

Retinotopic Patterns of Correlated Fluctuations in Visual Cortex Reflect the Dynamics of Spontaneous Perceptual Suppression

Tobias H. Donner,^{1,2,3} Dov Sagi,⁴ Yoram S. Bonneh,^{4,5} and David J. Heeger²

¹Department of Psychology, University of Amsterdam, 1018 XA Amsterdam, The Netherlands, ²Department of Psychology and Center for Neural Science, New York University, New York, New York 10003, ³Bernstein Center for Computational Neuroscience, Charité-Universitätsmedizin, 10115 Berlin, Germany, ⁴Department of Neurobiology, Brain Research, The Weizmann Institute of Science, Rehovot 76100 Israel, and ⁵Department of Human Biology, University of Haifa, Mount Carmel, Haifa 31905, Israel

While viewing certain stimuli, perception changes spontaneously in the face of constant input. For example, during “motion-induced blindness” (MIB), a small salient target spontaneously disappears and reappears when surrounded by a moving mask. Models of such bistable perceptual phenomena posit spontaneous fluctuations in neuronal activity throughout multiple stages of the visual cortical hierarchy. We used fMRI to link correlated activity fluctuations across human visual cortical areas V1 through V4 to the dynamics (rate and duration) of MIB target disappearance. We computed the correlations between the time series of fMRI activity in multiple retinotopic subregions corresponding to MIB target and mask. Linear decomposition of the matrix of temporal correlations revealed spatial patterns of activity fluctuations, regardless of whether or not these were time-locked to behavioral reports of target disappearance. The spatial pattern that dominated the activity fluctuations during MIB was spatially nonspecific, shared by all subregions, but did not reflect the dynamics of perception. By contrast, the fluctuations associated with the rate of MIB disappearance were retinotopically specific for the target subregion in V4, and the fluctuations associated with the duration of MIB disappearance states were target-specific in V1. Target-specific fluctuations in V1 have not previously been identified by averaging activity time-locked to behavioral reports of MIB disappearance. Our results suggest that different levels of the visual cortical hierarchy shape the dynamics of perception via distinct mechanisms, which are evident in distinct spatial patterns of spontaneous cortical activity fluctuations.

Introduction

Perception sometimes changes spontaneously between distinct interpretations of the same sensory input (Blake and Logothetis, 2002; Deco and Romo, 2008). For example, during “motion-induced blindness” (MIB), a salient visual target spontaneously disappears and reappears when surrounded by a moving mask (Fig. 1A) (Bonneh et al., 2001; Donner et al., 2008; Bonneh and Donner, 2011). Models of such “bistable” perceptual phenomena posit that the dynamics of perception are caused by intrinsic (independent of stimulus changes) fluctuations in neuronal activity across multiple stages of the visual cortical hierarchy (Blake and Logothetis, 2002; Wilson, 2003; Freeman, 2005; Kim et al., 2006; Moreno-Bote et al., 2007; Noest et al., 2007; Gigante et al., 2009). These fluctuations, in turn, are produced by the interplay

of adaption within specialized populations of neurons, interactions between these populations, and “noise.” According to these models, MIB occurs because of fluctuations in the activity of retinotopically organized neuronal populations processing the target across visual cortical areas.

It has long been known that neuronal activity in visual cortex fluctuates continuously in the face of constant sensory input (Heggelund and Albus, 1978; Arieli et al., 1996; Fiser et al., 2004; Fox and Raichle, 2007; Faisal et al., 2008; Ringach, 2009). These fluctuations are correlated across visual cortical areas at a “global” spatial scale (Leopold et al., 2003; Nir et al., 2006; Yeo et al., 2011). However, optical imaging in monkeys (Tsodyks et al., 1999; Kenet et al., 2003) and fMRI in humans (Heinzle et al., 2011) revealed also spatially specific fluctuations that reflect the retinotopically organized intrinsic anatomical connectivity of cortical networks.

The goal of the present study was to determine which spatial patterns of cortical fluctuations, the global or the retinotopically specific ones, are linked to the dynamics of bistable perception. To identify these patterns, we applied an alternative analysis approach to fMRI data recorded during MIB. Most previous studies of perceptual bistability used subjects’ behavioral reports of their perceptual changes to quantify the underlying fluctuations of activity in visual cortex (Leopold and Logothetis, 1996; Polonsky et al., 2000; Blake and Logothetis, 2002; Tong et al., 2006; Donner

Received July 16, 2012; revised Nov. 15, 2012; accepted Dec. 7, 2012.

Author contributions: T.H.D., D.S., Y.S.B., and D.J.H. designed research; T.H.D. performed research; T.H.D. and D.J.H. contributed unpublished reagents/analytic tools; T.H.D. analyzed data; T.H.D. and D.J.H. wrote the paper.

This research was supported by the German National Academy of Sciences Leopoldina (BMBF-LPD 9901/8-136, T.H.D.), the NIH (R01-EY016752, D.J.H.), the US-Israel Binational Science Foundation (2007224, D.J.H. and D.S.), and the Weizmann-NYU Demonstration Fund in Neuroscience (D.J.H. and D.S.). We thank Christopher Honey, Tomas Knapen, David A. Leopold, and Markus Siegel for comments.

Correspondence should be addressed to Tobias H. Donner, Department of Psychology, University of Amsterdam, Weesperplein 4, 1018 XA Amsterdam, The Netherlands, E-mail: T.H.Donner@uva.nl.

DOI:10.1523/JNEUROSCI.3388-12.2013

Copyright © 2013 the authors 0270-6474/13/332188-11\$15.00/0

et al., 2008; Sterzer et al., 2009). However, there is not necessarily an instantaneous, and one-to-one, mapping between neuronal activity in visual cortex and perceptual reports (Lamme, 2003). During bistable perception, some activity fluctuations may directly culminate in a reported perceptual transition, whereas others might lay the foundation for these transitions, or affect the stability of perceptual states, without being closely coupled to the reports (Lumer and Rees, 1999; Lee et al., 2007). By exploiting the intrinsic temporal correlations between fMRI time series, our analysis quantified ongoing cortical fluctuations, regardless of whether or not they were time-locked to behavioral reports of target disappearance. The analysis revealed spatial patterns of correlated fluctuations across multiple retinotopic subregions, corresponding to target and mask, in visual cortical areas V1 through V4. The spatial patterns whose fluctuations reflected the rate and duration of MIB both exhibited retinotopic specificity for the target subregion, but in distinct cortical areas.

Materials and Methods

The current study is based on a reanalysis of fMRI data from a previously published study of MIB (Donner et al., 2008). In the following, we focus on the procedures unique to the present reanalysis of those data.

Subjects and scanning sessions. Data were acquired from six healthy subjects with normal or corrected-to-normal vision (one female, age range: 25–35 years). The experiments were conducted with the written consent of each subject and in accordance with the safety guidelines for fMRI research, as approved by the University Committee on Activities Involving Human Subjects at New York University.

Each subject participated in multiple scanning sessions before the main experiments: one to obtain a high-resolution anatomical volume for cortical surface extraction, one to define retinotopically organized visual cortical areas, and one to identify the subregions of these areas corresponding to the target and mask locations (for details, see ROIs, below).

Experimental protocols. Subjects participated in several sessions of the main experiments involving three experimental conditions: MIB, fixation, and periodic stimulation. All six subjects participated in the MIB experiments (between 19 and 27 runs). Three of the subjects also participated in the fixation (between 10 and 15 runs) and periodic stimulation experiments (between two and four runs). Subjects were supine and viewed the screen through an angled mirror. The scanner room was dark.

During MIB, subjects fixated a central cross and reported the disappearance and reappearance of a static yellow contour (the “target”) surrounded by a moving blue random dot pattern (the “mask”) (Fig. 1A). Target and mask were separated by a black “protection zone.” The target was presented in one of the four visual field quadrants, which was selected individually for each subject to maximize the percentage of viewing time that the target was invisible. The size of the target was then adjusted to be as large as possible while maintaining a minimum MIB strength (20% of viewing time invisible). Each run lasted 180 s. Subjects reported their perception of the target by depressing a button with the middle (visible) or index finger (invisible) of the right hand. Thus, subjects continuously depressed a button with the same hand throughout each run.

During fixation runs, subjects fixated a central cross, which was presented on an otherwise black screen. Each run lasted 180 s.

During periodic stimulation, subjects fixated a central cross and viewed a block-alternation (15 cycles of 16.8 s, i.e., 0.0595 Hz) between two complementary diagonal arrangements of probe stimuli. Each of these probes was centered in one visual field quadrant, and each had the same spatial configuration and eccentricity as the MIB target, but flickered at 8 Hz. Each run lasted 252 s.

Perceptual dynamics. The perceptual dynamics of MIB target disappearance were quantified, separately for each run, by means of two different measures: rate and duration. MIB rate was computed as the number of target disappearances per minute. MIB duration was computed as the median duration of all target disappearances across a run. The median, rather than the mean, was used as a measure of central

tendency because the distributions of perceptual state durations were skewed (Donner et al., 2008, their Fig. 1). Both parameters (rate, duration) were used for a different median split of MIB runs: high- versus low-rate and long- versus short-duration. Each median split was done separately for each subject.

