

Early Visual Responses Predict Conscious Face Perception within and between Subjects during Binocular Rivalry

Kristian Sandberg^{1,2}, Bahador Bahrami^{1,2,3}, Ryota Kanai²,
Gareth Robert Barnes⁴, Morten Overgaard¹,
and Geraint Rees^{2,4}

Abstract

■ Previous studies indicate that conscious face perception may be related to neural activity in a large time window around 170–800 msec after stimulus presentation, yet in the majority of these studies changes in conscious experience are confounded with changes in physical stimulation. Using multivariate classification on MEG data recorded when participants reported changes in conscious perception evoked by binocular rivalry between a face and a grating, we showed that only MEG signals in the 120–320 msec time range, peaking at the M170 around 180 msec and the P2m at around 260 msec, reliably predicted conscious experience. Conscious perception could not only be decoded significantly better than chance from the sensors that showed the largest average difference, as previous studies sug-

gest, but also from patterns of activity across groups of occipital sensors that individually were unable to predict perception better than chance. In addition, source space analyses showed that sources in the early and late visual system predicted conscious perception more accurately than frontal and parietal sites, although conscious perception could also be decoded there. Finally, the patterns of neural activity associated with conscious face perception generalized from one participant to another around the times of maximum prediction accuracy. Our work thus demonstrates that the neural correlates of particular conscious contents (here, faces) are highly consistent in time and space within individuals and that these correlates are shared to some extent between individuals. ■

INTRODUCTION

There has been much recent interest in characterizing the neural correlates of conscious face perception, but two critical issues remain unresolved. The first is the time at which it becomes possible to determine conscious face perception from neural signals obtained after a stimulus is presented. The second is whether patterns of activity related to conscious face perception generalize meaningfully across participants, thus allowing comparison of the neural processing related to the conscious experience of particular stimuli between different individuals. Here, we addressed these two questions using MEG to study face perception during binocular rivalry. We also examined several more detailed questions, including which MEG sensors and sources were the most predictive, which frequency bands were predictive, and how to increase prediction accuracy based on preprocessing and preselection of trials.

The neural correlates of conscious face perception have only been studied in the temporal domain in a few recent EEG studies. The most commonly employed strategy in

those studies was to compare neural signals evoked by masked stimuli that differ in stimulus-mask onset asynchrony that results in differences in visibility of the masked stimulus (Harris, Wu, & Woldorff, 2011; Pegna, Darque, Berrut, & Khateb, 2011; Babiloni et al., 2010; Pegna, Landis, & Khateb, 2008; Liddell, Williams, Rathjen, Shevrin, & Gordon, 2004). However, because all but one of these studies (Babiloni et al., 2010) compared brief presentations with long presentations, the stimuli (and corresponding neural signals) differed not only in terms of whether or not they were consciously perceived but also in terms of their duration. Conscious perception of a stimulus was thus confounded by physical stimulus characteristics (Lumer, Friston, & Rees, 1998). Moreover, all of these earlier studies used conventional univariate statistics, comparing, for example, the magnitude of averaged responses between different stimulus conditions across participants. Such approaches are biased toward single strong MEG/EEG sources and may overlook distributed yet equally predictive information.

It remains controversial whether relatively early or late ERP/ERF components predict conscious experience. The relatively early components in question are the N170 found around 170 msec after stimulus onset and a later response

¹Aarhus University Hospital, ²University College London, ³Aarhus University, ⁴Institute of Neurology, London, United Kingdom

at around 260 msec (sometimes called P2 or N2, depending on the analyzed electrodes, and sometimes P300 or P300-like). The N170 is sometimes found to be larger for consciously perceived faces than for those that did not reach awareness (Harris et al., 2011; Pegna et al., 2011; Babiloni et al., 2010), yet this difference is not always found (Pegna et al., 2008; Liddell et al., 2004). Similarly, the P2/N2 correlated positively with conscious experience in one article (Babiloni et al., 2010) and negatively in others (Pegna et al., 2011; Liddell et al., 2004). Additionally, both the N170 (Pegna et al., 2008) and the P2/N2 (Pegna et al., 2011; Liddell et al., 2004) depend on invisible stimulus characteristics, suggesting that these components reflect unconscious processing (but see Harris et al., 2011).

Late components are found between 300 and 800 msec after stimulus presentation. Two studies point to these components (300–800 msec) as reflecting conscious experience of faces (Pegna et al., 2008; Liddell et al., 2004), yet these late components are only present when stimulus durations differ between conscious and unconscious stimuli and not when stimulus duration is kept constant across the entire experiment and stimuli are classified as conscious or unconscious by the participants (Babiloni et al., 2010).

Here, we therefore sought to identify the time range for which neural activity was diagnostic of the contents of conscious experience in a paradigm where conscious experience changed, but physical stimulation remained constant. We used highly sensitive multivariate pattern analysis of MEG signals to examine the time when the conscious experience of the participants viewing intermittent binocular rivalry (Leopold, Wilke, Maier, & Logothetis, 2002; Breese, 1899) could be predicted. During intermittent binocular rivalry, two different stimuli are presented on each trial—one to each eye. Although two different stimuli are presented, the participant typically reports perceiving only one image, and this image varies from trial to trial. In other words, physical stimuli are kept constant, but conscious experience varies from trial to trial. This allowed us to examine whether and when MEG signals predicted conscious experience on a per-participant and trial-by-trial basis. Consistent with previous studies using multivariate decoding, we collected a large data set from a relatively small number of individuals (Raizada & Connolly, 2012; Carlson, Hogendoorn, Kanai, Mesik, & Turret, 2011; Haynes, Deichmann, & Rees, 2005; Haynes & Rees, 2005), employing a case-plus-replication approach supplemented with group analyses where necessary.

Having established the temporal and spatial nature of the neural activity specific to conscious face perception by use of multivariate pattern analysis applied to MEG signals, we further sought to characterize how consistently this pattern generalized between participants. If the pattern of MEG signals in one participant was sufficient to provide markers of conscious perception that could be generalized to other participants, this would provide one way to compare similarities in neural processing related

to the conscious experience of particular stimuli between different individuals.

After having examined our two main questions, two methods for improving multivariate classification accuracy were also examined: stringent low-pass filtering to smooth the data and rejection of trials with unclear perception. Next, univariate and multivariate prediction results were compared with find correlates of conscious face perception that are not revealed by univariate analyses. This analysis was performed at the sensor level as well as on activity reconstructed at various cortical sources. In addition to these analyses, it was examined whether decoding accuracy was improved by taking into account information distributed across the ERF or by using estimates of power in various frequency bands.

METHODS

MEG signals were measured from healthy human participants while they experienced intermittent binocular rivalry. Participants viewed binocular rivalry stimuli (images of a face and a sinusoidal grating) intermittently in a series of short trials (Figure 1A) and reported their percept using a button press. This allowed us to label trials by the reported percept, yet time-lock analyses of the rapidly changing MEG signal to the specific time of stimulus presentation instead of relying on the timing of button press reports, which are both delayed and variable with respect to the timing of changes in conscious contents. The advantages of this procedure have been described elsewhere (Kornmeier & Bach, 2004).

Participants

Eight healthy young adults (six women) between 21 and 34 years (mean = 26.0 years, $SD = 3.55$ years) with normal or corrected-to-normal vision gave written informed consent to participate in the experiment. The experiments were approved by the University College London Research Ethics Committee.

Apparatus and MEG Recording

Stimuli were generated using the MATLAB toolbox Cogent (www.vislab.ucl.ac.uk/cogent.php). They were projected onto a 19-in. screen (resolution = 1024×768 pixels, refresh rate = 60 Hz) using a JVC D-ILA, DLA-SX21 projector. Participants viewed the stimuli through a mirror stereoscope positioned at approximately 50 cm from the screen. MEG data were recorded in a magnetically shielded room with a 275-channel CTF Omega whole-head gradiometer system (VSM MedTech, Coquitlam, BC, Canada) with a 600-Hz sampling rate. After participants were comfortably seated in the MEG, head localizer coils were attached to the nasion and 1 cm anterior (in the direction of the outer canthus) of the left and right tragus to monitor head movement during recording.

