Nonstimulated early visual areas carry information about surrounding context

Fraser W. Smith[a,b,1,2] and Lars Muckli[a,b]

[a]Centre for Cognitive Neuroimaging (CCNi), Institute of Neuroscience & Psychology, College of Medical, Veterinary & Life Sciences, and [b]School of Psychology, College of Science & Engineering, University of Glasgow, Glasgow G12 8QB, United Kingdom

Even within the early sensory areas, the majority of the input to any given cortical neuron comes from other cortical neurons. To extend our knowledge of the contextual information that is transmitted by such lateral and feedback connections, we investigated how visually nonstimulated regions in primary visual cortex (V1) and visual area V2 are influenced by the surrounding context. We used functional magnetic resonance imaging (fMRI) and pattern-classification methods to show that the cortical representation of a nonstimulated quarter-field carries information that can discriminate the surrounding visual context. We show further that the activity patterns in these regions are significantly related to those observed with feed-forward stimulation and that these effects are driven primarily by V1. These results thus demonstrate that visual context strongly influences early visual areas even in the absence of differential feed-forward thalamic stimulation.

linear classifier | natural visual scenes | cortical feedback | lateral interaction

It is well known that the majority of input that arrives to a specific neuron in the early visual system comes from other cortical neurons (either local or long-range connections). Such connections provide one way for prior knowledge and context to modulate the responses of neurons in early vision. However, studies that investigate the role of such mechanisms within natural vision are relatively rare. In the present experiments we devised a paradigm in which we analyzed the influence of surrounding context on visually nonstimulated parts of primary visual cortex (V1) and visual area V2. We set out to investigate whether blood oxygen level-dependent (BOLD) functional MRI (fMRI) activity in nonstimulated early visual regions carries information about a surrounding visual context. Our goal was to test the hypothesis that lateral and feedback connections modulate and possibly prime regions of visual cortex by transmitting relevant contextual information.

If this hypothesis is true, we would expect such mechanisms to work most productively when participants are presented with natural visual stimuli, because these stimuli contain a multitude of rich contextual associations across space and time (1). Therefore we presented participants with natural visual scenes in all the experiments reported here. Support for the above hypothesis is provided by recent demonstrations illustrating that V1 receives feedback even in regions that extend beyond the bottom-up stimulated area. For instance, the size of the activated region in retinotopic cortex corresponds to perceived size, not absolute size, in the context of a size illusion (2); moreover, the subjectively perceived apparent-motion illusion activates nonstimulated V1 on the motion trace (3), and object classification triggers discriminable activity at nonstimulated foveal representations (4). Other work has shown that feature-based attention spreads to nonstimulated regions of V2–V4 (5). Therefore contextual influences can extend beyond the range of feed-forward stimulated retinotopic space (6, 7). In the present work we used natural visual scenes to ask whether visually nonstimulated regions of V1 and V2 contain contextual information that permits decoding of the surrounding visual context.

We presented three different natural visual scenes (a car, a boat, or people; Fig. S1*A*) with the lower right quadrant oc-cluded by a uniform white field (Fig. 1*A*, occluded) in a block-design fMRI experiment. Participants maintained fixation and monitored the sequence of images for a change in color (whole image; random occurrence within/across blocks). We independently mapped the cortical representation of the occluded quarter-field in early visual areas V1 and V2, making sure to minimize the risk of any spillover effects from surrounding regions (Fig. 1*B*, Methods, Fig. S2 and *SI Methods*). We then trained a linear classifier to discriminate between the different scenes presented, based purely on the signal from these nonstimulated early visual regions (8–12). After training we used the classifier to decode which scene had been presented to observers in an independent set of test data. If the decoding accuracy of the classifier on the test set reliably exceeds chance we can conclude that activity in nonstimulated early visual regions does indeed discriminate the surrounding context. For the purposes of comparison, we also included blocks in which the full visual scene was presented [i.e., the cortical representation of the occluded quarter-field received visual scene stimulation (Fig. 1*A*, control)], which we refer to as the "natural stimulation" (control) condition.