MRI data acquisition. Data were acquired with a 3 T Allegra MRI scanner (Siemens Medical Systems) equipped with a transmit head coil (NM-011) and a four-channel phased-array receive surface coil (NMSC-021; both Nova Medical Inc.) positioned at the back of the head. We measured blood oxygenation level-dependent (BOLD) changes in MRI signal intensity using an echoplanar pulse sequence with the following parameters: repetition time, TR = 1.2 s; echo time, TE = 30 ms; flip angle, 72°; 64 × 64 matrix; voxel size, 3 × 3 × 3 mm; 22 slices oriented roughly perpendicular to the calcarine sulcus, covering the occipital lobe and part of the temporal and parietal lobes. In retinotopic mapping sessions we used the same imaging parameters with the following exceptions: TR = 1.5 s; flip angle, 75°; 27 slices. At the beginning of each session, we acquired an anatomical T1-weighted MPRAGE volume in the same slices as the functional volumes, but with twice the in-plane resolution (voxel size, 1.5 × 1.5 × 3 mm).

fMRI preprocessing. Data from the beginning of each functional run were discarded (10 volumes from each run of MIB and fixation; 14 volumes from each run of the retinotopic mapping and periodic stimulus alternation experiments) to minimize the effect of transient magnetic saturation, and to allow the hemodynamic response to reach steady-state baseline. We compensated for head movements within and across runs with standard procedures (Jenkinson et al., 2002). Within each region of interest (ROI; see ROIs below) and each run, the time series of each voxel was converted from arbitrary intensity units to percentage modulation (to compensate for distance from the coil and partial volume effects). We then averaged the time series across voxels in each ROI, high-pass filtered the average ROI time series (cutoff: 0.01 Hz), and finally converted the average ROI time series into units of SD (i.e., z-scores).

The anatomical volume from each session was aligned to a high-resolution anatomical volume (acquired in a different scanning session) by an automated robust image registration algorithm (Nestares and Heeger, 2000). The resulting alignment parameters were used to resample the functional data from the localizer runs of each scanning session to the image space of the high-resolution anatomy. Cortical surfaces were extracted from the high-resolution anatomy to visualize and define the ROIs on computationally flattened representations of the occipital cortex. The alignment parameters were also used to transform the ROIs from the coordinates of the high-resolution anatomical volume to the image space of each session. This enabled us to extract the mean time series from each ROI for each run in each scanning session.

ROIs. Retinotopic maps in human visual cortex were measured using standard procedures (Wandell et al., 2007) to identify visual field maps corresponding to V1, V2, V3, and V4 in each individual subject following the conventions of Wandell et al. (Wandell et al., 2007).

In a separate session, subjects completed eight to 10 runs of the periodic stimulation experiment described above. We used these data to identify the subregions V1–V4 that corresponded to the target in MIB experiments (Fig. 1B), as well as these subregions’ retinotopic complements for the other three visual field quadrants. To minimize any possible coregistration errors, we repeated between two and five of these target localizer runs within each session of the MIB experiments and further restricted the subregions (for details, see Donner et al., 2008).

We also performed three to five repeats of the periodic stimulation experiment in the fixation experiment sessions. These data constituted the periodic stimulation condition used for validating our analysis approach (see Fig. 5), but they were not used for restricting the ROIs.

To identify cortical subregions responsive to the moving mask, subjects completed two to four runs in which the mask pattern alternated periodically (15 cycles of 16.8 s) between moving and static dots. This allowed us to define the subregions corresponding to the mask in each of V1–V4. The target subregion was excluded from the mask subregion in each visual area (for details, see Donner et al., 2008).

For the MIB experiment, two subregions per cortical area (target and mask subregions) were analyzed, yielding eight ROIs in total. For the

periodic stimulation and fixation experiments, four subregions per cortical area (corresponding to the four visual field quadrants) were analyzed, yielding 16 ROIs in total.

Spatial modes of activity fluctuations. A multivariate linear decomposition, principal component analysis (PCA), identified spatial patterns (“spatial modes”) across visual area ROIs, for which the activity covaried over time. We implemented PCA by means of singular value decomposition (SVD) (Mitra and Pesaran, 1999; Friston and Büchel, 2004). Our analysis approach was motivated by the notion that cortical fMRI responses reflect an approximately linear mixture of multiple components (Cardoso et al., 2012), with possibly distinct underlying electrophysiological signals (Maier et al., 2008; Cardoso et al., 2012). We chose PCA rather than other linear decompositions (such as independent component analysis; Cole et al., 2010) because it generalized for comparing between different experimental conditions described below. For the analysis of the MIB data, we focused on the first spatial mode, for two reasons: (1) during MIB, the first spatial mode accounted for by far the greatest proportion of the variance in the data than the remaining modes (the eigenvalue for the first mode was more than six times larger than the eigenvalue for the second mode; Fig. 1E); and (2) this simplified the interpretation because the orthogonality of the spatial modes can obscure the interpretation of the higher modes.

For each run i , the ROI time series were combined into a single matrix M_i , with one row for each successive time point (TR) and one column for each ROI. For the main analyses, M_i was 140 TRs \times 8 ROIs (target and mask subregions for each of V1–V4, see ROIs above). We computed the covariance between the time series from each pair of ROIs (i.e., the columns of M_i) according to:

$$C_i = M_i^T M_i \quad (1)$$

where C_i was an 8×8 covariance matrix for run i . Because the variance of each z-scored ROI time series was 1, the covariance matrices were identical to the correlation matrices. The matrices C_i were then averaged across runs, separately for each subject. The resulting mean covariance matrices were used for two separate stages of analysis (two variants of PCA).

Spatial modes of the activity fluctuations within one condition. In the first stage of the analysis, we applied SVD to the covariance matrices from individual experimental conditions (e.g., all MIB runs), to decompose these matrices into multiple orthogonal spatial modes:

$$C = V \lambda V^T, \quad (2)$$

where C was an 8×8 covariance matrix (averaged across runs), with one element per ROI-pair, λ was an 8×8 matrix containing the eigenvalues of C on its diagonal (sorted in decreasing order), and V was an 8×8 matrix with the orthonormal eigenvectors of C in its columns. Superscript T denotes matrix transposition. Each of the eigenvectors (i.e., each column of V) defined a spatial mode (i.e., a pattern of correlated activity fluctuations across the 8 ROIs). We call each of the elements of the spatial modes a “spatial mode weight.” The first column of V was the first spatial mode and corresponded to the largest eigenvalue. The analysis was repeated separately for each individual subject.

The same V (i.e., the same set of spatial modes) could have also been obtained by applying SVD to the matrix of the time series data, M_i , rather than their covariance matrix. We describe it in terms of the decomposition of the covariance matrices, because this generalizes for comparing between different experimental conditions (see Spatial modes of the ratios of the activity fluctuations in two conditions, below).

We calculated the time series corresponding to a given spatial mode p , separately for each run i :

$$t_i = M_i p, \quad (3)$$

where vector t_i was a time series (140 TRs), M_i was a data matrix (140 TRs \times 8 ROIs), and p was a spatial mode (i.e., one of the columns of V). The fluctuation amplitude corresponding to a given spatial mode was quantified, separately for each run, as follows:

$$s_i = t_i^T t_i = p^T C_i p, \quad (4)$$

where s_i was the square of the scalar L^2 -norm of vector t_i . To convert the fluctuation amplitude into units of percentage variance accounted for, we normalized s_i by the total variance in C_i (i.e., the sum of its eigenvalues).

Spatial modes of the ratios of the activity fluctuations in two conditions. The second stage of the analysis identified spatial modes that accounted for most variance during one condition relative to another condition. We formalized this as a solution to the “generalized eigenvalue problem” (Friston and Büchel, 2004), finding a vector (spatial mode) p_1 that satisfied the following equation and maximized λ_1 :

$$C_1 p_1 = \lambda_1 C_2 p_1, \quad (5)$$

where C_1 and C_2 were the mean 8×8 covariance matrices (as defined above) for two conditions (e.g., high-rate and low-rate MIB). In other words, the spatial mode p_1 maximized the ratio of the variances it accounted for in C_1 and C_2 . We identified p_1 by decomposing the matrix $C_2^{-1} C_1$:

$$C_2^{-1} C_1 = V \lambda V^T, \quad (6)$$

where λ was an 8×8 matrix containing the eigenvalues of $C_2^{-1} C_1$ on its diagonal, and V was an 8×8 matrix with the eigenvectors of $C_2^{-1} C_1$ in its columns. If C_1 and C_2 were identical, then $C_2^{-1} C_1 = I$ (their product equaled the identity matrix) and $\lambda = I$. If C_1 was a copy of C_2 scaled up by factor k , then $C_2^{-1} C_1 = \lambda = kI$. In all realistic scenarios, the eigenvalues on the diagonal of λ differed, indicating that the spatial modes in the columns of V accounted for different proportions of the variance in C_1 relative to C_2 . The first column of V corresponding to the largest eigenvalue λ_1 was the first spatial mode p_1 (see Eq. 5). We performed two comparisons: high- versus low-rate MIB (based on target disappearance rate) and long- versus short-duration MIB (based on median target disappearance duration).

