

Feature-Specific Information Processing Precedes Concerted Activation in Human Visual Cortex

Pavan Ramkumar,¹ Mainak Jas,¹ Sebastian Pannasch,¹ Riitta Hari,^{1,2} and Lauri Parkkonen^{1,2,3}

¹Brain Research Unit and MEG Core, O.V. Lounasmaa Laboratory, School of Science, ²Advanced Magnetic Imaging Centre, School of Science, and

³Department of Biomedical Engineering and Computational Sciences, School of Science, Aalto University, FI-00076 Aalto, Espoo, Finland

Current knowledge about the precise timing of visual input to the cortex relies largely on spike timings in monkeys and evoked-response latencies in humans. However, quantifying the activation onset does not unambiguously describe the timing of stimulus-feature-specific information processing. Here, we investigated the information content of the early human visual cortical activity by decoding low-level visual features from single-trial magnetoencephalographic (MEG) responses. MEG was measured from nine healthy subjects as they viewed annular sinusoidal gratings (spanning the visual field from 2 to 10° for a duration of 1 s), characterized by spatial frequency (0.33 cycles/degree or 1.33 cycles/degree) and orientation (45° or 135°); gratings were either static or rotated clockwise or anticlockwise from 0 to 180°. Time-resolved classifiers using a 20 ms moving window exceeded chance level at 51 ms (the later edge of the window) for spatial frequency, 65 ms for orientation, and 98 ms for rotation direction. Decoding accuracies of spatial frequency and orientation peaked at 70 and 90 ms, respectively, coinciding with the peaks of the onset evoked responses. Within-subject time-insensitive pattern classifiers decoded spatial frequency and orientation simultaneously (mean accuracy 64%, chance 25%) and rotation direction (mean 82%, chance 50%). Classifiers trained on data from other subjects decoded the spatial frequency (73%), but not the orientation, nor the rotation direction. Our results indicate that unaveraged brain responses contain decodable information about low-level visual features already at the time of the earliest cortical evoked responses, and that representations of spatial frequency are highly robust across individuals.

Introduction

Neurons of the mammalian primary visual cortex (V1) are sensitive to low-level visual features, such as orientation (OR), and spatial frequency (SF). Neurons with similar response properties are organized close to each other at the fine-scale columnar level as well as in coarse-scale spatial maps (Chalupa and Werner, 2003). Most of the knowledge about the cortical organization of neurons selective to low-level visual features has been obtained from invasive animal electrophysiology (Hubel and Wiesel, 1963) and optical imaging (Bonhoeffer and Grinvald, 1991), whereas direct observations of homologous neurons in the human visual system are limited.

In particular, studies investigating the timing of stimulus-feature-specific information processing have followed an

“activation-based” approach. For instance, invasive recordings from the macaque visual cortex have characterized the temporal progression of activation by quantifying the timing of neuronal firing in different visual areas (Schmolesky et al., 1998; Lamme and Roelfsema, 2000). Likewise, noninvasive recordings with electroencephalography (EEG; Di Russo et al., 2002) or magnetoencephalography (MEG; Vanni et al., 2001) have quantified the earliest discernible evoked-response onsets or peaks. Although a detailed description of the temporal sequence of cortical activation is valuable, it does not present a complete picture of stimulus-feature-specific information processing. To complement activation-based approaches (Hari et al., 2010), we directly investigated the information content of unaveraged single-trial, early visual MEG responses using a multivariate decoding strategy.

Until recently, noninvasive functional imaging methods were considered to be too coarse to exploit the fine-scale organization of the human visual cortex. However, multivariate pattern analyses (MVPA) of functional magnetic resonance imaging (fMRI) data comprising activity from voxels larger than 25 mm³ in volume have enabled above chance-level decoding of the orientation of visual gratings from the human primary visual cortex (Haynes and Rees, 2005; Kamitani and Tong, 2005). Consequently, Kriegeskorte et al. (2006) proposed information-based brain mapping as a broad approach for functional neuroimaging to map brain regions containing information about the task or stimulus under study. In this approach, one seeks information about prespecified attributes of the task or stimulus, embedded in the activity patterns of functional imaging.

Received Aug. 8, 2012; revised Feb. 19, 2013; accepted March 20, 2013.

Author contributions: P.R., S.P., R.H., and L.P. designed research; P.R., M.J., and S.P. performed research; P.R., M.J., and S.P. contributed unpublished reagents/analytic tools; P.R., M.J., S.P., R.H., and L.P. analyzed data; P.R., M.J., S.P., R.H., and L.P. wrote the paper.

We were supported by the Academy of Finland, the Brain2Brain European Research Council Advanced Grant #232946 (R.H.), and the FP7-PEOPLE-2009-IEF program #254638 (S.P.). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of this manuscript.

The authors declare no competing financial interests.

This article is freely available online through the J. Neurosci. Author Open Choice option.

Correspondence should be addressed to Pavan Ramkumar, Brain Research Unit and MEG Core, O.V. Lounasmaa Laboratory, School of Science, Aalto University FI-00076 Aalto, Espoo, Finland. E-mail: pavan@neuro.hut.fi.

P. Ramkumar's present address: Sensorimotor Performance Program, Rehabilitation Institute of Chicago and Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL 60611.

S. Pannasch's present address: Institut für Psychologie III, Ingenieurpsychologie und Kognitive Ergonomie, Technische Universität D-01062 Dresden, Germany.

DOI:10.1523/JNEUROSCI.3905-12.2013

Copyright © 2013 the authors 0270-6474/13/337691-09\$15.00/0

In this study, encouraged by the success of information-based brain mapping with fMRI, we attempted to go beyond the activation-based quantifications of temporal information processing by (1) investigating whether it is possible to decode low-level visual features, such as OR, SF, and rotation direction (RD), from single-trial MEG responses using time-insensitive pattern classifiers, and if so, (2) quantifying how early and for how long the responses contain information allowing above-chance decoding using time-resolved pattern classifiers.

Materials and Methods

Subjects

Nine healthy volunteers (four females and five males; mean age 27 years, range 21–39 years) with normal or corrected-to-normal vision participated in the study after written informed consent. The recordings had prior approval by the Ethics Committee of the Helsinki and Uusimaa Hospital District (protocols No. 9-49/2000 and No. 95/13/03/00/2008, granted to N. Forss and R. Hari). Seven of nine subjects participated in all the main experiments (see the following section.).

Stimuli and experiments

Figure 1 shows the grating stimuli. In the main experiments, static or dynamic full-field annuli spanned 2–10° in the parafoveal visual field (Fig. 1*A,B*). The annuli were filled with linear sinusoidal gratings at two different ORs (45 and 135°) and two different SFs (0.33 and 1.33 cycles per degree [c/deg]). The dynamic stimuli comprised annuli with monotonically increasing (or decreasing) OR, i.e., gratings rotating clockwise or anticlockwise from 0 to 180° or vice versa over a 1 s duration.

In two control experiments with a single subject, SF classification performance was compared for gratings at three contrast levels and three phases: the 2–10° annuli comprised sinusoidal gratings of three different contrast levels (full, half, and one-fourth; Fig. 1*C*), or three different relative phases (0, one-fourth cycle phase shift, and half cycle phase shift; Fig. 1*D*).

For all stimuli except the contrast control experiment (see below, Control experiment 1: spatial frequency coding as a function of contrast), the luminance (mean intensity) as well as the Michelson contrast ($(I_{\max} - I_{\min}) / (I_{\max} + I_{\min})$) were identical.

