

Fundamentally Distributed Information Processing Integrates the Motor Network into the Mental Workspace during Mental Rotation

Alexander Schlegel, Dedeepya Konuthula, Prescott Alexander, Ethan Blackwood, and Peter U. Tse

Abstract

■ The manipulation of mental representations in the human brain appears to share similarities with the physical manipulation of real-world objects. In particular, some neuroimaging studies have found increased activity in motor regions during mental rotation, suggesting that mental and physical operations may involve overlapping neural populations. Does the motor network contribute information processing to mental rotation? If so, does it play a similar computational role in both mental and manual rotation, and how does it communicate with the wider network of areas involved in the mental workspace? Here we used multivariate methods and fMRI to study 24 participants

as they mentally rotated 3-D objects or manually rotated their hands in one of four directions. We find that information processing related to mental rotations is distributed widely among many cortical and subcortical regions, that the motor network becomes tightly integrated into a wider mental workspace network during mental rotation, and that motor network activity during mental rotation only partially resembles that involved in manual rotation. Additionally, these findings provide evidence that the mental workspace is organized as a distributed core network that dynamically recruits specialized subnetworks for specific tasks as needed. ■

INTRODUCTION

In a seminal experiment on the mental manipulation of visual imagery, Shepard and Metzler (1971) asked participants to mentally rotate visually presented 3-D objects to determine whether they matched other similar objects. Participants' RTs correlated tightly with the angle of rotation that would be necessary to align the two objects, suggesting that they had mentally rotated endogenous mental models of the objects in an analog fashion as if manually rotating a physical object through space. Subsequent behavioral research has explored other operations such as mental paper folding (Shepard & Feng, 1972), the generation and analysis of mental analog clocks (Paivio, 1978), and mental simulations of mechanical systems (Hegarty, 2004), a primary result being that volitional mental operations appear in many respects to resemble their corresponding physical operations. Other work has documented similar processes in domains such as mental time travel (Addis, Wong, & Schacter, 2007), creative synthesis of mental imagery (Finke & Slayton, 1988), and visuospatial reasoning (Blazhenkova & Kozhevnikov, 2010). Thus, the human brain appears to support a mental space analogous to the physical world in which mental models can be constructed, manipulated, and tested in a flexible and potentially analog manner.

Such abilities have been studied using several overlapping psychological constructs including working memory (Baddeley, 2003), mental imagery (Tong, 2013; Kosslyn, Behrmann, & Jeannerod, 1995), visuospatial ability (Uttal et al., 2013), mental models (Johnson-Laird, 2010; Hegarty, 2004), analogical reasoning (Bassok, Dunbar, & Holyoak, 2012), and the mental workspace (Logie, 2003). Generally, these terms denote the ability to work volitionally and flexibly with mental representations. Indeed, recent empirical work suggests that at least some of these constructs may actually refer to the same underlying neural mechanisms (Tong, 2013). Following Logie (2003), we will refer to the mental space in which these flexible cognitive processes occur as the mental workspace. According to this view, abilities such as working memory and mental rotation are both mental workspace processes, in that both rely on the flexible, volitional maintenance and manipulation of mental representations.

What is the neural basis of this mental workspace that appears to be so central to the human capacity for imagination? Traditional neural models of working memory and related processes posit an anatomically modular organization in which physically segregated regions implement component functions such as a "central executive" or a "visuospatial sketchpad" (Lee, Kravitz, & Baker, 2013; Postle, 2006; Baddeley, 2003; Haxby et al., 2001; Ishai, Ungerleider, & Haxby, 2000) in a manner analogous to

Dartmouth College


the functional specialization and anatomical segregation found in other bodily organs. In mental rotation paradigms specifically, previous studies have argued for the existence of anatomically distinct modules specialized for either representation or rotation of mental imagery (Ecker, Brammer, David, & Williams, 2006; Jordan, Heinze, Lutz, Kanowski, & Jäncke, 2001; Richter, Ugurbil, Georgopoulos, & Kim, 1997), for different types of mental rotation (Vingerhoets, de Lange, Vandemaele, Deblaere, & Achten, 2002; Kosslyn, Thompson, Wraga, & Alpert, 2001; Kosslyn, DiGirolamo, Thompson, & Alpert, 1998), and for specifically mental as opposed to perceptual processes (Just, Carpenter, Maguire, Diwadkar, & McMains, 2001; Barnes et al., 2000). However, recently developed information- and network-based neuroscientific methods suggest instead that the mental workspace and its component processes may be implemented in a fundamentally distributed manner across the cortex and some subcortical regions (Schlegel, Alexander, & Tse, 2016; Sporns, 2014; Schlegel et al., 2013; Turk-Browne, 2013; Bassett et al., 2010; van den Heuvel, Stam, Kahn, & Hulshoff Pol, 2009; Tononi, 2008; Rumelhart & McClelland, 1986) as opposed to within anatomically distinct modules. In particular, a recent study showed that information about both visual mental imagery and mental manipulations of that imagery is distributed among several regions across the cortex and that this information is shared in a common format via complex, hierarchical patterns of information flow (Schlegel et al., 2016). If, as these studies suggest, information is fundamentally distributed across the cortex during such high-level mental activity, then how and where does this information originate? Cognitive work such as that of Shepard and Metzler suggests the possibility that, to direct actions within the mental workspace, the brain may recruit existing neural circuitry that evolved for interactions with the physical world.

In fact, several neuroimaging studies have reported activation in various motor areas during mental rotation tasks (Langner et al., 2013; Zacks, 2008; Sack,

Lindner, & Linden, 2007; Michelon, Vettel, & Zacks, 2006; Vingerhoets et al., 2002; Kosslyn et al., 1998; Cohen et al., 1996). In addition, Kosslyn and colleagues (2001) found evidence that participants can be trained to simulate the mental rotation of objects as if they were rotated manually by hand. These findings support the idea that the mental workspace permits mental operations on endogenously constructed models as if they existed physically. However, these neuroimaging studies have given inconsistent accounts of the motor regions involved in mental rotation. Moreover, the univariate analyses used in previous studies reveal increases in cortical activation but are insensitive to the informational content of that activity and are thus difficult to interpret. For instance, processes such as attentional recruitment or motor preparation may occur during mental rotation while not being directly related to the mental rotations themselves. Moreover, recent work has shown that analyses based on univariate BOLD signal differences may miss information processing that does not entail overall changes in metabolic activity (Kohler et al., 2013). This is relevant to the present question, because past studies that did not find recruitment in particular motor regions during mental rotation may have used methods that were insensitive to information processing that existed in these regions but that did not lead to overall changes in BOLD signal response. Given the ambiguity and diversity of past findings, it is still unclear precisely what role motor processing may play, if any, in mental rotation.

In light of the above findings and an emerging view of the mental workspace as both highly flexible and fundamentally distributed, we hypothesized that the network underlying the core functionality of the mental workspace would recruit the motor network into a larger, dynamically constructed network to carry out mental rotation. Here, we define the motor network as the set of brain regions that are responsible for the planning, production, and monitoring of movements. To test the hypothesis that the role of the motor network in mental rotation is to simulate the execution of physical rotations

Figure 1. The eight figures used for mental rotation. Figures were shown during the prompt phase of the trial either as depicted above or flipped across the y axis, and unrotated, rotated 180° around the x axis, or rotated 180° around the z axis.



on imagined mental representations, we additionally investigated the relationship between information processing in the motor network during mental rotations and during corresponding physical hand (“manual”) rotations.

