Supporting Information

Smith and Muckli 10.1073/pnas.1000233107

SI Methods

Subjects. Six subjects (one male) participated in experiment 1. Seven different participants (four male) took part in experiment 2. An additional four participants (two male) took part in experiment 3, and a further two (one male) participants took part in experiment 4. There was no overlap of subjects across experiments. All subjects had normal or corrected to normal vision. All participants gave written, informed consent in accordance with procedures and protocols approved by the local ethics committee of the Faculty of Information and Mathematical Sciences of the University of Glasgow.

Stimuli and Design. Participants were presented with static natural visual scenes with the lower-right quadrant occluded (Fig. 1*A*) in either a block design (experiments 1, 3, and 4) or a rapid-event–related design (experiment 2). There were three individual scenes, a people scene, a car scene, and a boat scene (Fig. S1*A*). In experiments 1 and 2 the same scenes also were presented nonoccluded as a control. The entire stimulus spanned 22.5 × 18° of visual angle, and the occluded region spanned ≈11 × 9° of visual angle, approximately one-quarter of the full stimulus (Fig. 1*A*).

In experiment 1, each run comprised six sequences of stimulation with intervening fixation periods. Each stimulation sequence lasted 96 s, consisting of one presentation of each of the six trial types (three scenes, occluded or nonoccluded) for 12 s, with 12-s fixation at the beginning and the end of each such series. Within each 12-s block of stimulus presentation, the stimulus was flashed on and off (200 ms on/200 ms off) 30 times; we used this presentation cycle to maximize the signal-to-noise ratio (1). The trial order within each sequence was pseudorandom with the constraint that the occluded and nonoccluded versions of one scene could never follow one another directly. Participants were instructed to maintain fixation throughout each run. The task of the participant was to monitor the stream of frames for a change in color (randomly chosen frames were changed from gray-scale to red). With six sequences of 96 s (576 s total), the scanning time was less than 10 min per run. For two of the six subjects the experimental runs were kept independent from the mapping runs; for the other four subjects mapping and experimental conditions were interleaved in the same run.

In the independent mapping procedure, we localized the cortical representation of the occluded stimulus region (Fig. 1*B* and ref. 2). Participants viewed checkerboard stimuli in a traditional block design in either a target or a surround region. The procedure consisted of alternating 12-s blocks of fixation or stimulation. Within each stimulation block, participants viewed contrast-reversing checkerboard stimuli (4 Hz) in one of four locations [target right visual field (RVF), surround (RVF), target left visual field (LFV), surround (LVF)], each presented on six occasions. We initially included the LVF stimuli to provide an additional control; however, we subsequently decided not to use any regions defined from that visual field in the present analyses. To minimize the influence of any spillover of activity of neighboring (and therefore stimulated) areas, we included several precautions. First, the occluded stimulus began 0.5° (diagonally) from the center of fixation. Second, the surround-area mapping stimulus comprised the inner 1° (diagonally from fixation) of the nonstimulated region in the main experiment, and the target stimulus was further offset diagonally by 1° from the inner edge of the surround-area stimulus. Thus the target region was diagonally ≈2° away from the stimulated region. However, because we accept only vertices that show a significant target-area response with no significant surround-area response, our effective border is likely to be much greater. The target region spanned 9.6 × 7.5°.

For the final four participants, we optimized this cortical mapping procedure by including the cortical mapping blocks within each run of the main experiment and mapping only the RVF. Each mapping sequence consisted of presentation of checkerboard stimuli (parameters as above) in both the target RVF and the surround-area RVF position, each for 12 s, with 12 s of fixation between each stimulus and at the beginning and end of the sequence. Two such sequences were included in each run, at randomly chosen breaks between the main stimulus sequences, with the order of the mapping stimuli counterbalanced across repetitions. Scanning time was about 12 min per run; each participant completed four runs of this version of experiment 1.

In experiment 2 (rapid-event–related design), in each run each scene was presented 20 times occluded and 20 times nonoccluded as a control. Each 4-s trial consisted of the presentation of one image in a cycle of 200 ms on/200 ms off for three cycles, followed by just the fixation checkerboard for 2,800 ms. The participant's task was to perform one-back repetition detection across trials. Participants were instructed to maintain fixation throughout each experimental run. Trial order was pseudorandom with the constraint that the occluded and nonoccluded versions of the same scene could never follow one another directly. Each run comprised 120 trials (six trial types × 20 repetitions of each) preceded and followed by 20 s of fixation. Thus each run lasted less than 9 min. Seven participants performed either four or eight runs of the experiment in one or two scanning sessions (giving 80 or 160 trials per condition). All participants also performed the independent mapping experiment (once during each scanning session but as an independent run), as described above.

In experiment 3, we again used the successful block design of experiment 1 with the following changes: (*i*) only occluded trials were shown to participants, and (*ii*) the task was changed to detection of a change in the color of the fixation marker. We showed twice as many occluded trials to maximize our chances of finding the effect; instead of six trials per trial type per run (and six trial types), there were 12 trials per trial type per run (but only three trial types). Thus within each main sequence we presented each occluded scene twice in pseudorandom order (each trial type being constrained not to follow itself). In all other respects the design was identical to that of the version of experiment 1 that included mapping trials within the main runs. Each participant performed four runs. The paradigm in experiment 4 was exactly the same as in experiment 3, the only difference being in the preprocessing applied to the scene stimuli before the experiment was run. This processing (spectral normalization) allowed control of low-level image properties such as global luminance, contrast, and energy at different spatial frequencies and orientations within the naturally stimulated region of the image (image processing methods are described fully in the section, "Low-Level Image Control").