Projector delay estimation. Since we are interested in accurate temporal characterization of sensory stimulus discrimination, it is important to quantify the delay between the onset command from the stimulus system and the actual onset of the stimulus. The largest contribution to this delay is expected to be from the projector. We measured the projector delay by attaching a photodiode to the stimulus screen, and subsequently averaged the signal with respect to the stimulus trigger. The patches were $24.8 \times 24.8^\circ$ squares, presented 100 times each for 1 s, with an inter-stimulus interval (ISI) of 1 s. The MEG responses to these patches were averaged with respect to the stimulus trigger. To confirm that the decoding accuracy trace follows the onset of the stimulus faithfully, we attempted to distinguish between a black and a white stimulus patch by decoding the photodiode signal.

Experiment 1: static gratings with different orientation and spatial frequency. The aim of Experiment 1 was to identify the OR and SF of large circular sinusoidal gratings (i.e., to distinguish between gratings with one of two ORs and SFs) from single-trial MEG responses. Eight volunteers participated in this experiment (four females, four males, mean age 25.5 years, range 21–39; Subjects 1–8). Each grating (cf. Fig. 1*A*) was presented 100 times for 1 s. We presented the gratings in a random order with an ISI of 1 s.

Experiment 2: dynamic gratings with different rotation directions. In this experiment, we attempted to decode the direction of grating rotation




Figure 1. All stimuli used in the study were sinusoidal gratings within annuli spanning the full visual field at eccentricities between 2° and 10°. *A*, In the static-grating experiment, stimuli were of two distinct SFs (left, 0.33 c/deg; right, 1.33 c/deg) and two distinct ORs (top, 45°; bottom, 135°). *B*, In the rotating-grating experiment, gratings with an SF of 1.33 c/deg rotated either clockwise or anticlockwise in the range of 0–180°. *C*, In the cross-contrast decoding experiment, static gratings oriented at 135° were presented at full, half, and one-fourth contrasts. *D*, In the cross-phase decoding experiment, static gratings, oriented at 135°, were presented at zero, quarter-cycle, and half-cycle phase shifts.

(clockwise vs anticlockwise). Eight volunteers (four females, four males, mean age 27.6 years, range 25–39; Subjects 2–9) participated in this experiment. We presented dynamic gratings with an SF of 1.33 c/deg (Fig. 1*B*). Gratings rotated for 180° in 15 discrete frames, each 40 ms in duration. The first five frames (0–200 ms) represented a fade-in period: a horizontally oriented grating presented at monotonically increasing contrast (0, 0.2, 0.4, 0.6, and 0.8), and the following 15 frames (200–800 ms) consisted of gratings with uniformly increasing (or decreasing) OR, such that the sixth and the fifteenth frame were horizontal (corresponding to 0° to 180°). The last five frames (800–1000 ms) were used for fade out: the horizontal grating decreased in contrast (0.8, 0.6, 0.4, 0.2, and 0). The stimuli were constructed as videos presented at 25 frames per second. Each video was presented 100 times in a random order for 1 s, with an ISI of 1 s.

Control experiment 1: spatial frequency coding as a function of contrast. To eliminate the possibility that the decoder is mainly influenced by evoked responses to local contrasts rather than by spatial frequency, we measured the response to both SF gratings (0.33 and 1.33 c/deg) at three different contrasts (full, half, and one-fourth), pooled together the trials at all contrasts except the target contrast, and attempted to decode the SF at the target contrast (Fig. 1*C*). One male volunteer (aged 21 years) participated in this experiment. We presented static gratings oriented at 135° for 1 s with an ISI of 1 s (Fig. 1*C*). Each grating was shown 50 times at each contrast in a random order.

Control experiment 2: spatial frequency coding as a function of phase. To eliminate the possibility that the decoder is dominated by a response component specific to the phase of the grating, or by local contrasts at certain eccentricities, we varied the phase of each SF grating and attempted to decode the SF of a target phase stimulus based on the single-trial evoked responses to the other phase stimuli (Fig. 1*D*). One male volunteer (aged 21 years) participated in this experiment. As in the previous control experiment, we presented dynamic gratings oriented at 135° for 1 s with an ISI of 1 s. The gratings had a spatial frequency of either 0.33 c/deg or 1.33 c/deg. Each grating was presented at one of three different phase shifts (30°, 120°, or 210°; Fig. 1*D*) with respect to the gratings presented in the main experiments. Each grating was displayed 50 times at each phase in a random order.

Apparatus

All stimuli were presented with a Panasonic D7700 DLP projector on a back-projection screen placed 120 cm in front of the subject. The resulting viewing angle was 34.7° horizontally and 25.8° vertically. All stimuli were displayed with a size of $24.8 \times 24.8^\circ$ of visual angle, with the gratings themselves spanning eccentricities from 2 to 10°. Dynamic stimuli were presented at a frame-rate of 25 frames per second.

Cortical responses were recorded with a 306-channel neuromagnetometer (Vectorview; Elekta Oy) in a magnetically shielded room (Imedco AG) at the Brain Research Unit of the O.V. Lounasmaa Laboratory, Aalto University. The MEG signal was bandpass filtered through 0–330 Hz and sampled at 1000 Hz. At the beginning of each measurement, we acquired the head position of the subject relative to the MEG sensors using four small head position indicator (HPI) coils attached to the scalp. The coils were briefly energized to emit single frequencies between 300 and 330 Hz.

Concurrently with MEG data, we acquired eye movement data using an SR Eyelink 1000 (SR Research) infrared eye-tracking system (sampled at 1000 Hz) to ensure that our subjects fixated on the center of the screen within an interval of 200 ms before and after each stimulus was shown. Each trial started with the presentation of a fixation cross. For the detection of fixations, we used the algorithm provided by SR Research. The on-line parser analyzed the sample data in search for saccades; everything that was not detected as a saccade was marked as a fixation (or blink, if no sample data were available). Saccades were identified by deflections in eye position exceeding 0.1° , with a minimum velocity of 30° s^{-1} and a minimum acceleration of $8000^\circ \text{ s}^{-2}$, maintained for at least 4 ms. If the initial fixation at the fixation cross did not occur correctly (i.e., was <200 ms or >100 pixels away from the cross), the stimulus was not displayed. In general, subjects did not have any difficulty fixating at the center of the screen.

Data analysis

Decoding the photodiode trace. To estimate the delay of the projector with respect to the onset command from the stimulus program, we presented either black or white patches, and recorded the luminance of the stimulus screen using a photodiode. The photodiode signal was averaged across these trials with respect to the rising trigger corresponding to the stimulus onset command. Since the onset of the averaged signal lagged behind the trigger onset consistently by 36 ms (see Results, Projector delay and effects of filtering on latency), we used this constant delay during latency estimation. We then decoded the color of the stimulus patch from the photodiode signal in a time-resolved manner, with a 20 ms window shifted forward by 1 ms, using a linear support vector machine (SVM) classifier.

Preprocessing. The MEG data were preprocessed using temporal Signal Space Separation (Taulu and Simola, 2006) and then bandpass filtered between 0.1–40 Hz with a causal, second-order infinite impulse response, approximately linear-phase Butterworth filter. We compensated for the 36 ms lag of the stimulus onset with respect to the trigger signal from the stimulus system. Evoked responses from 300 ms preceding the stimulus onset up to 260 ms post offset (i.e., 1260 ms post onset) were extracted from the continuous data using the trigger signals and then baseline-corrected using a time window from -300 to -100 ms with respect to the trigger signal associated with the stimulus onset.

Analysis of filtering effects on latency estimation. To study the effects of filtering, if any, on the decoding latency and accuracy, we performed time-resolved decoding of SF on a single subject's MEG responses on both filtered and unfiltered data.

