NeuroImage 128 (2016) 385–397

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

# Representational dynamics of object recognition: Feedforward and feedback information flows

Erin Goddard [a,b,*], Thomas A. Carlson [a,b], Nadene Dermody [a], Alexandra Woolgar [a,b]

[a] Perception in Action Research Centre (PARC) and Department of Cognitive Science, Macquarie University, Sydney, NSW 2109, Australia
[b] ARC Centre of Excellence in Cognition and its Disorders (CCD), Macquarie University, Sydney, NSW 2109, Australia

ARTICLE INFO

*Article history:*
Received 18 June 2015
Accepted 5 January 2016
Available online 13 January 2016

*Keywords:*
Magnetoencephalography
Pattern classifier analysis
Feedback
Coarse-to-fine
Granger causality
Visual perception

ABSTRACT

Object perception involves a range of visual and cognitive processes, and is known to include both a feedfoward flow of information from early visual cortical areas to higher cortical areas, along with feedback from areas such as prefrontal cortex. Previous studies have found that low and high spatial frequency information regarding object identity may be processed over different timescales. Here we used the high temporal resolution of magnetoencephalography (MEG) combined with multivariate pattern analysis to measure information specifically related to object identity in peri-frontal and peri-occipital areas. Using stimuli closely matched in their low-level visual content, we found that activity in peri-occipital cortex could be used to decode object identity from ~80 ms post stimulus onset, and activity in peri-frontal cortex could also be used to decode object identity from a later time (~265 ms post stimulus onset). Low spatial frequency information related to object identity was present in the MEG signal at an earlier time than high spatial frequency information for peri-occipital cortex, but not for peri-frontal cortex. We additionally used Granger causality analysis to compare feedforward and feedback influences on representational content, and found evidence of both an early feedfoward flow and later feedback flow of information related to object identity. We discuss our findings in relation to existing theories of object processing and propose how the methods we use here could be used to address further questions of the neural substrates underlying object perception.



Visual object recognition and identification is an important task in everyday life, and the speed and accuracy with which we can identify objects are consistent with the visual system devoting considerable resources to this ecologically relevant process. Object perception involves a range of visual and cognitive processes, including a feedfoward flow of information along the 'ventral stream' of visual cortex (for example, see Tanaka, 1996; Grill-Spector et al., 2001), and also feedback from frontal and parietal areas such as prefrontal cortex. However, the way in which these areas interact to contribute to object perception remains unclear despite a growing experimental literature on the topic.

A range of theories of object perception (Bullier, 2001; Bar, 2003; Peyrin et al., 2010; Tapia and Breitmeyer, 2011; Horr et al., 2014; Hochstein and Ahissar, 2002) have hypothesized that very early top-down feedback, around 100–150 ms after stimulus onset (Bar et al., 2006), carries content regarding object identity from prefrontal cortex to the traditional 'bottom-up' dorsal and ventral visual pathways. These theories are based on results such as the speed with which humans can correctly respond to simple object categorisation tasks (for example, 120 ms after stimulus onset, Kirchner and Thorpe, 2006), combined with reports that early activity (130–150 ms after stimulus onset) from prefrontal sites can vary with object recognition (Thorpe et al., 1996; Bar et al., 2006).

The relative timing of feedforward and feedback flows of information in object perception has been suggested to depend on the spatial frequency content of the image, with the earliest information about object information coming from low spatial frequency (low-pass) image components (Bar et al., 2006; Chaumon et al., 2014; Fintzi and Mahon, 2014). This is broadly consistent with psychophysical results implying that low spatial frequency image components are processed prior to high spatial frequency components (Hughes et al., 1996; Parker et al., 1992, 1997; Schyns and Oliva, 1994; Neri, 2011). However, these effects are likely contingent on the relative usefulness of low and high spatial frequencies to the participant's task (De Gardelle and Kouider, 2010; Stein et al., 2014; Patai et al., 2013), which challenges the notion of clear segregation between rapidly propagated low-pass signals and slower 'High-pass' signals. Furthermore, neuroimaging studies providing evidence for differential processing of high and low spatial frequency stimuli in object processing (Bar et al., 2006; Chaumon et al., 2014; Fintzi and Mahon, 2014) have used stimuli varying in total spatial frequency content, and have not equated the stimulus types for low-level

* Corresponding author at: Department of Cognitive Science, Macquarie University, Sydney, NSW 2109, Australia.
*E-mail address:* erin.goddard@mq.edu.au (E. Goddard).

http://dx.doi.org/10.1016/j.neuroimage.2016.01.006
1053-8119/

properties such as overall luminance and contrast. This leaves open the possibility that their differential effects may be due to different image statistics across conditions.

Here we devised a new methodology to explore the timing of feedforward and feedback flows of object-related information based on magnetoencephalography (MEG) recordings. We applied a time resolved multivariate pattern classification analysis to magnetoencephalography (MEG) data (Carlson et al., 2011, 2013; Isik et al., 2014) and compared the object-related information in peri-occipital and peri-frontal areas at different time points. Using a novel extension of Granger causality analysis, we tested for evidence that the representational structure of object-related information in frontal regions predicted the representational structure of later responses in occipital areas.

Using this method, we compared the brain's processing of low and high spatial frequency object-related information. We used stimuli that had the same power at each spatial frequency in their Fourier amplitude spectra, and the same overall contrast, varying only the 'diagnostic' spatial frequencies (De Gardelle and Kouider, 2010). Stimuli in our 'Low-pass' condition had object signal in the low spatial frequencies while high spatial frequencies were phase-randomized, while for the 'High-pass' condition this was reversed.

## Methods

### *Participants*

Twelve participants (nine female, ten naïve to the purposes of the study) took part in an initial psychophysical experiment used to calibrate the visual images for their low and high spatial frequency content. Nine participants (five female, eight naïve to the purposes of the study) completed the second psychophysical experiment and the MEG experiment. All had normal or corrected to normal vision, and naïve participants were paid for their time. All participant recruitment and experiments were conducted with the approval of the Macquarie University Human Research Ethics Committee.

### *Visual stimuli*

Visual stimuli were generated and presented using Matlab (version R2013a) and routines from Psychtoolbox (Brainard, 1997; Pelli, 1997). In all experiments we used the same set of 24 images, which were selected from a set of 92 supplied by Nikolaus Kriegeskorte (described in Kriegeskorte et al., 2008). All images were segmented real world objects on a gray background. In both psychophysical and imaging experiments, participants judged whether each presented object was smaller or larger than a shoebox. We chose the 24 images such that for both the 'smaller than a shoebox' and 'larger than a shoebox' groups there were six animate and six inanimate objects.

Each image was 175 × 175 pixels, and subtended 15 degrees visual angle (dva) in each experiment. We converted each original color image to grayscale by setting the RGB coordinate of each pixel to the average of the R, G, and B coordinates of that pixel in the original image. In order to equate all images for their power at each orientation and spatial frequency, we set the amplitude matrix of each image to the average amplitude matrix of all images. To find the average amplitude matrix across the 24 images, we performed a two-dimensional discrete Fourier transform of each image, which yielded an amplitude and phase matrix for each image, and then for each point in amplitude matrix we used the average amplitude across the 24 images. We used the same amplitude matrix for every stimulus image, varying only the phase matrix that was used in the two-dimensional inverse discrete Fourier transform to generate a given stimulus.

The phase matrix of each stimulus was derived from one of the 24 images, with varying amounts of phase randomization introduced to the phase matrix. The four stimulus conditions, along with the pattern of phase randomization in each case, are illustrated in Fig. 1. Phase randomization was introduced to one or more of three spatial frequency bands: low (< 0.90 cycles/dva), medium (≥ 0.90 cycles/dva and ≤ 1.03 cycles/dva) and high (> 1.03 cycles/dva). In the 'Low-pass' condition, the phase of spatial frequencies in the high and medium bands was randomized, and in the 'High-pass' condition the low and medium bands were randomized. In the 'Strong signal' and 'Weak signal' stimulus conditions, the phase of all spatial frequencies in the medium band was randomized, along with varying proportions of the frequencies in the low and high bands, as detailed below.

Since every image contained at least some phase randomization, we were able to repeat the randomization process and generate different versions of the same image with the same object signal, such that a new image was used for every trial. Also, since all images had the same amplitude matrix, the objects could not be distinguished when phase randomization was complete across all spatial frequencies. This ensured that objects could not be identified based on the orientation/spatial frequency profile of the randomized images.

### *Psychophysical experiments*

We conducted two psychophysical experiments in order to measure the detectability of the objects in the different conditions. Stimuli were generated and displayed on a Dell OptiPlex 9010 desktop computer driving an AMD Radeon HD 7570 graphics card to draw stimuli to a 60 × 33 cm Samsung SyncMaster SA950 Full HD 3D LED monitor, refreshed at 120 Hz. Experiments took place in a darkened room and the monitor was viewed from a distance of .64 m.

In the first experiment, we included only the 'Low-pass' and 'High-pass' conditions, and measured the detectability of each object in these two conditions as a function of the amount of phase randomization. At the start of each session the participant chose the keys on a keyboard they would use for their responses ('smaller' and 'larger' than a shoebox) and after these responses the experiment commenced. Each trial began with a central fixation marker (a small gray cross) that was displayed on a black background for 250 ms, after which the stimulus image was displayed on a black background for 500 ms before being replaced by the fixation marker. Participants were given an unlimited amount of time to respond. Following the participant's response, they would receive feedback on their decision (displayed as 'correct' in green, or 'incorrect' in red) for 500 ms, and then the next trial would commence.

Each of the 12 participants completed 8 sessions of 15–20 minutes each, consisting either of only 'High-pass' or only 'Low-pass' stimuli, and including either the first 12 or the second 12 images in the set. The order of sessions was counterbalanced across participants. Each session included 12 randomly interleaved adaptive psychophysical staircases (one for each of the 12 images) (Kontsevich and Tyler, 1999) consisting of 30 trials each. The adaptive staircase set the degree of phase randomization on each trial in order to reliably estimate the detection threshold of each image (the point at which the participant was 75% correct in identifying whether the object was smaller or larger than a shoebox). At the completion of the 8 sessions, we had two estimates of detection threshold of each image in both the 'Low-pass' and 'High-pass' conditions. Results and stimuli from the first psychophysics experiment are included in the Supplementary Material.

