

Magnetoencephalographic Activity Related to Conscious Perception Is Stable within Individuals across Years but Not between Individuals

Kristian Sandberg^{1,2,5}, Gareth Robert Barnes², Geraint Rees^{2,3},
and Morten Overgaard^{1,4}

Abstract

■ Studies indicate that conscious perception is related to changes in neural activity within a time window that varies between 130 and 320 msec after stimulus presentation, yet it is not known whether such neural correlates of conscious perception are stable across time. Here, we examined the generalization across time within individuals and across different individuals. We trained classification algorithms to decode conscious perception from neural activity recorded during binocular rivalry using magnetoencephalography (MEG). The classifiers were then used to predict the perception of the same participants during different recording sessions either days or years later as well as between different participants. No drop in decoding accuracy

was observed when decoding across years compared with days, whereas a large drop in decoding accuracy was found for between-participant decoding. Furthermore, underlying percept-specific MEG signals remained stable in terms of latency, amplitude, and sources within participants across years, whereas differences were found in all of these domains between individuals. Our findings demonstrate that the neural correlates of conscious perception are stable across years for adults, but differ across individuals. Moreover, the study validates decoding based on MEG data as a method for further studies of correlations between individual differences in perceptual contents and between-participant decoding accuracies. ■

INTRODUCTION

The last two decades have seen an upsurge of empirical research into the neural correlates of conscious perception, yet it remains unknown whether candidate neural signatures of conscious perception determined by such earlier work are stable across time. Here, we addressed this question by measuring and comparing magnetoencephalography (MEG) signals associated with conscious perception that were acquired from the same individuals days or years apart.

We focused on one very consistent neural correlate of conscious perception, the so-called visual awareness negativity (VAN; Koivisto & Revonsuo, 2003). VAN refers to awareness-specific event-related EEG activity occurring within a time window of around 130–320 msec after stimulus presentation and has been observed in more than 30 independent EEG studies (Koivisto & Revonsuo, 2010), and at least three MEG experiments have reported activity corresponding to the VAN (Sandberg et al., 2013; Liu, Paradis, Yahia-Cherif, & Tallon-Baudry, 2012; Vanni, Revonsuo, Saarinen, & Hari, 1996). The VAN consists of

an early and a late part (Koivisto & Revonsuo, 2010), and in some experiments, these two parts of the VAN are observed as separate ERP/event-related field (ERF) components (e.g., Sandberg et al., 2013; Fahrenfort, Scholte, & Lamme, 2007).

Although there are differences in the interpretation of the VAN and its individual parts/components, studies with very different theoretical backgrounds consistently report VAN-like findings. Koivisto and Revonsuo (2010), for instance, interpret the VAN as reflecting recurrent processing in sensory areas associated with conscious experience in the theory of Lamme (2010), and research by Lamme's own group report activity around the first VAN component at 110–210 msec as the main correlate of visual awareness in masking tasks (Van Loon, Scholte, van Gaal, van der Hoort, & Lamme, 2012; Fahrenfort et al., 2007). In contrast, Dehaene and others report that signals around the second VAN component, peaking around 270 msec, correlate with subjective, graded ratings of visibility, although they consider later, bimodal responses as the correlates of conscious report (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Sergent, Baillet, & Dehaene, 2005; Sergent & Dehaene, 2004; Dehaene, Kerszberg, & Changeux, 1998). In this study, we remain agnostic regarding the theoretical interpretation of the VAN, but simply use it as the object of our analysis as one of the components that correlate with the content of consciousness.

¹Aarhus University Hospital, ²University College London, ³Wellcome Trust Centre for Neuroimaging, Institute of Neurology, London, UK, ⁴Aalborg University, ⁵Hammel Neurorehabilitation and Research Center

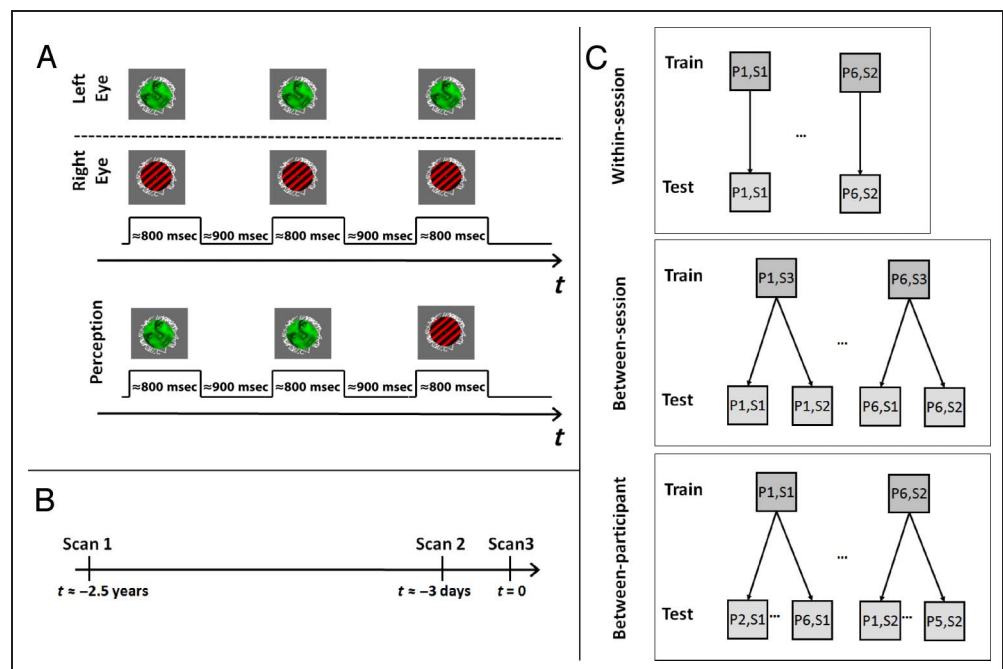
Several studies have examined the stability of different aspects of the EEG signal within individuals across time. Many such studies have focused on visual-evoked potentials (VEPs), typically evoked by a reversing checkerboard stimulus, which elicit two early components, an N75 (at around 75 msec) and a P100 (at around 100 msec). For latency, moderate-to-strong correlations of 0.3–0.8 is often found across 0.2–13 months, although for the P100 correlations as high as .93 have been observed (Sarnthein, Andersson, Zimmermann, & Zumsteg, 2009; Oken, Chiappa, & Gill, 1987; Hall, Rappaport, Hopkins, & Griffin, 1973). Also for amplitude, defined as the N75-P100 difference, moderate correlations of .4–.7 have been reported (Sarnthein et al., 2009; Schellberg, Gasser, & Köhler, 1987). One study reported that interindividual VEP differences were generally larger than intraindividual differences (Sarnthein et al., 2009), although for children aged 10–13 years, intraindividual differences across 10 months appear to be of the same magnitude as interindividual differences (Schellberg et al., 1987). Other studies have examined error-related negativity, a component found around 50 msec after the individual makes an error, and found a moderate-to-strong amplitude correlation across 2–6 weeks and a moderate correlation across 1.5–2.5 years (Weinberg & Hajcak, 2011).

Finally, studies have reported relatively high stability of the EEG power spectra related to resting as well as related to performing a working memory task across up to 40 months (Näpflin, Wildi, & Sarnthein, 2007, 2008).

Although these previous studies indicate that some degree of interindividual difference is to be expected and possibly that smaller intraindividual difference is to be expected, it is difficult to draw parallels between studies because of the difference in the examined components. Most of the reported studies were concerned with very early components, which are generally not considered to be modulated by conscious content. Furthermore, this study examines not only the components related to a certain visual stimulation but also endogenous, consciousness-specific modulation of these components, thus making it difficult to make predictions regarding the stability of activity in the VAN time window across years based on previous studies.