Feature extraction and classification. Once the data had been extracted into a 3D matrix of trials \times time points \times channels, we reshaped this matrix into a 2D form, gathering the spatiotemporal data from selected channels into a single feature vector corresponding to each trial. We selected only a subset of 40 channels from the parieto-occipital regions labeled "Occipital" by the MEG vendor.

For all decoding problems, we used linear SVM implemented in the MATLAB Bioinformatics toolbox (<http://www.mathworks.com/products/bioinfo>). Multiclass decoding problems were addressed using the voting method, i.e., selecting the majority class predicted by all possible pairwise SVMs.

From the extracted MEG signal features, we performed subjectwise decoding as follows. First, we shuffled the order of the trials to avoid any learning or adaptation effects. For time-resolved classifiers, we repeated this shuffling separately for each time window so that accidental biases resulting from the shuffling procedure would not affect all time windows. Second, we partitioned the data into five nonoverlapping sets, with each set containing an equal number of trials per class. Third, we performed

fivefold cross-validation, i.e., we trained classifiers on all but one partition and tested them on the remaining partition. We considered the mean accuracy across the folds to represent the generalization accuracy. Fourth, to estimate confidence intervals on the test-set accuracies, we used a Monte Carlo bootstrapping approach as follows. From the predictions on the test samples, we drew 400 predictions with replacement, computed the mean accuracy on those predictions, and repeated this 999 times. Across these repetitions, we reported the 2.5th and 97.5th percentile scores as an estimate of the 95% confidence interval.

Time-insensitive decoding. For the static-gratings experiment, we attempted to solve a four-class decoding problem, i.e., to identify both the OR (45° or 135°) and the SF (0.33 c/deg or 1.33 c/deg) from a single trial. Since the SVM is inherently a binary classifier, we adopted a standard multiclass SVM strategy, where we combined the outputs of pairwise two-class SVMs by a majority-voting method. To obtain error estimates on the classification accuracy, we performed fivefold cross-validation. From the predicted ORs and SFs, we computed confusion matrices. The entry at row i and column j of the confusion matrix represents the percentage of samples belonging to class i , labeled as class j by the classifier. Thus, the diagonal entries of the confusion matrix represent the classwise accuracy, and the off-diagonals represent misclassification. In addition to the four-class decoding problem, we also attempted to decode SF separately for each OR, and OR separately for each SF. Likewise, for the rotating gratings experiment, we attempted to decode the direction of rotation as clockwise or anticlockwise.

Time-resolved decoding using moving and growing windows. To understand which temporal features of the evoked responses were important for each decoding problem, we built two types of time-resolved pattern classifiers defined on windowed segments of the single-trial evoked response.

First, we built moving window classifiers using responses from 20 ms windows, which were gradually shifted in 1 ms steps over the duration of the evoked response from 300 ms before stimulus onset up to 260 ms after stimulus offset. Second, we built growing window classifiers, by starting with a 20 ms window ranging from -300 to -281 ms, and gradually increased the window size in 1 ms steps, up to 1260 ms after stimulus onset. For the static stimuli, we repeated the SF decoding for each OR, and the OR decoding for each SF separately and subsequently averaged these accuracies.

Statistical testing of decoding onset and offset times. To estimate the time instant at which decoding exceeds chance level, we adopted the following procedure. We set the empirical chance-level threshold as two SDs above the mean decoding accuracy in the baseline range of -300 to 0 ms. Although the features of the stimulus were unpredictable a priori, the rationale for using an empirical threshold, rather than a theoretical baseline of 50%, was to account for possible biases resulting from the small training set size.

For each time window, we performed a one-tailed t test (to the right) to ascertain when the decoding exceeded chance-level threshold (40 accuracies: 8 subjects \times 5 cross-validation folds). For the SF and OR decoding traces, we estimated these latencies both at stimulus onset and offset. We set a strict significance threshold of $\alpha = 0.00005$, recognizing that the neighboring time windows are not independent. In particular, given that a single time point is common to maximally 20 time windows, a conservative Bonferroni correction for multiple comparisons would still keep the corrected significance level at $\alpha = 0.001$.

To estimate confidence bounds on the latencies, we applied Monte Carlo bootstrapping. We sampled the 40 decoding accuracies 999 times with replacement, and applied the one-tailed t test as described above. We then reported their median and 95% confidence bounds.

To test the hypothesis that SF-specific information is available both earlier and for longer than OR-specific information in single-trial responses, we constructed a bootstrap distribution of 999 samples for each corresponding latency, and then compared these distributions using a nonparametric Kolmogorov–Smirnov test for equality of distributions.

Robustness analysis. To rule out that the decoding results obtained were due to specific biases in training or test sets resulting from the selection of the train, test split ratio of 80:20, we varied the ratio of

training to test set sizes from 5:95 to 80:20, in steps of 5%. For each split ratio, we randomized the training and test set samples 10 times to obtain error estimates on the classification accuracies. We then reported the mean and SD of the mean classification accuracies across the eight subjects.

Intersubject classification. Despite known differences in subject-specific responses due to differences in cortical folding, head shape, size, and position, we attempted to investigate whether a decoder that learned from seven subjects could predict OR or SF from the MEG responses of the eighth subject. To minimize intersubject variability, we aligned head positions using a method based on signal space separation (Taulu and Simola, 2006) by translating single-subject data to the average head position across all subjects and measurements. Once normalized, we performed a leave-one-subject-out test, where an SVM classifier learned from trials of all but one subject, and was subsequently tested on the held-out subject. We applied such a classifier to each two-class problem viz. static OR (for each SF), static SF (for each OR), and RD.

Cross-contrast and cross-phase decoding. For the control experiment, we trained classifiers to identify the SF of the grating by pooling training data across two contrasts (or phases) and testing on the third contrast (or phase, respectively).

Results

Projector delay and effects of filtering on latency

To accurately quantify the latencies at which the stimulus discrimination begins, the delay between the stimulus delivery command and the actual delivery of the stimulus must be accurately estimated and accounted for. We measured the projector delay by averaging the photodiode signal with respect to the trigger of the stimulus system. We found that the stimulus onset lagged behind the trigger by 36 ms (Fig. 2A).

In addition, we decoded the color of the stimulus patch (black vs white) from single trials of the photodiode signal using a time-resolved decoder. The timing of decoding accuracy onset agreed with the photodiode signal onset time (Fig. 2B).

The inset in Figure 2A shows that the rise time of our projector from baseline to full luminance was 2 ms (from 36 to 38 ms after the stimulus trigger). Since visual perception likely begins before the stimuli reach full luminance, we applied the conservative correction of 36 ms to all latencies.

To assess the effects of filtering on the decoding accuracy, we decoded spatial frequencies from one subject's filtered and unfiltered data in a time-resolved manner. Figure 2C shows the decoding accuracy traces obtained from filtered and unfiltered data. As the filtering did not introduce any delay, we chose not to apply any latency correction to the filtered data during subsequent analysis.

