

Relative Blindsight in Normal Observers and the Neural Correlate of Visual Consciousness

Hakwan C. Lau^{1,2}, Richard E. Passingham^{1,2}

¹ Wellcome Department of Imaging Neuroscience, University College London

² Department of Experimental Psychology, University of Oxford

In press in PNAS

Content: 20 pages of (including references and figure legends), 5 figures
Character count: 37181 (all text) + 180*(6+9+6+7+7) + 120*5*2 (figures) = 44681
Abbreviations: fMRI (functional magnetic resonance imaging),
mid-DLPFC (mid-dorsolateral prefrontal cortex)

Correspondence:
Hakwan Lau
Functional Imaging Laboratory
12 Queen Square, London
WC1N 3BG, United Kingdom
Tel: 44 20 7833 7486
Email: hakwan@gmail.com

Abstract

Using a novel paradigm based on metacontrast masking, we created experimental conditions where the subjective report of consciousness differs, but the objectively measured ability to discriminate visual targets does not. This allowed us to study the neural correlate of consciousness (NCC), while having performance levels carefully matched, for the first time in healthy human subjects. A comparison of the neural activity associated with these conditions, as measured by functional magnetic resonance imaging (fMRI), showed that conscious perception is associated with spatially specific activity in the mid-dorsolateral prefrontal cortex (area 46). Further analysis confirms that this activation is not only free from any performance confound, but is also not driven by differences in the timing of the physical stimuli. Our results suggest that the prefrontal cortex is important for the essentially subjective aspects of conscious perception.

Introduction

Blindsight refers to the phenomenon that after a lesion to the primary visual cortex, the subject can exhibit above-chance performance in detecting or discriminating the visual stimuli in a forced-choice setting, despite the lack of acknowledged consciousness of the stimuli (1, 2). In some instances, blindsight subjects can perform at an impressively high level of accuracy (above 80%) in the forced-choice task, even when the subjects believe that they are guessing. This potentially high level of performance makes blindsight an interesting case study for visual consciousness, as it indicates that it is consciousness but not the basic capacity to process information that is completely abolished. This dissociation of visual consciousness and performance also allows for the opportunity to compare between conditions where visual consciousness is present in one but not the other, while having performance levels matched (3). This approach is particularly promising in regard to the search of the neural correlate of consciousness (NCC) because the resulting difference between the conditions cannot be explained in terms of a mere difference in performance levels. This means that the result would be more likely to reflect some essential properties of consciousness per se, rather than simply due to a difference in the effectiveness of information processing in regard to the visual stimuli.

Here we try to apply this approach to normal human subjects, and use functional magnetic resonance imaging (fMRI) to uncover the NCC. Attempts to unequivocally demonstrate "blindsight" in normal observers have proved to be controversial, resulting in findings that are sometimes hard to replicate (4-6). However, this difficulty dissolves when we consider the fact that in brain imaging studies, a statistical comparison between two conditions is valid so long as the factor

in question differs in relative terms: we need not produce conditions in which visual consciousness is either completely present or absent. Instead of looking for a complete dissociation of performance and visual consciousness as in the case of blindsight, we set out to look for a relative difference in the level of visual consciousness in two conditions where performance levels are matched. We call this phenomenon "relative blindsight", because as in blindsight subjects, there is a dissociation between consciousness and performance, although here the dissociation is relative across conditions.

We demonstrate that these conditions can be created by using a novel psychophysical paradigm based on metacontrast masking. This relies on the fact that type II metacontrast masking (7) gives rise to a non-linear, U-shaped, masking function. In metacontrast masking, a figure that overlaps with the contour of the target is presented after the target. Discrimination performance for a target stimulus decreases and then increases as the temporal distance between the target and a metacontrast mask increases gradually; this distance is referred to as stimulus onset asynchrony (SOA). Figure 1 shows the schematic diagram for a visual discrimination task that involves metacontrast masking. Typically, when the mask is effective, there is a perceived decrease in the clarity and luminance of the target. The masking function characterizes the relationship between SOA and masking effectiveness. The u-shaped nature of the masking function (e.g. Figure 2, upper panel) implies that there are two SOA points at which the same performance level can be found. We conjectured that the subjective level of visual consciousness might not be exactly the same at these SOA points, and this was confirmed by a psychophysical study. Based on this behavioural finding, we conducted an fMRI study to uncover the neural correlate of visual consciousness.

Results

Behavioral Study

We studied both the performance and the subjective report for conscious perception in a metacontrast masking paradigm (figure 1), to demonstrate behaviorally the phenomenon of relative blindsight. Volunteers had to make forced-choice judgments as to whether the target was a square or a diamond. We assessed the level of consciousness by asking the participants, immediately after they made the forced-choice discrimination, to press keys indicating whether they actually saw the identity of the target or simply guessed what it was. Figure 2 (upper panel) showed the data for the 8 participants. At the SOA points 33ms and 100ms, the performance levels (i.e. accuracy rates for the square vs. diamond discrimination) were very similar. However, the subjective judgment of consciousness differed ($p=0.036$, 2 tailed; Wilcoxon signed-rank test), in that at the earlier SOA point volunteers were more likely to claim to have just guessed the answers. Therefore, we confirmed our hypothesis: the subjective level of consciousness can differ in the absence of a difference of performance levels.

fMRI Study

Fourteen volunteers participated in this experiment. The data for one participant were discarded because of unsatisfactory behavioral performance. For details of methodology please see the supporting online material. The task was the same as in the behavioral experiment (Figure 1). Inspection of the data in that experiment suggested that the temporal profile of the masking function differed across subjects. To maximize power, in the fMRI study we therefore recreated the conditions according to the individual temporal profile of each participant. Every participant was

tested with both a short SOA condition where the SOA was at 33 ms, and a long SOA condition where the SOA were computed online to ensure that the performance level in this condition was similar to that of the short SOA condition. To achieve this, we used an adaptive procedure to search for a long SOA time point for each participant such that performance level in both the short and long SOA conditions would be as similar as possible. This bypassed the problem of the variability of the temporal profile of the masking function mentioned above. The group average for the resulting long SOA value was 104 ms (s.d. = 40 ms).