The average detection thresholds across participants were used to generate low and high-pass versions of each object that were of approximately equal detectability. We found the maximum possible signal for which the signal in the low and high-pass images were equal multiples of the average detection threshold, and used these maximum matched signal values to define the 'Low-pass' condition and the 'High-pass' condition of the MEG experiment. Images in the 'Strong signal' condition were defined by setting the signal in both the low and high spatial frequency bands to these maximum matched values.

Finally, in the second psychophysical experiment, we calibrated the signal level in the 'Weak signal' condition individually for each of the nine participants who went on to complete the MEG experiment.

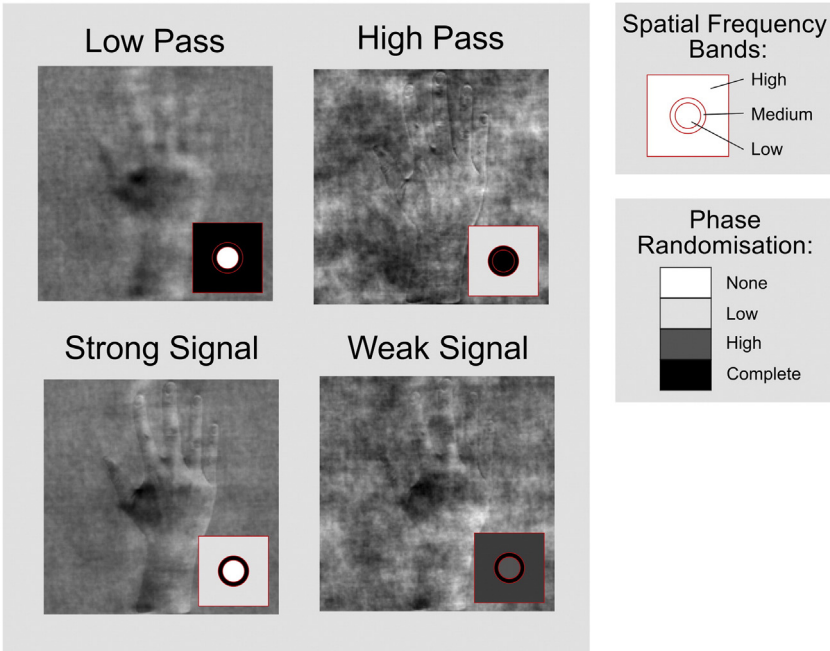


**Fig. 1.** An example stimulus image (of an open hand) in each of the four conditions. For each image condition the example stimulus is accompanied by an inset (not part of the stimulus) which schematically illustrates the phase randomization that was applied to the Fourier amplitude spectra of the original image. Based on the results of an initial psychophysical experiment, the degree of phase randomization in the high and low spatial frequency bands was matched so that images in the 'High-pass' and 'Low-pass' conditions were equally identifiable. For 23 of 24 images, the detection threshold for the 'High-pass' image was lower than for the 'Low-pass' image, so for each of these images the 'High-pass' condition had a low amount of phase randomization in the high spatial frequency band in order to equate their detectability with the 'Low-pass' condition. For the single image where the detection threshold for the 'Low-pass' image was lower than for the 'High-pass' image, this was reversed.

For each object, the ratio of signal in the low and high spatial frequency bands was fixed to the values determined in first experiment, and the overall signal in a given image was varied between 0 (totally randomized) and maximum (the 'Strong signal' condition). As before, we used interleaved adaptive psychophysical staircases, in this case estimating the detectability of each object as a function of the combined low and high spatial frequency signal amplitude. For each observer and each object the psychometric function that best fit these data was used to determine the signal level corresponding to 85% correct, and these images comprised the 'Weak signal' condition in the MEG experiment.

*MEG experimental design*

MEG data were collected with a whole-head MEG system (Model PQ1160R-N2; KIT, Kanazawa, Japan) consisting of 160 coaxial first-order gradiometers with a 50 mm baseline (Kado et al., 1999; Uehara et al., 2003). Prior to MEG measurements, five marker coils were placed on the participant's head and their positions and the participant's head shape were measured with a pen digitizer (Polhemus Fastrack, Colchester, VT, USA).

Each participant's MEG data were collected in a single session of approximately 90 minutes, and raw MEG data were recorded at 1000 Hz. While participants were lying in the scanner, stimuli were projected through a customized window by an InFocus IN5108 LCD back projection system (InFocus, Portland, Oregon, USA) located outside the Faraday shield, onto a screen above the participant that was viewed from a distance of 113 cm. On each trial, the timing was the same as that for the psychophysical experiments except for the inclusion of variable inter-trial interval. Each trial commenced with 250 ms of a fixation marker, after which the stimulus was presented for 500 ms then replaced with the fixation marker. Participants indicated their response using a Fiber Optic Response Pad (fORP, Current Designs, Philadelphia, PA, USA) with the Bimanual 4-Button Fiber Optic Response Pad. Participants were then provided with feedback and then the next trial commenced after an inter-trial interval that varied randomly between 900 ms and 1200 ms to avoid expectancy effects.

The MEG session was divided into six blocks, and the response buttons corresponding to 'smaller' and 'larger' than a shoebox were alternated at the beginning of each block. Each of the 24 objects was presented a total of 96 times each, 24 times in each of the four stimulus conditions ('Low-pass', 'High-pass', 'Strong signal' and 'Weak signal'), giving a total of 2304 trials per session. The order of presentation of each of these 2304 trials was balanced so that each object was equally likely to be preceded by and followed by every object, and each of the four conditions was equally likely to be preceded by and followed by every condition. This counterbalanced order was split into six blocks of 384 trials each, to which the last trial of the preceding block was added to the beginning, and the first trial of the next block was added to the end. Data from these first and last trials were discarded.

For the MEG experiments, the average response accuracy was consistent with the object related task being easiest in the 'Strong signal' condition (97.65% correct), hardest in the 'Weak signal' condition (81.81%), and intermediate in the 'Low-pass' (96.14%) and 'High-pass' (89.22%) conditions. A one-way analysis of variance (ANOVA) revealed a significant main effect of condition on accuracy ($F_{(3,32)} = 16.07$, $p < 0.01$). Post-hoc comparisons with Dunn-Sidak adjustment for multiple comparisons showed that accuracy was significantly lower in the 'Weak signal' condition than in any other condition ($p < 0.05$), and significantly higher in the 'Strong signal' condition than in the 'High-pass' condition ($p < 0.05$). Accuracy was higher in the 'Low-pass' condition than in the 'High-pass' condition, but this difference was not significant ($p > 0.05$). Reaction times followed the same trend as the accuracy data, with slightly shorter median reaction times when accuracy was higher ('Strong signal': 1.71 s; 'Weak signal': 1.79 s; 'Low-pass': 1.73 s; 'High-pass': 1.77 s) but a one-way ANOVA on the logarithm of the reaction times did not reveal a significant effect of condition of reaction time ($F_{(3,32)} = 1.26$, $p = 0.31$).

*MEG data analysis: Data reduction*

Raw MEG data were imported into Matlab (version R2013a) and preprocessed with notch filtering (to remove 50 Hz modulation). Data corresponding to each trial were extracted by taking the data from −100 ms to +1200 ms relative to stimulus onset. In order to compare the MEG signal in frontal and peri-occipital sensors, for each participant we created three data sets, the first included data from all 160 sensors, and the second and third data sets included data from the frontal and peri-occipital sensors. Using the measured location of each sensor, we found the most anterior and the most posterior sensor for each participant and the distance between them ($d$). We then defined the frontal sensors as those within $d/4$ of the most frontal sensor, and the peri-occipital sensors as those within $d/4$ of the most posterior sensor.

All subsequent analyses were carried out for each of the three data sets described above. Firstly, each data set, comprising 1300 ms of data for each of 2304 trials and up to 160 sensors, was reduced using principal components analysis. Data from the first n components that accounted for 99% of the variance were retained; data from remaining components were discarded. Following this data reduction, the 1000 Hz data were downsampled to 200 Hz using the Matlab 'decimate' function.

*MEG data analysis: Classifier analyses*

In order to measure the similarity/dissimilarity of brain responses to the different objects in different stimulus conditions, we used linear discriminant analysis (LDA) to examine the extent to which brain activity could be used to predict the stimulus on any given trial. We included only those trials where participants correctly identified the object as smaller or larger than a shoebox. For each possible pair of the 24 objects we trained a classifier to discriminate between the two objects, then tested the classifier on separate data. Initially, we included data from all four conditions, and tested the classification accuracy using 10-fold cross-validation. The classification rule was learnt using 90% of trials, and then the accuracy of this rule was tested on the remaining 10% of trials. This process was repeated for each of 10 partitions of the data, such that all data were included in the test set once, and no data were ever used in both the training and test set (leave-one-out train-and-test). In order to measure how classification accuracy evolved over time we repeated this process at each time point in the 1300 ms window and at each time point we averaged classifier accuracy across all possible pairs of the 24 objects. The leave-one-out train-and-test classification analysis was then repeated for each stimulus condition separately, resulting in a separate time course of classifier accuracy for each of the 'Strong signal', 'Weak signal', 'Low-pass' and 'High-pass' conditions.

We also tested cross-condition classification accuracy by training the classifier on data from one stimulus condition (e.g. 'Strong signal') and testing on data from every other condition. We performed this analysis to test the similarity of the information that was used by the classifier in each condition. Since by definition the training and test data did not overlap, we used all training data to learn a single classification rule that was applied to data from each trial of the test data in the remaining conditions.

In addition to training and testing classifiers to discriminate each pair of objects (object identity), we also grouped objects according to different category classes. We divided the objects according to animacy (animate versus inanimate) and size (smaller versus larger than a shoebox), collapsing across stimulus conditions, and in both cases trained and tested classification accuracy using 10-fold cross-validation, as described above. In both cases (animacy and size), there were equal numbers of objects on either side of the category boundary, and these two category boundaries were orthogonal. We also devised a third, pseudo-random 'category' boundary, which divided the set of 24 objects into two groups with equal numbers of animate and inanimate objects, and objects that are smaller and larger than a shoebox. The classifier performance based on this pseudo-random category provided a baseline against which to compare the classifiers based on meaningful object categories. At each time point, we performed a paired $t$-test comparing the between-subject classifier performance for the meaningful category (e.g. animacy) with the pseudo-random category. In both cases we used a false discovery rate (FDR) correction to control for multiple comparisons across time points (Genovese et al., 2002).