In a variation of Shepard and Metzler’s classic paradigm, we recruited 24 right-handed participants for an initial behavioral session and a subsequent fMRI scanning session in which they completed a series of trials involving either the mental rotation of 3-D cube assemblages (Figure 1) or corresponding empty-handed manual rotations. Figure 2C provides a visual schematic of the experimental trial design. In each mental rotation trial, participants mentally rotated a presented stimulus figure by 90° in one of four rotation directions (Figure 2A, B). Two of these rotations occurred along the x axis (called “forward” and “backward” rotations), and two of these rotations occurred along the z axis (called “left” and “right” rotations). Thus, these four rotations shared the hierarchical relationship shown in Figure 2B, such that rotations in a particular direction were most similar to other rotations along the same direction, moderately similar to rotations along the same axis but in the opposite direction, and least similar to rotations along a different axis.

In manual rotation trials, participants merely rotated their empty right hand in analogous directions. Because of the hierarchical relationship among the rotation directions, we could use multivariate decoding methods to evaluate whether rotation-specific information processing in a set of cortical and subcortical ROIs matched the informational structure of the rotation operations themselves, thus providing a strong test of the functional role of each ROI in the network during mental rotation. We additionally used a newly developed ROI cross-classification analysis to evaluate whether the information carried by each network node was shared among all nodes (Schlegel et al., 2016), as would be expected if information processing in the mental workspace is fundamentally distributed.

We used two strategies to evaluate the hypothesis that the motor network’s role in mental rotation is related to its function during physical motor actions. First, we used a two-group training design similar to that used by Kosslyn and colleagues (2001) to evaluate whether the neural similarity of mental and manual rotations could be manipulated. In an initial behavioral session, participants were randomly assigned to one of two training groups without their knowledge and subsequently completed 100 training trials. Interleaved on half of the training session trials, participants in the first “nonmotoric” training group were shown an animation of the stimulus figure being rotated; in the subsequent fMRI session they were told to “imagine the mental rotations as an internal movie playing in your head.” Instead of the animations, participants in the second “motoric” training group were provided with physical wooden replicas of the stimulus figures that they could rotate manually; in the fMRI session, they were told to “imagine rotating your mental image as you did the physical model.” Our second strategy to evaluate the role of the

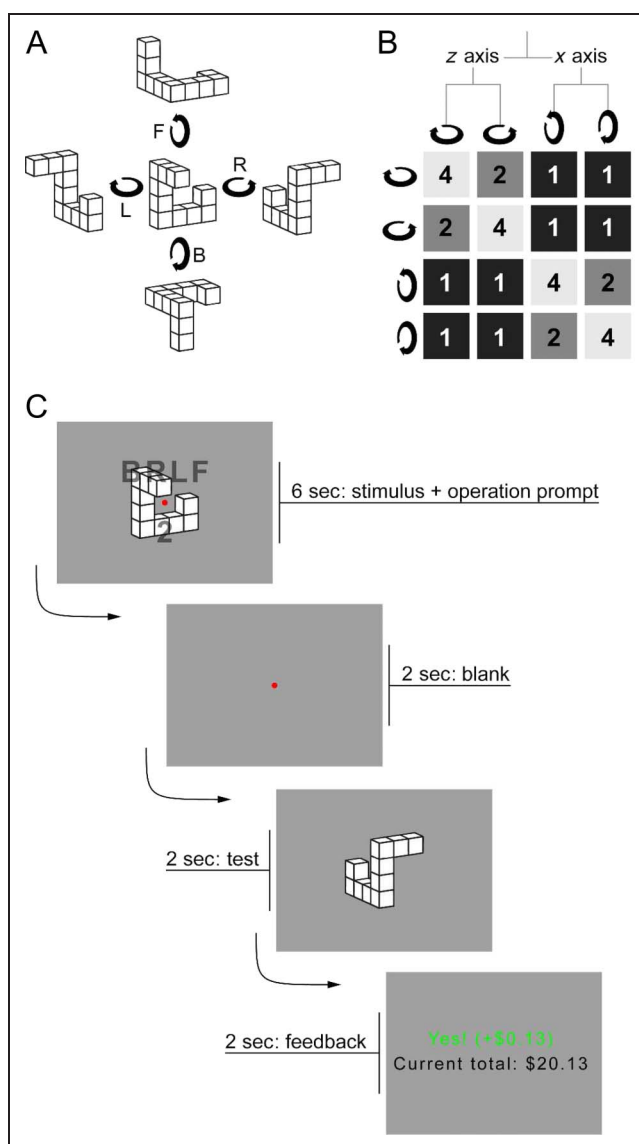


Figure 2. Experimental design. (A) The four 90° mental rotations used in the experiment. Left (“L”) and right (“R”) rotations occurred along the z axis, whereas forward (“F”) and backward (“B”) rotations occurred along the x axis. B. The four rotation directions are related in a two-level hierarchy, such that trials involving a given rotation are most similar to other trials involving the same rotation, moderately similar to trials involving opposite rotations along the same axis, and least similar to trials involving rotations along a different axis. This similarity structure is encoded in the model similarity matrix presented. (C) A schematic of the trial design. In Screen 1, the stimulus figure and mental rotation for the current trial were presented. In Screen 2, only a fixation dot was shown. During the first 8 sec of the trial, participants performed the indicated mental rotation and a concurrent fixation task to ensure that they did not move their hands in accordance with the mental rotation. In Screen 3, a test figure appeared, and participants indicated whether this figure matched the result of the mental rotation. In Screen 4, participants were given feedback about their response.

motor network was to perform a cross-classification analysis comparing the fMRI data from mental rotation and manual rotation trials to assess whether information processing in the motor network during mental rotation

trials resembles information processing during analogous manual rotation trials.

METHODS

Participants

Twenty-four participants (11 women, aged 18–24 years) with normal or corrected-to-normal vision gave informed written consent according to the guidelines of the Committee for the Protection of Human Subjects at Dartmouth College before participating. All were right-handed according to the Edinburgh Handedness Inventory (Oldfield, 1971). Participation consisted of two sessions: one behavioral session in which participants were trained in the task and a subsequent fMRI scanning session.