MRI Data Acquisition. Participants viewed the images through a set of high-quality visual display goggles (Nordic Neurolab). MRI data were collected with a 3-T Siemens Tim Trio System with a 12-channel head coil and integrated parallel imaging techniques (IPAT factor: 2). Blood oxygen level-dependent (BOLD) signals were measured with an echo-planar imaging sequence (echo time: 30 ms, repetition time: 1,000 ms, field of view: 210 mm, flip angle: 62°, 10% gap, 16–18 axial slices). We pushed for higher spatial resolution throughout the experiment and began with voxel sizes of 2.5 × 2.5 × 4.4 mm (experiment 1) and 3 × 3 ×

4 mm (experiment 2) but improved the resolution for later subjects to 2.3 × 2.3 × 3 mm (experiment 1) and 2.5 × 2.5 × 4.4 mm (experiment 2). All participants were recorded with resolution of 2.3 × 2.3 × 3 mm in experiment 3 and with a resolution of 2.5 × 2.5 × 3 mm in experiment 4. The slices were positioned to maximize coverage of occipital regions. A high-resolution 3D anatomical scan (3D MPRAGE, 1 × 1 × 1-mm resolution) was recorded in the same session as the functional scans. For participants who performed two sessions of the main experiment, separate anatomical scans were recorded in each scanning session (a faster scan making use of IPAT was used in the second session) to facilitate realignment of the functional data.

MRI Data Processing. The following steps were carried out independently for each participant in each experiment. Functional data for each run (main runs and mapping runs) were corrected for slice time and 3D motion, temporally filtered (high-pass filtered at 0.01 Hz and linearly detrended), and spatially normalized into the Talairach space with Brain Voyager QX (Brain Innovation). We then projected the anatomical data onto a flattened surface representation and overlaid the functional data (2–4).

The mapping data (either the main runs inclusive of mapping or the independent mapping run, depending on the participant) then were subjected to a standard general linear model (GLM) with one predictor for each condition. We defined a patch of primary visual cortex (V1) and a patch of visual area (V2) from the contrast of target area minus surround area (Fig. 1*B*). Within these patches we selected for further analyses only the vertices that met the following profile: significant positive effect for target area alone ($t > 1.65$, uncorrected $P < 0.05$) and, crucially, no significant effect for surround area alone (absolute $t < 1.65$, uncorrected $P > 0.05$). With this definition we sought to minimize the potential for any spillover activity from stimulated regions affecting the signal in our selected vertices. Full details on the selected vertices are given in Tables S1–S3.

We then extracted the time course from each selected vertex and applied a GLM to estimate response amplitudes on a single-block basis (experiments 1, 3, and 4) or a single-trial basis (experiment 2), independently per run. The resulting beta weights estimate peak activation for each single block or trial of stimulation, assuming a standard 2γ model of the hemodynamic response function. We used these estimates (beta weights) from the target vertices as the input to the pattern classifier.

Multivariate Pattern Classification Analysis. We performed the pattern classifier analysis independently for occluded and control trial types (experiments 1 and 2). Thus we trained a linear classifier [Linear Discriminant Analysis (LDA) or Linear Support Vector Machine (SVM)] to learn the mapping between a set of multivariate observations of brain activity and the particular scene (people, car or boat) that had been presented (5–9). We then tested the classifier on an independent set of test data. We trained the classifiers with a set of single-block (or single-trial) brain-activity patterns (beta values). We tested the classifiers either on independent single blocks (single trials) or on the average brain-activity pattern for each stimulus class in the independent set of test data.

In building our classifiers we chose randomly from the set of designated target vertices (i.e., those showing an effect for the target area but not for the surround area), initially pooled across V1 and V2, sampling 30 times for each of several different vertex set sizes (1, 5, and then to 70 vertices in increments of 5, giving 15 different set sizes) (5). In subsequent analyses we performed the analyses independently for each visual region (choosing vertices randomly, again in set sizes from 1–30). Note that for experiment 2, using our explicit criteria (target area response > 1.65, surround area response < 1.65), some participants had fewer than 30 valid vertices within either V1 or V2 (1 subject in V1, 2 subjects in V2; see Tables S2 and S3 for full details). These participants were included in the classifier analyses to the extent possible (i.e. if a participant had 12 vertices defined for V1, then data from that participant would go into the subject averages only for vertex set sizes of 1, 5, and 10). To assess the performance of our classifiers, we used an *n*-fold leave one run out cross-validation procedure; thus our models were built from $n - 1$ runs and were tested on the independent *n*th run (repeated for the *n* different possible partitions of the runs in this scheme) (6, 10, 11). Note that the random sampling procedure gives an estimate of the variance associated with classifier performance for each given vertex level across the target vertices, providing a good indication of whether the signal in the region of interest can discriminate among the three scenes presented for the specific vertex level considered. Independent classifiers were constructed for occluded and control trials for each participant, and we report average classifier performance across participants. The same procedure was followed for both experiments (the only differences being the number of training and test examples). We also applied the same basic method in further generalization analyses where the classifier was trained on one type of trial (e.g., occluded) and then was tested on the other type of trial [e.g., control (8, 12)]. Note, however, that for these latter analyses we used all available data for each trial type either to train or test the classifier (because we generalize here across trial types), and therefore there is no run-specific partitioning for these analyses.

