Neuropsychologia 51 (2013) 1132–1141

Contents lists available at SciVerse ScienceDirect

Neuropsychologia

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

Functional connectivity in a frontoparietal network involving the dorsal anterior cingulate cortex underlies decisions to accept a hypothesis

Jennifer C. Whitman [a,b], Paul D. Metzak [a,b], Katie M. Lavigne [a,b], Todd S. Woodward [a,b,*]

[a] Department of Psychiatry, University of British Columbia, Vancouver, BC, Canada
[b] BC Mental Health and Addictions Research Institute, Vancouver, BC, Canada

ARTICLE INFO

Article history:
Received 19 July 2012
Received in revised form
24 December 2012
Accepted 24 February 2013
Available online 7 March 2013

Keywords:
Anterior cingulate cortex
Functional connectivity
Constrained principal component analysis
fMRI
Decision making
Probabilistic reasoning

ABSTRACT

Optimally interpreting our situations and experiences frequently requires comparing the evidence supporting conflicting hypotheses and deciding which to accept. This decision is comparable to an "Aha!" moment reached during insightful problem solving. We used a probabilistic reasoning task to investigate the neural activity underlying these processes. Participants rated the probability that a given focal hypothesis, rather than its alternative, was true. Decisions to accept the focal hypothesis elicited a stronger signal than decisions to reject it in a network involving the dorsal anterior cingulate cortex (dACC) and functionally connected frontal, parietal, and occipital regions. Follow-up analyses suggested that this was not simply a higher overall level of activation within the dACC or other individual regions of the network, but reflected a stronger signal for the network as a whole. This result is discussed in terms of functional connectivity between the dACC and other brain regions as a possible mechanism for coherence between components of a mental representation.

1. Introduction

Effectively judging the validity of a hypothesis is fundamental to success in many aspects of life, such as social interaction, economic decision making, and voter choice. Hypothesis judgment is also essential to more basic cognitive functions, such as memory recognition or interpreting a visual scene. Effectively judging a hypothesis often requires comparing it to at least one alternative. In order to decide whether the hypothesis being judged (the focal hypothesis) is more probable than an alternative, one must assess the strength of evidence supporting each hypothesis, compare those strengths, and decide which hypothesis to accept. In the current study, we used functional magnetic resonance imaging (fMRI) to establish a biological basis for these fundamental aspects of hypothesis judgment.

In order to investigate the functional brain networks involved in judging hypotheses, we used a probabilistic reasoning task with objectively quantifiable evidence (Beach, 1968; Moritz, Woodward, & Lambert, 2007; Speechley, Whitman, & Woodward, 2010; Whitman & Woodward, 2011, 2012). In a typical version of this paradigm, the participant is presented with an item of a given color (the relevant color) drawn from one of two containers. The participant rates the probability that it was drawn from one particular container (the focal hypothesis) rather than a second container (the alternative hypothesis). The container with the most items of the relevant color is the most probable origin of the item drawn. An advantage of this paradigm is that it allows precise control over the strength of supporting or refuting evidence.

Judging the validity of a hypothesis ultimately involves a decision to accept a focal hypothesis which is considered more coherent with the available evidence than its alternative (provided that decisions are made accurately). When cognitive representations of evidence-hypothesis information form a coherent mental construct, that construct is considered stable and salient (Köhler, 1929; Metzger, 2006), which may translate to a stronger and more stable pattern of activity in the underlying neural network. In the context of our hypothesis comparison task, such coherence and stability might result in a stronger fMRI signal for the network underlying mental representations of evidence-hypothesis matches.

The decision to accept a focal hypothesis due to sufficient coherence with the evidence could be considered a type of "Aha!" moment. To the extent that this is true, the dorsal anterior cingulate cortex (dACC) could be expected to be active under these circumstances. During insightful problem solving and difficult perceptual recognition tasks, the dACC is shown to be active in response to "Aha!" moments, that is to say, when information relevant to interpreting a problem is reorganized into

* Corresponding author at: Room A3-A116, BC Mental Health and Addictions Research Institute—Translational, Research Building, 3rd Floor, 938 W. 28th Avenue, Vancouver, British Columbia, Canada V5Z 4H4.
Tel.: +1 604 875 2000x4724; fax: +1 604 875 3871.
E-mail address: Todd.S.Woodward@gmail.com (T.S. Woodward).

0028-3932/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.neuropsychologia.2013.02.016

a gestalt, or coherent mental construct, of the solution. Other brain regions involved include the frontal eye fields (FEF), the dorsolateral prefrontal cortex (DLPFC) and parietal regions including the intraparietal sulcus (IPS) and superior and inferior parietal lobules (Aziz-Zadeh, Kaplan, & Iacoboni, 2009; Kounios & Beeman, 2009; Luo, Niki, & Phillips, 2004; Ploran et al., 2007). In the current study, we assessed functional connectivity between the dACC and these other brain regions triggered by recognition of a match between the evidence and the focal hypothesis.

We made two predictions regarding the expected functionally connected frontoparietal network involving the dACC. First, based on our view of the dACC and other functionally connected frontal and parietal regions being involved in recognizing coherence between aspects of an emerging mental construct in an "Aha!" moment, we predicted stronger activity in the underlying neural network when the focal hypothesis was accepted (because it was coherent with the evidence) than when it was rejected. Second, we expected this network to be more strongly activated in the hypothesis comparison task than the less cognitively complex evidence assessment control task, as dACC-based brain networks are known to be responsive to cognitive demands (Duncan & Owen, 2000).

2. Material and methods

2.1. Participants

Forty-six volunteers (26 females, 20 males) with a mean age of 25.0 years (SD=5.2) participated in the experiment. They received $10 per hour and were reimbursed for transportation expenses. All participants were right-handed. Participants were recruited via posters on the University of British Columbia (UBC) campus and in community centers in the greater Vancouver area, and via postings on online bulletin boards. Ethical approval was provided by the UBC Clinical Research Ethics Board. Participants were excluded from participating if they could not safely undergo an MRI scan, if they had experienced any head injuries resulting in loss of consciousness for more than 20 min, if they suffered from epilepsy, encephalitis, or meningitis, or if they or an immediate family who member suffered from a psychotic illness (e.g. schizophrenia or bipolar disorder).

2.2. Procedure

On each trial of the task, participants were presented with a scene depicting three lakes, two of which were upstream from the third (see Fig. 1). At the beginning of each trial, an animated series of images was displayed, depicting a single fish, either black or white, breaking the surface, jumping along an inverted U-shaped path (parabolic) for 140 ms, then disappearing again below the surface. We will refer to the color of this jumping fish as the *relevant color*. The color of this fish was also specified throughout the remaining duration of the trial within the question adjacent to the rating scale. This ensured that participants would be aware of the current relevant color even if they had not seen the fish jump. The populations of 100 fish in each of the two upstream lakes then became visible. Aside from the jumping fish, no other fish were ever visible in the downstream lake. The positions of the black and white fish within each lake were randomized over trials, so that any two trials with identical ratios of black to white fish would not be identical in appearance. On hypothesis comparison trials, participants were told that any fish appearing in the downstream lake originated in either the left-hand upstream lake or the right-hand upstream lake. They were required to rate the probability that the jumping fish came from one particular lake (the focal lake) rather than the alternative lake. The assignment of the left-hand and right-hand lakes as focal and alternative hypotheses was randomized across trials. On the evidence assessment trials, participants reported the percentage of fish of the relevant color in both lakes together.