The behavioral results are shown in the lower panel of Figure 2. The participants were on average 68.27 % correct (s.d. 4.54 %) in the short SOA condition, and 70.22 % (s.d. 5.24 %) correct in the long SOA condition. The two measures did not differ significantly ($p = 0.213$, 2 tailed; Wilcoxon signed-rank test). We also used the criterion-free measure, d' , to assess performance, and found that they did not differ significantly across the conditions either ($p=0.26$, paired t-test, 2-tailed): they were 0.99 and 1.12 for the short and long SOA conditions respectively (s.d. 0.25 and 0.31 respectively). Also, the reaction times for visual discrimination in both conditions did not differ significantly ($p = 0.858$, paired t-test, 2-tailed): they were 723 ms and 717 ms for the short and long SOA conditions respectively (s.d. = 180 ms and 192 ms respectively). The reaction times for both correct and incorrect trials were considered together so that this reflected the difference between all trials in the two task conditions, as they were all analyzed in the fMRI analysis. These results suggest that not only was the performance of identification in the two conditions matched, the task difficulty as measured by reaction times was very similar as well.

However, as predicted, the levels of subjective consciousness differed significantly ($p = 0.007$, 2 tailed; Wilcoxon signed-rank test). The 'seen' rates were

49.17 % and 60.13 % for the short and long SOA conditions respectively (s.d. 23.17 % and 19.79 % respectively). All but one participant showed the effect in the same direction, and 7 of the 13 participants showed a significant effect when their data were analyzed individually ($p < 0.05$, chi-square test).

In the long SOA condition there was a significantly higher level of subjective awareness as compared to the short SOA condition, and this was associated with a significant increase in activity in the left mid-dorsolateral prefrontal cortex (mid-DLPFC, coordinate = -46, 48, 14). This analysis included both the correct and incorrect trials. The significance level of this activation was assessed by directly estimating the probability that the area under curve for the haemodynamic time course is bigger in the long SOA than in the short SOA condition. This Bayesian approach has the advantage of bypassing the problem of multiple statistical comparisons. The probability of activation was 0.97. The location of this activation and the time courses of activity for the two conditions are shown in figure 3. This activation lies specifically in Brodmann's area 46, the macaque monkey homologue of which lies around the principal sulcus (8).

We have also performed a conventional analysis with standard parameters, fitting the canonical haemodynamic response function (HRF) to the data, and assessed the significance by using classical frequentist statistics (the t -test). This approach is less sensitive because the canonical haemodynamic response function peaks at 5-6 s, and the present data do not support this assumption (figure 3). Nonetheless, the activation in the left mid-DLPFC survived a threshold of $p < 0.001$ (uncorrected). This suggests that the activation was not due to any idiosyncratic features of the initial analysis.

Although the performance levels did not differ significantly between the two conditions, the averaged value was higher in the long SOA condition than in the short SOA condition by about 2%. To ensure that the activation in the mid-DLPFC was not due to this difference in performance, we performed another analysis, excluding the 3 participants who showed the biggest magnitude of superior performance in the long SOA condition. This resulted in a reversal of the trend for performance: the averaged performances in the short SOA condition and in the long SOA condition in the remaining 10 participants were 69.58% correct and 69.33% correct respectively (they did not differ significantly, $p = 0.890$). Similarly, the trend reversed for the averaged d -primes for the short SOA and long SOA conditions, which were now 1.07 and 1.06 respectively (they did not differ significantly, $p=0.940$). However, the difference in subjective level of consciousness remained: the ‘seen’ rates for the short and long SOA conditions were 55.04% and 62.79% respectively (difference significant, $p = 0.034$). The activation in the left mid-DLPFC also remained significant (probability > 0.95).

We also tried to test whether parietal activity differed in the two SOA conditions, perhaps by a degree that was only marginally below threshold. We lowered the threshold to probability > 0.9 for the comparison between the two SOA conditions and found no parietal activity. We also tested this using the canonical HRF approach, and at a liberal threshold of $p < 0.01$ uncorrected we did not find any parietal activity.

We performed another analysis to test whether the activity in the left mid-DLPFC merely reflected the physical difference in the stimuli (the SOA), or genuinely reflected a difference in the subjective level of consciousness. We took the activity in the mid-DLPFC for each subject, at the peak of the activation for each

subject, for the short SOA condition, because the SOA in this condition was the same for every subject. The variability (standard deviations) in the peak location for the subjects were 6 mm, 10 mm, 10 mm in the x, y and z directions respectively. We tested whether the level of activation correlated with the level of subjective consciousness across subjects. We found a correlation indicating that the higher the activity in the left mid-DLPFC, the higher the percentage of trials the subjects declared that they were conscious of the identity of the stimuli ($r = 0.512$). Because we predicted the direction of the correlation based on the results of the comparison between the short and long SOA conditions, a one-tail statistical test is justified. Under this test the correlation was significant ($p=0.037$). This is plotted in figure 4.

To identify the brain areas where activity reflected discrimination performance in general, we also analyzed the fMRI results by comparing correct and incorrect trials, combining short SOA and long SOA trials. We found widespread activations in the ventral prefrontal, premotor, parietal, temporal cortices. There were also activations in the basal ganglia and cerebellum which may relate to the fact that the reaction times for correct trials (683 ms) were quicker than for incorrect trials (815 ms) ($p=0.0001$, t-test, 2-tailed). The cortical activations were similar to those previous identified as parts of the NCC (9), as shown in figure 5. These include activations in the right ventrolateral prefrontal cortex (VLPFC, coordinate = 36, 40, 2, probability of activation = 0.98), right intra-parietal sulcus (IPS, coordinate = 42, -58, 38, probability of activation = 0.99) and middle temporal gyrus bilaterally (MTG, coordinates = 50, -36, -8 and -48, -36, -6, probability of activation = 1.00 and 0.99 respectively). However, at these coordinates, the probabilities that the activity reflected the difference in the subjective level of consciousness (i.e. long SOA > short

SOA) were in general low: 0.45 (right VLPFC), 0.59 (right IPS), 0.35 (left MTG) and 0.48 (right MTG).

Discussion

Inspired by the phenomenon of blindsight, we have demonstrated a phenomenon which we call “relative blindsight”: subjects can be performing equally well in two conditions in terms of their accuracy and reaction times in a forced-choice discrimination task, and yet, show a relative difference in the frequency of their acknowledged consciousness for identity of the stimulus. The stimuli used in the two conditions are similar, with only a subtle difference in the temporal distance between the target and the metacontrast mask (the SOA). This allowed for a comparison of different levels of consciousness even when performance level is controlled. This comparison revealed specific brain activation in the left mid-DLPFC. This corresponds to Brodmann's area 46 (8).