For each classification performance time course where we report the onset of above chance classifier performance, we tested whether the between-subject mean ($n = 9$) was above chance using a one-tailed $t$-test, and applied an FDR correction to control for multiple comparisons across time points. We defined the onset of above chance classification performance as the start of the first 10 consecutive time points (50 *ms* total) for which classifier performance was above chance ($p < 0.05$). To estimate the 95% confidence intervals of these onsets, we created 1000 bootstraps of the data, each time randomly sampling the individual subject means (with replacement) a total of nine times (corresponding to our nine subjects). After applying the same statistical testing to estimate the onset of above-chance classifier performance for each bootstrapped sample, we defined the 95% confidence intervals as the lower and upper limits of the central 95% of these onset estimates.

*MEG data analysis: Classifier accuracy curve fitting*

In order to compare the time course of classier accuracy across location (peri-frontal versus peri-occipital), participant, and stimulus condition, we reduced the time course of classifier accuracy in each case to the accuracy at the 'early' and 'late' peaks, by fitting a simple curve to the data. For each plot of classifier accuracy over time we used the Matlab '*nlinfit*' function to find the best-fitting model ($\mathbf{M}(t)$) of the curve.

$$\mathbf{M}(t) = a[g(t, \alpha_1, \beta_1) + b(g(t + \alpha_1, \alpha_2, \beta_2))] \quad (1)$$

The model fit as a function of time ($\mathbf{M}(t)$) is a weighted sum of two gamma probability density functions ($g(t)$), which are defined as

$$g(t, \alpha, \beta) = \left[\beta^{\alpha} t^{(\alpha-1)} e^{-\beta t}\right] / \Gamma(\alpha) \text{ where } \alpha > 0 \text{ and } \beta > 0 \quad (2)$$

the free parameters are $a$ and $b$ (scalars), and the offset ($\alpha_1$ and $\alpha_2$) and dispersion ($\beta_1$ and $\beta_2$) of the peaks of the two gamma probability density functions.

Using the parameters of $\mathbf{M}(t)$ that best fit the average classifier accuracy based on the entire data set (shown for an example participant in Fig. 3A) we constrained all other model fits for that participant so that each of the offset parameters ($\alpha_1$ and $\alpha_2$) was within 10 ms of the fit based on the entire data set. This ensured that we could sensibly reduce the data to the location and amplitude of their 'early' and 'late' peaks, and compare across conditions.

Since the fitted curve is always greater than zero, where classification performance is at chance this curve-fitting method will identify local positive peaks in the noisy time course of classification performance. To find a baseline based on chance performance against which to compare the measured peaks, for each participant in each condition we took the time course of classifier accuracy from before stimulus onset (−100 to 0 ms), fit the same model ($\mathbf{M}(t)$) to this noise data, and took the maximum peak of this fitted model. Across participants, the average baseline estimate was d' = 0.045, with standard error 0.003. In results Figs. 4 and 5, the dotted horizontal lines indicate this baseline.

## MEG data analysis: Granger analysis of feedforward/feedback information flows

We summarized the classifier performance for each pair of images into dissimilarity matrices (DSMs) for each time point, for both peri-

frontal and peri-occipital sensors. Each DSM was a 24 × 24 matrix, where each cell in the DSM was defined as the classification accuracy for a single pair of images. The diagonal axis of these matrices was nominally zeros, and the matrix is by definition symmetric about the diagonal axis, so for all correlation values calculated below we included only the triangular part of the matrix above the diagonal.

These matrices capture more about the representation of object information than the average classification accuracy, and similar analyses have been applied previously to fMRI and electrophysiology data to compare representations of object information across brain areas, and across species (Kriegeskorte et al., 2008), and have also been applied to MEG data to compare the evolution of object representations over time (Carlson et al., 2013). Here we used this measure of object representational space to test for evidence that object-related information was passed from peri-occipital to peri-frontal brain regions (feedforward) and from peri-frontal to peri-occipital regions (feedback).

Specifically, we tested for Granger causal relationships between multivariate measures of object-related information (the DSMs) from the peri-frontal and peri-occipital sensors. The logic of Granger causality is that time series X 'Granger causes' time series Y if X contains information that helps predict the future of Y better than information in the past of Y alone (for a recent review on the application of Granger causality to neuroscience, see Friston et al., 2013). Since we were particularly interested in the evolution of these relationships over time we performed a sliding-window analysis of a simplified (special case) of Granger causality, using the partial correlations in Eqs. (3) and (4) to define 'Feedforward' (*FF*) and 'Feedback' (*FB*) information flows at each time point ($t$).

$$FF(t) = \rho\mathrm{DSM}_{(front,t)}\mathrm{DSM}_{(back,t-125)}.\mathrm{DSM}_{(front,t-125)} \quad (3)$$

$$FB(t) = \rho\mathrm{DSM}_{(back,t)}\mathrm{DSM}_{(front,t-125)}.\mathrm{DSM}_{(back,t-125)} \quad (4)$$

where $\mathrm{DSM}_{(loc,t)}$ is the DSM based on the sensors at location *loc* at time *t* *ms* post stimulus onset, and $\mathrm{DSM}_{(loc,t-125)}$ is the DSM based on the sensors at location *loc*, averaged across all time points from *t*-150 *ms* to *t*-100 *ms* post stimulus onset. A schematic illustration of this analysis is found in Fig. 6. In a review of electrophysiological estimates of visual response latencies, Lamme and Roelfsema (2000) report that for macaque, the mean visual response latencies in V1 and prefrontal cortex are 72 *ms* and 141 *ms,* respectively, although these response latencies will be slower for humans, due to their larger head size. We chose 100 *ms* as a generous estimate of the time taken for responses to be propagated from occipital to prefrontal areas and vice versa, and averaged across 100 *ms* to 150 *ms* so that our analysis was not dependent on an exact estimate of these propagation times.

We report the results of this analysis in terms of the difference between the feedforward and feedback information flows (*FF-FB*). To assess whether this difference was significantly above or below chance, we generated a null distribution of this difference at every time point by performing the same analysis on 1000 bootstraps of data from each subject where the exemplar labels were randomly shuffled for each of the DSMs used in Eqs. (3) and (4).

## Results

Using MEG, we measured participant's brain responses while they viewed a series of objects and made a simple judgment about the object's identity. Each object was presented in each of four stimulus conditions (described above, and illustrated in Fig. 1). Two of the stimulus conditions had signal in both the high and low spatial frequency bands, with either a low level of noise ('Strong signal' condition) or a high level of noise ('Weak signal' condition) added to all spatial frequencies. In the 'Low-pass' condition, signal in the low spatial frequency range was retained, while high spatial frequency components were randomized, and images the 'High-pass' condition had signal in the high spatial frequency range while low spatial frequency components were randomized. In all analyses of MEG data we included only data from trials where participants responded correctly.

### *Object identity, object category and stimulus condition decoding*

First we asked what information about object category and individual exemplar identity could be extracted from the MEG data. To address this question we used a series of classification analyses, collapsing across data from all stimulus conditions. In the first analysis, classifiers were trained to discriminate each pair of exemplars; the average classification across each pair of exemplars is plotted in Fig. 2A. Next, the exemplars were grouped into categories of animate or inanimate (Fig. 2B, upper plot), and smaller or larger than a shoebox (Fig. 2B, lower plot). These category boundaries both divided the 24 exemplars into two groups of 12 exemplars, and these two boundaries were orthogonal. Finally, we trained classifiers to discriminate between trials of different stimulus conditions. We trained a classifier to discriminate stimuli in the 'Low-pass' from the 'High-pass' condition (Fig. 2C, upper plot) and to discriminate stimuli in the 'Strong-signal' condition from the 'Weak-signal' condition (Fig. 2C, lower plot). These analyses were carried out for data from the peri-frontal and peri-occipital sensors separately.

The evolution of classification performance over time followed a similar pattern in each case and above chance classification performance in peri-occipital areas always preceded that in peri-frontal areas (where present) by 130–165 *ms*, ($p < 0.01$ in each case, based on bootstrapped difference values). For peri-occipital areas, the onset of above chance performance (taken as the start of the first 50 *ms* interval where classification performance was significantly above chance at all time points) was slightly earlier for individual exemplar classification (80 *ms* after stimulus onset) than for category boundaries (100 *ms*). The small differences in onset times between the exemplar decoding and the category decoding approached but did not reach significance when the bootstrapped distribution of differences was compared to zero ($p = 0.063$ for animacy versus exemplar, $p = 0.12$ for size versus exemplar). Classification performance was stronger overall when the classifier was trained to discriminate pairs of individual exemplars, so the slightly earlier onset of above chance classifier performance in this case might reflect the earlier accumulation of a stronger signal, rather than a difference in the time at which relevant information is present in peri-occipital areas. For peri-frontal areas, the onset of above chance classifier performance did not vary significantly between decoding object identity and decoding object category ($p = 0.34$ for animacy, $p = 0.40$ for size).

Separating responses to object category from those encoding individual object identity is complicated by the fact that classification performance could be based on either or both of these responses for both exemplar and category decoding. However, the fact that exemplar classification was more accurate than classification of object category suggests that when the classifier is trained to discriminate pairs of objects it cannot be using only the object's animacy or size (task related) category. By using the classification of objects according to a pseudo-random category rule we found the baseline level of classification performance we could expect if the classifier learnt groups of individual exemplars rather than using categorical object information present in the neural data. This baseline is shown in Fig. 2B as green and orange dashed lines for peri-occipital and peri-frontal classifiers, respectively. For both the animacy and size (task related) categories, the classifier performance exceeded this baseline for both peri-occipital and peri-frontal regions, although for peri-frontal regions there were no 50 *ms* periods where classification performance was significantly greater than for the random category (the criterion for classification 'onset'). For peri-occipital regions, the onset of classifier performance exceeding that of a random category was significantly ($p < 0.01$) later than the onset of