Results

Fig. 2*A* shows the cross-validated performance of two linear classifiers (Support Vector Machine - SVM and Linear Discriminant Analysis - LDA) in decoding which scene was presented when the cortical representation of the target region received visual-scene stimulation (control condition) as a function of the number of randomly selected vertices entering the classifier (pooled here over V1 and V2). The corresponding event-related average of the fMRI signal is shown in Fig. S1*B*. The BOLD fMRI signal of 70 vertices was sufficient to diagnose correctly which scene was presented with, on average, $96 \pm 2.2\%$ accuracy [SVM classifier: average level prediction, $t(5) = 28.3$, $P = 5.2 \times 10^{-7}$; single-block prediction, $81 \pm 2.8\%$, $t(5) = 17.2$, $P = 6.11 \times 10^{-6}$], greatly above the level of chance (33%). More importantly, Fig. 2*B* shows the performance of the classifiers in decoding when the target region received no visual-scene stimulation (occluded condition). Performance reaches a maximum accuracy of $65 \pm 5.2\%$ here [SVM classifier: average level prediction, $t(5) = 6.16$, $P = 0.0008$; single-block prediction, $50 \pm 5.1\%$, $t(5) = 3.22$, $P = 0.012$], again well above chance (33%). Thus a linear classifier can decode reliably which scene has been shown to participants even when no visual-scene information is presented to these regions of early visual cortex. We note that

NEUROSCIENCE

Fig. 1. Experimental design and cortical mapping procedure. (*A*) The two conditions of visual stimulation used in the current experiments. In occluded trials (*Left*) the lower-right quadrant of the image was occluded with a uniform white field. In control trials (*Right*) there was no occlusion. The same three scenes were shown in both occluded and control conditions (experiments 1 and 2). The participant always had to fixate the central fixation marker (very small checkerboard) and either detect a change in frame color (experiment 1) or perform a one-back repetition task (experiment 2). Note that the black bars highlighting the occluded region are for display purposes only. The whole image spanned 22.5 × 18°, with the occluded region spanning ≈11 × 9°. (*B*) Cortical mapping procedure used to map the retinotopic representation of the occluded region. In a standard block-design protocol, participants were presented with contrast-reversing checkerboards in either the target region (green) or the surround region (red). We defined a patch of V1 and a patch of V2 from the contrast of target stimulus minus surround stimulus (shown here on the inflated cortical surface [left hemisphere] for one representative participant). Any vertex taken as representing the occluded portion of the main stimulus had to meet two additional constraints: a significant effect for the target stimulus alone (target > baseline; t > 1.65) and, crucially, a nonsignificant effect for the surround stimulus alone (surround stimulus = baseline; absolute t < 1.65).

Fig. 2. Pattern-classification analysis for experiments 1 and 2. (*A–D*) Classifier performance for experiment 1 (block design). *A* shows performance (percentage correct) for 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 (Av.) and single-block prediction (SB). Note that vertices are pooled across V1 and V2 in this analysis. Performance is averaged across participants (error bars represent one SEM). Chance performance is indicated by the dark green bar at 33%. *B* shows the same data for occluded trials. *C* and *D* show the results for the trial-type generalization analyses. *E–H* show classifier performance for experiment 2 (rapid-event–related design). *E–H* are arranged as in *A–D* but report performance for average (Av.) and single-trial (ST) prediction.

the difference in performance between the two types of classifier (SVM vs. LDA) is greatest as the number of vertices approaches the number of training examples (here 54), leading to overfitting in the case of LDA, a well-known problem in machine learning (13). The SVM classifier, however, is a regularized classifier and therefore does not suffer from this problem, maintaining high performance as the number of vertices increase (13, 14).