One possible criticism of blindsight is that after the brain lesion, subjects adopt an extremely conservative criterion for declaring the stimuli to be consciously perceived. Here, however, the two conditions were intermixed within the same experiment and presented to the same subjects, and also the stimuli were similar and presented at the same spatial location. Even if the difference in the frequency of acknowledged consciousness could be formulated in terms of a difference in criterion setting, this difference is not trivial: one still needs to explain why the same individual would set the criteria differently in the different conditions. It seems that the most likely explanation seems to be that the subjects are actually more frequently conscious of the identity of the stimuli in the long SOA condition than in the short SOA condition.

It could be argued that the activity in the mid-DLPFC merely reflects the difference in the physical properties (i.e. the SOA) of the stimuli in the two conditions. This criticism, of course, also holds for many previous studies of the neural correlates of consciousness (NCC). However, a previous fMRI study of metacontrast masking suggests that stimulus timing driven effects can only be found in the early visual areas, but not in the prefrontal cortex (10). We therefore think it is unlikely that the activation in the left mid-DLPFC is simply driven by the SOA difference. So as to check, we have also performed an across-subjects correlation analysis to show that, given the same SOA, the level of activity in the left mid-DLPFC was correlated with the percentage of trials where the subjects claim to be conscious of the identity of the stimuli. This means the results we found in the prefrontal cortex cannot be explained by a stimulus confound.

Given that the activation is unlikely to be due to a stimulus confound, it is tempting to generalize the results to other stimuli involving different masking methods. However, to unequivocally demonstrate this is not an easy task because the way we achieved the matching of performance levels depends on the special property of metacontrast masking that it gives rise to a u-shaped masking function. One reason to speculate that the present result is specific to metacontrast masking is that the mid-DLPFC has been formally characterized (11) as a final converging point for information from both the dorsal and ventral visual streams (12, 13). It has been proposed that the mechanism of metacontrast masking depends critically on the interaction between the magnocellular and parvocellular pathways, which roughly maps to the dorsal and ventral streams respectively (7). However, an imaging study of blindsight suggest that the prefrontal cortex may be involved in consciousness in general (14).

Sahraie et al (14) have reported results related to the present study from an experiment on a blindsight subject (known as GY). The subject has a lesion to the primary visual cortex that affects roughly half of his visual field, stimuli presented to which yield no phenomenal visual awareness. The authors presented to this “blind field” slowly moving (3 deg/s) stimuli of which the subject was unaware, and found that the subject could nonetheless discriminate the direction of the horizontal movement at slightly above 80% correct. This visual stimulation was associated with a lack of significant activation in Brodmann’s area 46. However, when the speed of the movement of the stimuli was increased to 20 deg/s, the subject reported a sense of awareness even though the visual presentation was to the “blind field” (a phenomenon known as type II blindsight, see ref 1 and 2), and the performance of discrimination was above 90% correct. This visual stimulation was also associated with a significant activation of Brodmann’s area 46. The performance levels for discrimination task in these conditions were different, but could be considered roughly matched, because they were both well above chance. This suggests that activity in area 46 may be important for the awareness of visual stimuli, even when performance level is not an important contributing factor. Incidentally, when visual stimuli were presented to the unimpaired field of the blindsight subject, there was also significant activation in area 46. Along with the currently reported findings, these results are compatible with the dominant theoretical views that awareness depends critically on anterior regions of the brain (15-17), and possibly only a limited number of neurons located in a restricted region (18).

As in our result, a recent study on perceptual decision making (19) also found an activation that was specific to the left mid-DLPFC. The authors suggest that the activation was associated with the computational mechanism underlying a perceptual

decision task, in such a way that the intensity of the activation reflects the contrast of the strength of the signal of the competing stimuli. When the contrast was high, the stimulus was less ambiguous, and the participants were presumably more conscious of its identity. Also, they performed better when the contrast was high. It was found that in this situation there was more activity in the left mid-DLPFC. The similarity between that result and the present one is intriguing because it highlights the possibility that conscious report is essentially related to a computational mechanism that is general for perceptual decisions. However, the critical difference between the two studies is that performance was controlled in the present study, which means that even if conscious reports are related perceptual decisions, they do not necessarily depend on the same mechanism that drives basic visual discrimination performance.

If the conscious reports studied in this experiment depend on a decision mechanism, one could argue that they do not reflect the phenomenal experiences *per se*, but only reflect higher cognitive aspects of consciousness. One might therefore seek to look for correlates of phenomenal experiences in lower or early cortical areas such as the visual cortex. However, the present study shows that performance could be an important confound in the study of consciousness, so one must be cautious in ensuring that the earlier correlates do not simply reflect performance. In many visual tasks, performance is also characterized by reports, only they are usually forced-choice reports of properties of the stimuli. Our conscious reports (the Seen vs Guess judgements) are however reports of the experiences themselves. The phenomenal experiences themselves are not directly observable, but studying their immediate subjective reports is still closer than studying the reports of something else (i.e. the stimuli).

Although activity in the parietal cortex was found to be related to performance in general, it did not significantly differ between the conditions where the subjective criteria for conscious perception differed. It could be argued that the lack of significant activity in the parietal cortex could be due to a lack of statistical power. However, we failed to find activity even when we lowered the threshold to a very liberal level ($p < 0.01$ uncorrected). This seems to differ from the results of many previous NCC studies, which have found parietal activity to be as significant as the prefrontal activity, if not more so (9, 20-25). However, it is possible that some of those previously reported activations in the parietal cortex may reflect performance, as typically a difference in consciousness level is associated with a difference in performance. Future work should aim at further clarifying this issue.

It has been argued that a lesion of the prefrontal cortex does not affect visual consciousness (26). However, in most typical experiments we assess objectively-measured performance levels in a visual task, and we argue that it is the subjective reports of consciousness that is mediated by the mid-DLPFC. Latto and Cowey (27) have made bilateral lesions on the posterior lateral surface (i.e. frontal eye fields; Brodmann's area 8) of the prefrontal cortex in the macaque monkey brain, and they report a significant increase of luminance threshold. This suggests that the prefrontal cortex may indeed be essential to some aspects of visual consciousness. This issue could be further clarified by giving the task used in the present study to patients with lesions that include the mid-DLPFC, and we are currently carrying out this test. This and related future work should be able to contribute further to our understanding of the role of the prefrontal cortex in subjective conscious perception.