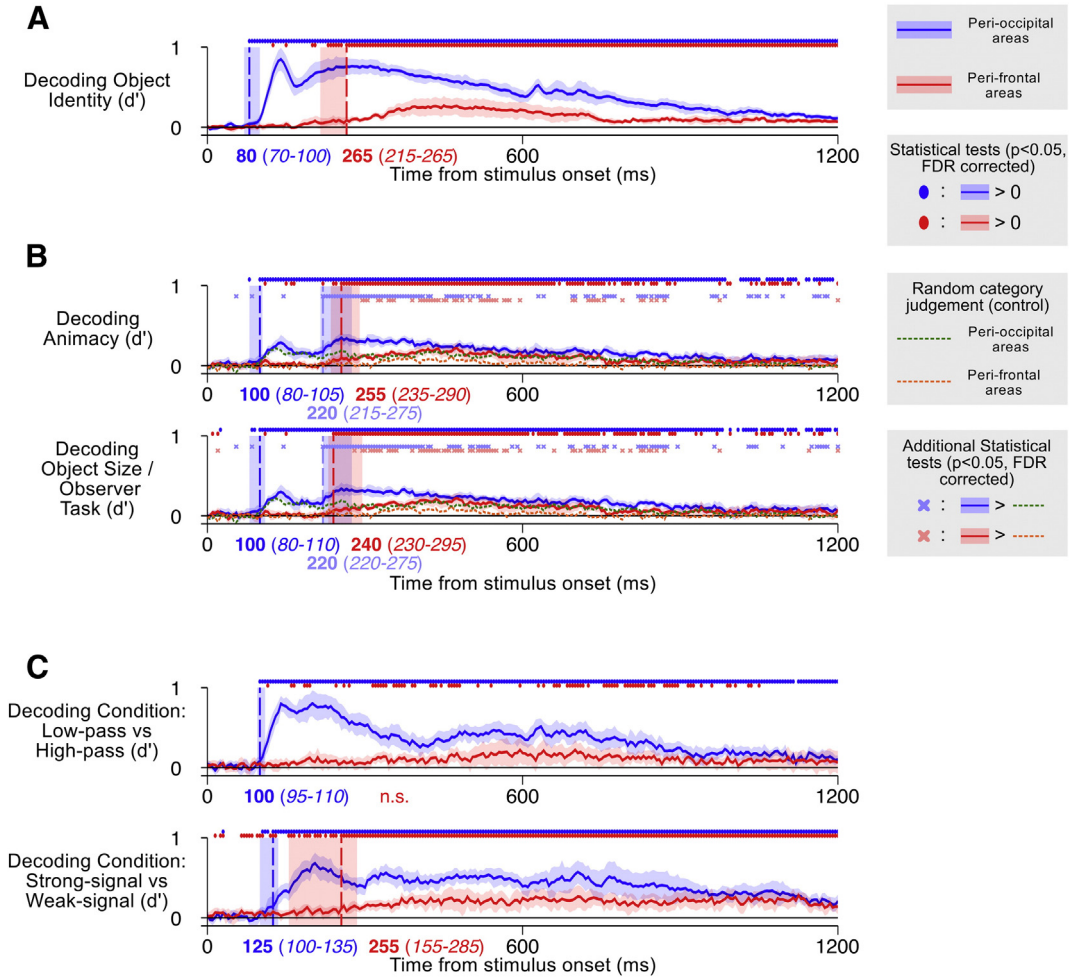


**Fig. 2.** Classification performance over time. Classifiers were trained and tested on a decision boundary based either on object identity, ignoring stimulus condition (**A** and **B**) or on stimulus condition ('Strong signal', 'Weak signal', 'Low-pass' and 'High-pass'), ignoring object identity (C). In **A**, classifiers were trained and tested on exemplar identity, averaged across each pair of exemplars. In **B**, classifiers were trained to assign objects to one of two categories, either animate versus inanimate objects (upper plot) or to perform the participant's task: choosing whether objects were larger or smaller than a shoebox, (lower plot). In **C**, data were split into two halves and the classifier was trained to discriminate object condition, either 'Low-pass' versus 'High-pass' (upper plot) or 'Strong-signal' versus 'Weak-signal' (lower plot). In all plots the blue and red filled lines show classifier accuracy based on data from peri-occipital and peri-frontal sensors, respectively, and shaded error bars indicate +/− 1.96 standard deviations (95% confidence intervals) of the between-subject mean ($n = 9$). Blue and red ovals indicate time points at which classification performance was significantly above 0 ($p < 0.05$, one-sided $t$-test, FDR corrected). For the category decision rules in **B** we also compared classifier performance on these category judgments to its performance when trained and tested on an arbitrary 'category' rule, where exemplars from the animacy and task categories were equally represented. Classifier performance that exceeded that found for this arbitrary category decision rule ($p < 0.05$, paired one-sided $t$-test, FDR corrected) is indicated by the light blue and light red crosses. Vertical lines (times labeled in bold type of corresponding color) show the onset of above chance classifier performance: the start of first 50 ms interval of significant classifier performance (where present). Vertical shaded error bars indicate the 95% confidence intervals of these onsets (range printed in italicized type).

above chance exemplar decoding. These results imply that at later times (~220 ms after stimulus onset) classifier performance was based on neural representations at the peri-occipital sensors where category information was an important feature, such that responses to within-category objects were more similar than to between-category objects.

Finally, we show the classifier performance when trained to discriminate the stimulus condition: 'Low-pass' versus 'High-pass' condition (Fig. 2C, upper plot) and 'Strong-signal' versus 'Weak-signal' condition (Fig. 2C, lower plot). These judgments are orthogonal to the classification of object identity and object category, and so do not provide any direct evidence of object-related information in the neural response. Instead, these results offer insight into differences in how the brain is processing stimuli from different conditions, which had exactly the same power at each spatial frequency, but varied in the signal strength at each spatial frequency. For both peri-occipital and peri-frontal sensors, the classifier performance was above chance in both comparisons. For peri-occipital sensors, the onset of decoding object condition was 20–45 *ms* later than the onset of decoding object identity; decoding of 'Strong-signal' versus 'Weak-signal' was significantly later ($p < 0.01$), while for 'Low-pass' versus 'High-pass' the difference approached significance ($p < 0.055$). The fact that decoding of object condition tended to be later than decoding object identity is consistent with the stimulus conditions being closely matched in their low-level image statistics. After the first information about object identity is present in the neural signal, suppression of irrelevant spatial frequencies and attention to the image signal could be used by the classifier to discriminate trials in the 'Low-pass' condition from those in the 'High-pass' condition. When the classifier was trained to discriminate 'Strong-signal' from 'Weak-signal' trials, the image signal is in the same spatial frequency bands. However, the classifier performance may be based on the subject attending to a greater extent or making more effort when the image signal was weak, or on a neural correlate of uncertainty.

### *Classification of object identity across stimulus conditions*

Next we considered the effect of stimulus condition ('Strong signal', 'Weak signal', 'Low-pass' or 'High-pass') on the decoding of object identity. To reduce our data set and facilitate comparisons between conditions, we used a curve fitting procedure to summarize the average classification of pairs of exemplars over time. For each participant, we

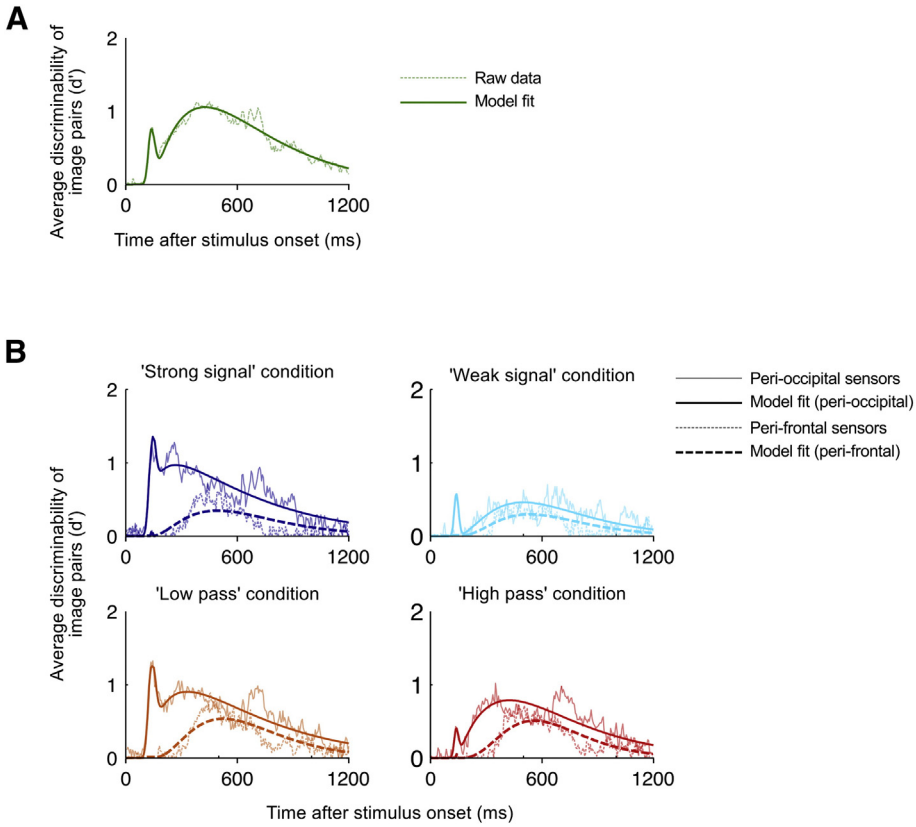


**Fig. 3. A** and **B** show the time course of average classifier accuracy for an example participant, averaged across the classifier's ability to discriminate each possible pair of objects, using d' as a measure of the discriminability of the image pairs. In **A**, the classifier was trained and tested on data from all four conditions across all sensors. In **B** the classifiers were trained and tested on data from each of the four conditions separately, and were restricted to data either from the peri-occipital sensors (filled lines) or peri-frontal sensors (dashed lines). In each case average classifier performance is overlaid by the best fitting model (a weighted sum of two gamma probability density functions, see text for details).

initially performed the classifier analysis on all their data, including each of the four stimulus conditions and MEG data from all 160 sensors. The time course of classifier accuracy based on the entire data set is shown in Fig. 3A for an example participant. We fit this curve with a weighted sum of two gamma probability density functions, and used the fitted parameters of this curve to restrict the model fits of other conditions for that participant (see Methods for details). The average early peak time for the peri-occipital sensors ranged from 138 *ms* ('Low-pass' condition) to 141 *ms* ('Weak signal' condition), while the average time of the late peak ranged from 227 *ms* ('Strong signal' condition) to 372 *ms* ('Weak signal' condition). For peri-frontal sensors, the average time of the late peak ranged from 403 *ms* ('Strong signal' condition) to 469 *ms* ('Weak signal' condition). In each case the time of the late peak covaried with the amplitude of the peak, consistent with a weaker signal taking longer to accumulate, and so we restricted our analyses to the amplitudes of the fitted peaks. The average amplitudes of these peaks across participants ($n = 9$) are plotted in Fig. 4.