Task

During each trial, participants performed one of four mental rotations on one of eight figures derived from Shepard and Metzler's (1971) original stimulus set (Figure 1). All rotations were 90°; two rotations were along the x axis (called "forward" and "backward" rotations), and two rotations were along the z axis (called "left" and "right" rotations; see Figure 2A). Each trial lasted 12 sec and consisted of three phases: the task prompt and operation phase (8 sec), the test phase (2 sec), and the feedback phase (2 sec). Figure 2C presents a visual schematic of the following trial description: At the beginning of the prompt/operation phase, a randomly chosen figure from the stimulus set, 8° of visual angle in size, was shown centrally. The figure was shown either as depicted in Figure 1 or flipped across the y axis, and additionally either unrotated or rotated 180° along either the x or z axes. Superimposed on this figure in partially transparent text were two prompts: above, a randomly permuted sequence of the letters L, R, F, and B, and below, an integer from 1 to 4. The integer indicated the position of the letter in the above sequence that denoted the mental rotation to carry out on the current trial (e.g., the integer "3" shown below the sequence "BLFR" would denote the "F" and indicate that the current trial called for a forward rotation). The trial's rotation was indicated in this way to equate the visual stimuli across the four mental rotation conditions. Had each rotation's corresponding prompt letter appeared alone on each trial, the visual stimulus would then have differed systematically between conditions and created a possible visual confound in the subsequent multivariate classification analyses (described below). One could argue that increased attention was still directed to the indicated letter and thus may have led to systematic differences in visual representational processing between conditions. However, our confusion matrix-based classifier performance measure (described below) served as a control for this possibility, as it was sensitive to a particular structure of relationships

between the rotation directions that did not occur between the letter stimuli.

The figure and rotation direction stimuli remained on screen for 6 sec and were replaced by a blank screen for 2 sec. The participant was instructed to perform the indicated mental rotation on the presented figure and to construct as vivid a mental image of the output as possible during this 8-sec period. Additionally, a red fixation dot appeared centrally during this phase of the trial. The fixation dot blinked blue on average once every 2 sec, and the participant was instructed to press the "up" button on a four-button box held in the right hand whenever this color change occurred. This fixation task was used to ensure that participants were (a) awake and attending to the fixation point and (b) not moving their hand to mimic the mental rotation being performed.

After the prompt/operation phase, a test figure appeared on the screen for 2 sec. On half of trials, the test figure was the initial prompt figure as it would appear after having undergone the indicated rotation ("correct" figure); on the other half of trials, the test figure was a y -axis-flipped version of the initial prompt figure that had undergone the same rotation ("mirror image" figure). Participants were instructed to indicate within the 2 sec that the test figure was present on screen whether it was the correct figure ("left" button) or the mirror image figure ("right" button).

Finally, a feedback screen indicated whether the participant made the correct response and, during the fMRI session, the current reimbursement amount. As an incentive to attend carefully during the approximately 1.5-hr fMRI session, participants gained \$0.125 for each correct response and lost \$0.625 for each incorrect response, with a baseline, minimum reimbursement of \$20 and a maximum of \$40.

Each 5-min 28-sec run of the fMRI session consisted of 16 trials (four trials of each rotation type in counter-balanced order), with 8 sec of rest occurring between each trial. The fMRI session consisted of 10 runs of mental rotation trials followed by three runs of analogous hand rotation trials, in which physical rotations of the empty right hand were performed instead of mental rotations. Hand rotation trials matched the design of the mental rotation trials except that no figures were shown, no fixation task or test response was required, and participants merely rotated their right hand continuously according to the prompt until the word "Stop" appeared at the time at which the feedback screen appeared during mental rotation trials. Before the hand rotation runs, videos were shown to the participant to demonstrate proper hand rotations in each of the four directions. Hand rotations resembled the motor actions that would be performed if a physical object was rotated in the same manner as the mentally rotated figures. Even though left and right rotations and forward and backward rotations, respectively, involved back and forth hand rotations along the same axis, instructions and the videos made

clear that more emphasis was to be placed on motion in the indicated direction (e.g., more emphasis on the forward phase of rotation during forward rotation trials). Participants were not told about the hand rotation runs until they occurred to avoid biasing participants to imagine their hands playing any role during preceding mental rotation runs.

Training

During the initial behavioral session, participants were instructed in the task and completed 100 practice trials. The prompt/operation phase of practice trials was self-paced: Participants viewed the prompt stimulus for as long as desired and indicated with a key press when they were ready for the test phase. In half of the practice trials, the prompt stimulus was accompanied by a guide to assist participants in performing the mental rotation, and in the other half of the trials, the prompt occurred without a guide. Guide and no-guide trials were interleaved. Without their knowledge, participants were divided randomly into two training groups (12 participants in each group). In the nonmotoric training group, the guide was a looping animation shown below the prompt stimulus that depicted the figure undergoing the indicated rotation. In the motoric training group, the guide was a physical, wooden model that matched the prompted figure and that participants held and rotated manually.

MRI Acquisition

MRI data were collected using a 3.0-T Philips Achieva Intera (Andover, MA) scanner with a 32-channel sense head coil located at the Dartmouth Brain Imaging Center. One T1-weighted structural image was collected using a magnetization-prepared rapid acquisition gradient-echo sequence (8.176 msec repetition time [TR]; 3.72 msec echo time; 8° flip angle; 240 × 220 mm field of view; 188 sagittal slices; 0.9375 × 0.9375 × 1 mm voxel size; 3.12 min acquisition time). T2*-weighted gradient EPI scans were used to acquire functional images covering the whole brain (2000 msec TR, 20 msec echo time; 90° flip angle, 240 × 240 mm field of view; 3 × 3 × 3.5 mm voxel size; 0 mm slice gap; 35 slices).

MRI Data Preprocessing

High-resolution anatomical images were processed using the FreeSurfer image analysis suite (Dale, Fischl, & Sereno, 1999). fMRI data were motion and slice-time corrected, temporally high pass filtered with a 100-sec cutoff, and spatially smoothed with a 6-mm FWHM Gaussian kernel, all using FSL (Smith et al., 2004). Data from each run were concatenated temporally for each participant after aligning each run using FSL's FLIRT tool and demeaning each voxel's time course. For the ROI classification (described below), data were prewhitened for each ROI separately