Materials and Methods

Behavioural Study

Eight people participated in this study. At the beginning of the experiment, a fixation cross (0.5°) was presented at the centre of a computer screen (20-inch LCD monitor, 60 Hz refresh rate), at which the participants were instructed to fixate their gaze whenever the cross was present. The disappearance of the cross marked the beginning of a trial, and also served as a ready signal. Five hundred ms afterwards, in most trials (where $SOA > 0$, see next paragraph), either a square (1°) or a diamond (i.e. a square of the same size rotated by 45 degrees) was presented at the centre of the screen (duration = 33 ms). This was followed by variable blank period, and then a metacontrast mask (2° , Figure 1, main text, duration = 50 ms), which overlaps with part of the contour of the target, but not the target itself. Then, the screen went blank until 850 ms after the beginning of the trial. In the other trials (where $SOA < 0$), the order of presentation of the mask and the target (square or diamond) was reversed.

After the variable blank period, matched for all SOA conditions, a question appeared at 3° above the centre of the screen (font = Helvetica, font size set at maximum height = 1°), which either read “Diamond or Square?” or “Square or diamond?” (counter-balanced across participants). In the former case, the participant pressed the left key if a diamond was thought to be presented and the right key if a square was thought to be presented; the opposite was required in the case of “Square or diamond?” After the response, a second question appears at 3° degree below the centre of the screen (same font and font size), which read “Seen or Guessed?” The participant pressed the left key if the target was actually consciously perceived, and the right key if it was not, and that the previous responding was based on intuition. After the response, the fixation cross reappeared. The trial ended at 3650 ms after the beginning. If the participants failed to respond to both questions after this time, the

question disappeared, the fixated reappeared, and the trial was discarded. There was an additional inter-trial interval of 350 ms, such that including its following interval, each trial took exactly 4000 ms in total. The participants were told that accuracy, but not speed, was important, so long as they were fast enough to avoid failing to respond within the 2800 ms limit in a significant amount of trials.

The SOA, stimulus onset asynchrony, was the difference between the onset of the target and the mask, and they are -50 ms, -33 ms, -17 ms, 0 ms, 17 ms, 33 ms, 50 ms, 67 ms, 83 ms, 100 ms, 117ms, 133 ms. These time points were selected to accommodate for the refresh rate limitation of the projector used in the fMRI scanner for the main experiment. The projector operates at 60 Hz. There were 80 trials for each time point, and they were presented in a randomized order. The participants were given a chance to take a break for as long as they wished after every 40 trials.

fMRI Study

Fourteen right-handed people participated in this study (nine female, mean age = 24.2, s.d. = 2.4 years, one was discarded from the analysis; see below). They all gave informed consent, and were screened for MRI safety before they participated in this experiment. The task was similar to the one used in the pilot study, but only backward masks ($SOA > 0$) were used, so that the mask always followed the target. The participants were given a short training (about 3 mins) outside of the scanner before the fMRI experiment began, to make sure that they understood and were comfortable with the tasks.

In the scanner, the participants viewed the stimuli on a projector screen through inverting mirrors, and they responded using an MRI-compatible button box. An adaptive staircasing procedure (28) was used to ensure that every participant gave

similar performance levels in the two experimental conditions. In the short SOA condition, the SOA was set at 33 ms. The SOA of the long SOA condition was 100 ms at the beginning of the experiment, and was continually updated throughout the experiment according to a 2-down-1-up rule: the SOA decreased by one step (16.6 ms) after 2 successive correct responses, and increased by one step after an incorrect response. This procedure ensures that the target performance level would be at about 71%, but it requires that there is a monotonic relationship between SOA and performance. This was ensured in by constraining that the SOA in the long SOA condition could not fall below 50 ms. To ensure that the short SOA condition matched the long SOA condition, a similar staircasing procedure was used which computed the size of the stimuli online throughout the experiment. Varying size instead of SOA for the short SOA condition has the advantage that this SOA would be fixed at 33 ms and thus would never be too close to that of the long SOA condition. Also, by making sure that even at $SOA = 33$ ms the performance would be at about 71%, we ensured that the stimuli were sufficiently small so that the masking function cannot be flat at the ceiling level for all SOAs. The size of the target started as 1 degree wide, as in the pilot study, but it decreased by one step (0.05°) after 2 successive correct responses, and increased by one step after an incorrect response; the size of the mask was scaled accordingly. To ensure that the sizes of the stimuli for the two conditions were the same, a long SOA trial always used the same value obtained by the staircasing procedure from the previous short SOA trial. This procedure similarly ensures that the target performance level for the short SOA condition would be at about 71%. The minimum target size was set at 0.25 degree. The maximum was at 3 degree. One participant remained at the ceiling level for most of the experiment, and her overall

performance was therefore far from 71% in the short SOA condition. The data for this participant were therefore discarded from the analysis.

To maximise the efficiency for detecting haemodynamic differences between the two conditions, the trials were organised such that 6 trials of the same condition would be presented in succession after the presentation of a series of 6 trials of the other condition. (The fMRI data analysis was however event-related.) In the train of 6 short SOA trials, the size of stimuli could vary, and in the train of 6 long SOA trials the SOA could vary. These variations were based on the staircasing procedure described above and could last throughout the experiment, in order to ensure the performance levels were constant across conditions and time. However, in practice, the parameters stabilized quickly in the first few minutes. Participants in total performed 240 trials for each condition (short SOA and long SOA). These were divided into 4 runs, between which the participant remained inside the scanner, but could communicate with the experimenter. Within each run, there was also a compulsory break of 26 s given to the participant after every 40 trials, during which the participant was instructed to relax. The entire fMRI experiment lasted for about 45 mins for each participant, including the breaks between experimental runs.

Further details for the procedures of imaging data acquisition and analysis can be found in the supplementary online material.

