

# Neural Correlates of Stimulus Reportability

Oliver J. Hulme, Karl F. Friston, and Semir Zeki

## Abstract

■ Most experiments on the “neural correlates of consciousness” employ stimulus reportability as an operational definition of what is consciously perceived. The interpretation of such experiments therefore depends critically on understanding the neural basis of stimulus reportability. Using a high volume of fMRI data, we investigated the neural correlates of stimulus reportability using a partial report object detection paradigm. Subjects were presented with a random array of circularly arranged *disc-stimuli* and were cued, after variable delays (following stimulus offset), to report the presence or absence of a disc at the cued location, using variable motor

actions. By uncoupling stimulus processing, decision, and motor response, we were able to use signal detection theory to deconstruct the neural basis of stimulus reportability. We show that retinotopically specific responses in the early visual cortex correlate with stimulus processing but not decision or report; a network of parietal/temporal regions correlates with decisions but not stimulus presence, whereas classical motor regions correlate with report. These findings provide a basic framework for understanding the neural basis of stimulus reportability without the theoretical burden of presupposing a relationship between reportability and consciousness. ■

## INTRODUCTION

The fact that consciousness is private poses a fundamental problem for studying its neural basis. How can experimenters know what their subject is experiencing? Most investigators make assumptions about which behaviors signify the occurrence of which conscious contents. For example, if a subject can make a voluntary report about “X” then the subject is “conscious of X.” The logic is that by measuring the correlation between a report of “X” and the subject’s brain states, one can infer the neural correlates of the conscious experience of “X.” Stimulus reportability thus acts as an implicit “bridging principle.” Most experiments using this principle deploy a strategy best characterized by the maxim “keep the stimulus constant, change the percept.” In principle, this paradigm has the elegant property of dissociating the neural correlates of stimulus processing from those underlying its perception. Examples include binocular rivalry (Haynes, Deichmann, & Rees, 2005; Lumer, Friston, & Rees, 1998; Tong, Nakayama, Vaughan, & Kanwisher, 1998; Leopold & Logothetis, 1996; Logothetis & Schall, 1989), bistable figures (Kleinschmidt, Buchel, Zeki, & Frackowiak, 1998), detection tasks (Rees & Heeger, 2003), masking (Haynes, Driver, & Rees, 2005), and the attentional blink (Sergent, Baillet, & Dehaene, 2005), whereby the same stimulus can result in two or more conscious contents. These studies often show activation in fronto-parietal and posterior sensory cortices, although there is disagreement about which correlates with conscious perception.

We propose that the interpretations of these experiments are vulnerable because of the bridging principle they adopt. If one operationally defines a conscious state with a given objective marker such as a motor report, then one will be guaranteed to find the neural correlates of that marker, in this instance, the neural correlates of motor report. If the conceptual bridge between the marker and consciousness is erroneous, then making conclusions about the neural correlates of consciousness (NCC) will be also be erroneous (Chalmers, 1998). Here we investigate the neural correlates of stimulus reportability itself, without the burden of prior conceptual commitments to its relationship with consciousness. According to different theoretical or philosophical perspectives certain components of reportability may or may not be dissociable from consciousness (or components of consciousness). Once one has a framework for the neural basis of reportability, then one can speculate about which components relate to consciousness. Whether one believes reportability is dissociable from or integral to consciousness does not change the basic empirical findings of this study.

To investigate the neural basis of stimulus reportability, we used a partial report paradigm (Sperling, 1960), incorporating object detection (where the objects are simple “discs”). This requires subjects to view an array of discs and then report on the cued location. Because the cue occurs after stimulus offset, the report is dependent on iconic memory (Sperling, 1960). The task can be characterized as involving three stages: stimulus processing, decision, and motor report. We found the neural correlates of each using the following manipulations.

---

University College London

First, by varying the cue delay, we manipulated performance such that at short delays the report is coupled to stimulus presence, whereas at long delays the two are decoupled. Second, by varying the hand used to report the presence or absence of the stimulus, we decoupled the decision from the motor act used to report it. With this approach, we show that retinotopically specific responses in the early visual cortex correlate with stimulus processing but not with decision or report, that activity in a network of parietal/temporal regions correlates with decisions but not stimulus presence, whereas activity in classical motor regions correlates with the motor act of reporting. Which of these components relates to consciousness is considered from different theoretical perspectives, but we argue that without resolving these issues one should be cautious in interpreting neural correlates of reportability as being equivalent to the NCC.

