

Executed and Observed Movements Have Different Distributed Representations in Human aIPS

Ilan Dinstein,¹ Justin L. Gardner,^{1,2} Mehrdad Jazayeri,³ and David J. Heeger^{1,2}

¹Center for Neural Science and ²Department of Psychology, New York University, New York, New York 10003, and ³Physiology and Biophysics Department, University of Washington, Seattle, Washington 98195

How similar are the representations of executed and observed hand movements in the human brain? We used functional magnetic resonance imaging (fMRI) and multivariate pattern classification analysis to compare spatial distributions of cortical activity in response to several observed and executed movements. Subjects played the rock–paper–scissors game against a videotaped opponent, freely choosing their movement on each trial and observing the opponent's hand movement after a short delay. The identities of executed movements were correctly classified from fMRI responses in several areas of motor cortex, observed movements were classified from responses in visual cortex, and both observed and executed movements were classified from responses in either left or right anterior intraparietal sulcus (aIPS). We interpret above chance classification as evidence for reproducible, distributed patterns of cortical activity that were unique for execution and/or observation of each movement. Responses in aIPS enabled accurate classification of movement identity within each modality (visual or motor), but did not enable accurate classification across modalities (i.e., decoding observed movements from a classifier trained on executed movements and vice versa). These results support theories regarding the central role of aIPS in the perception and execution of movements. However, the spatial pattern of activity for a particular observed movement was distinctly different from that for the same movement when executed, suggesting that observed and executed movements are mostly represented by distinctly different subpopulations of neurons in aIPS.

Key words: motor control; classification; fMRI; movement selectivity; mirror neurons; mirror system; movement perception

Introduction

The anterior intraparietal sulcus (aIPS) is often described as an intermediate sensory-motor area of the brain that integrates visual and motor neural processes. It is thought to contain many types of multimodal neurons including mirror neurons, which are visuomotor neurons that respond both when executing a particular movement and when observing someone else performing the same movement (Gallese et al., 1996). These neurons have been proposed to play a central role in perception of observed movements, because different subsets of mirror neurons respond selectively to particular movements so as to form a unique neural representation for each movement (Rizzolatti and Craighero, 2004). Direct evidence for the existence of movement-selective mirror neurons in IPS has been shown only in macaque monkeys (Fogassi et al., 1998, 2005). A few fMRI studies in humans, however, have tested whether human cortical areas exhibit movement-selective responses; a critical feature of neural populations involved in the perception of movement. These studies gauged movement-selectivity by assessing fMRI adaptation (rep-

etition suppression) during repeated observation (Shmuelof and Zohary, 2005; Hamilton and Grafton, 2006; Dinstein et al., 2007; Chong et al., 2008) and execution of movements (Dinstein et al., 2007; Chong et al., 2008). It is not known from these studies whether executed and observed movements give rise to distributed movement-specific patterns of activity in human aIPS, let alone whether these patterns are similar across execution and observation of the same movement, as may be expected for responses dominated by mirror neuron activity.

Here, we used multivariate pattern classification to assess whether aIPS and several other cortical areas responded with reproducible movement-selective response patterns during observation and/or execution of particular hand movements. fMRI is incapable of measuring the response selectivity of single neurons, because each fMRI voxel reflects the pooled responses of many neurons, which may have different selectivities. Nonetheless, assuming that there are neurons selective for particular observed and executed movements, it may be possible to associate each movement with a unique spatial pattern of fMRI responses. The distribution of movement-selective neurons within each fMRI voxel may be uneven such that the neural activity within particular voxels would be consistently larger (biased) for one movement compared with the others. This would generate distinct spatial patterns of fMRI responses associated with different movements, enabling decoding of movement identity from the distributed pattern of responses across voxels. This approach has been used successfully to characterize cortical activity selective

Received July 27, 2008; revised Sept. 16, 2008; accepted Sept. 18, 2008.

This work was supported by National Institutes of Health Grants R01-MH69880 (D.J.H.) and F31-MH080457 (I.D.), the Weizmann–New York University Demonstration Fund in Neuroscience (I.D.), a Burroughs Wellcome Fund Career Award in the Biomedical Sciences (J.L.G.), and a Helen Hay Whitney Foundation fellowship (M.J.).

Correspondence should be addressed to Ilan Dinstein, Center for Neural Science, New York University, 4 Washington Place, New York, NY 10003. E-mail: ilan@cns.nyu.edu.

DOI:10.1523/JNEUROSCI.3585-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/2811231-09\$15.00/0

for object categories (Haxby et al., 2001; Williams et al., 2007), grating orientations (Kamitani and Tong, 2004, 2005; Haynes and Rees, 2005), and visual motion directions (Kamitani and Tong, 2006) in different visual cortical areas. This approach, however, has not been used in the motor and somatosensory domains, nor has it been used to assess response selectivity for particular movements whether observed or executed.

To apply this approach to movement-selectivity, subjects in our experiments played the rock–paper–scissors (RPS) game, freely choosing which game movement to execute on each trial and, after a short delay, observing the movement of a videotaped opponent. Following Kamitani and Tong (2005), we performed a multivariate pattern classification analysis in aIPS and several motor and/or visual cortical areas that responded strongly during movement execution and/or observation. Validating the classification approach for motor as well as sensory responses, the identities of executed movements were correctly classified from fMRI responses in several areas of motor cortex and observed movements were classified from responses in visual cortex. Both observed and executed movements were classified from responses in left and right aIPS, highlighting this area's role in the perception and execution of movement. In aIPS, movement-specific response patterns that enabled accurate within modality decoding did not enable accurate decoding across modalities (i.e., decoding observed movement identity from a classifier trained on responses to executed movements and vice versa). This dissociation between the representations of observed and executed movements suggests that observed and executed movements are represented by subpopulations of neurons that are distributed differently in aIPS.

Materials and Methods

Subjects. Five healthy subjects (2 males) between the ages of 25 and 33 participated in two scanning sessions, which included a high-resolution anatomical scan and 10 runs of the RPS experiment. All subjects had normal or corrected-to-normal vision, provided written informed consent, and were paid for their participation in the study. The Tel-Aviv Sourasky Medical Center in Israel and the University Committee on Activities Involving Human Subjects at New York University approved the experimental procedures.

Visual stimuli and motor response. Stimuli were presented via an LCD projector and custom optics onto a rear-projection screen in the bore of the MRI scanner. Subjects were supine and viewed the projected stimuli through an angled mirror, which also prevented them from seeing their own hands. Subjects executed hand movements above their pelvis where they rested their hand between movements. These movements were videotaped using a video camera installed in the MRI room.