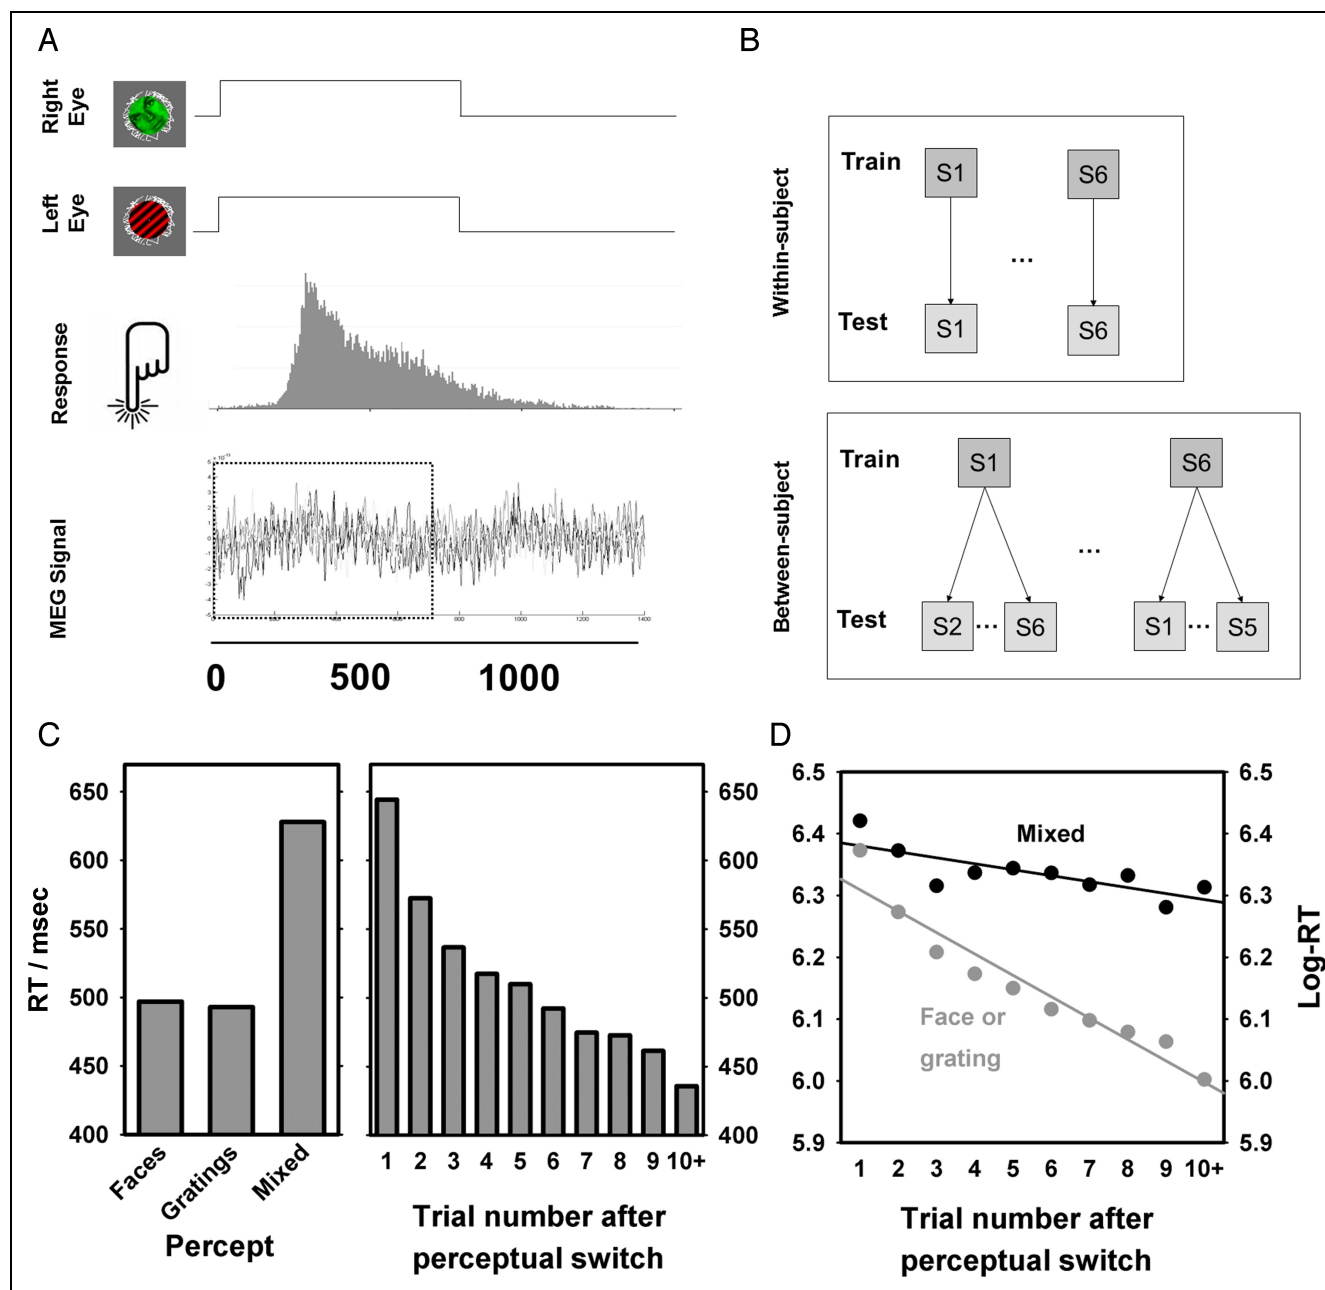


Figure 1. Experimental design and results. (A) Experimental design. Rivaling stimuli (face/grating) were presented for trials lasting ~800 msec separated by blank periods of ~900 msec. Stimuli were dichoptically presented to each eye and rotated in opposite directions at a rate of 0.7 rotations per second. Participants reported which of the two images they perceived with a button press as soon as they saw one image clearly. If perception did not settle, or if the perceived image changed during the trial, the participant reported mixed perception with a third button press. (B) Classification procedure. SVMs were trained to distinguish neuromagnetic activity related to conscious face and grating perception for each participant. The SVMs were then used to decode the perception of (1) the same participant on different trials (top) and (2) each of the other participants (bottom). (C) Left: RT as a function of perceptual report. Right: RT as a function of trial number after a perceptual switch. (D) RT as a function of time after a perceptual switch by perception. The decrease in RT for nonmixed perception indicates that perception on average is clearer far from a perceptual switch than immediately after. Trials for which the same percept has been reported at least 10 times are hereafter referred to as “stable” whereas other trials are referred to as “unstable.”

Stimuli

A red Gabor patch (contrast = 100%, spatial frequency = 3 cycles/degree, standard deviation of the Gaussian envelope = 10 pixels) was presented to the right eye of the participants, and a green face was presented to the

left eye (Figure 1A). To avoid piecemeal rivalry where each image dominates different parts of the visual field for the majority of the trial, the stimuli rotated at a rate of 0.7 rotations/sec in opposite directions, and to ensure that stimuli were perceived in overlapping areas of the visual field, each stimulus was presented within an annulus

(inner/outer $r = 1.3/1.6$ degrees of visual angle) consisting of randomly oriented lines. In the center of the circle was a small circular fixation dot.

Procedure

During both calibration and experiment, participants reported their perception using three buttons each corresponding to either face, grating, or mixed perception. Participants swapped the hand used to report between blocks. This was done to prevent the classification algorithm from associating a perceptual state with neural activity related to a specific motor response. To minimize perceptual bias (Carter & Cavanagh, 2007), the relative luminance of the images was adjusted for each participant until each image was reported equally often ($\pm 5\%$) during a 1-min-long continuous presentation.

Each participant completed six to nine runs of 12 blocks of 20 trials, that is, 1440–2160 trials were completed per participant. On each trial, the stimuli were displayed for approximately 800 msec. Each trial was separated by a uniform gray screen appearing for around 900 msec. Between blocks, participants were given a short break of 8 sec. After each run, participants signaled when they were ready to continue.

Preprocessing

Using SPM8 (www.fil.ion.ucl.ac.uk/spm/), data were down-sampled to 300 Hz and high-pass filtered at 1 Hz. Behavioral reports of perceptual state were used to divide stimulation intervals into face, grating or mixed epochs starting 600 msec before stimulus onset and ending 1400 msec after. Trials were baseline-corrected based on the average of the 600 msec prestimulus activity. Artifacts were rejected at a threshold of 3 pT. On average 0.24% ($SD = 0.09$) of the trials were excluded for each participant because of artifacts.

ERF Analysis

Traditional, univariate ERF analysis was first performed. For this analysis, data were filtered at 20 Hz using a fifth-order Butterworth low-pass filter, and face and grating perception trials were averaged individually using SPM8.

Source Analysis

Sources were examined using the multiple sparse priors (MSP; Friston et al., 2008) algorithm. MSP operates by finding the minimum number of patches on a canonical cortical mesh that explain the largest amount of variance in the MEG data, this tradeoff between complexity and accuracy is optimized through maximization of model evidence. The MSP algorithm was first used to identify

the electrical activity underlying the grand-averaged face/grating contrast maps at a short time window around the M170 and the P2m (100–400 msec after stimulus onset). Afterwards, the MSP algorithm was used to make a group-level source estimation based on template structural MR scans using all trials (over all conditions) from all eight participants. The inverse solution restricts the sources to be the same in all participants but allows for different activation levels. This analysis identified 33 sources activated at stimulus onset (see Table 1). Activity was extracted on a single trial basis across the 33 sources for each scan of each participant and thus allowed for analyses to be performed in source space.

Multivariate Prediction Analysis

Multivariate pattern classification of the evoked responses was performed using the linear support vector machine (SVM) of the MATLAB Bioinformatics Toolbox (Mathworks). The SVM decoded the trial type (face or grating) independently for each time point along the epoch. Classification was based on field strength data as well as power estimates in separate analyses.

Conscious perception was decoded within and between participants. For within-subject training/testing, 10-fold cross-validation was used (Figure 1B). For between-subject training/testing, the SVM was trained on all trials from a single participant and tested on all trials of each of the remaining participants. The process was repeated until data from all participants had been used to train the SVM (Figure 1B).

To decrease classifier training time (for practical reasons), the SVM used only 100 randomly selected trials of each kind (200 in total). As classification accuracy cannot be compared between classifiers trained on different numbers of trials, participants were excluded from analyses if they did not report 100 of each kind of analyzed trials. The number of participants included in each analysis is reported in the Results section.

In addition to the evoked response analysis, a moving window discrete Fourier transform was used to make a continuous estimate of signal power in selected frequency bands over time: theta = 3–8 Hz, alpha = 9–13 Hz, low beta = 14–20 Hz, high beta = 21–30 Hz, six gamma bands in the range of 31–90 Hz, each consisting of 10 Hz (Gamma 1, for instance, would thus be 31–40 Hz) but excluding the 50-Hz band. The duration of the moving window was set to accommodate at least three cycles of the lowest frequency within each band (e.g., for theta [3–8 Hz], the window was 900 msec).

Statistical Testing

All statistical tests were two-tailed. Comparisons of classification accuracies were performed on a within-subject

Table 1. Sources

Source	Area	Name	<i>x</i>	<i>y</i>	<i>z</i>
1	Occipital lobe	IV1	-2	-96	5
2		rV1	12	-98	-1
3		lvOCC1	-16	-94	-18
4		rvOCC1	21	-96	-17
5		lvOCC2	-14	-80	-13
6		rvOCC2	15	-80	-12
7		ldOCC	-18	-81	40
8		rdOCC	19	-82	40
9	OFA	lOFA	-38	-80	-15
10		rOFA	39	-80	-15
11	Face-specific	lpSTS1	-54	-63	9
12		rpSTS1	53	-63	13
13		lpSTS2	-55	-50	23
14		rpSTS2	54	-49	18
15		lpSTS3	-59	-33	10
16		rpSTS3	55	-34	7
17		IFFA	-53	-51	-22
18		rFFA	52	-52	-22
19	Parietal	lSPL1	-40	-37	60
20		rSPL1	36	-37	60
21		lSPL2	-33	-65	49
22		rSPL2	36	-64	46
23		lSPL3	-41	-35	44
24		rSPL3	39	-36	44
25	Motor	IPC	-54	-12	15
26		rPC	54	-11	13
27	Frontal	laMFG1	-40	18	27
28		raMFG1	38	18	26
29		laMFG2	38	41	19
30		lOFC1	-24	7	-18
31		rOFC1	22	8	-19
32		lOFC2	-43	31	-16
33		rOFC2	41	35	-15

The 33 sources judged to be most active across all trials independently of perception/stabilization across all participants. Sources were localized using MSPs to solve the inverse problem. Source abbreviations: V1 = striate cortex; OCC = occipital lobe; IT = inferior temporal cortex; SPL = superior parietal lobule; PC = precentral cortex; MFG = middle frontal gyrus. Navigational abbreviations: l = left hemisphere; r = right hemisphere; p = posterior; a = anterior; d = dorsal; v = ventral.

basis using the binomial distributions of correct/incorrect classifications. To show the reproducibility of the within-subject significant effects across individuals, we used the cumulative binomial distribution,

$$\Pr(X \leq x) = \sum_{i=0}^x \binom{n}{i} p^i (1-p)^{n-i} \quad (1)$$

where *n* is the total number of participants, the within-subject significant criterion is *p* (= .05), *x* is the number of participants that reach this criterion, and $\binom{n}{i}$ is the binomial coefficient.