Averaging and visualizing spatial modes. For visualizing the spatial modes, we averaged them across subjects and reorganized them into matrices with dimensions “cortical area” \times “subregion.” The overall sign of the spatial modes (i.e., eigenvectors) is arbitrary. Thus, the spatial modes of different subjects might have opposite signs and cancel each other out during averaging. We implemented the following procedure to check for such sign flips. Each subject’s spatial mode was projected onto a “reference” vector, which was the average of the “raw” (i.e., unaligned) individual spatial modes. If there had been a flip in the sign of any of the individual spatial modes, its projection would have been negative. In that case, the sign of the spatial mode was flipped back, multiplying its elements by -1 . After aligning the individual spatial modes in that way (if necessary), the spatial modes were again averaged together. While this procedure enabled alignment of the spatial modes in case of sign flips, it was not needed for the main results of this article, since there were (by chance) no sign flips for any of the individual spatial modes featured here.

Characterizing temporal fluctuations of spatial modes. To test for differences in the strength of the fluctuations of the spatial modes between conditions (see Figs. 2, 4), we first quantified the fluctuation strength in terms of percentage variance explained, within each individual run. We then compared the fluctuation strengths between conditions by means of a two-sided permutation test (Efron and Tibshirani, 1998), separately for each individual subject. In another analysis (see Fig. 5E), we calculated the time series t_i corresponding to each spatial mode (Eq. 3), and characterized their temporal structure by means of spectral analysis (Welch’s averaged periodogram method). For each of a number of successive time windows (overlap: 95%), the fast Fourier transform was computed with 256 points (spectral resolution of 0.0033 Hz). We calculated the power spectra from 0.0065 Hz to 0.4167 Hz (Nyquist frequency at 1.2 s TR). The resulting power spectra were averaged across runs and subjects.

Testing the overall retinotopic specificity of spatial modes (retinotopic specificity index). A scalar “retinotopic specificity index” quantified the spatial modes’ overall retinotopic structure, across the four cortical visual areas. This index was computed based on the group average spatial modes shown in Figures 1, 3, and 5. The spatial mode weights were extracted from each retinotopic subregion. We then averaged the spatial

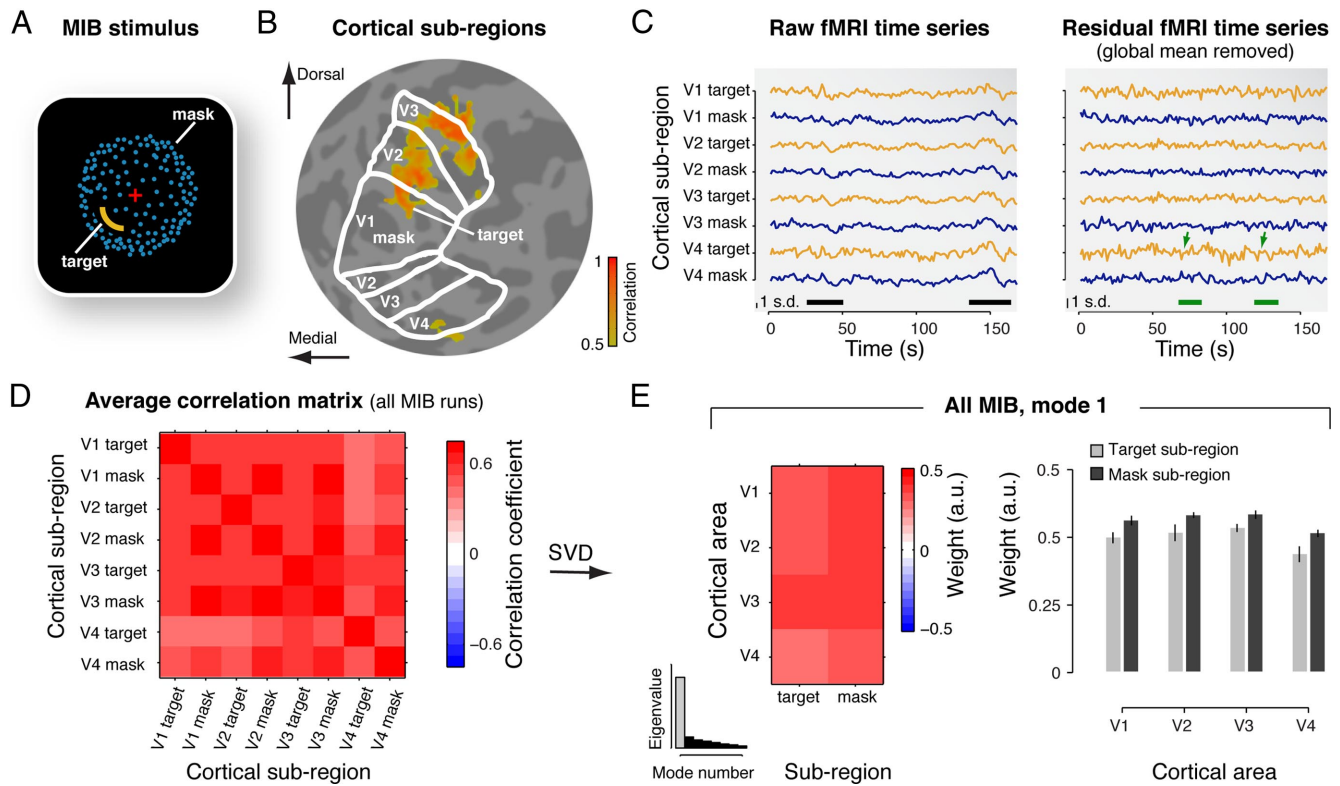


Figure 1. Correlation structure of cortical activity during MIB. **A**, Snapshot of the dynamic MIB stimulus. A salient yellow target was surrounded by a moving dot pattern (blue), which appeared as a rotating sphere. The target was presented in different visual field quadrants for different subjects. While viewing this stimulus, subjects repeatedly experienced and reported the spontaneous disappearance and reappearance of the target. **B**, Cortical subregions. Separate localizer protocols delineated the subregions corresponding retinotopically to the MIB target and the surrounding mask, within each of V1, V2, V3, and V4 (highlighted for V1). The figure shows an example map of cortical responses to a flickering probe stimulus presented at the spatial location of the MIB target. Colors represent correlation between measured activity and stimulus alternations (threshold: $r > 0.5$). The map is superimposed on a flattened representation of the subject’s occipital lobe. The borders of areas V1, V2, V3, and V4 are indicated as white outlines. The probe in the lower left visual field quadrant evoked responses in the dorsal subregions of right hemisphere visual areas V1–V3, and in the corresponding subregion of area V4. **C**, Example fMRI time series during one run of MIB. Left, Raw fMRI time series of the eight cortical subregions (yellow traces, target subregions; blue traces, mask subregions). The black bars at the bottom mark epochs of strong global fluctuations shared by all subregions. Right, Residual time series after removing the global mean across all subregions. The green bars and arrows mark epochs of anti-correlated fluctuations in the target and mask subregions of V4 (compare traces for target and mask subregions). Vertical bars in bottom left corner of each panel, 1 SD of fMRI image intensity modulation. Note that the scales are different in the two panels. **D**, Correlation matrix showing pairwise correlation coefficients, for each pair of ROIs, averaged across all MIB runs and the six subjects. **E**, Spatial mode, averaged across subjects. Left, Map of the spatial mode weights. The spatial modes are vectors, whose elements are reorganized as matrices (4 cortical areas \times 2 subregions) for illustration. Right, Bar graph of the spatial mode weights. Error bars indicate SEM, $N = 6$. Inset, histogram of the eigenvalues. Gray bar, eigenvalue corresponding to the first spatial mode.

mode weights across all pairs of retinotopically corresponding subregions, and, separately, across all pairs of retinotopically noncorresponding subregions. Both averaging procedures comprised three subsequent steps: (1) averaging the spatial mode weights across the two subregions within each pair, (2) discarding the sign, and (3) averaging across pairs. Discarding the sign in step (2) accounted for the different signs of different subregion pairs (e.g., target and mask subregions). We then computed the difference between the averages for the pairs of corresponding subregions and for the pairs of noncorresponding subregions. Finally, this difference was normalized by the maximum absolute spatial mode weight of all subregions to express the retinotopic specificity index in units of percentage. We used a permutation procedure (Efron and Tibshirani, 1998) to assess the statistical significance of the index. Repeating the above procedure 1000 times after randomly permuting the labels of the subregions yielded a null distribution of indices, against which we compared the observed index to obtain a p -value.

Testing the target specificity in individual cortical areas. To assess the target-specificity separately for each visual area (V1 through V4), we computed the difference between the spatial mode weights of mask and target subregions within each area (Fig. 3). The mask-target differences were computed based on the spatial modes of individual subjects and tested against zero by means of two-sided t tests across the six subjects. The t tests were chosen for two reasons: (1) the quantification of retinotopic specificity should take into account both the sign and the magni-

tude of the difference between mask and target subregions; and (2) linear combinations of spatial mode weights are likely to be normally distributed. Indeed, χ^2 tests did not reject the hypothesis that the empirically observed differences came from a normal distribution.