Acknowledgements

We thank Chris Frith, Alan Cowey and John-Dylan Haynes for discussions, and Ray Dolan and Sara Bengtsson for comments on an earlier version of the manuscript. This work is supported by the Wellcome Trust.

Reference

1. Weiskrantz, L. (1986) *Blindsight. A case study and implications*. (Oxford University Press, Oxford, U.K.).
2. Weiskrantz, L. (1999) *Consciousness Lost and Found* (Oxford University Press, Oxford).
3. Weiskrantz, L., Barbur, J. L. & Sahraie, A. (1995) *Proc. Natl. Acad. Sci. U. S. A* 92, 6122-6126.
4. Kolb, F. C. & Braun, J. (1995) *Nature* 377, 336-338.
5. Morgan, M. J., Mason, A. J. & Solomon, J. A. (1997) *Nature* 385, 401-402.
6. Robichaud, L. & Stelmach, L. B. (2003) *Psychon. Bull. Rev.* 10, 206-209.
7. Breitmeyer BG (1984) *Visual Masking: an Integrative Approach* (Oxford University Press).
8. Petrides, M. & Pandya, D. N. (1999) *Eur. J. Neurosci.* 11, 1011-1036.
9. Rees, G., Kreiman, G. & Koch, C. (2002) *Nat. Rev. Neurosci.* 3, 261-270.
10. Haynes, J. D., Driver, J. & Rees, G. (2005) *Neuron* 46, 811-821.
11. Young, M. P. (1992) *Nature* 358, 152-155.
12. Ungerleider, L. G. & Haxby, J. V. (1994) *Curr. Opin. Neurobiol.* 4, 157-165.
13. Goodale, M. A. & Milner, A. D. (1992) *Trends Neurosci.* 15, 20-25.
14. Sahraie, A., Weiskrantz, L., Barbur, J. L., Simmons, A., Williams, S. C. & Brammer, M. J. (1997) *Proc Natl Acad Sci U S A* 94, 9406-11.
15. Crick, F. & Koch, C. (1998) *Cereb. Cortex* 8, 97-107.
16. Crick, F. & Koch, C. (2003) *Nat. Neurosci.* 6, 119-126.
17. Dehaene, S., Sergent, C. & Changeux, J. P. (2003) *Proc Natl Acad Sci U S A* 100, 8520-5.
18. Crick, F. & Koch, C. (2000) in *The Neuronal Correlate of Consciousness*, ed. Metzinger, T. (MIT Press, Cambridge, MA), pp. 103-110.
19. Heekeren, H. R., Marrett, S., Bandettini, P. A. & Ungerleider, L. G. (2004) *Nature* 431, 859-862.
20. Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B. & Riviere, D. (2001) *Nat. Neurosci.* 4, 752-758.
21. Lumer, E. D. & Rees, G. (1999) *Proc. Natl. Acad. Sci. U. S. A* 96, 1669-1673.

22. Beck, D. M., Rees, G., Frith, C. D. & Lavie, N. (2001) *Nat Neurosci* 4, 645-50.
23. Pessoa, L. & Ungerleider, L. G. (2004) *Cereb. Cortex* 14, 511-520.
24. Marois, R., Chun, M. M. & Gore, J. C. (2000) *Neuron* 28, 299-308.
25. Marois, R., Yi, D. J. & Chun, M. M. (2004) *Neuron* 41, 465-472.
26. Pollen, D. A. (1999) *Cereb. Cortex* 9, 4-19.
27. Lattin, R. & Cowey, A. (1971) *Brain Res.* 30, 1-24.
28. Levitt, H. L. (1971) *Journal of the Acoustical Society of America* 49, 467-477.

Legends for Figures

Figure 1 – Visual Discrimination Task With Metacontrast Masking

The stimuli were presented on a black background. The mask overlaps with part of the contour of the target, without leaving gaps nor does it overlap with the target spatially. After the presentation of the target and the mask, the participants were first asked to decide whether a diamond or a square was presented. Then, they had to indicate whether they actually saw the target, or that they simply guessed the answer. Shown in the brackets are the durations of each stimulus.

Figure 2 – Behavioral Results

The upper panel shows results from the behavioural study, which employed the procedure described in figure 1, except that it also included trials where the mask was presented before the target (paracontrast masking). The data were averaged over 8 participants, and the error bars represent one standard error of the mean. Note that at the SOAs where the performance levels (% correct) were the same (e.g. 33 ms and 100 ms), the awareness levels (% seen) differed significantly. The lower panels show the behavioral results from the fMRI study. The data were averaged over 13 participants, and the error bar represents one standard error of the mean. We managed to create two conditions (short SOA and long SOA) that differed significantly in terms of subjective awareness but not in terms of discrimination accuracy.

Figure 3 – Activity in the mid-DLPFC reflects visual consciousness (long SOA > short SOA)

The activity in this area is higher in the long SOA condition than in the short SOA condition, despite the fact that the two conditions did not yield different

discrimination accuracy. There were however more trials where the stimuli were classified by the participants as 'seen' in the long SOA condition than in the short SOA condition. This was the only area that was found to be activated in this comparison; parietal activity did not differ significantly even at liberal thresholds. The plot shows the time courses of the haemodynamic activity in the two conditions, as measured by the percentage signal change from baseline. The data were averaged over 13 participants and the error bar represents one standard error of the mean. Marked by asterisk is the timepoint at which the signal differed significantly as reflected by a paired t-test.

