Cortical networks involved in visual awareness independent of visual attention

Taylor W. Webb[a], Kajsa M. Igelström[a], Aaron Schurger[b], and Michael S. A. Graziano[a,1]

[a]Department of Psychology, Princeton University, Princeton, NJ 08544; and [b]Cognitive Neuroimaging Unit, NeuroSpin Research Center, Commissariat a l'Energie Atomique (CEA)-Saclay, 91191 Gif-sur-Yvette, France

Edited by Michael E. Goldberg, Columbia University College of Physicians, New York, NY, and approved October 19, 2016 (received for review July 13, 2016)

It is now well established that visual attention, as measured with standard spatial attention tasks, and visual awareness, as measured by report, can be dissociated. It is possible to attend to a stimulus with no reported awareness of the stimulus. We used a behavioral paradigm in which people were aware of a stimulus in one condition and unaware of it in another condition, but the stimulus drew a similar amount of spatial attention in both conditions. The paradigm allowed us to test for brain regions active in association with awareness independent of level of attention. Participants performed the task in an MRI scanner. We looked for brain regions that were more active in the aware than the unaware trials. The largest cluster of activity was obtained in the temporoparietal junction (TPJ) bilaterally. Local independent component analysis (ICA) revealed that this activity contained three distinct, but overlapping, components: a bilateral, anterior component; a left dorsal component; and a right dorsal component. These components had brain-wide functional connectivity that partially overlapped the ventral attention network and the frontoparietal control network. In contrast, no significant activity in association with awareness was found in the banks of the intraparietal sulcus, a region connected to the dorsal attention network and traditionally associated with attention control. These results show the importance of separating awareness and attention when testing for cortical substrates. They are also consistent with a recent proposal that awareness is associated with ventral attention areas, especially in the TPJ.

awareness | consciousness | attention | temporoparietal junction | TPJ

A major goal of the scientific study of consciousness is to understand which brain regions form the substrate of subjective awareness. This goal has often been approached by comparing the effects of a stimulus on the brain when subjects are aware vs. unaware of the stimulus. A number of elegant paradigms have been designed for this purpose, with the aim of making the stimulus as similar as possible in the two conditions (1, 2), such that any neural differences that emerge are more reasonably attributed to a difference in awareness than to a difference in the stimuli themselves. Using this approach, studies have pointed to a frontoparietal network, including the dorsolateral prefrontal cortex, the middle and inferior frontal gyri, the intraparietal sulcus, the superior parietal lobule, and a number of regions within the temporoparietal junction (TPJ) (3–9), although some questions remain regarding the specific role of each of these regions (9–11).

One potential concern is the conflation of attention and awareness. During the last 15 y, it has become well established that awareness and attention can be separated. It is possible for people to attend to a visual stimulus, as measured by standard attention tasks such as the spatial Posner cuing task, while reporting no subjective awareness of the stimulus (12–15). Attention and awareness are, however, closely linked. They covary under normal conditions (16–18). Unless care is taken to explicitly control for this confound, significantly more attention may be drawn to a stimulus when subjects are aware of it. This potential confound is especially important because many of the frontoparietal regions identified as playing a role in awareness have also been implicated in attention. A network including areas in the intraparietal sulcus and superior parietal lobule, called the dorsal attention network, has been shown to be involved in top-down control of attention (19, 20), as well as in bottom-up attention driven by salience (20, 21). A network including the TPJ, middle frontal gyrus, and inferior frontal gyrus, often called the ventral attention network, has been suggested to play a role in bottom-up attention driven by salience (22), or in redirection of attention to a behaviorally relevant stimulus (20, 23). A network including the TPJ and dorsolateral prefrontal cortex, called the frontoparietal control network, has been suggested to be involved in top-down executive control (24, 25). These three networks overlap substantially with brain regions identified as involved in awareness. Given the correlation between attention and awareness under normal conditions, which activations obtained in these networks reflect awareness, and which reflect attention?

In the present study, we used a behavioral paradigm that can manipulate visual awareness while controlling visual attention (26). We used a Posner cueing paradigm, a standard way to measure bottom-up spatial attention drawn to a visual cue (27). The cue was masked with metacontrast masking, a commonly used method to manipulate visual awareness (2). In one condition, the timing of the mask rendered the cue perceptually visible to the participants (aware trials), and in another condition, the timing of the mask rendered the cue perceptually invisible (unaware trials). Crucially, we measured the amount of attention that was drawn to the cue. We found that spatial attention drawn to the cue was not significantly different between the two conditions. Participants performed the task in an MRI scanner while brain activity was measured. The results were analyzed to find brain areas more active in the aware condition than in the unaware condition.

Results

Fig. 1 shows the behavioral paradigm (26). It is briefly summarized here and described in detail in the *Materials and Methods*. While the subject fixated centrally, a briefly presented (50 ms) cue

Significance

Do specific areas of the brain participate in subjective visual experience? We measured brain activity in humans using fMRI. Participants were aware of a visual stimulus in one condition and unaware of it in another condition. The two conditions were balanced for their effect on visual attention. Specific brain areas were more active in the aware than in the unaware condition, suggesting they were involved in subjective awareness independent of attention. The largest cluster of activity was found in the temporoparietal junction (TPJ), supporting a recent suggestion that the TPJ plays a role in computations concerning awareness.

Author contributions: T.W.W., K.M.I., A.S., and M.S.A.G. designed research; T.W.W., K.M.I., and M.S.A.G. performed research; T.W.W., K.M.I., and M.S.A.G. analyzed data; and T.W.W. and M.S.A.G. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. Email: graziano@princeton.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1611505113/-/DCSupplemental.