RPS experiment. The RPS experiment was based on the popular game rock–paper–scissors (http://en.wikipedia.org/wiki/Rock_Paper_Scissors). Subjects played the game against a videotaped opponent whose arm and hand were visible in the frame. Movement execution and observation alternated such that subjects executed one of the three movements and then observed the opponent's movement several seconds later (Fig. 1). During movement execution trials subjects were prompted with the word "GO" (presented for 1.5 s) to execute their freely chosen movement. Subjects were asked to raise their hand 10–15 cm, perform their movement, and return to relaxation by the time the "GO" disappeared from the screen. Observation trials consisted of the virtual opponent's response, presented in a 1.5 s video clip. Execution and observation trials were each followed by a blank intertrial interval of 3–6 s. Using randomized intertrial intervals enabled us to separate fMRI responses of movement observation and movement execution. Each experimental run/game included 36 execution trials and 36 observation trials, which lasted an average of 6 s each. We generated a different sequence of opponent movements for each of the games by randomly ordering video clips of the three movements. Subjects played a total of 10 games in two separate scanning sessions and were videotaped throughout the experiment

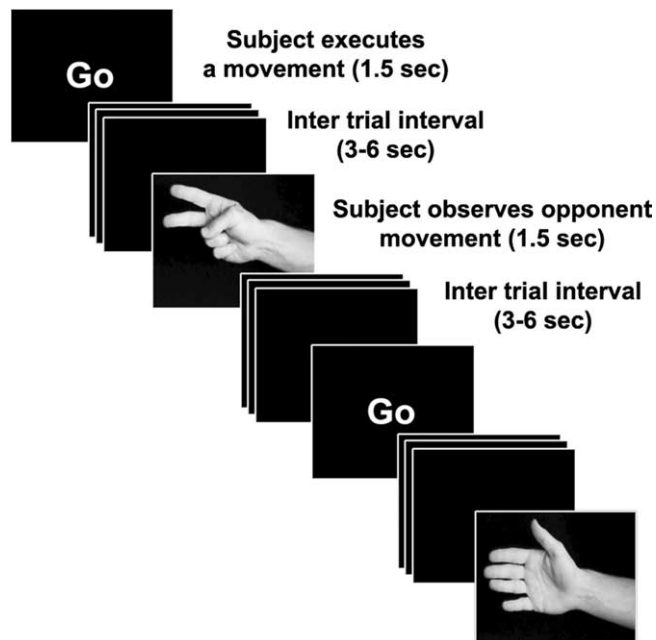


Figure 1. Experimental design. Subjects were prompted by a visual stimulus (the word "Go") to execute their game movement within a 1.5 s window. Execution was followed by 3–6 s of rest/blank after which subjects observed the opponent's movement. After another interval of 3–6 s, subjects were again prompted to execute a movement. Each game/run consisted of 36 execution and 36 observation trials with a total length of 7.5 min. Each subject participated in two scanning sessions performed on separate days and played five games in each session.

to extract their movement choices. Subjects were told in advance that they would receive extra monetary compensation if they won more trials than they lost over each scanning session.

Because subjects freely chose their movement on every trial, the overall number of rocks, papers, and scissors movements across the 10 games varied from subject to subject (rock = 118 ± 4.6 , paper = 113 ± 2.7 , scissors = 129 ± 4.9 , mean and SD across the 5 subjects).

fMRI acquisition. Functional and anatomical images of the brain were acquired with a Siemens 3T Trio MRI scanner equipped with a birdcage head coil used for RF transmit and receive. Blood oxygenation level-dependent (BOLD) contrast was obtained using a $T2^*$ sensitive echo planar imaging (EPI) pulse sequence: repetition time (TR) = 1500 ms, echo time (TE) = 30 ms, flip angle = 75° , 24 slices, $3 \times 3 \times 3$ mm voxels, field of view (FOV) = 192 mm. High-resolution anatomical volumes were acquired with a T1 weighted 3D-MPRAGE pulse sequence ($1 \times 1 \times 1$ mm).

Preprocessing, segmentation, and flattening. fMRI data were processed using standard tools available in the *Brain Voyager* software package (R. Goebel, Brain Innovation, Maastricht, The Netherlands). The first two images of each functional scan were discarded. Functional scans were then subjected to 3D motion correction, which involved resampling with trilinear interpolation, and temporal high-pass filtering with a cutoff frequency of six cycles per scan. Functional images were aligned with the high-resolution anatomical volume using trilinear interpolation. Anatomical and functional images were transformed to the Talairach coordinate system (Talairach and Tournoux, 1988), again using trilinear interpolation. This involved defining eight anatomical landmarks in each of the subjects, which included the anterior, posterior, dorsal, ventral, right, and left borders of the brain as well as the anterior and posterior commissures. The relative locations of these landmarks were aligned to fit the Talairach template. Other than the trilinear interpolation during these transformations, no additional spatial smoothing was performed. The cortical surface was reconstructed from high-resolution anatomical images, separately for each subject. The procedure included segmenting the gray and white matter, inflating the gray matter, cutting along several medial locations including the calcarine sulcus, unfolding the cortical

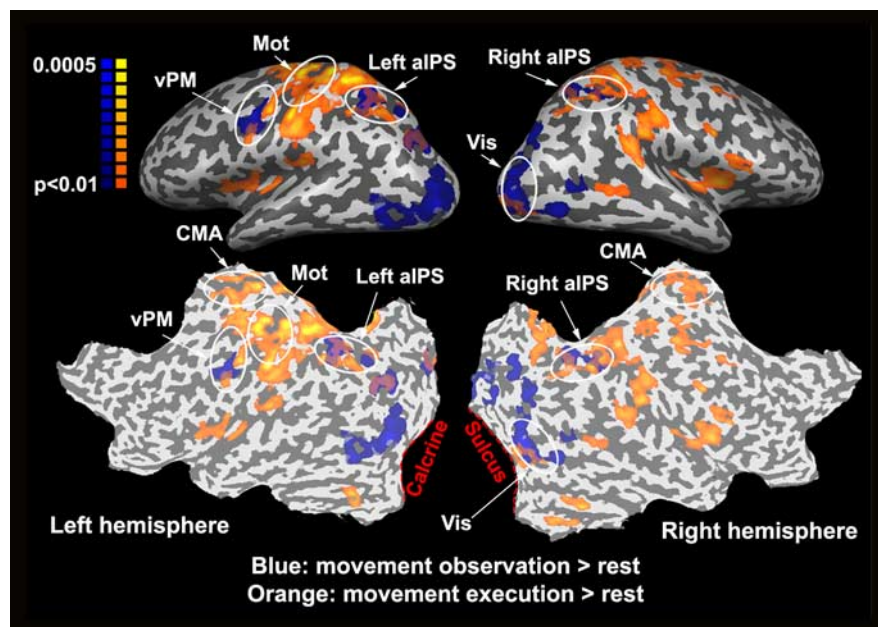


Figure 2. ROIs. Orange, Cortical areas that exhibited larger responses during movement execution trials than rest. Blue, Cortical areas that exhibited larger responses during movement observation trials than rest. Results calculated across all subjects are displayed on inflated (top) and flattened (bottom) hemispheres of one subject. We defined several ROIs for each of the subjects individually: bilateral cingulate motor area (CMA), left ventral premotor cortex (vPM), left primary motor and somatosensory cortex (Mot), left and right anterior intraparietal sulcus (aIPS), and visual cortex (Vis). The general location of the ROIs is outlined in white ellipses (the actual ROIs were defined individually for each subject).