To test the stability of the VAN over extended periods of time, we used intermittent binocular rivalry (BR; Figure 1A). During BR, two images are presented dichoptically, and perception alternates spontaneously between the two images. When BR stimuli are presented intermittently, that is, for short durations of less than a second separated by a short break with no visual stimulation, the

Figure 1. Experimental design. (A) Rivaling stimuli (face/grating) were presented in trials of ~800 msec separated by blank periods of ~900 msec. Stimuli were dichoptically presented to the eyes of the participant 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. (B) Time line of recording sessions. Session 1 was performed around 2.5 years before Session 2, which was performed 2–5 days before Session 3. (C) Classification procedure. SVMs were trained to distinguish MEG 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 (P) on different trials from the same recording session (S; top), (2) trials from the same participant but from a different recording session (middle), and (3) trials from recording sessions from each of the other participants (bottom).



participant typically reports only one percept per trial, thus allowing examination of percept-specific signals that are time-locked to stimulus onset.

We examined whether multivariate support vector machines (SVM) could decode the percept that a participant reported on any given trial when trained and tested on different data sets, that is, when the training data came from the same participant, but was gathered days or even years apart from the testing data, or when the training and testing data came from different individuals (Figure 1B, C). The degree to which prediction accuracy did (or did not) fall when generalizing across data sets in this way served as an index of the stability of the multivariate signal that correlated with conscious perception at baseline, either over time within the same individual or across individuals at the same time point. We further explicitly examined the latency and amplitude of ERF components during the VAN as well as the activated sources to interpret changes in decoding accuracy more fully.

The MEG of eight participants was recorded in Session 1, and the analyses reported here were based on an in-depth examination of four participants who were available 2.5 years later. The full analysis of all eight participants of recording Session 1 is reported elsewhere (Sandberg et al., 2013; see Methods for details). We emphasize that the drop in between-participant prediction accuracy reported in this study of four participants is highly similar to the drop previously reported for the full group of eight participants, thus indicating that the current four participants are representative of the group. We further emphasize that significant differences observed with a relatively small sample of participants reflect that the effect size is larger (and thus more relevant) than if a larger sample of participants were needed to obtain the same level of significance (Friston, 2012).

METHODS

As mentioned above, comprehensive analyses of the data from Session 1 are reported elsewhere (Sandberg et al., 2013). The goal of this previous study was to examine the MEG correlates of the content of consciousness during intermittent BR. The main analysis was decoding using an SVM classifier, and decoding was based on field strength amplitude (using various filters) as well as power estimates across a wide range of power spectra, and it was established that field strength amplitude was the most informative signal. Furthermore, the potential advantage of feeding the SVM with data across multiple time points was examined, yet no such advantage was found. Analyses were performed at the sensor level as well as on data reconstructed at a number of sources distributed across cortex achieving similar results in the two conditions. Additionally, successful decoding across participants was achieved, yet a relatively large drop in decoding accuracy was observed.

On the basis of the results of this previous study, decoding using field strength data at single time points was chosen as the main method of analysis in this study. This study expands on the previous by examining the stability of consciousness-specific modulation of neural activity across time for the first time and by performing an in-depth analysis of the cause of the drop in between-participant decoding. To accomplish this, the approach was supplemented by three new analyses. Specifically, we estimated the power and latency of the components previously found to predict conscious content, and we used multivariate Bayesian (MVB) modeling to estimate the source composition of these components. All these methods are explained in-depth below.

Participants

Eight healthy adults (six women) with normal or corrected-to-normal vision gave written informed consent to participate in the experiment and performed recording Session 1. Four of these (three women) were available 2.5 years later for recording Sessions 2 and 3. Data for these four participants were analyzed here. The ages of the participants was between 26 and 32 years (mean = 29 years, $SD = 2.5$ years). 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. Head localizer coils were attached to the nasion and 1 cm anterior of the left and right outer canthus to monitor head movement during the recording sessions and to establish head position for source analysis.

Stimuli

A red Gabor patch (contrast = 100%, spatial frequency = 3 cycles/degree, SD of the Gaussian envelope = 10 pixels) was presented to the right eye of the participants, and a green, female face was presented to the left eye (Figure 1A). To avoid piecemeal rivalry, the stimuli rotated at a rate of 0.7 rotations per second 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

The MEG of the participants was recorded on three separate occasions. Recording 1 was at around $t = -2.5$ years, Recording 2 was at around $t = -3$ days, and Recording 3 was at $t = 0$. For all recordings, the following procedure was used.

First, the stereoscope was calibrated by adjusting the mirrors until the fixation circles around the stimuli fused. Second, 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. During both these calibrations runs and the actual 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.

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 (this stimulation period was calibrated individually for each participant so that the percept had time to form and did not switch during the stimulation period—for all participants, the stimulus duration was between 750 and 850 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 downsampled 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 200 msec before stimulus onset and ending 600 msec after. Trials containing artifacts were rejected at a threshold of 3 pT.

Source Space Activity Reconstruction

To minimize confounds caused by differences in the head position of the participant in the MEG system between recording sessions, some analyses were performed on data reconstructed in source space. The source reconstruction took into account the head position of the participants, which was measured at the beginning and end of each run using localizer coils as described above. The mean difference in position between start and end head position was 3.4 mm ($SD = 2.1$). We identified the sources that were most active 0–400 msec after stimulus onset

using the multiple sparse priors (MSP) algorithm (Friston, Harrison, et al., 2008). 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 performed a group-level reconstruction based on template structural MR scans using all trials (labeled identically) from all 12 recording sessions (4 participants \times 3 sessions/participant). The 35 most active sources were identified (see Table 1). SPM was then used to reconstruct the activity of the correctly (face/grating) labeled trials across the 35 sources for each recording, and these data sets were used in the analyses below. Note that the use of generic anatomy (head model and cortical mesh) means that differences in source reconstruction will be due primarily to differences in sensor level data (rather than differences in cortical folding over individuals). For this reason, the use of generic models will result in more similar source estimates over subjects than the use of models based on individual MRIs, and larger differences between individuals than within individuals (between sessions) cannot be an artifact of the source estimation process.

Multivariate Prediction Analysis

Multivariate pattern classification of the evoked responses was performed using the linear SVM of the MATLAB Bioinformatics Toolbox (Mathworks, Natick, MA). The SVM attempted to decode the trial type (face or grating) independently for each time point along the epoch. Classification was based on 2–10 Hz filtered data, as the components of interest in the 130–320 msec time windows are < 10 Hz. Application of a 1–20 Hz filter or a 1-Hz high-pass filter only does not change MEG results qualitatively (Sandberg et al., 2013).