(a white spot on a black background) was used to draw attention to the left or right of fixation. The cue was followed by metacontrast masks, white rings whose inner diameter matched the outer diameter of the cue. These masks appeared either 50 ms after cue onset or 100 ms after cue onset. These intervals were selected with the aim of making participants unaware of the cue in one condition (50 ms) and aware of the cue in the other condition (100 ms). The masks always appeared on both sides of fixation to prevent biasing attention to one side. Following the masks, 180 ms after cue onset, a target stimulus was added to one of the mask rings, either on the right or the left of fixation. The location of the cue did not predict the location of the target. The target consisted of a line through the mask ring. Participants were asked to indicate as quickly as possible, by button press, whether the line was tilted left or right. Participants performed this task with a mean accuracy of 84% (SD = 12%) and a mean latency of 712 ms (SD = 124 ms). Participants were then asked to indicate, by button press, whether they had been aware of a cue on that trial (awareness probe). (See *Supporting Information* for a discussion of separating the cue from the target and a discussion of oddball effects.)

The responses on catch trials (one third of trials on which no cue was presented) were analyzed to determine the false-positive rate, the rate at which participants indicated that a cue was present when none was. This rate was low (mean = 18%, SD = 14%), suggesting participants were generally not guessing about the presence of the cue. Responses on the remaining two-thirds of trials in which a cue was presented were analyzed to determine the true-positive rate, the rate at which participants correctly indicated a cue was present. The true-positive rate was compared with the false-positive rate to calculate d′, a measure of sensitivity to the presence or absence of the cue. The results showed that participants were not aware of the cue on the majority of trials with a 50-ms cue/mask interval (27% true-positive rate; SD = 15%; d′ = 0.29), and were aware of the cue on the majority of trials with a 100-ms cue/mask interval (80% true-positive rate; SD = 13%; d′ = 2.01). A paired t test confirmed a highly significant difference between the two cue/mask interval conditions ($t = 13.7$; $P = 7 \times 10^{-13}$). The metacontrast masking manipulation thus successfully separated most trials into "aware" and "unaware."

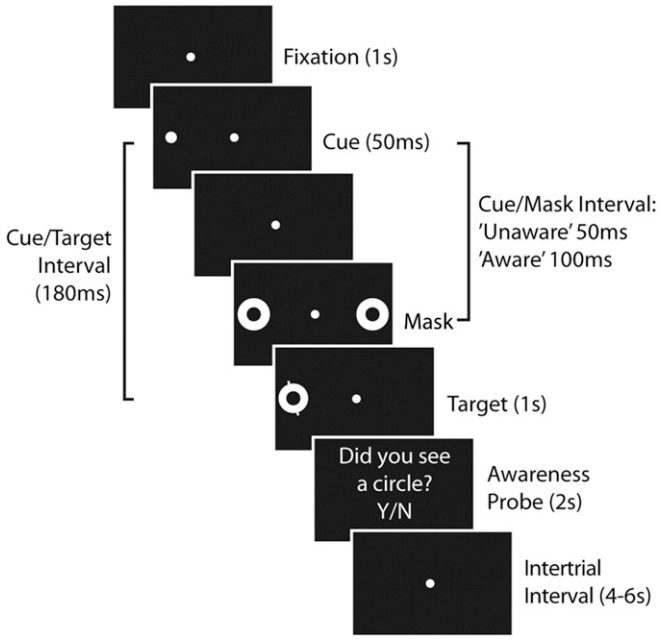

Fig. 1. Behavioral paradigm. In long cue/mask interval trials ("aware" condition), the mask was timed to allow participants to see the cue. In short cue/mask interval trials ("unaware" condition), the mask was timed to prevent participants from seeing the cue. The tilted discrimination target was presented on the same side as the cue (spatially matching as shown here) or on the opposite side as the cue (mismatching). After indicating the tilt of the target, participants were probed whether or not they had seen the cue.

The logic of the Posner task is as follows. Stimulus-driven spatial attention is drawn to the cue and lingers briefly at that location. This, in turn, speeds processing of a subsequent target at that location and slows processing of targets at different locations. Thus, target response latencies should be faster when the target appears at the same location as the cue (aligned trials), and slower when the target appears at the opposite location as the cue (misaligned trials). A standard measure of attention drawn by the cue is Δt = (latency when target and cue are misaligned − latency when target and cue are aligned). To the extent that this measure is significantly above zero, the cue can be said to have pulled spatial attention to one side at the expense of the other side. This was the case for both aware and unaware trials (aware condition: Δt = 25 ms; SD = 28, significantly above 0; $t = 4.5$; $P = 0.0001$; unaware condition: Δt = 17 ms, SD = 31, significantly above 0; $t = 2.8$; $P = 0.01$). These two measures were not significantly different from each other (planned comparison, paired $t = 1.3$; $P = 0.21$), indicating that the cue did not draw significantly more attention in the aware condition than in the unaware condition (see *Supporting Information* for further analysis of attention effects). There was also no significant difference in discrimination accuracy between the aware and unaware conditions (planned comparison, paired $t = 0.65$; $P = 0.52$), indicating that the two conditions were balanced for task difficulty.

A standard, general linear model (GLM) analysis was performed on the fMRI data. The signal was modeled using the cue onset as an event convolved with a gamma function. For details, see *Materials and Methods*. Fig. 2*A* shows the result of the contrast (aware condition − unaware condition). The largest area of activity was in the bilateral TPJ. Smaller areas of activity were scattered through the left prefrontal cortex, cingulate cortex, occipitotemporal cortex, left insula, and left precuneus. Table 1 gives coordinates for the peaks of activity (see *Supporting Information* for brain sections).