Table 1. Mean ROI location and size across subjects

ROI name	Talairach coordinates			ROI size
	x	y	z	
Mot	−36 (1.5)	−26 (3.9)	51 (2.4)	678 (152)
Vis	21 (2.3)	−89 (2.3)	2 (0.6)	613 (78)
Left aIPS	−32 (4)	−53 (4.4)	45 (2.5)	496 (152)
Right aIPS	30 (1.6)	−52 (3.2)	48 (3.6)	484 (106)
vPM	−46 (3.8)	−5 (3.8)	36 (3.2)	507 (74)
CMA	−1 (2)	−14 (4.3)	50 (1)	840 (200)

Talairach coordinates (mean x, y, and z center of mass) are listed for each ROI along with the SD in parentheses. Mean ROI sizes (number of functional voxels) are listed along with SD in parentheses.

surface, and flattening. Inflated and flattened cortical maps were used to visualize cortical responses during the games (Fig. 2) and to identify regions of interest (ROIs).

ROI definition. Statistical parameter mapping (SPM) analysis (Friston et al., 1994) was used to identify brain areas that were active during movement execution and/or observation. In short, we constructed a general linear model (GLM) for the underlying neural responses to executed and observed movement trials. The regression matrix contained a row for each time point in the experiment where neural activity was modeled as either “on” = 1 or “off” = 0. The matrix also contained two columns, one for observed movement trials and the other for executed movement trials. The expected neural activity model (each column of the matrix) was convolved with a canonical hemodynamic response function (HRF) to create a model of the expected fMRI responses (Boynton et al., 1996). We used linear regression to estimate response amplitudes (β values) for each voxel in each of the two conditions, solving an equation of the form $y = Ax$, where the vector y was the measured time course during one scan, the vector x contained the estimated response amplitudes, and A was the regression matrix described above. Data were combined across subjects for visualization purposes (Fig. 2) using a random-effects analysis; response amplitudes were computed separately for each subject and then a paired t test (4 degrees of freedom) was used to determine significant voxel-by-voxel response differences across subjects (Friston et al., 1999).

Six ROIs were selected, for each subject individually, by overlaying the

statistical parameter map from each subject on their high-resolution anatomical scan and choosing all active voxels within a radius of 15 mm around particular anatomical landmarks. A false discovery rate (FDR) of 0.01 was used to threshold the statistical parameter maps. FDR is a method of correcting for multiple comparisons by controlling for the expected proportion of false positives among suprathreshold voxels (Genovese et al., 2002) rather than for the rate of false positives among all voxels as done by the stricter Bonferroni method.

Visual cortex (Vis) was defined by selecting voxels around the right occipital pole that responded more during movement observation than rest. Primary motor and somatosensory cortex (Mot) was defined by selecting voxels around the left “hand knob” landmark in the central sulcus that responded more during movement execution than rest. The cingulate motor area (CMA) was defined by selecting voxels bilaterally around the middle of the cingulate sulcus (adjacent to the medial end of the central sulcus) that responded more during movement execution than rest. Left and right aIPS were defined by selecting voxels around the middle of the anterior arm of the IPS that responded more during both movement execution and observation than rest (using a conjunction analysis of observation > rest and execution > rest). Ventral premotor cortex (vPM) was defined by selecting voxels in the left hemisphere around the junction of the inferior frontal sulcus and precentral sulcus that responded more during both movement execution and observation than rest (same as aIPS). See Table 1 for details about ROI coordinates and sizes.

HRF estimation. A deconvolution analysis was used to estimate a subject-specific hemodynamic response function (HRF) for each ROI, separately for observation trials and execution trials. Responses were averaged across all voxels in each ROI. A GLM was then used to estimate the HRFs. The GLM matrix contained 26 columns: 13 columns to estimate the observation HRF and another 13 columns to estimate the execution HRF. The HRFs were estimated across all movement trials regardless of movement identity so as not to bias the shape or amplitude of the HRF to one movement type. These subject-specific HRFs were used for the multivariate pattern classification analysis (as described next), although similar results were obtained using a canonical HRF (Boynton et al., 1996) instead of the subject-specific HRFs (compare Fig. 3 and supplemental Fig. 4, available at www.jneurosci.org as supplemental material).

Multivariate pattern classification analysis. fMRI response amplitudes were estimated, using linear regression, separately for each voxel in each ROI, and separately for each observation and execution trial. This regression matrix contained a row for each time point and a separate column for each execution trial and each observation trial (a total of $36 \times 2 = 72$ columns per game). Each column, therefore, had only a single time point of expected neural activity, which was modeled as “on” = 1, whereas all other time points were modeled as “off” = 0. Each column was convolved with the appropriate subject-specific HRF (or, as noted above, by the canonical HRF for the results reported in supplemental Fig. 4, available at www.jneurosci.org as supplemental material).

Following Kamitani and Tong (2005), the spatial patterns of fMRI responses were used to calculate three binary discrimination functions, one for each pair of movements (rock–paper, rock–scissors, and paper–scissors). The distributed response pattern of each movement trial can be

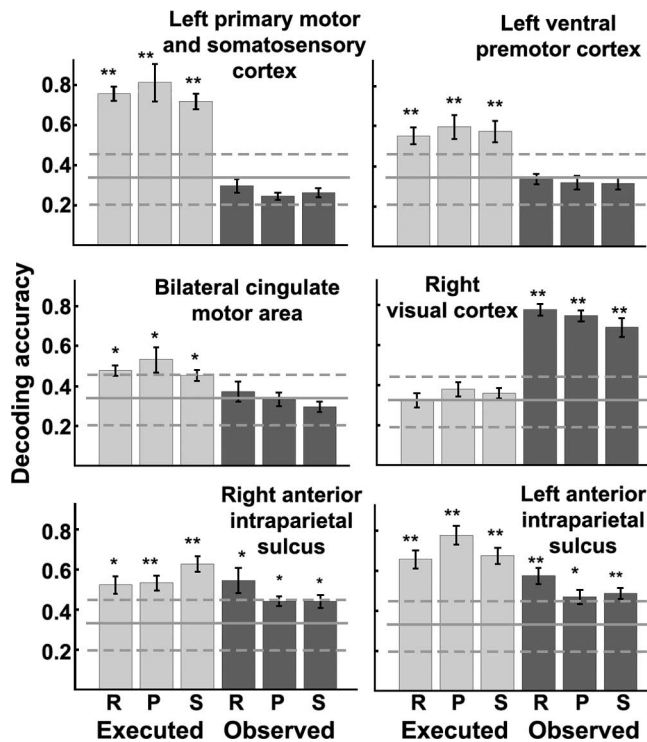


Figure 3. Classifier decoding accuracy. Light gray, Decoding accuracy for executed movements. Dark gray, Decoding accuracy for observed movements. Error bars, SEM across subjects. Asterisks, Statistical significance assessed with one-tailed *t* tests across subjects ($*p < 0.05$; $**p < 0.005$). Solid lines, Chance classification accuracy level (0.33). Dashed lines, Fifth and ninety-fifth percentiles for chance classification, estimated with randomization tests separately for each ROI (see Materials and Methods and supplemental Fig. 1, available at www.jneurosci.org as supplemental material).