We attempted to decode conscious perception both within participants across time and between participants. For all analyses, 100 randomly selected trials of each kind (face/grating perception) were used to ensure a balanced training scheme. For within-participant training/testing, two analyses were performed. First, we trained and tested the classifier using data from the same recording session for Sessions 1 and 2. This formed a baseline against which between-session decoding could be compared. For this within-session decoding, 10-fold cross-validation was used to ensure that the classifier was never trained and tested on the same data. Next, for between-session training/testing, the SVMs were trained on all 100 trials from Session 3 for each participant separately and tested on all trials of Sessions 1 and 2 (separately) for that participant. This allowed us to compare between-session decoding accuracy between training/testing data gathered days versus years apart. For between-participant training/testing, the SVM was trained on all 100 trials from a single participant and tested on all 100 trials of each of the remaining participants. The process was repeated until data from all participants

Table 1. The 35 Sources Found to Be Most Active across All Trials Independently of Perception across All Participants

Source No.	Area	Name	x	y	z
1	Occipital	IV1	-2	-96	5
2		lvOCC1	-15	-93	-16
3		rvOCC1	21	-96	-16
4		lvOCC2	-18	-80	-11
5		rvOCC2	18	-80	-10
6		lvOCC3	-26	-70	-8
7		rvOCC3	26	-71	-6
8		ldOCC	-18	-80	39
9		rdOCC	21	-81	40
10		ldOCC2	-8	-80	29
11		rdOCC2	10	-79	28
12	Face specific	lOFA	-39	-80	-15
13		rOFA	39	-81	-16
14		lpSTS1	-52	-64	-13
15		rpSTS1	51	-64	-9
16		lpSTS2	-54	-37	26
17		rpSTS2	54	-38	26
18		laIT	-33	-20	-31
19		raIT	33	-19	-30
20		lpIT	-56	-9	-24
21		rpIT	60	-9	-19
22	Parietal	lSPL1	-38	-38	60
23		rSPL1	36	-37	60
24		lSPL2	-32	-66	50
25		rSPL2	35	-66	46
26		lSPL3	-21	-53	59
27		rSPL3	19	-54	61
28	Precentral	lPC	-53	-9	13
29		rPC	53	-10	13
30	Prefrontal	raMFG	45	17	27
31		laMFG	-45	20	30
32		lPFC	-30	27	36
33		rPFC	31	30	37
34		lOFC	-37	37	-16
35		rOFC	35	38	-13

Sources were localized using MSPs to solve the inverse problem. Source abbreviations: V1 = striate cortex; OCC = occipital lobe; OFA = occipital face area; 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.

had been used to train the SVM. For data from Sessions 1 and 2, we thus obtained decoding accuracies for when the classifier was trained on data from the same session, on data from a separate session, and on data from different participants.

Increasing Signal-to-Noise Ratio

During intermittent BR, participants typically report perceiving the same stimuli on many consecutive trials (Leopold, Wilke, Maier, & Logothetis, 2002), and the signal-to-noise ratio is slightly lower immediately after the report of a perceptual switch than when perception has stabilized (Sandberg et al., 2013). For this reason, we used only trials for which the same percept had been reported five or more times (trials with so-called “stable” perception). On average, participants reported stable face perception on 40.6% of the trials ($SD = 11.8$) and stable grating perception on 28.8% of the trials ($SD = 17.2$). Trials with mixed perception (13.2%, $SD = 10.6$), unstable face perception (7.8%, $SD = 4.7$), and unstable grating perception (9.6%, $SD = 7.2$) made out the rejected 30.6% ($SD = 20.5$) of the trials.

ERF Analysis

Traditional, univariate ERF analysis was performed to examine differences in the timing and amplitude of the percept-specific activity across sessions/participants. For these analyses, data were band-pass filtered at 1–20 Hz and averaged individually for the relevant conditions using SPM8. This filter was similar to that used in ERP studies of ambiguous perception employing an intermittent presentation paradigm (Kornmeier & Bach, 2004, 2005). Results remained qualitatively unchanged when a 2–10 Hz filter was used as in the decoding analyses.

The peak latencies of percept-specific activity were defined as the peaks in the VAN in the 100–400 msec time window. For all recordings, a clear VAN peak was observed around two face-specific components (the M170 around 180 msec and P2m around 270 msec). These face-grating activity differences (i.e., VAN peaks) were used in the analyses of latency and amplitude differences. The topographies of the VANs are plotted in Figure 2 for illustrative purposes. Note, however, that the topography is dependent upon head position in relation to the sensors, and differences in topography should not be interpreted as differences in the underlying signal. Head position does not impact upon measures of peak latency, and this analysis is thus relatively noise-free. Head position does, however, impact to some extent upon measures of amplitude differences. To minimize the impact on this, amplitude differences were measured at the sensor for which the largest difference was observed rather than at the same sensor for all recordings. This solution is not perfect, and for this reason, we expected measures of amplitude differences

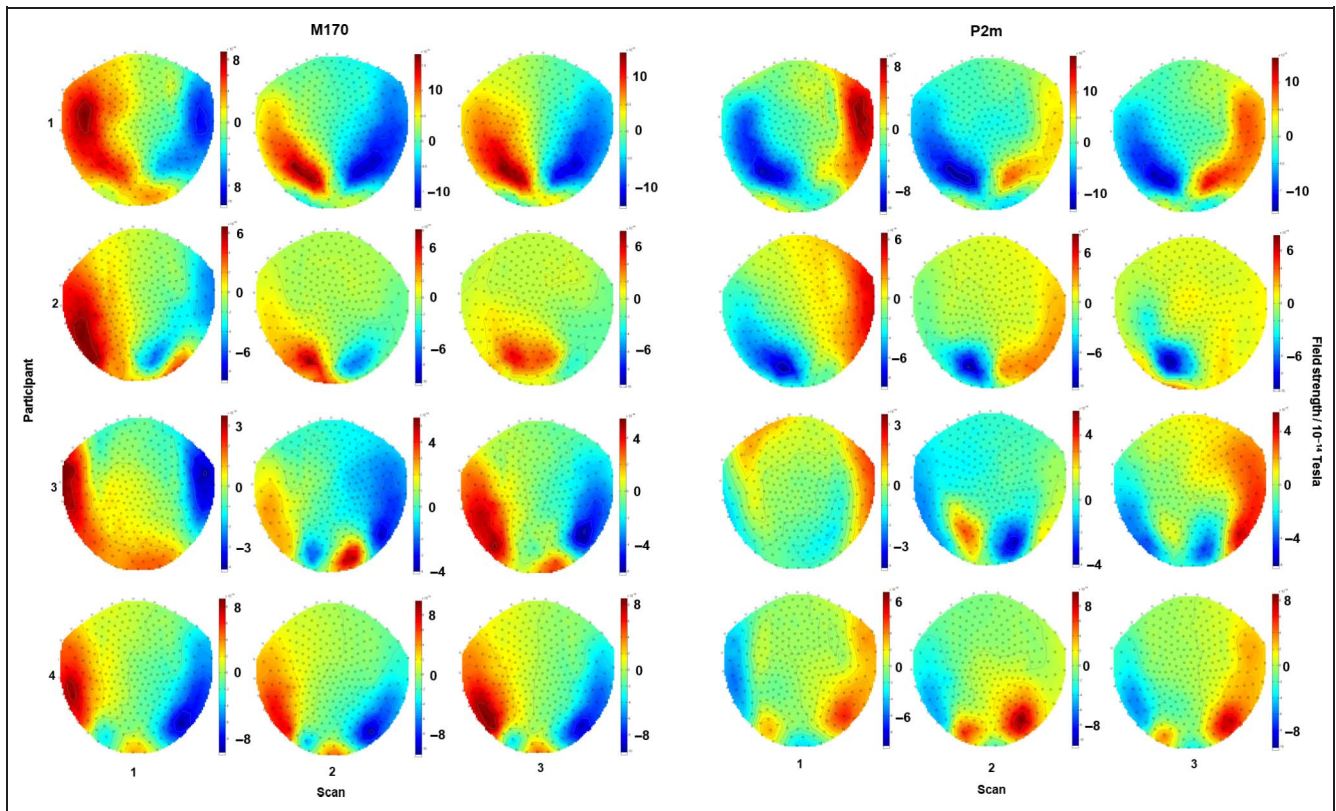


Figure 2. Topography of the face-grating perception contrast at the two peaks in difference. The first peak (left) is observed around the face-specific M170, and the second peak (right) is observed around face-specific P2m.

to be slightly noisier than the other measures. We return to this issue in the Discussion. For illustration, ERFs around the peak sensors are plotted in Figure 3. Note that these are not obtained at the same sensor across recordings and are thus difficult to compare. The source space multivariate between-session and between-participant decoding provides a better measure of difference/similarity.