One potential concern with these results is that the metacontrast masking manipulation did not perfectly separate trials into "aware" and "unaware" conditions. To address this concern, we performed a separate analysis in which the data set was restricted to the 80% of long cue/mask interval trials on which participants reported being aware of the cue, and the 73% of short cue/mask interval trials on which participants reported being unaware of the cue. We refer to these as "confirmed aware" and "confirmed unaware" trials. Fig. 2*B* shows the result of this more restrictive analysis (confirmed aware trials − confirmed unaware trials). The results closely resemble those shown in Fig. 2*A*, but with fewer scattered areas of activity outside the TPJ.

Local Independent Component Analysis of TPJ activity. The activity shown in Fig. 2 covers a broad region of the TPJ. To fractionate the TPJ into more specific components, we used a local independent component analysis (ICA) analysis. ICA is a blind source-separation technique that has previously been used to parcellate fMRI data into spatially overlapping, but maximally independent, spatiotemporal components (28–34). We previously showed that local ICA can reliably parcellate the TPJ into spatiotemporal components, each functionally connected to distinct, brain-wide networks (32). Some of these subdivisions have also been observed in other studies, using data-driven parcellation of the TPJ (35, 36).

Using the fMRI data from the present experiment, we performed a local ICA, using the same methods described in our previous studies (32–34). The ICA was performed within a cortical mask that included the TPJ, a margin of surrounding cortex, and cortical areas dorsal to the TPJ that include the intraparietal

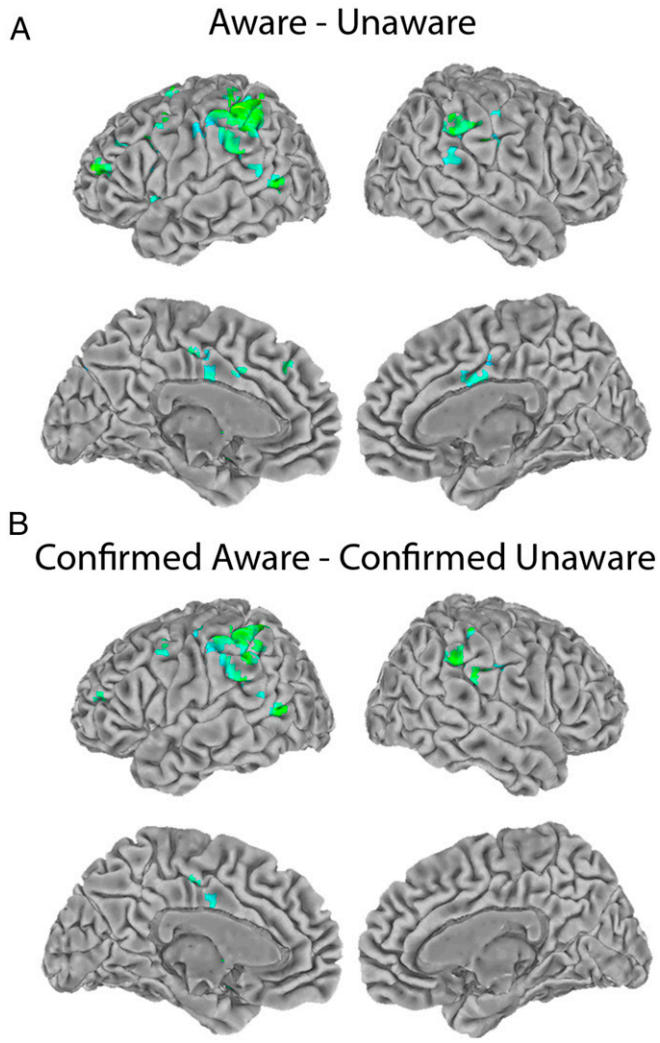

Fig. 2. Group fMRI results. Regression coefficients (positive only) for voxels that pass a threshold of $P < 0.01$ corrected for multiple comparisons adjusted for a 5-voxel minimum cluster size. (*A*) Group result for the contrast (aware condition − unaware condition). (*B*) Group result for the contrast (confirmed aware trials − confirmed unaware trials).

sulcus. The mask is described in detail in the *Materials and Methods*. By restricting the results to a region of interest, local ICA allows for a more fine-grained parcellation of fMRI data than a whole-brain ICA would (29, 32–34).

The local ICA decomposed the fMRI data into 20 independent components (ICs) within the mask, each IC defined by its pattern of activity over time. The ICs were further analyzed (see *Materials and Methods* for details) to determine which ICs were significantly more active in the aware condition than in the unaware condition. Essentially, the ICs were treated like voxels in a standard GLM analysis. Because 20 ICs were tested in this manner, the analysis was Bonferroni corrected for 20 comparisons.

The analysis revealed that three TPJ components responded significantly more to the aware condition than to the unaware condition. These components included a left-lateralized, dorsal IC ($F = 39.2$; $P < 0.05$ Bonferroni corrected for 20 comparisons) consistent with the location of left TPJd (dorsal TPJ) in our previous work; a right-lateralized, dorsal IC ($F = 11.4$; $P < 0.05$ Bonferroni corrected for 20 comparisons) consistent with the location of right TPJd in our previous work, and a bilateral, anterior IC ($F = 11.5$; $P < 0.05$ Bonferroni corrected for 20 comparisons) consistent with the location of TPJa (anterior TPJ) in our previous work. No other ICs responded significantly to this comparison.

Winner-take-all spatial maps for the significant ICs are shown in Fig. 3*A*. The right TPJd component has some bilateral representation. It includes a large area on the right and a smaller area on the left. The left TPJd is confined to the left hemisphere. The TPJa is approximately equally represented in both hemispheres. Regression coefficients for the aware and unaware conditions for each component are shown in Fig. 3*B*. Fig. 4 shows the time courses for these components.