Time-insensitive decoding of visual features

To examine whether single-channel evoked responses showed differences between different stimulus conditions, we visually in-

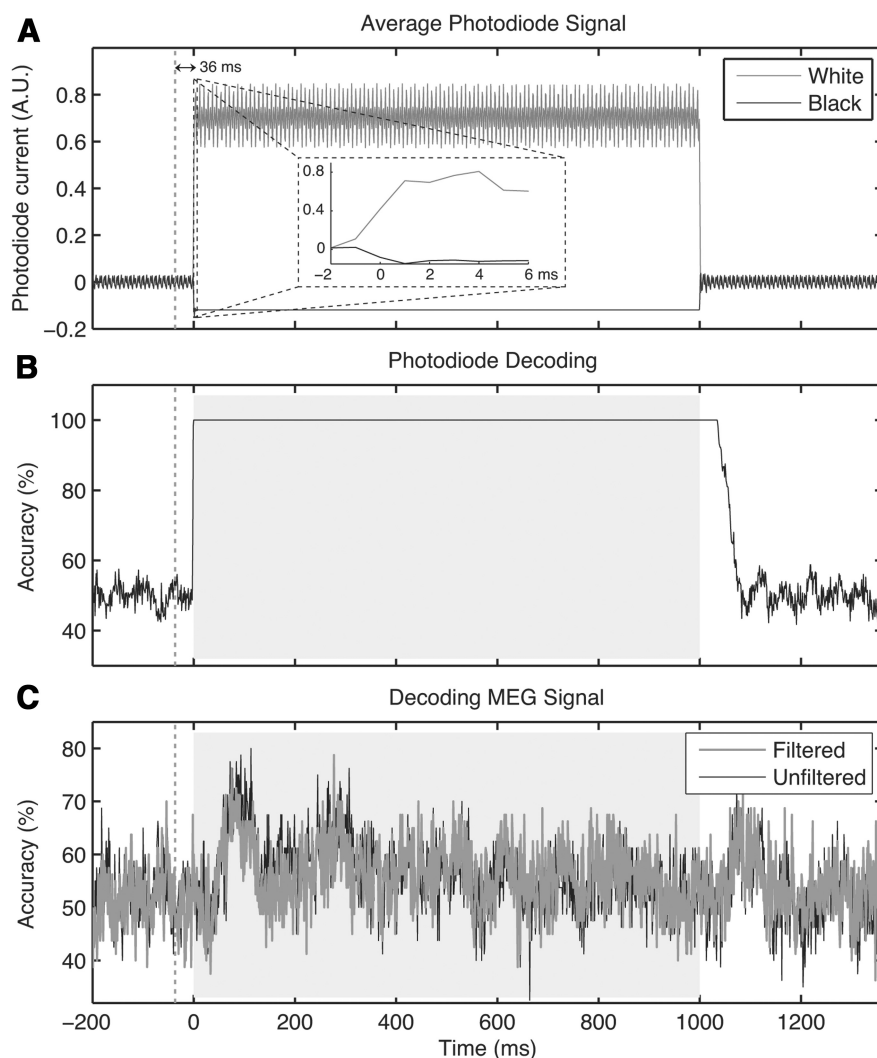


Figure 2. *A*, The average photodiode responses to 100 trials each of black (black trace) and white (gray trace) stimulus patches were averaged with respect to the stimulus trigger (gray vertical dashed line). The photodiode signal begins to change 36 ms after the stimulus trigger, as indicated by the arrow on the top left. The dotted lines indicate the region near the photodiode signal onset, from which the signal in the inset is displayed. The inset shows the 2 ms rise time to maximum luminance. *B*, Time-resolved decoding of the stimulus patch color from the single-trial photodiode responses. *C*, Time-resolved decoding of SF from a single subject's single-trial filtered (gray) and unfiltered (black) MEG responses.

spected the trial-averaged evoked responses. Figure 3 shows representative evoked responses in a single subject at a parieto-occipital planar gradiometer, averaged across 100 trials for each condition. The responses are shown for two spatial frequencies at a given orientation (Fig. 3A) and for two orientations at a given spatial frequency (Fig. 3B). Although the differences are most prominent at the main peaks of the evoked responses at ~ 70 and ~ 100 ms, the responses start to differ already at ~ 50 ms. It is important to note that our multivariate classifier used information from several channels and was thus more sensitive than any single channel to the subtle differences between the responses to different stimulus categories. We quantified the earliest peak latency of the evoked response by noting the time stamp of the maximum response amplitude within 100 ms after stimulus onset. Pooled across subjects, channels, and static gratings, we estimated this peak latency to be 71.4 ± 4.4 ms.

To examine whether single-trial responses contain sufficient information to decode low-level visual features, we presented large sinusoidal gratings for 1 s (Fig. 1) and applied decoders to the entire

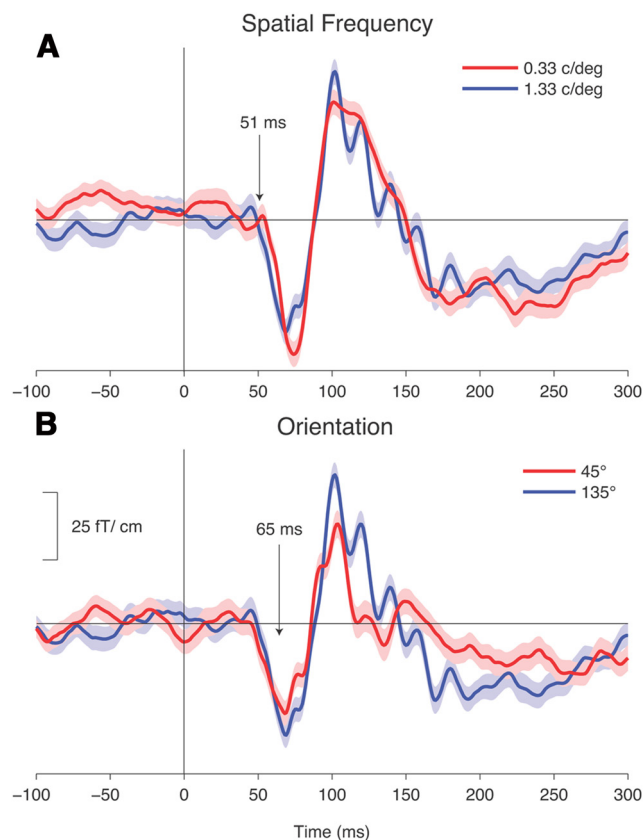


Figure 3. Evoked responses to static gratings. Responses averaged across 100 trials for a representative subject's parieto-occipital planar gradiometer. **A**, Low (0.33 c/deg; red) and high (1.33 c/deg; blue) spatial frequencies for gratings oriented at 135°. **B**, Gratings oriented at 45° (red) and 135° (blue), both at 1.33 c/deg. Shaded boundaries show SEMs. Arrows show the median onset of above-chance decoding accuracy (see Results, Time-resolved decoding of visual features and Fig. 5).

epoch from stimulus onset up to 260 ms after stimulus offset, from a selection of parieto-occipital planar gradiometer channels.

In Experiment 1 (see Materials and Methods, Experiment 1: static gratings with different orientation and spatial frequency), we attempted to simultaneously decode OR and SF of static sinusoidal gratings from unaveraged responses (Fig. 1A), representing two ORs and two SFs (chance level 25%). We used 80% of the trials for training an SVM classifier and the remaining 20% of the trials for predicting the OR and SF. We repeated this across five random cross-validation folds to estimate the variability of the classifier. Figure 4A shows the mean classification accuracies and their bootstrapped 95% confidence intervals across five cross-validation folds for all eight subjects. The accuracy across subjects and validation folds was on average 64.2%.

Figure 4B shows the mean OR classification accuracy as a function of SF classification accuracy for each subject, along with 95% confidence intervals (CI) bootstrapped across the cross-validation folds.

In general, the decoding accuracy was a third better for SFs than ORs (mean and 95% CIs: 91.3% and [88.5%, 93.7%], vs 66.8% and [62.3%, 71.3%]; $p < 0.0001$).

Figure 4C shows the confusion matrix averaged across subjects, with each column representing one predicted category and each row representing one actual category, summarizing how the test samples were classified or misclassified. As indicated by the dark blue entries, ORs tend to get confused with each other more

often than do SFs. To test the hypothesis that SFs can be decoded better for one OR versus another, we performed a two-sided Wilcoxon rank-sum test to compare the medians of the mean SF decoding accuracies for each OR, across the five cross-validation folds and eight subjects. The result did not reach the threshold of statistical significance ($p = 0.35$).