MVB Model Testing

To test for differences regarding where in the cortex activity varied as a function of perception across recordings/participants, MVB models (Friston, Chu, et al., 2008) were made individually for all 35 sources for all 12 data sets. These models show the relative importance of one site compared with another in explaining activity differences within data set, and the ranking of sites was compared across recording sessions and participants.

Statistical Testing

All statistical tests were two-tailed. For analyses comparing decoding accuracy across conditions (e.g., comparing within-session to between-session accuracy), decoding accuracy was averaged across participants and compared using all samples within the time window where VAN is found (140–310 msec, 52 time points). As the various

conditions had unequal variances (within-participant decoding averages were based on 200 values per participant per time point, whereas between-participant decoding were based on 600), nonparametric Kruskal–Wallis tests were used for analyses. These were followed up by Student’s *t* tests that are robust against unequal variances when sample sizes are equal (as was the case).

For analyses comparing timing as well as amplitude of the ERF components between sessions and participants, nonparametric Kruskal–Wallis tests were used. Kruskal–Wallis tests were also used to examine the null hypotheses that the same sources explained the activity difference between face and grating perception across sessions and participants.

RESULTS

As mentioned above, MEG signals were recorded on three separate occasions from each of four participants while they viewed BR between face/grating stimuli and reported their perception (Figure 1). Recording Session 1 was at around $t = -2.5$ years, Session 2 was at around $t = -3$ days, and Session 3 was at $t = 0$. Perception was decoded from field strength data using SVMs, which were trained either on data from the same recording session, from a different recording session from the same participant, or from a recording session of a different participant.

The procedure allowed us to compare similarities in percept-specific activity within individuals across time (days/years) and between individuals using within-session decoding as a baseline of maximum decoding performance. All decoding analyses were performed on data reconstructed at the source level.

Within-participant Decoding of Conscious Perception

SVM classification accuracies for the three conditions (within-session, between-session, and between-participant decoding) for recording Sessions 1 and 2 can be seen in Figure 4. For Session 1, within-session accuracy was

61.1% (95% CI [60.2%, 62.0%]), between-session accuracy was 59.9% (95% CI [59.2%, 60.6%]), and between-subject accuracy was 53.7% (95% CI [53.0%, 54.4%]). For Session 2, the accuracies were 62.3% (95% CI [(61.2%, 63.3%]), 61.2% (95% CI [60.4%, 62.1%]), and 50.9% (95% CI [50.1%, 51.8%]), respectively. Kruskal–Wallis tests showed no difference in decoding accuracy around the VAN between recording sessions in general [$\chi^2(1) = 0.4, p = .54$] but showed highly significant differences across the three training/testing conditions [$\chi^2(2) = 191, p = .0001$]. Accuracy was 1.10% (95% CI [0.22%, 2.00%], $p = .0149$) lower when testing across sessions compared with within sessions and 8.27% (95% CI [7.46%, 9.08%], $p < .0001$) lower when testing across participants compared


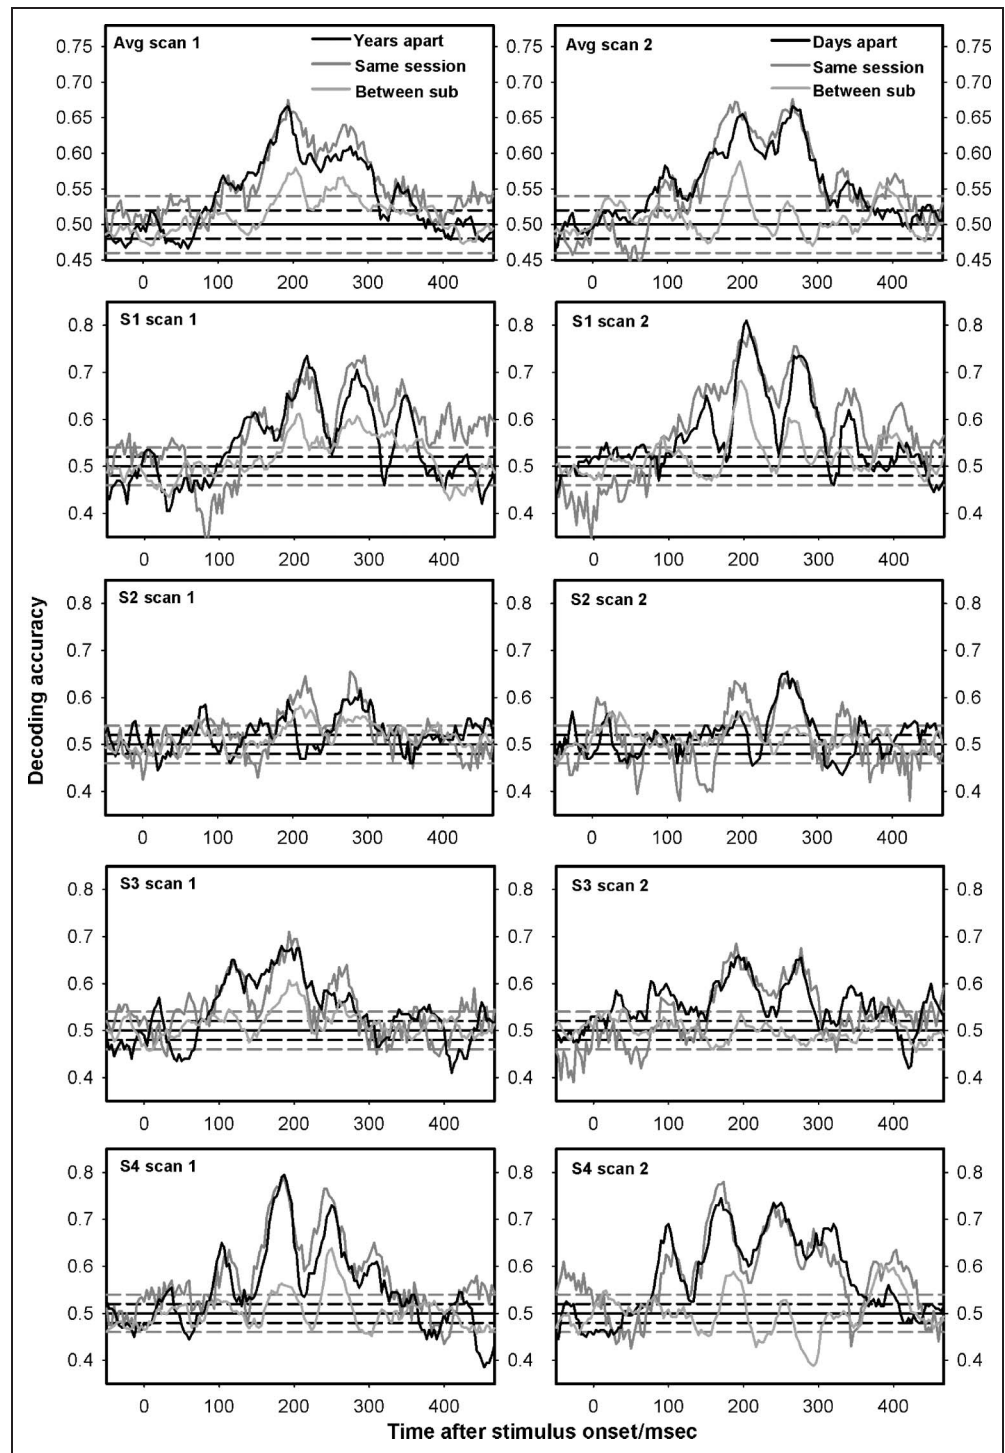


Figure 3. The ERF related to each percept, face (black curve), and grating (gray curve), at the sensor showing the largest difference in Figure 2. Note that the largest difference typically corresponds to suppression of the face-specific components during grating perception. This is a replication of a previous finding (Sandberg et al., 2013).