We performed a functional connectivity analysis (Fig. 5) on the three TPJ components identified in the ICA-based regression. The bilateral TPJa component was functionally connected to the bilateral anterior insula, bilateral precuneus, bilateral cingulate gyrus, bilateral inferior frontal gyrus, and bilateral middle temporal gyrus, a pattern consistent with the connectivity of TPJa found in our previous studies (32–34) and partially overlapping the known connectivity of the ventral attention network (20, 22, 23), although the ventral attention network is typically right biased. It also shared similarities with a previously described "salience" network and a "cingulo-opercular" network (37, 38) (see *Supporting Information* for quantitative comparison). The left TPJd component was functionally connected to the left dorsolateral prefrontal cortex, left medial prefrontal cortex, left precuneus, left middle temporal gyrus, and left insula, with smaller corresponding clusters on the right side, matching the pattern of connectivity obtained in our previous studies of the TPJd (32–34) and partially overlapping the known connectivity of the frontoparietal control network (24, 25). The right TPJd component was functionally connected to the right dorsolateral prefrontal cortex, right medial prefrontal cortex, right inferior frontal gyrus, right precuneus, and right insula, with smaller corresponding clusters on the contralateral side. This pattern showed some similarity to the frontoparietal control network and to the ventral attention network.

Involvement of the Dorsal Attention Network. We examined the functional connectivity pattern for each of the 20 components identified by the ICA and found one component that showed connectivity matching the dorsal attention network (see *Supporting Information* for quantitative comparison). As expected, this component was located bilaterally, extending over both banks of the intraparietal sulcus. It was functionally connected to the bilateral superior parietal lobule, bilateral frontal eye fields, bilateral supplementary eye fields, and bilateral middle and inferior temporal gyri (Fig. 5*D*), closely

Table 1. MNI coordinates for areas of activation in Fig. 2*A*

Area	Volume	Peak		
		LR	PA	IS
Left TPJ	390	−55	−46	57
Cingulate gyrus	46	−1	8	40
Left DLPFC	35	−46	49	20
Left MFG	30	−49	11	37
Right TPJ	29	66	−32	41
Left MTG	21	−58	−65	0
Right TPJ	21	63	−29	25
Left insula	15	−40	3	7
Left precuneus	13	−4	−85	42
Left insula	12	−40	−15	−7
Left TPJ	11	−64	−53	14
Left MFG	10	−28	−3	72

Volume (number of voxels) and MNI coordinates (mm) of the peak of activity for clusters of activation in Fig. 2*A* with 10 voxels or more. Cingulate gyrus cluster includes activity in both left and right cingulate gyri. DLPFC, dorsolateral prefrontal cortex; IS, inferior-superior; LR, left-right; MFG, middle frontal gyrus; MTG, middle temporal gyrus; PA, posterior-anterior.

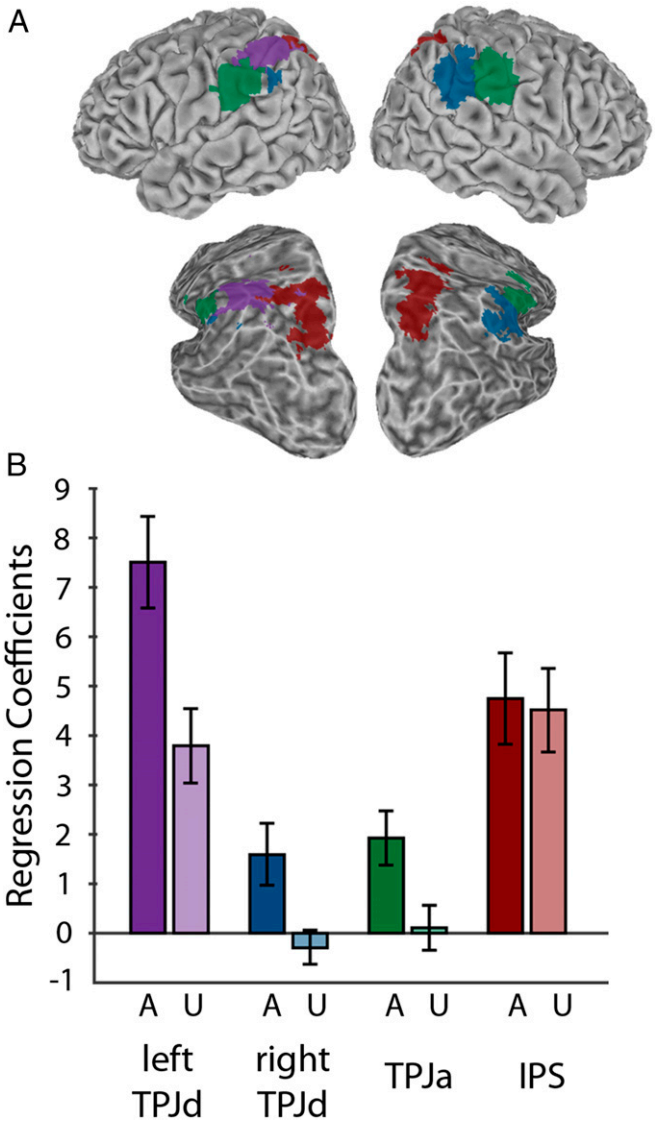

Fig. 3. Group ICA results. (*A*) Winner-take-all maps showing location of Left TPJd (purple); Right TPJd (blue); TPJa (green); intraparietal sulcus (IPS) (red). A partially inflated view from a posterior, dorsal, lateral angle is also shown to better reveal the inside of the intraparietal sulcus. (*B*) Regression coefficients for the aware (A) and unaware (U) conditions. Error bars show SE among subjects.

matching the dorsal attention network, as identified by previous task-based studies of attention networks (19, 20) and previous parcellations of cortical networks (39, 40). The winner-take-all spatial map for this component is shown in Fig. 3*A*.