The classifier analysis for experiments 3 and 4 proceeded in a similar manner, although for this analysis we report single-participant performance rather than subject averages because we had a smaller number of participants in these experiments.

The LDA classifier was implemented in MATLAB by the Statistics Toolbox function Classify. The linear SVM was implemented using the LIBSVM toolbox (13), with default parameters (notably C = 1). Note that the activity of each vertex was normalized (separately for training and test data) within a range of −1 to 1 before input to the SVM.

Permutation Test. To be sure that single-trial classification performance statistically exceeded chance in the occluded condition of experiment 2, we performed an additional permutation test (14, 15) at the group level. Note that this effect was the weakest we report here as assessed by a standard one-sample *t* test against chance; $t(6) = 2.06$, $P = 0.043$ [39% correct against chance (33%)].

This test involved permuting the relationship between the stimulus labels and the data (independently for each run) 100 times for each subject and computing classifier performance as before (i.e., with *n*-fold cross-validation and the same set of vertex samples as for the observed data) for each permutation of the labels. This computation gave a permutation distribution for each subject to which the matching observed value was appended. We then sampled 10^4 combinations of sample size seven (we had seven subjects), one value from each subject-specific distribution each time, and found the probability of these samples having a mean equaling or exceeding the observed value of 39% correct (15).

Eye Tracking. Eye movements were recorded for all participants in experiment 3 by an eye- tracking camera with a 60-Hz sample rate integrated into the goggle display system (Nordic Neurolab) in concert with the ViewPoint software. We detrended the time series from each run in segments of 96 s, because this time corresponded to one sequence of the block design (including fixation periods at the beginning and at the end). We then transformed the data into units of degrees of visual angle and classified as an eye movement any succession of samples whose radius exceeded 1.5° of visual angle for a duration of 150 ms (4). We then tested whether either the mean number of saccades or the mean position on the *x* or y

axis differed across the three conditions in this experiment. We report the results of this analysis in Fig. S4.

Low-Level Image Control (Experiment 4). In experiment 4, we investigated the contribution of basic low-level stimulus features to the context effect reported. The experiment was identical to that of experiment 3 except that the scene stimuli first were preprocessed as follows: The amplitude spectrum of each complete scene was replaced with the mean across the three scenes, thereby controlling exactly the global luminance, contrast, and the energy present at each spatial frequency and orientation in the complete image. This replacement was achieved by taking a 2D fast Fourier transform (FFT) of the image, thus decomposing it into the amplitude and phase spectrum, finding the average amplitude spectrum across the scenes, and then performing an inverse 2D FFT to combine the average amplitude spectrum across scenes with the phase spectrum of each image. Of course, because we then occluded the lower-right portion of the image, the control of such properties was no longer perfect; however, at least in terms of global contrast and luminance, the stimulated areas remained well matched across the three scenes (contrast: mean 0.24, range 0.013; luminance: mean 0.4, range 0.024 on a 0–1 scale).

We did not attempt to normalize the scene with the occluded field overlaid, for two reasons. First, edge artifacts would be present because of the abrupt change in contrast change from the naturally stimulated region to the uniform white occluded field. Second, normalizing in such a manner would lead to the low-level statistics within the occluded region changing across the three scenes. Obviously the latter situation must be avoided in the present paradigm. The argument we put forward here is simply that the possible contribution of global contrast, luminance, and energy at different spatial frequencies and orientations (in the naturally stimulated region) to explaining any possible context effect in experiment 4 is much reduced as compared with the earlier experiments.

Weight Analysis. To shed some light on the possible role of spillover activity/spreading activity and lateral interaction (as opposed to cortical feedback) in explaining the observed effects, we investigated the relationship between the weights of each classifier solution in the occluded condition and the univariate *t*-values for the two independent mapping conditions (target area and surround area; Fig. 1*B*). The logic is that if high (absolute) weights are correlated with high *t* values for the surround mapping stimulus, then that correlation might suggest a possible influence of spillover signal (or spreading activity) in our classifiers. On the other hand, a positive correlation with the target mapping condition would indicate that the more important vertices in the classifier decision function are those with a strong signal to the target stimulus, as we would expect to be the case (purely on signal-to-noise considerations).

Independently for each participant in experiment 1, we extracted the absolute value of the weights of the SVM and LDA linear classifiers for the analysis incorporating both V1 and V2 in the occluded condition. We extracted these weights for the specific vertex set size leading to asymptotic performance (defined as the maximal number of vertices for SVM but as a vertex set size of 25 for LDA, where it achieves maximal performance), for each sampling iteration (30), cross-validation fold (four), and binary classification problem (three). The absolute value of the weight at each vertex for a given instantiation of the classifier indicates the relative influence of each vertex on the classifier's solution (e.g., 8, 16). We note, however, that in the SVM highly correlated vertices, which necessarily have lower weight, still may contribute to the solution (because the SVM is a regularized classifier; ref. 17). This possibility must be borne in mind when considering the analysis using the SVM weights, especially for relatively larger vertex set sizes (where we might expect more correlated vertices).