Each trial allowed 6 s for the participant to make a rating. All responses were made by moving a slider up or down a vertical response scale using button presses on a LUMItouch fiber-optic response device (Lightwave Medical, Vancouver, British Columbia, Canada). The entire scale was 160 pixels in length. All responses were made with the dominant (right) hand. The two outer response buttons served to move the cursor either up or down ten pixels (the index finger moved the cursor down), while the two inner response buttons served to move the cursor either up or down two pixels (the middle finger moved the cursor down).

The Evidence Assessment and Hypothesis Comparison tasks were performed in alternating blocks of 14 trials each, with 4 blocks per functional run. There were 7 conditions within the Hypothesis Comparison task, corresponding to different percentages of fish of the relevant color in the focal and alternative lakes. If the percentage was 20% in the focal lake, it was either 20% or 50% in the alternative lake. If it was 50% in the focal lake, it was 20%, 50%, or 80% in the alternative lake. If it was 80% in the focal lake, it was either 50% or 80% in the alternative lake. Thus, the 7 conditions matched the structure of a Likert scale: the evidence strongly favored accepting the focal hypothesis (50% focal vs. 20% alternative), weakly favored it (80% focal vs. 50% alternative), was neutral (80% vs. 80%, 50% vs. 50%, or 20% vs. 20%), weakly favored rejecting the focal hypothesis (50% focal vs. 80% alternative), or strongly favored rejecting the focal hypothesis (20% focal vs. 50% alternative). There were also 7 conditions in the Evidence Assessment task that used the same visual displays, but these required the participants to rate the percentage of fish of the relevant color in both lakes combined. For each of those 14 conditions, there were 2 trials per run with an inter-trial interval (ITI) of 2 s and 2 trials per run with an ITI of 8 seconds. During the ITI, three empty lakes were displayed, i.e. without fish or a response scale. Halfway through each experimental run, a 30-s rest break occurred, during which the words "Take a 30 second break" were displayed on a dark gray screen. The total duration of each run was 740 s (370 scans).

2.3. Image acquisition

Imaging was performed at the University of British Columbia's MRI Research Centre on a Philips Achieva 3.0 T scanner with Quasar Dual gradients (with peak strength of 80 mT/m and maximum slew rate of 200 T/m/s). The participant's head was firmly secured using a custom head holder. Functional image volumes were collected with T2*-weighted gradient echo spin pulse sequences (TR=2000 ms, TE=30 ms, flip angle 90°, 36 slices, 3 mm thick, 1 mm gap, sense factor 2, matrix is 80×80 reconstructed at 128, FOV=240 mm × 240 mm × 143 mm, measured voxel is 3 mm × 3 mm × 3 mm, actual bandwidth per pixel is 53.6 Hz) effectively covering the whole brain (143 mm axial extent). Each participant completed one structural scan and two functional runs of 370 scans each.

2.4. Image preprocessing

Functional images were reconstructed offline, and the scan series was realigned and motion corrected using the method implemented in SPM5. Translation and rotation corrections did not exceed 2 mm or 2.5° for any of the participants. Parameters for spatial normalization into the modified MNI space used in SPM5 were determined using mean functional images constructed from the realigned images of each participant and scan series. The normalized functional images were smoothed with an 8 mm full width at half maximum Gaussian filter. Data were normalized to the EPI template using an affine transformation and voxels of $4\times4\times4$ mm^3. Any artifacts resulting from head movement were removed via regression, with the regressors being the six head movement parameters output by SPM5 during image realignment.

2.5. Functional connectivity analysis

To characterize how the activity of functionally connected networks differed between experimental conditions, we used a multivariate analysis technique that identifies brain regions showing temporally correlated activation (i.e., functional networks). Constrained Principal Component Analysis for fMRI (fMRI-CPCA) combines multivariate regression analysis and principal component analysis to derive networks from the portion of the BOLD signal that is explained by the timing of task events. CPCA differs from other approaches to examining correlations in activation among regions in that it (1) identifies functional networks that are based on *task related* covariance/correlation in blood-oxygen level dependent (BOLD) signal, as opposed to being based on any source of covariance/correlation in the time course, (2) estimates the hemodynamic response (HDR) for each network, and (3) quantifies the effect of experimental manipulations on each functional network.

The details of the fMRI-CPCA method are presented elsewhere (Metzak, Feredoes et al., 2011; Metzak, Riley et al., 2011; Woodward et al., 2006). Briefly, after variance in the BOLD signal attributable to the task has been separated from that not attributable to the task, the dominant patterns of inter-correlation between voxels over time are used to derive functional networks. For the comprehensive CPCA theory and proofs please see previously published work (Hunter & Takane, 2002; Takane & Hunter, 2001; Takane & Shibayama, 1991). The fMRI-CPCA application is available on-line, free of charge (www.nitrc.org/projects/fmricpca). We now briefly present matrix equations for the current application of fMRI-CPCA. This application of CPCA involved preparation of two matrices: Z and G.

2.6. Preparation of Z

The first matrix, Z, contained the BOLD time series of each voxel, with one column per voxel and one row per scan. Each column contained normalized and smoothed activations over all scans, at first for each subject separately, but then vertically concatenated to form a final Z matrix comprised of stacked individual Z matrices.