To test for significant differences in the classifier performance between stimulus conditions, we performed a repeated-measures within-subjects analysis of variance (ANOVA) for both peri-occipital and peri-frontal sensors. For peri-occipital sensors this indicated a significant main effect of train and test condition on classifier accuracy, for both early ($F_{(3,24)} = 69.2$, $p = 0.001$), and late ($F_{(3,24)} = 27.3$, $p < 0.001$) peaks. As shown in Fig. 4, pairwise comparisons revealed that classifier performance was significantly better in the 'Strong signal' condition than in the 'Weak signal' condition, for both early and late peaks ($p < 0.01$), and better in the 'Low-pass' than 'High-pass' condition for the early peak ($p < 0.01$) but not the late peak ($p = 0.09$). For all pairwise comparisons we used the Dunn-Sidak correction for multiple comparisons. For peri-frontal sensors, there was no significant effect of train and test condition on classifier accuracy for either the early ($F_{(1.5,11.8)} = .25$, $p = .72$[1]) or late ($F_{(1.5,12.1)} = 1.2$, $p = .32$[1]) peaks.

*Generalization of classification performance across stimulus conditions*

Next we asked whether object information was represented in a similar way across stimulus condition. To address this we tested how well the classification decision boundaries generalized from one stimulus condition to another (Fig. 5). Specifically, for each participant we trained a classifier on all data from one condition (for example the 'Low-pass' condition), and then used this decision boundary to classify each trial from another condition (for example the 'High-pass' condition). The simplest example of this analysis is that shown in Fig. 5A, where we trained on data from the 'Low-pass' condition and tested on data from the 'High-pass' condition, and vice versa. In the unlikely scenario that the high and low spatial frequency information contributed to separate neural representations of object identity, there would be little or no generalization of the classification rule from the training to the test condition, which would result in low accuracy even where classification performance within each condition is high. Instead, as shown in Fig. 5A, the classification accuracy when training on 'Low-pass' and testing on 'High-pass' or vice versa was similar to that found when training and testing within the 'High-pass' condition, suggesting that signal from high and low spatial frequencies produces qualitatively similar representations of object identity. This does not mean that there is no information in the neural response about which spatial frequency is

[1] Degrees of freedom corrected using Greenhouse–Geisser Epsilon after Mauchly's test indicated that the assumption of sphericity had been violated; early: $\chi^2$ "15.0, p ă 0.05; late: $\chi^2$ "11.6, p ă 0.05.

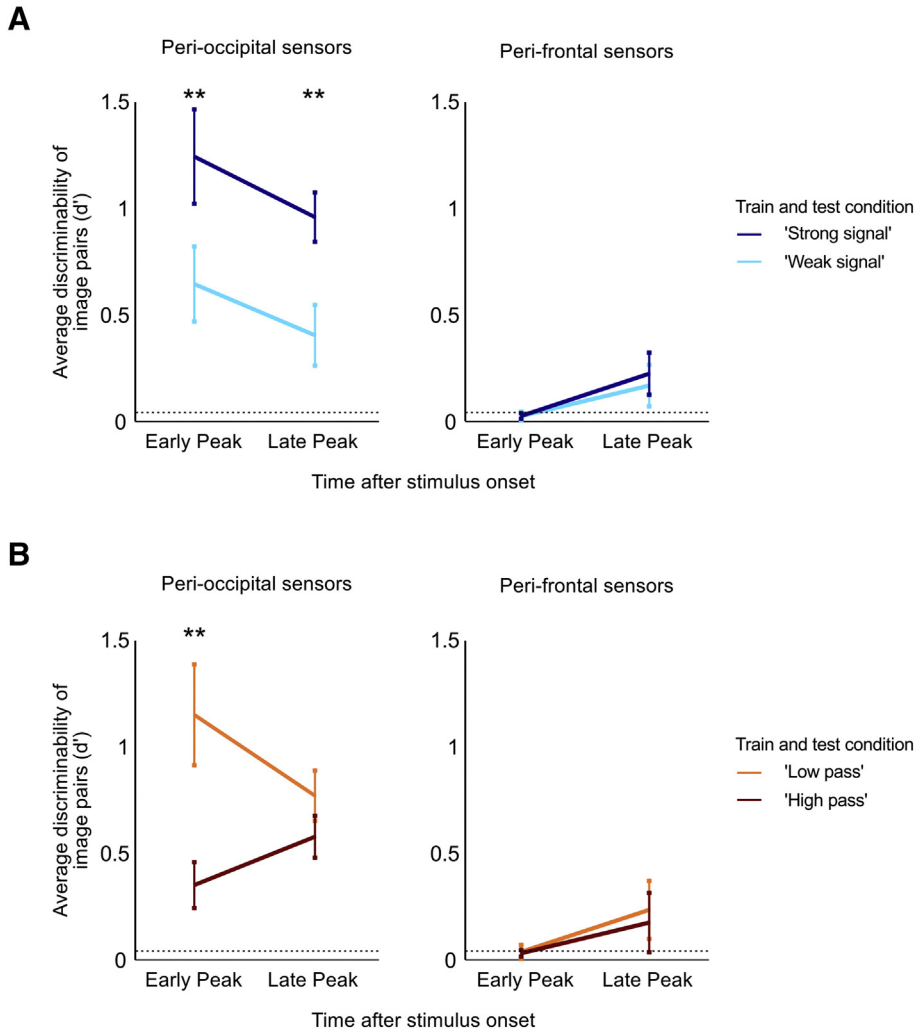


**Fig. 4.** Average within-condition classifier performance ($n = 9$) at both early and late peaks of the curves that were fit individually to the time courses of classifier performance over time for each participant (see illustration in Fig. 3 and text for details). Classifier performance based on data from peri-occipital sensors is shown in plots on the left, based on the peri-frontal sensors is shown in plots on the right. **A**: 'Strong signal' and 'Weak signal' stimulus conditions **B**: 'Low-pass' and 'High-pass' signal conditions. Error bars indicate 95% confidence intervals of the between-subject means, and dotted horizontal lines indicate expected peaks based on chance performance (see Methods for details). Asterisks (**) show where pairwise comparisons revealed a significant difference between the stimulus conditions ($p < 0.01$).

carrying signal: as shown in Fig. 2C and discussed above, classification of whether a stimulus belonged to the 'Low-pass' or 'High-pass' condition was above chance, but it means that there is also information in the signal that generalizes over the different spatial frequencies.

Is there any evidence that low-pass and high-pass signal produce qualitatively different representations of object identity? As a further test for evidence of this we used cross-classification performance to test whether the object-related information used by the classifier in the 'Strong signal' and 'Weak signal' conditions showed a bias towards low or high spatial frequencies. We trained classifiers on data from either the 'Strong signal' or 'Weak signal' condition, and used these classification boundaries to classify data from the 'Low-pass' and 'High-pass' conditions. If the object representation is qualitatively similar across conditions, then the cross-condition classification performance should depend only on the strength of that signal in the training and test conditions (as summarized in Fig. 4). Alternatively, if brain activity at the early peak had a bias (for example) for information contained in the low spatial frequencies, we would expect that a classifier trained on data from the broadband stimulus conditions ('Strong signal' and 'Weak signal') would show better generalization when tested on data from the 'Low-pass' condition than when tested on data from the 'High-pass' condition. Importantly, to signify a real bias, this difference in generalization must be beyond that predicted by the generally higher classifier performance within the 'Low-pass' compared with the 'High-pass' condition.

Results of the cross-condition classification analysis are shown in Fig. 5B. We conducted a series of analyses of covariance to test the significance of the effect of training and test stimulus condition on classification accuracy when controlling for the covariates of classification accuracy within the training and test conditions. At the early peak, there was no significant difference in classification accuracy between the training and test conditions for peri-occipital ($F_{(3,16)} = 12.50$, $p = .10$) and peri-frontal ($F_{(3,16)} = 0.89$, $p = .47$) sensors. Similarly, at the late peak there was no significant difference in classification accuracy between the training and test conditions for peri-occipital ($F_{(3,16)} = 2.07$, $p = .14$) and peri-frontal ($F_{(3,16)} = 1.34$, $p = .30$) sensors.

In summary, for each of the cross-condition classification analyses (Fig. 5A and B) our results are consistent with the classifier relying on a qualitatively similar representation of object identity. This does not exclude the possibility that spatial frequency dependent representations exist, but implies that the dominant representations of object identity that are learnt by the classifier are not highly spatial-frequency specific, even when the classifier is trained on brain responses to 'Low-pass' or 'High-pass' stimuli.