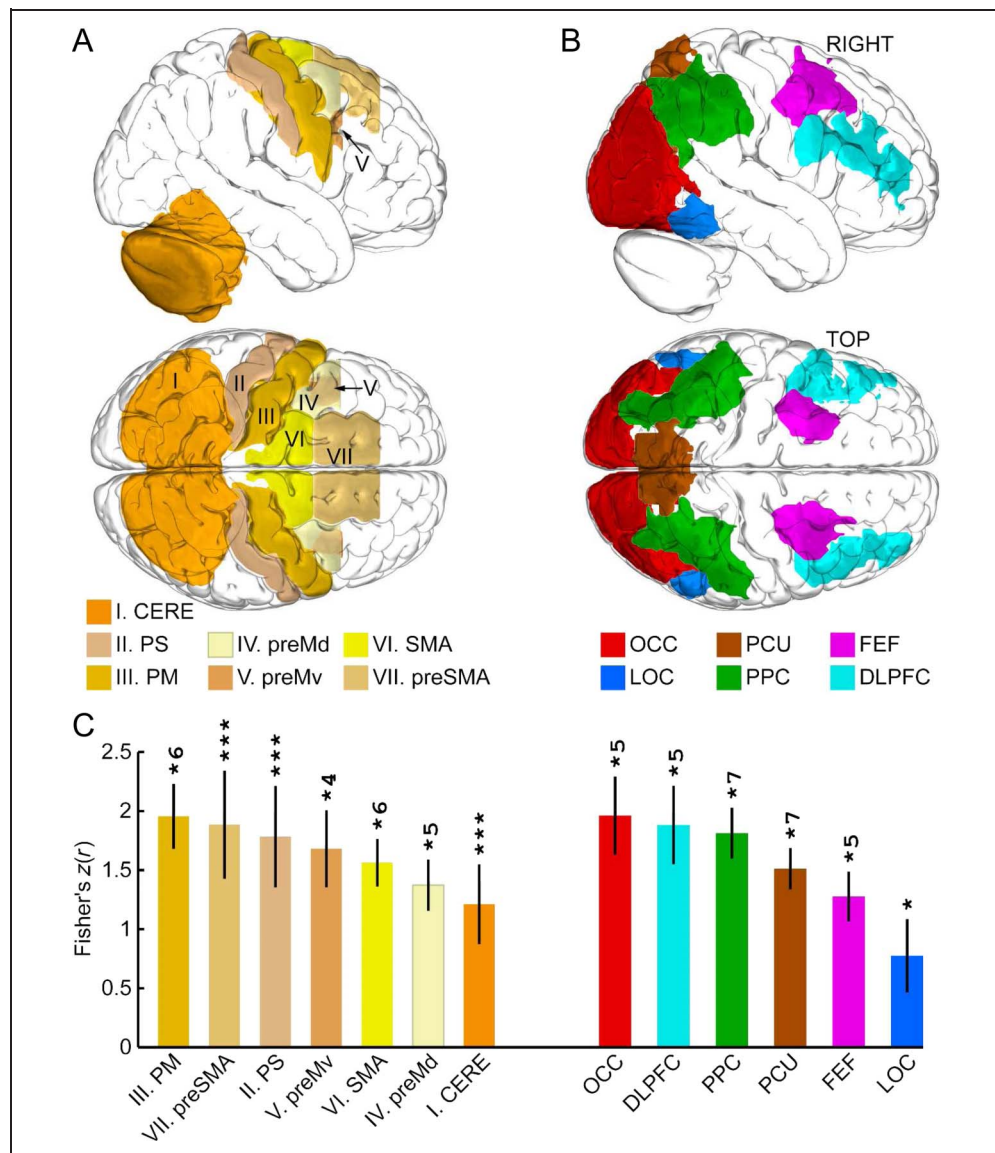
using FSL's MELODIC tool (i.e., principal components were extracted using MELODIC's default dimensionality estimation method with a minimum of 10 components per ROI).

ROI Classification

For each of the 13 ROIs, we used PyMVPA (Hanke et al., 2009) to perform a spatiotemporal multivariate classification analysis between the four mental rotation directions. Five of these ROIs (lateral occipital cortex, posterior parietal cortex, precuneus, dorsolateral pFC, and FEFs), along with three additional ROIs used in a control classification analysis (medial-temporal lobe, medial pFC, and thalamus) were functionally defined, bilateral masks in Montreal Neurological Institute (MNI) space that were then transformed into each participant's native functional space. In a previous study, these five ROIs were defined by their greater recruitment during tasks involving the mental manipulation of visual imagery over simple maintenance of visual imagery (Schlegel et al., 2013). Additionally, these ROIs, along with an occipital (OCC) ROI that was defined anatomically for each participant, supported information about specific manipulations of visual imagery. The OCC ROI was defined in each participant's native anatomical space using the following labels from FreeSurfer's cortical parcellation: inferior occipital gyrus and sulcus, middle occipital gyrus and sulci, superior occipital gyrus, cuneus, occipital pole, superior occipital and transverse occipital sulci, and anterior occipital sulcus (all bilateral). The remaining seven motor network ROIs were defined anatomically using the following FreeSurfer labels (again all bilateral): cerebellar cortex, primary somatosensory cortex (postcentral gyrus), primary motor cortex (precentral gyrus, central sulcus, precentral sulcus [inferior and superior parts]), dorsal premotor cortex (posterior third of the middle frontal gyrus, lateral half of the posterior third of the superior frontal sulcus), ventral premotor cortex (inferior frontal sulcus, opercular part of the inferior frontal gyrus), SMA (posterior third of the superior frontal gyrus, medial half of the posterior third of the superior frontal sulcus), pre-SMA (middle third [in the posterior-anterior direction] of the superior frontal gyrus). In a postprocessing step for each participant, voxels that were initially shared between multiple ROIs were assigned to only one ROI using the following, descending order of precedence: ventral premotor cortex, dorsal premotor cortex, SMA, pre-SMA, primary motor cortex, primary somatosensory cortex, cerebellar cortex, dorsolateral pFC, FEF, posterior parietal cortex, precuneus, lateral occipital cortex, OCC. The ROIs shown in Figure 3A, B were created as described above but using the MNI template brain for visualization.

For the spatiotemporal multivariate classification, we used a linear support vector machine classifier and leave-one-trial-out cross validation. Because we only considered correct-response trials, a nonuniform number of

Figure 3. ROI multivariate classification results. (A) Seven bilateral motor network ROIs used in the analyses are shown on an MNI template brain. CERE = cerebellar cortex; PS = primary somatosensory cortex; PM = primary motor cortex; preMd = dorsal premotor cortex; preMv = ventral premotor cortex. (B) Six bilateral ROIs found in previous studies (Schlegel et al., 2013, 2016) to mediate processing in the mental workspace and used in the analyses here. **OCC** = occipital cortex; LOC = lateral occipital cortex; PCU = precuneus; PPC = posterior parietal cortex; DLPFC = dorsolateral pFC. (C) Results of four-way classification analyses in each ROI. Correlations between resulting confusion matrices and the similarity structure in Figure 2B are Fisher's Z-transformed. Error bars are jackknife-corrected standard errors of the mean (see Methods). Asterisks indicate significance in a one-tailed jackknife *t* test comparing Fisher's Z-transformed correlations to zero across subjects ($*p \leq .05$; $***p \leq .001$; $^{(n)}p \leq 1 \times 10^{-n}$). Results are FDR-corrected for multiple comparisons across the 13 ROIs.



trials existed for each condition and participant (35.4 trials per condition on average; see Table S1 for details). Even though these differences were small, we ensured that they could not affect the classification results by including a target balancing step in our cross-validation procedure. In this step, each classification fold was performed 10 times using random, balanced samples of the training data, and the results for that fold were averaged across the 10 bootstrapped folds. For each classification, we used the spatiotemporal pattern of prewhitened BOLD signal data from the first 5 TRs of each correct response trial, shifted by 1 TR to account for the hemodynamic response function delay inherent in fMRI data. We shifted by only 1 TR to include as much trial data as possible. Prewhitening reduced each ROI's voxel-based pattern to an average of 72.3 data features ($SEM = 3.69$). Thus, each classification used spatiotemporal patterns of, on average, 361 dimensions ($SEM = 18.5$). Each feature dimension was z-scored by run before classification

to reduce between-run differences in signal that may have occurred due to scanner or physiological noise.