Prediction accuracy for each power envelope was averaged across a 700-msec time window after stimulus presentation (211 sampling points) for each participant. Histogram inspection and Shapiro–Wilk tests showed that the resulting accuracies were normally distributed. One-sample *t* tests (*n* = 8) were used to compare the prediction accuracy level of each power band to chance (0.5). Bonferroni correction for 10 comparisons was used as 10 power bands were analyzed.

RESULTS

EEG research points to the N170 and the component sometimes called the P2 as prime candidates for the correlates of conscious face perception (following convention, we shall call these M170 and P2m hereafter) but later sustained activity around 300–800 msec may also be relevant. To search for predictive activity even earlier than this, activity around the face-specific M100 was also examined. Before analyses, trials with unclear perception were identified and excluded from subsequent analyses.

Identification of Unclear Perception Based on Behavioral Data

Analyses were optimized by contrasting only face/grating trials on which perception was as clear as possible. Participants generally reported perception to be unclear in two ways, both of which have been observed previously (see Blake, 2001). First, participants reported piecemeal rivalry where both images were mixed in different parts of the visual field for the majority of the trial. Such trials were not used in the MEG analyses. Second, participants sometimes experienced brief periods (<200 msec) of fused or mixed perception at the onset of rivalry. Participants were not instructed to report this initial unclear perception if a stable image was perceived after a few hundred milliseconds to keep the task simple. To minimize the impact of this type of unclear perception on analyses, we exploited the phenomenon of stabilization that occurs during intermittent rivalry presentations, which will be explained below.

On average, participants reported face perception on 45.5% ($SD = 15.1$) of the trials, grating perception on 42.6% ($SD = 16.1$), and mixed perception on 11.9% ($SD = 10.6$). Mean RT across participants ($n = 8$) was 516 msec ($SD = 113$) overall, and the frequency histogram of the data in Figure 1A shows the variance in RT. Average RT was 497 msec ($SD = 112$) for face perception, 493 msec ($SD = 134$) for grating perception, and 628 msec ($SD = 117$) for mixed perception, reflecting a longer decision-making time when perception was unclear (Figure 1C).

During continuous rivalry, the neural population representing the dominant image strongly inhibits the competing neural population, but as adaptation occurs, inhibition gradually decreases until perception switches after a few seconds (Noest, Van Ee, Nijs, & Van Wezel, 2007; Wilson, 2003, 2007; Freeman, 2005). In contrast, during intermittent presentation, adaptation does not easily reach the levels at which inhibition decreases significantly while at the same time the percept-related signal stays high possibly because of increased excitability of the dominant neurons (Wilson, 2007) or increased subthreshold elevation of baseline activity of the dominant neurons (Noest et al., 2007). Behaviorally, this results in a high degree of stabilization, that is, the same image being perceived on many consecutive trials, and a swift inhibition of the nondominant image is thus to be expected on such stabilized trials. This should result in minimization of the brief period of fused or mixed perception, causing a faster report of the perceived image. We hypothesized that stabilization-related perceptual clarity builds up gradually across trials following a perceptual switch and tested this by examining RTs. If the hypothesis is correct, a negative correlation between RT and trial number counted from a perceptual switch would be expected for face/grating, but not for mixed perception. In other words, when stabilization increases across time, perceptual clarity is expected to increase and RT to decrease. When perception remains mixed, no such effect is expected, although participants press the same response button on consecutive trials.

As can be seen in Figure 1D, log-transformed RT did indeed correlate negatively with time after a perceptual switch for face/grating perception ($r = -0.39$, $p < .001$), but not for mixed perception ($r = -0.11$, $p = .37$). This gradual build-up of stabilization-related perceptual clarity was confirmed in additional MEG analyses to be reported elsewhere (Sandberg et al., submitted). On the basis of both these findings, we analyzed only MEG trials for which participants had reported at least 10 identical percepts. We refer to these as “stable trials.” A similar criterion was used by Brascamp et al. (2008). After artifact rejection and rejection of unstable trials, on average 396 face perception and 393 grating perception trials remained per participant.

The impact of rejection of unstable trials on decoding accuracy is reported in the Appendix: Improving Decoding

Accuracy section. Please note that results remain highly significant without rejection of these trials.

Univariate ERF and Source Differences

We first examined which ERF components varied with conscious perception. We calculated a face/grating contrast using stable trials, and as shown in Figure 2A, activity related to face perception differed clearly from that related to grating perception particularly at two time points, 187 msec (M170) and 267 msec (P2m), after stimulus presentation. The three face-specific peaks, the M100, M170, and P2m are shown in Figure 2B, C. Figure 2D shows that the difference at 187 msec was localized almost exclusively to temporal sensors.

The electrical activity underlying the grand-averaged face/grating contrast maps was estimated using the MSP algorithm, and the solution explained 97% of the variance in the MEG signals for the period from 100 to 400 msec after stimulus onset. The posterior probability map, showing those cortical locations with 95% probability of having nonzero current density at $t = 180$ msec (the time of maximal activity difference) is plotted in Figure 2E. The activity pattern was strikingly consistent with activation of the face-processing network (Haxby, Hoffman, & Gobbini, 2000) with the right occipital face area (OFA) indicated as the largest source.

Within-subject Decoding of Conscious Perception

To determine the times when MEG activity accurately predicted conscious experience, multivariate SVM classifiers were trained to decode perception on each trial. To demonstrate that results remained significant without any preselection of trials, classifiers were first trained on 1–20 Hz filtered data from 100 randomly selected trials of each kind (face/grating), thus including both stable and unstable trials.

Conscious perception was predicted at a level significantly above chance in the 120–300 msec time window with average classification performance peaking at around 180 and 260 msec after stimulus onset (Figure 3A, C–J) (the third, smaller peak at around 340 msec was not observed for all participants and was not replicated in the between-subject analyses). Activity after 350 msec only predicted conscious experience to a very small degree or not at all. The temporal positions of the two peaks in classification performance corresponded well with the M170 and the P2m. On the basis of the binomial distribution of correct/incorrect classifications, classification accuracy was above chance at the $p < .05$ level at 187 msec for all eight participants and at 270 msec for seven of eight participants. The probability of finding significantly above chance within-subject prediction accuracies for seven or eight of the total eight participants in this case-study-plus-replication design by chance was $p = 6.0 \times 10^{-9}$ and $p = 3.9 \times 10^{-11}$, respectively (uncorrected for comparisons over latencies).