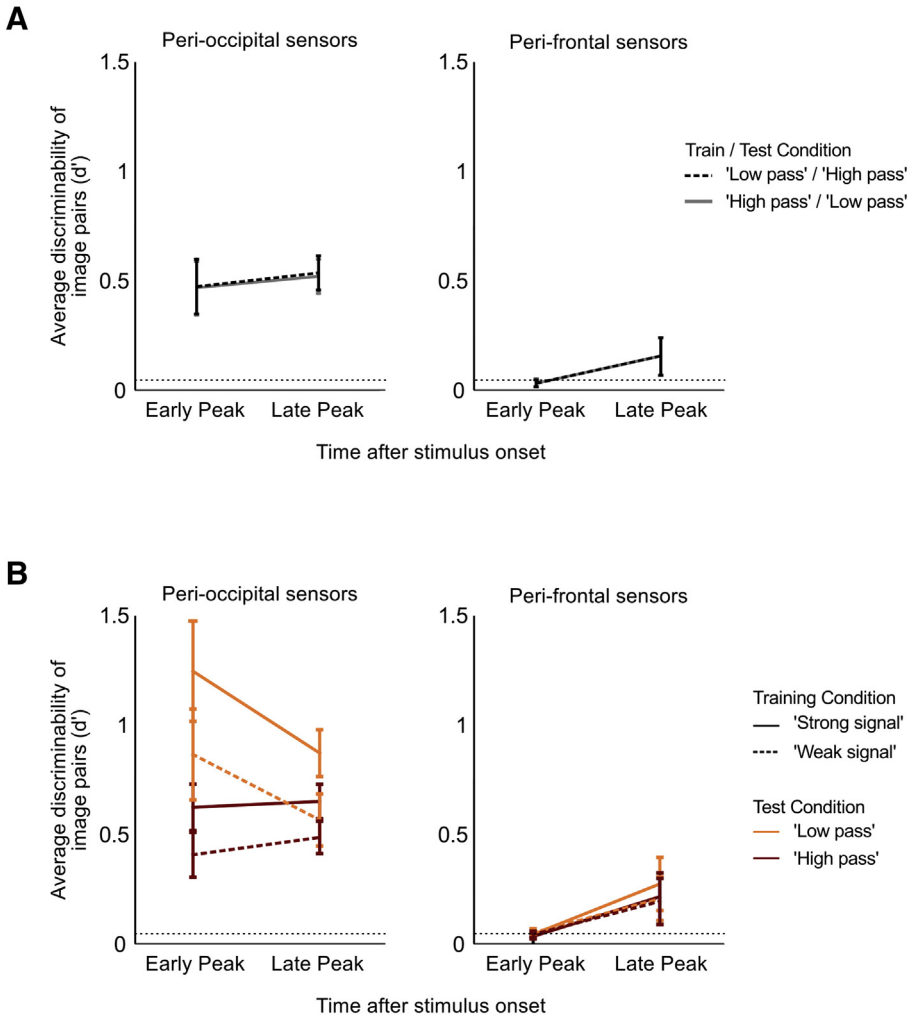


**Fig. 5.** Average cross-condition classifier accuracy when generalizing from one stimulus condition to another ($n = 9$). Plotted here is the average classification accuracy when a classifier is trained to discriminate each image pair using data from one stimulus condition, then tested on data from another stimulus condition, conventions as in Fig. 4. **A**: Generalization across different spatial frequency conditions. The classifier was trained on data from the 'Low-pass' condition and tested on data from the 'High-pass' condition (and vice versa). As seen in the plot, when the training and test conditions were swapped the classification performance was near identical; since this was found for all cases of generalization, in **B** only one direction of generalization is reported for each case. **B**: Generalization from the conditions with broadband spatial frequency content (the 'Strong signal' and 'Weak signal' conditions) to the 'Low-pass' and 'High-pass' conditions. An analysis of covariance (see text) did not reveal any significant differences between the conditions once the differing signal strengths of the conditions were factored out.

*Feedforward and feedback flows of object representations*

Finally, we tested for evidence of feedforward and feedback flows of information during the evolution of object representations. Using dissimilarity matrices (DSMs) as a measure of the object representational space at a given time, we quantified the Granger causal relationships between the evolution of DSMs based on peri-occipital sensors, and based on the peri-frontal sensors. Briefly, this quantifies the degree to which the DSM based on peri-occipital sensors predicted the DSM based on peri-frontal sensors at a later time, and vice versa. These methods are illustrated in Fig. 6, and the calculations of feedforward and feedback information flows are specified in Eqs. (3) and (4).

Results of this analysis are plotted in Fig. 7, for the average across all conditions (Fig. 7A) and for individual stimulus conditions (Fig. 7B–E). Across all stimulus conditions, feedforward flows of information were significantly greater than feedback flows of information within the first 500 *ms* post stimulus onset, that is, while the stimulus was visible. The earliest time at which the feedback flow of information was significantly greater than the feedforward flow was at ~300 ms post stimulus onset, for the 'Low-pass' condition.

Based on the Bar (2003); Bar et al. (2006) model of object perception, we predicted that there would be more evidence of feedback in the 'Low-pass' than the 'High-pass' condition. We found more time points with evidence of feedback flows of information in the 'Low-pass' than the 'High-pass' condition, but not at the early time (<120 ms after stimulus onset) predicted by the Bar (2003); Bar et al. (2006) model. Comparing feedforward and feedback flows across conditions is complicated by the fact that the estimates will vary with both the degree to which the signals are fed forward and fed back, and also on the signal strength. For instance, the greater feedback in the 'Strong signal' compared to the 'Weak signal' condition is likely due to the lower signal strength in the 'Weak signal' condition, rather than necessarily showing there is less feedback in the 'Weak signal' case. The estimates of feedforward and feedback information flows will be affected by noise in the dissimilarity matrices (DSMs) when signal strength is lower, as in the case of the 'Weak signal' condition. Nonetheless, we believe this approach provides a useful new method for analyzing the connectivity between brain regions in terms of the information they carry about a specific feature of interest (in this case, object identity).

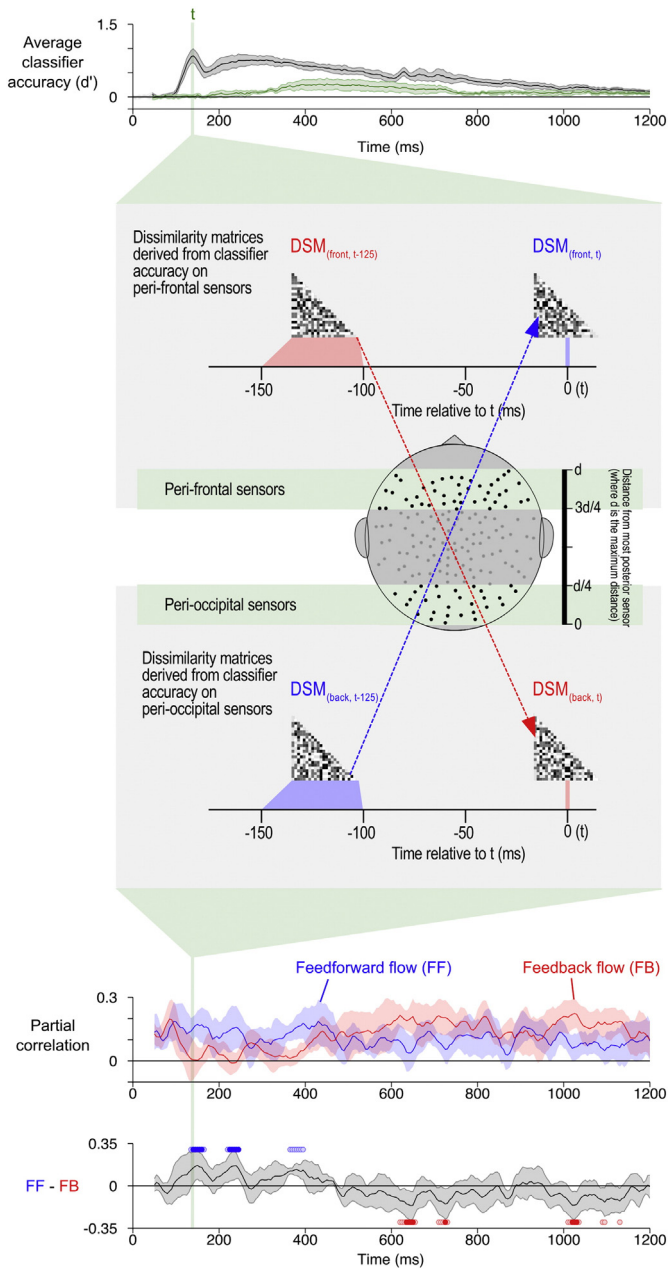


**Fig. 6.** Illustration of the method for calculating simple Granger causal relationships between the peri-frontal and peri-occipital sensors, dubbed feedfoward and feedback information flows. In a 'sliding window' analysis, at each time point ($t$) we compared the dissimilarity matrices (DSMs) derived from classifier accuracy on the peri-frontal sensors and the peri-occipital sensors. To estimate the 'feedforward flow' of information (shown in blue) we took the partial correlation between the DSM based on peri-frontal sensors at time $t$ ($DSM_{(front,t)}$), and the average DSM based on the peri-occipital sensors from time $t$-150 ms and $t$-100 ms ($DSM_{(back,t-125)}$), discounting the correlation with the average DSM based on the peri-frontal sensors from time $t$-150 ms and $t$-100 ms ($DSM_{(front,t-125)}$), as given in Eq. (3). We used the same principle with the same time points to estimate the 'feedback flow' of information (shown in red, and specified in Eq. (4)). This analysis yielded an estimate of the feedforward and feedback flows of information for each time point, for each participant. The plot at the top shows the average classifier accuracy across participants (classifier performance based on peri-occipital and peri-frontal sensors is shown in black and green, respectively). The plots at the bottom show the estimated feedforward and feedback information flows (upper plot), and the difference between them (lower plot). In each case the shaded error bars indicate the 95% confidence intervals of the between-subject mean. In the lower plot, times at which the difference was significantly positive or negative are indicated with circles, filled with blue and red, respectively. Light circles show time points where the difference was significant at $p < 0.01$, (uncorrected) and dark circles indicate $p < 0.05$ (FDR correction for multiple comparisons across time).

## Discussion

What neural mechanisms underlie our ability to rapidly and accurately identify visual objects? Previous work has suggested that low and high spatial frequency stimulus information make differential contributions to object perception. In particular, a range of studies (Bar, 2003; Bar et al., 2006; Kveraga et al., 2007; Chaumon et al., 2014; Fintzi and Mahon, 2014) have found evidence consistent with low-pass object information being rapidly projected (within 130 *ms* of stimulus onset) to frontal areas where it is fed back to occipital object-selective areas. However, no study to date has tested for evidence that object identity information is present in the frontal regions at this early time after stimulus onset. Here we applied multivariate pattern classification analysis to MEG data (Carlson et al., 2013; Isik et al., 2014) to test for evidence of object identity information in both peri-occipital and peri-frontal areas, and to characterize the time course of how this information evolves.

Importantly, unlike previous studies, we used stimuli that had the same power at each spatial frequency, and the same overall contrast, varying only the 'diagnostic' spatial frequencies (De Gardelle and Kouider, 2010). We equated our stimulus conditions for these low level features to facilitate better comparisons between conditions, given that the early response of the occipital cortex to visual objects, as measured using EEG and/or MEG, covaries with low-level stimulus statistics (Martinovic et al., 2008; Clarke et al., 2013).

### *Object related information in peri-frontal and peri-occipital regions*

Across all conditions, there was an early (138–140 *ms*) and a late (227–372 *ms*) peak in classifier performance when the classifier was trained on data from the peri-occipital sensors. In peri-frontal sensor data we did not find any evidence of early (<200 *ms*) neural encoding of object identity, but we did find evidence of object identity information at a later time, peaking at 403–469 *ms* after stimulus onset. This result is seemingly inconsistent with the previous findings of Bar et al. (2006), who reported the left OFC (located within our broadly defined 'peri-frontal' sensors) showed differential activity to recognized versus not recognized objects as early as 130 *ms* after stimulus onset.