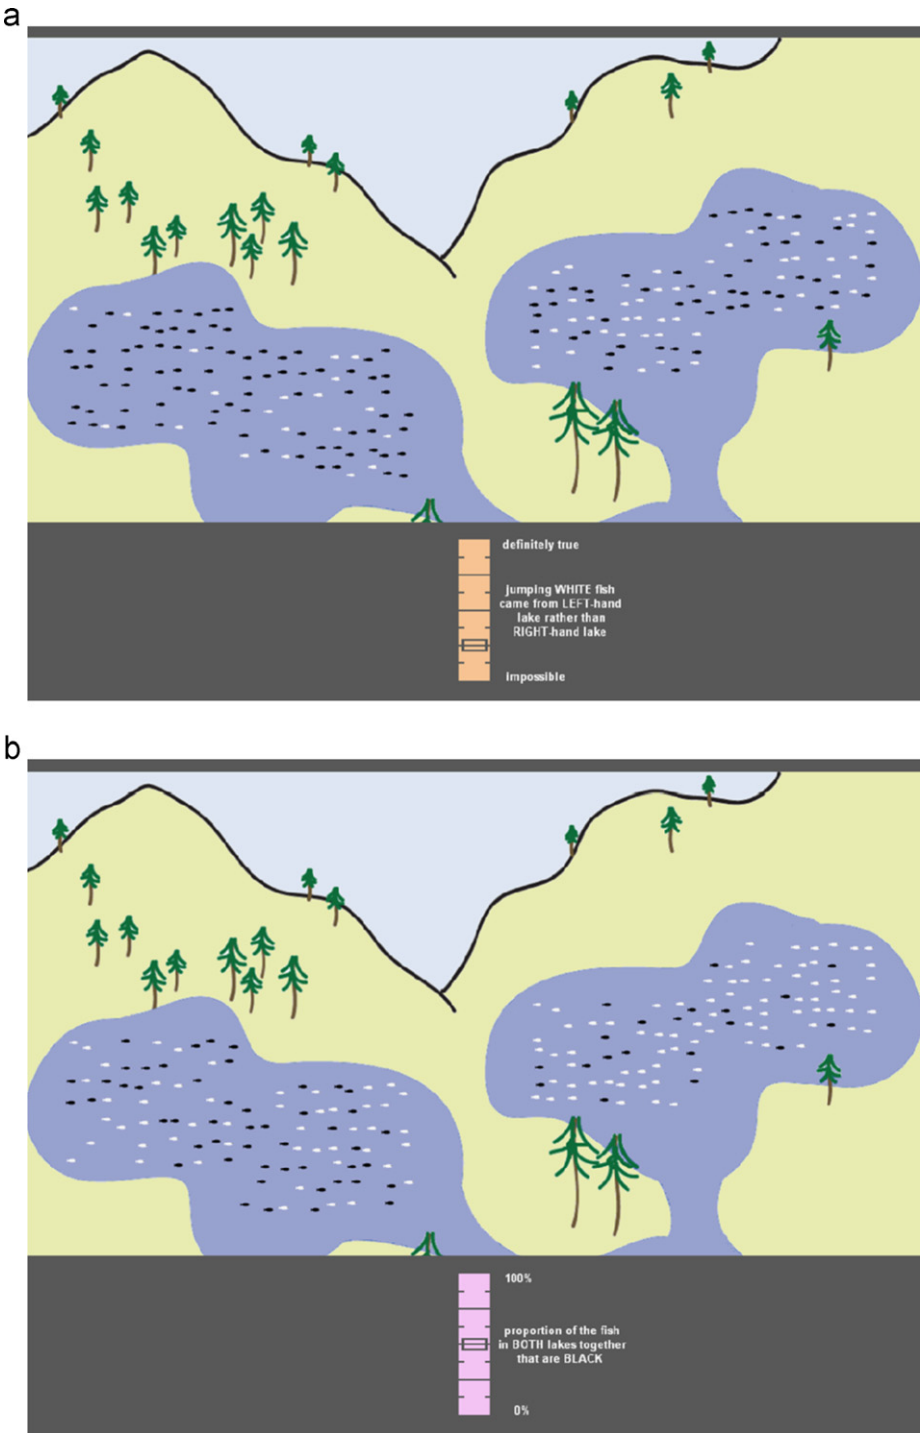

Fig. 1. The displays presented while ratings were made during (a) the hypothesis comparison task, and (b) the evidence assessment task. Each trial began with a single fish (either black or white) jumping in the downstream lake. The contents of the two upstream lakes then became visible, and a rating scale appeared in the downstream lake. On hypothesis comparison trials, participants rated the probability that the jumping fish migrated downstream from the left-hand upstream lake (the focal hypothesis) rather than the right-hand upstream lake (the alternative hypothesis). On evidence assessment control trials, participants assessed the percentage of fish in both upstream lakes together that were of the relevant color (the color of the jumping fish).

Prior to concatenation, linear and quadratic trends over the functional runs were removed from the BOLD signal (to correct for scanner drift) using multivariate multiple regression, as was variance related to head movement parameters.

2.7. Preparation of G

The second matrix, *G*, or the design matrix, contained finite impulse response (FIR) models of the expected BOLD response to the timing of stimulus presentations. Since the model of the BOLD response is applied to scans, the number of rows in *G* is equivalent to the number of rows in *Z*. FIR models estimate the increase in BOLD signal at specific post-stimulus scans relative to all other scans. The value 1 is placed in rows of *G* for which BOLD signal amplitude is to be estimated, and the value 0 in all other rows ("mini boxcar" functions). The time points for which a basis function was specified in the current study were the 1st to 10th scans following stimulus presentation. Since the repetition time (TR) for these data was 2 s, this resulted in estimating BOLD signal over a 20 s window, with the start of the first time point corresponding to encoding stimulus onset. We chose to model a 20 s window because the full HDR can occasionally last up to 20 s in some individuals and some brain regions, although the main peak occurs

mostly within the first 10 s post-stimulus (Muthukumaraswamy, Evans, Edden, Wise, & Singh, 2012; Wager, Keller, Lacey, & Jonides, 2005). We wished to be able to clearly describe the entire HDR time-course. Given that we used a FIR model rather than a modeled canonical HDR, the analysis did not force the detection of an HRF lasting a full 20 s. Rather, the shape of the time-course identified in our FIR regression analysis was data-driven.

In this analysis we created a G matrix that would allow us to estimate subject-and-condition specific effects by inserting a separate FIR basis set for each condition and for each individual subject. The columns in this subject-and-condition based G matrix code 10 post-stimulus time points, 14 experimental conditions (7 conditions within the hypothesis comparison task and 7 matched conditions within the evidence assessment control task), and 46 participants, resulting in 6440 columns ($10 \times 14 \times 46 = 6440$). In the results, the difference between accepting or rejecting the focal hypothesis did not vary as a function of whether the evidence was strong or weak, so all results are collapsed across the strong and weak levels of evidence. Trials on which the evidence for the focal and alternative hypotheses was equated were also excluded from the ANOVAs in the results section. These trials were excluded because they required button pressing in two-thirds of evidence assessment trials but not in the evidence comparison trials, and therefore would have confounded contrasts between the comparison and assessment conditions with motor activity.

2.8. Matrix equations

The matrices of the BOLD time series and experimental design are taken as input, with BOLD in Z being predicted from the FIR model in G. In order to achieve this, multivariate least-squares linear regression was carried out, whereby the BOLD time series (Z) was regressed onto the design matrix (G):

$$Z = GC + E, \tag{1}$$

where $C = (G'G)^{-1}G'Z$ using least squares regression. This analysis yielded condition-specific regression weights in the C matrix (i.e., regression weights specific to the experimental conditions as defined by the design matrix). The condition-specific regression weights are often referred to (in conventional fMRI analyses) as beta images. GC contains variability in Z that was predictable from the design matrix G, that is to say, variability in Z that was predictable from the timing of stimulus presentations. For the analysis presented here, the G matrix was standardized for each run separately.

The next step involved applying singular value decomposition to extract components representing functional networks from GC:

$$UDV' = GC, \tag{2}$$

where U is the matrix of left singular vectors; D is the diagonal matrix of singular values; V is the matrix of right singular vectors. Each column of V can be overlaid on a structural brain image to allow visualization of the neural regions involved in each functional network. In the current application of CPCA, we orthogonally rotated (Metzak et al., 2011) and rescaled the V matrix prior to display, so that a rotated *loading matrix* is displayed.

2.9. Predictor weights

To interpret the components with respect to the conditions represented in G, we produced *predictor weights* (Hunter & Takane, 2002) in matrix P. These are the weights that would be applied to each column of the matrix of predictor variables (G) to create U ($U = G \times P$). The values in P indicate the importance of each column in the G matrix to the network(s) represented by the component(s), so are essential for relating the resultant components to the experimental conditions of interest represented in G. This approach estimates a HDR shape for each individual and each condition separately.