The result of each cross-validation was a 4×4 confusion matrix that represented a summary record of the classifier's predicted labels relative to the true target labels across all cross-validation folds. A perfect classifier would yield a confusion matrix with nonzero values only along the diagonal, because cells along the diagonal represent instances in which the target and predicted labels were the same. However, because we used mental rotations that shared a specific hierarchical similarity relationship (see Figure 2B), we expected the classifier to make a specific pattern of confusions among the brain activity patterns in ROIs that were involved in carrying out those mental rotations. For example, we expected that the classifier would confuse a left rotation with a right rotation (both along the *z* axis) more often than it would confuse a left rotation (*z* axis) with a forward rotation (*x* axis), but only if the information processing underlying the brain

activity patterns was related specifically to the mental rotations that were performed. Thus, our measure of classifier performance was the correlation between the confusion matrix resulting from the cross-validation and the matrix form of the rotation-direction similarity structure shown in Figure 2B. Note that because we used correlation as our measure, the absolute numerical values of this model similarity matrix are not important. Only the pattern of relative magnitudes of values matter for the correlation calculation, in this case signifying that a trial involving a particular rotation direction is most highly related to trials with the same rotation, moderately related to trials with opposite rotation directions along the same axis, and least related to trials with rotations along a different axis. This confusion matrix correlation measure has been used successfully in previous studies to probe the complex structure of information processing in the mental workspace (Schlegel et al., 2013, 2016). It is additionally more sensitive than classification accuracy because it also takes into account confusions between conditions that result from the hierarchical relationship between the rotations.

We used a jackknife procedure to perform random-effects analyses evaluating the significance of the correlations (Schlegel et al., 2016; Miller, Patterson, & Ulrich, 1998). In the case of noisy estimates such as individual subject confusion matrices, jackknife analyses can provide cleaner results without biasing statistical significance. In a jackknifed analysis with N subjects, N grand means of the data (in this case, confusion matrices) are calculated, each with one subject left out. The correlation between each of these grand mean confusion matrices and the model similarity structure was then calculated, and a one-tailed t test evaluated whether the Fisher's Z -transformed correlations were positive (i.e., whether there was a significant correlation between confusion matrices and the model similarity structure across participants). Because the jackknife procedure reduces the variance between subjects artificially, a correction must be applied to the t -statistic calculation; specifically, the sample standard deviation between correlations is multiplied by the square root of $(N - 1)$. This jackknife method was developed for studying ERP latency differences but applies generally to estimates which may be noisy on a single-subject level.

ROI Cross-classification Analysis

To assess whether information about mental rotations was shared in a common format between areas, we performed a cross-classification analysis in which a classifier was trained on data from one ROI and tested on data from a second ROI (see Schlegel et al., 2016).

A technical challenge to cross-classifying between ROIs is that each ROI exists initially as an incompatible voxel-based feature space (i.e., each ROI consists of a different number of voxels [feature dimensions], and there is no meaningful mapping between the voxels of each ROI).

Thus, cross-classification between two ROIs first requires that their data be transformed into a common feature space. To do this, we first used FSL's MELODIC tool to transform each ROI's data from voxel space to 50 principal component signals using PCA. After this step, each ROI's pattern had the same dimensionality, but those patterns' features would be unlikely to correspond. Therefore, for each pair of ROIs, these component signals were matched pairwise as follows to maximize the total similarity between component signals: First, the correlation distance ($1 - |r|$) between each pair of components was calculated, yielding a 50×50 correlation distance matrix. Next, the rows and columns of this matrix were reordered using the Hungarian algorithm to minimize the matrix trace (Kuhn, 1955). The components meeting along the diagonal of this reordered, trace-minimized matrix defined the pairwise matching. If two components were matched by this procedure but were anticorrelated, one component was negated to produce positively correlated component pairs.

We performed this matching procedure for each fold of the cross-validation independently, excluding test data to avoid artificially inflating the similarity between training and testing patterns. Once this procedure was completed, data from the two ROIs shared a common feature space, that is, the two feature spaces had the same dimensionality and corresponding features in the two spaces were maximally similar. Cross-classification could then proceed by training the classifier on data from one ROI and testing it on data from the other ROI. Each ROI served both as the training set and as the testing set in separate cross-validation folds, with confusion matrices averaged between the two folds. Other than the feature-space transformation and difference in training and testing data sets, the ROI cross-classification was conducted exactly as described for the ROI classification above.

Mental/Manual Rotation Cross-classification

To assess whether motor involvement in mental rotation resembled motor activity during physical rotation of the hands, we performed a cross-classification analysis for each ROI in which we trained a classifier on data from the mental rotation trials and tested the classifier on data from the manual rotation trials and vice versa. Mental rotation trials were given the same labels as the corresponding manual rotation trials (e.g., trials in which forward mental rotations were carried out were given the same label as trials in which a forward hand rotation was prompted). The classification analysis was performed and evaluated identically to the ROI classification analysis described above except for the difference between training and testing datasets. Note that each cross-validation involved only two folds in this analysis (train on mental rotation and test on manual rotation, train on manual rotation and test on mental rotation), but the same 10-subfold target balancing procedure was used to ensure that training

data were balanced. As in the previous analyses, confusion matrices were averaged across folds.

ROI BOLD Signal Comparison of Training Groups

To assess whether the two different training procedures induced differential brain activity that reflected different cognitive strategies employed during mental rotation, for each ROI we conducted a two-tailed unpaired *t* test across participants comparing trial-related mean BOLD activity between the training groups.

We initially used FSL's FEAT tool to perform a first-level whole-brain GLM for each participant in which we defined boxcar predictors for correct response and incorrect response trials. The resulting voxel-wise beta-weights for the correct response predictor, representing the average change in BOLD signal in each voxel during correct response mental rotation trials compared with rest, were then averaged across each ROI. This procedure yielded a single mean trial-related activity estimate for each participant and ROI. For each ROI, these values were then partitioned by training group and used to perform a two-tailed unpaired *t* test.

Whole-brain BOLD Signal Comparison of Training Groups

In a more exploratory variant of the ROI-based BOLD signal comparison of training groups described above, we performed a whole-brain gray matter-only BOLD signal comparison using FSL's permutation-based randomize tool with 5000 permutations. The gray matter mask used to restrict the analysis was derived from FreeSurfer's gray matter segmentation of the MNI template brain. The input data to randomize were the correct response beta-weight volumes resulting from the first-level GLM analysis described above (one volume for each participant). The design matrix supplied to randomize defined a single predictor that differentiated between nonmotoric and motoric participants. *t* contrasts were defined for nonmotoric > motoric and motoric > nonmotoric.

ROI Classification Comparison of Training Groups

To assess whether patterns of mental rotation-related activity differed between the two training groups, we used a procedure similar to that used in the ROI-based BOLD signal comparison described above to compare the results of the ROI classification analyses between the groups. In this case, our inputs to the unpaired *t* tests were the classification results for each participant and ROI, specifically the Fisher's Z-transformed correlations between each participant's confusion matrix resulting from the four-way classification and the model mental rotation similarity structure.

RESULTS

Performance accuracy was high during the fMRI session (mean correct response rate was 88.5% [*SEM* = 1.02%] across all participants and conditions), indicating that participants had little difficulty carrying out the instructed mental rotations. A one-way ANOVA showed no significant differences in the correct response rate between conditions [$F(3, 92) = 1.33, p = .270$], confirming that the difficulty was well matched between rotation conditions (see Supplemental Table S1 for behavioral results separated by rotation direction). We additionally found no significant behavioral differences between the two training groups [correct response rate for nonmotoric training group: 86.2% (*SEM* = 2.87%); for motoric training group: 90.9% (*SEM* = 1.64%); $t(22) = -1.42, p = .170$], and no significant training group-by-condition interaction for correct response rate [$F(3, 66) = 2.13, p = .104$].