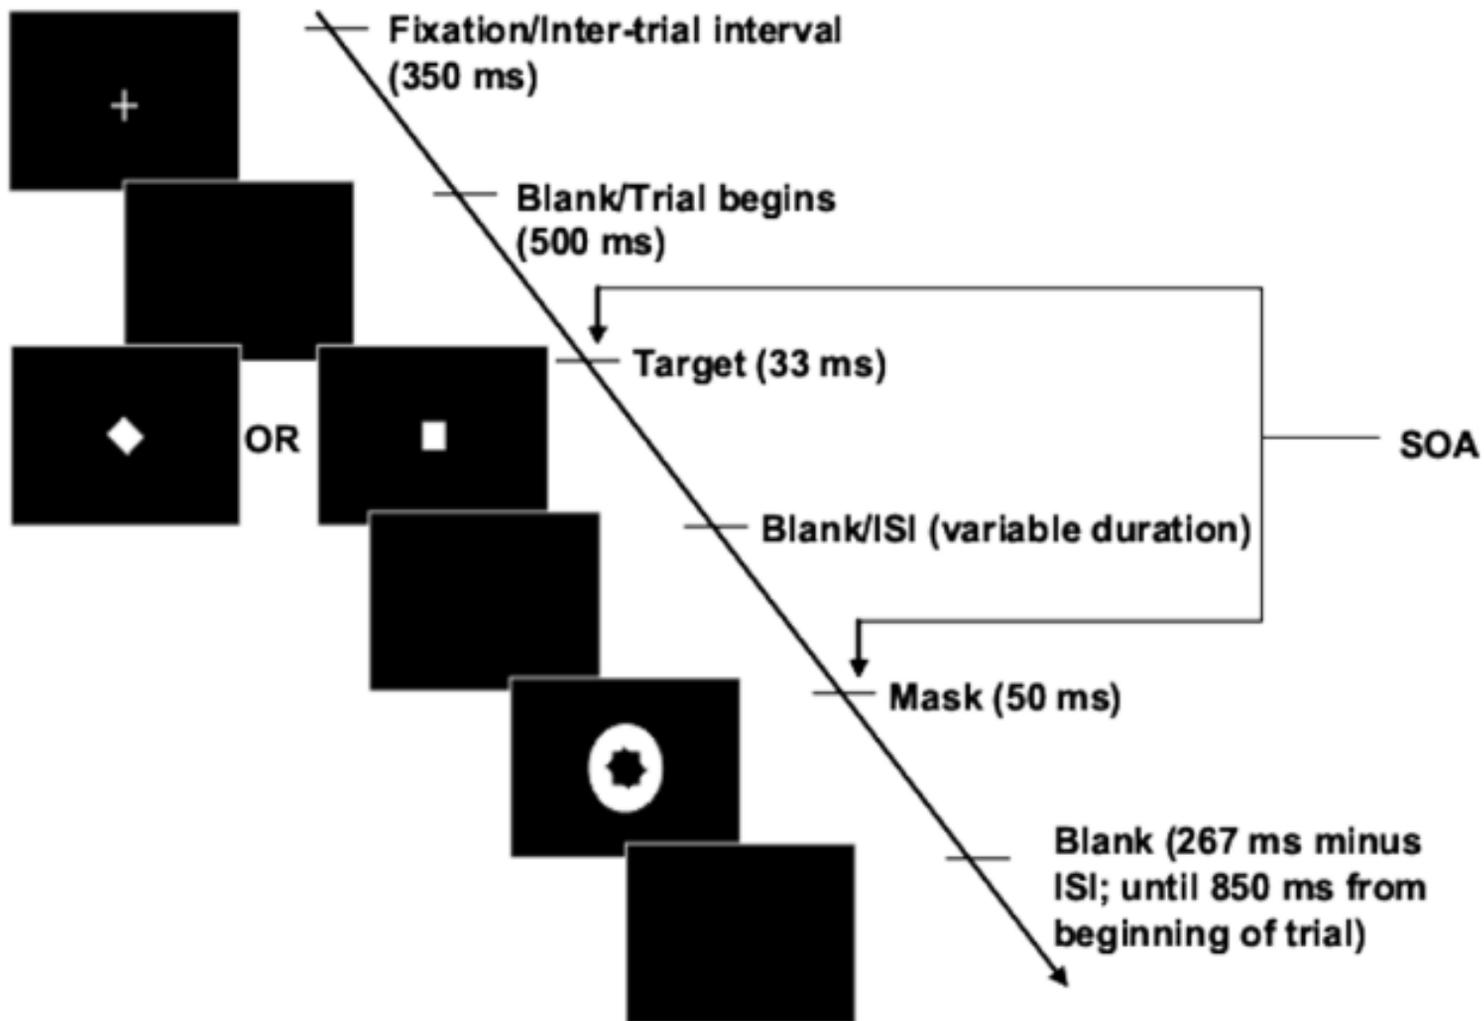
Figure 4 – Activity in the mid-DLPFC across subjects under the same SOA

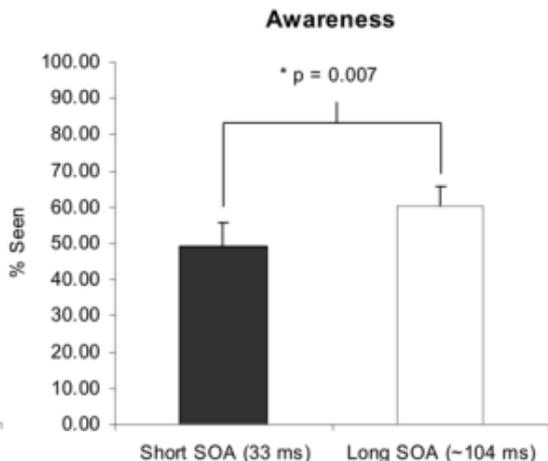
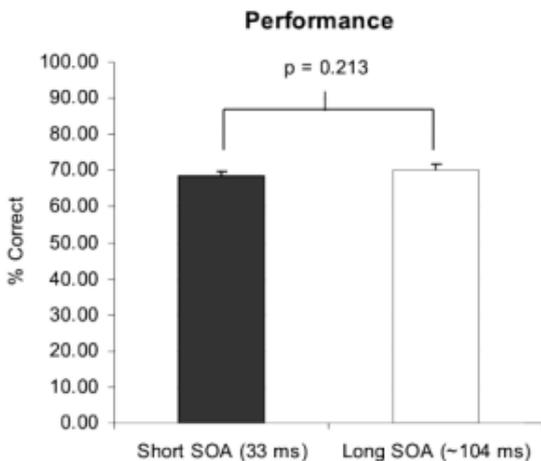
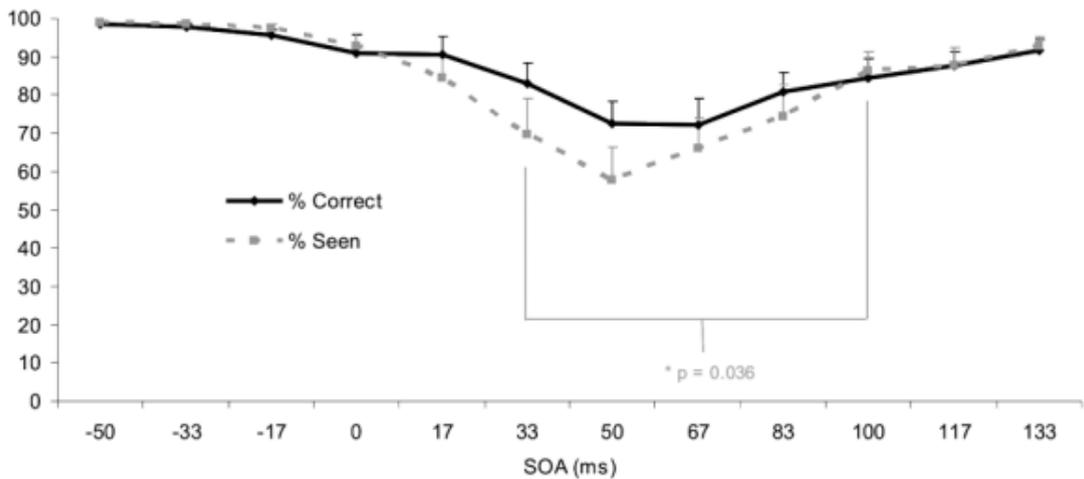
This correlation was performed to test if the activation shown in figure 3 merely reflects a stimulus confound (long vs short SOA). Each point represents a single subject, and the data were taken from the Short SOA condition where the SOA was the same for each subject at 33 ms. As shown in the plot, the higher the activity in the mid-DLPFC, the more likely the identity of the stimulus was declared as consciously seen. This correlation was weak ($r = 0.512$, $p=0.037$ 1-tailed), but it shows that the fact that activity in the mid-DLPFC reflects the subjective criterion for conscious perception cannot be explained simply in terms of SOA.