Results

In a previous report, we quantified the mean response modulations, time-locked to behavioral reports of target disappearance and target reappearance, within individual visual cortical areas (Donner et al., 2008). Here, we instead used a multivariate linear decomposition technique to identify spatial patterns (spatial modes) across visual cortical areas V1–V4, for which the activity covaried during MIB, regardless of the timing of subjects’ behavioral reports (see Materials and Methods). We first defined two retinotopic subregions, corresponding to MIB target and mask, within each cortical area (Fig. 1A,B). We then correlated the fMRI time series corresponding to each pair of subregions (Fig. 1C, left), separately for each 3 min run of MIB (between 19 and 27 runs, depending on the subject), and averaged the resulting matrix of correlations across runs (Fig. 1D). In our analyses, the correlation matrices were identical to covariance matrices because the fMRI time series were z-scored. The correlation matri-

ces were then decomposed in two stages. In the first stage, we decomposed the average correlation matrices from all MIB runs, to identify the spatial modes that accounted for most of the variance (fluctuation) in the multivariate fMRI time series. Each spatial mode corresponded to a pattern of correlated activity fluctuations that were distributed across the cortical subregions. The spatial modes were vectors with one number per subregion. We call each of these numbers a “spatial mode weight” because it quantified the subregion’s contribution to the spatial pattern of activity fluctuations. In the second analysis stage, we decomposed the correlation matrix from half of the MIB runs (e.g., high-rate of MIB), multiplied with the inverse of the correlation matrix from the other half of the MIB runs (e.g., low-rate MIB). This identified spatial modes that accounted for most variance in the first condition relative to the second condition. If, for example, a spatial mode occurred in both conditions and fluctuated equally strongly in both, it was discarded by this analysis. If, by contrast, a spatial mode fluctuated more strongly in the first condition, it was revealed by the analysis. While the first analysis stage characterized the general spatial structure of correlated fluctuations in cortical activity (Fox and Raichle, 2007; Ringach, 2009; Cole et al., 2010), we found that the second stage was critical for linking these fluctuations to the dynamics of perception.

The Results are organized as follows. First, we characterize the overall correlation structure and the first spatial mode during MIB. Second, we present the spatial modes that specifically reflected the dynamics (rate and duration) of MIB target disappearance. Third, we validate our analysis approach by showing that it yielded the expected spatial modes in two conditions, for both of which we expected to find retinotopic specificity, based on previous studies using similar protocols (Ringach, 2009; Heinze et al., 2011): fixation of an otherwise blank screen, and periodic alternations between two complementary spatial stimulus configurations.

Nonspecific first spatial mode during MIB

Activity during MIB was dominated by a strong, global (retinotopically nonspecific) component, which was, however, not linked to the dynamics of MIB. The time series of the eight cortical subregions exhibited several epochs of high-amplitude fluctuations in activity, which were shared by all cortical subregions (Fig. 1C, left, marked by the black bars). Similarly strong and global fluctuations were found in all runs and subjects. Consequently, the average correlation matrix was largely homogenous, with similar correlation coefficients across all pairs of subregions (Fig. 1D).

Global fluctuations of high amplitude are commonly observed in fMRI measurements during “rest” (Cole et al., 2010; see also Spatial modes of activity during fixation and periodic visual stimulation below). Some portion of these fluctuations may be caused by non-neuronal physiological factors (e.g., respiration) (Birn et al., 2006; Bianciardi et al., 2009b). However, electrophysiological measurements of cortical population activity have shown fluctuations (of similar strength and global spatial structure) that accounted for the fMRI signal fluctuations (Leopold et

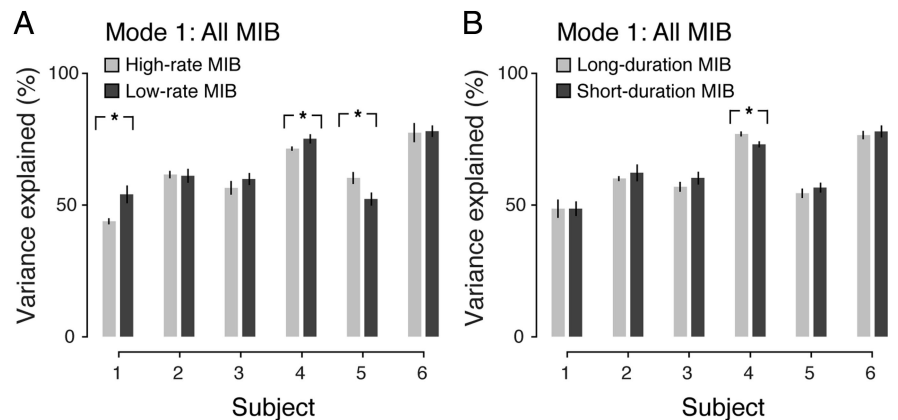


Figure 2. Percentage variance accounted for by first spatial mode during MIB. The first spatial mode extracted from all MIB runs (i.e., global component, Fig. 1E) was used to determine percentage variance accounted for in different subsets of MIB runs. **A**, Percentage variance accounted for in high- versus low-rate MIB runs. **B**, Percentage variance accounted for in long- versus short-duration MIB runs. Error bars indicate SEM across runs; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$, two-sided permutation tests.

al., 2003; Schölvinck et al., 2010). We, therefore, infer that the global signal fluctuations observed here also reflected, at least in part, neuronal activity.

The first spatial mode provided by the first stage of our linear decomposition analysis reflected the global component of the ongoing activity fluctuations evident in Figure 1C (left). The spatial mode weights of all subregions had the same sign (red in Fig. 1E), which accounted for by far the largest proportion of variance in the multivariate time series (Fig. 1E, compare first bin in inset histogram with the other bins). Note that the sign of a cortical subregion in the spatial mode (i.e., blue or red) is irrelevant; what matters, is whether different subregions have identical or different signs. While there was a trend for retinotopic specificity (i.e., different weights for target and mask subregions), this was not statistically significant. We tested the retinotopic specificity of the spatial mode in two complementary approaches (see Materials and Methods). First, we quantified the overall specificity (regardless of cortical area) in the group average spatial modes in terms of a scalar retinotopic specificity index. This index would have been larger than 0% if the spatial mode weights were more similar for retinotopically corresponding than for noncorresponding subregions. The index was not significantly different from that expected by chance ($p = 0.99$, permutation test, $N = 8$ subregions). Second, we tested the difference between the spatial mode weights of mask and target subregions, separately for each cortical area (V1 through V4). These were not significant in any of V1 through V4, albeit approaching significance for V4.

In each subject, the time series of the first spatial mode (see Materials and Methods, Eq. 3) was strongly correlated with the mean time series across all subregions in each individual subject ($p < 10^{-3}$, permutation test based on randomly shuffling individual runs). Thus, the first spatial mode of activity during MIB indeed closely reflected the global component of fluctuations evident in the raw fMRI time series.

The global component (i.e., the first spatial mode) accounted for a substantial fraction of the total variance in cortical activity (between 50% and ~80% depending on the subject), but it was not consistently linked to the perceptual dynamics of MIB (Fig. 2). We used two complementary parameters to summarize the dynamics of MIB: the rate of target disappearance events and duration of target invisible states (Table 1). In two subjects, the first spatial mode fluctuated more for low-rate than for high-rate MIB runs, whereas it showed the opposite effect in another subject, and no difference in

Table 1. Statistics of MIB dynamics

| Subject | 1 | 2 | 3 | 4 | 5 | 6 |
|---|----------------|-----------------|------------------|-----------------|------------------|-------------------|
| Disappearance rate (no. of reports min ⁻¹) | 3.7 (2.3, 6.0) | 9.3 (4.3, 12.0) | 12.3 (7.7, 15.3) | 9.7 (6.8, 10.7) | 14.0 (7.3, 14.9) | 21.7 (12.7, 22.3) |
| Typical disappearance duration (s) | 1.6 (0.9, 2.1) | 1.7 (1.0, 4.0) | 1.3 (0.6, 1.6) | 1.7 (1.1, 3.2) | 0.8 (0.6, 1.3) | 0.6 (0.5, 0.8) |
| Percentage of trials in same group for median splits based on MIB rate and duration | 64.3 | 63.2 | 71.4 | 25.9 | 41.7 | 39.1 |

Numbers in first row are medians (2.5, 97.5 percentiles) of disappearance rate across runs. Numbers in second row are medians (2.5, 97.5 percentiles) of typical disappearance duration across runs. Typical disappearance duration for each run is computed as the median duration of all individual target invisible states in that run.