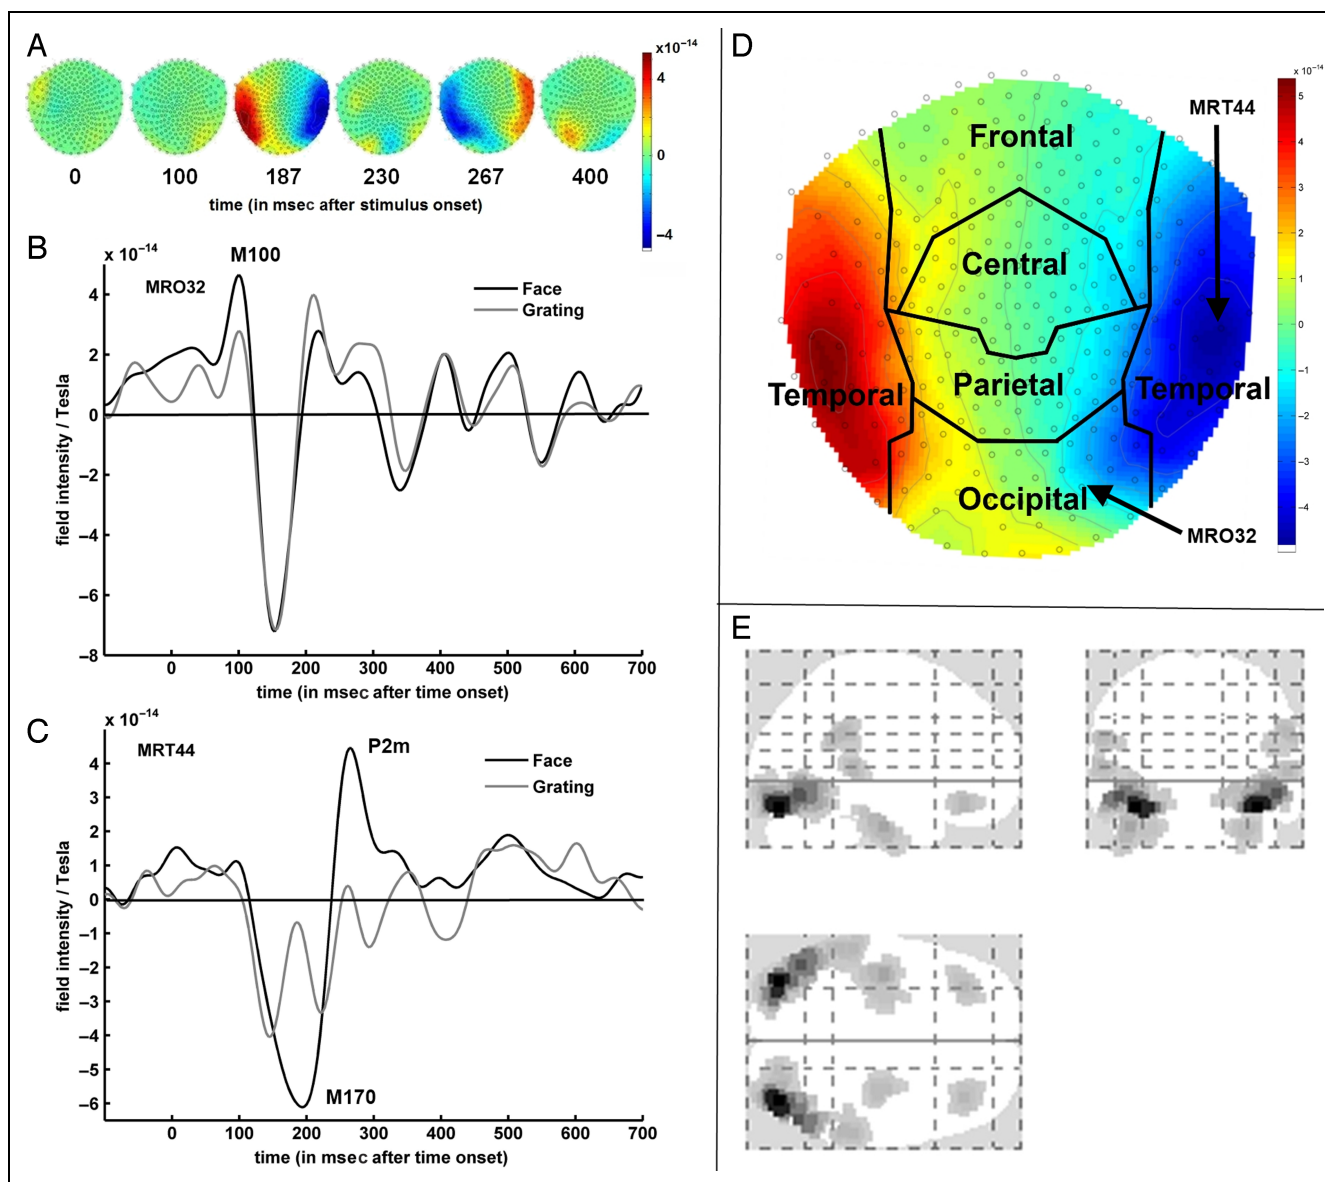


Figure 2. Univariate analyses on averaged field strength data (stable trials). (A) Topographic maps showing face/grating contrast. The largest differences were found at 187 and 267 msec after stimulus onset. (B) Activity at the sensor for which the largest M100 difference was found (MRO32). Generally, only small differences were observed. (C) Activity at the sensor for which the largest M170 and P2m difference was found (MRT44). Notice that face-related activity is larger than grating-related at both peaks. (D) Map of sensor location. (E) Posterior probability map of estimated cortical activity underlying the average difference between face and grating perception in the 100–400 msec time window using the MSP algorithm. The gray–black scale shows the regions of the cortical surface with greater than 95% chance of being active. The solution explains 97% of the measured data. The image is plotted at $t = 180$ msec, the peak latency at the peak source location (38, –81, –17). The activity pattern was consistent with activation of the face processing network (Haxby et al., 2000).

At no time point around the M100 were significant within-subject differences found for more than two participants, giving a combined $p = .057$, thus indicating that little or no group differences between face and grating perception were present at the M100. Overall, the main predictors of conscious perception thus appeared to be the M170 (at 187 msec) and to a slightly lesser extent the P2m (at 270 msec).

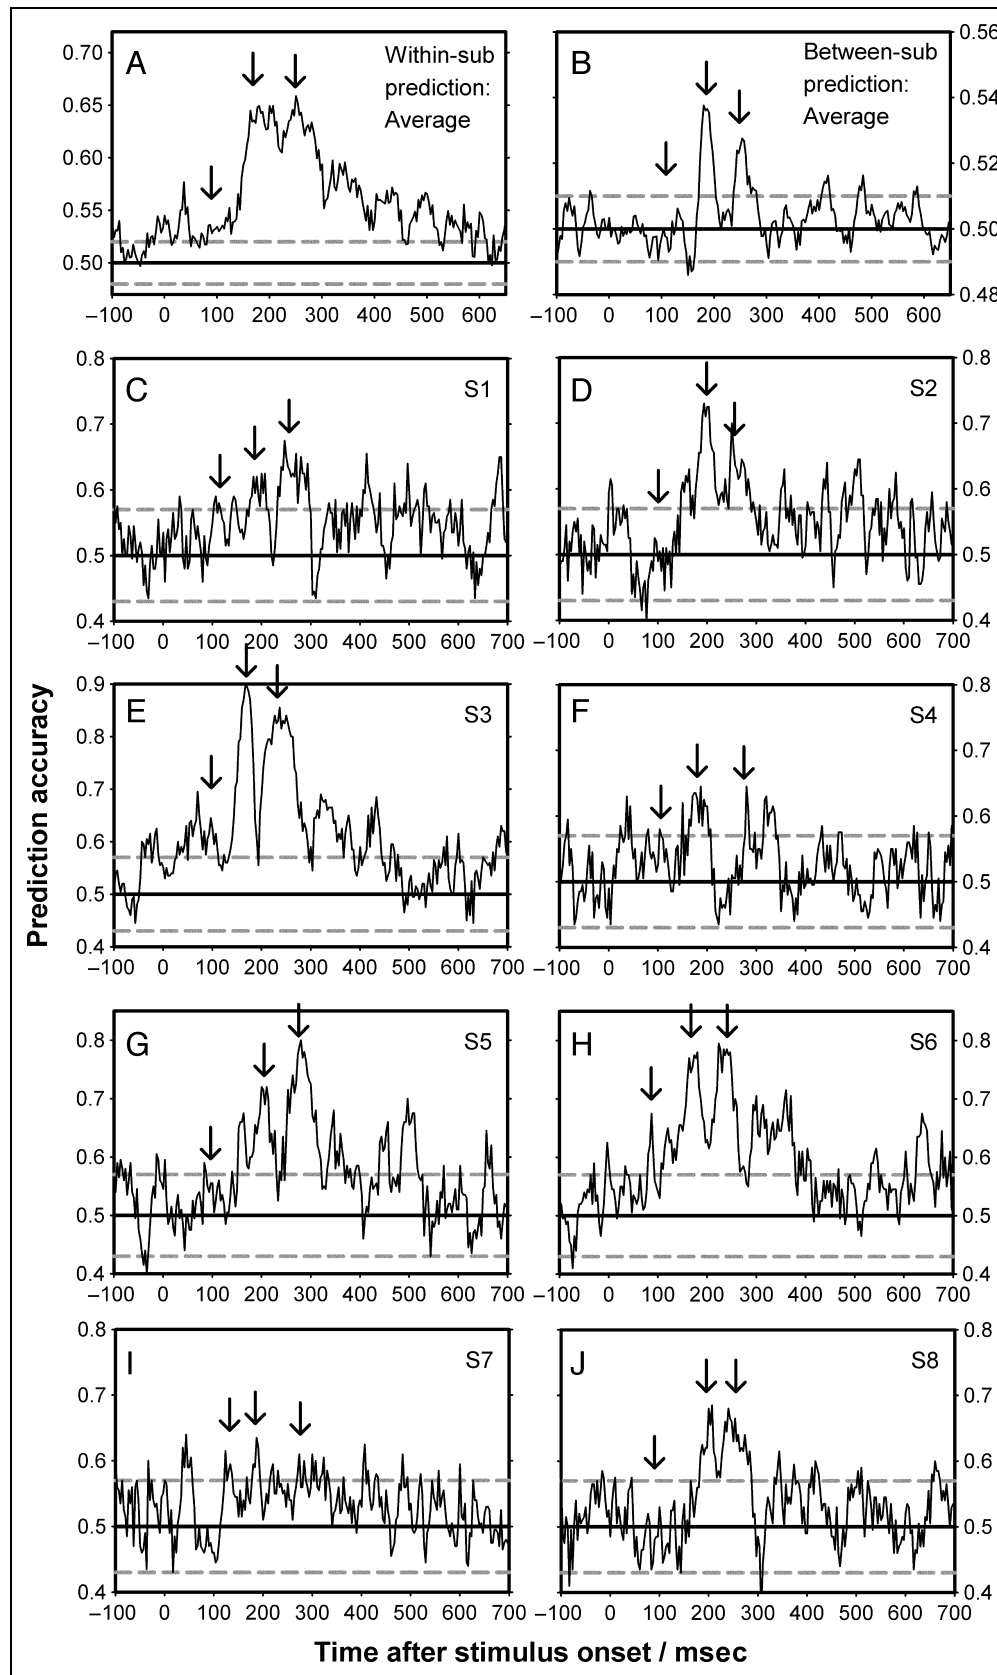
Having determined that conscious experience could be predicted within participants in the 120–300 msec time range, SVM classifiers were trained on data from one par-

ticipant to decode the conscious content of a different participant (Figure 1B, bottom).

Between-subject Decoding of Conscious Perception

For between-subject decoding, peaks were observed around the M170 and the P2m, but no above-chance accuracy was observed around the M100 (Figure 3B). Accuracy was significantly above chance for seven of eight participants at 180 msec and for five of eight participants

Figure 3. Prediction accuracy across time using all trials. Average prediction accuracy for all trials (stable and unstable) across participants is plotted based on the single trial, 1–20 Hz filtered MEG field strength data as a function of time. An SVM was trained to predict reported perception (face vs. grating) for each time point. The dotted gray line indicates the threshold for which a binomial distribution of the same number as the total number of trials the prediction is performed upon is different from chance (uncorrected). (A) Average within-subject prediction accuracy for all eight participants is plotted (i.e., classification accuracy when the SVM was trained and tested on data from the same participant). Notice the two clear peaks (the M170 at 187 msec and the P2m at 267 msec) indicated by the second and third arrows. The first arrow indicates the expected timing of the M100. (B) Average between-subject prediction accuracy for all between-subject tests across time (i.e., classification accuracy when the SVM was trained and tested on data from different participants). (C–J) Prediction accuracy for each individual participant for the within-subject predictions.



at 250 msec. The probability of observing these within-participant repeated replications were $p = 6.0 \times 10^{-9}$ and $p = 1.5 \times 10^{-5}$, respectively. No significant differences were found around the M100.

Overall, the M170 was thus found to be the component that predicted conscious experience most accurately and significantly both within and between individuals, closely followed by the P2m. Before initiating further analyses, we examined how different analysis parameters might change decoding accuracy as described below.

We hypothesized that decoding accuracy could be increased in two ways: by rejecting trials for which perception was not completely clear and by applying a more stringent filter to the data. Participant's reports (see Results) suggested that the probability of clear perception on a given trial increased the further away the trial is from a perceptual switch. We thus tested classifiers trained on stable versus unstable trials and on 1–300 Hz, 1–20 Hz, and 2–10 Hz filtered data. This analysis is reported in the Appendix: Improving decoding accuracy and showed that the best results were obtained using 2–10 Hz filtered data from stable trials. Please note that this should not be taken as an indication that higher frequencies are considered noise in a physiological sense, simply that the ERF components in the present experiment may be viewed as half cycles of around 3–9 Hz and that the temporal smoothing

of a 10-Hz low-pass filter may have minimized individual differences in latency of the M170 and P2m.