This intraparietal component did not respond significantly more in the aware condition than in the unaware condition ($F = 0.2$; $P > 0.05$ Bonferroni corrected for 20 comparisons). The bar graphs in Fig. 3*B* show regression coefficients for the aware and unaware conditions for the left TPJd, right TPJd, TPJa, and the intraparietal component. The awareness manipulation had an effect on the activity in the three TPJ components, but not in the intraparietal component.

Fig. 4 shows the time courses for all four components. Again, the awareness manipulation affected activity in the TPJ components, but not the intraparietal component.

Discussion

Previous studies suggested that a broad frontoparietal network may be involved in subjective awareness (3–9). Which parts of this network participate specifically in subjective awareness, and which parts are involved in the closely related process of attention? Because the unaware and aware stimuli in our task had balanced effects on attention, the present experiment provides an opportunity to disambiguate these processes. The results suggest that more ventral regions of this network, especially in the TPJ, may participate in subjective awareness. More dorsal regions in the banks of the intraparietal sulcus did not respond in association with awareness, despite their established involvement in attention.

When local ICA (32–34) was used to specify subcomponents of the TPJ, three specific subcomponents showed greater activity in the aware than in the unaware conditions. These subcomponents included the left TPJd, the right TPJd, and a bilateral component in the TPJa. Their functional connectivity partly resembled the previously reported connectivity of the frontoparietal control network and the ventral attention network (20, 22–25). The functional connectivity of these three TPJ components overlapped with the smaller foci of awareness-related activity obtained in our GLM analysis throughout the prefrontal, cingulate, and temporal cortex (Fig. 2; see *Supporting Information* for quantitative comparison). Thus, although the largest area of activity obtained in the present study was in the TPJ, it was likely acting in the context of larger cortical networks.

In contrast, a component in the banks of the intraparietal sulcus did not show significantly greater activity in the aware than in the unaware conditions. This component was functionally connected with the dorsal attention network identified by previous studies (19, 20, 39, 40).

These results are consistent with a recent proposal that subjective awareness is associated with the TPJ (41–43). In that proposal, awareness serves as an internal model to help the brain monitor its own state of attention. Awareness and attention therefore normally work together. However, when they are separated, as in the present task, then awareness should be more closely associated with the TPJ than with the dorsal attention network.

These results also have implications for the clinical syndrome of hemispatial neglect. Neglect can include a mixture of different

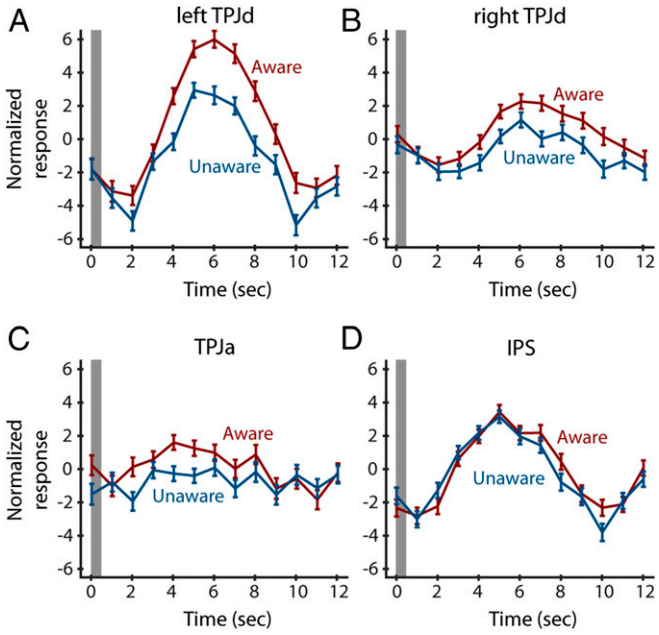

Fig. 4. Time courses for the independent components obtained in the left TPJd (*A*), right TPJd (*B*), TPJa (*C*), and the intraparietal sulcus (*D*). *y* axis shows arbitrary units of MR activity provided as an output of the ICA method. Gray bar shows time of cue presentation. Error bars show SE.

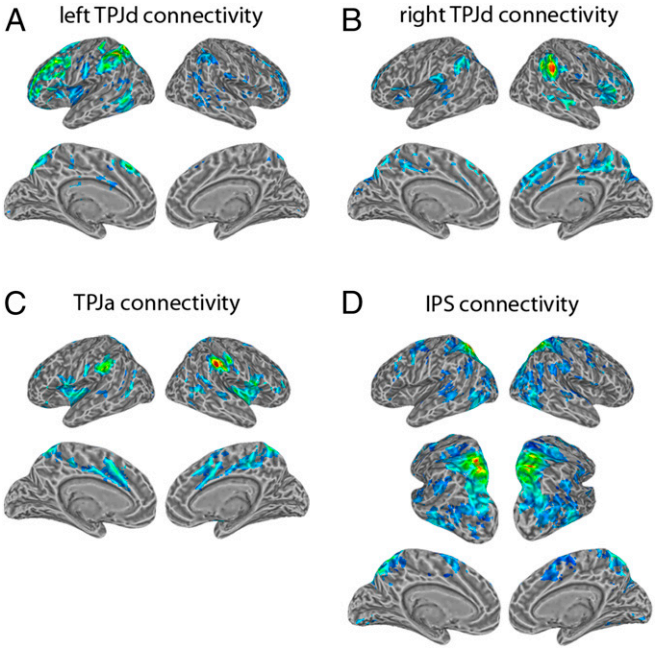

Fig. 5. Functional connectivity patterns for the independent components obtained in the left TPJd (*A*), right TPJd (*B*), TPJa (*C*), and the intraparietal sulcus (*D*). The time series for each component served as a seed for its functional connectivity analysis. Shown are regression coefficients (positive only) for voxels that pass a threshold of $P < 0.01$ corrected for multiple comparisons adjusted for a 5-voxel minimum cluster size. A partially inflated view from a posterior, dorsal, lateral angle is also shown to better reveal the inside of the intraparietal sulcus.

symptoms in different combinations caused by damage to a range of brain structures, including the TPJ (44). Neglect has not typically been probed with standard paradigms for distinguishing attention from awareness, but the present results suggest it may be worth testing for a dissociation between awareness-related and attention-related neglect caused by damage to different, adjacent networks.