We Pearson-correlated the absolute value of the weight for each vertex with the corresponding *t* values from the target-area and the surround-area mapping conditions. We used the absolute and not the signed weights because we were interested in finding a relationship between the overall importance of each vertex (its contribution to the classifier solution) and the mapping *t* values, not in the direction in which each vertex is important (e.g., a higher beta value for scene 1 vs. scene 2). We then averaged the resulting correlation values (30 sampling iterations × four cross-validation runs × three binary classification problems) for each participant and performed a two-tailed *t* test to test if the mean correlation across participants was significantly different from zero.

This analysis revealed a significant positive correlation between the LDA weights and the target mapping condition [$r = 0.14$, $t(5) = 2.7$, $P = 0.042$] and a nonsignificant but suggestive negative correlation with the surround-area mapping condition [$r = -0.08$, $t(5) = 2.2$, $P = 0.081$]. As expected, the more important vertices in the LDA classifier were associated with higher responses to the target region and, if anything, with lower responses to the surround area. For the SVM, the results were reversed, with a significant negative relation with the target mapping condition [$r = -0.18$, $t(5) = 4.25$, $P < 0.01$] and a nonsignificant positive correlation with the surround area [$r = 0.06$, $t(5) = 1.59$, $P = 0.17$]. Thus this pattern of correlations suggests that the more important vertices for the LDA classifier do not necessarily match the more important vertices for the SVM classifier, at least when measured using absolute and not signed values. Note that the same pattern held when we investigated the SVM classifier with a vertex set size of 25, thus matching the LDA classifier in terms of feature dimensions (with target area: $r = -0.16$, $P = 0.0252$; with surround area: $r = 0.065$, $P = 0.168$). It is important to realize that the discrepancy between the two classifiers observed here might reflect the smaller weight values necessarily given to correlated vertices in the SVM (17).

Importantly, the broad pattern of correlation results also was replicated in experiment 2 with equal vertex set sizes (70) as input: LDA with target area ($r = 0.20$, $P = 0.006$), LDA with surround area ($r = -0.091$, $P = 0.027$), SVM with target area ($r = -0.14$, $P = 0.003$), and SVM with surround area ($r = 0.052$, $P = 0.156$). We did not assess such correlations for experiments 3 and 4 because of the lower number of participants in these experiments.

Because the interpretation of SVM weights is arguably more complex than that of LDA (17), and because both classifiers found above-chance levels of decoding in almost all tests in every experiment (the sole exception was the LDA classifier in experiment 2 in one type of cross-generalization; Fig. 2*G* and Table S4), we believe that the pattern of results observed with the LDA classifier demonstrates that classification of surrounding visual context is possible with a set of vertices that have no significant positive relation to surround-area mapping activity but instead have a significant positive relation to the target. We argue that this pattern speaks against a spreading activity (or spillover) explanation of our data.

1. Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
2. Muckli L, Kohler A, Kriegeskorte N, Singer W (2005) Primary visual cortex activity along the apparent-motion trace reflects illusory perception. *PLoS Biol* 3:e265.
3. Kriegeskorte N, Goebel R (2001) An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. *Neuroimage* 14: 329–346.
4. Weigelt S, Kourtzi Z, Kohler A, Singer W, Muckli L (2007) The cortical representation of objects rotating in depth. *J Neurosci* 27:3864–3874.
5. Cox DD, Savoy RL (2003) Functional magnetic resonance imaging (fMRI) "brain reading": Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* 19:261–270.
6. Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.

7. Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
8. Formisano E, De Martino F, Bonte M, Goebel R (2008) "Who" is saying "what"? Brain-based decoding of human voice and speech. *Science* 322:970–973.
9. Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Science* 10:424–430.
10. Duda RO, Hart PE, Stork DG (2001) *Pattern Classification* (John Wiley & Sons, New York), 2nd Ed.
11. Hastie T, Tibshirani R, Friedman J (2001) *The Elements of Statistical Learning* (Springer, New York).
12. Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
13. Chang CC, Lin CJ (2001) LIBSVM: A library for support vector machines. Available at http://www.csie.ntu.edu.tw/~cjlin/libsvm. Accessed March 10, 2010.
14. Golland P, Fischl B (2003) Permutation tests for classification: Towards statistical significance in image-based studies. *Inf Process Med Imaging* 18:330–341.
15. Chen Y, et al. (2010) Cortical surface-based searchlight decoding. *Neuroimage*, 10.1016/j.neuroimage.2010.07.035.
16. Mourão-Miranda J, Bokde AL, Born C, Hampel H, Stetter M (2005) Classifying brain states and determining the discriminating activation patterns: Support Vector Machine on functional MRI data. *Neuroimage* 28:980–995.
17. Pereira F, Mitchell T, Botvinick M (2009) Machine learning classifiers and fMRI: A tutorial overview. *Neuroimage* 45 (1, Suppl)S199–S209.