2.10. Statistical test of component reliability and impact of experimental manipulations

As is explained above, predictor weights are produced for each combination of post-stimulus time point, condition, and participant. These weights can be used to statistically test whether the network-based BOLD response differed from zero during post-stimulus time, and to confirm that these values are reflecting a HDR shape (Metzak, Feredoes et al., 2011; Metzak, Riley et al., 2011). The impact of the experimental conditions on the network-based estimated HDR can also be tested statistically. Specifically, in this experiment, we sought to test: (1) whether the amplitude of the HDR differed as a function of whether the focal hypothesis was accepted or rejected, and (2) whether it differed between the hypothesis comparison task and the evidence assessment control task. In the first case, this would be reflected by a significant interaction between Time Point (post-stimulus time) and Decision (Accept Focal vs. Reject Focal) for the estimated network-based HDR measure (i.e., the predictor weights). Omitting the predictor weight representing the first point of post-stimulus time (which was adjusted to zero in all conditions for the purposes of display and data analysis), this analysis was carried out as a 9×2 within-subjects ANOVA for each component, with the factors of Time Point and Decision as within-subject factors. Selecting "repeated" contrasts for the within-subjects factor of Time Point allowed significance tests to be restricted to adjacent time points, such that complex 9×2 interactions (e.g., between Time Point and Decision) were broken down into 8 different 2×2 interactions involving adjacent levels of the Time Point. Inspection of the relative size of the p values for these 8 different 2×2 interactions can pinpoint the time points responsible for the 9×2 interactions (e.g., the Accept vs. Reject pairwise comparison increases significantly from the 5 s to 7 s post-stimulus time points). Tests of sphericity were carried out for all ANOVAs. Greenhouse–Geisser adjusted degrees of freedom are reported where violations of sphericity affected the interpretation of results; otherwise, the original degrees of freedom are reported. Since our significance testing is carried out at the level of subject-specific HDR estimates, use of bootstrapping to produce Z-map images is not required; therefore, point estimates (from orthogonally rotating and rescaling the V matrix of component loadings) are overlaid on structural brain images for depiction of the spatial arrangement of the functional networks.

2.11. Follow-up analysis of contributions by individual clusters to a network time course

In one functionally connected network for which the strength of the HDR differed between experimental conditions, we performed a follow-up analysis to examine how different regions of the network contributed to that effect. This clarified whether the difference between conditions in that network was simply due to a higher overall level of activation within any single region of the network (e.g., the dACC), or reflected a stronger signal for the functionally connected network as a whole. This involved exploring whether each region of the network showed stronger signal during one condition than during another. It was achieved as follows. We restricted the analysis to voxels with the most dominant 20% of component loadings (the values overlaid on the brain images, rotated V in Eq. (2)). For each cluster visible in the network, we multiplied those component loadings for voxels within the cluster of interest by the corresponding predicted scores (GC in Eq. (1)) to create cluster-specific component scores (one score per functional scan, analogous to rotated U in Eq. (2)). These were then correlated with the model (G in Eq. (1)) to estimate cluster-specific predictor weights. ANOVAs performed on these predictor weights provided a post-hoc description of the significant effects identified in the full-network analysis.

2.12. Follow-up analysis of contributions by peak and off-peak voxels to a network time course

In addition to exploring how individual clusters within a network contributed to the network time course, we also assessed the contributions of peak and off-peak voxels. In this analysis, we included voxels from every cluster within the network, but adjusted the percentage cutoff for the most dominant loadings. We first estimated predictor weights including the top 20% of component loadings for each cluster. The percentage for this cutoff was then lowered in increments of 5% to include more voxels, with the analysis being repeated at each cutoff, until a cutoff was reached for which the effect was statistically significant. Note that the initial analysis on full-network predictor weights, described in the sections above, reflects a percentage cutoff of 100% (i.e., all brain voxels were included, and the contribution of each voxel was weighted by how strongly it loaded onto the component of interest).

3. Results

Behavioral responses are displayed in Table 1. Most of the ratings in the hypothesis comparison task and percentage estimates in the evidence assessment control task were quite close to the mathematically normative ratings. The exceptions were the ratings made when the focal and alternative lakes contained 20% and 50% fish of the relevant color, or vice versa. These ratings were significantly closer to the mid-point of the rating scale than were the mathematically normative ratings. The magnitude of the deviation from the mathematically normative rating did not depend, $t(45) = 1.58$, *ns*, on whether the evidence favored the focal hypothesis, $t(45) = 5.58$, $p < 0.001$, or its alternative $t(45) = 4.04$, $p < 0.001$. Participants were simply more cautious than necessary when presented with strong evidence, regardless of which hypothesis that evidence favored. As mentioned above, for all analyses of fMRI data, only trials on which the participants moved the cursor on the response scale were included.

Table 1
Ratings of the relative probability that the focal hypothesis rather than its alternative is true, expressed as a percentage of the total height of the rating scale.

Percentage of fish of the relevant color		Hypothesis comparison task		Evidence assessment (control) task	
Focal lake (f)	Alternative lake (a)	Probability rating: mean (SD)	Mathematical norm: $100 \times f/(f+a)$	Percentage estimate: mean (SD)	Mathematical norm: $(f+a)/2$
20	20	48.2 (4.2)	50.0	19.7 (10.7)	20.0
20	50	34.8 (10.5)*	28.6	34.1 (6.0)	35.0
50	20	63.8 (9.3)*	71.4	34.2 (6.4)	35.0
50	50	50.5 (3.4)	50.0	51.0 (6.0)	50.0
50	80	39.1 (11.1)	38.5	66.0 (8.2)	65.0
80	50	65.4 (11.0)	61.5	66.5 (6.4)	65.0
80	80	51.8 (3.38)	50.0	79.8 (6.2)	80.0

* Indicates a rating deviating significantly from the mathematical norm at $p < 0.001$, which remains significant after a Bonferroni correction for multiple comparisons.

Table 2
Cluster volumes for most extreme 20% of Component 1 loadings, with anatomical descriptions, MNI coordinates, and Brodmann's area for the peaks within each cluster. Only clusters > 25 mm^3 are presented here. All of the loadings for this component were positive.

Cortical regions	**Cluster volume (mm^3)**	**Brodmann's area for peak locations**	**MNI coordinate (*X Y Z*) for peak locations**		
Cluster 1 (bilateral)	308,032				
Superior parietal		7	32	−60	52
Middle occipital		18	−32	−96	0
Fusiform gyrus		18	−28	−88	−16
Superior parietal lobe		7	−24	−68	52
Precuneus		7	28	−72	44
Inferior occipital gyrus		18	32	−92	−12
Cingulate gyrus/dACC/supplementary motor area		32	4	16	48
Lingual gyrus		17	−12	−100	−8
Middle occipital gyrus		19	36	−88	12
Middle frontal gyrus		6	32	0	60
Fusiform gyrus		19	−44	−76	−16
Cerebellum (right)		n/a	28	−68	−20
Lateral occipital		19	32	−80	28
Superior parietal lobule		40	−36	−48	48
Lateral occipital gyrus		19	−28	−80	28
Lateral occipital gyrus		37	48	−64	−16
Cerebellum (left)		n/a	36	−56	−28
Supramarginal gyrus		40	52	−36	48
Precentral gyrus		9	−52	8	32
Precentral gyrus		4	−36	−20	64
Superior frontal gyrus/supplementary motor cortex		6	−4	0	72
Cluster 2:	7168				
Right precentral gyrus		44	56	12	32
Cluster 3:	5632				
Right superior frontal gyrus/frontal pole		46	40	44	28
Cluster 4:	960				
Right thalamus		n/a	12	−16	8
Cluster 5:	768				
Left middle frontal gyrus		9	−44	28	28
Cluster 6:	320				
Right insula/orbitofrontal cortex		47	32	20	0
Cluster 7:	256				
Right Cerebellum		n/a	24	−40	−48
Cluster 8:	64				
Left temporal pole		22	−52	12	−8
Cluster 9:	64				
Right thalamus		n/a	12	−4	12

This allowed us to match the hypothesis comparison and evidence assessment tasks in terms of motor activity.