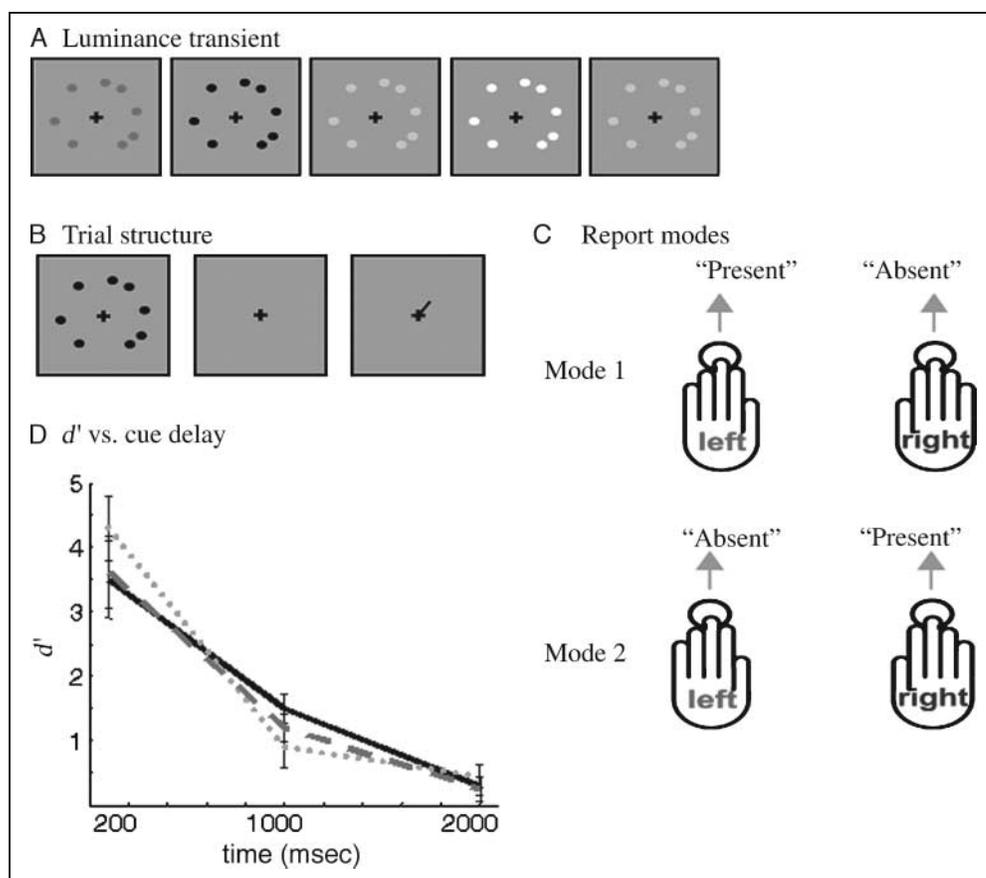
## METHODS

### Subjects and Task

Three subjects, all experienced psychophysical observers (2 men, mean age = 23 years, all right-handed) and with normal vision, were tested and scanned. We used a case-study design because the degrees of freedom were more

powerfully deployed within subjects, that is, we acquired a large number of sessions from three subjects (as opposed to three sessions from a large number of subjects). All subjects gave informed consent in accordance with the Declaration of Helsinki, and the Ethics Committee of the National Hospital for Neurology and Neurosurgery, London, UK, granted ethical approval for the study. This stimulus program was realized using Cogent 2000 developed by the Cogent 2000 team at the FIL and the ICN and Cogent Graphics developed by John Romaya at the LON and Wellcome Trust Centre for Neuroimaging ([www.vislab.ucl.ac.uk/cogent\\_2000.php](http://www.vislab.ucl.ac.uk/cogent_2000.php)). At the onset of each trial, a disc array consisting of eight discs appeared for 200 msec, positioned randomly at 8 out of 16 radial positions. To prevent the random occurrence of several adjacent discs, thus making it easier to chunk them together into a higher-level shape, no more than one disc could appear in each octant. Each disc ( $0.5^\circ$  diameter) was defined by a sinusoidal luminance transient (Figure 1A), which was parameterized to prevent afterimages (see Supplementary Methods 1). After a variable delay, a central cue (a line extending from the fixation cross to  $1^\circ$  eccentricity) was presented for 100 msec (Figure 1B). The orientation of the cue line pointed pseudorandomly toward one of 16 disc positions, indicating which position the subject was to report on via button press (yes, disc present vs. no, disc absent).

**Figure 1.** (A) Schematic depiction of a disc array defined by a 200-msec duration luminance transient. (B) Depicts the basic structure of a trial. Following the cueing stage, the subject makes their response. (C) Depicts the two response modes, each used to answer the same question of whether the disc was present or absent at the cued location. (D) Plot of the variation in  $d'$  as a function of cue delay, acquired during fMRI data acquisition. Gray dots, gray bars, and a black solid line indicate Subjects 1, 2, and 3, respectively. Error bars depict SEM across stimulus position.



Subjects were given a 500-msec time window in which to respond immediately after the cue offset. All discs were presented parafoveally at an eccentricity of  $3^\circ$  and a distance of 50 cm (63 cm in scanner), while the subject fixated a central cross. Psychophysics outside of the scanner was performed on a Sony high-resolution CRT monitor with a refresh rate of 100 Hz. For each subject, there was a parameter-optimization session, using a staircasing procedure (Macmillan & Creelman, 1996) in which stimulus contrast (Michelson average contrast for peak deviation in the transient =  $0.85 \pm 0.13$ ; Michelson, 1927) and cue delay were adjusted in order to normalize performance (see Supplementary Methods 2 and Supplementary Table 1 in supplementary material).