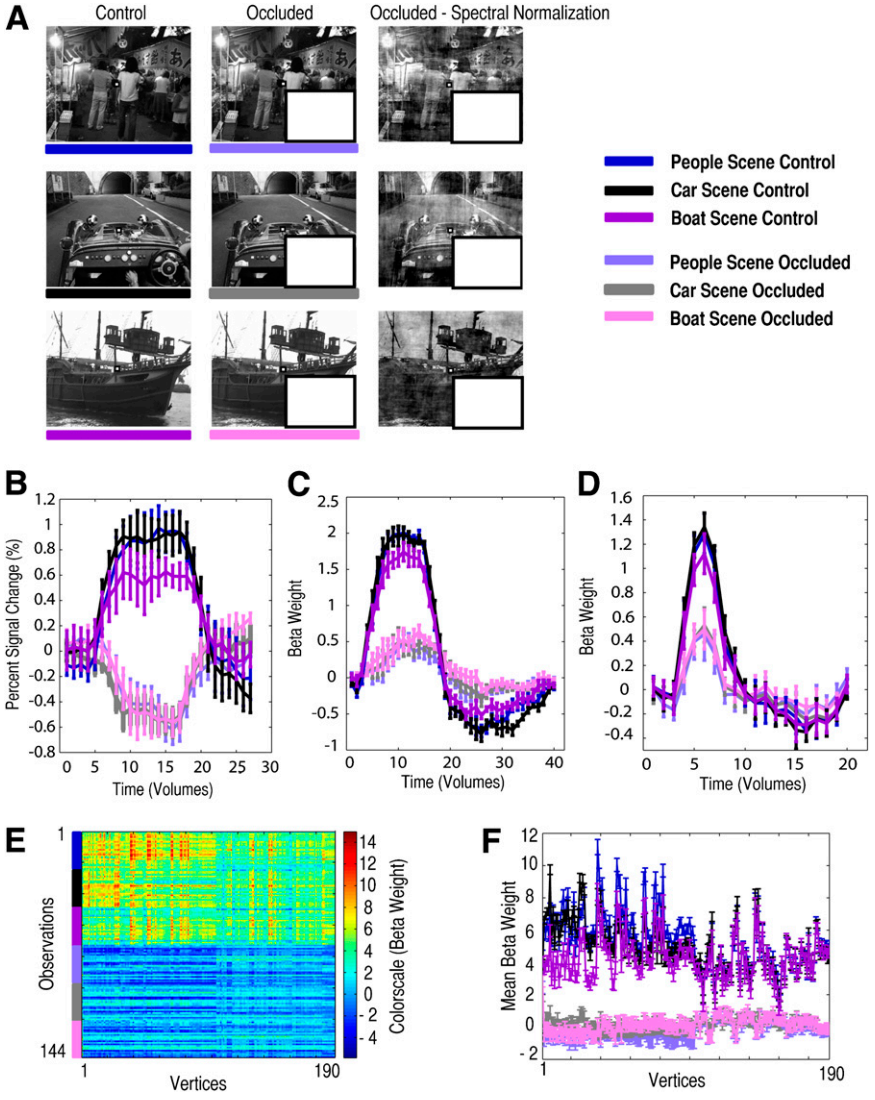

Fig. S1. Stimuli, time courses, and response-pattern estimates. (*A*) The three visual scenes used in the present experiments; as shown in the control condition (experiments 1 and 2) (*Left*), the occluded condition (experiments 1–3) (*Middle*), and the occluded scenes as presented in the low-level image-control experiment (experiment 4, spectral normalization) (*Right*). (*B–D*) Univariate time course and GLM-based hemodynamic response functions (HRF) in experiments 1 and 2. *B* shows the event-related average for each condition in experiment 1 (pooled across V1/V2 and averaged across participants). *C* shows the HRF for each condition in experiment 1, estimated with deconvolution. (*D*) HRF estimate for experiment 2 (rapid-event–related design). As shown in the legend, control (nonoccluded) trials are coded in darker colors (blue, black, and purple); occluded trials are coded in the lighter corresponding colors. The positive beta weights in the occluded condition (*C* and *D*) must be considered relative to the contrast change that takes place between fixation (gray screen) and presentation of a scene in the occluded condition (white occluder; *Methods*). Note the clear distinction between control and occluded trials in experiments 1 and 2 (*B–D*). Note also the large degree of overlap between the three scenes in the occluded compared with the control condition. (*E* and *F*) Single-subject response-pattern estimates. *E* shows the full observations by vertices single-block beta-weight matrix for a representative participant in experiment 1 (pooled across V1/V2). Rows of this matrix correspond to the estimated activation pattern for a given block of stimulation. All patterns input to the classifiers are sampled from this matrix. The top half of the matrix represents activation patterns for control trials; the bottom shows the activation patterns for occluded trials (see colored labels at left edge). Within each half-panel, observations are organized according to each specific scene (see color code at left). *F* shows the average activation pattern for the three scenes in both the control and the occluded conditions (singe-subject data taken from *E*). Error bars reflect one SE across observations.