Moreover, in the Appendix, we also report an analysis of the predictive ability of power in various frequency bands (Appendix: Decoding using power estimations). This analysis shows that the low frequencies dominating the ERF components are the most predictive, yet prediction accuracy was never better than for analyses based on the evoked field strength response. The following analyses are thus performed on 2–10 Hz filtered data from the six participants who reported at least 100 trials of stable face/grating perception.

Identification of Predictive Sensors

One advantage of multivariate decoding over univariate analyses is the sensitivity to distributed patterns of information. We therefore examined which group of sensors was most predictive of conscious face perception independently of whether these sensors showed the largest grand average difference.

Identification of predictive sensors was based on the standard CTF labeling of sensors according to scalp areas as seen in Figure 2D. First, the number of randomly selected sensors distributed across the scalp required to decode perception accurately around the most predictive

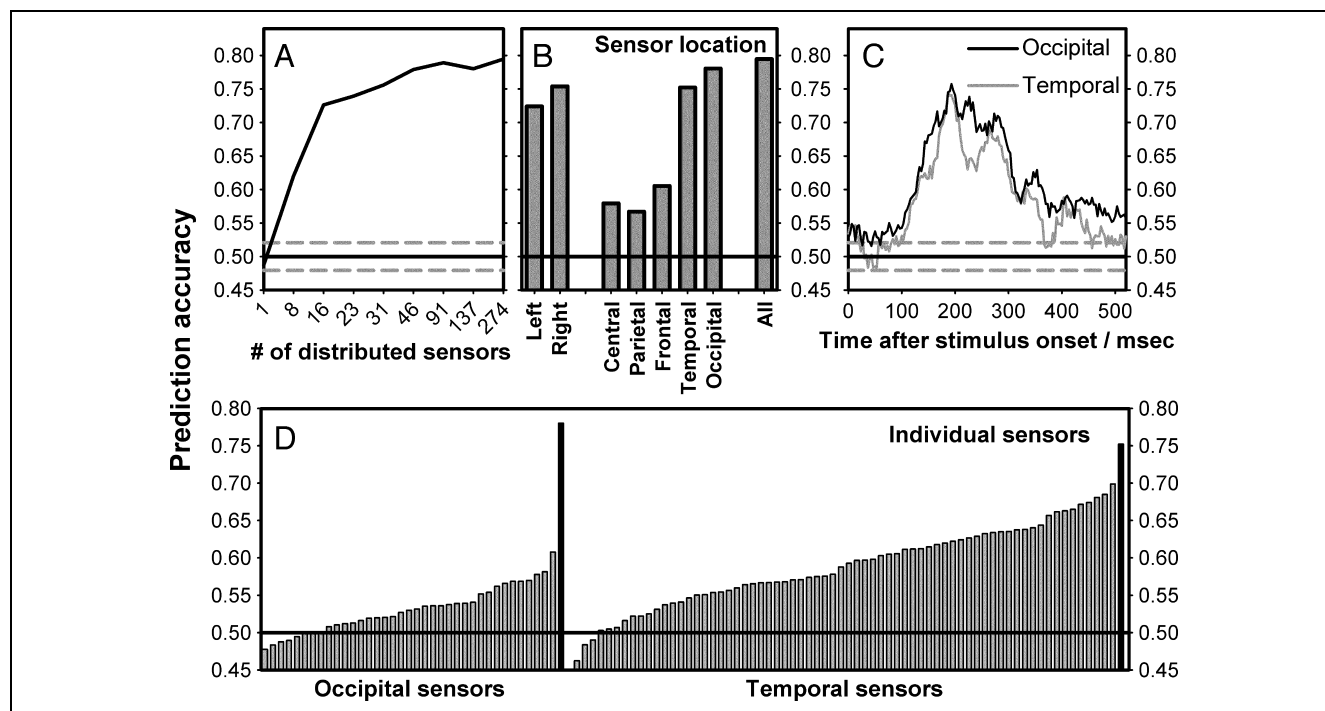


Figure 4. Predictability by sensor location (stable trials). Six participants had enough trials to train the classifiers on stable trials alone. The figure plots prediction accuracy based on 2–10 Hz filtered data from these participants. Dotted gray line represents the 95% binomial confidence interval around chance (uncorrected). (A) Prediction accuracy as a function of the number of randomly selected sensors from all scalp locations. (B) Group-level prediction accuracy as a function of sensor location. Left/right indicate that classifier is trained on left/right hemisphere sensors, respectively. Other sensor locations can be seen in Figure 2D. (C) Average prediction accuracy for within-subject tests across time when classifier is trained/tested using occipital and temporal sensors, respectively. (D) Prediction accuracy at the time of the M170 when the classifier is trained on single sensors (i.e., univariate classification) or all sensors (multivariate classification) in occipital/temporal locations. Each gray bar plots accuracy for a single sensor. Black bars plot group-level performance.

sensors, the group level classification (black bar) is much greater than that of the single best sensor, whereas this is not the case for temporal sensors. In fact, a prediction accuracy of 74.3% could be achieved using only 10 occipital sensors with individual chance-level performance (maximum of 51.3%).

Just as multivariate classification predicted conscious face perception at sensors that were at chance individually, it is possible that perception might be decoded using multiple time points for which individual classification accuracy was at chance. It may also be possible that the information at the P2m was partially independent from the information at the M170, causing joint classification accuracy to increase beyond individual classification. For these reasons, we examined classification accuracy when the SVM classifiers were trained on data from multiple time points. The formal analysis is reported in Appendix: Decoding using multiple time points and shows that including a wide range of time points around each peak (11 time points, 37 msec of data) does not improve decoding accuracy. Neither does inclusion of information at both time points in a single classifier, and finally, decoding of consciousness perception is not improved above chance using multiple time points individually at chance.

Decoding in Source Space

Our finding that signals from single time points at the sensors close to visual areas of the brain were the most predictive does not necessarily mean, however, that the activity at these sensors originates from visual areas. To test this, analyses of sources are necessary. Therefore, activity was reconstructed at the 33 sources that were most clearly activated by the stimuli in general (i.e., independently of conscious perception), and decoding was performed on these data. The analysis was performed on 2–10 Hz filtered data from stable trials using the six participants who had 100 or more stable trials with reported face/grating perception.

First, decoding accuracy was examined across time when classifiers were trained/tested on data from all sources (Figure 5A). Next, classifiers were trained on groups of sources based on cortical location (see Table 1). Comparisons between the accuracies achieved by each group of sources may only be made cautiously as the number of activated sources differs between areas, and the classifiers were thus based on slightly different numbers of features. The occipital, the face-specific, the frontal, and the parietal groups, however, included almost the same number of sources (8, 8, 7, and 6, respectively). Overall, Figure 5 (A, B) shows that for all sources, decoding accuracy peaked around the M170 and/or the P2m and that conscious perception could be predicted almost as accurately from eight occipital or face-specific sources as from all 33 sources combined. This was not found for any other area.

Decoding accuracy was also calculated for the individual sources at the M170 (Figure 5C) and the P2m (Figure 5D) using the individual peaks of each participant (see Figure 3). The single most predictive source with an accuracy of 64% at the M170 and 59% at the P2m was the right OFA—a face-sensitive area in the occipital lobe. The majority of the remaining predictive sources were found in occipital and face-specific areas with the exception of a ventral medial prefrontal area and possibly an area in the superior parietal lobe around the P2m. The peak classification accuracies for groups of sources (black bars in Figure 5C, D) were also the highest for occipital and face-specific sources, yet when combined the sources in other areas also became predictive above chance. Overall, it appeared that the most predictive sources were in the visual cortex, although information in other areas also predicted conscious perception. Generally, little or no difference was observed regarding which sources were predictive at the M170 and at the P2m.

DISCUSSION

Two unresolved major questions were presented in the Introduction. The first was the question of which temporal aspects of the MEG signal are predictive of conscious face perception.

M170 and P2m Predict Conscious Face Perception

Multivariate classification on binocular rivalry data demonstrated that activity around the face-specific M170 and P2m components differed on a single trial basis, depending on whether a face was perceived consciously or not. Perception was predicted significantly better than chance from temporal sensors showing large average activity differences, and around these sensors group-level decoding accuracy was dependent on the single best sensor used. Additionally, perception could be decoded as well or better when using occipital sensors that showed little or no mean activity differences between conscious perception of a face or not. At these locations, perception was predicted as accurately when using sensors that were individually at chance as when using all temporal sensors, thus showing a difference that was not revealed by univariate analyses. No predictive components were found after 300 msec, thus arguing against activity at these times predicting conscious experience.

Interestingly, the event-related signal related to conscious face perception found in the masking study using identical durations for “seen” and “unseen” trials (Babiloni et al., 2010) appeared more similar to that found in the present experiment than to those found in other EEG masking experiments. This indicates that when physical stimulation is controlled for, very similar correlates of conscious face perception are found across paradigms. In neither experiment were differences found between

late components (in fact, no clear late components are found).

MEG/EEG Sensor and Source Correlates of Visual Consciousness

Our findings appear to generalize to not only to conscious face perception across paradigms but also to visual awareness more generally. For example, Koivisto and Revonsuo (2010) reviewed around 40 EEG studies using different experimental paradigms and found that visual awareness correlated with posterior amplitude shifts around 130–320 msec, also known as visual awareness negativity, whereas later components did not correlate directly with awareness. Furthermore, they argued that the earliest and most consistent ERP correlate of visual awareness is an amplitude shift around 200 msec, corresponding well with the findings of this study.

Nevertheless, other studies have argued that components in the later part of the visual awareness negativity around 270 msec (corresponding to the P2m of this study) correlate more consistently with awareness and that the fronto-parietal network is involved at this stage and later (Del Cul, Baillet, & Dehaene, 2007; Sergent, Baillet, & Dehaene, 2005). In this study, the same frontal and parietal sources were identified, but little or no difference was found in the source estimates at the M170 and the P2m, and in fact, the frontoparietal sources were identified already at the M170. At both the M170 and the P2m, however, occipital and later face-specific source activity was more predictive than frontal and parietal activity, and early activity (around the M170) was much more predictive than late activity (>300 msec). One reason for the difference in findings, however, could be that these studies, Del Cul et al. and Sergent et al., examined having any experience versus having none (i.e., seeing vs. not seeing), whereas our study examined one conscious content versus another (but participants perceived something consciously on all trials).