Figure 4. Average decoding accuracy across participants. SVMs were trained to decode perception. Decoding accuracy is plotted for conditions where the SVM training and testing data came from the same participants and the same recording session (dark gray lines), for conditions where training and testing data came from the same participant but were gathered days (right, black line) or years (left, black line) apart, and finally for conditions where training and testing data came from different participants (light gray lines). Horizontal, solid black lines represent chance. Horizontal, dotted black and gray lines represent the 95% binomial confidence interval around chance (uncorrected) for between-participant and within-participant decoding, respectively.



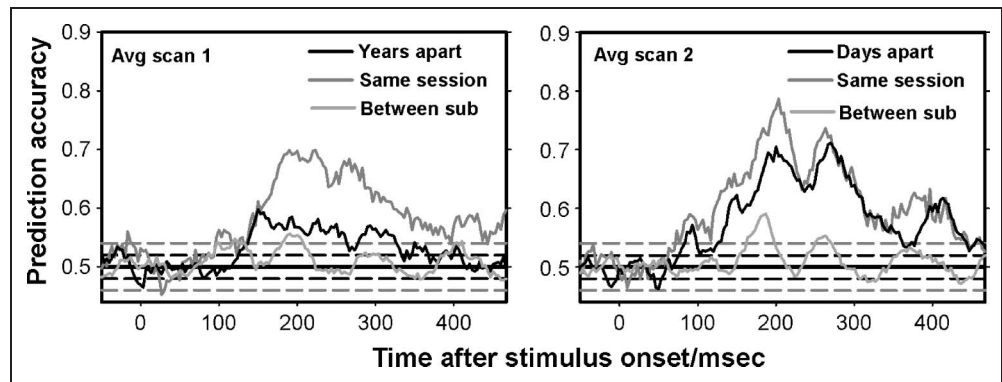
with across sessions. Additionally, decoding accuracy was only 0.14% (95% CI [-0.66%, 0.93%], $p = .72$) lower when testing across years rather than days. In other words, prediction accuracy across 2.5 years was most likely maintained and certainly decreased no more than 1%. Decoding accuracy was found to be above chance for all three conditions in the time interval around the VAN [$t(103) > 7.6$ and $p < .0001$ in all cases].

As mentioned above, all decoding analyses were performed using data reconstructed at the source level. The

advantage of source-based versus sensor-based decoding is presented in Figure 5. In sensor space, a large drop in decoding accuracy was observed across years. As no such drop was not observed in source space, the sensor space difference should not be taken as a sign of difference in neural activity, but rather a trivial difference in head position demonstrating the necessity of source space activity reconstruction.

The VAN is most often found around two ERF components around the 130–320 msec time range (Koivisto &

Figure 5. Average decoding accuracy across subjects when classifiers were trained and tested on sensors data instead of source data. SVMs were trained to decode perception. Decoding accuracy is plotted for conditions where the SVM training and testing data came from the same participants and the same recording session (dark gray lines), for conditions where training and testing data came from the same participant but were gathered days (right, black line) or years (left, black line) apart, and finally for conditions where training and testing data came from different participants (light gray lines). Horizontal, solid black lines represent chance. Horizontal, dotted black and gray lines represent the 95% binomial confidence interval around chance (uncorrected) for between-participant and within-participant decoding, respectively. Note the large drop in classification accuracy across years compared with when analyses were performed in source space (Figure 1). The topography thus differed across years, whereas the sources did not, indicating that head position in relation to the sensors was a large artifact for sensor space decoding. For this reason, analyses in the main text were performed on data reconstructed in source space.




Revonsuo, 2010). For face processing, these two components are the M170 and the P2m (Sandberg et al., 2013). The drop in decoding accuracy for between-participant versus within-participant classification could be caused by changes in the latency, amplitude, as well as the sources of these two components. All these aspects are examined in the following.

Variability in Latency

The latencies of the M170 and the P2m are plotted in Figure 6 (the topographies around the peak are plotted in Figure 2, and the percept-related ERFs are plotted in Figure 3). Kruskal–Wallis tests revealed that the latency of both components varied across participants [M170: $\chi^2(3) = 9.4, p = .025$; P2m: $\chi^2(3) = 8.9, p = .031$], whereas no such difference was found within participants across time [M170: $\chi^2(2) = 0.1, p = .97$; P2m: $\chi^2(2) = 1.0, p = .59$].

Figure 6. Latency of M170 and P2m. The average peak latencies for the M170 (left) and the P2m (right) are plotted for all data sets for all participants. Notice the apparent variability in latency across participants and the apparent stability of latency within participants.




Variability in Amplitude

The maximum amplitude differences between face and grating perception around the M170 and the P2m are plotted in Figure 7. Kruskal–Wallis tests revealed that the amplitude differences of both components varied across participants [M170: $\chi^2(3) = 9.5, p = .024$; P2m: $\chi^2(3) = 7.8, p = .050$], whereas no clear difference was found within participants across time [M170: $\chi^2(2) = 0.8, p = .67$; P2m: $\chi^2(2) = 2.0, p = .37$].

Variability of Sources

MVB models were made individually at the M170 and the P2m for all 35 sources for all recordings. The model evidence is plotted in Figure 8 for the M170 and in Figure 9 for the P2m. Conclusions should not be based on the absolute values of the model evidence, but within a data set, the model evidence provides an estimate of how well the activity at a particular source predicts perception

Figure 7. Amplitude differences of the M170 and P2m between face and grating perception. The average peak amplitude difference for the M170 (left) and the P2m (right) are plotted for all recording sessions for all participants. Notice that the amplitude difference varies more between participants than within participants.



compared with the other sources. For each data set, MVB modeling was thus used to rank the importance of the sources, and these rankings were compared within and between participants using Kruskal–Wallis tests. Again, large differences were found around both components across participants [M170: $\chi^2(3) = 14.0, p = .0029$; P2m: $\chi^2(3) = 20.7, p = .0001$], and no clear, significant differences were found within participants across time [M170: $\chi^2(2) = 3.4, p = .18$; P2m: $\chi^2(2) = 2.3, p = .31$].

DISCUSSION

Between-participant Differences but No Across-time Differences

We found that conscious perception could be decoded around the time window of the VAN (140–310 msec after stimulus onset). Decoding accuracy decreased little or not at all across years compared with across days (<1%), implying that the VAN-related neural correlates of conscious perception were very stable across time within individuals. Follow-up analyses revealed that the consciousness-specific modulation of the two main ERF components of the experiment, the M170 and the P2m, remained stable over time in terms of latency, amplitude, and sources.