This psychophysical characterization phase allowed accurate plotting of a performance–cue delay decay function, assessing how performance decreased with increasing cue delay. To check that the results would not be corrupted by stimulus arrays simply being invisible, we additionally asked subjects to rate, on a scale of 1–10, how visible each disc array onset was for delays of 100, 1000, and 2000 msec. Visibility was rated as 10 for all subjects with no effect of delay. Each subject performed 10 sessions of 600 trials, making a total of 6000 trials. The hand used to report presence or absence via button press was pseudorandomized (see Figure 1C). For instance, for the first session, the subject might be instructed to report “present” with the left hand and “absent” with the right, whereas in another session the converse would be instructed. This not only controlled for any possible bias between the hand used to make the report and the actual meaning of the report but also was important in allowing us to dissociate the motor act of reporting from the decision in the subsequent fMRI experiment. Subjects were trained and tested for 6 hr prior to scanning.

## Scanning

The same stimuli were used as in the psychophysical phase (see above) with the exception that only three delays were used—short (200 msec), intermediate (900–1100 msec), and long (2200 msec). The contrast levels for the disc array were recalibrated prior to every sequence of scanning sessions to account for any differences in the display system used in the scanner, and to take account of any learning effects that may have occurred. A staircasing procedure was again applied during the scanner sessions, adjusting contrast (absolute magnitude of luminance transient) to maintain performance between 92.5% and 97.5% for the short delay (see Supplementary Table 2, supplementary material). This enforced consistent performance by limiting variations within a session, for instance, due to tiredness. The starting contrast at the beginning of each session was identical within subjects, set according to the psycho-

physical calibration phase. The experiment consisted of eight sessions of 600 trials each, making a total of 4800 trials per subject. Each session was approximately 30 min in duration with total acquisition time per subject exceeding 7 hr. Subjects were informed via a notch on the fixation cross which hand they were to use to report. All stimuli were presented using COGENT 2000 Graphics ([www.vislab.ucl.ac.uk/cogent](http://www.vislab.ucl.ac.uk/cogent)) running in MATLAB (Mathworks; [www.mathworks.com](http://www.mathworks.com)). Visual stimuli were projected via an LCD projector onto a transparent screen positioned over the subject’s head and were viewed through a tilted mirror fixed to the head coil. Subjects were instructed to fixate a central cross, and eye tracking was used to ensure correct fixation. Neither stimulus nor cue onset induced saccades.

Scanning was performed using a 3-T Siemens Allegra fMRI scanner with a head-volume coil (Siemens, Erlangen, Germany). A gradient EPI sequence was used to maximize blood oxygen level dependent contrast (TE = 65 msec, repeat time = 2.47 sec). Each brain image was acquired in a descending sequence comprising 38 axial slices (each 2 mm thick with 1 mm interleaved gaps), such that the volume consisted of  $64 \times 64$  voxels and gave near whole-brain coverage (excluding the most ventral part of the cerebellum in some subjects). Each of the eight sessions yielded an average of 730 volumes, with the first five being discarded to allow for T1 equilibration effects.

We analyzed the data using SPM2 software (Wellcome Department of Imaging Neuroscience, London, UK; [www.fil.ion.ucl.ac.uk/spm](http://www.fil.ion.ucl.ac.uk/spm)). The EPI images were realigned spatially and normalized to the MNI template provided in SPM2. Due to spatial resolution issues pertaining to the induced responses of small regions of the visual cortex, we analyzed spatially unsmoothed data. To finesse the application of Gaussian field theory, which assumes spatial smoothness, we supersampled all functional images to  $1\text{-mm}^3$  voxels. All data were temporally filtered with a high-pass filter with a cutoff period of 128 sec. Global changes in activity were removed by proportional scaling. The experiment employed an event-related design, and each trial was therefore modeled as a compound event with a stick (i.e., delta) function convolved with SPM2’s canonical hemodynamic response function. The analysis used a series of fixed-effects models on the four-way factorial design, at the individual subject level. The factors consisted of signal detection category (4 levels: hits, miss, false-alarm, correct-rejection; see Supplementary Figure 1), cue location (16 levels), report mode (2 levels: right, left), and cue delay (3 levels: short, intermediate, and long). Additionally we specified regressors to model the “uncued” discs (locations of discs which were not cued) at each of the 16 locations and two regressors to model the button presses. The coordinates of all activations are reported in MNI coordinates. Where anatomical regions are reported, this was achieved by reference to the anatomy toolbox in SPM2.