Overall, this study appears to support the conclusion that the most consistent correlate of the contents of visual awareness is activity in sensory areas at around 150–200 msec after stimulus onset. Prediction of conscious perception was no more accurate when taking information across multiple time points (and peaks) into account than when training/testing the classifier on the single best time point.

Between-subject Classification

The second question of our study was whether the conscious experience of an individual could be decoded using a classifier trained on a different individual. It is important to note that between-subject classifications of this kind do not reveal neural correlates of consciousness that generally distinguish a conscious from an unconscious

state or whether a particular, single content is consciously perceived or not, but they do allow us to make comparisons between the neural correlates of particular types of conscious contents (here, faces) across individuals.

The data showed that neural signals associated with specific contents of consciousness shared sufficient common features across participants to enable generalization of performance of the classifier. In other words, we provide empirical evidence that the neural activity distinguishing particular conscious content shares important temporal and spatial features across individuals, which implies that the crucial differences in processing are located at similar stages of visual processing across individuals. Nevertheless, generalization between individuals was not perfect, indicating that there are important interindividual differences. Inspecting Figure 3, for instance, it can be seen that the predictive time points around the M170 varied with up to 40 msec between participants (from ~170 msec for S3 to ~210 msec for S2). At present, it is difficult to conclude whether these differences in the neural correlates indicate that the same perceptual content can be realized differently in different individuals or whether they indicate subtle differences in the perceptual experiences of the participants.

Methodological Decisions

The results of the present experiment were obtained by analyzing the MEG signal during binocular rivalry. MEG signals during binocular rivalry reflect ongoing patterns of distributed synchronous brain activity that correlate with spontaneous changes in perceptual dominance during rivalry (Cosmelli et al., 2004). To detect these signals associated with perceptual dominance, the vast majority of previous studies have “tagged” monocular images by flickering them at a particular frequency that can subsequently be detected in the MEG signals (e.g., Kamphuisen, Bauer, & Van Ee, 2008; Srinivasan, Russell, Edelman, & Tononi, 1999; Brown & Norcia, 1997; Lansing, 1964). This method, however, impacts on rivalry mechanisms (Sandberg, Bahrami, Lindelov, Overgaard, & Rees, 2011) and causes a sustained frequency-specific response, thus removing the temporal information in the ERF components associated with normal stimulus processing. This not only biases the findings but also makes comparison between rivalry and other paradigms difficult. To avoid this, yet maintain a high signal-to-noise ratio (SNR), we exploited the stabilization of rivalrous perception associated with intermittent presentation (Noest et al., 2007; Leopold et al., 2002; Orbach, Ehrlich, & Heath, 1963) to evoke signals associated with a specific (stable) percept and time locked to stimulus onset. Such signals proved sufficient to decode spontaneous fluctuations in perceptual dominance in near real-time and in advance of behavioral reports. We suggest that this general presentation method may be used in future ambiguous perception

experiments when examining stimulus-related differences in neural processing.

Potential Confounds

There were two potential confounds in our classification analysis: eye movements and motor responses. These are, however, unlikely to have impacted on the results as source analysis revealed that at the time of maximum classification, sources related to visual processing were most important for explaining the differences related to face and grating perception. Additionally, the fact that the motor response used to signal a perceptual state was swapped between hands and fingers every 20 trials makes it unlikely that motor responses were assigned high weights by the classification algorithm. Nevertheless, our findings of prediction accuracy slightly greater than chance for power in high-frequency bands may conceivably have been confounded by some types of eye movements.

Although we may conclude that specific evoked activity (localized and distributed) is related to conscious experience, this should not be taken as an indication that induced oscillatory components are not important for conscious processing. Local field potentials, for instance, in a variety of frequency bands are modulated in monkeys by perception during binocular rivalry (Wilke, Logothetis, & Leopold, 2006).

Apart from potential confounds in the classification analyses, it could be argued that the use of rotating stimuli alters the stimulus-specific components. The purpose of rotating the stimuli in opposite directions was to minimize the amount of mixed perception throughout the trial (Haynes & Rees, 2005). It is possible, and remains a topic for further inquiries, whether this manipulation affects the mechanisms of the rivalry process, for instance, in terms of stabilization of perception. Inspecting the ERF in Figure 2, it is nevertheless clear that we observed the same face-specific components as are typically found in studies of face perception as reported above. Our M170 was observed slightly later than typically found (peaking at 187 msec). This has previously been observed for partially occluded stimuli (Harris & Aguirre, 2008), and the delay in this study might thus be because of binocular rivalry in general or rotation of the stimuli. The impact of rotating the stimuli upon face-specific components thus appears minimal.

Conclusion

In this study, participants viewed binocular rivalry between a face and a grating stimulus, and prediction of conscious face perception was attempted based on the MEG signal. Perception was decoded accurately in the 120–300 msec time window, peaking around the M170 and again around the P2m. In contrast, little or no above-chance accuracy was found around the earlier M100 component. The find-

ings thus argue against earlier and later components correlating with conscious face perception.

In addition, conscious perception could be decoded from sensors that were individually at chance performance for decoding, whereas this was not the case when decoding using multiple time points. The most informative sensors were located above the occipital and temporal lobes, and a follow-up analysis of activity reconstructed at the source level revealed that the most predictive single sources were indeed found in these areas both at the M170 and the P2m. Nevertheless, conscious perception could be decoded accurately from parietal and frontal sources alone, although not as accurately as from occipital and later ventral stream sources. These results show that conscious perception can be decoded across a wide range of sources, but the most consistent correlates are found both at early and late stages of the visual system.

The impact of increasing the number of temporal features of the classifier was also examined. In contrast to including more spatial features, more temporal features had little or no impact on classification accuracy. Furthermore, the predictive strength of power estimation was examined across a wide range of frequency bands. Generally, the low frequencies contained in the evoked response were the most predictive and the peak time points of classification accuracy coincided with the latencies of the M170 and the P2m. This indicates that the main MEG correlates of conscious face perception are the two face-sensitive components, the M70 and the P2m.

Finally, the results showed that conscious perception of each participant could be decoded above chance using classifiers trained on the data of each of the other participants. This indicates that the correlates of conscious perception (in this case, faces) are shared to some extent between individuals. It should be noted, though, that generalization was far from perfect, indicating that there are significant differences as well for further exploration.

APPENDIX

Improving Decoding Accuracy

We hypothesized that decoding accuracy could be increased in two ways: by rejecting trials for which perception was not completely clear and by applying a more stringent filter to the data. Participant's reports (see Results) suggested that the probability of clear perception on a given trial increased the further away the trial is from a perceptual switch. Classifiers were thus trained and tested on unstable perception (Trials 1–9 after a switch) and stable perception (Trial 10 or more after a switch) separately and decoding accuracies were compared. Five participants reported 100 trials of all kinds (stable/unstable faces/gratings) required for training the classifier, and the analysis was thus based on these. Figure A1a shows that analyzing stable trials as compared

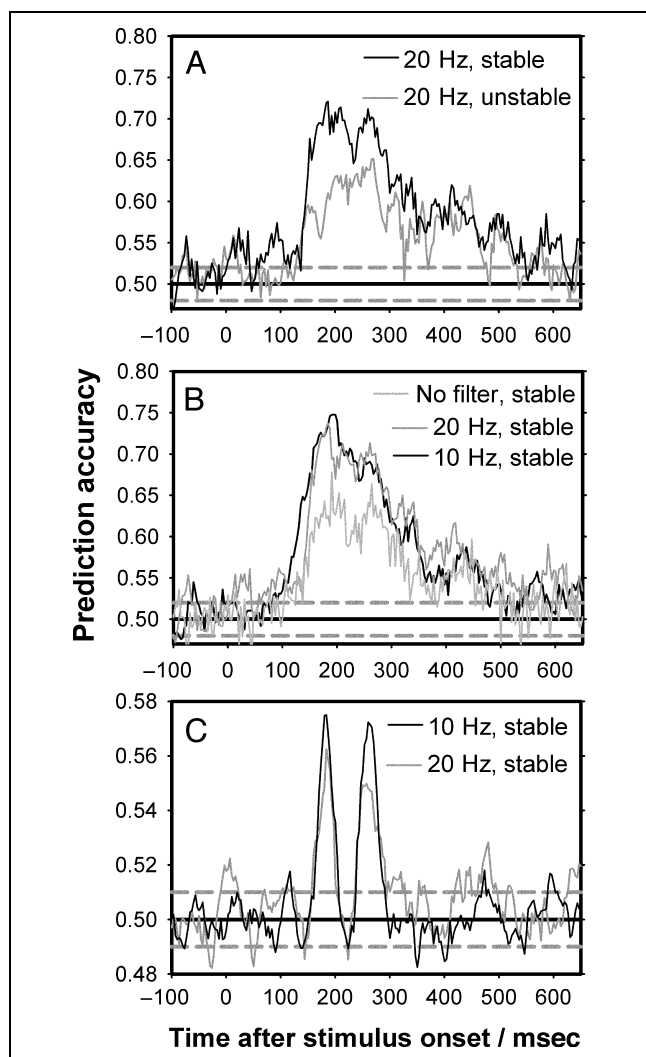


Figure A1. Improvements to prediction accuracy by filtering and trial selection. The figure plots the impact of using stable trials only as well as filtering the data. Dotted gray line represents the 95% binomial confidence interval around chance (uncorrected). (A) Prediction accuracy for stable and unstable trials, respectively. The comparison is based on the five participants who reported enough trials of all conditions (stable/unstable faces/gratings) to train the classifiers. (B, C) Within-subject (B) and between-subject (C) prediction accuracy for data that has not been low-pass filtered compared with data low-pass filtered at 20 and 10 Hz, respectively. This analysis was based on stable trials, and the data reported are from the analysis of the six participants reporting enough stable face and grating trials to train the classifier.

with unstable trials results in a large improvement in classification accuracy of around 10–15% around the M170 (~187 msec), 5–8% around the P2m (~260 msec), and similarly 5–8% around the M100 (~93 msec). Significant improvements in classification accuracy was found for at least three of five participants for all components (cumulative $p = .0012$, uncorrected).

Some components analyzed (M100, M170, and P2m) had a temporal spread of around 50–130 msec (see Figure A1a–c), yet the classifiers were trained on single time

points only in the analyses above. This makes classification accuracy potentially vulnerable to minor fluctuations at single time points. Such fluctuations could reflect small differences in latency between trials as well as artifacts and high-frequency processes that the classifier cannot exploit, and analyses based on field strength data may thus be improved if the impact of these high-frequency components and trial-by-trial variation is minimized. There are two methods to do this: classification may either use several neighboring time points or a low low-pass filter may be applied before analysis to temporally smooth the data.