described mathematically as a vector containing the response amplitudes of all the voxels in an ROI. The response to each movement trial was, therefore, a data point in a multidimensional space where the dimensionality was the number of voxels in the ROI. The discrimination function was a hyperplane in this space that best separated the data points of the two groups. We used a linear discriminant, similar to Fisher linear discriminant (FLD) analysis, to compute the hyperplane which optimally separated the responses if the response variability in each voxel (across trials) was normally distributed, and statistically independent across voxels. We note that with more voxels than trials, inversion of the full covariance matrix necessary for FLD would be unstable because the covariance matrix would be underdetermined. We therefore disregarded covariances between voxels by setting all off-diagonal elements in the covariance matrix to 0. We then calculated the output of the classifier using the vector orthogonal to the discrimination hyperplane with the following equation:

$$g(\mathbf{x}) = \mathbf{w}_0 + \sum_{i=1}^n \mathbf{x}_i \mathbf{w}_i,$$

where i indexes over voxels, \mathbf{x}_i is a vector of fMRI response amplitudes, \mathbf{w}_i is a vector of the computed weights that is orthogonal to the discrimination hyperplane, and \mathbf{w}_0 is a scalar such that $g(\mathbf{x}) < 0$ for trials from one movement category and $g(\mathbf{x}) > 0$ for trials from the other movement category.

We used “leave one out” validation to test the accuracy of the classification. We calculated (or “trained”) the three binary classifiers while leaving out one of the exemplars and tested whether the left out exemplar was accurately assigned to the proper movement category (“decoding”). Before calculating the classification functions we averaged groups of 8 trials from the same movement category to improve the signal-to-noise

ratio of the spatial patterns of responses. This decreased the number of exemplars we had for each movement type, but also decreased the amount of noise inherent in each exemplar. The training and decoding process was repeated for each of the available exemplars and the whole process was also repeated 100 times while averaging the movements in different randomly selected groups of 8 trials. An exemplar was classified as a particular movement if it successfully passed two binary discriminations (e.g., a rock exemplar was correctly assigned to the rock category only if it passed both the rock–paper and the rock–scissors discrimination functions). Exemplars that were misclassified by either or both of the two relevant discrimination functions were considered errors. This yielded an accuracy of 33% when testing the classifier on randomly generated numbers (supplemental Fig. 5, available at www.jneurosci.org as supplemental material). Accuracy was computed separately for each subject and then averaged across subjects.

To assess statistical significance of decoding accuracy, we performed two statistical tests. First we used a *t* test (4 degrees of freedom) in each of the ROIs to test whether decoding accuracy of each trial type was significantly higher than chance. This *t* test had the advantage of treating intersubject differences as a random factor. The *t* test, however, assumed that the chance distribution of decoding accuracies was normally distributed, which may not be the case. We, therefore, also performed a randomization test that was identical to the classification analysis described above except that we randomly shuffled the identities of the observed/ executed movements before training the classifier. We ran this analysis 300 times for each subject separately, reshuffling the identity of the movements each time. This yielded a distribution of 1500 decoding accuracies expected by chance in each of the ROIs, with a mean of ~ 0.33 and 5th through 95th percentiles of ~ 0.2 to ~ 0.46 when collapsed across subjects (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). To assess statistical significance, we tested whether decoding accuracy was above the 95th percentile.

We obtained almost identical results using a 3-way maximum likelihood classifier. The Matlab (Mathworks) “classify” function with the “diaglinear” option estimated the response variances and then classified the left out exemplar. This analysis, like that described above, assumed that each voxel’s responses were normally distributed and that voxels were statistically independent from one another. A “leave one out” procedure was again used for the classification, we again averaged groups of 8 trials from the same movement category to improve the signal-to-noise ratio of the response patterns, and we performed both a randomization statistical test and a *t* test to determine if the classifier performance was significantly above chance.

To determine whether decoding accuracy was affected by transforming to Talairach space, we reanalyzed data from two of the subjects in aIPS ROIs that were defined in the coordinates of the original EPI slices. The resulting decoding accuracies were almost identical to those calculated in Talairach space. Hence, the spatial blurring inherent in the Talairach transformation had negligible effects on the results (although the data were still resampled with trilinear interpolation by the motion compensation procedure).

Pattern correlation analysis. In a complementary analysis, we also assessed the similarity of spatial response patterns across different movement types using correlation (following Haxby et al., 2001). We followed the same steps described for the classification analysis to estimate fMRI response amplitudes in each voxel of each ROI, but here we averaged across all trials of the same category from each game/scan, generating a total of 10 exemplars for each movement category (e.g., 10 executed scissors patterns) for each subject. We also computed the average response pattern across all movement types in each game for execution and observation trials separately, generating 10 execution and 10 observation averages. The average game responses were projected out of each category exemplar (e.g., the projection of the executed scissors exemplar onto the execution mean was subtracted from the exemplar). This step ensured that there was no correlation between each exemplar and the average of all three. Finally, we computed correlation coefficients among exemplars of the same category (e.g., executed rocks) and exemplars of different categories (e.g., executed rocks and executed papers), averaging the category specific correlations computed from all game pairs.