The scree plot of singular values sorted by magnitude suggested extraction of two components (corresponding to two functional networks). The predictor weights associated with each component followed a time course reflecting an HDR shape. Correspondingly, a significant main effect of Time Point was present for each component. The percentages of task-related variance accounted for by the first and second rotated components were 19.3% and 9.8%, respectively. The brain regions comprising each functional network are portrayed in Tables 2 and 3 and Fig. 2. Note that only the strongest 20% of loadings are displayed in those figures, for the sake of clearly displaying the dominant features of the spatial pattern of each component. However, all of the predictor weights analyzed below represent a weighted aggregate of activity across 100% of brain voxels, with the only exception being the paragraph

Table 3
Cluster volumes for most extreme 20% of Component 2 loadings, with anatomical descriptions, MNI coordinates, and Brodmann's area for the peaks within each cluster. Only clusters > 25 mm^3 are presented here. Positive and negative loadings are presented in the top and bottom sections of the table, respectively.

Cortical regions	Cluster volume (mm^3)	Brodmann's area for peak locations	MNI coordinate (*X Y Z*) for peak locations		
Positive loadings					
Cluster 1:	7808				
Left occipital pole		17	−28	−100	0
Cluster 2:	7680				
Right occipital pole		18	28	−96	0
Cluster 3:	2624				
Superior parietal lobule		7	28	−64	48
Cluster 4:	768				
Left superior & inferior parietal lobules		40	−44	−40	44
Cluster 5:	704				
Left superior & inferior parietal lobules		7	−24	−68	48
Cluster 6:	64				
Left middle & superior frontal gyri		6	−28	−4	52
Negative loadings					
Cluster 1 (bilateral):	94,784				
Ventromedial prefrontal cortex		11	0	52	−12
Ventromedial Prefrontal Cortex		32	−4	48	0
Paracingulate gyrus		10	0	52	4
Middle & superior frontal gyri		32	−24	28	44
Middle & superior frontal gyri		9	−20	36	44
Superior frontal gyrus		8	−4	40	52
Superior frontal gyrus		44	−40	16	48
Cluster 2 (right hemisphere):	77,696				
Superior temporal gyrus		22	56	−60	16
Middle temporal gyrus		21	60	−8	−24
Orbitofrontal cortex/inferior frontal gyrus		47	32	32	−20
Orbitofrontal cortex/inferior frontal gyrus		45	52	32	−4
Superior temporal gyrus		48	60	−4	−4
Temporal pole		38	36	24	−28
Temporal pole		20	36	16	−44
Temporal pole		36	32	4	−24
Parietal operculum		48	48	−32	20
Parahippocampal gyrus		20	28	−8	−24
Hippocampus		36	28	−8	−24
Cluster 3 (left hemisphere):	58,048				
Middle temporal gyrus & lateral occipital cortex		39	−48	−76	32
Postcentral gyrus		20	−60	−12	20
Orbitofrontal cortex		47	−28	32	−16
Inferior temporal gyrus		37	−64	−52	−8
Orbitofrontal cortex		38	−48	24	−12
Temporal pole		21	−48	8	−36
Inferior frontal gyrus		45	−52	24	12
Inferior frontal gyrus		47	−28	12	−24
Parahippocampal gyrus		38	−32	4	−20
Cluster 4 (bilateral):	53,056				
Precuneus/posterior cingulate		4	40	−24	60
Cluster 5:	8152				
Right postcentral gyrus		4	40	−24	60
Cluster 6:	4928				
Right cerebellum		n/a	24	−84	−36
Cluster 7:	3200				
Left parahippocampal gyrus		37	−32	−32	−20
Left hippocampus		30	−24	−20	−20
Left parahippocampal gyrus & left hippocampus		36	−20	−8	−28
Cluster 8:	3136				
Left cerebellum		n/a	−24	−84	−36
Cluster 9:	768				
Right subcallosal cortex		25	4	12	−12
Cluster 10:	128				
Left insula		20	−40	−12	−12
Cluster 11:	128				
Right precentral gyrus		6	12	−20	76
Cluster 12:	64				
Right precentral gyrus		4	1	−24	68

describing our follow-up analysis, in which we explored how predictor weight time courses changed as they were calculated for the voxels with the top 5% of loadings, those with the top 10%, and so forth. Component 1 involved activation in a network involving the dACC and bilateral parietal and lateral occipital regions. Component 2 involved deactivation in the precuneus, ventromedial prefrontal cortex, and bilateral middle temporal gyri; areas which overlap substantially with the default mode network (Buckner, Andrews-Hanna, & Schacter, 2008; Fox et al., 2005), as well as activation in lateral occipital clusters. Note that, as Component 2 is dominated by negative loadings, a strongly positive predictor weight for Component 2 mostly reflects