Given the temporal extent of the three analyzed components (50–130 msec), they can be seen as half cycles of waves with frequencies of 4–10 Hz (i.e., around 100–250 msec). For this reason, we compared classification accuracies for nonfiltered data, 1–20 Hz filtered data, and 2–10 Hz filtered data. We used only stable trials. Six participants had 100 stable trials or more of each kind (face/grating) and were thus included in the analysis.

Figure A1b shows the differences between the three filter conditions for within-subject decoding. Improvement in decoding accuracy was found comparing no filter and the filtered data. Comparing unfiltered and 1–20 Hz filtered data at the M170 and P2m, differences of 5–10% were found around both peaks, and around the M100 a difference of around 5% was found. Decoding accuracy was significantly higher for five of six participants at the 187 msec (cumulative probability of $p = 1.9 \times 10^{-6}$, uncorrected) and for four of six participants at 260 msec (cumulative probability of $p = 8.7 \times 10^{-5}$, uncorrected), but only for two of six participants at 90 msec (cumulative probability of $p = .03$, uncorrected). The largest improvement of applying a 20-Hz low-pass filter was thus seen for the two most predictive components, the M170 and the P2m. The only impact of applying a 2–10 Hz filter instead of a 1–20 Hz filter was significantly increased accuracy for two participants at 187 msec, but decreased for one.

As between-subject ERF variation is much larger than within-subject variation (Samthein, Andersson, Zimmermann, & Zumsteg, 2009), we might expect that the most stringent filter mainly improved between-subject decoding accuracy. Figure A1c shows a 2–3% improvement of using a 2–10 Hz compared with a 1–20 Hz filter at the M170 and the P2m and a <1% improvement at the M100. This improvement was significant for two participants at the 180 and 260 msec (cumulative $p = .03$, uncorrected), for both, and one participant around the M100 at 117 msec (cumulative $p = .27$, uncorrected).

Overall, the best decoding accuracies were achieved using stable trials and filtered data. Numerically better and slightly more significant results were achieved using 2–10 Hz filtered data compared with 1–20 Hz filtered data. Importantly, using this more stringent filter did not alter the time points for which conscious perception

could be decoded—it only improved accuracy around the peaks.

Decoding Using Power Estimations

Power in several frequency bands (for all sensors) was also used to train SVM classifiers. This analysis revealed

that theta band power was the most highly predictive of perception followed by alpha power (Figure A2). Again the data were the most informative at around 120–320 msec after stimulus onset. Power estimates in the higher-frequency bands related to both face and grating perception (40–60 Hz) and possibly also some related to face perception alone (60–80 Hz) could be used

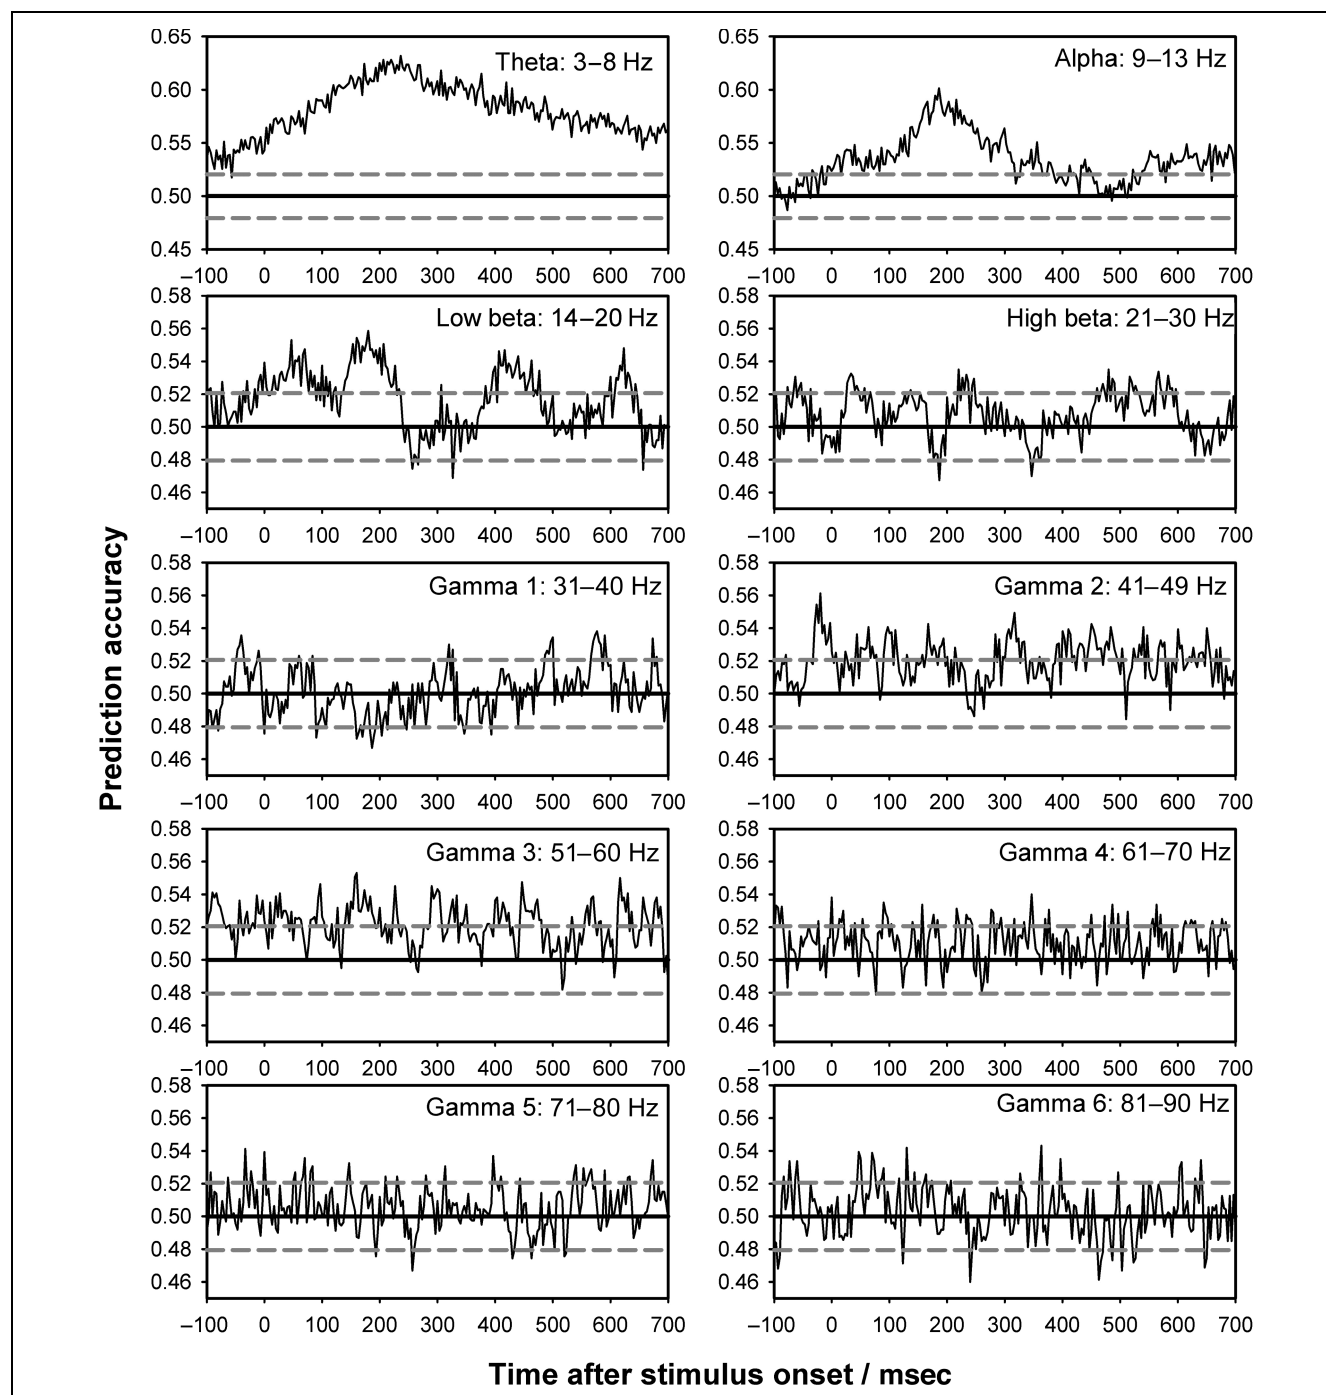


Figure A2. Prediction accuracy across time for various frequencies (stable trials). Six participants had enough trials to train the classifiers on stable trials alone. The figure plots the data from these participants. The dotted gray line indicates the threshold for which a binomial distribution of the same number as the total number of trials the prediction is performed upon is different from chance (uncorrected). Average prediction accuracy is plotted across participants based on estimates of power in different frequency bands as a function of time. SVMs were trained to predict reported perception (face vs. grating) for each time point.

to predict perception significantly better than chance (Duncan et al., 2010; Engell & McCarthy, 2010). In these bands, the prediction accuracy did not have any clear peaks (Figure A2).

Using Bonferroni correction, average prediction accuracies across participants across the stimulation period were above chance in the theta ($t(7) = 4.4, p = .033$), gamma 2 (40–49 Hz) ($t(7) = 4.9, p = .017$), and gamma 3 (51–60 Hz) ($t(7) = 4.2, p = .038$) bands. Without Bonferroni correction, alpha ($t(7) = 3.2, p = .0151$), low beta ($t(7) = 3.7, p = .0072$), high beta ($t(7) = 3.1, p = .0163$), gamma 4 (61–70 Hz) ($t(7) = 3.3, p = .0123$), and gamma 5 (71–80 Hz) ($t(7) = 2.4, p = .0466$) were also above chance.

The classification performance based on the moving window spectral estimate was always lower than that based on the field strength. Also, spectral classification was optimal for temporal frequencies dominating the average evoked response (inspecting Figure 2B, C, it can be seen, for instance, that for faces, the M170 is half a cycle of a 3–4 Hz oscillation). Taken together, this suggests that the predictive information was largely contained in the evoked (i.e., with consistent phase over trials) portion of the single trial data.