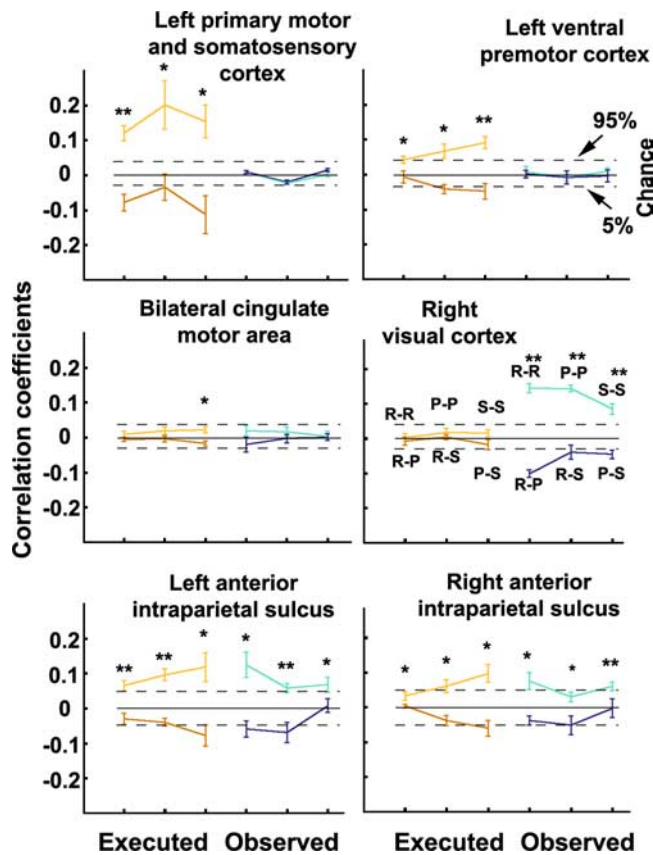


Figure 4. Correlation of response patterns across games. Light orange, Correlations for same executed movement type: rock–rock (R–R), paper–paper (P–P), and scissors–scissors (S–S). Dark orange, Correlations for different executed movement types: rock–paper (R–P), rock–scissors (R–S), and paper–scissors (P–S). Light blue, Correlations for same observed movement type. Dark blue, Correlations for different observed movement types. Error bars, SEM across subjects. Asterisks, Statistical significance assessed with one-tailed *t* tests across subjects (**p* < 0.05; ***p* < 0.005). Solid horizontal lines, Chance correlation level (0). Dashed horizontal lines, Fifth and ninety-fifth percentiles for chance correlations, estimated with randomization tests separately for each ROI (see Materials and Methods and supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

Here, too, *t* tests and randomization tests were used to determine whether correlation values were significantly higher than zero in each of the ROIs. For the randomization tests, we performed a correlation analysis identical to the one described above except that the identities of the observed/executed movements were shuffled. We ran this analysis 6000 times in each subject, reshuffling the identity of the movements each time. This yielded a distribution of the correlation accuracies expected by chance, with a mean of 0 and a 95th percentile of ~0.045 (supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

Voxel bias maps. The weights calculated for the three binary discrimination functions, one for each pair of movements (rock–paper, rock–scissors, and paper–scissors), are informative in showing which of the voxels responded consistently more strongly during the execution or observation of one of the movements compared with the other two. This enabled us to create maps of the voxel “biases” for each of the subjects and each of the ROIs. We assigned each voxel with a value of 0, 1, 2, or 3 corresponding to no consistent bias, rock-bias, paper-bias, or scissors-bias. For example, we assigned a voxel with the rock-bias value only if it had rock biased weights in both of the rock–paper and rock–scissors discrimination functions.

Results

Subjects played 10 RPS games in two separate scanning sessions performed on different days (five games in each session). The subjects were visually prompted to execute their freely chosen

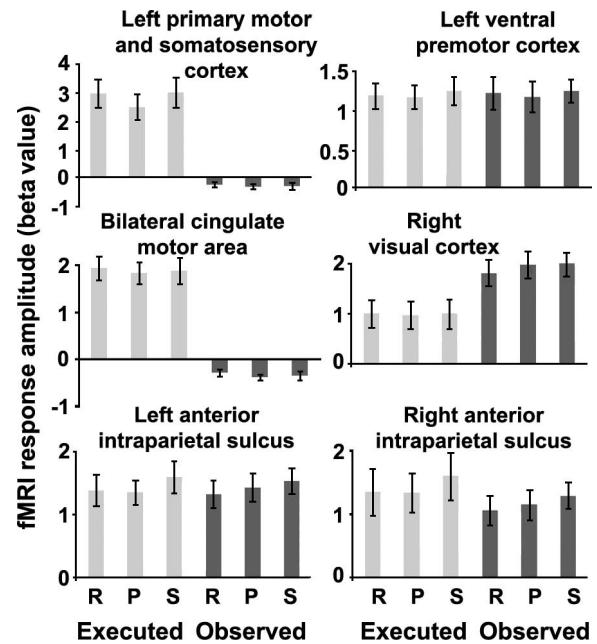


Figure 5. fMRI responses to the three RPS movements, averaged across voxels in each ROI and across subjects. Each of the three movements evoked similar (statistically indistinguishable) mean fMRI response amplitudes. R, Rock; P, paper; S, scissors. Error bars, SEM.

movements on movement execution trials, which were followed by movement observation trials that contained the virtual opponent’s movements. The RPS game was used to keep subjects engaged while executing and observing many trials of only three movement types.

A whole brain SPM analysis showed that the typical visual and motor cortical areas responded robustly during trials where movements were executed or observed (Fig. 2, orange and blue, respectively), regardless of movement identity. We focused our analyses on commonly described candidate human “mirror system” areas, which responded during both movement observation and execution: left and right anterior intraparietal sulcus (aIPS) as well as left ventral premotor cortex (vPM). We hypothesized that because analogous areas of monkey cortex have been shown to contain single neurons selective for particular executed and observed movements (Gallese et al., 1996; Fogassi et al., 1998, 2005; Umiltà et al., 2001), they might enable accurate decoding of both observed and executed movement identity. Three control ROIs were also identified: left primary motor and somatosensory cortex (Mot), bilateral cingulate motor area (CMA), and right visual cortex (Vis). We expected the motor ROIs to exhibit accurate decoding of executed, but not observed, movements and the opposite from the visual ROI. These ROIs were identified in each subject separately and consisted of ~600 functional voxels that were active during either movement execution, movement observation, or both (see Materials and Methods and Table 1 for details).

A multivoxel pattern classifier was used to assess whether the distributed patterns of activity in each ROI were unique and reproducible for execution and/or observation of each particular movement. The statistical significance of the results was assessed in two ways. First, we performed a *t* test across subjects to determine whether the decoding accuracy was significantly greater than chance (33%). Second, we performed a randomization test in which movement identities were randomly shuffled to generate a distribution of decoding accuracies expected by chance,

separately for each ROI (see Materials and Methods for details). Decoding accuracies, according to the null hypothesis (randomized movement identities), had a mean of 0.33 and a 95th percentile ranging between 0.45 and 0.47 (supplemental Fig. 2, available at www.jneurosci.org as supplemental material), as indicated by the dashed horizontal lines in Figure 3.

The spatial response patterns in each ROI enabled above-chance decoding of observed movement identity, executed movement identity, or both (Fig. 3). Decoding accuracy for executed movements (collapsed across the three movement types) was above chance in Mot (76%), left vPM (57%), and CMA (49%), whereas accuracy for observed movements in the same areas was indistinguishable from chance. Conversely, decoding accuracy only for observed movements was above chance in Vis (72%). Decoding accuracy in left aIPS was significantly above chance for all three executed and all three observed movements. In right aIPS, accuracy was above chance for all three executed movements and for one of the three observed movements. The other two observed movements were significantly above chance as assessed by the *t* test, but just below the 95th percentile of the randomization analysis. Note that decoding accuracy was computed across all 10 scans acquired during two different sessions; movement-specific fMRI response patterns were, therefore, stable enough across days to support classification.