What is the source of this contextual information? To investigate the nature of the information driving successful classifier performance in the absence of visual-scene stimulation, we explored the relation between the activity patterns observed in the two different types of trial (control and occluded) presented to observers (11, 15). We trained a new classifier on one set of trials (e.g., control) and tested its generalization performance for the other type of trial (e.g., occluded). If the classifier could decode above chance in this situation, that performance would be evidence for some degree of commonality in the activity patterns generated in these regions of early visual cortex in trials where visual-scene stimulation is present (control) vs. absent (occluded). We found our classifiers trained on control trials still could discriminate above chance when tested with occluded trials (Fig. 2*C*), reaching 70 ± 8% accuracy with 70 vertices [SVM classifier: average level prediction, $t(5) = 4.59$, $P = 0.0029$; single-block prediction, $47 ± 3.7\%$, $t(5) = 3.74$, $P = 0.0067$; chance: 33%]. Similarly if we trained on occluded trials and tested with control trials (Fig. 2*D*), the classifiers also could perform well above chance [SVM classifier: average level prediction, $79 ± 8.9\%$ accuracy, $t(5) = 5.19$, $P = 0.0017$; single block prediction, $57 ± 4.1\%$, $t(5) = 5.71$, $P = 0.0012$; chance: 33%]. The linear classifiers thus show that the activity patterns elicited in the presence vs. absence of visual-scene stimulation are similar to a nontrivial degree. The higher performance in trial-type generalization (Fig. 2 *C* and *D*) than in within-trial generalization in the occluded
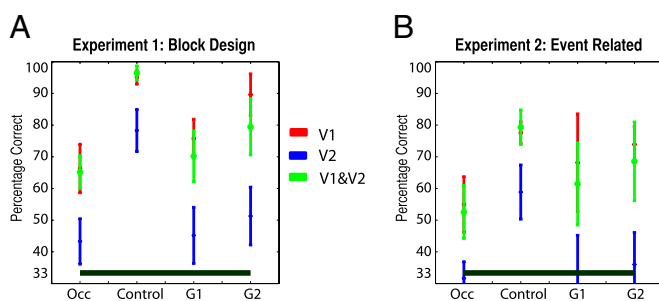
condition (Fig. 2*B*) might result from the greater number of trials used to train and test the classifier in the trial-type generalization (72 for training and testing in the trial-type generalization vs. 54 and 18, respectively, in the within-trial generalization).

We have shown that nonstimulated regions of early visual cortex carry information about surrounding context and furthermore that the activity patterns in these nonstimulated regions are similar in occluded and control trials. The previous experiment used a block-design paradigm that leads to robust pattern estimates due to averaging over many volumes for each stimulus presentation but it is unclear how what happens during a block of stimulation (in this case, 12 s) compares with any one trial of perception (i.e., the standard design in visual cognition). Within the duration of a block, for instance, it is unclear what other cognitive processes the subject is engaged in; therefore, it is beneficial if the result can be replicated in single-trial conditions. To increase the generalizability of our results, we conducted a second experiment with a rapid-event–related design (4-s intertrial interval), with a different task (one-back repetition detection) and seven new participants. Fig. 2 *E–H* shows the results of this experiment. In control trials (Fig. 2*E*), the linear classifier reached $79 ± 5.5\%$ accuracy [SVM classifier: average level prediction, $t(6) = 8.41$, $P = 7.71 × 10^{-5}$; single-trial prediction, $48 ± 1.5\%$, $t(6) = 9.76$, $P = 3.32 × 10^{-5}$; chance: 33%]. Most importantly, the classifier again successfully decoded which scene had been shown in occluded trials (Fig. 2*F*), reaching an accuracy of $53 ± 8.4\%$ [SVM: average level, $t(6) = 2.30$, $P = 0.031$]. Note that the classifier also reliably decoded occluded trials based purely on single-trial data [SVM: $39 ± 2.8\%$, $t(6) = 2.06$, $P = 0.043$; chance: 33%]. To corroborate this result, the weakest we report here, we conducted a group-level permutation test (*SI Text*) that gave $P = 1 × 10^{-4}$. Thus, with this different fMRI design, task, and participants, the linear classifier successfully decoded which occluded scene was presented on the basis of just 1 s of stimulation. We

note, however, that the advantage of the SVM over LDA is much reduced in experiment 2 compared with experiment 1; in fact, in experiment 2, LDA outperforms or is equal to the SVM. The reasons underlying this pattern are not entirely clear. One likely influential factor is that, because of the greater number of training observations in experiment 2 (a minimum of 180) compared with experiment 1, overfitting is much less of a problem for the LDA classifier. A second factor is the greater intersubject variability for the SVM than for LDA in experiment 2. In any case, taken together, our results provide strong evidence that a linear classifier indeed can discriminate between the different scenes presented, even when the target region receives no visual-scene stimulation.

Using the event-related data, we tested the performance of the linear classifier in generalizing across trial types, and again we found that the classifiers could generalize successfully (Fig. 2 *G* and *H*). The classifier trained on occluded trials and tested with control trials had an average level prediction accuracy of 69% ± 12.4% [SVM classifier: $t(6) = 2.84$, $P = 0.015$; chance: 33%] and a single-trial prediction accuracy of 42 ± 3.3% [SVM: $t(6) = 2.75$, $P = 0.017$]. On the other hand, the classifier trained on control trials and tested with occluded trials had an average level prediction accuracy of 61 ±13% [SVM: $t(6) = 2.17$, $P = 0.037$] and a single-trial prediction accuracy of 41 ± 3.2% [SVM: $t(6) = 2.45$, $P = 0.025$]. Thus with this experiment we have replicated successfully the main pattern of results obtained in experiment 1.