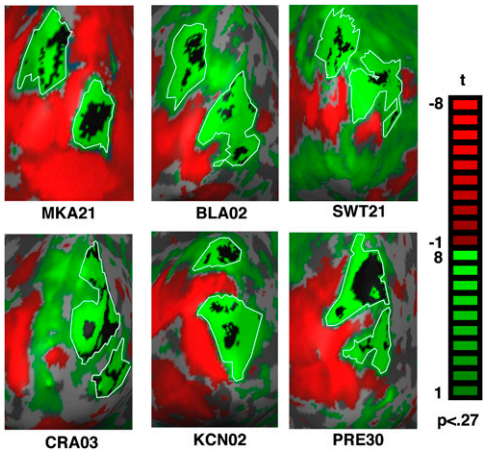

Fig. S2. Definition of V1 and V2 and selected vertices in experiment 1. Each panel shows a close-up view of the cortical reconstruction of the left hemisphere, centered on early visual cortex, for one participant in experiment 1. The color map reflects the contrast of target minus surround area from the mapping data (green, target > surround area; red, surrounding area > target; *Methods*). Overlaid on this map, within the defined V1 and V2 regions (approximately delineated with white outlines), is the location of the selected vertices input to the pattern classification analyses (shown in black) (*Methods* and Fig. 1*B*). Note that such vertices met two additional criteria, in addition to target > surround area: a significant effect for target alone and a nonsignificant effect for the surround area alone (*Methods* and Fig. 1*B*).

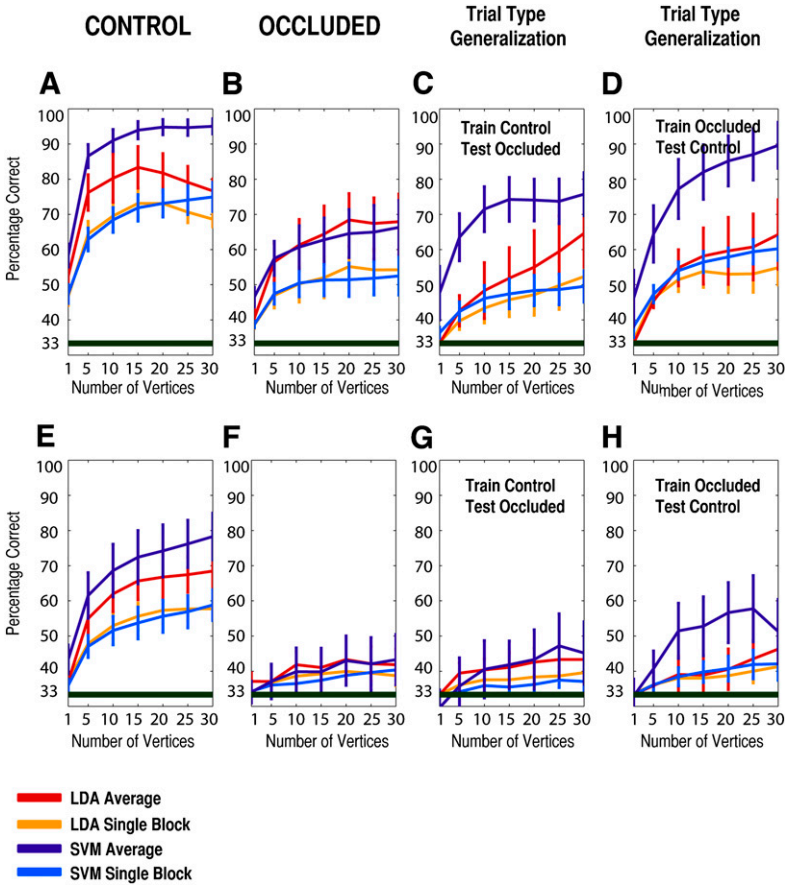

Fig. S3. Pattern-classification performance for experiment 1 split by visual area. (*A–D*) Classifier performance for V1. *A* shows performance (percentage correct) for the two linear classifiers LDA and SVM in decoding which scene was presented in control trials as a function of the number of vertices entering the classifier, for both average level and single-block prediction for vertices in V1. Performance is averaged across participants (error bars represent 1 SEM). Chance performance is indicated by the dark green bar at 33%. *B* shows the same information for occluded trials. *C* and *D* show the results for the trial-type generalization analyses (see main text). (*E–H*) Classifier performance for V2, arranged as for V1.