## RESULTS

Attentional fluctuations were therefore monitored through the staircasing procedure (see Methods), which was applied at all times throughout the scanning sessions. Running average of percent correct averaged over 50 trials showed very low variance over the whole session. This is evident from the fact that the staircasing program only initiated a few transitions in contrast, in order to keep subject performance within its designated range. We therefore conclude that, over time, performance did not fluctuate, and thus, if attentional fluctuations contribute to performance, there were no strong attentional fluctuations. To visualize the decay in performance as a function of cue delay, we plotted  $d'$  for each delay and stimulus position. As described previously for tasks of this type, the decay (see Figure 1D) appears to be approximately exponential in nature (Lu, Neuse, Madigan, & Doshier, 2005; Cohen & Dehaene, 1998; Tiitinen, May, Reinikainen, & Näätänen, 1994). At the short delay, each subject performed with an accuracy above  $\sim 95\%$ , resulting in very few false alarms or misses, whereas at the long delay, each subject's performance was close to chance (average accuracy 56.1%). It is thus the intermediate delay (900–1100 msec), where each subject is at threshold (defined here as 75% correct), which is of principal interest. Analysis of fMRI responses was therefore restricted to this intermediate delay condition, ensuring a balanced sampling of all signal detection categories and additionally avoiding potentially confounding effects of cue delay. The use of the short, intermediate, and long delay conditions was necessary to prevent cue predictability and

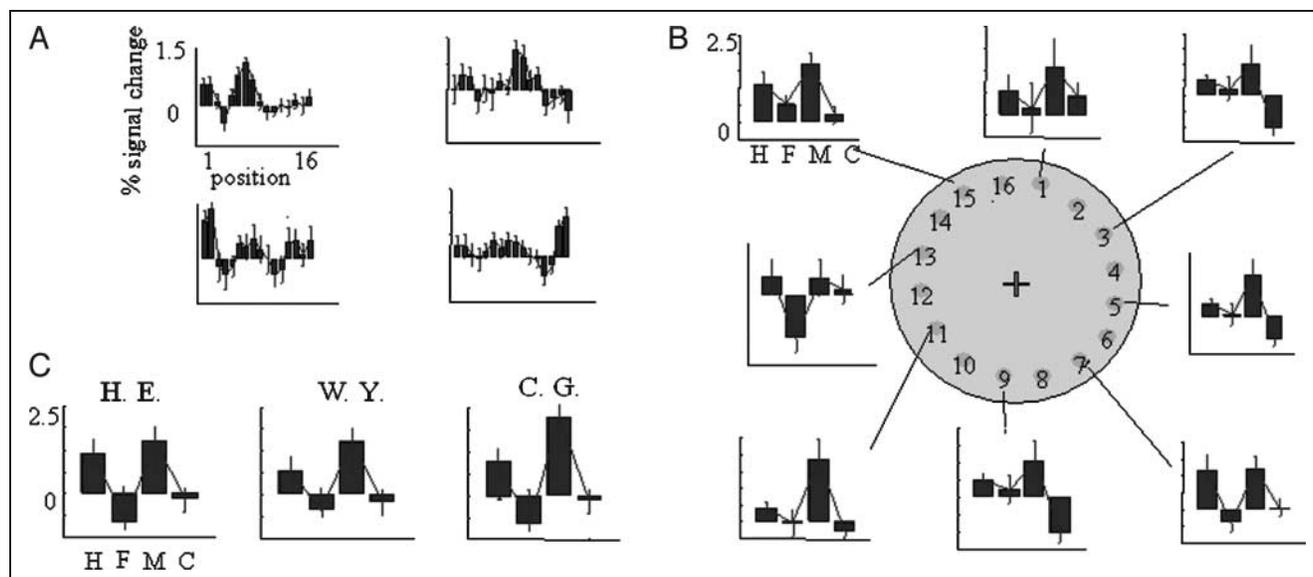
manipulate the reportability of the stimulus from near-perfect (short delay) to near-chance (long delay) reportability. The effect of delay on performance (indexed by  $d'$ ) was not coupled with an effect on decision criteria because delay had little effect on decision criteria which were approximately zero for all delays and for all subjects.

SPM analyses were then carried out to identify regions that were significantly selective for the three components of the task—the stimulus, the decision, and the motor report.

### Stimulus-specific Effects

We first located brain regions with retinotopically selective responses to each stimulus position, that is, those that topographically “map” visual space. As predicted, an  $F$ -contrast testing for the significant main effect of uncued locations showed significant activity in V1 and V2 ( $p < .001$ , uncorrected), namely, early visual cortical areas (see Supplementary Figure 2 for more details). This test was to define a region of interest in which we performed  $t$  contrasts to test for the main effect of each individual uncued stimulus location. Figure 2A shows the resulting parameter estimates for four exemplar voxels, illustrating spatial selectivity profiles.

To inquire whether these regions correlated with the stimulus rather than the perceptual decision or motor response, we used a signal detection theoretic approach (Ress & Heeger, 2003). This involves computing a “signal detection profile” for each voxel, that is, the proportion of hits, misses, false alarms, and correct rejects.



**Figure 2.** (A) Shows position selectivity for different loci in V1/V2. Criteria for V1/V2 loci are the most significant voxel for the simple main effect for a single uncued stimulus position. The figure depicts four voxels for illustration, one taken from each quadrant of retinotopic space. Plots show parameter estimates for all 16 stimulus positions for Subject H. E. (B) Upper: Graphs display signal detection profiles for each V1 loci (and therefore stimulus position) plotted around a depiction of stimulus space (odd-numbered positions only for figure clarity). (C) The average of all of these loci is plotted for each subject. For B and C, H = hits; F = false alarms; M = misses; C = correct rejects.

**Figure 3.** Depicts all areas that conform to the signal detection profile for a “yes” decision area as formalized by the contrast yes versus no collapsing across all other factors.