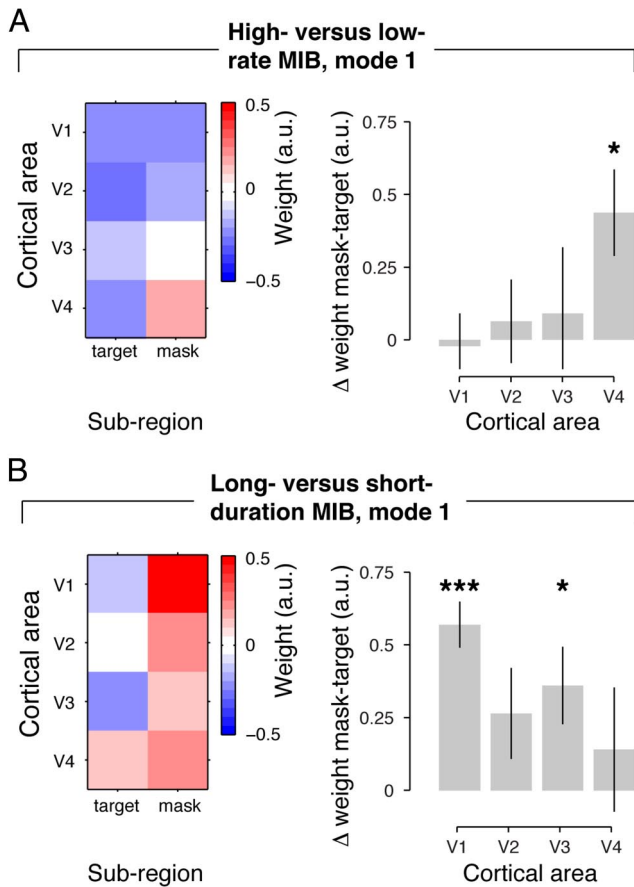


Figure 3. Spatial modes corresponding to differences in the dynamics of MIB. **A**, Difference between high- and low-rate MIB runs (median split based on disappearance rate). Left, Map of first spatial mode (same format as in Fig. 1E). Right, Retinotopic specificity for individual cortical areas. Retinotopic specificity was computed as the mean difference in spatial mode weights for mask and target subregions. **B**, Long- versus short-duration MIB (median split based on median duration of target disappearance). Same format as **A**. Error bars indicate SEM across $N = 6$ subjects; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$, two-sided t tests.

the remaining three (Fig. 2A). Combined across subjects, there was no evidence for any difference in fluctuation strength between high-rate and low-rate MIB runs ($p = 0.71$, two-sided permutation test). Likewise, the first spatial mode showed no difference in fluctuation strength between long- and short-duration MIB runs in five of the subjects (Fig. 2B) and combined across subjects ($p = 0.63$, two-sided permutation test).

Specific spatial modes reflecting differences in the dynamics of MIB target disappearance

In contrast to the global component, two retinotopically specific spatial modes exhibited fluctuations that were reliably associated with the rate or the duration of MIB target disappearance (Figs. 3, 4). These two spatial modes were revealed by the second stage of

our analysis. Although both contained some specificity for the target subregion, their detailed spatial structures were distinctly different from one another (Fig. 3). In the first spatial mode for high- versus low-rate MIB (Fig. 3A), the target-specificity was statistically significant only in V4 ($t_5 = 2.94$, $p < 0.05$), but not in V1 ($t_5 = -0.20$, $p = 0.85$) or in V2 ($t_5 = 0.44$, $p = 0.68$) or V3 ($t_5 = 0.40$, $p = 0.71$). By contrast, the first spatial mode for long- versus short-duration MIB (Fig. 3B) exhibited significant target-specificity in V1 ($t_5 = 7.19$, $p < 0.01$), but not in V4 ($t_5 = 0.66$, $p = 0.54$). Some degree of target-specificity was also evident in V2 and V3 (V2: $t_5 = 1.70$, $p = 0.14$; V3: $t_5 = 2.71$, $p < 0.05$; pooled across V1–V3: $t_5 = 5.41$, $p < 0.01$).

The opposite signs for target and mask subregions evident in the spatial modes are consistent with competitive neuronal interactions within cortical areas (Carandini and Heeger, 2012). Competitive interactions between nonoverlapping stimuli are well established for area V4 (Reynolds et al., 1999), and they may play a key role in the spontaneous perceptual alternations in phenomena like MIB (Blake and Logothetis, 2002). However, when discarding the global component, the opposite signs are, in principle, also consistent with no correlation between the target- and mask-specific subregions (Murphy et al., 2009; Cole et al., 2010).

These retinotopically specific spatial modes accounted for a smaller fraction of the variance than the global component, but they were reliably linked to perception (Fig. 4). The variance explained ranged from $\sim 15\%$ to $\sim 35\%$, and from $\sim 5\%$ to $\sim 55\%$, across subjects for the high- versus low-rate and the long- versus short-duration, respectively. Since several inevitable noise sources other than neuronal variability (e.g., scanner noise, residual head motion, vascular noise, respiration) contribute to the total variance of the fMRI signal, these numbers represent a lower bound for the true percentage of neuronal variability explained by these spatial modes. Importantly, the high- versus low rate spatial mode consistently accounted for a larger fraction of the variance during runs with high disappearance rate than runs with low disappearance rate (Fig. 4A). Conversely, the long- versus short-duration spatial mode consistently accounted for a larger fraction of the variance during runs with long disappearance duration than runs with short disappearance duration (Fig. 4B).

It is notable that target specificity was evident in the first spatial modes distinguishing between high- and low-rate MIB (Fig. 3A) and long- and short-duration MIB (Fig. 3B), despite the fact that the overall correlation structure was dominated by the global component (Fig. 1D). The correlation matrices for the separate median splits were also dominated by the global component (data not shown). Since the global component did not fluctuate more strongly during high- than low-rate MIB, nor during long- than short-duration MIB (Fig. 2), it was discarded by the second stage of our analysis. Indeed, removing a proxy of the global component (the mean across all eight time series) from the raw fMRI time series by linear regression (Donner et al., 2008) typically revealed epochs of salient anti-correlated fluctuations be-

tween the target and mask subregion in the residual time series (see area V4 in the right panel of Fig. 1C, marked with green bars and arrows). This highlights the importance of the second stage of our analysis for revealing spatial models associated with the dynamics of MIB.

The retinotopic specificity of V1 modulations during MIB has not yet been discovered based on conventional report-triggered averaging procedures (Donner et al., 2008; Hsieh and Tse, 2009; Schölvinck and Rees, 2010). Using report-triggered averaging, we previously identified a spatial pattern of cortical modulations during MIB (Donner et al., 2008), which closely resembled the high-versus low-rate spatial mode identified by the present analysis. During target disappearance reports, the fMRI responses decreased specifically in the target subregion of V4, followed (a few seconds later) by a retinotopically nonspecific decrease in activity across all of V1 and V2 (Donner et al., 2008). The similarity between the results from both analysis approaches may have reflected the fact that the “high-rate” runs contained many report-related transients, which consequently contributed strongly to the correlation structure. To test this idea, we repeated the present analysis on the residual fMRI time series after removing (by linear regression) the event-related modulations time-locked to reports (Donner et al., 2008), and found no evidence for retinotopic specificity (particularly not in V4) in the high- versus low-rate mode. The long- versus short-duration mode was, however, nearly unchanged and retained highly significant specificity in V1 ($t_5 = 4.99$, $p < 0.01$) after removing the report-related modulations. We conclude, therefore, that the high- versus low-rate mode, but not the long- versus short-duration mode, largely reflected fluctuations coupled to behavioral report.

There was no strong and systematic relationship between the rate and duration of target disappearance. The correlation between both parameters across the runs was not statistically significant in four of the six subjects, and the individual correlation coefficients ranged from negative to positive (mean: 0.14). If the two median splits had been orthogonal, the percentage of trials in the same (upper or lower) bin would have been 50%. By contrast, if the median splits had been perfectly correlated (anti-correlated), that percentage would have been 100% (0%). In line with the correlation coefficients, the actual percentages of overlap lay in between these two extremes, but generally closer to the “orthogonal scenario” (Table 1). When averaged across subjects, the percentage of trials in the same bin was 50.93%. Note that the rate of disappearance events and the duration of target-invisible states may become decoupled through a third parameter, the duration of target-visible states. A longer invisible duration on some runs does not necessarily lead to a lower disappearance rate: if the visible duration shortens correspondingly, the rate remains constant. Conversely, a higher disappearance rate on some runs does not necessarily lead to a shorter invisible duration: if the visible duration shortens correspondingly, the invisible duration remains constant. In sum, the two parameters used here, the rate and duration of MIB target disappearance, quantified two largely separate processes governing the perceptual dynamics.

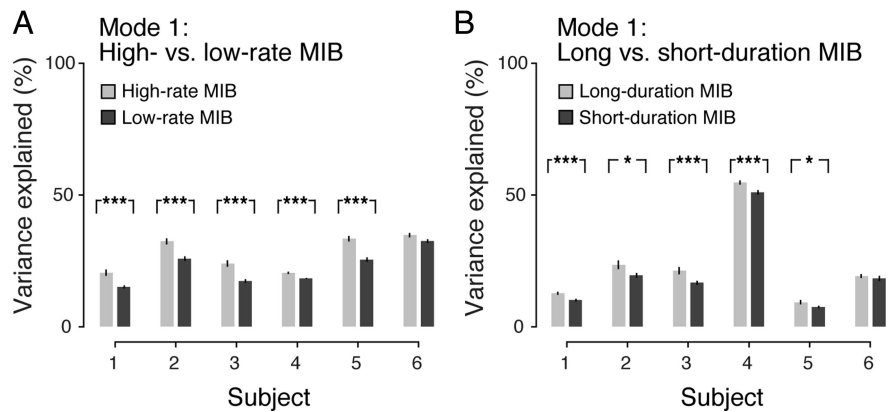


Figure 4. Percentage variance accounted for by spatial modes associated with the dynamics of MIB. The first spatial modes from the second analysis stage (Fig. 3) were used to determine percentage variance accounted for in different subsets of MIB runs. **A**, Percentage variance accounted for by first spatial mode for high- versus low-rate MIB (Fig. 3A). **B**, Percentage variance accounted for by first spatial mode for long- versus short-duration MIB (Fig. 3B). Error bars indicate SEM; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$, two-sided permutation tests).