In contrast, decoding accuracy dropped drastically when training and testing classifiers across individuals as was previously reported in Sandberg et al. (2013). Note that even with only four of the original eight participants, statistical testing was consistently significant. In this study, further analyses revealed significant differences in the consciousness-specific modulation of the M170 and the P2m in terms of latency, amplitude, and sources. Crucially, the between-participant drop in decoding accuracy cannot be caused by (a) technical artifacts related to classifying across different recording sessions (as within-participant decoding did not show this drop) and (b) artifacts of the

source activity reconstruction process (as differences were found between individuals in sensor as well as source space). It is also unlikely that the results are caused artifacts related to (consistent) eye movements as these would be localized to the inferior part of the frontal cortex and previous analyses of Recording Session 1 (Sandberg et al., 2013) demonstrated that peak classification accuracy can be obtained when classification is based on extrastriate visual cortex activity alone. Similarly, occipital and temporal lobe sources were consistently found to be the most predictive of conscious perception in that study (Sandberg et al., 2013).

The results are also compatible with the previously mentioned study showing much higher intraindividual stability of VEPs for simple visual stimuli across months compared with interindividual stability (Sarnthein et al., 2009). We note again, however, that here we did not simply demonstrate that the ERPs were stable, but rather that the consciousness-specific modulation of these components (given equivalent physical stimulation) was comparable within participants across years. In other words, the endogenous suppression of the nonperceived image during BR remains highly comparable across years in the normal, healthy human brain.

At least two different interpretations appear valid for the finding of interindividual differences in the neural correlates of conscious contents, and the study thus opens up the possibility for further studies into this domain. The results could be taken to indicate that conscious contents are represented differently across individuals (i.e., the weighted contribution of cortical sources may differ between individuals having the same conscious experience), or alternatively that different individuals have slightly different perceptual experiences of the same object (for instance, it might be expected that individual differences in the degree of endogenous suppression of irrelevant information could lead to differences in perceived visual clarity). These interpretations could be examined in further