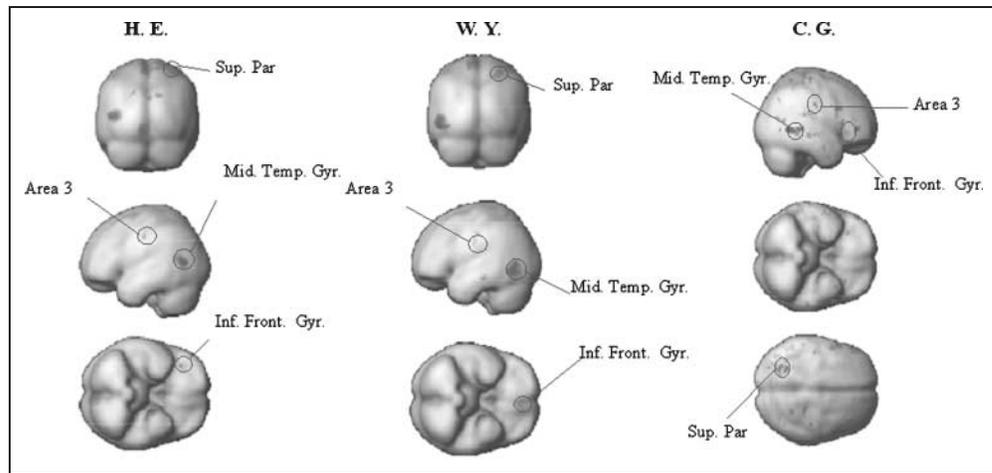


Figure 2B shows the signal detection profile for the most selective V1 voxel for each of the odd-numbered positions (for Subject H. E.). These signal detection profiles were acquired for all 16 V1 loci and averaged to give a mean V1 signal detection profile for each subject (Figure 2C). The important point to note here is that the pooled responses are significantly higher for hits and misses than they are for false alarms or correct rejections [ $t(15) = 2.78, p < .05$ ], showing that these regions reflect stimulus presence as opposed to decision (the same comparison for hits vs. misses was not significant at the same threshold).

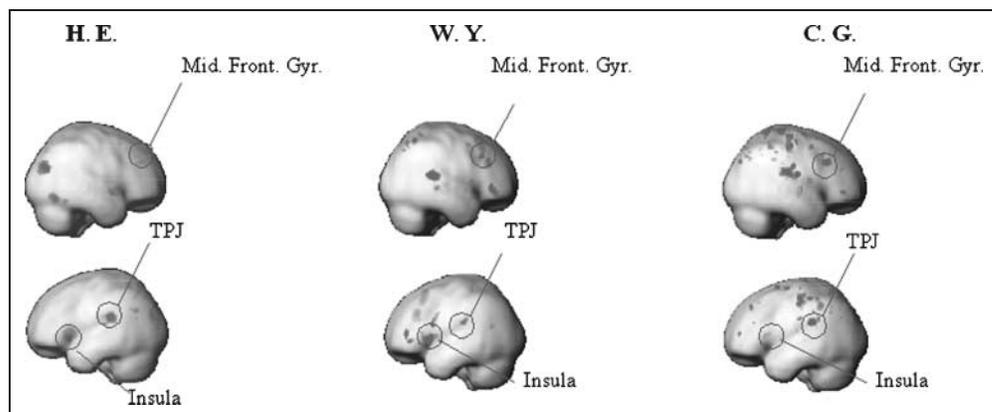
### Decision Effects

To establish which areas are selectively more active when subjects decide “yes” (hits and false positives) than when they decide “no” (misses and correct rejections), independently of stimulus processing and report, we performed a  $t$  contrast “yes > no” for the intermediate delay, averaging across all cue positions and report modes. This constitutes a main effect of subjective detection, and revealed a common pattern of activation in all three subjects that included the superior parietal

cortex, Brodmann’s area 3, and the middle temporal gyrus, as well as the inferior prefrontal cortex. The SPMs for this contrast are displayed in Figure 3 using a threshold of  $p = .001$  (uncorrected), although all regions mentioned in the text survived an adjustment to  $p < .05$ . Parameter estimates for these regions are shown in Supplementary Figure 3.

To find areas that are more active when subjects decide “no,” independently of stimulus presence and report, we performed the converse of the “yes > no” contrast (i.e., “no > yes”). In all three subjects, this revealed a different set of areas (Figure 4), which included the middle frontal gyrus, the insula, and the temporoparietal junction (TPJ). Parameter estimates for these regions are shown in Supplementary Figure 4. In summary, there is a wide network of areas, spanning frontal, parietal, and temporal lobes, which are differentially active according to the decision. Performing the same tests for position-specific decision effects (e.g., “yes > no” for Position 1) in position-specific voxels identified previously in V1/V2 revealed no significant effects at  $p < .01$  (uncorrected). In the same voxels, we tested for decision-specific effects, where controlling for differences in visual contrast, by contrasting hits versus misses (stimulus

**Figure 4.** Depicts all areas that conform to the signal detection profile for a “no” area, as formalized by the contrast no versus yes collapsing across all other factors.



present in both) and false alarms versus correct rejection (stimulus absent for both), revealed significant differences ( $F$  test,  $p < .01$  uncorrected) for only a few stimulus positions (2, 1, and 1 in Subjects 1–3, respectively). Together, the position-specific analyses show that decision-specific effects are not strongly expressed in early visual areas for this paradigm.