One possible reason for the discrepancy between our results and those of Bar et al. (2006) is that our study explicitly tests for object-identity information by using a classification analysis, whereas Bar et al. (2006) reported differences in the average response of the OFC when a stimulus is or is not recognized. It could be that the object-related responses found in previous studies (Thorpe et al., 1996; Bar et al., 2006) are associated with the detection of object, without carrying specific information about object identity.

Alternatively, it is possible that object identity information was present in the front of the brain at an early time but too weak to be detected. The fact that we find significant classifier performance based on the peri-frontal sensors at a later time demonstrates that our method is sensitive enough to reveal object identity information in the peri-frontal sensors, but our stimulus conditions may not have been optimal for revealing any early object identity information in prefrontal cortex. For example, object identity may be decodable from the early activity of prefrontal cortex when there is a shorter stimulus duration or in the presence of masking, as used in previous work (Bar et al., 2006; Thorpe et al., 1996).

### *Comparison of object-identity information across different conditions*

We compared brain responses to the different stimulus conditions using the average discriminability of each image pair when the classifier was trained on recordings of the peri-occipital and peri-frontal sensors. We found significant differences between conditions for the peri-occipital but not the peri-frontal sensors, even though classification accuracy based on the peri-frontal sensors was above chance at the

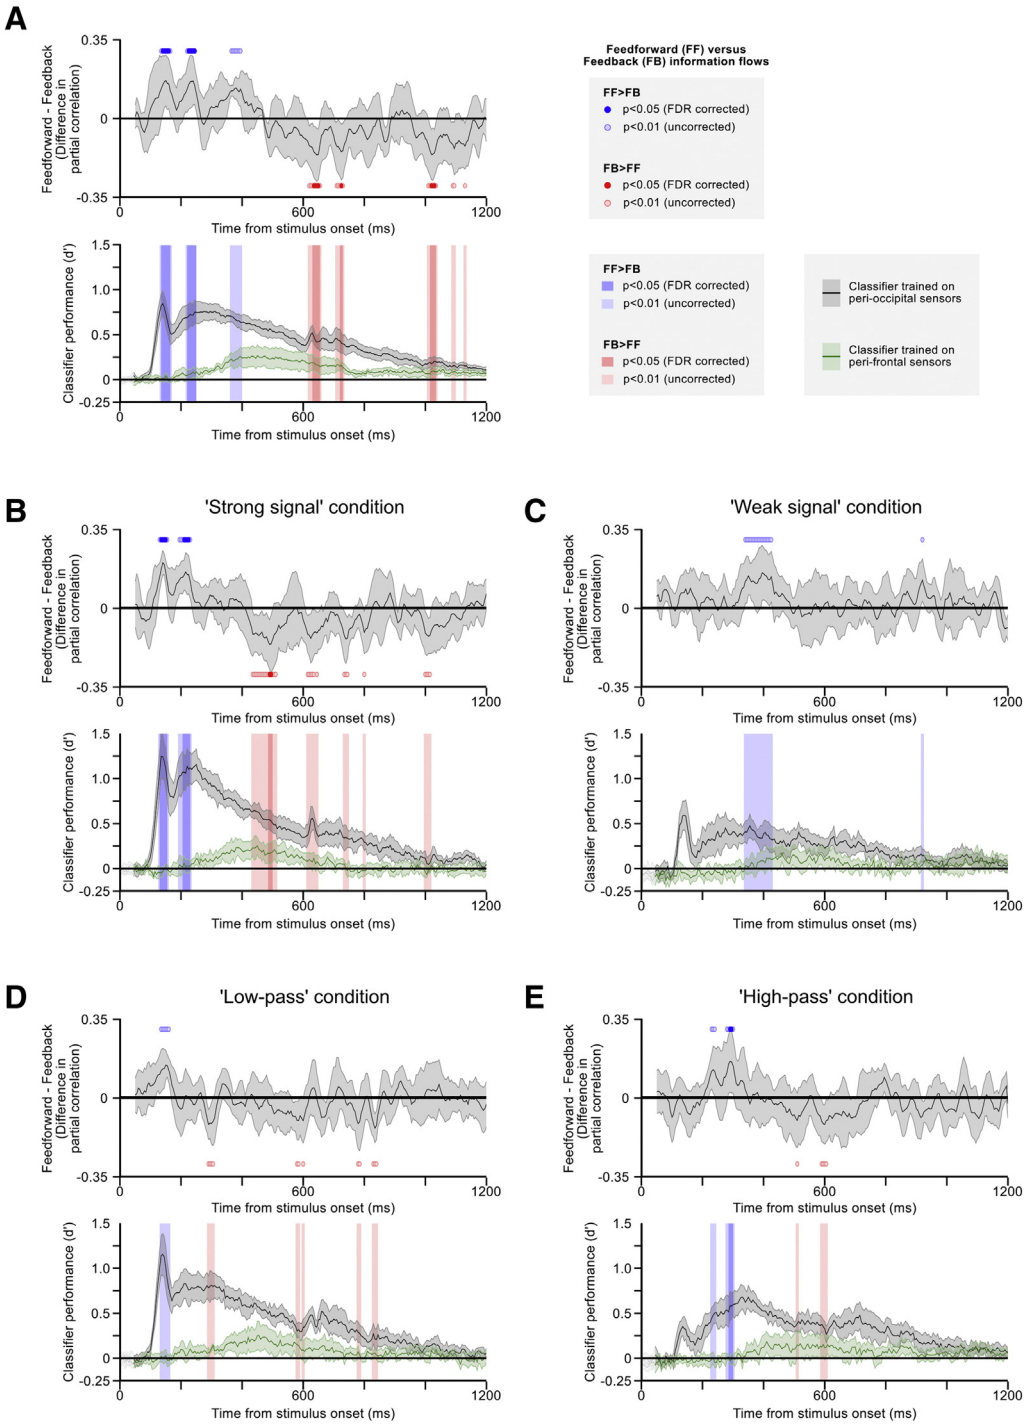


**Fig. 7.** Difference between feedforward (*FF*) and feedback (*FB*) flows of information (upper plots) and average classifier performance (lower plots) across all conditions (**A**), and for the 'Strong signal' (**B**), 'Weak signal' (**C**), 'Low-pass' (**D**) and 'High-pass' (**E**) conditions. Partial correlations were used to estimate *FF* and *FB* flows of information (see Fig. 6), and the difference between these flows is plotted in the upper plots of **A**–**E**, with shaded gray error bars indicating the 95% confidence intervals of the between-subject mean. The filled circles indicate times at which this difference was significantly above or below zero (as assessed by comparison with a bootstrapped null distribution) for two statistical thresholds: $p < 0.01$ (uncorrected) is shown in lighter circles, and the more conservative $p < 0.05$ (FDR corrected for multiple comparisons across time points) is shown in darker circles. In the lower plots of **A**–**E**, average classifier performance is plotted for classifiers based on the peri-occipital sensors (black) and peri-frontal sensors (green), with shaded error bars indicating the 95% confidence intervals of the between-subject mean. For these plots of average classifier performance, times at which the *FF-FB* difference was significantly above or below zero are replotted, in blue where *FF-FB* > 0 and red where *FF-FB* < 0.

late peak (403–469 *ms*) for all four stimulus conditions. This interaction might mean that responses to the different conditions are more similar in frontal regions; however, since the classifier performance was lower overall for peri-frontal than for peri-occipital sensors we cannot exclude the possibility that the interaction is caused by the weaker signal making it harder to detect between-condition differences.

For peri-occipital sensors, at both early and late peaks, the average classification performance was significantly better with data from the 'Strong signal' than the 'Weak signal' condition, and classification performance was stronger at the earlier than the later peak for both conditions (Fig. 4A). When the stimulus signal was restricted to either low or high spatial frequencies, the classification performance was

significantly better for the 'Low-pass' condition than the 'High-pass' condition at the early peak, but at the late peak there was no significant difference between these conditions (Fig. 4B).

Overall, classification performance at the late peak reflected task performance better than at the early peak, although both were well before average reaction time (late peak: 227–372 *ms*, median reaction time: 1750 *ms*). The main difference was driven by the 'High-pass' condition: participant performance was better for 'High-pass' than for 'Weak signal' stimuli, but at the early peak, classifier performance for the 'Weak signal' condition exceeded that of the 'High-pass' condition. At the late peak, classifier performance on these two conditions was reversed, reflecting participant performance, implying that object representation decoded at the later peak better reflected the 'final' object representation on which the participant's task decision was based.

*Evidence of coarse-to-fine processing for the peri-occipital sensors*

Interestingly, the 'High-pass' condition for peri-occipital sensors was the only condition where classifier decoding was better at the later peak than at the early peak. This result implies that the high spatial frequency information about object identity took longer to become available than the low spatial frequency object information. This interpretation is consistent with the fact that classifier performance in the 'Low-pass' condition was higher than the 'High-pass' condition at the early peak but at the later peak the classifier performance was similar for these conditions. Critically, all our experimental stimuli had the same Fourier amplitude spectra, meaning that they were matched in terms of their spatial frequency content. This ensured that any differences between the 'Low-pass' and 'High-pass' conditions did not reflect the presence or absence of power in the image at the relevant spatial frequencies. Instead, the differences between these conditions reflect differences between the ways in which image signal at these spatial frequencies is processed, and how it contributes towards object perception.

These findings are consistent with previous psychophysical results that suggest a coarse-to-fine bias in human scene and object perception. Coarse-to-fine processing has been used to describe a broad range of physiological and perceptual phenomena, and it remains unclear how or whether these observed effects are related (Hegd´e, 2008), but in each case refers to the general precedence of low over high spatial frequencies in visual perception. For instance, when participants are shown a rapid sequence of scene or face images containing a narrow range of spatial frequencies, they more readily integrate information when the filtered images are presented in a coarse-to-fine than 'fine-to-coarse' order Schyns and Oliva (1994); Parker et al. (1992, 1997). For simple grating stimuli, Hughes et al. (1996) found that a low frequency distracter interfered with the processing of the high frequency target when presented up to 100 *ms* after the onset of the target, but the high frequency distracter did not interfere with the processing of the low frequency target. Each of these results implies that the coarse-to-fine presentation more closely reflects the order in which the visual system extracts information under natural viewing conditions.