Spatial modes of activity during fixation and periodic visual stimulation

Fluctuations of cortical activity during fixation were also (like MIB) dominated by a strong global component, and contained retinotopically specific structure (Fig. 5). We analyzed four subregions within each cortical area, one per visual field quadrant. In these analyses, we focused on the first and the second spatial mode, for two reasons: (1) during fixation, the homogenous structure of mode 1 (no systematic differences in spatial modes weights between cortical subregions) for fixation, combined with the large number of subregions (16), imposed no specific constraints on the structure of mode 2 (Fig. 5B); (2) during periodic stimulation, mode 2 had an eigenvalue almost as large as mode 1 (Fig. 5D).

Just as observed during MIB (Fig. 1D), the correlation matrix during fixation was largely homogenous (Fig. 5A). Accordingly, all subregions had the same sign in the first spatial mode (red in Fig. 5B), which accounted for the largest proportion of variance in the multivariate time series (Fig. 5B, first bin in inset histogram). The retinotopic specificity index for this spatial mode was not significant ($p = 0.99$, permutation test; $N = 8$ subregions).

But the fluctuations also exhibited retinotopically specific structure, which was evident in the second spatial mode that accounted for the second largest proportion of variance (Fig. 5B). Corresponding subregions in different visual cortical areas tended to have the same sign in the spatial modes. This was most pronounced for the subregions corresponding to the lower right visual field quadrant (blue throughout the first column denoted “LR” in the map of the second spatial mode). The other three subregions were less homogenous, but showed a similar trend. Consequently, this spatial mode had a retinotopic specificity index that was significantly larger than expected by chance (index: 15.93%, $p < 0.05$, permutation test, $N = 16$ subregions). This is consistent with the conclusion from a number of previous optical imaging and fMRI studies, that the spatial structure of spontaneous activity fluctuations during rest is shaped by the intrinsic (e.g., retinotopically specific) anatomical connectivity of the visual cortex (Ringach, 2009).

The spatial and temporal structure of cortical activity fluctuations also reflected activity evoked by periodic visual stimulation (Fig. 5C,D). Two diagonally arranged visual field quadrants were stimulated in phase with one another (in anti-phase with the

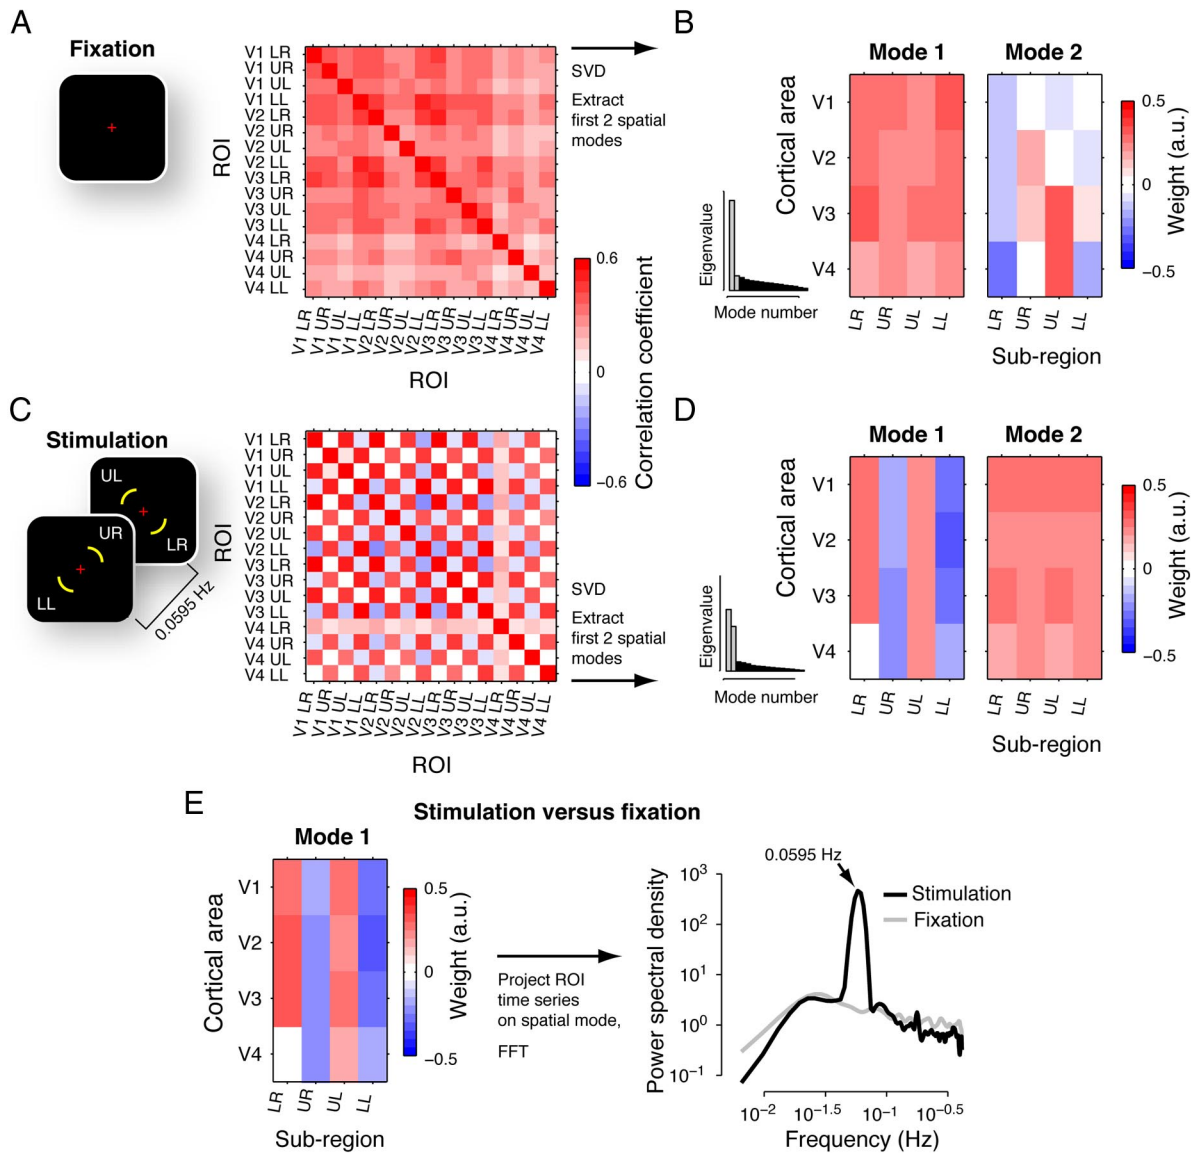


Figure 5. Spatial modes of activity during fixation and periodic stimulation. **A**, Correlation matrix for fixation. Left, Stimulus during fixation condition. Right, Correlation matrix, averaged across subjects. Within each of V1, V2, V3, and V4, four ROIs were defined as the subregions that corresponded retinotopically to a portion of each of the four visual field quadrants (see stimulus snapshot in **C**): LR, lower right quadrant of the visual field; UR, upper right; UL, upper left; LL, lower left. **B**, Maps for the first two spatial modes during fixation, averaged across subjects, are shown along with histograms of the average eigenvalues (corresponding eigenvalues highlighted in gray). The spatial modes were reorganized as matrices (4 cortical areas \times 4 subregions) for illustration. **C**, Correlation matrix for periodic visual stimulation. Left, Snapshots of stimuli during periodic stimulation. Visual field locations in white are indicated for illustration only. Right, Correlation matrix, averaged across subjects. Same format as **A**. **D**, First two spatial modes during periodic stimulation. Same format as **B**. **E**, Difference between stimulation and fixation. Left, Map of the first spatial mode, averaged across subjects. Right, Power spectra of the spatial mode's time series during fixation and stimulation. Note the prominent peak at the stimulation frequency (absent for fixation).

other two quadrants) at a fixed temporal frequency (0.0595 Hz). Thus, within V1–V3, the dorsal part of the left hemisphere was stimulated in phase with the ventral part of the right hemisphere and the ventral part of the left hemisphere was stimulated in phase with the dorsal part of the right hemisphere. This spatial pattern of fMRI responses was revealed by our analysis; correlations were positive for the subregions that were stimulated simultaneously, and negative for the subregions receiving anti-phase stimulation (Fig. 5C). This structure of positive and negative correlations was evident in the first spatial mode (Fig. 5D), for which the retinotopic specificity index was significant (index: 48.02%, $p < 10^{-3}$, permutation test, $N = 16$ subregions).