Materials and Methods

Subjects. Twenty-five subjects (14 women; age range, 18–48; normal or corrected-to-normal vision) participated in the experiment. All subjects provided informed consent, and all experimental procedures were approved by the Princeton Institutional Review Board.

Behavioral Paradigm. Stimuli were projected with the Hyperian MRI Digital Projection System (Psychology Software Tools) at the end of the scanner bore. Each subject lay face-up on the scanner bed, with foam surrounding the head to reduce head movements and earplugs to reduce noise. All stimuli were developed and presented with the MATLAB psychophysics toolbox (45).

Fig. 1 shows the behavioral paradigm. Each trial began with a white central fixation point on a black background. Participants were instructed to fixate during the trial. After 1 s, the cue period began and lasted three refresh cycles (~50 ms). Throughout the cue period, the cue (a white spot 1.1° in diameter) was presented 6° to the left of fixation (1/3 of trials), to the right of fixation (1/3 of trials), or not presented (1/3 of trials).

The cue was followed by a mask. In the long cue/mask interval condition (1/2 of trials), a cue/mask interval of three refresh cycles (~50 ms) was inserted between the cue and the mask. Thus, the time from cue onset to mask onset was ~100 ms, which was intended to allow the cue to be subjectively visible. The mask then remained on the screen for five refresh cycles (~80 ms). In the short cue/mask interval condition (1/2 of trials), no time was inserted between the cue and the mask. The time from cue onset to mask onset was thus ~50 ms, which was intended to render the cue subjectively invisible. The mask then remained on the screen for eight refresh cycles (~130 ms). In both conditions, the time from cue onset to mask offset was 11 refresh cycles (~180 ms). The mask consisted of two white metacontrast rings (2) with an outer diameter of 1.5° and an inner diameter of 1.1°, presented 6° to the left and right of fixation. The mask therefore did not preferentially attract bottom-up attention to one side.

After the mask period, a target was added to one of the already-present circles that composed the mask. The target was either on the left (1/2 of trials) or right (1/2 of trials) of fixation. The location of the cue did not predict the location of the target. The interval between cue onset and target onset was 11 refresh cycles (~180 ms). The target consisted of one white line segment extending from the top of the mask ring and one white line segment extending from the bottom of the mask ring. The segments were collinear, forming an implied line through the ring. The line was tilted toward the left (1/2 of trials) or toward the right (1/2 of trials) by 3°. Participants were required to discriminate the orientation of the line, using a button box with the right hand. The target lasted 1 s, and participants were instructed to respond as quickly as possible during that period. The 20% of trials with incorrect or no response during the 1-s response window were excluded from the behavioral analysis. Trials with a latency <300 ms were also excluded from the behavioral analysis because the reaction time was too short to be a plausible response to the discrimination task.

On each trial, after the 1-s target period, a question appeared on the screen: "Did you see a circle? Y/N." Participants were instructed that this question referred to the cue stimulus presented at the beginning of some trials. The question remained on the screen for 2 s, and participants were instructed to respond during that time, using a button box with the right hand.

After the question period, a 4–6-s intertrial interval began, during which the fixation point was present. Participants were instructed to maintain fixation during the intertrial interval.

All trial types were randomly interleaved. Each participant performed practice trials followed by nine runs of 24 trials each (216 trials total). Target orientation (tilted left or right), target location (left or right), cue condition (left, right, or no cue), and cue/mask interval (long or short) were randomized and counter-balanced within each run.

MRI Data Collection and Preprocessing. MRI images covering the whole cortex were acquired with a 20-channel receiver head coil on a Siemens Skyra scanner. Functional imaging used a gradient echo, echoplanar pulse sequence with a 64 × 64 matrix [35 axial slices, 3 mm thick, field of view (FOV), 192 × 192 mm; repetition time (TR), 2 s; echo time (TE), 30 ms; flip angle (FA), 77°; in-plane resolution, 3 × 3 mm]. Functional images were aligned with a high-resolution anatomical scan (MPRAGE) taken at the end of the session (FOV, 256 × 224 mm; TR, 2.3 s; TE, 2.98 ms; FA, 9°; 256 × 224 matrix; 1 mm^3 resolution).

Preprocessing was done with AFNI (46) and FSL (47) software packages. The functional data were slice time-corrected, motion corrected (to the image acquired closest in time to the anatomical scan), and detrended (linear and quadratic) with AFNI. Single-session ICA was applied to each subject's unsmoothed functional data, using the MELODIC toolbox in FSL (28), and components that represented noise were regressed out using the FSL tool fsl_regfilt (28, 48). The following spatial or temporal features were considered to represent noise: 1, spatial association with white matter, ventricles, or background voxels; 2, a lack of cluster formation; 3, large spikes in the time course; 4, high-frequency noise; or 5, temporal sawtooth patterns likely to reflect aliasing of cardiac or respiratory signals exceeding the Nyquist frequency (48). To ensure that all neural activity would remain untouched for local ICA on the group level, ICs that appeared to contain a mixture of noise and signal were not filtered out in the denoising step. The denoised data were spatially smoothed with a Gaussian kernel (full-width half-maximum 4 mm), using AFNI.

