyzed patients. *IEEE Trans Rehabil Eng* 8:190–193.
- Chen XY, Wolpaw JR (1995) Operant-conditioning of H-reflex in freely moving rats. *J Neurophysiol* 73:411–415.
- Cox RW, Jesmanowicz A (1999) Real-time 3D image registration for functional MRI. *Magn Reson Med* 42:1014–1018.
- Cox RW, Jesmanowicz A, Hyde JS (1995) Real-time functional magnetic resonance imaging. *Magn Reson Med* 33:230–236.
- deCharms RC, Christoff K, Glover GH, Ludlow D, Pauly JM, Soneji D, Gabrieli JDE (2004) Learned regulation of spatially localized brain activation using real-time fMRI. *NeuroImage* 21:436–443.
- deCharms RC, Maeda F, Glover GH, Ludlow D, Pauly JM, Soneji D, Gabrieli JDE, Mackey SC (2005) Control over brain activation and pain learned by using real-time functional MRI. *Proc Natl Acad Sci USA* 102:18626–18631.
- Ehrsson HH, Geyer S, Naito E (2003) Imagery of voluntary movement of fingers, toes, and tongue activates corresponding body-part-specific motor representations. *J Neurophysiol* 90:3304–3316.
- Eickhoff S, Stephan KE, Mohlberg H, Grefkes C, Fink GR, Amunts K, Zilles K (2005) A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *NeuroImage* 25:1325–1335.
- Fetz EE (1969) Operant conditioning of cortical unit activity. *Science* 163:955–958.
- Galbicka G (1994) Shaping in the 21st century: moving percentile schedules into applied settings. *J Appl Behav Anal* 27:739–760.
- Gembris D, Taylor JG, Schor S, Frings W, Suter D, Posse S (2000) Functional magnetic resonance imaging in real time (FIRE): sliding-window correlation analysis and reference-vector optimization. *Magn Reson Med* 43:259–268.
- Gerardin E, Sirigu A, Lehericy S, Poline JB, Gaymard B, Marsault C, Agid Y, Le Bihan D (2000) Partially overlapping neural networks for real and imagined hand movements. *Cereb Cortex* 10:1093–1104.
- Gottfried JA, O'Doherty J, Dolan RJ (2002) Appetitive and aversive olfactory learning in humans studied using event-related functional magnetic resonance imaging. *J Neurosci* 22:10829–10837.
- Green PR, Gentle L, Peake TM, Scudamore RE, McGregor PK, Gilbert F, Dittich WH (1999) Conditioning pigeons to discriminate naturally lit insect specimens. *Behav Processes* 46:97–102.
- Hawkins RD, Clark GA, Kandel ER (2006) Operant conditioning of gill withdrawal in *Aplysia*. *J Neurosci* 26:2443–2448.
- Hefferline RF, Keenan B, Harford RA (1959) Escape and avoidance conditioning in human subjects without their observation of the response. *Science* 130:1338–1339.
- Keil A, Muller MM, Gruber T, Wienbruch C, Elbert T (2001) Human large-scale oscillatory brain activity during an operant shaping procedure. *Cogn Brain Res* 12:397–407.
- Leonardo M, Fieldman J, Sadato N, Campbell G, Ibanez V, Cohen L, Deiber MP, Jezzard P, Pons T, Turner R, LeBihan D, Hallett M (1995) A functional magnetic resonance imaging study of cortical regions associated with motor task execution and motor ideation in humans. *Hum Brain Mapp* 3:83–92.
- Ljungberg T, Apicella P, Schultz W (1992) Responses of monkey dopamine neurons during learning of behavioral reactions. *J Neurophysiol* 67:145–163.
- Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA (2004) Cognitive control signals for neural prosthetics. *Science* 305:258–262.
- Neuper C, Scherer R, Reiner M, Pfurtscheller G (2005) Imagery of motor actions: differential effects of kinesthetic and visual-motor mode of imagery in single-trial EEG. *Cogn Brain Res* 25:668–677.
- Nyberg L, Eriksson J, Larsson A, Marklund P (2005) Learning by doing versus learning by thinking: an fMRI study of motor and mental training. *Neuropsychologia* 44:711–717.
- O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal

- difference models and reward-related learning in the human brain. *Neuron* 38:329–337.
- Posse S, Fitzgerald D, Gao KX, Habel U, Rosenberg D, Moore GJ, Schneider F (2003) Real-time fMRI of temporolimbic regions detects amygdala activation during single-trial self-induced sadness. *NeuroImage* 18:760–768.
- Samejima K, Ueda Y, Doya K, Kimura M (2005) Representation of action-specific reward values in the striatum. *Science* 310:1337–1340.
- Schwartz MS (1995) *Biofeedback—a practitioner's guide*. New York: Guilford.
- Skinner BF (1938) *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner BF (1953) *Science and human behavior*. New York: MacMillan.
- Small W (1901) An experimental study of the mental processes of the rat. *Am J Psychol* 12:206–239.
- Stephan KM, Fink GR, Passingham RE, Silbersweig D, Ceballosbaumann AO, Frith CD, Frackowiak RSJ (1995) Functional anatomy of the mental representation of upper extremity movements in healthy subjects. *J Neurophysiol* 73:373–386.
- Stinear CM, Byblow WD, Steyvers M, Levin O, Swinnen SP (2006) Kinesthetic, but not visual, motor imagery modulates corticomotor excitability. *Exp Brain Res* 168:157–164.
- Svartdal F (1995) When feedback contingencies and rules compete—testing a boundary-condition for verbal control of instrumental performance. *Learn Motiv* 26:221–238.
- Thorndike E (1911) *Animal intelligence*. New York: Macmillan.
- Weiskopf N, Veit R, Erb M, Mathiak K, Grodd W, Goebel R, Birbaumer N (2003) Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *NeuroImage* 19:577–586.
- Weiskopf N, Scharnowski F, Veit R, Goebel R, Birbaumer N, Mathiak K (2004) Self-regulation of local brain activity using real-time functional magnetic resonance imaging. *J Physiol (Paris)* 98:357–373.
- Yacubian J, Glascher J, Schroeder K, Sommer T, Braus DF, Buchel C (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J Neurosci* 26:9530–9537.
- Yoo S-S, Jolesz FA (2002) Functional MRI for neurofeedback: feasibility study on hand motor task. *NeuroReport* 13:1377–1381.
- Yoo SS, O'Leary HM, Fairney T, Chen NK, Panych LP, Park H, Jolesz FA (2006) Increasing cortical activity in auditory areas through neurofeedback functional magnetic resonance imaging. *NeuroReport* 17:1273–1278.
- Yousry TA, Schmid UD, Alkadhi H, Schmidt D, Peraud A, Buettner A, Winkler P (1997) Localization of the motor hand area to a knob on the precentral gyrus: a new landmark. *Brain* 120:141–157.