A complementary correlation analysis revealed similar results. We computed a single spatial response pattern for each of the six trial types per game/run, and correlated different and same movement types across pairs of games (Fig. 4). Statistical significance of the correlations was again assessed using both a *t* test (across subjects) and a randomization test. Chance correlations, according to the null hypothesis (randomized movement identities), had a mean of 0 and a 95th percentile ranging between 0.035 and 0.056 (supplemental Fig. 2, available at www.jneurosci.org as supplemental material), as indicated by the dashed horizontal lines in Figure 4. Response patterns in Mot and in vPM, evoked by the same executed movements, were significantly correlated across games, whereas response patterns evoked by different executed movements and by observed movements were not. In contrast with the classification results, bilateral CMA did not exhibit any significant correlations. In visual cortex (Vis), same executed movements were not correlated, but same observed movements were. Left and right aIPS exhibited significant correlations both for same executed movements and same observed movements. The correlation values, based on responses computed independently from the 10 separate games/runs, could be dramatically increased by averaging across multiple games although the signal-to-noise ratio remained similar (supplemental Fig. 3, available at www.jneurosci.org as supplemental material). The correlation analysis is affected equally by all participating voxels, as opposed to the classification analysis which assigns a weight to each voxel such that the contribution of “noisy” voxels is minimized. The results of the correlation analysis, therefore, may be weaker than those of the classification analysis depending on the proportion of “noisy” voxels included in each ROI.

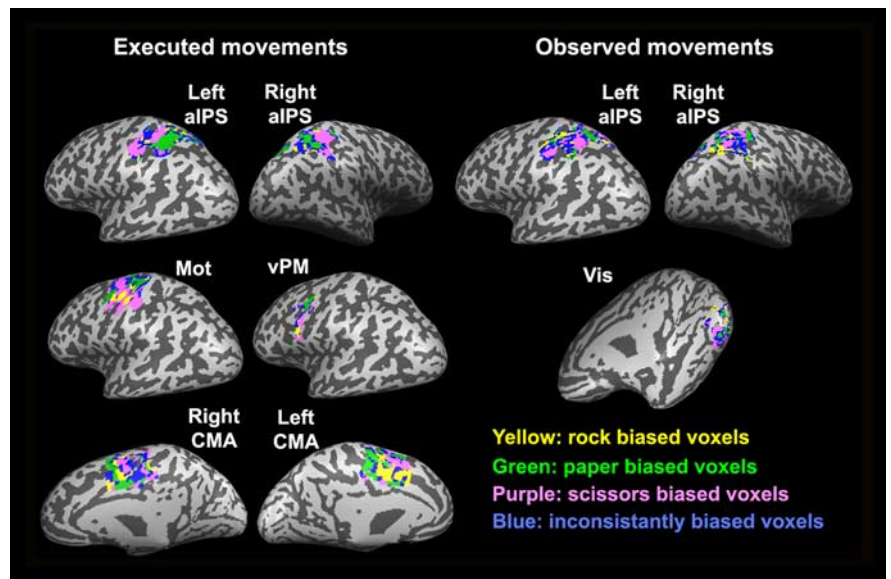


Figure 6. Classification weights from a single typical subject. Left, Executed movements. Right, Observed movements. Green, Voxels with a bias for the rock movement. Yellow, Voxels with a bias for the paper movement. Purple, Voxels with a bias for the scissors movement. Blue, Voxels with no consistent bias to any of the movements.

Accurate classification was attributable to information available in the spatially distributed pattern of responses, rather than the overall response amplitudes throughout each ROI. Averaging responses across all voxels in any of the ROIs showed no difference in the response amplitudes for the three movements, whether executed or observed (Fig. 5). Furthermore, plots of voxel classification weights (illustrating their bias for a particular observed or executed movement) in each of the ROIs (Fig. 6) showed that movement biased voxels were intermingled forming a distributed response pattern in each ROI. Figure 6 shows an example of the classification weights from a single subject; the weights were similarly intermingled in all subjects although each had a different idiosyncratic distribution of weights.

We compared the distributed response patterns of observed and executed movements in left and right aIPS, which responded with similar overall amplitude during execution and observation trials (Fig. 5). The classifier was trained using trials from one modality and decoding accuracy was tested using trials from the other modality. Decoding accuracy was indistinguishable from chance in both cases (Fig. 7, left column). A complementary correlation analysis showed that correlation coefficients of response patterns to the same movement when observed and executed were indistinguishable from zero (Fig. 7, right column). As noted above, classification accuracy and correlation coefficients in aIPS were significantly above chance within each modality. Area aIPS, therefore, exhibited a clear dissociation of within- versus between-modality classification and correlation; consistent and unique spatial patterns of activity were exhibited by each movement category, but these distributed response patterns were different for observation and execution.

To further compare the distributed representations of observed and executed movements in aIPS, we also performed a six-way classification analysis (Fig. 8), in which each movement type was compared with all other movement types regardless of modality (e.g., executed scissors was compared with executed rock and executed paper as well as with observed rock, paper, and scissors). In this analysis, a trial was assigned to a particular category only if it successfully passed all five relevant pairwise discrimination functions.

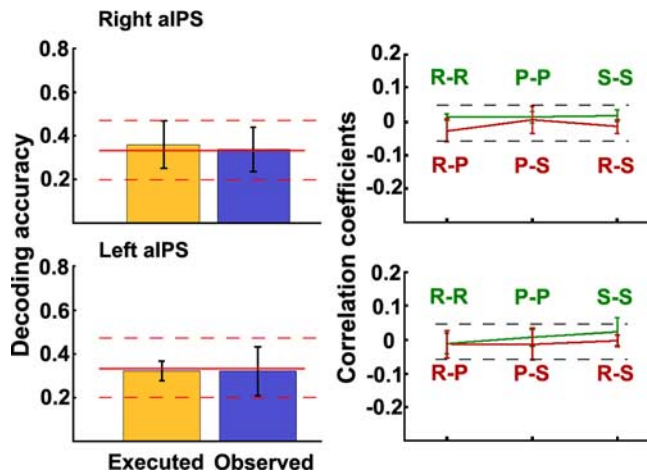


Figure 7. Cross-modal decoding accuracy and cross-modal correlation of response patterns. Top row, Right aIPS. Bottom row, Left aIPS. Left column, Decoding accuracy. Orange, Average decoding accuracy for executed movement identity when training classifier on observed movements. Blue, Average decoding accuracy for observed movement identity when training classifier on executed movements. Right column, Correlation of response patterns. Green, Correlations in response patterns for observed and executed rock (R–R), paper (P–P), and scissors (S–S). Red, Correlations for different observed and executed movement types: rock–paper (R–P), rock–scissors (R–S), paper–scissors (P–S). Error bars, SEM across subjects. Dashed horizontal lines, Fifth and ninety-fifth percentiles for chance decoding accuracy and chance correlation coefficients, estimated with randomization tests separately for each ROI (see Materials and Methods).