GLM Analysis of fMRI Data. Statistical analyses were performed using multiple regression within the framework of the GLM (49) with AFNI. An event-related design was used in which cue onset was the relevant event and differing responses to aware vs. unaware cues were analyzed. Separating the response to the cue onset from the response to other elements of the trial such as fixation onset or target onset was not necessary in this design, as these other trial elements were balanced across the aware and unaware conditions (*Supporting Information*). The blood oxygenation level dependent response was modeled by convolving the stimulus timing (cue onset) with a gamma function. Regressors of noninterest were included to account for head motion and linear drift in scanner signal. These regressions produced statistical maps at the individual subject level for two conditions of interest: long cue/mask interval trials on which a cue was presented (aware condition), and short cue/mask trials on which a cue was presented (unaware condition). Subject-level statistical maps were spatially normalized to FSL's

Montreal Neurological Institute (MNI)-152 template, using AFNI (50). A mixed-effects analysis was then computed at the group level (awareness condition as fixed effect, subject as random effect), using AFNI (51), producing the contrast (aware condition − unaware condition). Coefficients for this contrast were thresholded using a Monte Carlo simulation to achieve a corrected significance of $P < 0.01$ adjusted for a cluster size of 5 adjacent voxels.

Local ICA of TPJ Activity. The fMRI data from all subjects were temporally concatenated and then subjected to probabilistic ICA, using the MELODIC toolbox in FSL (28). ICA decomposition was applied to the voxels within a region of interest mask containing the TPJ and the intraparietal sulcus. The mask was constructed from the standard surface cvs_avg35_inMNI152 in Freesurfer, using mri_label2vol to combine multiple labels from the aparc.a2009s atlas into one mask (G_pariet_inf-Supramar, G_pariet_inf-Angular, G_temp_sup-Plan_tempo, G_temp_sup-Lateral, G_temp_sup-G_T_transv, S_interm_prim-Jensen, S_temporal_sup, S_temporal_transverse, S_intrapariet_and_P_trans) and trimming temporal cortex voxels anterior to the postcentral sulcus. The fMRI data were decomposed into 20 ICs. ICs were thresholded at Z = 4 for the creation of winner-take-all maps for the figures.

The time courses from all 20 ICs were separated into time courses for each individual subject and entered into a multiple regression analysis, using AFNI. This analysis was similar to the voxel-wise GLM analysis described in the previous section, except that instead of using the time courses of voxels as dependent variables, the time courses of ICs were used. This analysis produced coefficients for the aware and unaware conditions, one for each of 25 subjects, for each of 20 ICs. These coefficients were then entered into a mixed-effects analysis at the group level (awareness condition as fixed effect, subject as random effect), using MATLAB (MathWorks).

Functional connectivity analysis was performed to identify the brain-wide networks to which each IC was connected. IC time courses were used as the independent variables in multiple regression analyses at the individual subject level, using AFNI. These regressions produced brain-wide statistical maps for each of 25 subjects. Subject-level statistical maps were spatially normalized to FSL's MNI-152 template, using AFNI (50), and entered in a mixed-effects analysis at the group level (subject as random effect) using AFNI (51), producing group-level statistical maps. These maps were thresholded using a Monte Carlo simulation to achieve a corrected significance of $P < 0.01$ adjusted for a cluster size of 5 adjacent voxels.

ACKNOWLEDGMENTS. We thank Alexandra Reblando and Hope Kean for assistance in data collection. This work was funded by the Princeton Neuroscience Institute Innovation Fund.