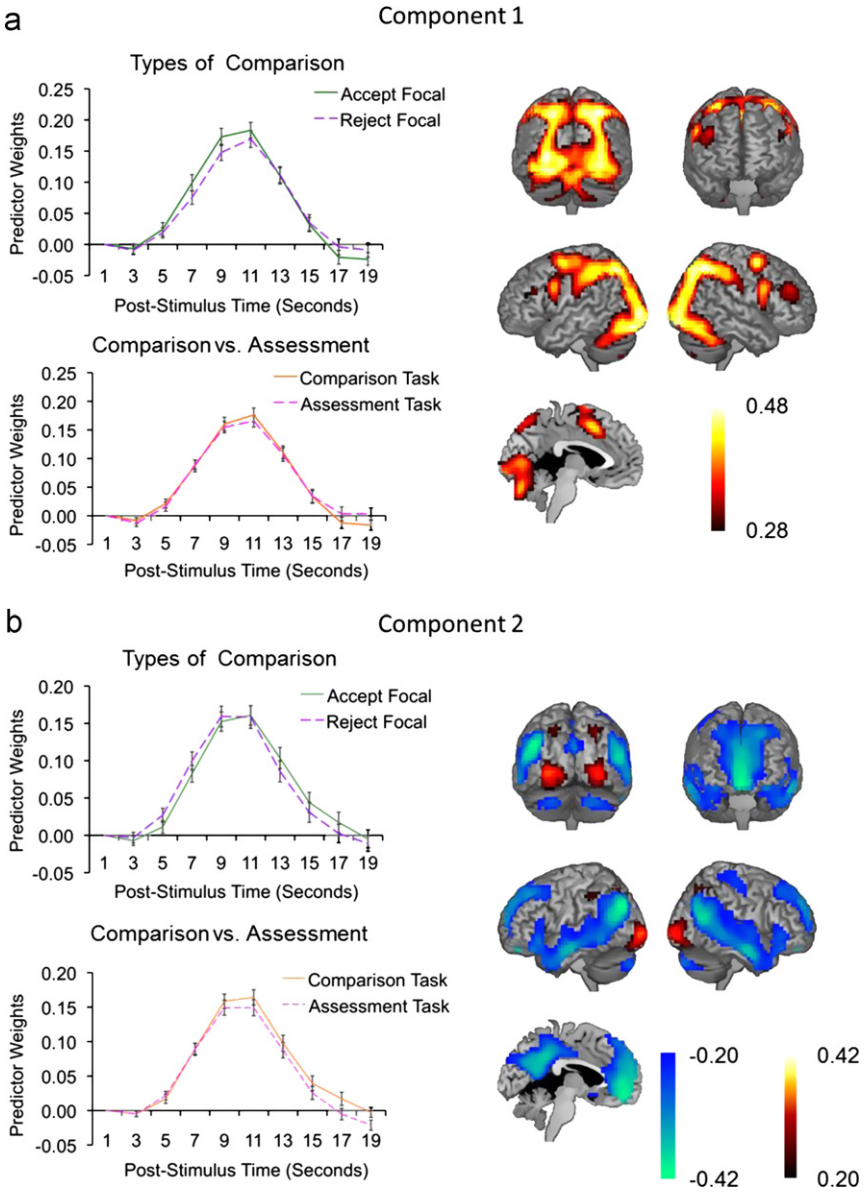

Fig. 2. Functional networks identified by CPCA, with the strongest 20% of component loadings shown. Note that predictor weight time courses are based on a whole-brain network, although only the top 20% are displayed here. Component 1 (a) consisted of activations (positive loadings) within a dACC-based network. This network was more active when the evidence supported accepting the focal hypothesis than when it supported rejecting it. Component 2 (b) was dominated by deactivations (negative loadings) in regions widely reported to belong to the default mode network, but also included activations (positive loadings) in two parietal regions and two lateral occipital regions. Note that predictor weights can be interpreted as correlations, so a strongly positive predictor weight in the time course for Component 2 indicates that the pattern of deactivations in regions of the default mode network is more intense for that post-stimulus time point and experimental condition.

activation reductions. Specifically, it reflects an aggregate of deactivations in the large number of voxels with negative loadings and concurrent activations in the small number of voxels with positive loadings. Thus, a strongly positive predictor weight for a given condition and post-stimulus time point reflects strong activation reduction of the default mode network.

3.1. Accept focal vs. reject focal

In order to investigate how each component varied as a function of acceptance or rejection of the focal hypothesis, we performed 2×9 ANOVAs with factors of Decision (Accept Focal vs. Reject Focal) and Time Point. A Decision $\times$ Time Point ANOVA carried out on the Component 1 predictor weights revealed no main effect of Decision, $F(1,45)=0.15$, *ns*, a significant main effect of Time Point, $F(8,360)=96.27$, $p<0.001$, and a significant Decision $\times$ Time Point interaction, $F(8,360)=3.23$, $p=0.001$. The post-hoc analysis examining this significant interaction using only adjacent time points indicated that the pair-wise contrast of Accept > Reject increased substantially from 5 to 7 s post-stimulus ($p=0.01$), reflecting a higher HDR peak in the Accept Focal condition relative to the Reject Focal condition.

A Decision $\times$ Time Point ANOVA carried out on the Component 2 predictor weights revealed no main effect of Decision, $F(1,45)=0.02$, *ns*, a significant main effect of Time Point, $F(8,360)=89.37$, $p<0.001$, and a significant Decision $\times$ Time Point interaction $F(8,360)=2.50$, $p<0.05$. The follow-up analysis examining this significant interaction using only adjacent time points indicated that the pair-wise contrast of Accept > Reject increased substantially from 11 to 13 s post-stimulus ($p=0.05$), reflecting a slightly right-shifted HDR peak in the Accept Focal condition relative to the Reject Focal condition.

When Components 1 and 2 were entered into a combined analysis, there was a significant Decision $\times$ Component

interaction, $F(1,45)=4.24$, $p<0.05$, reflecting the fact that the mean difference in predictor weights between the Accept Focal and Reject Focal conditions was greater for Component 1 ($M=0.017$) than for Component 2 ($M=0.001$). As can be seen in Fig. 2, the predictor weight time course for Component 1 (the dACC-based network) reached higher amplitude in the Accept Focal (evidence–hypothesis match) condition than in the Reject Focal (evidence–hypothesis mismatch) condition. In contrast, for Component 2 (the default-mode network) there was no clear amplitude difference, but there was a slightly right-shifted HDR peak in the Accept Focal condition.

A follow-up analysis was carried out to clarify whether the difference between focal hypothesis acceptance and rejection in the dACC-based network was simply due to a higher overall level of activation within any single region of the network (e.g., the dACC), or reflected a stronger signal for the functionally connected network as a whole. This involved exploring whether each region of the dACC-based network showed stronger signal during focal hypothesis acceptance than rejection. Cluster-specific predictor weights were estimated as described in the Methods section. We first performed this analysis including only voxels corresponding to the top 20% of component loadings, as are displayed in Fig. 2A, for each individual cluster. The Decision × Time Point interaction was not significant for any individual cluster, including the dACC, indicating that the significant interaction reported above is not specific to any one particular brain region within the network. We also performed this analysis at the 20% cutoff for all displayed clusters simultaneously and found no significant interaction. We then carried out this analysis including all voxels that passed a threshold increasing by 5% increments (i.e., 20%, 25%, 30%, etc.) of the dominant loadings. The Decision × Time Point interaction reached significance, $F(8,360)=2.42$, $p=0.05$, once the threshold was set at the dominant 45% of loadings. No new clusters emerged at this threshold; rather, the newly included voxels formed concentric rings around the existing clusters. It was not possible to create cluster-specific predictor weights at the 45% cutoff because as the clusters grew in size their edges met, resulting in one contiguous cluster identified per cerebral hemisphere.

The above results suggest that a pattern spanning much of the brain, not just the peaks displayed in Fig. 2, is important to representing the difference between conditions. It should be noted that the spatial patterns of positive and negative peaks were the same for the two components. The difference between them lay in the off-peak voxels. In Component 1 (the dACC-based network), 95.76% of voxels had positive loadings. In Component 2 (the default mode network), 80.23% of voxels had negative loadings. In sum, the off-peak voxels accounted for both the differences in spatial pattern between brain networks and also for the differences in cognitive processing. These analyses are congruent with an account holding that the stronger signal during focal hypothesis acceptance reflects coordinated activity across a network of brain regions, and emphasizes the importance of the contribution of non-peak brain regions to this effect (i.e., the shape of the cluster as well as the peak amplitude and location).