We have shown in two independent experiments that non-stimulated regions of early visual cortex do indeed carry information about the surrounding context. Furthermore, we have shown that the activity patterns in these nonstimulated regions are similar in occluded and control trials. What is the contribution of each visual area to the performance we observe? To explore this question we reran our classifier analyses independently for patches of V1 and V2 (*SI Methods*). Fig. 3*A* shows the asymptotic performance (defined as performance with the maximum number of vertices) obtained by the SVM classifier (average level prediction) for each visual area (V1, V2, or pooled) for the three types of analysis (control, occluded, or trial type generalization) of experiment 1. Fig. 3*B* shows the same data for experiment 2. (Figs. S3 and S4 give the full classifier results for both experiments split by visual area, and Tables S1–S3 give vertex details.) It is clear from these plots that V1 greatly outperforms V2 across both experiments and each different analysis, reaching or even exceeding the performance obtained by pooling across V1 and V2, with a smaller number of vertices. Thus, V1 responses are the primary determinant of the ability of our classifiers to discriminate successfully between the three scenes, regardless of the

specific trial type (control, occluded, or trial type generalization analyses) considered.

What kind of mechanism could be responsible for these effects? One plausible explanation would be the presence of an autoassociative memory-based process active at the level of early visual areas (16–18). As the experiment progresses, relevant early visual neurons would develop a strong expectation of the structure that is present in the occluded region of the stimuli, obtained from presentation of that information in control trials. Therefore in occluded trials, the three-quarters of the stimulus that is presented would be a very good cue to retrieve the remaining (occluded) part of the stimulus.

To test this possibility, we conducted a third experiment in which participants were presented only with occluded trials (i.e., they never saw the actual visual-scene information of the occluded region). We present the results of this experiment in Fig. 4 *A–D*. For three of the four participants, clear discrimination can be seen (SVM average level prediction). For the remaining participant the SVM classifier (average level) does not perform well; however, the LDA classifier suggests that discrimination ability is present within this subject also. Thus the effect we report clearly goes beyond autoassociative recall based on memories of the actual visual-scene information of the occluded region. (See also Fig. S5 for plots of subjects eye gaze position in this experiment.)

One additional explanation is that information regarding the distribution of low-level visual features (e.g., luminance and contrast, energy at each spatial frequency, and orientation) in the stimulated area is transmitted to the nonstimulated region. To address this possibility, we ran two more participants on a modified version of experiment 3 in which we explicitly controlled these low-level image properties within the naturally stimulated region (Fig. 4*E* and *SI Methods*, "Low-Level Image Control"). Perfor-
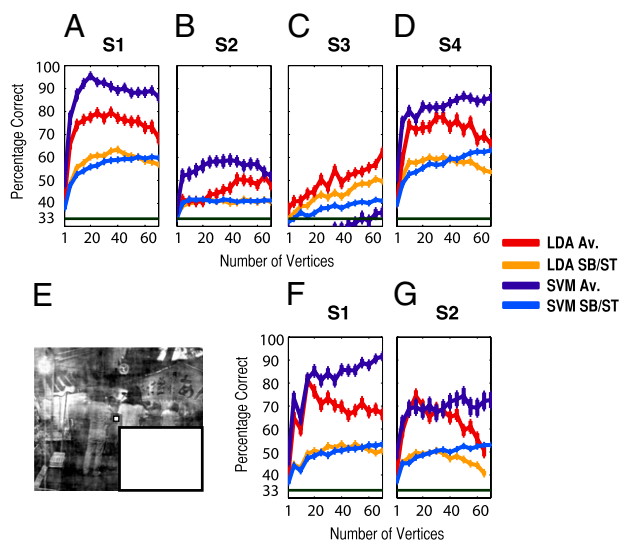
Fig. 4. Pattern classification analysis for experiments 3 and 4. (*A–D*) Classifier performance per participant for experiment 3. *A* shows performance (percentage correct) for the two linear classifiers LDA and SVM in decoding which scene was presented in occluded trials as a function of the number of vertices entering the classifier, for both average level (Av.) and single-block prediction (SB), for subject 1. Note that vertices are pooled across V1 and V2 in this analysis. Performance is averaged across sampling iterations and cross-validation cycles (error bars represent one SEM). Chance performance is indicated by the dark green bar at 33%. *B–D* show the same information for the remaining participants in experiment 3. *E* shows an example of a scene stimulus shown in experiment 4 (*SI Methods*, "Low-Level Image Control"). *F* and *G* show performance data for the two participants who took part in experiment 4.