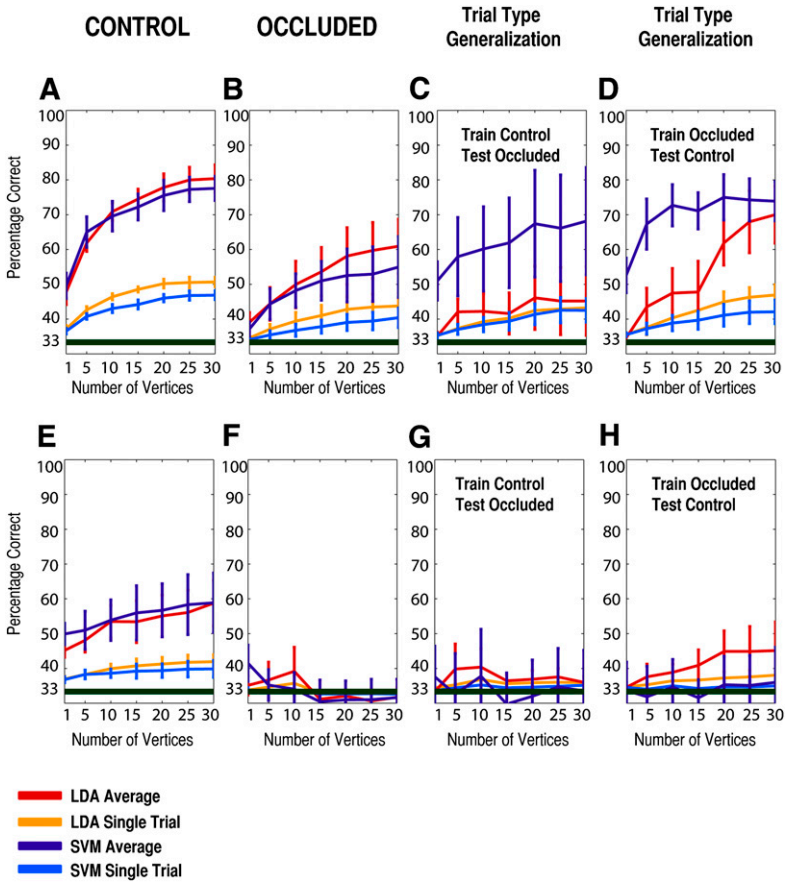

Fig. S4. Pattern classification analysis for experiment 2 split by visual area. Data presentation is as in Fig. S2.

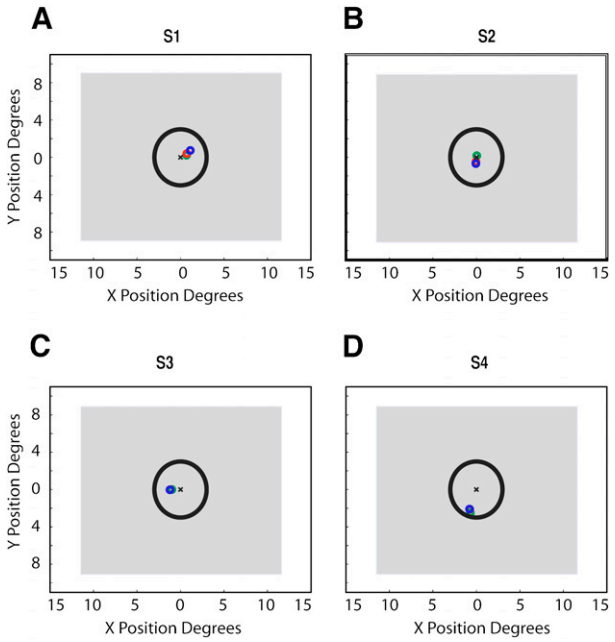

Fig. S5. Mean eye gaze position for each participant and condition in experiment 3. (*A*) The mean eye gaze position for each condition for S1. Each colored circle represents the mean position collapsed across all scanning sessions for one condition (one scene). The black circle represents a radius of 3° outward from fixation, and the shaded gray region represents the actual visual extent of the scenes presented. *B–D* show the same data for the remaining three participants. There were no significant differences across the three scenes shown in terms of the mean eye position on the *x* [$F(2,6) < 1$, $P = 0.66$], or *y* axis [$F(2,6) < 1$, $P = 0.98$], or the mean number of saccades [$F(2,6) = 1.13$, $P = 0.38$]. Note the same pattern holds if subject 4 is excluded from these analyses. Thus eye movements are not responsible for the effects we observe.

Table S1. Statistics for the selected vertices in experiments 1–3 (vertices pooled over V1 and V2)

Experiment	Subject	Vertices (*n*)	MaxT	MinT	MaxS	MinS	nUSigO	nUSigC	USigO	USigC
Experiment 1	1	281	21.38	6.05	1.63	−1.60	32	153	0.21	0.14
	2	354	19.78	2.68	1.64	−1.65	86	120	0.57	0.43
	3	194	31.37	9.66	1.65	−1.62	6	105	0.55	0.01
	4	151	27.50	5.30	1.64	−1.49	0	110	0.67	0.00
	5	199	29.59	12.91	1.64	−1.60	29	67	0.21	0.30
	6	179	27.71	9.87	1.65	−1.65	59	108	0.05	0.00
Experiment 2	1	129	24.25	12.52	1.59	−1.56	0	19	0.33	0.07
	2	187	40.57	3.08	1.64	−1.65	4	105	0.41	0.00
	3	177	29.45	5.62	1.65	−1.63	16	37	0.48	0.55
	4	159	15.28	3.46	1.55	−1.65	18	66	0.67	0.08
	5	100	10.12	1.86	1.58	−1.49	0	79	0.34	0.01
	6	326	11.00	1.84	1.64	−1.63	102	202	0.06	0.02
	7	271	14.05	3.05	1.64	−1.64	66	151	0.11	0.03
Experiment 3	1	86	24.33	6.57	1.64	−1.65	39		0.04	
	2	111	24.44	4.58	1.62	−1.48	31		0.35	
	3	76	25.51	4.89	1.65	−1.65	1		0.38	
	4	244	30.26	3.56	1.63	−1.64	165		0.00	

MaxT, maximum *t* values for the target mapping stimulus; MinT, minimum *t* values for the target mapping stimulus; MaxS, maximum *t* values for the surround mapping stimulus; MinS, minimum *t* values for the surround mapping stimulus; nUniSigO, number of vertices showing significant univariate discrimination between the three scenes for occluded trials; nUniSigC, number of vertices showing significant univariate discrimination between the three scenes for control trials; USigO, *P* value from a test of univariate discrimination pooled across all vertices for occluded trials; USigC, *P* value from a test of univariate discrimination pooled across all vertices for control trials. Note that nUniSigC and USigC are not defined for participants in experiment 3.