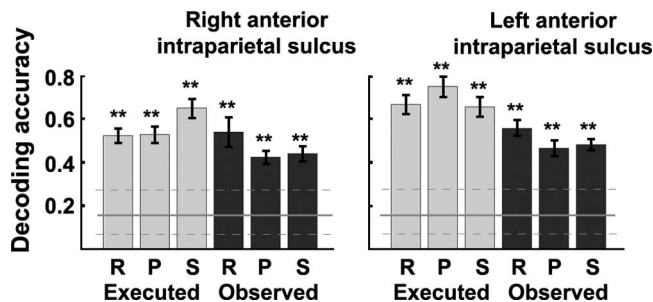


Figure 8. Six-way classifier decoding accuracy (same format as Fig. 3). Light gray, Decoding accuracy for executed movements. Dark gray, Decoding accuracy for observed movements. Error bars, SEM across subjects. Asterisks, Statistical significance assessed with one-tailed *t* tests across subjects (***p* < 0.005). Solid lines, Chance classification accuracy level (0.165). Dashed lines, Fifth and ninety-fifth percentiles for chance classification, estimated with randomization tests separately for each ROI (see Materials and Methods).

Chance accuracy was 16.5% as computed with an equivalent randomization analysis where movement identities were randomly shuffled among the six categories. Despite the addition of three “cross-modal” discrimination functions, decoding accuracies, for both observed and executed movement identities, were almost identical to those obtained with the three-way classifier (Fig. 3). This means that executed movements were rarely misclassified as observed movements and vice versa. Rather, errors, when they occurred, were attributable to misclassifying movement identity within the visual or motor modality (>99.9% of errors) and not across the modalities (<0.1%). This result confirms that distinctly different fMRI response patterns were generated by executed and observed movements in aIPS.

Discussion

Overlapping visual and motor movement selectivity in aIPS

In this study we used classification and correlation to compare fMRI response patterns associated with different observed and/or

executed hand movements. Left and right aIPS exhibited reproducible response patterns during both execution and observation of the RPS movements (Figs. 3, 4). Although both areas responded with indistinguishable overall response amplitudes to all three movements (Fig. 5), consistent differences in the voxel-by-voxel spatial pattern of responses to each movement type enabled us to decode the identities of executed movements (70 and 53%, left and right aIPS, respectively), and observed movements (51 and 46%). These results do not prove that there are single neurons selective for movement type in aIPS. Rather, we assume that such neurons are present based on monkey electrophysiology studies (Gallese et al., 1996; Fogassi et al., 1998, 2005; Umiltà et al., 2001) and suggest that the movement-specific activity patterns reported here were generated by large subpopulations of such neurons.

There was a clear dissociation in aIPS of within-modality (motor or visual) and between-modality (motor to visual or visual to motor) classification. Specifically, between-modality decoding accuracy of movement identity was indistinguishable from chance (Fig. 7), unlike the within-modality decoding. This combination of results suggests one of two possible interpretations, which are not mutually exclusive. The first is that the activity patterns in aIPS were generated by distinct subpopulations of movement-selective motor neurons (active during movement execution) and movement-selective visual neurons (active during movement observation), which were distributed differently throughout left and right aIPS. A second possibility is that the activity patterns were generated by a single population of visuo-motor neurons, but that these neurons exhibited different selectivities during observation and execution such that they consistently responded with one activity pattern to the observation of a particular movement and with a different activity pattern to the execution of that same movement.

In a previously published study, we reported an analogous dissociation in aIPS when using an fMRI adaptation (repetition suppression) protocol to identify cortical areas that adapted during repeated observation and execution of movements (Dinstein et al., 2007). fMRI adaptation studies rely on the assumption that strongly stimulated neurons decrease their firing rates (adapt/habituate) when stimulated repeatedly, resulting in smaller amplitude fMRI responses (Grill-Spector and Malach, 2001; Huk and Heeger, 2002; Boynton and Finney, 2003; Fang et al., 2005; Grill-Spector, 2006). In the adaptation study, area aIPS exhibited both visual and motor adaptation, but a lack of cross-modal (visual to motor or motor to visual) adaptation. Visual and motor movement-selective responses were, therefore, evident in aIPS, but there was no evidence for movement-selective responses across the two modalities, which might have been expected from a neural population with the same movement-selectivity across the visual and motor modalities (mirror neurons). fMRI classification provides a complementary technique for studying response selectivity, which relies on the assumption that neural populations with different selectivities are distributed unevenly such that unique fMRI response patterns can be associated with particular stimuli or tasks. Thus, converging evidence from the adaptation and classification techniques indicates that aIPS exhibits movement-selective responses during both observation and execution of movement, but that the representations of observed and executed movements are fundamentally different from one another.

Lateralization of movement execution and perception

Although we did not design our experiments to study the lateralization of movement perception and execution, a notable aspect of our results is that movement-selective cortical responses were evident bilaterally in aIPS when executing a movement with the right-hand and when observing someone else performing a movement with their right-hand. Previous studies have reported lateralized response differences between left and right aIPS during movement execution, observation, and imitation (Decety et al., 2002; Aziz-Zadeh et al., 2006). Yet these studies based their conclusions on small response amplitude differences between the entire left and right aIPS. These differences in response amplitudes might have different neural underpinnings that may or may not be related to movement perception or execution (e.g., differences in effort/arousal related to imitating with the dominant versus nondominant hand). Using the classification approach it was possible to decode movement identity from activity patterns of either hemisphere, revealing a clear bilateral representation of executed and observed movements.

Distributed movement-selective visual responses

The visual cortex ROI enabled accurate decoding of observed movement identity, but not of executed movement identity (Figs. 3, 4). We did not perform retinotopic mapping to identify visual cortical areas in our subjects, yet we assume that this ROI contained multiple visual areas surrounding the occipital pole including V1, V2, V3, and V4. The three movement video clips used in the game differed in many low-level visual features including the amount of visual motion and the distribution of local luminance and contrast levels. We, therefore, suggest that the reproducible response patterns evident in visual cortex were attributable to the consistent responses of large neural populations selective for these visual features, as previously reported when using low-level visual stimuli (Haynes and Rees, 2005; Kamitani and Tong, 2005, 2006) or static real-life images (Kay et al., 2008). Note that these visual areas were active not only during movement observation, but also during movement execution (Fig. 5) because subjects were prompted to execute their chosen movements by a visual cue (the word "Go"). However, as expected, this robust fMRI response did not enable accurate decoding of executed movement identities.