Using reverse correlation of stimulus and single-unit recordings in macaque, Mazer et al. (2002) and Bredfeldt and Ringach (2002) found that the preferred spatial frequency of the majority of V1 cells shifts from lower to higher frequencies over the course of their stimulus-evoked response. Applying reverse correlation methods to human perception, Neri (2011) found evidence of an early non-linear process with coarse spatial tuning followed by a linear process with fine spatial tuning. However, as noted by Neri (2011) and Mareschal et al. (2006), it is unclear whether these behavioral effects can be attributed to the response properties of single cells as revealed with electrophysiology.

The findings of the present study suggest there is a coarse-to-fine progression in the population-level of the neural response to objects in peri-occipital areas: object representations based on high spatial frequency information take longer to evolve than those based on low spatial frequency information, suggesting that this information is slower to propagate through the early visual system. However, cross-classification analyses (Fig. 5A) demonstrated that the object identity information learnt by the classifier in the 'Low-pass' condition could be used to accurately discriminate stimuli from the 'High-pass' condition, and vice versa, at both early and late peaks. This cross-condition classification performance suggests that signal from high and low spatial frequencies results in qualitatively similar representations of object identity, even though low-spatial frequency signals are propagated more quickly by precortical stages of the visual system.

*Evidence for feedforward and feedback information flows*

By applying a simplified Granger causality analysis to our multivariate measure of object representation, we were also able to measure the extent to which the object-identity information from the peri-occipital sensors could predict later object representations of the peri-frontal sensors, and vice versa (Fig. 7). We used this analysis to test for evidence of feedforward and feedback flows of information. For all conditions, times of significantly greater feedforward than feedback occurred within the first 500 *ms* (that is, while the stimulus was still present). This result is consistent with the visual system continuing to receive 'bottom-up' information about object identity while the stimulus is present. After the offset of the stimulus, the interactions between frontal areas and early visual areas are dominated by 'top-down' information flows.

We found the earliest feedback in the 'Low-pass' condition, which is consistent with the faster processing of low spatial frequency information, as discussed above. The onset of this feedback (~300 *ms* after stimulus onset) occurred while the stimulus was still present, at a time when there was evidence of feedforward flows in other conditions ('High-pass' and 'Weak-signal' conditions). The onset of feedback in the 'Low-pass' condition is still not as early as that predicted by Bar et al. (2006) (130–180 *ms*). However, a limitation of the present analysis for evaluating feedback in the first few hundred milliseconds is that feedforward flows of information were dominant while the stimulus was present, suggesting that future work with shorter stimulus presentations may better elucidate the role of feedback immediately after stimulus onset.

This analysis of feedforward and feedback information offers a useful new tool for investigating the flow of relevant information between brain areas, rather than a simpler analysis of their connectivity which may or may not reflect connectivity related to the processing of a feature of interest.

## Conclusions

We used MEG to explore the representational dynamics of visual object processing, and the impact of stimulus spatial frequency on these dynamics. Our study confirms that low and high spatial frequency information about object identity is processed differently by the visual system: information about object identity that is carried by low spatial frequency signals is more readily available in peri-occipital areas at an early time (~140 *ms* after stimulus onset), consistent with previous behavioral studies demonstrating the coarse-to-fine progression of visual perception. Our new method for measuring feedforward and feedback information flows offers a new approach for testing whether brain regions are functionally connected and communicating specific information about a feature of interest. We believe this approach could be applied to a variety of experimental questions, and may be useful in future work to further reveal the role of prefrontal cortex in object perception.

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.neuroimage.2016.01.006.

## Acknowledgments

This project was funded under an ARC DECRA (DE120100898), Future Fellowship (FT120100816), and Discovery Project (DP160101300). We

thank Erin Martin and Erika Contini for assistance with data collection, and Simon Kwok for advice on Granger causality.

## References


Bar, M., 2003. A cortical mechanism for triggering top-down facilitation in visual object recognition. J. Cogn. Neurosci. 15, 600–609.

Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmid, A.M., Schmidt, A.M., Dale, A.M., Hamalainen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R., Halgren, E., 2006. Top-down facilitation of visual recognition. Proc. Natl. Acad. Sci. U. S. A. 103, 449–454.

Brainard, D.H., 1997. The psychophysics toolbox. Spat. Vis. 10, 433–436.

Bredfeldt, C.E., Ringach, D.L., 2002. Dynamics of spatial frequency tuning in macaque V1. J. Neurosci. 22, 1976–1984.

Bullier, J., 2001. Integrated model of visual processing. Brain Res. Brain Res. Rev. 36, 96–107.

Carlson, T.A., Hogendoorn, H., Kanai, R., Mesik, J., Turret, J., 2011. High temporal resolution decoding of object position and category. J. Vis. 11 (9), 1–7.

Carlson, T., Tovar, D., Alink, A., Kriegeskorte, N., 2013. Representational dynamics of object vision: The first 1000 ms. J. Vis. 13 (1), 1–19.

Chaumon, M., Kveraga, K., Barrett, L.F., Bar, M., 2014. Visual predictions in the orbitofrontal cortex rely on associative content. Cereb. Cortex 24 (11), 2899–2907.

Clarke, A., Taylor, K.I., Devereux, B., Randall, B., Tyler, L.K., 2013. From perception to conception: How meaningful objects are processed over time. Cereb. Cortex 23, 187–197.

De Gardelle, V., Kouider, S., 2010. How spatial frequencies and visual awareness interact during face processing. Psychol. Sci. 21, 58–66.

Fintzi, A.R., Mahon, B.Z., 2014. A bimodal tuning curve for spatial frequency across left and right human orbital frontal cortex during object recognition. Cereb. Cortex 24, 1311–1318.

Friston, K., Moran, R., Seth, A.K., 2013. Analysing connectivity with Granger causality and dynamic causal modelling. Curr. Opin. Neurobiol. 23, 172–178.

Genovese, C.R., Lazar, N.A., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. NeuroImage 15, 870–878.

Grill-Spector, K., Kourtzi, Z., Kanwisher, N., 2001. The lateral occipital complex and its role in object recognition. Vis. Res. 41, 1409–1422.

Hegd´e, J., 2008. Time course of visual perception: Coarse-to-fine processing and beyond. Prog. Neurobiol. 84, 405–439.

Hochstein, S., Ahissar, M., 2002. View from the top: Hierarchies and reverse hierarchies in the visual system. Neuron 36, 791–804.

Horr, N.K., Braun, C., Volz, K.G., 2014. Feeling before knowing why: The role of the orbitofrontal cortex in intuitive judgments–an MEG study. Cogn. Affect. Behav. Neurosci. 14, 1271–1285.

Hughes, H.C., Nozawa, G., Kitterle, F., 1996. Global precedence, spatial frequency channels, and the statistics of natural images. J. Cogn. Neurosci. 8, 197–230.

Isik, L., Meyers, E.M., Leibo, J.Z., Poggio, T., 2014. The dynamics of invariant object recognition in the human visual system. J. Neurophysiol. 111, 91–102.

Kado, H., Higuchi, M., Shimogawara, M., Haruta, Y., Adachi, Y., Kawai, J., Ogata, H., Uehara, G., 1999. Magnetoencephalogram systems developed at KIT. IEEE Trans. Appl. Supercond. 9, 4057–4062.

Kirchner, H., Thorpe, S.J., 2006. Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. Vis. Res. 46, 1762–1776.

Kontsevich, L.L., Tyler, C.W., 1999. Bayesian adaptive estimation of psychometric slope and threshold. Vis. Res. 39, 2729–2737.

Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., Bandettini, P.A., 2008. Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron 60, 1126–1141.

Kveraga, K., Ghuman, A.S., Bar, M., 2007. Top-down predictions in the cognitive brain. Brain Cogn. 65, 145–168.

Lamme, V.A., Roelfsema, P.R., 2000. The distinct modes of vision offered by feedforward and recurrent processing. Trends Neurosci. 23, 571–579.

Mareschal, I., Dakin, S.C., Bex, P.J., 2006. Dynamic properties of orientation discrimination assessed by using classification images. Proc. Natl. Acad. Sci. U. S. A. 103, 5131–5136.

Martinovic, J., Gruber, T., MM, M.¨l., 2008. Coding of visual object features and feature conjunctions in the human brain. PLoS One 3, e3781.

Mazer, J.A., Vinje, W.E., McDermott, J., Schiller, P.H., Gallant, J.L., 2002. Spatial frequency and orientation tuning dynamics in area V1. Proc. Natl. Acad. Sci. U. S. A. 99, 1645–1650.

Neri, P., 2011. Coarse to fine dynamics of monocular and binocular processing in human pattern vision. Proc. Natl. Acad. Sci. U. S. A. 108, 10726–10731.

Parker, D.M., Lishman, J.R., Hughes, J., 1992. Temporal integration of spatially filtered visual images. Perception 21, 147–160.

Parker, D.M., Lishman, J.R., Hughes, J., 1997. Evidence for the view that temporospatial integration in vision is temporally anisotropic. Perception 26, 1169–1180.

Patai, E.Z., Buckley, A., Nobre, A.C., 2013. Is attention based on spatial contextual memory preferentially guided by low spatial frequency signals? PLoS One 8, e65601.

Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. Spat. Vis. 10, 437–442.

Peyrin, C., Michel, C.M., Schwartz, S., Thut, G., Seghier, M., Landis, T., Marendaz, C., Vuilleumier, P., 2010. The neural substrates and timing of top-down processes during coarse-to-fine categorization of visual scenes: A combined fMRI and ERP study. J. Cogn. Neurosci. 22, 2768–2780.

Schyns, P., Oliva, A., 1994. From blobs to boundary edges: Evidence for time- and spatial-scale-dependant scene recognition. Am. Psychol. Soc. 5, 195–200.

Stein, T., Seymour, K., Hebart, M.N., Sterzer, P., 2014. Rapid fear detection relies on high spatial frequencies. Psychol. Sci. 25, 566–574.

Tanaka, K., 1996. Inferotemporal cortex and object vision. Annu. Rev. Neurosci. 19, 109–139.

Tapia, E., Breitmeyer, B.G., 2011. Visual consciousness revisited: magnocellular and parvocellular contributions to conscious and nonconscious vision. Psychol. Sci. 22, 934–942.

Thorpe, S., Fize, D., Marlot, C., 1996. Speed of processing in the human visual system. Nature 381, 520–522.

Uehara, G., Adachi, Y., Kawai, J., Shimogawara, M., Higuchi, M., Haruta, Y., Ogata, H., Kado, H., 2003. Multi-channel SQUID systems for biomagnetic measurement. IEICE Trans. Electron. E86-C, 43–54.