3.2. Hypothesis comparison vs. evidence assessment

In order to compare component activity in the Hypothesis Comparison task to that in the Evidence Assessment task, we used 2 × 9 ANOVAs with factors of Task (Hypothesis Comparison vs. Evidence Assessment) and Time Point (9 time points displaying how the BOLD response progressed over 2–20 s after stimulus presentation). An ANOVA carried out on the Component 1 predictor weights revealed no main effect of Task, $F(1,45)=0.01$, *ns*, a significant main effect of Time Point, $F(8,360)=122.93$, $p<0.001$, and no Task × Time Point interaction, $F(3.71,167.14)=2.08$, *ns*. A Task × Time Point ANOVA carried out on the Component 2 predictor weights revealed no main effect of Task, $F(1,45)=0.51$, *ns*, a significant main effect of Time Point, $F(8,360)=119.09$, $p<0.001$, and no Task × Time Point interaction, $F(8,360)=1.16$, *ns*. As can be seen in Fig. 2A, the time courses of the two tasks overlap very closely on Components 1 and 2.

3.3. Follow-up analysis testing for lateralization effects

The tasks used in this study involved randomization designed to ensure that the left and right lakes corresponded to the focal hypothesis equally often. However, we felt it prudent to confirm empirically that no left/right imbalances occurred. To this purpose, we compared the cluster-specific predictor weights between left and right hemisphere parietal-occipital clusters. There were no main effects of hemisphere and no interactions of hemisphere with any other factors of interest in either the Hypothesis Comparison task or the Evidence Assessment control task.

4. Discussion

In order to understand the neural basis of comparing evidence for conflicting hypotheses, we used a probabilistic reasoning paradigm to investigate functional brain networks engaged in deciding which hypothesis to accept. As these decisions are comparable to "Aha!" moments, when coherence between the correct interpretation and the available evidence is recognized, we expected a dACC-based network to be involved. As predicted, we found more activity within a dACC-based network when the focal hypothesis was accepted (an evidence–hypothesis match) than when it was rejected (an evidence–hypothesis mismatch). This is consistent with a role for the dACC and its associated functional network in recognizing a coherent mental construct, or triggering recognition of a match between the evidence and the focal hypothesis. However, the results emphasize that the entire dACC-based network is associated with this effect, not only any individual region. There was also a delayed BOLD response in the default-mode network for hypothesis acceptance, but no overall difference in amplitude.

The CPCA method was optimal for identifying the above pattern of increased activity in the dACC-based network during focal hypothesis acceptance. This is primarily because fmri-CPCA is designed to identify responses to experimental manipulations of task conditions, rather than identifying spontaneous activity. It is the initial regression phase of CPCA that sets it apart from other component extraction methods and makes it advantageous for identifying task-related functional network changes. This regression constrains the analysis to the small portion variance in brain activity attributable to task performance. This ensures that the results of the subsequent component extraction step are dominated by task-related activity rather than spontaneous activity. This order of operations is particularly important in cases where the spatial configuration of a network changes in response to cognitive demands. While CPCA uses regression to predict brain activity (the criterion) from task timing (the predictor), thereby constraining spatial network configurations to variance predictable from cognitive demands. This is advantageous relative to methods that treat both of these as criterion variables and searches for common factors, such as partial least squares (PLS; McIntosh, Chau, & Protzner, 2004), which treats both of these as criterion variables and searches for common factors (Metzak, 2011). An advantage of CPCA over seed-based connectivity methods is that the lack of a requirement to select a seed region

makes the analysis more data-driven. Furthermore, PCA identifies the dominant pattern of intercorrelations in a dataset in a more efficient manner than performing a large number of pairwise correlation analyses. In sum, the results of this study demonstrate the efficacy of constraining an analysis to variance attributable to manipulations of interest prior to employing component analysis.

Although we focus here on the role of the dACC in our task, the role of the other regions in the network merits some discussion. The dACC-based network also included bilateral activity in the precentral gyrus. This was stronger in the left hemisphere, which is consistent with the use of the right hand to respond. We also found strong activity in bilateral parietal and lateral occipital clusters, which is consistent with the use of large visual displays in our paradigm. Finally, we found weak right-lateralized DLPFC activity. As other reports of the role of the DLPFC in perceptual decision-making indicate that the left DLPFC plays the dominant role (Heekeren, Marrett, & Ungerleider, 2008; Philiastides, Auksztulewicz, Heekeren, & Blankenburg, 2011; Rahnev, Lau, & de Lange, 2011), our DLPFC activity is unlikely to reflect the same function. As can be seen in Fig. 2, it also exhibits noticeably weaker loadings than the dACC or other clusters in Component 1. The network formed by all of these regions represents a variant of the 'Task-Positive Network' widely reported to be involved in spatial attention, working memory, and a wide variety of other tasks (Fox, Zhang, Snyder, & Raichle, 2009). In the current study we focus on the role of the dACC in this network because it exhibits the strongest loadings outside of visual and motor areas.

Our follow-up analyses suggested that the stronger signal from the dACC-based network in response to evidence–hypothesis matches (Accept Focal) depended on an interconnected network-wide pattern. In other words, it depended on more cortical regions than only the peak locations and their amplitude. These analyses indicate that the stronger signal during focal hypothesis acceptance reflects coordinated activity across a network of brain regions, rather than increased activity within any individual cluster peak or subset of clusters peaks. They also suggest that activity throughout each network cluster (e.g., cluster shape), rather than only at the peaks, underlies the stronger signal in response to evidence–hypothesis matches. This demonstrates the importance of a functional connectivity approach to fMRI data analysis based on identifying network patterns, in that it may be able to detect condition differences that could be missed with univariate analyses.

A secondary goal was to investigate activity in the hypothesis comparison task over and above that in the evidence assessment control task. As the hypothesis comparison task involves more cognitive steps than the evidence assessment control task, we expected it to recruit the dACC-based network more strongly, given that the dACC-based network is known to be responsive to cognitive demands (Duncan & Owen, 2000). Contrary to our predictions, we found no between-condition differences in network activity. Assuming that a number of cognitive operations do differ between these conditions, we can conclude that our methods were not sensitive to them. That is not to say that differences in brain activity do not exist, as certain cognitive operations clearly present in the experimental condition are absent in the tightly matched control condition. One possibility is that the BOLD signal may simply be too temporally coarse to detect these differences. The hemodynamic response to any single cognitive event takes several seconds to reach its peak (Boynton, Engel, Glover, & Heeger, 1996; Friston et al., 1998). Therefore the responses to adjacent rapid cognitive events involved in complex cognitive processes, such as those differing between evidence assessment and comparison, may be merged in the measured BOLD signal Whitman, Ward & Woodward (2013). Consequently, the fMRI modality may not be sensitive to subtle differences between two series of complex cognitive events, such as those involved in evidence assessment and comparison here.

This study is subject to a number of limitations. As the results are dependent on the analysis of functional networks, they may not replicate across other non-connectivity-based analysis methods. Moreover, it may be that other functional brain networks accounting for smaller portions of variance in the BOLD signal also play a role in hypothesis comparison, but were not detected here. Furthermore, the results depict a correlational relationship between differences in network activity and the experimental manipulations of interest, leaving the question of causality open. Finally, the data are limited by the low temporal resolution of fMRI and slow time course of the BOLD response, and other modalities sensitive to more precise time scales (such as MEG) may be beneficial in this regard Whitman, Ward & Woodward (2013).