Spatially distributed motor control

Several motor system areas enabled accurate decoding of executed movement identity (Figs. 3, 4). Although it is not surprising that different populations of motor and somatosensory neurons respond selectively during the execution of particular movements, almost all of the evidence for such selectivity comes from electrophysiology studies in nonhuman primates (e.g., selectivity for movement direction, Georgopoulos et al., 1986). The reason for this is that the neural populations active during the execution of different movements are intermingled within these cortical motor areas such that the overall fMRI response amplitudes to different movements are indistinguishable (Fig. 5). Our ability to accurately distinguish among responses to different hand movements, whose execution required the recruitment of similar muscle groups in different combinations, validates the use of classification techniques for the study of response selectivity in the motor and somatosensory domains.

Comparing distributed patterns of cortical responses, as was done in the current study for observed versus executed movements, may enable various dissociations between the cortical representations of movement dynamics, kinematics, goals/inten-

tions, and symbolic value. For instance, if neural populations in area aIPS indeed code for different movement goals as has been proposed (Fogassi et al., 2005; Hamilton and Grafton, 2006), it may be possible to accurately decode movement goal across movements with the same goal and different kinematics using cortical response patterns from aIPS, but not from primary motor cortex. The ability to perform such experiments with human subjects, while simultaneously scanning the entire brain, offers many advantages for the study of motor control.

Movement selectivity, mirror neurons, and movement perception

Mirror neurons were discovered over 10 years ago in two cortical areas of the macaque monkey, F5 and IPL (Gallese et al., 1996; Fogassi et al., 1998, 2005; Umiltà et al., 2001), that are thought to be analogous to human vPM and aIPS, respectively. Mirror neurons were defined functionally as visuomotor neurons that responded selectively to a particular preferred movement (e.g., grasping, placing, or manipulating) whether the monkey executed it or observed someone else performing it. Several theories have proposed that mirror neurons act as a mapping mechanism between the observation of an action and its execution so that when someone observes a movement, they simulate performing that movement (Rizzolatti and Craighero, 2004). This simulation then allows them to assign their own associated intentions, goals, emotions, and social values to the person whom they observed. Note that for this mechanism to work, it is critical that the observed movement be mapped onto the particular neural circuits used to execute it; otherwise improper intentions will be inferred.

Area vPM is commonly considered a candidate human mirror system area and indeed this area responded with equal strength during both the observation and the execution of movements in our study (Figs. 2, 5). We, therefore, expected this area to exhibit accurate decoding for both the identities of executed and observed movements. Yet our analyses revealed accurate decoding only for the identities of executed movements and not for identities of observed movements (Figs. 3, 4). Because this is a negative result with several possible explanations (e.g., the distribution of visually selective responses in this area may be restricted to small subregions of vPM beyond our voxel size resolution) we will not discuss it further.

Had the responses in area aIPS during movement execution and observation been dominated by the activity of mirror neurons, one would have expected a similar distributed response pattern during the observation and execution of the same movement. This does not seem to be the case. Cortical activity in aIPS exhibited visual and motor movement-selectivity, enabling accurate decoding of both observed movement identity and executed movement identity. This suggests that neural populations in aIPS may indeed play a central role in the perception of movement as proposed by the theories above. The distinct spatial response patterns to the observation and execution of the same movement, however, show that these responses were not generated by a single population of movement-selective mirror neurons, but rather from distinct subpopulations of visual and motor neurons, or from "promiscuous" visuomotor neurons that responded selectively to one observed movement and to a different executed movement. Although there may very well be movement-selective mirror neurons in human aIPS, our results suggest that they make up only a minority of the neurons active during observation or execution of movement and do not dominate the fMRI responses of this area.

References

- Aziz-Zadeh L, Koski L, Zaidel E, Mazziotta J, Iacoboni M (2006) Lateralization of the human mirror neuron system. *J Neurosci* 26:2964–2970.
- Boynton GM, Finney EM (2003) Orientation-specific adaptation in human visual cortex. *J Neurosci* 23:8781–8787.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221.
- Chong TT-J, Cunnington R, Williams MA, Kanwisher N (2008) fMRI adaptation reveals mirror neurons in human inferior parietal cortex. *Curr Biol*, in press.
- Decety J, Chaminade T, Grèzes J, Meltzoff AN (2002) A PET exploration of the neural mechanisms involved in reciprocal imitation. *Neuroimage* 15:265–272.
- Dinstein I, Hasson U, Rubin N, Heeger DJ (2007) Brain areas selective for both observed and executed movements. *J Neurophysiol* 98:1415–1427.
- Fang F, Murray SO, Kersten D, He S (2005) Orientation-tuned fMRI adaptation in human visual cortex. *J Neurophysiol* 94:4188–4195.
- Fogassi L, Gallese V, Fadiga L, Rizzolatti G (1998) Neurons responding to the sight of goal-directed hand/arm actions in the parietal area PF (7b) of the macaque monkey. *Soc Neurosci Abstr* 24:257.5.
- Fogassi L, Ferrari PF, Gesierich B, Rozzi S, Chersi F, Rizzolatti G (2005) Parietal lobe: from action organization to intention understanding. *Science* 308:662–667.
- Friston KJ, Jezzard P, Turner R (1994) Analysis of functional MRI time-series. *Hum Brain Mapp* 1:153–171.
- Friston KJ, Holmes AP, Price CJ, Büchel C, Worsley KJ (1999) Multisubject fMRI studies and conjunction analyses. *Neuroimage* 10:385–396.
- Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. *Brain* 119:593–609.
- Genovese CR, Lazar NA, Nichols T (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15:870–878.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
- Grill-Spector K (2006) Selectivity of adaptation in single units: implications for fMRI experiments. *Neuron* 49:170–171.
- Grill-Spector K, Malach R (2001) fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol (Amst)* 107:293–321.
- Hamilton AF, Grafton ST (2006) Goal representation in human anterior intraparietal sulcus. *J Neurosci* 26:1133–1137.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
- Huk AC, Heeger DJ (2002) Pattern-motion responses in human visual cortex. *Nat Neurosci* 5:72–75.
- Kamitani Y, Tong F (2004) Pattern recognition of orientation-selective fMRI signals in the human visual cortex [Abstract]. *J Vis* 4:186.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
- Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
- Rizzolatti G, Craighero L (2004) The mirror-neuron system. *Annu Rev Neurosci* 27:169–192.
- Shmuelof L, Zohary E (2005) Dissociation between ventral and dorsal fMRI activation during object and action recognition. *Neuron* 47:457–470.
- Talairach J, Tournoux P (1988) Co-planar stereotaxic atlas of the human brain. New York: Thieme Medical.
- Umiltà MA, Kohler E, Gallese V, Fogassi L, Fadiga L, Keysers C, Rizzolatti G (2001) I know what you are doing. A neurophysiological study. *Neuron* 31:155–165.
- Williams MA, Dang S, Kanwisher NG (2007) Only some spatial patterns of fMRI response are read out in task performance. *Nat Neurosci* 10:685–686.