ROI Classification Analysis

We defined 13 ROIs for each subject that we evaluated for mental rotation-specific information processing using a multivariate classification analysis (Figure 3A, B). Seven of these ROIs were anatomically defined regions of the motor network, and six were previously shown to form part of a cortex-wide network of regions that mediate mental workspace processes (Schlegel et al., 2013, 2016; see Methods for details on how each ROI was defined). The classification analysis used a standard cross-validation procedure (Norman, Polyn, Detre, & Haxby, 2006). Briefly, in each fold of the cross-validation, a linear support vector machine classifier was initially trained by presenting it with a set of brain activity patterns derived from individual correct-response mental rotation trials along with the directions of the rotations performed on those trials. In a subsequent testing step, the classifier was presented with a holdout sample activity pattern without a rotation-direction label and its ability to correctly label the pattern based on the previous training step was evaluated. Our measure of classifier performance was the correlation between the confusion matrices resulting from the cross-validation procedure and the model similarity structure in Figure 2B (see Methods).

We conducted this procedure for each ROI and participant individually and then assessed the information content within each ROI by performing an across-subject random effects analysis to determine whether a significant correlation existed between that ROI's confusion matrices and the model rotation-direction similarity structure. Results of this analysis are presented in Figure 3C, showing that each of our 13 ROIs supported robust information processing related to mental rotation (all results are false discovery rate [FDR]-corrected for multiple comparisons across the 13 ROIs; see Figure 4 for confusion matrices for each ROI). This result may seem surprising according to traditional models of functional localization, because it




Figure 4. Individual mean confusion matrices from the ROI classification analysis. Compare to model similarity structure in Figure 2B. Values represent percentage of cross-validation folds in which each target/prediction combination occurred. Color scaling for visualization was performed separately for each confusion matrix, because only relative values matter for the correlation analysis. Matrix elements are ordered as in Figure 2B, and abbreviations are as in Figure 3.

indicates that areas as seemingly unrelated to the rotation directions as occipital cortex and primary somatosensory cortex carry information about specific mental rotations. However, this finding is consistent with previous results suggesting that information processing in the mental workspace is fundamentally distributed in the sense that traditional anatomical boundaries of functionality break down in these high level mental processes (e.g., Schlegel et al., 2016). In particular, this analysis establishes robustly that information processing related directly to mental rotation occurs throughout the motor network.

To confirm that our results were specific to the 13 mental workspace and motor network ROIs we studied, we additionally performed ROI classification analyses as above but on three control ROIs: medial-temporal lobe, medial pFC, and thalamus. In a previous study of the mental workspace, these regions showed differences in activity depending on whether participants manipulated or simply maintained mental representations, but activity in these areas was not specific to the type of mental manipulation performed. As expected, none of the classifications in these three control regions were significant (Supplemental Table S2).

ROI Cross-classification Analysis

We next sought to assess whether the processing of mental rotations is truly distributed across the 13 regions of this network. An alternative possibility is that each of the 13 regions plays a role in mental rotation, but that processing in each area is functionally isolated as would be expected in the case of anatomically modular functional localization. Investigating this question also allowed us to evaluate whether the motor network plays a separate role or becomes integrated into the larger mental workspace network during mental rotation. We used a recently developed ROI cross-classification analysis to assess these alternatives. In this analysis a classifier is trained on data from one ROI and tested on data from a different ROI (Schlegel et al., 2016). A successful cross-classification would provide evidence that information is shared between the two ROIs. An unsuccessful cross-classification would leave open the possibility that the two ROIs represent information in a nonoverlapping manner.

We performed this cross-classification analysis for each ROI pair, with results shown in Figure 5 (all results FDR-corrected across the 78 ROI pairs). Each arc represents a successful cross-classification, indicating that information associated with mental rotations is shared between that pair of ROIs. Connections within the motor or core mental workspace subnetworks are shown in light orange and light blue, respectively, whereas connections across these two subnetworks are shown in dark blue. We could successfully cross-classify between most pairs of ROIs, suggesting that information processing related to mental rotations is shared in a distributed manner across the network. In particular, a robust set of connections exist both within and between the motor network and other mental workspace regions, suggesting that the motor network becomes tightly integrated into the greater mental workspace network during mental rotation.

Mental/Manual Rotation Cross-classification

What role does the motor network play in mental rotation? To assess the possibility that the mental workspace recruits the motor network to simulate mental rotations as if they were manual rotations of physical objects, we performed a cross-classification analysis within each ROI in which we trained a classifier on data from mental rotation trials and tested it on data from manual rotation trials and vice versa. A successful cross-classification in a given ROI would imply that mental and manual rotations share overlapping neural implementations within that ROI. Other than the difference in training and testing data sets and the different number of cross-validation folds (data in this analysis were partitioned by mental/manual rotation condition rather than by trial), the classification analysis was performed and evaluated exactly as in the ROI classification analysis. Results of the cross-classification for each ROI are presented in Figure 6 (all results





Figure 5. ROI cross-classification results. Arcs indicate pairs of ROIs in which cross-classification was successful, that is, in which a one-tailed, jackknife t test of Fisher's Z -transformed correlations between confusion matrices and the model similarity structure in Figure 2B was greater than zero across participants. Results are FDR-corrected across the 78 ROI pairs. Arc thickness scales with Z -transformed correlation. Connections within the motor network are colored light orange, those within the core mental workspace network are colored light blue, and those crossing between the subnetworks are colored dark blue. Abbreviations are as in Figure 3.

FDR-corrected across the 13 ROIs). Three ROIs showed significant informational similarity between mental and manual rotations. One of these ROIs, primary motor cortex, was in the motor network, and two ROIs, posterior parietal cortex and precuneus, were in the core mental workspace network. Two additional ROIs showed significant cross-classification results that did not pass multiple comparisons correction (SMA and dorsolateral pFC). Thus,

Figure 6. Manual/mental cross-classification results. Results of four-way classification analyses in which the classifier was trained using data from mental rotation trials and tested using data from manual rotation trials, and vice versa. Correlations between resulting confusion matrices and the similarity structure in Figure 2B are Fisher's Z -transformed. Error bars are jackknife-corrected standard errors of the mean (see Methods). Asterisks indicate significance in a one-tailed jackknife t test

comparing Fisher's Z -transformed correlations to zero across subjects ($^{\#}p \leq .05$ before FDR correct; $*p \leq .05$; $**p \leq .01$). Results are FDR-corrected for multiple comparisons across the 13 ROIs. Abbreviations and ordering are the same as in Figure 3.



some of the tested ROIs appear to share overlapping implementations of mental and manual rotations, whereas others may support each process in a distinct manner.

Between-group Differences in Mental Rotation

We reported above that the nonmotoric and motoric training groups did not show significant differences in behavioral performance. However, as Kosslyn et al. (2001) suggest, participants in different training groups may still have employed different cognitive strategies that would lead to differences in information processing when performing mental rotation. We conducted several analyses to evaluate this possibility.