1. Kim CY, Blake R (2005) Psychophysical magic: Rendering the visible 'invisible'. *Trends Cogn Sci* 9(8):381–388.
2. Breitmeyer B, Öğmen H (2006) *Visual masking* (Oxford Univ. Press, New York).
3. Lumer ED, Friston KJ, Rees G (1998) Neural correlates of perceptual rivalry in the human brain. *Science* 280(5371):1930–1934.
4. Dehaene S, et al. (2001) Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci* 4(7):752–758.
5. Rees G, Kreiman G, Koch C (2002) Neural correlates of consciousness in humans. *Nat Rev Neurosci* 3(4):261–270.
6. Naghavi HR, Nyberg L (2005) Common fronto-parietal activity in attention, memory, and consciousness: Shared demands on integration? *Conscious Cogn* 14(2):390–425.
7. Carmel D, Lavie N, Rees G (2006) Conscious awareness of flicker in humans involves frontal and parietal cortex. *Curr Biol* 16(9):907–911.
8. Bor D, Seth AK (2012) Consciousness and the prefrontal parietal network: Insights from attention, working memory, and chunking. *Front Psychol* 3:63.
9. Tallon-Baudry C (2012) On the neural mechanisms subserving consciousness and attention. *Front Psychol* 2:397.
10. Lau HC (2008) Are we studying consciousness yet. *Frontiers of Consciousness: Chichele Lectures*, eds Weiskrantz L, David M (Oxford Univ. Press, Oxford, UK), pp 245–258.
11. Tsuchiya N, Wilke M, Frässle S, Lamme VA (2015) No-report paradigms: Extracting the true neural correlates of consciousness. *Trends Cogn Sci* 19(12):757–770.
12. McCormick PA (1997) Orienting attention without awareness. *J Exp Psychol Hum Percept Perform* 23(1):168–180.
13. Koch C, Tsuchiya N (2007) Attention and consciousness: Two distinct brain processes. *Trends Cogn Sci* 11(1):16–22.
14. Kentridge RW, Nijboer TC, Heywood CA (2008) Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia* 46(3):864–869.
15. Norman LJ, Heywood CA, Kentridge RW (2013) Object-based attention without awareness. *Psychol Sci* 24(6):836–843.
16. Posner MI (1994) Attention: The mechanisms of consciousness. *Proc Natl Acad Sci USA* 91(16):7398–7403.
17. Merikle PM, Joordens S (1997) Parallels between perception without attention and perception without awareness. *Conscious Cogn* 6(2–3):219–236.
18. De Brigard F, Prinz J (2010) Attention and consciousness. *Wiley Interdiscip Rev Cogn Sci* 1(1):51–59.
19. Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci* 23:315–341.
20. Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3(3):201–215.
21. Geng JJ, Mangun GR (2009) Anterior intraparietal sulcus is sensitive to bottom-up attention driven by stimulus salience. *J Cogn Neurosci* 21(8):1584–1601.
22. Downar J, Crawley AP, Mikulis DJ, Davis KD (2002) A cortical network sensitive to stimulus salience in a neutral behavioral context across multiple sensory modalities. *J Neurophysiol* 87(1):615–620.
23. Corbetta M, Patel G, Shulman GL (2008) The reorienting system of the human brain: From environment to theory of mind. *Neuron* 58(3):306–324.
24. Liston C, Matalon S, Hare TA, Davidson MC, Casey BJ (2006) Anterior cingulate and posterior parietal cortices are sensitive to dissociable forms of conflict in a task-switching paradigm. *Neuron* 50(4):643–653.
25. Vincent JL, Kahn I, Snyder AZ, Raichle ME, Buckner RL (2008) Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *J Neurophysiol* 100(6):3328–3342.
26. Webb TW, Kean HH, Graziano MSA (2016) Effects of awareness on the control of attention. *J Cogn Neurosci* 28(6):842–851.
27. Posner MI (1980) Orienting of attention. *Q J Exp Psychol* 32(1):3–25.
28. Beckmann CF, Smith SM (2004) Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans Med Imaging* 23(2):137–152.
29. Keck IR, et al. (2006) Region of interest based independent component analysis. *Neural Information Processing* (Springer, Berlin), pp 1048–1057.
30. Calhoun VD, Kiehl KA, Pearlson GD (2008) Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Hum Brain Mapp* 29(7):828–838.
31. Kiviniemi V, et al. (2009) Functional segmentation of the brain cortex using high model order group PICA. *Hum Brain Mapp* 30(12):3865–3886.
32. Igelström KM, Webb TW, Graziano MSA (2015) Neural processes in the human temporoparietal cortex separated by localized independent component analysis. *J Neurosci* 35(25):9432–9445.
33. Igelström KM, Webb TW, Graziano MSA (April 12, 2016) Functional connectivity between the temporoparietal cortex and cerebellum in autism spectrum disorder. *Cereb Cortex*, 10.1093/cercor/bhw079.
34. Igelström KM, Webb TW, Graziano MSA (2016) Topographical organization of attentional, social and memory processes in the human temporoparietal cortex. *eNeuro* 3:e0060.
35. Mars RB, et al. (2012) Connectivity-based subdivisions of the human right "temporoparietal junction area": Evidence for different areas participating in different cortical networks. *Cereb Cortex* 22(8):1894–1903.
36. Bzdok D, et al. (2013) Characterization of the temporo-parietal junction by combining data-driven parcellation, complementary connectivity analyses, and functional decoding. *Neuroimage* 81:381–392.
37. Dosenbach NU, et al. (2007) Distinct brain networks for adaptive and stable task control in humans. *Proc Natl Acad Sci USA* 104(26):11073–11078.
38. Seeley WW, et al. (2007) Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci* 27(9):2349–2356.
39. Fox MD, Corbetta M, Snyder AZ, Vincent JL, Raichle ME (2006) Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proc Natl Acad Sci USA* 103(26):10046–10051.
40. Yeo BT, et al. (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* 106(3):1125–1165.
41. Graziano MSA, Kastner S (2011) Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cogn Neurosci* 2(2):98–113.
42. Graziano MSA (2013) *Consciousness and the Social Brain* (Oxford Univ. Press, Oxford, UK).
43. Kelly YT, Webb TW, Meier JD, Arcaro MJ, Graziano MS (2014) Attributing awareness to oneself and to others. *Proc Natl Acad Sci USA* 111(13):5012–5017.
44. Vallar G, Perani D (1986) The anatomy of unilateral neglect after right-hemisphere stroke lesions. A clinical/CT-scan correlation study in man. *Neuropsychologia* 24(5):609–622.
45. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10(4):433–436.
46. Cox RW (1996) AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29(3):162–173.
47. Jenkinson M, Beckmann CF, Behrens TEJ, Woolrich MW, Smith SM (2012) FSL. *Neuroimage* 62(2):782–790.
48. Kelly RE, Jr, et al. (2010) Visual inspection of independent components: Defining a procedure for artifact removal from fMRI data. *J Neurosci Methods* 189(2):233–245.
49. Friston KJ, et al. (1994) Statistical parametric maps in functional imaging: A general linear approach. *Hum Brain Mapp* 2(4):189–210.
50. Talairach J, Tournoux P (1988) Co-planar stereotaxic atlas of the human brain (Thieme Medical Publishers, New York).
51. Chen G, Saad ZS, Nath AR, Beauchamp MS, Cox RW (2012) FMRI group analysis combining effect estimates and their variances. *Neuroimage* 60(1):747–765.
52. Ollinger JM, Shulman GL, Corbetta M (2001) Separating processes within a trial in event-related functional MRI I. The Method. *Neuroimage* 13(1):210–217.
53. Stevens AA, Skudlarski P, Gatenby JC, Gore JC (2000) Event-related fMRI of auditory and visual oddball tasks. *Magn Reson Imaging* 18(5):495–502.