In Experiment 2 (see Materials and Methods, Experiment 2: dynamic gratings with different rotation directions), we attempted to decode the RD (clockwise vs anticlockwise; range 0–180°; Fig. 1B) from single-trial responses to dynamic gratings. The RDs were classified well above chance (mean and 95% CIs: 81.8% and [76.1%, 86.4%]).

Time-resolved decoding of visual features

Having established that information about low-level visual features can be decoded from single trials using the entire response epoch, we investigated the role of temporal features. To this end, we performed time-resolved decoding from signals in 20 ms windows that were shifted across the epoch in 1 ms steps (moving windows) or in windows growing forward in time by 1 ms steps (growing windows), from 300 ms before the stimulus onset up to 260 ms after the stimulus offset.

Figure 5 shows the accuracies averaged across eight subjects for classifiers trained on moving (left) and growing (right) windows. The time points—corresponding to leading edges of windows at which above-chance decoding performance was obtained—are indicated by black bars below the accuracy traces in Figure 5A,C,E.

For the SF and OR decoders, Figure 5, A and C, show that the accuracy (1) quickly exceeds chance level after stimulus onset; (2) reaches a peak ~70 and 90 ms, respectively; (3) drops down to near chance level after the stimulus onset-locked transient evoked responses; and (4) shows another transient increase and fall after the responses to stimulus offset, with peaks at ~1110 and 1120 ms.

Table 1 lists the median onset and offset latencies of successful decoding following onsets and offsets of the stimuli, along with 95% CIs estimated using Monte Carlo bootstrapping.

In the following, we report median estimates of decoding onset and offset latencies at stimulus onset and offset.

The SF decoder gave significantly above-chance performance ($p < 0.00005$, uncorrected) as early as 51 ms; i.e., already in the 32–51 ms window (Fig. 5A, inset), and the above-chance performance continued until the 243–262 ms post onset window. In comparison, the OR decoder gave above-chance performance ($p < 0.00005$, uncorrected) as early as 65 ms (46–65 ms window onwards; Fig. 5C, inset) and remained above chance until the 174–193 ms window.

Likewise, for the offset response, SF-specific information was available for ~160 ms, between the 1041–1060 ms and 1200–1219 ms windows, and OR-specific information was available for only ~15 ms from the 1079–1098 ms to the 1104–1123 ms windows.

Together, these latency estimates suggest that SF-specific information was available earlier and lasted longer than OR-specific information during both stimulus onset and offset. To test this hypothesis statistically, we estimated distribution of latencies using bootstrapping and then compared the distributions of decoding onset and offset latencies for SF and OR. The corresponding SF and OR latencies were found to be statistically significantly different ($p < 0.00001$; Kolmogorov–Smirnov test).

Finally, RD-specific information was available from the 279–298 ms post onset window (Fig. 5E, inset), or at least as early as 98

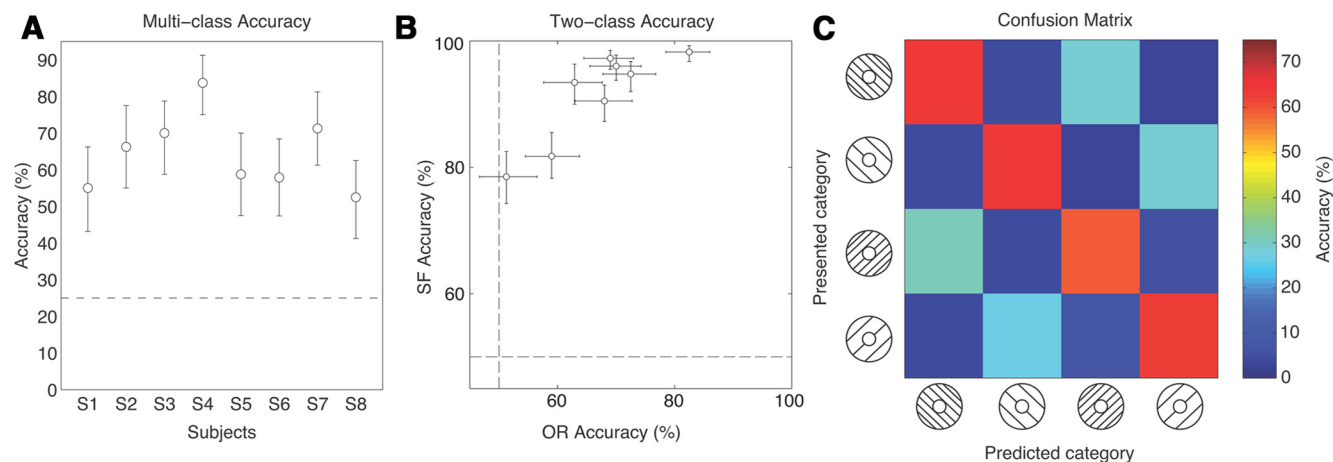


Figure 4. Time-insensitive decoding. **A**, Subjectwise classification accuracies for the four-class decoding problem: identification of both SF and OR from a closed set. The chance level (25%) is given as a dashed line. Error bars indicate bootstrapped 95% CIs. **B**, Decoding accuracies for the two-class problem: OR (x-axis) versus SF (y-axis). Each circle represents one subject; error bars indicate bootstrapped 95% CIs. **C**, Confusion matrix for the four-class decoding problem, viz. prediction of OR and SF. Individual confusion matrices estimated from each subject and cross-validation fold were averaged. Entries in each row show percentage of trials predicted as the category corresponding to the respective column. Chance level is 25%. The categories—indicated by schematics depicting the stimulus features—correspond to left-oriented (135°) high SF, left-oriented low SF, right-oriented (45°) high SF, and right-oriented low SF.