The results of the current experiment help to clarify the precise role of a dACC-based network in the recognition of the "Aha!" moment, or in triggering recognition of a match between the evidence and the focal hypothesis. A number of other roles have been attributed to the dACC, such as monitoring the environment for conditions that may require adjustments in control over the course of action (Behrens, Woolrich, Walton, & Rushworth, 2007; Paus, 2001; Woodward, Metzak, Meier, & Holroyd, 2008), detection of conflict (Walsh, Buonocore, Carter, & Mangun, 2011), detection of errors (Carter et al., 1998), or detection of surprise (Egner, 2011; Egner & Hirsch, 2005). The moment a decision is made (i.e., an insightful "Aha!" moment) a given pattern of coherence between aspects of a mental representation is recognized. Thus, these data support a general role for a dACC-based network in recognition that a change in mental set is required (Woodward, Metzak, Meier, & Holroyd, 2008).

Optimally interpreting our situations and experiences frequently requires comparing the evidence supporting conflicting hypotheses and deciding which to accept, with the final decision stage being comparable to an "Aha!" moment reached during insightful problem solving. The results suggested that this involves a stronger signal for a dACC-based network as a whole, and that functional connectivity between the dACC and other brain regions is a possible mechanism for coherence between components of a mental representation. This helps clarify the role of the dACC in the wide variety of tasks which involve judging and comparing hypotheses, such as social interaction, economic decision making, voter choice, perceptual decisions, and evaluating scientific research.

References

Aziz-Zadeh, L., Kaplan, J. T., & Iacoboni, M. (2009). "Aha!": The neural correlates of verbal insight solutions. *Human Brain Mapping*, *30*(3), 908–916.

Beach, L. R. (1968). Probability magnitudes and conservative revision of subjective probabilities. *Journal of Experimental Psychology*, *77*(1), 57–63.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221.

Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, *16*(13), 4207–4221.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network—Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1–38.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*(5364), 747–749.

Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, *23*(10), 475–483.

Egner, T. (2011). Surprise! A unifying model of dorsal anterior cingulate function? *Nature Neuroscience*, *14*(10), 1219–1220.

Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience*, *8*(12), 1784–1790.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678.

Fox, M. D., Zhang, D., Snyder, A. Z., & Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of Neurophysiology*, *101*, 3270–3283.

Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D., & Turner, R. (1998). Event-related fMRI: Characterizing differential responses. *Neuroimage*, *7*(1), 30–40.

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, *9*, 407–479.

Hunter, M. A., & Takane, Y. (2002). Constrained principal component analysis: Various applications. *Journal of Educational and Behavioral Statistics*, *27*(2), 105–145.

Köhler, W. (1929). *Gestalt psychology*. New York: H. Liveright.

Kounios, J., & Beeman, M. (2009). The aha! moment: The cognitive neuroscience of insight. *Current Directions in Psychological Science*, *18*(4), 210–216.

Luo, J., Niki, K., & Phillips, S. (2004). Neural correlates of the 'Aha! reaction'. *Neuroreport*, *15*(13), 2013–2017.

McIntosh, A. R., Chau, W. K., & Protzner, A. B. (2004). Spatiotemporal analysis of event-related fMRI data using partial least squares. *Neuroimage*, *23*(2), 764–775.

Metzak, P. D., Feredoes, E., Takane, Y., Wang, L., Weinstein, S., Cairo, T., et al. (2011). Constrained principal component analysis reveals functionally connected load-dependent networks involved in multiple stages of working memory. *Human Brain Mapping*, *32*(6), 856–871.

Metzak, P. D., Riley, J., Wang, L., Whitman, J. C., Ngan, E. T. C., & Woodward, T. S. (2011). Decreased efficiency of task-positive and task-negative networks during working memory in schizophrenia. *Schizophrenia Bulletin*, *38*(4), 803–813.

Moritz, S., Woodward, T. S., & Lambert, M. (2007). Under what circumstances to patients with schizophrenia jump to conclusions? A liberal acceptance account. *British Journal of Clinical Psychology*, *46*, 127–137.

Muthukumaraswamy, S. D., Evans, C. J., Edden, R. A., Wise, R. G., & Singh, K. D. (2012). Individual variability in the shape and amplitude of the BOLD-HRF correlates with endogenous GABAergic inhibition. *Human Brain Mapping*, *33*(2), 455–465.

Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognition interface. *Nature Reviews Neuroscience*, *2*(6), 417–424.

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology*, *21*, 980–983.

Ploran, E. J., Nelson, S. M., Velanova, K., Donaldson, D. I., Petersen, S. E., & Wheeler, M. E. (2007). Evidence accumulation and the moment of recognition: Dissociating perceptual recognition processes using fMRI. *Journal of Neuroscience*, *27*(44), 11912–11924.

Rahnev, D., Lau, H., & de Lange, F. P. (2011). Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *The Journal of Neuroscience*, *31*(29), 10741–10748.

Speechley, W. J., Whitman, J. C., & Woodward, T. S. (2010). The contribution of hypersalience to the "jumping to conclusions" bias associated with delusions in schizophrenia. *Journal of Psychiatry & Neuroscience*, *35*(1), 7–17.

Takane, Y., & Hunter, M. A. (2001). Constrained principal component analysis: A comprehensive theory. *Applicable Algebra in Engineering, Communication and Computing*, *12*, 391–419.

Takane, Y., & Shibayama, T. (1991). Principal component analysis with external information on both subjects and variables. *Psychometrika*, *56*(1), 97–120.

Wager, T. D., Keller, M. C., Lacey, S. C., & Jonides, J. (2005). Increased sensitivity in neuroimaging analyses using robust regression. *Neuroimage*, *26*(1), 99–113.

Walsh, B. J., Buonocore, M. H., Carter, C. S., & Mangun, G. R. (2011). Integrating conflict detection and attentional control mechanisms. *Journal of Cognitive Neuroscience*, *23*(9), 2211–2221.

Whitman, J. C., & Woodward, T. S. (2011). Evidence affects hypothesis judgments more if accumulated gradually than if presented instantaneously. *Psychonomic Bulletin & Review*, *18*(6), 1156–1165.

Whitman, J. C., & Woodward, T. S. (2012). Self-selection bias in hypothesis comparison. *Organizational Behavior and Human Decision Processes*, *118*(2), 216–225.

Whitman, J.C., Ward, L.M., & Woodward, T.S. (2013). Patterns of cortical oscillations organize neural activity into whole-brain functional networks evident in the fMRI BOLD signal. *Frontiers in Human Neuroscience*, 7, 80.

Woodward, T. S., Cairo, T. A., Ruff, C. C., Takane, Y., Hunter, M. A., & Ngan, E. T. C. (2006). Functional connectivity reveals load dependent neural systems underlying encoding and maintenance in verbal working memory. *Neuroscience*, *139*(1), 317–325.

Woodward, T. S., Metzak, P. D., Meier, B., & Holroyd, C. B. (2008). Anterior cingulate cortex signals the requirement to break inertia when switching tasks: A study of the bivalency effect. *Neuroimage*, *40*(3), 1311–1318.