First, we conducted a univariate analysis similar to that used by Kosslyn et al. (2001) to assess whether training induced differences in mental rotation-related brain activity. We initially restricted our analysis to the 13 ROIs from the previous analyses. For each ROI and participant, we calculated the mean BOLD signal activity change during mental rotation trials. For each ROI, we then performed a two-tailed, unpaired t test to assess whether these mean mental rotation-related activity levels differed between the two groups. No ROI showed a significant difference in activity after FDR correction across the 13 ROIs (see Supplemental Table S3 for results). We next conducted an analogous but more exploratory whole-brain analysis to identify regions of the cortex that showed differences in activity between the two groups. No voxels were significant in this analysis after FDR correction.

Although we found no univariate differences in brain activity between the two training groups, our more sensitive multivariate analysis might still show that information processing differed between the groups. To assess this possibility, we performed two-tailed, unpaired t tests as above but compared the ROI classification results, the ROI cross-classification results, and the mental/manual cross-classification results between the two groups. Each set of t tests was FDR-corrected independently, and none

of these tests showed significant differences between the two groups after correction (see Tables S4–S6). Thus, we failed to replicate the findings of Kosslyn et al. (2001), because none of our multiple analyses found a behavioral or neuronal difference between the two groups because of the training manipulation.

DISCUSSION

Here we investigated the role of the motor network during mental rotation and its integration into the wider mental workspace. We found that the motor network supported robust information processing related directly to mental rotation and that this processing became dynamically integrated with a distributed, cortex-wide neural network underlying the mental workspace. These findings support a model of the mental workspace as consisting of a flexible core network that can dynamically recruit domain-specific subnetworks for specific functions, much like a general contractor would employ specialists as needed for specific jobs.

Each of the seven motor network ROIs that we tested carried information about specific mental rotations. This result held even in primary somatosensory cortex, a region better known for its role in mediating peripheral sensation. Although perhaps surprising, previous studies of mental rotation have found increases in activity during mental rotation in this and several other areas of the cortex (Zacks, 2008). The present results move beyond this previous work by showing that activity in each of these regions is specific to the mental rotations that participants performed. Thus, many areas in and beyond the motor network appear to play a functional role in carrying out mental rotations.

Not only do regions throughout the cerebral and cerebellar cortex support information specific to mental rotation, as revealed by our ROI-based multivariate classification analysis, but this information additionally appears to be shared throughout a widely distributed network. Our ROI cross-classification analysis found that many pairs of ROIs in the network that we studied shared information in the sense that a classifier could use information from one ROI to make a mental rotation-related prediction based on information from a different ROI. This information sharing held true both within the motor network, suggesting that several subregions of the motor network become tightly integrated during mental rotation, and also between the motor network and a core network of regions underlying the mental workspace. Such widely distributed information about mental rotations and the associated dense pattern of information sharing speak against strict functional or anatomical modularity among cortical regions. In particular, they suggest that information processing subserving mental rotation entails a breakdown in the anatomical modularity argued for by traditional models of working memory such as Baddeley's (2003) that derive from functional localization methods

which merely measure mean BOLD signal levels in isolation and are insensitive to distributed patterns of information processing.

How could a large and widely dispersed array of neurons implement the kind of distributed information sharing observed in this study? A recent neurophysiology study by Siegel, Buschman, and Miller (2015) found results compatible with ours in an overlapping set of regions (middle temporal area, visual area 4, inferior temporal cortex, lateral intraparietal area, pFC, and FEF) in two monkeys that performed a task involving selective attention to either motion or color and a subsequent choice between one of two saccades. Their findings suggest that, although task-relevant information may be generated initially in domain-specific cortical regions (e.g., middle temporal area for motion or pFC for behavioral choice), information about many task variables soon becomes widely distributed among many regions of the cortex. Along these lines, recent calls have been made to envision the architecture of the brain as simultaneously supporting both highly distributed, parallel processing strategies at multiple levels of cognition and modular, serial operations taking place within highly interconnected networks (Singer, 2013). Our ROI cross-classification results suggest that a common representational format may underlie the interregional communication and coordination that would be required within such a distributed system.

Our findings are consistent with a model of the mental workspace that involves a domain general core network that can recruit other specialized subnetworks (e.g., the visual cortex or motor network) for specific tasks as needed. In particular, we found that the motor network was recruited and tightly integrated into the wider core mental workspace network during mental rotation. Consistent with the proposal that the motor network's role is to simulate rotations of imagined objects as if they existed physically, we found that information processing in some regions of the network resembled information processing that occurred during actual physical hand rotations. However, in other regions both within and outside the motor network, we found no similarity between mental and manual rotations. We also did not find that training participants to think of mental rotations as simulations of manual manipulations of physical objects had any effect on subsequent neural activity. The reason for our nonreplication of the effect reported by Kosslyn et al. (2001) is unclear. However, this difference in results may underscore the flexibility of the mental workspace that could allow different participant groups to implement the same functions (e.g., mental rotation) using widely different strategies. In summary, our results suggest that, although the motor network may contribute specialized action-related functionality to the mental workspace during mental rotation, its constituent nodes are also recruited in novel ways for processing that is unique to purely mental simulations.

Much of the last two decades of cognitive neuroscience research has been concerned with assigning functions to localized regions of the cortex in what has been described as a kind of “neophrenology” (Uttal, 2003). However, recent studies such as ours and that of Siegel and colleagues (Schlegel et al., 2013, 2016; Siegel et al., 2015) and recent work focusing on the brain as a densely connected network (Sporns, 2014; Bassett & Gazzaniga, 2011) suggest instead that high-level cognition and possibly cognition generally may entail fundamentally distributed processing with a concomitant breakdown of local specialization of function. Furthermore, these findings suggest that distributed informational processing may coexist with functionally localized processing, either on different timescales or at different levels of informational organization. These new models may hint at a level of neural information processing that could form the basis of conscious activity similar to that of the Global Workspace Theory proposed by researchers such as Baars and Dehaene (Baars, 2002; Dehaene & Naccache, 2001), while remaining consistent with localized accounts proposed by Zeki and Bartels (1999). Future work should investigate the range of cognitive processes that entail dynamically distributed processing such as that described here. Is this kind of fundamentally distributed information processing unique to high-level mental functions, or might new methodological advances reveal that distributed processing is the rule rather than the exception for the brain?

Acknowledgments

A. S. would like to thank Patrick Cavanagh, David Kraemer, and Frank Tong for serving on his PhD dissertation committee, of which this study was a part. This study was funded by NSF Graduate Research Fellowship 2012095475 to A. S. and Templeton Foundation grant 14316 to P. U. T. All data and code are available online at www.dartmouth.edu/~petertse. Supplementary information can be found at www.alexschlegel.com/research/mental_rotation/.

Reprint requests should be sent to Alexander Schlegel, Department of Psychological and Brain Sciences, H. B. 6207, Moore Hall, Dartmouth College, Hanover, NH 03755, or via e-mail: schlegel@gmail.com.

REFERENCES

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, *45*, 1363–1377.

Baars, B. J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Sciences*, *6*, 47–52.

Baddeley, A. D. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, *4*, 829–839.

Barnes, J., Howard, R. J., Senior, C., Brammer, M., Bullmore, E. T., Simmons, A., et al. (2000). Cortical activity during rotational and linear transformations. *Neuropsychologia*, *38*, 1148–1156.