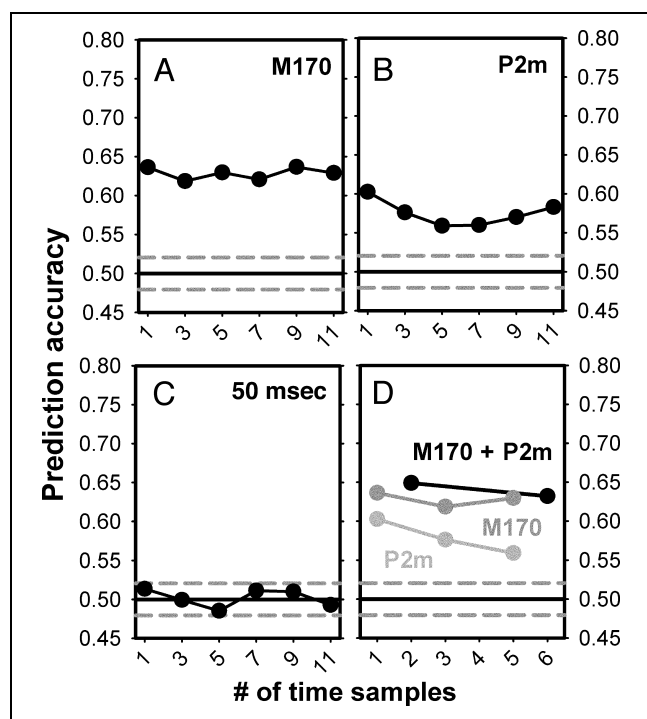


Figure A3. Prediction based on multiple time points (stable trials). Six participants had enough trials to train the classifiers on stable trials alone. The figure plots the data from these participants. Classifiers were trained/tested on 1 Hz high-pass filtered data from 16 randomly distributed sensors. (A–C) Prediction accuracy as a function of the number of neighboring time samples used to train the classifier around the M170 peak (A), the P2m peak (B), and 50 msec after stimulus onset (C). No improvement was found at the peaks nor at 50 msec when classifier baseline accuracy was close to chance. (D) Prediction accuracy when classifiers were trained on data around both peaks combined versus each peak individually.

Decoding Using Multiple Time Points

The potential benefit of including multiple time points when training classifiers was examined. As multiple time points increase the number of features drastically, the SVM was trained on a subset of sensors only. For these analyses, 16 randomly selected sensors giving a performance of 72.6% when trained on a single time point were used (see Figure 4A). As the temporal smoothing of low-pass filter would theoretically remove any potential benefit of using multiple time points for time intervals shorter than one cycle of activity, these analyses were performed 1 Hz high-pass filtered data. Here, the sampling frequency of 300 Hz is thus the maximum frequency.

We tested the impact of training on up to 11 time points (37 msec) around each peak (M170 and P2m) and around a time point for which overall classification accuracy was at chance (50 msec). At 50 msec, the signal should have reached visual cortex, but a 37-msec time window did not include time points with individual above-chance decoding accuracy. We also tested the combined information around the peaks. As seen in Figure A3, the inclusion of more time points did not increase accuracy, and the use of both peaks did not increase accuracy beyond that obtained at the M170 alone. This may indicate that the contents of consciousness (in this case, rivalry between face and grating perception) are determined already around 180 msec.

Acknowledgments

This work was supported by the Wellcome Trust (G. R. and G. R. B.), the Japan Society for the Promotion of Science (R. K.), the European Commission under the Sixth Framework Programme (B. B., K. S., M. O.), the Danish National Research Foundation and the Danish Research Council for Culture and Communication (B. B.), and the European Research Council (K. S. and M. O.). Support from the MINDLab UNIK initiative at Aarhus University was funded by the Danish Ministry of Science, Technology, and Innovation.

Reprint requests should be sent to Dr. Kristian Sandberg, Cognitive Neuroscience Research Unit, Aarhus University Hospital, Noerrebrogade 44, Building 10G, 8000 Aarhus C, Denmark, or via e-mail: krisand@rm.dk.

REFERENCES

- Babiloni, C., Vecchio, F., Buffo, P., Buttiglione, M., Cibelli, G., & Rossini, P. M. (2010). Cortical responses to consciousness of schematic emotional facial expressions: A high-resolution EEG study. *Human Brain Mapping, 31*, 1556–1569.
- Blake, R. (2001). A primer on binocular rivalry, including current controversies. *Brain and Mind, 2*, 5–38.
- Brascamp, J. W., Knapen, T. H. J., Kanai, R., Noest, A. J., Van Ee, R., Van den Berg, A. V., et al. (2008). Multi-timescale perceptual history resolves visual ambiguity. *PLoS One, 3*, e1497.
- Breese, B. B. (1899). On inhibition. *Psychological Monographs, 3*, 1–65.
- Brown, R. J., & Norcia, A. M. (1997). A method for investigating binocular rivalry in real-time with the steady-state VEP. *Vision Research, 37*, 2401–2408.

- Carlson, T. A., Hogendoorn, H., Kanai, R., Mesik, J., & Turret, J. (2011). High temporal resolution decoding of object position and category. *Journal of Vision, 11*, 9.1–9.17.
- Carter, O., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PLoS One, 2*, e343.
- Cosmelli, D., David, O., Lachaux, J.-P., Martinerie, J., Garnero, L., Renault, B., et al. (2004). Waves of consciousness: Ongoing cortical patterns during binocular rivalry. *Neuroimage, 23*, 128–140.
- Del Cul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology, 5*, e260.
- Duncan, K. K., Hadjipapas, A., Li, S., Kourtzi, Z., Bagshaw, A., & Barnes, G. (2010). Identifying spatially overlapping local cortical networks with MEG. *Human Brain Mapping, 31*, 1003–1016.
- Engell, A. D., & McCarthy, G. (2010). Selective attention modulates face-specific induced gamma oscillations recorded from ventral occipitotemporal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 30*, 8780–8786.
- Freeman, A. W. (2005). Multistage model for binocular rivalry. *Journal of Neurophysiology, 94*, 4412–4420.
- Friston, K. J., Harrison, L., Daunizeau, J., Kiebel, S., Phillips, C., Trujillo-Barreto, N., et al. (2008). Multiple sparse priors for the M/EEG inverse problem. *Neuroimage, 39*, 1104–1120.
- Harris, A. M., & Aguirre, G. K. (2008). The effects of parts, wholes, and familiarity on face-selective responses in MEG. *Journal of Vision, 8*, 4.1–4.12.
- Harris, J. A., Wu, C.-T., & Woldorff, M. G. (2011). Sandwich masking eliminates both visual awareness of faces and face-specific brain activity through a feedforward mechanism. *Journal of Vision, 11*, 3.1–3.12.
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*, 223–233.
- 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., & Rees, G. (2005). Predicting the stream of consciousness from activity in human visual cortex. *Current Biology: CB, 15*, 1301–1307.
- Kamphuisen, A., Bauer, M., & Van Ee, R. (2008). No evidence for widespread synchronized networks in binocular rivalry: MEG frequency tagging entrains primarily early visual cortex. *Journal of Vision, 8*, 4.1–4.8.
- Koivisto, M., & Revonsuo, A. (2010). Event-related brain potential correlates of visual awareness. *Neuroscience & Biobehavioral Reviews, 34*, 922–934.
- Kornmeier, J., & Bach, M. (2004). Early neural activity in Necker-cube reversal: Evidence for low-level processing of a Gestalt phenomenon. *Psychophysiology, 41*, 1–8.
- Lansing, R. W. (1964). Electroencephalographic correlates of binocular rivalry in man. *Science (New York, N.Y.), 146*, 1325–1327.
- Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience, 5*, 605–609.
- Liddell, B. J., Williams, L. M., Rathjen, J., Shevrin, H., & Gordon, E. (2004). A temporal dissociation of subliminal versus supraliminal fear perception: An event-related potential study. *Journal of Cognitive Neuroscience, 16*, 479–486.
- Lumer, E. D., Friston, K. J., & Rees, G. (1998). Neural correlates of perceptual rivalry in the human brain. *Science (New York, N.Y.), 280*, 1930–1934.
- Noest, Van Ee, R., Nijs, M. M., & Van Wezel, R. J. (2007). Percept-choice sequences driven by interrupted ambiguous stimuli: A low-level neural model. *Journal of Vision, 7*, 1–14.
- Orbach, J., Ehrlich, D., & Heath, H. (1963). Reversibility of the Necker cube: I. An examination of the concept of “satiation of orientation”. *Perceptual and Motor Skills, 17*, 439–458.
- Pegna, A. J., Darque, A., Berrut, C., & Khateb, A. (2011). Early ERP modulation for task-irrelevant subliminal faces. *Frontiers in Psychology, 2*, 88.1–88.10.
- Pegna, A. J., Landis, T., & Khateb, A. (2008). Electrophysiological evidence for early non-conscious processing of fearful facial expressions. *International Journal of Psychophysiology, 70*, 127–136.
- Raizada, R. D. S., & Connolly, A. C. (2012). What makes different people’s representations alike: Neural similarity space solves the problem of across-subject fMRI decoding. *Journal of Cognitive Neuroscience, 24*, 868–877.
- Sandberg, K., Bahrami, B., Lindelov, J. K., Overgaard, M., & Rees, G. (2011). The impact of stimulus complexity and frequency swapping on stabilization of binocular rivalry. *Journal of Vision, 11*, 1–10.
- Sandberg, K., Barnes, G., Bahrami, B., Kanai, R., Overgaard, M., & Rees, G. (submitted). Distinct MEG correlates of conscious experience, perceptual reversals, and stabilization during binocular rivalry.
- Sarnthein, J., Andersson, M., Zimmermann, M. B., & Zumsteg, D. (2009). High test–retest reliability of checkerboard reversal visual evoked potentials (VEP) over 8 months. *Clinical Neurophysiology, 120*, 1835–1840.
- 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.
- Srinivasan, R., Russell, D. P., Edelman, G. M., & Tონoni, G. (1999). Increased synchronization of neuromagnetic responses during conscious perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 19*, 5435–5448.
- Wilke, M., Logothetis, N. K., & Leopold, D. A. (2006). Local field potential reflects perceptual suppression in monkey visual cortex. *Proceedings of the National Academy of Sciences, U.S.A., 103*, 17507–17512.
- Wilson, H. R. (2003). Computational evidence for a rivalry hierarchy in vision. *Proceedings of the National Academy of Sciences, U.S.A., 100*, 14499–14503.
- Wilson, H. R. (2007). Minimal physiological conditions for binocular rivalry and rivalry memory. *Vision Research, 47*, 2741–2750.