Fig. 3. Asymptotic performance of the SVM classifier in experiments 1 and 2. Asymptotic classifier performance (average level of prediction) for each visual area (V1, V2, and pooled) and each different analysis [occluded trials, control trials, generalization 1 (G1: train control, test occluded), and generalization 2 (G2: train occluded, test control)] for experiments 1 (*A*) and 2 (*B*). Note that "asymptotic" is defined as performance for the maximum number of vertices (here 70 for V1 and V2 pooled, 30 for each area considered independently). Error bars show one SEM across participants. Occ, occluded.

NEUROSCIENCE

mance clearly remained well above chance (Fig. 4 *F* and *G*) despite our control of these low-level image properties, demonstrating that the observed contextual sensitivity is not carried by these specific low-level visual properties.

Finally, to what extent could spreading activity (lateral connections) or spillover activity account for the present results? We performed an additional analysis investigating the weights from the pattern-classification analyses to address this question (*SI Methods*, "Weight Analysis"). Specifically we correlated the absolute value of the weights from the SVM and LDA classifiers leading to maximum performance on occluded trials with the *t* values of both the target and the surround mapping conditions (Fig. 1*B*). The absolute value of the weight at each vertex indicates the relative influence of each vertex to the classifier's solution (e.g., refs. 11, 19). The logic is that a positive correlation of high (absolute) weights with high *t* values for the surround mapping stimulus might suggest a possible influence of spillover signal (or spreading activity). On the other hand, a positive correlation with the target mapping stimulus would indicate that the more important vertices in the classifier decision function are those with a strong signal to the target stimulus, as we might expect to be the case purely on signal-to-noise considerations. Although results differed between LDA and SVM classifiers (*SI Methods*, "Weight Analysis"), the results with the LDA classifier confirmed that it is possible to decode the surrounding context with a set of weights that have no significant positive relationship with the surround mapping stimulus (and, in fact, may have a suggestive negative or negative correlation) and a significant positive relation to the target mapping stimulus (experiments 1 and 2; see also Table S4). The results with the LDA classifier thus argue against interpreting our data in terms of simple spillover activity.

Discussion

Across four independent experiments we have shown that non-stimulated early visual regions (V1 and V2) carry information about a surrounding visual context. We have shown, moreover, that the activity patterns in the nonstimulated early visual regions are related significantly to the activity observed in those regions under feed-forward stimulation (experiments 1 and 2). We have shown further that this effect is driven largely by V1 (experiments 1 and 2), that it does not depend on memory of complete images (experiment 3), and that it is not carried by several basic low-level visual properties (experiment 4). Finally, with the classifier weight analysis we have provided some reassurance that simple spreading activity or spillover activation from feed-forward–stimulated neighboring regions is unlikely to account for the results we report here, at least for the LDA classifier.

What type of mechanism could account for all experiments we report? This question concerns how information can be transmitted from regions representing the surrounding context to regions representing the nonstimulated quadrant. Bayesian models of human vision (20–22) offer one possible explanation in which the surrounding visual context would bias cortical feedback to the nonstimulated early visual areas. Such feedback could be spatially precise, resulting in something like "filling in," or spatially diffuse, leading, for example, to a general expectation of some property (i.e., categorical: human bodies or cars) in the target region (1, 23). Although we did not set out to test such predictive accounts directly, our paradigm could be adapted to test for prediction error signals in the target region of V1 after presenting appropriate context inducing stimuli in surrounding areas of the visual field (24, 25).

In principle, however, the type of visual information that is transmitted to the target region could concern low-level visual features (such as distributions of contrast, luminance, energy, spatial frequencies, orientations, and so forth) or, as mentioned earlier, higher-level visual features (such as contours, categories). Although we demonstrated in experiment 4 that several basic low-level visual features (global luminance and contrast, energy at each

spatial frequency, and orientation) were not responsible for generating the observed context sensitivity, many other low-level features (e.g., local and global distribution of orientations) remain to be tested before we can exclude a mechanism based on signals from the surrounding context carrying information about distributions of low-level visual information (local or global) to the nonstimulated areas. Further, although we have demonstrated a significant degree of "similarity" between the responses present in the target region as a result of feed-forward or surrounding area-only stimulation (experiments 1 and 2), this similarity does not imply that the signals necessarily reflect identical information. Thus the nature of the information that is carried by the signal from the surrounding context to the nonstimulated area still requires further study.