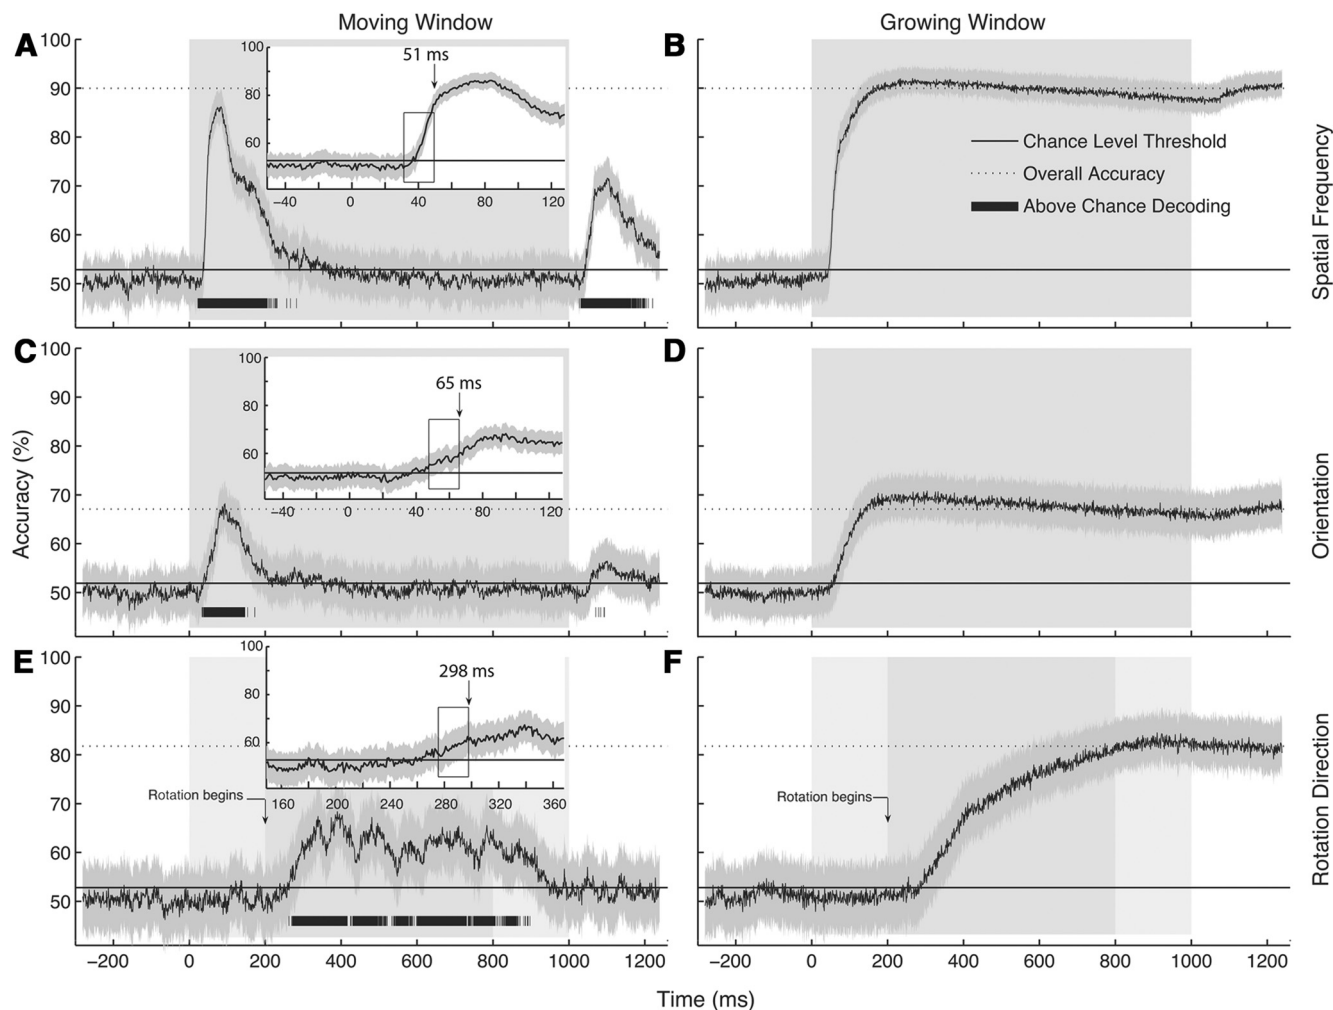


Figure 5. Time-resolved decoding. Performance of time-resolved classifiers using windows moving or growing in 1 ms time steps is shown. The moving and growing window traces for SF (**A**, **B**), OR (**C**, **D**), and RD (**E**, **F**) are shown, along with chance-level threshold as solid lines and overall accuracies from time-insensitive classifiers as dotted lines. The bounds on each trace indicate bootstrapped 95% CIs across eight subjects. Subjectwise accuracy traces were themselves obtained by averaging across five cross-validation folds. Black bars at the bottom of each trace show the duration in which above-chance decoding ($p < 0.00005$) was obtained. Insets in **A**, **C**, and **E** show the accuracy traces around the stimulus or rotation onset. Arrows within the insets indicate the leading edge of the windows at which chance level was first exceeded. The gray patches in **A–D** show the stimulus time course. The light gray patches for **E** and **F** between 0–200 ms and 800–1000 ms indicate the periods of contrast fade-in and fade-out of the grating, and the dark gray patches indicate the duration of rotation. The arrows in **E** and **F** show the onset of stimulus rotation.

Table 1. Periods of successful decoding

Stimulus feature	Stimulus onset at 0 ms		Stimulus offset at 1000 ms	
	Decoding onset	Decoding offset	Decoding onset	Decoding offset
SF	51 (50–53)	262 (247–287)	1060 (1057–1067)	1219 (1202–1239)
OR	65 (58–72)	193 (174–261)	1098 (1090–1106)	1123 (1115–1179)
RD	298 (284–313)			841 (772–863)

The numbers refer to latencies (in milliseconds) at which decoding emerged above (decoding onset) and returned back to (decoding offset) chance level at stimulus onset and offset. The values were obtained using a one-tailed *t* test against an empirical baseline estimated from the [−300 ms, 0] interval (see Material and Methods). Each entry gives the median and the 95% CI from the bootstrap distribution.

Table 2. Classification accuracies for intersubject analysis

Subject	Spatial frequency		Orientation		Rotation direction
	135°	45°	1.33°	0.33°	
1	70.5	74	53.5	53.5	—
2	75	70.5	54	53	56.5
3	66.5	73	64	58.5	61
4	75	68	45.5	55	56
5	75.5	78	53.5	50	49
6	69	64	59	53	56.5
7	74.7	69.5	53.7	55.8	54
8	82.5	79	55.5	57	53
9	—	—	—	—	58
Mean ± SD	73.6 ± 4.9	72.0 ± 5.1	54.2 ± 5.0	54.5 ± 2.7	55.5 ± 3.6

Each row corresponds to the decoding accuracies on one subject's data obtained with a classifier trained on data from all other subjects. The columns give the classification accuracies (percentage of trials correctly classified) for each stimulus feature (two-class problem). Chance level is 50%.

ms after the onset of rotation that occurred after 200 ms. This information continued to be available up to 822–841 ms after stimulus onset. No RD-specific information was available after stimulus offset.

Intersubject classification

Next, we investigated intersubject differences in single-trial responses by attempting to predict for each subject the OR and SF of static gratings, as well as the direction of rotation of the rotating gratings by training a classifier on the remaining seven subjects (see Materials and Methods, Intersubject classification). Table 2 shows the individual classification accuracies for SF (for each OR), OR (for each SF), and RD. As expected due to intersubject variability, mean accuracies are worse than those obtained when decoders were trained and tested separately on data from individual subjects (see above, Time-insensitive decoding of visual features). Nevertheless, SFs were classifiable above chance (73.6 ± 4.9 (SD) for gratings oriented at 135° and 72.0 ± 5.1 (SD) for gratings oriented at 45°).

Robustness analysis

In supervised machine learning, with a very small test-set size, it is possible to overestimate the generalization performance of the classifier. To study how many training samples (each sample comprising 40 parieto-occipital planar gradiometer signals of 1.2 s duration) were needed for robust learning, we attempted to classify SF (separately for each OR) and OR (separately for each SF) by varying the proportion of samples in the training set from 5% to 80% in steps of 5%.

Figure 6 shows the classification accuracies for the SF (A), OR (B), and RD (C) as a function of the training set proportion. As expected, the decoding performance gradually increases with increasing training-set proportion. Above-chance accuracies were obtained already with 5% of the data for SF, OR, and RD.




Figure 6. Robustness analysis. Classification performance for the static and dynamic gratings as a function of the percentage of samples in the training set are shown. Accuracies are shown as a function of the number of training samples used to train the classifier. The proportion of samples in the training set was varied from 5 to 80%, and the remaining trials were used for testing. Accuracies are shown for (A) SFs, (B) ORs, and (C) RDs. Error bars indicate SD of the mean classification accuracies across the eight subjects.

Table 3. Classification accuracies for cross-contrast and cross-phase decoding

Target contrast	Accuracy	Target phase	Accuracy
1	75	0	81
Half	79	¼	73
One-fourth	78	½	71

A spatial frequency decoder was trained on responses to gratings with contrasts (phase) other than the target contrast (phase) and tested on responses to gratings with the target contrast (phase).