Bassett, D. S., & Gazzaniga, M. S. (2011). Understanding complexity in the human brain. *Trends in Cognitive Sciences*, *15*, 200–209.

Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2010). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences, U.S.A.*, *108*, 7641–7646.

Bassok, M., Dunbar, K. N., & Holyoak, K. J. (2012). Introduction to the special section on the neural substrate of analogical reasoning and metaphor comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 261–263.

Blazhenkova, O., & Kozhevnikov, M. (2010). Visual-object ability: A new dimension of non-verbal intelligence. *Cognition*, *117*, 276–301.

Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., et al. (1996). Changes in cortical activity during mental rotation. A mapping study using functional MRI. *Brain*, *119*, 89–100.

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis: I. Segmentation and surface reconstruction. *Neuroimage*, *9*, 179–194.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

Ecker, C., Brammer, M. J., David, A. S., & Williams, S. C. (2006). Time-resolved fMRI of mental rotation revisited—Dissociating visual perception from mental rotation in female subjects. *Neuroimage*, *32*, 432–444.

Finke, R. A., & Slayton, K. (1988). Explorations of creative visual synthesis in mental imagery. *Memory & Cognition*, *16*, 252–257.

Hanke, M., Halchenko, Y. O., Sederberg, P. B., Hanson, S. J., Haxby, J. V., & Pollmann, S. (2009). PyMVPA: A python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics*, *7*, 37–53.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.

Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, *8*, 280–285.

Ishai, A., Ungerleider, L. G., & Haxby, J. V. (2000). Distributed neural systems for the generation of visual images. *Neuron*, *28*, 979–990.

Johnson-Laird, P. N. (2010). Mental models and human reasoning. *Proceedings of the National Academy of Sciences, U.S.A.*, *107*, 18243–18250.

Jordan, K., Heinze, H. J., Lutz, K., Kanowski, M., & Jäncke, L. (2001). Cortical activations during the mental rotation of different visual objects. *Neuroimage*, *13*, 143–152.

Just, M. A., Carpenter, P. A., Maguire, M., Diwadkar, V., & McMains, S. (2001). Mental rotation of objects retrieved from memory: A functional MRI study of spatial processing. *Journal of Experimental Psychology: General*, *130*, 493–504.

Kohler, P. J., Fogelson, S. V., Reavis, E. A., Meng, M., Guntupalli, J. S., Hanke, M., et al. (2013). Pattern classification precedes region-average hemodynamic response in early visual cortex. *Neuroimage*, *78*, 249–260.

Kosslyn, S. M., Behrmann, M., & Jeannerod, M. (1995). The cognitive neuroscience of mental imagery. *Neuropsychologia*, *33*, 1335–1344.

Kosslyn, S. M., DiGirolamo, G. J., Thompson, W. L., & Alpert, N. M. (1998). Mental rotation of objects versus hands: Neural mechanisms revealed by positron emission tomography. *Psychophysiology*, *35*, 151–161.

- Kosslyn, S. M., Thompson, W. L., Wraga, M., & Alpert, N. M. (2001). Imagining rotation by endogenous versus exogenous forces: Distinct neural mechanisms. *NeuroReport*, *12*, 2519–2525.
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*, 83–97.
- Langner, R., Sternkopf, M. A., Kellermann, T. S., Grefkes, C., Kurth, F., Schneider, F., et al. (2013). Translating working memory into action: Behavioral and neural evidence for using motor representations in encoding visuo-spatial sequences. *Human Brain Mapping*, *35*, 3465–3484.
- Lee, S.-H., Kravitz, D. J., & Baker, C. I. (2013). Goal-dependent dissociation of visual and prefrontal cortices during working memory. *Nature Neuroscience*, *16*, 997–999.
- Logie, R. H. (2003). Spatial and visual working memory: A mental workspace. *Psychology of Learning and Motivation*, *42*, 37–78.
- Michelon, P., Vettel, J. M., & Zacks, J. M. (2006). Lateral somatotopic organization during imagined and prepared movements. *Journal of Neurophysiology*, *95*, 811–822.
- Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*, 99–115.
- Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, *10*, 424–430.
- Oldfield, R. C. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, *9*, 97–113.
- Paivio, A. (1978). Comparisons of mental clocks. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 61–71.
- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*, 23–38.
- Richter, W., Ugurbil, K., Georgopoulos, A., & Kim, S. G. (1997). Time-resolved fMRI of mental rotation. *NeuroReport*, *8*, 3697–3702.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. Cambridge, MA: MIT Press.
- Sack, A. T., Lindner, M., & Linden, D. E. J. (2007). Object- and direction-specific interference between manual and mental rotation. *Perception & Psychophysics*, *69*, 1435–1449.
- Schlegel, A., Alexander, P., & Tse, P. U. (2016). Information processing in the mental workspace is fundamentally distributed. *Journal of Cognitive Neuroscience*, *28*, 295–307.
- Schlegel, A., Kohler, P. J., Fogelson, S. V., Alexander, P., Konuthula, D., & Tse, P. U. (2013). Network structure and dynamics of the mental workspace. *Proceedings of the National Academy of Sciences, U.S.A.*, *110*, 16277–16282.
- Shepard, R. N., & Feng, C. (1972). A chronometric study of mental paper folding. *Cognitive Psychology*, *3*, 228–243.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*, 701–703.
- Siegel, M., Buschman, T. J., & Miller, E. K. (2015). Cortical information flow during flexible sensorimotor decisions. *Science*, *348*, 1352–1355.
- Singer, W. (2013). Cortical dynamics revisited. *Trends in Cognitive Sciences*, *17*, 616–626.
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., et al. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, *23*(Suppl. 1), S208–S219.
- Sporns, O. (2014). Contributions and challenges for network models in cognitive neuroscience. *Nature Neuroscience*, *17*, 652–660.
- Tong, F. (2013). Imagery and visual working memory: One and the same? *Trends in Cognitive Sciences*, *17*, 489–490.
- Tononi, G. (2008). Consciousness as integrated information: A provisional manifesto. *Biological Bulletin*, *215*, 216–242.
- Turk-Browne, N. B. (2013). Functional interactions as big data in the human brain. *Science*, *342*, 580–584.
- Uttal, D. H., Meadow, N. G., Tipton, E., Hand, L. L., Alden, A. R., Warren, C., et al. (2013). The malleability of spatial skills: A meta-analysis of training studies. *Psychological Bulletin*, *139*, 352–402.
- Uttal, W. R. (2003). *The new pbrenology: The limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press.
- van den Heuvel, M. P., Stam, C. J., Kahn, R. S., & Hulshoff Pol, H. E. (2009). Efficiency of functional brain networks and intellectual performance. *Journal of Neuroscience*, *29*, 7619–7624.
- Vingerhoets, G., de Lange, F. P., Vandemaele, P., Deblaere, K., & Achten, E. (2002). Motor imagery in mental rotation: An fMRI study. *Neuroimage*, *17*, 1623–1633.
- Zacks, J. M. (2008). Neuroimaging studies of mental rotation: A meta-analysis and review. *Journal of Cognitive Neuroscience*, *20*, 1–19.
- Zeki, S., & Bartels, A. (1999). Toward a theory of visual consciousness. *Consciousness and Cognition*, *8*, 225–259.