The spatial mode that fluctuated most strongly during periodic stimulation relative to fixation (second stage of the analysis) also revealed the spatiotemporal structure of the stimulus-evoked

responses. Again, the structure of positive and negative correlations evoked by the periodic stimulation was evident in the first spatial mode (Fig. 5E), for which the retinotopic specificity index was significant (index: 47.16%, $p < 0.001$, permutation test; $N = 16$ subregions). Further, the time series of the spatial mode identified by this analysis exhibited a prominent peak at the stimulation frequency during stimulation, but not during fixation (Fig. 5E). Together, the results from the fixation and periodic stimulation experiments validate our approach for characterizing the correlated fluctuations in cortical activity.

Discussion

It has long been known that neuronal activity in visual cortex continuously fluctuates, even in the face of constant sensory input (Heggelund and Albus, 1978; Arieli et al., 1996; Fiser et al.,

2004; Fox and Raichle, 2007; Faisal et al., 2008). While the spatial structure of fluctuations of population activity is predominantly global across the visual cortex (Leopold et al., 2003; Nir et al., 2006; Yeo et al., 2011), previous fMRI studies revealed also retinotopically specific components (Haynes et al., 2005; Lauritzen et al., 2009; Freeman et al., 2011; Heinzle et al., 2011). But how are these fluctuations related to the spontaneous dynamics of perception? We addressed this question by linking spatial modes of fluctuations of activity in visual cortex to the statistics of the spontaneous perceptual suppression dynamics in MIB (rate and duration of disappearance states), regardless of the timing of subjects' behavioral reports of the perceptual changes. The rationale behind this approach was that some (perhaps many) of the activity fluctuations contributing to the perceptual dynamics may not be correlated in time with the behavioral reports (Lumer and Rees, 1999; Lee et al., 2007).

Our findings can be summarized as follows. Fluctuations of activity during constant (or absent) visual input (MIB and fixation) were dominated by a global component, which was homogeneous across all of V1–V4 (Figs. 1, 5A,B). This global component exhibited fluctuations of large amplitude (accounting for at least 50% of the variance in the fMRI measurements), which were, however, not systematically related to the perceptual dynamics of MIB (Fig. 2). By contrast, the spatial modes that were reliably linked to perception were retinotopically specific for the target and exhibited fluctuations of smaller amplitude (Figs. 3, 4). The spatial mode corresponding to high versus low rate of target disappearance was target-specific in V4, but not V1 (Fig. 3A), whereas the spatial mode corresponding to long versus short duration of target disappearance was target-specific in V1, but not V4 (Fig. 3B). These target-specific V1 fluctuations were decoupled in time from disappearance reports, which may explain why they had been overlooked by previous studies of MIB focusing on cortical fluctuations time-locked to perceptual reports (Donner et al., 2008; Hsieh and Tse, 2009; Schölvinck and Rees, 2010).

By linking V1 activity to the stability of perceptual suppression, our results shed new light on the role of V1 in bistable perception (Leopold and Logothetis, 1996; Wilke et al., 2006; Lee et al., 2007; Donner et al., 2008; Maier et al., 2008; Libedinsky et al., 2009). Our results suggest that different levels of the visual cortical hierarchy shape the dynamics of MIB via different mechanisms. One implication of this is that realistic computational models of MIB need to contain at least two processing levels. The same conclusion had previously been reached for models of binocular rivalry, based on other empirical results (Blake and Logothetis, 2002; Wilson, 2003; Freeman, 2005), suggesting that this may be a general principle underlying all bistable perceptual phenomena (Gigante et al., 2009).

It is unlikely that the target-specific V1 fluctuations were caused by the subjects' motor responses, or ongoing changes in arousal. First, spatially selective modulations in visual cortex are inconsistent with such nonspecific effects. Second, and more importantly, if due to motor responses, the target-specific V1 fluctuations should have been correlated in time with behavioral reports (and should have reflected the rate of these reports), contrary to what was observed here and in our previous study using report-triggered averaging (Donner et al., 2008).

It is also unlikely that the target-specific V1 fluctuations reflected changes in top-down, selective attention. Top-down attention modulates spatially specific correlations between visual cortical areas (Haynes et al., 2005; Lauritzen et al., 2009) and the dynamics of MIB (Schölvinck and Rees, 2009). Attention may

even be necessary for instigating normal bistable perceptual dynamics (Zhang et al., 2011). But it seems unlikely to find a strong attentional modulation in V1 but none in V4 (Kastner and Ungerleider, 2000). Therefore, the structure of the spatial mode associated with the disappearance duration is inconsistent with top-down attention.

What, then, underlies the target-specific V1 fluctuations, and their link to duration of perceptual suppression? One possible scenario is that the V1 fluctuations were caused by fixational eye movements and, consequently, affected the duration of perceptual suppression states by changing the cortical adaption state. The target representation in V1 is gradually weakened by adaptation, which is inherited by the target representation in V4. Transients in V1 cause a rapid recovery from adaptation, strengthening the V4 target representation and shortening perceptual suppression states. V1 transients may be generated by microsaccades (Martinez-Conde et al., 2004; Hsieh and Tse, 2009; Bonneh et al., 2010), which shift the target representation to a neighboring, nonadapted location in the retinotopic map (Martinez-Conde et al., 2004). Fewer microsaccades lead to fewer target-specific V1 transients, longer adaptation of the target representation and longer perceptual suppression. Microsaccades elicit retinotopically specific transients in V1 by driving neurons processing the static target more strongly than neurons processing the mask (Bair and O'Keefe, 1998; Martinez-Conde et al., 2000). Because of the sluggish hemodynamics, fewer target-specific V1 transients during runs with long disappearance states produce target-specific fMRI fluctuations with larger amplitudes. Because the target-specific V1 fluctuations were neither correlated with the target-specific V4 fluctuations, nor with behavioral reports, they must have affected perceptual state durations without immediately culminating in a change in the V4 target representation and report, for example by accumulating slowly in V4 (Gigante et al., 2009). While speculative, future studies using high-resolution monitoring of microsaccades during fMRI, and adopting our analysis approach, could test (specific aspects of) this scenario.

The scenario described in the preceding paragraph identified fixational eye movements and adaptation as factors that may drive fluctuations of the target representation in V1 and contribute to the duration of perceptual suppression states. But stimulus-independent activity fluctuations may have occurred in V1 for many reasons other than fixational eye movements. They may have emerged locally within early visual cortex, or they may have been mediated by corticocortical (Moore and Armstrong, 2003; Haynes et al., 2005; Ekstrom et al., 2008; Siegel et al., 2008; Gregoriou et al., 2009; Lauritzen et al., 2009) or thalamo-cortical (Wilke et al., 2009; Saalman et al., 2012) long-range interactions.

Our results have broader implications for the use of spontaneous fMRI signal fluctuations to identify so-called "resting-state" networks in the human brain (Biswal et al., 1995; Fox and Raichle, 2007; Vincent et al., 2007; Honey et al., 2009; Cole et al., 2010; Sadaghiani et al., 2010). The fluctuations assessed by the majority of these studies occur at coarse spatial scales spanning multiple cortical areas. Such large-scale fluctuations reflect, at least to some extent, corticocortical connections (Honey et al., 2009) and they have been linked to behavioral state (Bianciardi et al., 2009a) and perceptual performance (Fox and Raichle, 2007; Sadaghiani et al., 2010). Our present results put the focus on the fluctuations of specific cortical subpopulations within individual cortical areas. Although smaller in amplitude, these fluctuations reliably index the dynamics of perception.