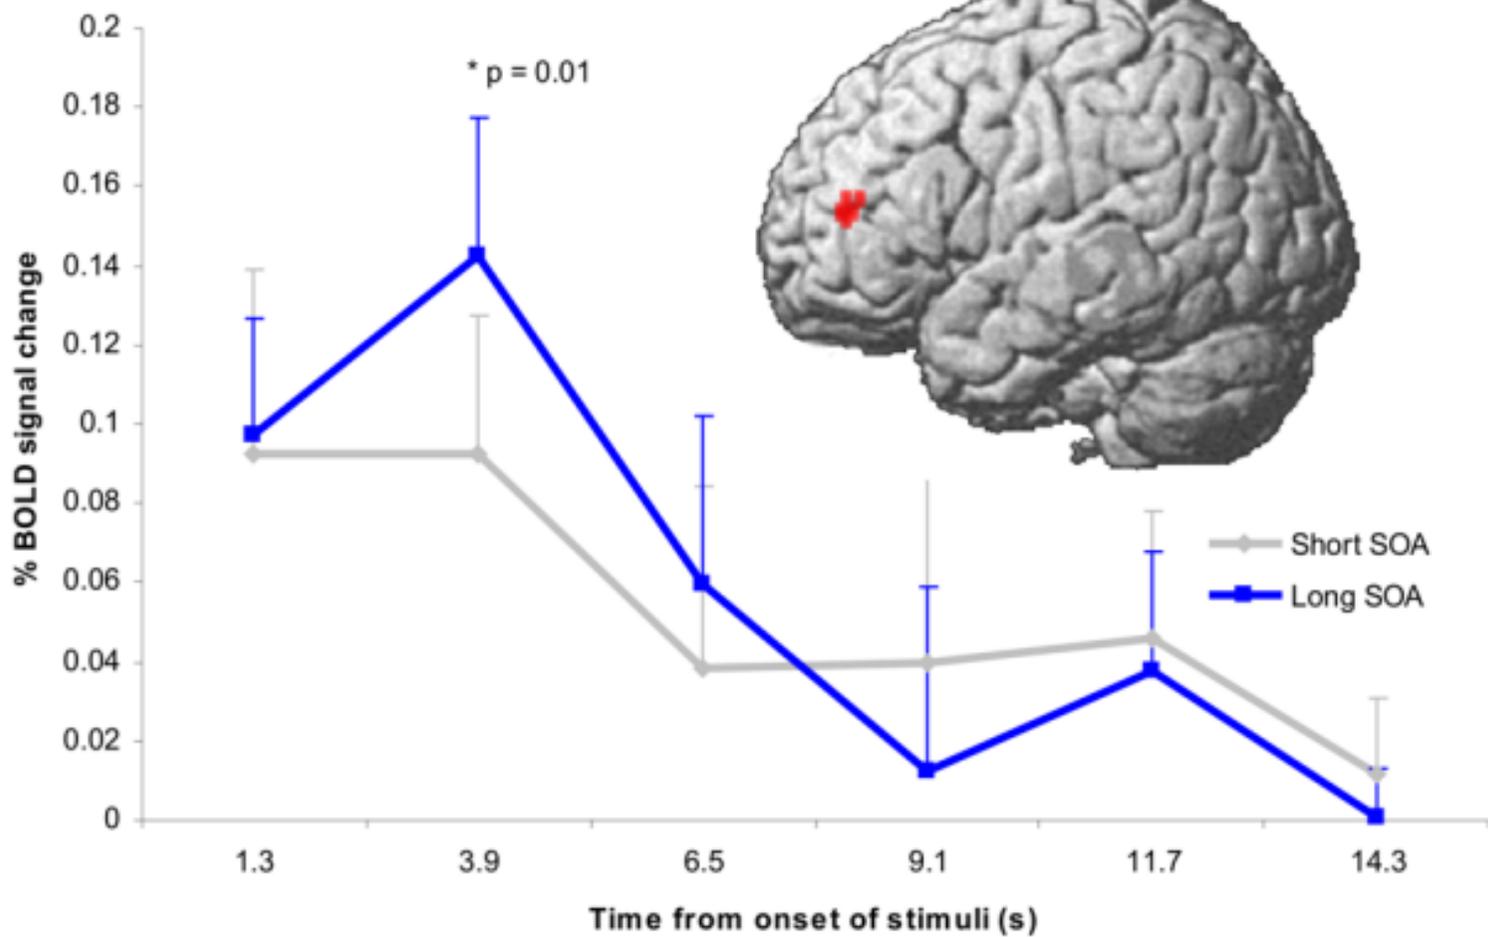
Figure 5 – Activations reflecting performance in general (correct trials > incorrect trials)