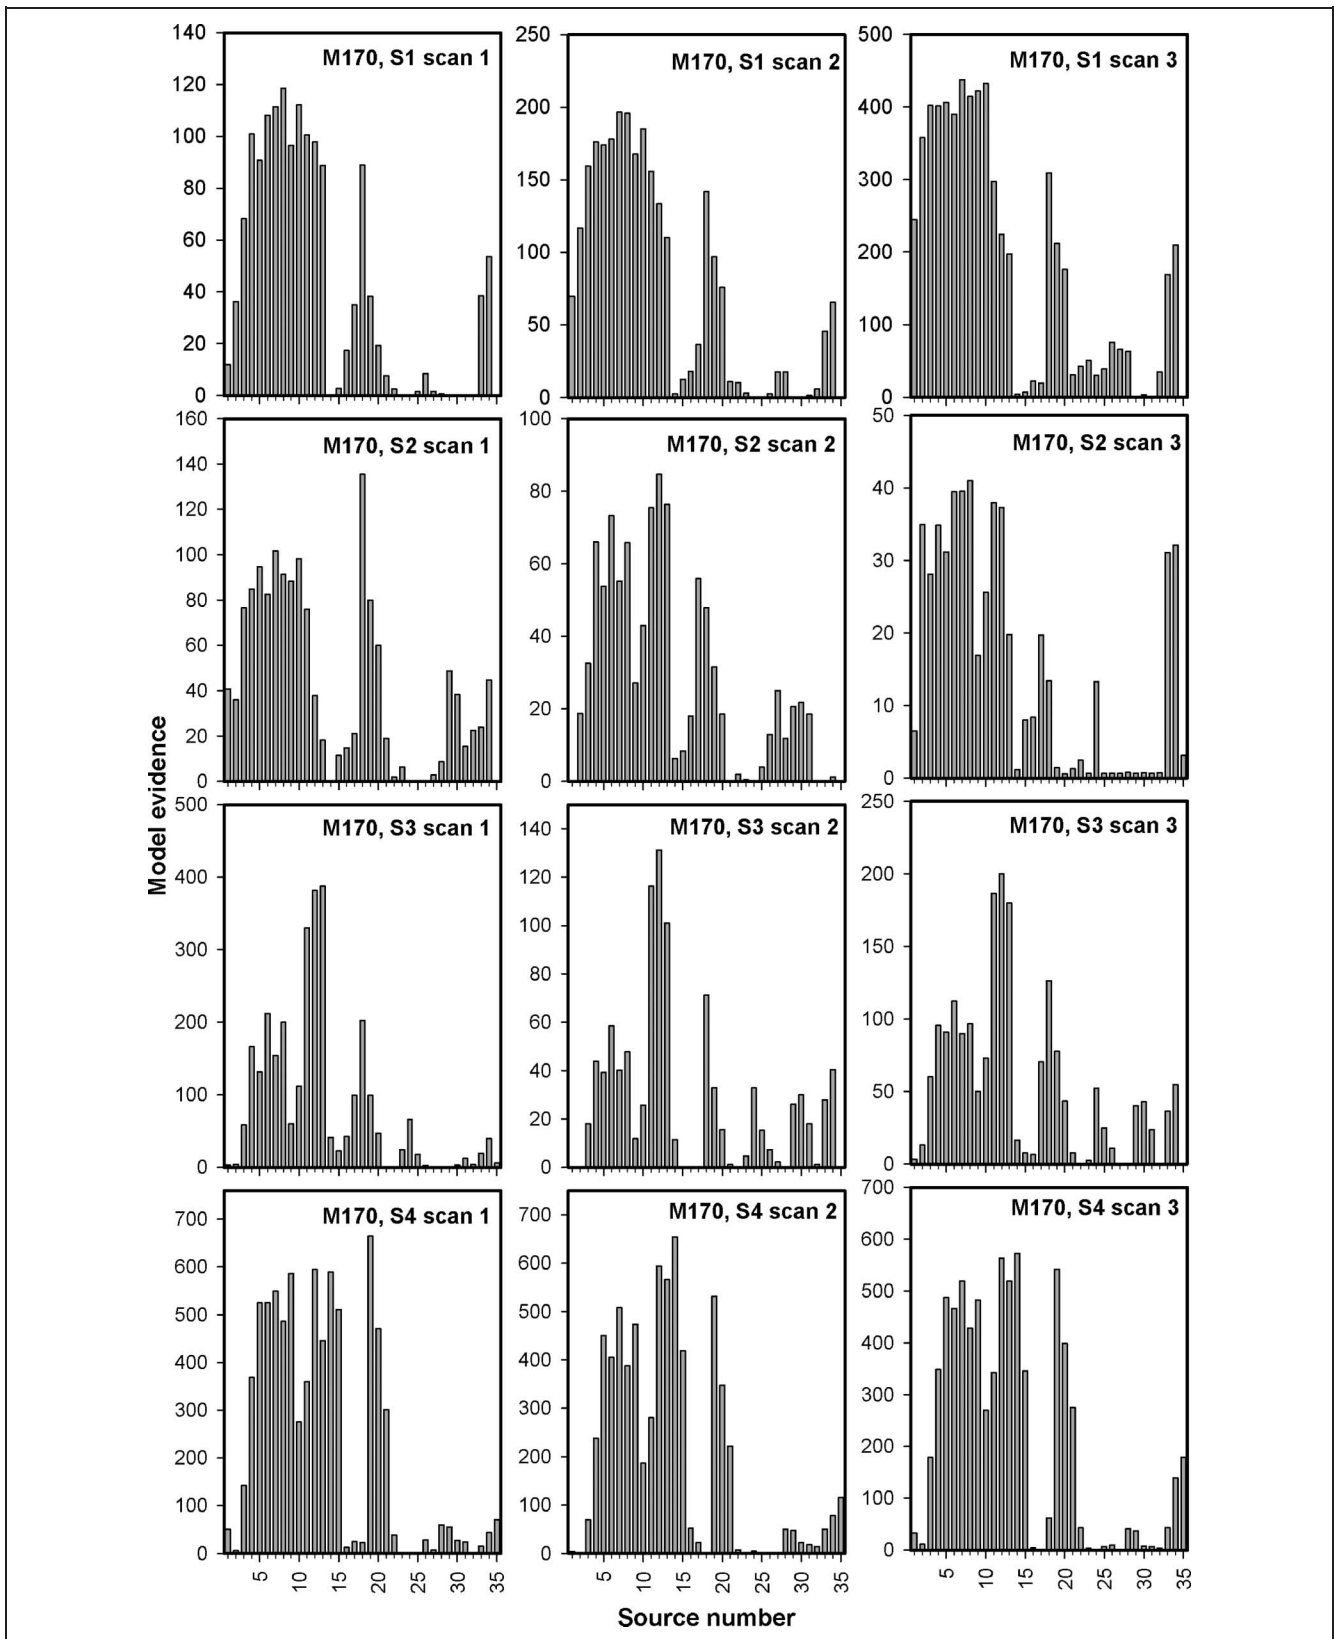


Figure 8. MVB model evidence across sources for the M170. MVB model evidence is plotted for the M170 for all sessions for all participants. Notice the larger similarity in the pattern of source evidence across data sets from the same participant (horizontal direction) than between participants (vertical direction).

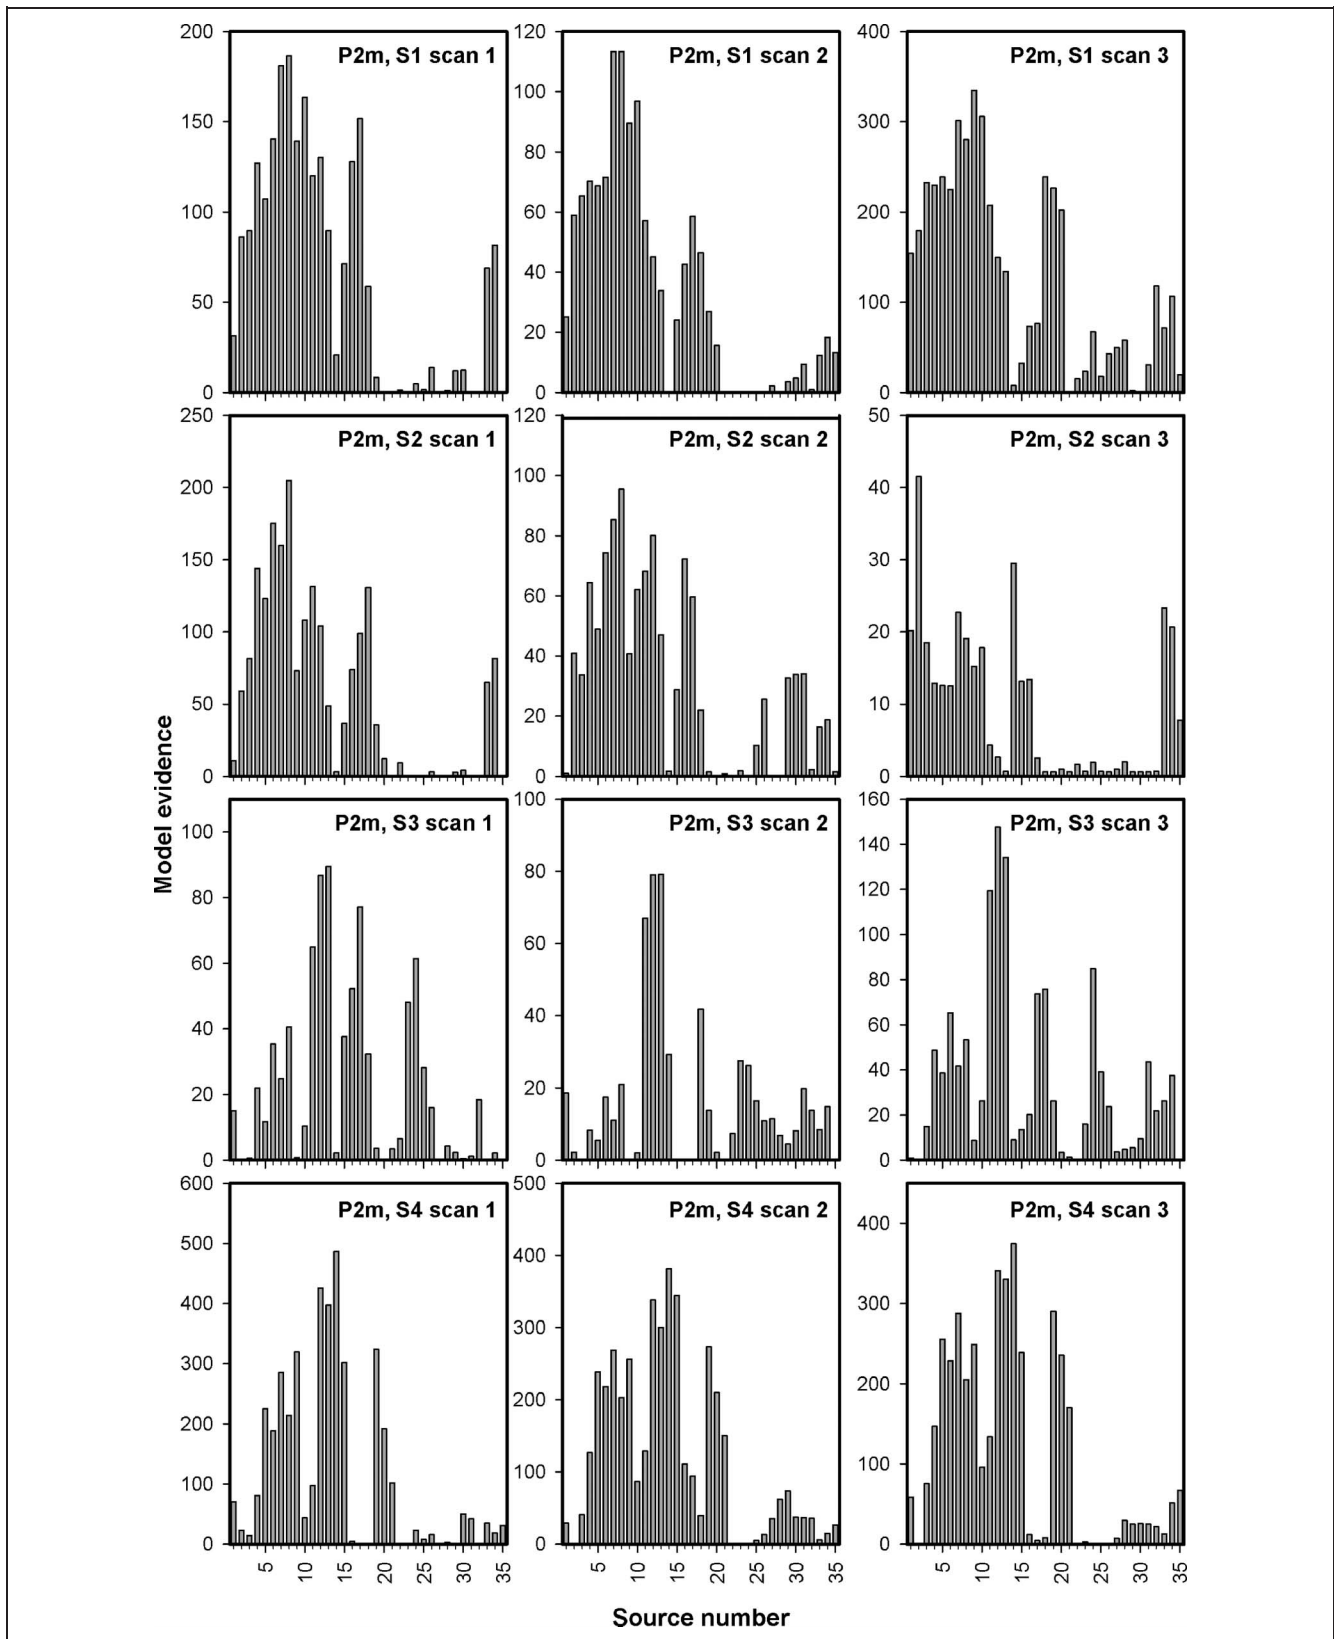


Figure 9. MVB model evidence across sources for the P2m. MVB model evidence is plotted for the P2m for all sessions for all participants. Notice the larger similarity in the pattern of source evidence across data sets from the same participant (horizontal direction) than between participants (vertical direction).

studies by correlating individual differences in detailed perceptual reports to interindividual decoding accuracy.