Table S2. Statistics for the selected vertices of the V1 patch in experiments 1 and 2

Experiment	Subject	Vertices (*n*)	MaxT	MinT	MaxS	MinS	nUSigO	nUSigC	USigO	USigC
Experiment 1	1	35	16.92	6.05	1.57	−1.58	19	20	0.01	0.02
	2	106	13.62	4.07	1.64	−1.63	46	50	0.24	0.36
	3	103	31.37	15.36	1.65	−0.36	6	103	0.41	0.00
	4	92	27.50	5.30	1.64	−1.44	0	92	0.61	0.00
	5	81	29.59	12.91	1.64	−1.48	26	36	0.16	0.17
	6	84	27.71	9.87	1.62	−1.49	57	83	0.02	0.00
Experiment 2	1	16	24.25	12.79	1.58	−1.54	0	9	0.92	0.01
	2	90	40.57	8.29	1.61	−1.64	2	85	0.65	0.00
	3	93	29.45	6.74	1.65	−1.53	13	29	0.87	0.48
	4	101	14.28	3.75	1.55	−1.65	18	15	0.42	0.62
	5	98	10.12	1.86	1.58	−1.49	0	77	0.35	0.01
	6	314	11.00	1.84	1.64	−1.63	102	190	0.06	0.02
	7	182	13.84	3.05	1.64	−1.63	37	72	0.17	0.07

Table S3. Statistics for the selected vertices of the V2 patch in experiments 1 and 2

Experiment	Subject	Vertices (*n*)	MaxT	MinT	MaxS	MinS	nUSigO	nUSigC	USigO	USigC
Experiment 1	1	246	21.38	8.22	1.63	−1.60	13	133	0.28	0.17
	2	248	19.78	2.68	1.63	−1.65	40	70	0.84	0.16
	3	91	23.11	9.66	1.64	−1.62	0	2	0.85	0.80
	4	59	14.60	8.01	1.62	−1.49	0	18	0.93	0.13
	5	118	26.85	18.86	1.64	−1.60	3	31	0.25	0.37
	6	95	27.25	13.04	1.65	−1.65	2	25	0.22	0.13
Experiment 2	1	113	23.90	12.52	1.59	−1.56	0	10	0.24	0.11
	2	97	19.44	3.08	1.64	−1.65	2	20	0.33	0.07
	3	84	24.85	5.62	1.63	−1.63	3	8	0.20	0.37
	4	58	15.28	3.46	1.51	−1.63	0	51	0.92	0.00
	5	2	8.53	8.34	1.58	1.49	0	2	0.11	0.02
	6	12	5.33	1.96	1.12	−0.88	0	12	0.27	0.00
	7	89	14.05	3.94	1.63	−1.64	29	79	0.09	0.00

Table S4. Mean performance of the SVM and LDA classifiers in experiments 1 and 2

		Experiment 1 (block design)				Experiment 2 (rapid-event–related design)			
Classifier		ST	AV	pST	pAV	ST	AV	pST	pAV
SVM	Occluded	0.50	0.65	1.18×10^{-2}	8.23×10^{-4}	0.39	0.53	4.26×10^{-2}	3.07×10^{-2}
	Control	0.81	0.96	6.11×10^{-6}	5.18×10^{-7}	0.48	0.79	3.32×10^{-5}	7.72×10^{-5}
	G1	0.47	0.70	6.70×10^{-3}	2.95×10^{-3}	0.41	0.61	2.49×10^{-2}	3.66×10^{-2}
	G2	0.57	0.79	1.16×10^{-3}	1.75×10^{-3}	0.42	0.69	1.66×10^{-2}	1.49×10^{-2}
LDA	Occluded	0.48	0.60	6.19×10^{-3}	1.89×10^{-3}	0.42	0.60	1.65×10^{-2}	4.06×10^{-3}
	Control	0.71	0.82	1.02×10^{-6}	2.76×10^{-6}	0.51	0.80	1.52×10^{-5}	3.56×10^{-5}
	G1	0.45	0.55	9.19×10^{-3}	2.41×10^{-3}	0.43	0.45	1.73×10^{-2}	1.26×10^{-1}*
	G2	0.51	0.65	2.57×10^{-3}	7.93×10^{-4}	0.46	0.74	1.24×10^{-3}	2.51×10^{-3}

All analyses are for maximal number of vertices (70), except LDA in experiment 1 (peak performance for 25 vertices). AV, average prediction; ST, single block or single trial prediction (experiment 1 or experiment 2); pAV; *P* value from one-tailed *t* test across subjects against chance for average performance; pST, *P* value from one-tailed *t* test across subjects against chance for single-block or single-trial performance.
*This is the only analysis that is nonsignificant.