Cross-contrast and phase decoding

Finally, we studied in one subject whether the representation of SF is invariant to contrast and phase of the stimulus by training and testing cross-contrast and cross-phase SF decoders, respectively. Specifically, a pattern classifier was trained on responses to two of three possible contrasts (or phases) and tested on the responses to the left-out contrast (or phase). Table 3 shows that the obtained accuracies were above chance for all possible combinations of training and test sets, and that they were comparable to the SF decoding accuracies for a single contrast and phase (see above, Time-insensitive decoding of visual features). Therefore, our main results cannot be attributed to responses to local contrast edges.

Discussion

We obtained four main results from the multivariate decoding analysis of single-trial MEG signals. First, we demonstrated that three low-level visual features are encoded robustly in single-trial MEG responses. Second, we showed that these visual features could be successfully decoded from 20 ms epochs; the information about visual features is available earlier than evoked-response peaks, and information about SFs is available earlier and for a longer duration than ORs. Third, we demonstrated that despite intersubject differences, it was possible to predict (above chance level) the SF from the data of any of the subjects using classifiers trained on the data of the seven other subjects, suggesting a highly robust representation of SF across subjects. Last, a control experiment confirmed that SF information can be decoded regardless of contrast or phase of the stimulus.

Representation of visual information in large neuronal populations

Tuning properties of single neurons in V1 to characteristic features of visual stimuli are well established (Hubel and Wiesel, 1963). Several studies have also investigated the temporal characteristics of these tuning properties (Ringach et al., 1997; Mazer et al., 2002). However, until very recently (Graf et al., 2011; Berens et al., 2012), the representation of stimulus features in populations comprising tens of neurons had not been experi-

mentally demonstrated. Our study complements these recent reports by demonstrating visual feature representation at the very different spatial scale of tens of thousands of neurons, in humans.

Evidence for rapid visual information processing

Over the past decade, several studies have provided neural evidence regarding the speed of information processing in the visual system (Thorpe et al., 1996; VanRullen and Thorpe, 2001; Bisley et al., 2004; Hung et al., 2005; Liu et al., 2009; Thorpe, 2009).

Judging from spike-onset latencies, layer 4C of macaque V1 receives the earliest input at 44 ms after stimulus onset (Schmolesky et al., 1998). Furthermore, a recent time-resolved decoding study of spike counts from a macaque V1 population revealed that orientation can be read out at latencies as early as 30 ms (Berens et al., 2012). Our results show that in the human brain visual features can be decoded above chance as early as 50 ms, i.e., around the time when the main visually evoked MEG responses of human visual cortex typically start to emerge (Vanni et al., 2001). By showing that stimulus features can be reliably decoded from the cortex as early as ~50 ms, our study complements this body of work, and implies that the human visual system has inferred low-level visual features at or before this time. Although a good correspondence exists between the timing of single-neuron firing in primate V1 and the onset of MEG signal decoding in our study, we may be dealing with very different phenomena at different scales, and therefore our results must be interpreted with caution.

Time-resolved encoding and decoding models are readily applicable to study the dynamics of pre-attentive stimulus discrimination versus conscious decision making (Lamme and Roelfsema, 2000; VanRullen and Thorpe, 2001). However, since our subjects were not asked to perform an overt discrimination task, we cannot tell apart the influence of stimulus discrimination processes, variable signal-to-noise ratio (SNR) of the responses, and conscious perceptual judgment on the decoding results.

Dynamics of spatial frequency and orientation processing

V1 population firing rates show selectivity to different low-level features at different times (Lamme et al., 1999). Thus, the dynamics of population-level signals may carry important information about visual feature processing.

It is worth noting that in our data, as well as in the population firing-rate data presented by Lamme et al. (1999), stimulus-feature decoding for static stimuli was successful only during the transient onset and offset responses. In contrast, in recent attempts to decode orientation of gratings from MEG signals (Duncan et al., 2010; Koelewijn et al., 2011), gamma-band responses contained static stimulus-specific information from 200 ms onward. Hence in future studies, time-resolved decoding could be used to elucidate how stimulus-feature selectivity evolves as a function of time and what aspects of visual information are reflected in transient evoked responses versus longer lasting gamma-band responses.

Our estimates of decoding onset and offset latencies (see Results, Time-resolved decoding of visual features; Fig. 5A, C; Table 1) suggest that SF-specific information is present 1.6 and 6.4 times as long as OR-specific information at stimulus onset and offset, respectively. It is likely that the contrast between responses to the two ORs and to the two SFs might differ and emerge (above the noise level) at different times, and that these unequal contrast-to-noise ratios could explain the timing differences between SF and OR decoding.

However, these timing differences could be explained by two other possibilities, which are speculative at present. One possibil-

ity is that neurons tuned to SF and OR exhibit different rates of neuronal adaptation, suggested by monkey V1 recordings (McLelland et al., 2010). Another possibility is that broader temporal selectivity of SFs than ORs may be explained by latency differences between high and low SFs. In particular, the difference between SF and OR decoding might derive from different contributions from magnocellular (M) and parvocellular (P) inputs as the faster M pathway is known to be sensitive to lower SFs, and single cells in V1 receive projections from both pathways. Furthermore, Mazer et al. (2002) found significant latency differences in spikes for high and low SFs from single V1 neurons in behaving macaques, but no similar effect for orientation was reported.

Why are spatial frequencies and orientations classifiable?

Our results with static and dynamic gratings suggest that SF- as well as OR-selective activity is robustly observable from single-trial MEG data. It is well known from early electrophysiology and optical imaging studies that the primary visual cortex consists of orientation hypercolumns that are organized in a pinwheel-like structure (Hubel and Wiesel, 1963), or SF-selective cells that are arranged in alternating bands (Tootell et al., 1981; Sirovich and Uglesich, 2004). Since it is unlikely that our sensor-level MEG signals would reflect activity due to fine columnar structure, our data could be taken to suggest that several columns sensitive to a certain SF or OR in our relatively large stimuli were activated in a very synchronized manner, thereby producing distinct spatio-temporal fingerprints for different stimulus features.

The nonuniform distribution of SF-selective neurons with respect to eccentricity (Tootell et al., 1981), or differential SF-tuning in different visual cortical areas (Henriksson et al., 2008), may give rise to these discriminative fingerprints. Similarly, a recent fMRI study (Freeman et al., 2011) suggested that groups of V1 neurons are tuned to radial orientations, in addition to having a pinwheel-like columnar structure (Hubel and Wiesel, 1963). Given that our large gratings span a wide range of eccentricities (from 2 to 10°), it is possible that our MEG data contain such orientation-specific fingerprints. However, measurements of responses to a wider range of orientations as well as explicit source modeling are required to confirm these interpretations.

Decoding challenges for MEG versus fMRI

Despite the complementary nature of EEG/MEG signals with respect to fMRI, few attempts have been made to decode low-level stimulus-specific information from EEG/MEG signals. Duncan et al. (2010) applied MVPA to single-trial MEG signals, both evoked responses from 0 to 300 ms poststimulus and frequency spectra from 300 to 2300 ms poststimulus, to decode at ~70% accuracy between two oblique orientations of gratings presented for 2.5 s to either the left or right visual field. Koelewijn et al. (2011) found significant differences in the evoked MEG responses and grating-induced MEG gamma responses between oblique versus cardinal gratings, suggesting that noninvasive measurements contain discriminative information about the orientation of gratings. Although recent decoding studies have shown that the gamma band is informative of task or stimulus parameters (Fuentemilla et al., 2010; Polanía et al., 2012; Jafarpour et al., 2013), we restricted our analysis to the evoked responses and low-pass filtered the data at 45 Hz in this study.