References

- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871. [CrossRef Medline](#)
- Bair W, O'Keefe LP (1998) The influence of fixational eye movements on the response of neurons in area MT of the macaque. *Vis Neurosci* 15:779–786. [Medline](#)
- Bianciardi M, Fukunaga M, van Gelderen P, Horowitz SG, de Zwart JA, Duyn JH (2009a) Modulation of spontaneous fMRI activity in human visual cortex by behavioral state. *Neuroimage* 45:160–168. [CrossRef Medline](#)
- Bianciardi M, Fukunaga M, van Gelderen P, Horowitz SG, de Zwart JA, Shmueli K, Duyn JH (2009b) Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study. *Magn Reson Imaging* 27:1019–1029. [CrossRef Medline](#)
- Birn RM, Diamond JB, Smith MA, Bandettini PA (2006) Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *Neuroimage* 31:1536–1548. [CrossRef Medline](#)
- Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34:537–541. [CrossRef Medline](#)
- Blake R, Logothetis N (2002) Visual competition. *Nat Rev Neurosci* 3:13–21. [CrossRef Medline](#)
- Bonnef YS, Donner TH (2011) Motion induced blindness. *Scholarpedia* 6:3321. [CrossRef](#)
- Bonnef YS, Cooperman A, Sagi D (2001) Motion-induced blindness in normal observers. *Nature* 411:798–801. [CrossRef Medline](#)
- Bonnef YS, Donner TH, Sagi D, Fried M, Cooperman A, Heeger DJ, Arieli A (2010) Motion-induced blindness and microsaccades: cause and effect. *J Vis* 10(14):22. [CrossRef](#)
- Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62. [CrossRef Medline](#)
- Cardoso MM, Sirotin YB, Lima B, Glushenkova E, Das A (2012) The neuroimaging signal is a linear sum of neurally distinct stimulus- and task-related components. *Nat Neurosci* 15:1298–1306. [CrossRef Medline](#)
- Cole DM, Smith SM, Beckmann CF (2010) Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. *Front Syst Neurosci* 4:8. [Medline](#)
- Deco G, Romo R (2008) The role of fluctuations in perception. *Trends Neurosci* 31:591–598. [CrossRef Medline](#)
- Donner TH, Sagi D, Bonnef YS, Heeger DJ (2008) Opposite neural signatures of motion-induced blindness in human dorsal and ventral visual cortex. *J Neurosci* 28:10298–10310. [CrossRef Medline](#)
- Efron B, Tibshirani R (1998) An introduction to the bootstrap, 1st CRC reprint. Boca Raton, FL: Chapman and Hall/CRC.
- Ekstrom LB, Roelfsema PR, Arsenault JT, Bonmassar G, Vanduffel W (2008) Bottom-up dependent gating of frontal signals in early visual cortex. *Science* 321:414–417. [CrossRef Medline](#)
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303. [CrossRef Medline](#)
- Fiser J, Chiu C, Weliky M (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431:573–578. [CrossRef Medline](#)
- Fox MD, Raichle ME (2007) Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci* 8:700–711. [CrossRef Medline](#)
- Freeman AW (2005) Multistage model for binocular rivalry. *J Neurophysiol* 94:4412–4420. [CrossRef Medline](#)
- Freeman J, Donner TH, Heeger DJ (2011) Inter-area correlations in the ventral visual pathway reflect feature integration. *J Vis* 11(4):15. [Medline](#)
- Friston KJ, Büchel C (2004) Functional connectivity. In: *Human brain function*, Ed 2 (Frackowiak RJS, ed). San Diego, Academic.
- Gigante G, Mattia M, Braun J, Del Giudice P (2009) Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Comput Biol* 5:e1000430. [CrossRef Medline](#)
- Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324:1207–1210. [CrossRef Medline](#)
- Haynes JD, Tregellas J, Rees G (2005) Attentional integration between anatomically distinct stimulus representations in early visual cortex. *Proc Natl Acad Sci U S A* 102:14925–14930. [CrossRef Medline](#)
- Heggelund P, Albus K (1978) Response variability and orientation discrimination of single cells in striate cortex of cat. *Exp Brain Res* 32:197–211. [Medline](#)
- Heinze J, Kahnt T, Haynes JD (2011) Topographically specific functional connectivity between visual field maps in the human brain. *Neuroimage* 56:1426–1436. [CrossRef Medline](#)
- Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P (2009) Predicting human resting-state functional connectivity from structural connectivity. *Proc Natl Acad Sci U S A* 106:2035–2040. [CrossRef Medline](#)
- Hsieh PJ, Tse PU (2009) Microsaccade rate varies with subjective visibility during motion-induced blindness. *PLoS One* 4:e5163. [CrossRef Medline](#)
- Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17:825–841. [CrossRef Medline](#)
- Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci* 23:315–341. [CrossRef Medline](#)
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425:954–956. [CrossRef Medline](#)
- Kim YJ, Grabowecy M, Suzuki S (2006) Stochastic resonance in binocular rivalry. *Vis Res* 46:392–406. [CrossRef Medline](#)
- Lamme VA (2003) Why visual attention and awareness are different. *Trends Cogn Sci* 7:12–18. [CrossRef Medline](#)
- Lauritzen TZ, D'Esposito M, Heeger DJ, Silver MA (2009) Top-down flow of visual spatial attention signals from parietal to occipital cortex. *J Vis* 9(13):18 11–14. [CrossRef](#)
- Lee SH, Blake R, Heeger DJ (2007) Hierarchy of cortical responses underlying binocular rivalry. *Nat Neurosci* 10:1048–1054. [CrossRef Medline](#)
- Leopold DA, Logothetis NK (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379:549–553. [CrossRef Medline](#)
- Leopold DA, Murayama Y, Logothetis NK (2003) Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging. *Cereb Cortex* 13:422–433. [CrossRef Medline](#)
- Libedinsky C, Savage T, Livingstone M (2009) Perceptual and physiological evidence for a role for early visual areas in motion-induced blindness. *J Vis* 9(1):14 1–10. [CrossRef](#)
- Lumer ED, Rees G (1999) Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc Natl Acad Sci U S A* 96:1669–1673. [CrossRef Medline](#)
- Maier A, Wilke M, Aura C, Zhu C, Ye FQ, Leopold DA (2008) Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nat Neurosci* 11:1193–1200. [CrossRef Medline](#)
- Martinez-Conde S, Macknik SL, Hubel DH (2000) Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. *Nat Neurosci* 3:251–258. [CrossRef Medline](#)
- Martinez-Conde S, Macknik SL, Hubel DH (2004) The role of fixational eye movements in visual perception. *Nat Rev Neurosci* 5:229–240. [CrossRef Medline](#)
- Mitra PP, Pesaran B (1999) Analysis of dynamic brain imaging data. *Biophys J* 76:691–708. [CrossRef Medline](#)
- Moore T, Armstrong KM (2003) Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421:370–373. [CrossRef Medline](#)
- Moreno-Bote R, Rinzel J, Rubin N (2007) Noise-induced alternations in an attractor network model of perceptual bistability. *J Neurophysiol* 98:1125–1139. [CrossRef Medline](#)
- Murphy K, Birn RM, Handwerker DA, Jones TB, Bandettini PA (2009) The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? *Neuroimage* 44:893–905. [CrossRef Medline](#)
- Nestares O, Heeger DJ (2000) Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med* 43:705–715. [CrossRef Medline](#)
- Nir Y, Hasson U, Levy I, Yeshurun Y, Malach R (2006) Widespread functional connectivity and fMRI fluctuations in human visual cortex in the absence of visual stimulation. *Neuroimage* 30:1313–1324. [CrossRef Medline](#)
- Noest AJ, van Ee R, Nijs MM, van Wezel RJ (2007) Percept-choice sequences driven by interrupted ambiguous stimuli: a low-level neural model. *J Vis* 7(8):10. [CrossRef](#)
- Polonsky A, Blake R, Braun J, Heeger DJ (2000) Neuronal activity in human primary visual cortex correlates with perception during binocular rivalry. *Nat Neurosci* 3:1153–1159. [CrossRef Medline](#)

- Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci* 19:1736–1753. [Medline](#)
- Ringach DL (2009) Spontaneous and driven cortical activity: implications for computation. *Curr Opin Neurobiol* 19:439–444. [CrossRef Medline](#)
- Saalmann YB, Pinsk MA, Wang L, Li X, Kastner S (2012) The pulvinar regulates information transmission between cortical areas based on attention demands. *Science* 337:753–756. [CrossRef Medline](#)
- Sadaghiani S, Hesselmann G, Friston KJ, Kleinschmidt A (2010) The relation of ongoing brain activity, evoked neural responses, and cognition. *Front Syst Neurosci* 4:20. [Medline](#)
- Schölvinck ML, Rees G (2009) Attentional influences on the dynamics of motion-induced blindness. *J Vis* 9(1):38 1–9. [CrossRef](#)
- Schölvinck ML, Rees G (2010) Neural correlates of motion-induced blindness in the human brain. *J Cogn Neurosci* 22:1235–1243. [CrossRef Medline](#)
- Schölvinck ML, Maier A, Ye FQ, Duyn JH, Leopold DA (2010) Neural basis of global resting-state fMRI activity. *Proc Natl Acad Sci U S A* 107:10238–10243. [CrossRef Medline](#)
- Siegel M, Donner TH, Oostenveld R, Fries P, Engel AK (2008) Neuronal synchronization along the dorsal visual pathway reflects the focus of spatial attention. *Neuron* 60:709–719. [CrossRef Medline](#)
- Sterzer P, Kleinschmidt A, Rees G (2009) The neural bases of multistable perception. *Trends Cogn Sci* 13:310–318. [CrossRef Medline](#)
- Tong F, Meng M, Blake R (2006) Neural bases of binocular rivalry. *Trends Cogn Sci* 10:502–511. [CrossRef Medline](#)
- Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286:1943–1946. [CrossRef Medline](#)
- Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME (2007) Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447:83–86. [CrossRef Medline](#)
- Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. *Neuron* 56:366–383. [CrossRef Medline](#)
- Wilke M, Logothetis NK, Leopold DA (2006) Local field potential reflects perceptual suppression in monkey visual cortex. *Proc Natl Acad Sci U S A* 103:17507–17512. [CrossRef Medline](#)
- Wilke M, Mueller KM, Leopold DA (2009) Neural activity in the visual thalamus reflects perceptual suppression. *Proc Natl Acad Sci U S A* 106:9465–9470. [CrossRef Medline](#)
- Wilson HR (2003) Computational evidence for a rivalry hierarchy in vision. *Proc Natl Acad Sci U S A* 100:14499–14503. [CrossRef Medline](#)
- Yeo BT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL (2011) The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* 106:1125–1165. [CrossRef Medline](#)
- Zhang P, Jamison K, Engel S, He B, He S (2011) Binocular rivalry requires visual attention. *Neuron* 71:362–369. [CrossRef Medline](#)