### Report Effects

To investigate which areas are selectively more active when subjects make a report, independently of the actual presence of the stimulus or the decision as to whether it is present or absent, we performed the  $t$  contrasts “left > right hand response” and its converse. The left response occurred for all trials in Mode 1 where subjects responded yes (“hits and false alarms”), and all trials in Mode 2 where they responded no (misses and correct rejections), and vice versa for the right response. These contrasts revealed, in all three subjects, consistent activations of area 4 (primary motor cortex), area 6 (premotor cortex), the TPJ, and the cerebellum (see Supplementary Figure 5). In addition, there was prefrontal (2/3) and parietal activity (2/3) in some subjects, but not all. Performing the same tests for position-specific report effects (e.g., “left- > right-hand response” for Position 1) in position-specific V1/V2 voxels revealed no significant effects at  $p < .01$ , uncorrected. It should be noted that the left- > right-hand contrast would not reveal areas which do not show a gross difference in hand preference, and it may be possible that differences exist also at a finer scale than can be resolved with this fMRI sequence.

### DISCUSSION

Our aim was to determine the neural correlates of stimulus reportability without any prior assumptions about its relationship to consciousness. We performed an experiment that dissociated three component stages leading to stimulus report, namely, stimulus processing, decision, and motor response. Our strategy was twofold.

First, we introduced a delay between the stimulus and the cue for the partial report task, which decreased performance below ceiling and dissociated the stimulus-bound sensory processing from the decision. The effect that the delay had on performance was indicative of reportability (in this task) being primarily dependent on iconic memory not on working memory. Partial report tasks for different features and tasks have different decay functions. A standard way to think about these functions is that they are determined by two processes, a fast decaying iconic memory for a large number of items and a much slower decaying working memory for a small number of items (e.g., see Lu et al., 2005). Exponential

temporal decay functions with half-lives on the order of less than a few seconds are generally recognized as being indicative of iconic memory (Lu et al., 2005; Landman, Spekreijse, & Lamme, 2003; Dehaene & Naccache, 2001; Sperling, 1960). Working memory is thought to have a much longer half-life on the order of 10's of seconds with some estimates suggesting the error rate peaks at  $\sim 20$  sec (Bugmann & Bapi, 2000; Ploner, Gaymard, Rivaud, Agid, & Pierrot-Deseilligny, 1998). Most partial report tasks involving letters almost fully decay within  $\sim 800$  msec (see Landman et al., 2003 for a discussion), with subjects above chance beyond that due to working memory. Letter recall is more difficult in the sense that, for each item, there is a chance performance is  $1/26$ . However, in our experiment, the stimuli are abstract discs and chance performance for detection is  $1/2$ . We are not aware of any other partial report stimuli that have used these stimuli before, however, the closest is that of Landman et al. (2003), which use a binary discrimination of oriented bars. They report that memory performance has a high capacity and remains intact for at least 1200 msec (1500 msec in experienced observers). It should be noted that another possible cause of the decay in reportability over time is due to the decay of the positional information such that, over time, mislocalizations of stimuli are more frequent.

One question, however, is whether decision-making, attentional, or other cognitive factors could also degrade performance over time. It is unlikely that the observers' ability to make an accurate decision would degrade monotonically over time, with this degradation process time-locked to the stimulus and, after  $\sim 2$  sec, the subject is almost unable to make any accurate decision. We know of no evidence or precedent for this hypothesis to be true, and it has long been dismissed as an explanation for partial report results. In addition, we calculated decision criteria (see Supplementary Table 2) for each delay and found no effect of delay on criteria. This perhaps indicates that it is the persistence of signal strength, not decisional factors, which is the cause of degraded performance. Another possibility is that, somehow, the observers' ability to attend to a particular position in space would degrade over time and, after  $\sim 2$  sec, the subject is unable to attend to anywhere at all, resulting in close-to-chance performance. Because we know that transient attentional shifts can occur every 150 to 250 msec (Carrasco & McElree, 2001; Nakayama & MacKeben, 1989) and endogenous shifts every  $\sim 250$  msec (Theeuwes, Godijn, & Pratt, 2004), then this is an unlikely explanation. Note that we are not arguing that attentional and decisional factors do not effect trial-by-trial performance, but rather that these factors do not systematically change as a function of delay. We would, of course, expect many sources of noise including those that could be bracketed as external (e.g., photon noise) or internal (attentional or decisional) to contribute to the trial-by-trial fluctuations in the subject's response. Also, the noise can corrupt the

present versus absent dimension of the signal but also the positional information. If the fidelity of the positional accuracy drifts over time due to noise, then this could also account for the degraded performance.

Secondly, to uncouple the decision from the motor report, we manipulated the hand used to make the report. Together, this allowed us to infer about the neural basis of stimulus processing, decision-making, and motor reporting independently. Below we discuss our three main findings before considering their implications for the search for the NCC.