It is important to point out one key difference between MEG and fMRI with respect to decoding: unlike fMRI, which offers relatively uniform SNR throughout the brain, the sensitivity of MEG to different parts of the cortex is highly nonuniform be-

cause of MEG's differential depth (better sensitivity for superficial than deep sources) and orientation sensitivity (better sensitivity for tangential vs radial currents with respect to the skull; source current orientation in the tangential plane directly affects the orientation of the field pattern). As an example of this effect, even a small change in current orientation may be detected with MEG (Hari et al., 1996). Thus, these nonuniform sensitivities of MEG might prove to be assets rather than shortcomings.

In summary, the implications of this study for noninvasive electrophysiology and low-level vision are as follows: The differences between the temporal accuracy traces of the time-resolved classifiers suggest interesting differences between the processing of spatial frequency and orientation of visual stimuli. The decodability of visual information nearly 20 ms before the peak of the cortical visual-evoked response was unexpected and certainly merits further investigation. Although earlier studies have reported time-resolved decoding of object-level information from electrophysiological data (Hung et al., 2005; Liu et al., 2009; Carlson et al., 2011), our decoding study is the first to present a detailed characterization of the rapid neural processing of visual features using a noninvasive technique in humans. Analogous to information-based brain mapping with fMRI (Kriegeskorte et al., 2006), the time-resolved decoding approach presented here is suitable to unravel, in addition to the type of information processed in the brain, the exact temporal structure and dynamics of the processing.

References

- Berens P, Ecker AS, Cotton RJ, Ma WJ, Bethge M, Tolias AS (2012) A fast and simple population code for orientation in primate V1. *J Neurosci* 32:10618–10626. [CrossRef Medline](#)
- Bisley JW, Krishna BS, Goldberg ME (2004) A rapid and precise on-response in posterior parietal cortex. *J Neurosci* 24:1833–1838. [CrossRef Medline](#)
- Bonhoeffer T, Grinvald A (1991) Iso-orientation domains in cat visual cortex are arranged in *pinwheel*-like patterns. *Nature* 353:429–431. [CrossRef Medline](#)
- Carlson TA, Hogendoorn H, Kanai R, Turret J (2011) High temporal resolution decoding of object position and category. *J Vis* 11(10): pii: 9. [CrossRef Medline](#)
- Chalupa LM, Werner JS (2003) *The visual neurosciences*. Cambridge, MA: MIT.
- Di Russo F, Martínez A, Sereno MI, Pitzalis S, Hillyard SA (2002) Cortical sources of the early components of the visual evoked potential. *Hum Brain Mapp* 15:95–111. [CrossRef Medline](#)
- Duncan KK, Hadjipapas A, Li S, Kourtzi Z, Bagshaw A, Barnes G (2010) Identifying spatially overlapping local cortical networks with MEG. *Hum Brain Mapp* 31:1003–1016. [Medline](#)
- Freeman J, Brouwer GJ, Heeger DJ, Merriam EP (2011) Orientation decoding depends on maps, not columns. *J Neurosci* 31:4792–4804. [CrossRef Medline](#)
- Fuentemilla L, Penny WD, Cashdollar N, Bunzeck N, Düzel E (2010) Theta-coupled periodic replay in working memory. *Curr Biol* 20:606–612. [CrossRef Medline](#)
- Graf AB, Kohn A, Jazayeri M, Movshon JA (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat Neurosci* 14:239–245. [CrossRef Medline](#)
- Hari R, Nagamine T, Nishitani N, Mikuni N, Sato T, Tarkiainen A, Shibasaki H (1996) Time-varying activation of different cytoarchitectonic areas of the human SI cortex after tibial nerve stimulation. *Neuroimage* 4:111–118. [CrossRef Medline](#)
- Hari R, Parkkonen L, Nangini C (2010) The brain in time: Insights from neuromagnetic recordings. *Ann NY Acad Sci* 1191:89–109.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691. [CrossRef Medline](#)
- Henriksson L, Nurminen L, Hyvärinen A, Vanni S (2008) Spatial frequency tuning in human retinotopic visual areas. *J Vis* 8(10):5. 1–13. [CrossRef Medline](#)
- Hubel DH, Wiesel TN (1963) Shape and arrangement of columns in cat's striate cortex. *J Physiol* 165:559–568. [Medline](#)
- Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. *Science* 310:863–866. [CrossRef Medline](#)
- Jafarpour A, Horner AJ, Fuentemilla L, Penny WD, Düzel E (2013) Decoding oscillatory representations and mechanisms in memory. *Neuropsychologia* 51:772–780. [CrossRef Medline](#)
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685. [CrossRef Medline](#)
- Koelewijn L, Dumont JR, Muthukumaraswamy SD, Rich AN, Singh KD (2011) Induced and evoked neural correlates of orientation selectivity in human visual cortex. *Neuroimage* 54:2983–2993. [CrossRef Medline](#)
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868. [CrossRef Medline](#)
- Lamme VA, Roelfsema PR (2000) The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci* 23:571–579. [CrossRef Medline](#)
- Lamme VA, Rodriguez-Rodriguez V, Spekreijse H (1999) Separate processing dynamics for texture elements, boundaries and surfaces in primary visual cortex. *Cereb Cortex* 9:406–413. [CrossRef Medline](#)
- Liu H, Agam Y, Madsen JR, Kreiman G (2009) Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex. *Neuron* 62:281–290. [CrossRef Medline](#)
- Mazer JA, Vinje WE, McDermott J, Schiller PH, Gallant JL (2002) Spatial frequency and orientation tuning dynamics in area V1. *Proc Natl Acad Sci U S A* 99:1645–1650. [CrossRef Medline](#)
- McLelland D, Baker PM, Ahmed B, Bair W (2010) Neuronal responses during and after the presentation of static visual stimuli in macaque primary visual cortex. *J Neurosci* 30:12619–12631. [CrossRef Medline](#)
- Polanía R, Paulus W, Nitsche MA (2012) Noninvasively decoding the contents of visual working memory in the human prefrontal cortex with high-gamma oscillatory patterns. *J Cogn Neurosci* 24:304–314. [Medline](#)
- Ringach DL, Hawken MJ, Shapley R (1997) Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387:281–284. [CrossRef Medline](#)
- Schmolesky MT, Wang Y, Hanes DP, Thompson KG, Leutgeb S, Schall JD, Leventhal AG (1998) Signal timing across the macaque visual system. *J Neurophysiol* 79:3272–3278. [Medline](#)
- Sirovich L, Uglyesich R (2004) The organization of orientation and spatial frequency in the primary visual cortex. *Proc Natl Acad Sci U S A* 101:16941–16946. [CrossRef Medline](#)
- Taulu S, Simola J (2006) Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Phys Med Biol* 51:1759–1768. [CrossRef Medline](#)
- Thorpe SJ (2009) The speed of categorization in the human visual system. *Neuron* 62:168–170. [CrossRef Medline](#)
- Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. *Nature* 381:520–522. [CrossRef Medline](#)
- Tootell RB, Silverman MS, De Valois RL (1981) Spatial frequency columns in primary visual cortex. *Science* 214:813–815. [CrossRef Medline](#)
- Vanni S, Tanskanen T, Seppä M, Uutela K, Hari R (2001) Coinciding early activation of human primary visual cortex and anteromedial cuneus. *Proc Natl Acad Sci U S A* 98:2776–2780. [CrossRef Medline](#)
- VanRullen R, Thorpe SJ (2001) The time course of visual information processing: from early perception to decision making. *J Cogn Neurosci* 13:454–461. [CrossRef Medline](#)