VLPFC: ventrolateral prefrontal cortex; IPS: intraparietal sulcus; MTG: Middle temporal gyrus. Activations were found in these areas when correct trials were

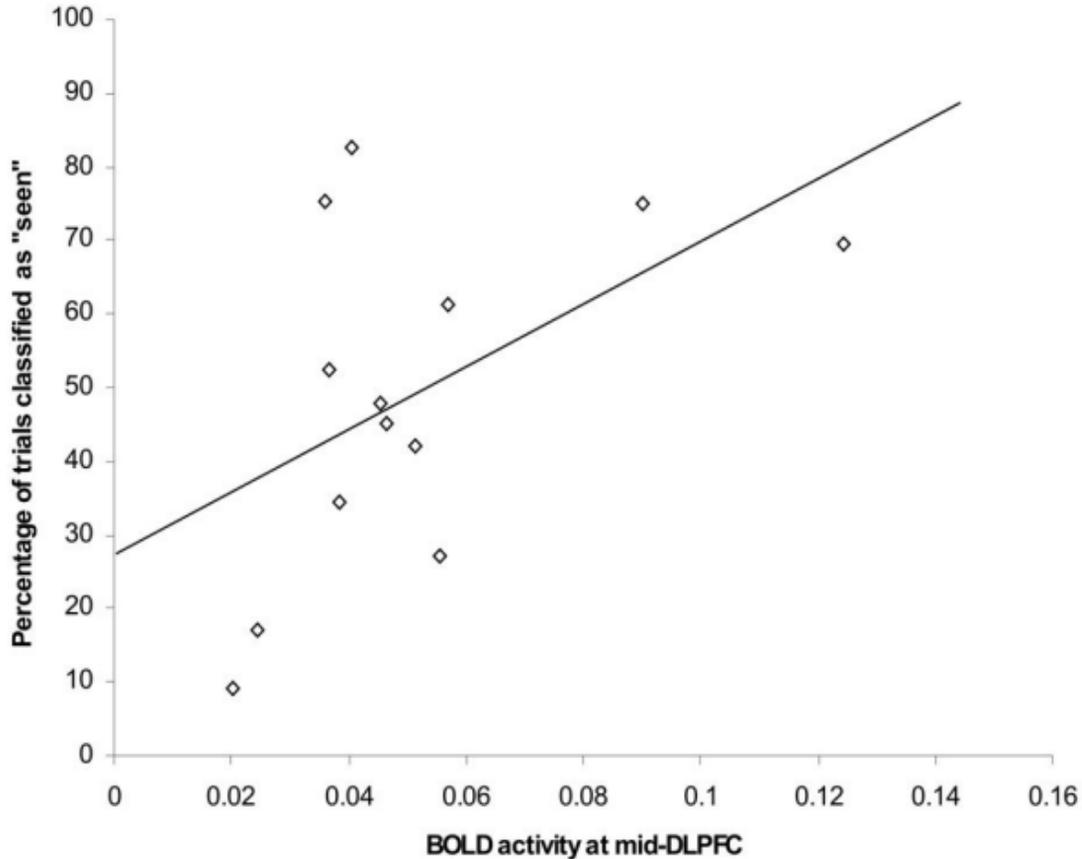
compared against incorrect trials, combining all trials both different SOA conditions. They resemble the pattern of activations in a so-called frontal-parietal network, typically reported in previous studies of visual consciousness.





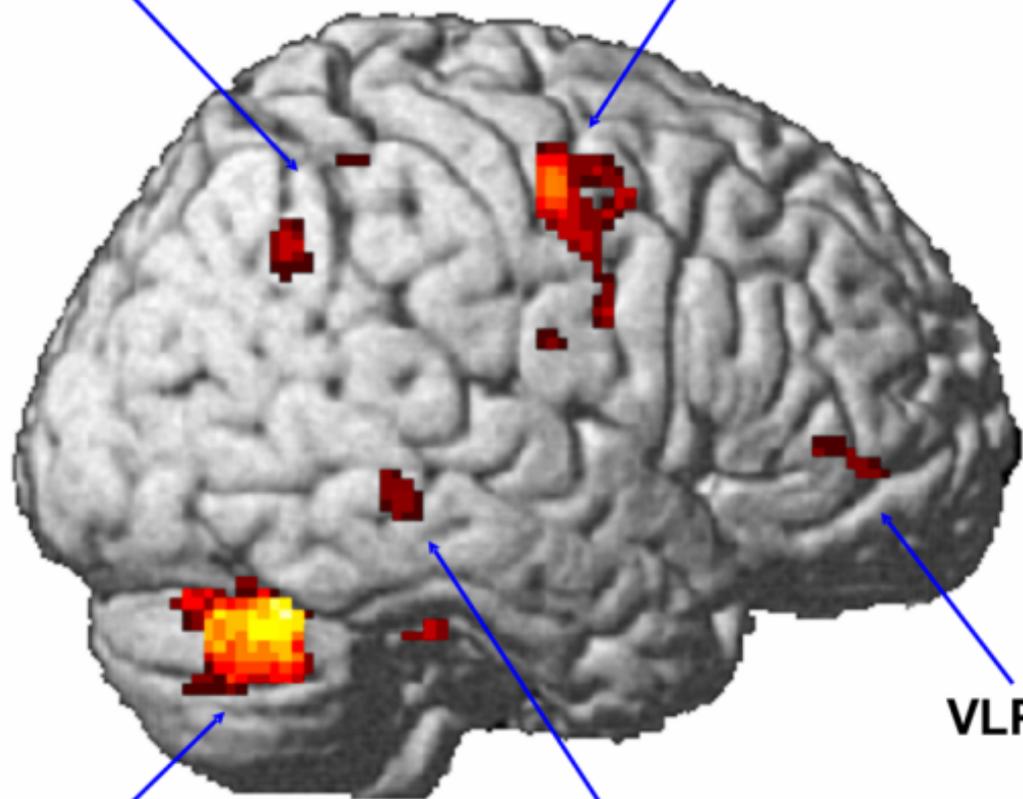


Across-subjects correlation for the Short SOA condition



IPS

Premotor



VLPFC

Cerebellum

MTG