First, we found that the early visual cortex (V1/V2) is selectively activated in a position-specific manner according to the retinotopic locations of the uncued stimulus. Using the uncued disc regressors to generate independent localizers for position-selective responses, we defined voxels in the early visual cortex that were selective for each disc position, and plotted the signal detection profiles for each of the 16 different cued loci independently (see Figure 2B). The most important result is that the average signal detection profile (across loci) for V1/V2 shows a response that does not reflect the decision or the motor report (Figure 2C). Instead, it conforms to a simple stimulus processing model, where retinotopic responses are driven by the contrast of the visual transient. In other words, the signal detection profile of V1/V2 loci shows that each locus responded to the presence of the stimulus, irrespective of the decision or the motor report. Second, we found that two distinct distributed networks correlate with decisions independently of stimulus presence or motor report and that the distribution of these networks depends on the decision made (“yes” or “no”). Across all three subjects, a decision that the stimulus was present correlated with activity in the superior parietal cortex, area 3, the middle temporal cortex, and the ventral prefrontal cortex. By contrast, a network of insula, middle frontal, TPJ, and superior parietal regions was more active when subjects decided that the stimulus was absent at the cued location. Third, we found that a network of motor regions, including the primary motor cortex, the premotor cortex, the TPJ, and the cerebellum, correlated with the motor act of reporting independently of all other factors.

The finding that V1 correlates with stimulus presence, independently of decision, may appear at first to contradict a previous study by Ress and Heeger (2003) (henceforth, R + H). This study reported that V1 activity correlates with decisions and not stimulus presence (Ress & Heeger, 2003), by showing that the response to hits and false alarms was greater than that to misses and correct rejections. In contrast, our results show that the response to hits and misses is greater than to false alarms and correct rejections. Our study used a similar signal detection approach, under conditions in which psychophysical performance for a visual detection task was at threshold. The key difference between our experiment and theirs, which leads to a difference in

interpretation, lies in the way in which threshold performance is achieved. In the R + H paradigm, performance was degraded by attenuating the contrast of the stimulus. As R + H point out, early noise in the visual system up to V1 would predict overlapping noise and signal + noise distributions in V1 (i.e., the distributions of neuronal activity under stimulus-absent and stimulus-present conditions, respectively). When a subject attempts to impose a decision criterion upon those overlapping distributions, erroneous reports are inevitable. The recorded signal detection profile that R + H find is in accord with a simple prediction made from selective averaging of activity above and below the decision criteria, such that activity is higher when subjects decide “signal is present” than when they decide “signal is absent.” This interpretation predicts that areas beyond V1, which are presumably accessing this information, would reflect a similar pattern of activity, as all regions would inherit the same noisy distributions. Fronto-parietal regions critical for deciding and responding to the stimulus may be expected to display similar response profiles as the visual cortex (in this context). This is supported by another study (Pins & Ffytche, 2003), which found that a contrast of hits versus misses invoked greater activity in both V1 and fronto-parietal regions.

We, on the other hand, degraded reportability performance by increasing the delay between the stimulus and the cue. Thus, performance was affected by variability not in the stimulus-bound transient, but in the activity supporting either the iconic memory trace or its attentional access. The fact that our analysis did not pick up any decision effects in V1/V2 or, indeed in any visual area, implies that the change in signal detection profiles from a *stimulus-specific* profile to a *decision-specific* profile occurs at some point beyond V1/V2. According to this proposal, a decision criterion imposed on an overlapping distribution beyond V1/V2 is then propagated to the regions accessing the decision and performing the appropriate report. We characterize this descriptive account as being consistent with a *late-noise* model. Our results thus address a weakness in many previous studies that failed to dissociate stimulus-specific processing from decisions and motor reports, and identify networks of areas whose activations should be interpreted as reflecting *components* of report rather than as general markers of consciousness. We argue that performing this dissociation depends critically on the type of task, which in turn can determine whether performance is limited by early or late noise. However, as stated before, we make no a priori assumption about which component of reportability signifies consciousness. We conclude by discussing our results with reference to two competing theories of consciousness. According to the global neuronal workspace hypothesis (GNW) (Dehaene & Naccache, 2001), a sensory area will not contribute to perception until it is mobilized by the GNW through an “ignition” process; only those stimuli which are se-

lected by attention are “ignited,” and thus, consciously perceived. Accordingly, only stimuli at cued locations are consciously perceived, whereas those stimuli that are uncued remain “preconscious” (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006). Under this framework, the position-selective activity reported in V1/V2 for the uncued locations corresponds to a preconscious correlate. An opposing stance is that of the phenomenal versus access (PA) theory of consciousness (see Block, 1990, 2005; Lamme, 2003). This theory proposes a distinction between phenomenal consciousness, which encompasses the qualitative aspects of consciousness and access consciousness, which encompasses one’s ability to access phenomenal experiences, such that they are “poised for voluntary action and reason.” PA theory would propose that subjects have a simultaneous phenomenal experience of all stimuli in the disc array but that they only have access consciousness of those stimuli that are cued and, therefore, attended. This supports a commonly held intuition that subjects in partial report tasks such as this one perceive all elements shortly after stimulus presentation, and that this is associated with very high reportability for the short delay condition (as in Figure 1), but that increasing the delay of the cue, while decreasing performance, does not affect the original perception. Under this framework, our data could be interpreted as showing that V1/V2 correlates with phenomenal consciousness (consistent with previous formulations of micro-consciousness theory; Zeki & Bartels, 1999), whereas activity in fronto-parietal regions correlates with access consciousness (defined by the decision).