For both this study and previous studies of interindividual EEG/MEG differences, the anatomical structure of the participants is a potential confound. In particular, individual differences in anatomy (e.g., cortical folding) means that the same cortical area may have different orientation and distance from the sensors over subjects. It is thus likely that at least part of the differences in the MEG signal across individuals is because of this individual mapping from sources to sensors. It is difficult to assess the impact of this confound, and it is likely to have a greater effect on the estimate of the source level amplitudes than latencies. One way to partially address this potential issue, in future studies, is to use combined MEG–EEG to sample from an increased measurement subspace (Sharon, Hämäläinen, Tootell, Halgren, & Belliveau, 2007).

Source Space versus Sensor Space

Parts of the results were based on decoding using classification algorithms on source space data, and parts were based on conventional statistics in sensor space. The combination of both methods was used to obtain converging evidence and reject some potential criticisms. The most important reason for basing the analyses on source space data was to avoid confounding the results by larger differences in head position in relation to the sensors across years than days. Indeed, our sensor space compared with source space decoding revealed that this confound was present. Within-participant topographical map differences across years were thus entirely caused by the trivial finding that participants positioned their head very similarly on sessions close in time but not on sessions years apart. This confound, however, could not impact upon the estimation of the M170 and P2m component latencies, and for this reason, this analysis was performed in sensor space to include data that had not been modified by source space reconstruction. A slight impact of head position might be expected upon the difference in activity between the component amplitude for face and grating perception at the peak latencies as the signal-to-noise ratio is influenced by the distance between the source and the sensor, which of course varies with head position. To minimize the impact of this confound, the difference was measured at the sensor showing the largest difference, although this might not be the same sensor for every data set.

Decoding versus Traditional ERF Measures

Decoding accuracy was used to obtain a meaningful objective measure of quantitative differences between data sets—that is, a measure that would indicate whether the information considered important to predict the conscious experience of one individual at a given time is also the information that predicts the conscious of that individual

at a different time or the conscious experience of a different individual. By allowing the multivariate classifier to select a weighted combination of sources, the choice of relevant information was as close to optimal as possible and could not be biased by subjective choices regarding which sensors/sources and time points to compare.

Conclusion

Taken together, the experiment demonstrates that the neural correlates of conscious perception generalize very well across years within individuals, but not across individuals. The large stability of the latency, amplitude, and the relative importance of a large number of sources argue against spontaneous changes in how specific conscious contents are represented (on the general scale that MEG examines). This study thus indicates that, once a brain has found a way to process and consciously represent an object, it will continue to use this even across years (presumably until some drastic event such as intensive training or neural damage forces a change).

Acknowledgments

This work was supported by the Wellcome Trust (G. R., G. R. B., and K. S.) and the European Research Council (K. S. and M. O.).

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: krissand@rm.dk.

REFERENCES

- Carter, O., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PLoS One*, *2*, e343.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, *10*, 204–211.
- Dehaene, S., Kerszberg, M., & Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences, U.S.A.*, *95*, 14529–14534.
- Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of Cognitive Neuroscience*, *19*, 1488–1497.
- Friston, K. J. (2012). Ten ironic rules for non-statistical reviewers. *Neuroimage*, *61*, 1300–1310.
- Friston, K. J., Chu, C., Mourão-Miranda, J., Hulme, O., Rees, G., Penny, W., et al. (2008). Bayesian decoding of brain images. *Neuroimage*, *39*, 181–205.
- 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.
- Hall, R. A., Rappaport, M., Hopkins, H. K., & Griffin, R. B. (1973). Peak identification in visual evoked potentials. *Psychophysiology*, *10*, 52–60.
- Koivisto, M., & Revonsuo, A. (2003). An ERP study of change detection, change blindness, and visual awareness. *Psychophysiology*, *40*, 423–429.

- 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.
- Kornmeier, J., & Bach, M. (2005). The Necker cube—An ambiguous figure disambiguated in early visual processing. *Vision Research*, *45*, 955–960.
- Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, *1*, 204–220.
- Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience*, *5*, 605–609.
- Liu, Y., Paradis, A.-L., Yahia-Cherif, L., & Tallon-Baudry, C. (2012). Activity in the lateral occipital cortex between 200 and 300 msec distinguishes between physically identical seen and unseen stimuli. *Frontiers in Human Neuroscience*, *6*, article 211.
- Näpflin, M., Wildi, M., & Sarnthein, J. (2007). Test-retest reliability of resting EEG spectra validates a statistical signature of persons. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, *118*, 2519–2524.
- Näpflin, M., Wildi, M., & Sarnthein, J. (2008). Test-retest reliability of EEG spectra during a working memory task. *Neuroimage*, *43*, 687–693.
- Oken, B. S., Chiappa, K. H., & Gill, E. (1987). Normal temporal variability of the P100. *Electroencephalography and Clinical Neurophysiology*, *68*, 153–156.
- Sandberg, K., Bahrami, B., Kanai, R., Barnes, G. R., Overgaard, M., & Rees, G. (2013). Early visual responses predict conscious face perception within and between subjects during binocular rivalry. *Journal of Cognitive Neuroscience*, *25*, 969–985.
- 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.
- Schellberg, D., Gasser, T., & Köhler, W. (1987). The intraindividual reproducibility of flash-evoked potentials in a sample of children. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, *5*, 135–143.
- 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.
- Sergent, C., & Dehaene, S. (2004). Neural processes underlying conscious perception: Experimental findings and a global neuronal workspace framework. *Journal of Physiology, Paris*, *98*, 374–384.
- Sharon, D., Hämäläinen, M. S., Tootell, R. B. H., Halgren, E., & Belliveau, J. W. (2007). The advantage of combining MEG and EEG: Comparison to fMRI in focally stimulated visual cortex. *Neuroimage*, *36*, 1225–1235.
- Van Loon, A. M., Scholte, H. S., van Gaal, S., van der Hoort, B. J. J., & Lamme, V. A. F. (2012). GABA(A) agonist reduces visual awareness: A masking-EEG experiment. *Journal of Cognitive Neuroscience*, *24*, 965–974.
- Vanni, S., Revonsuo, A., Saarinen, J., & Hari, R. (1996). Visual awareness of objects correlates with activity of right occipital cortex. *NeuroReport*, *8*, 183–186.
- Weinberg, A., & Hajcak, G. (2011). Longer term test-retest reliability of error-related brain activity: Long-term reliability of the ERN. *Psychophysiology*, *48*, 1420–1425.