The disagreement between the two theories comes down to an interpretation of what conscious status potentially reportable stimuli attain. GNW proposes they are “preconscious” and PA claims they are “phenomenal.” The debate about whether the activity we record in V1/V2 reflects preconscious or phenomenal processes, and whether the decision network activity reflects access consciousness or just conscious perception, is one that requires further experimentation and argument. Ultimately, we argue that until this issue can be resolved, it is more appropriate for neuroscience to restrict the interpretations of NCC experiments to the operational marker used; not what they interpret it as standing for. Even when interpreting NCC data purely in terms of the operational marker, it is crucial to delineate the neural correlates of its components. In this study, we show how this might be achieved.

### Acknowledgments

This work was done through a programme grant from the Wellcome Trust, London to S. Z.; O. H. held an MRC Ph.D. scholarship. We thank Stewart Shipp, Louise Whiteley, and Barrie Roulston for their helpful comments.

Reprint requests should be sent to Oliver J. Hulme, Wellcome Trust Centre for Neuroimaging, University College London,

Gower Street, London, WC1E6BT, UK, or via e-mail: o.hulme@ucl.ac.uk.

### REFERENCES

- Block, N. (1990). Consciousness and accessibility. *Behavioral and Brain Sciences*, *13*, 596–598.
- Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, *9*, 46–52.
- Bugmann, G., & Bapi, R. S. (2000). Modelling relative recency discrimination tasks using a stochastic working memory model. *Biosystems*, *58*, 195–202.
- Carrasco, M., & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences, U.S.A.*, *98*, 5363–5367.
- Chalmers, D. (1998). On the search for the neural correlate of consciousness. In A. K. S. Hameroff & A. Scott (Eds.), *Toward a science of consciousness: II. The second Tucson discussions and debates*. Cambridge, MA: MIT Press.
- Cohen, L., & Dehaene, S. (1998). Competition between past and present. Assessment and interpretation of verbal perseverations. *Brain*, *121*, 1641–1659.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. R. M., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, *10*, 204–211.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.
- Haynes, J.-D., Deichmann, R., & Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. *Nature*, *438*, 496–499.
- Haynes, J.-D., Driver, J., & Rees, G. (2005). Visibility reflects dynamic changes of effective connectivity between V1 and fusiform cortex. *Neuron*, *46*, 811–821.
- Kleinschmidt, A., Buchel, C., Zeki, S., & Frackowiak, R. S. (1998). Human brain activity during spontaneously reversing perception of ambiguous figures. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *265*, 2427–2433.
- Lamme, V. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*, 12–18.
- Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*, 149–164.
- Leopold, D. A., & Logothetis, N. K. (1996). Activity changes in early visual cortex reflect monkeys’ percepts during binocular rivalry. *Nature*, *379*, 549–553.
- Logothetis, N. K., & Schall, J. D. (1989). Neuronal correlates of subjective visual perception. *Science*, *245*, 761–763.
- Lu, Z.-L., Neuse, J., Madigan, S., & Doshier, B. A. (2005). Fast decay of iconic memory in observers with mild cognitive impairments. *Proceedings of the National Academy of Sciences, U.S.A.*, *102*, 1797–1802.
- Lumer, E. D., Friston, K. J., & Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. *Science*, *280*, 1930–1934.
- Macmillan, N. A., & Creelman, C. D. (1996). *Detection theory: A user’s guide*. Hillsdale, NJ: Erlbaum.
- Michelson, A. (1927). *Studies in optics*. Chicago: University of Chicago Press.
- Nakayama, K., & MacKeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research*, *29*, 1631–1647.
- Pins, D., & Ffytche, D. (2003). The neural correlates of conscious vision. *Cerebral Cortex*, *13*, 461–474.

- Ploner, C. J., Gaymard, B., Rivaud, S., Agid, Y., & Pierrot-Deseilligny, C. (1998). Temporal limits of spatial working memory in humans. *European Journal of Neuroscience*, *10*, 794–797.
- Ress, D., & Heeger, D. J. (2003). Neuronal correlates of perception in early visual cortex. *Nature Neuroscience*, *6*, 414–420.
- Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience*, *8*, 1391–1400.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, *74*, Whole issue.
- Theeuwes, J., Godijn, R., & Pratt, J. (2004). A new estimation of the duration of attentional dwell time. *Psychonomic Bulletin & Review*, *11*, 60–64.
- Tiitinen, H., May, P., Reinikainen, K., & Näätänen, R. (1994). Attentive novelty detection in humans is governed by pre-attentive sensory memory. *Nature*, *372*, 90–92.
- Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, *21*, 753–759.
- Zeki, S., & Bartels, A. (1999). Toward a theory of visual consciousness. *Consciousness and Cognition*, *8*, 225–259.