One additional explanation of our findings concerns the possible allocation of attention to expected features in the occluded region. An intriguing report (5) showed that feature-based attention spreads even to nonstimulated parts of the visual field. The authors showed that attended direction of motion (45° vs. 135°) in the stimulated visual field (left or right) could be decoded from ipsilateral and thus nonstimulated early visual regions, specifically V2, V3, and V4. Based on this report, one could propose that attention is allocated to certain locations in the occluded region where participants expect to find object features. These locations would differ across the three different scenes presented to participants and therefore could be responsible for the classification results that we observe. Note that this account is not independent of the predictive accounts given earlier but shifts the emphasis to neural mechanisms that implement visual attention.

At a neurophysiological level, it is known that most of the input received by any given V1 neuron comes from intra- and interareal projections, with the minority coming from thalamic afferents (26, 27). As such, monosynaptic lateral connections (i.e., single-synapse intracortical projections) might play a role in explaining how contextual information is promulgated from an area of cortex that receives feed-forward stimulation to an area that does not (a distance here of 2° on the fixation diagonal), possibly by transmitting scene-specific information across the low-contrast summation fields of V1 neurons (28, 29). These connections, however, are unlikely to account for all the activity in nonstimulated V1, because they are not wide enough to span the entire nonstimulated region. Furthermore, laterally transmitted effects would be predicted to be strongest close to the border of the occluded region, a prediction that is not consistent with the results of the classifier weight analyses (*SI Methods*, "Weight Analysis"). Therefore interareal projections (i.e., cortical feedback) are likely to be critical in explaining the present effects.

A potential mechanism through which cortical feedback might contribute to the observed context effect would be by direct connections between mirror-symmetric parts of the visual field, as reported in primate middle temporal visual area (30). Analysis of spontaneous BOLD fMRI fluctuations in anesthetized primates also has revealed correlations between mirror-symmetric parts of the visual field that are pronounced along the horizontal meridian (31). Such mirror-symmetric sensitivity in higher visual areas certainly could play a role in explaining where such cortical feedback originates. Although at present we cannot disentangle the contribution from cortical feedback and lateral interaction, we believe that the fact that such information is transmitted at all from stimulated to nonstimulated regions highlights the ubiquitous role that context might play in shaping the activity of all neurons in early vision. Indeed, several authors recently have pointed out that, to increase our understanding of V1 beyond the best current models, a deeper consideration of the role of context (both spatial and temporal) and of how such contextual information is transmitted is required (32–34).

In the present experiments we found that V1 was the primary driving force for the context effect. Indeed, there was barely any evidence for context sensitivity in V2 (Fig. 3). Although in most

cases significant classification performance in V1 co-occurs with similar performance in V2 (9, 10 experiment 1, 15; refs. 35 and 36), these studies typically reflect above-threshold stimulation protocols. Interestingly, one exception (10, experiment 2) involved decoding subjectively invisible stimulus orientation and led to significant decoding only in V1, mirroring the present results. It has been shown, furthermore, that V1 contains more reliable information for binary image reconstruction than V2 (36). Taken together, these results make the asymmetry that we found in performance across V1 and V2 less surprising.

Much recent work has demonstrated that a multitude of important visual information can be extracted from the fMRI BOLD response in early visual areas when taken as a multivariate quantity, allowing discrimination between visual features such as orientation and motion direction (9, 35), discrimination between scene categories (37), discrimination of the identity of natural images (38), and even allowing the reconstruction of arbitrary binary contrast images (36). Here, extending these lines of research, we have shown that discrimination between individual scenes is possible even in nonstimulated regions of the primary visual cortex.

Taken together, we believe that our data demonstrate the crucial role of context, carried by the physiological mechanisms of feedback and/or lateral interaction, in driving activity even in nonstimulated early visual areas. We further believe these data are consistent with theoretical views concerning the importance of predictive codes in the visual system (16, 20, 39, 40).

Methods

Subjects. Six subjects took part in experiment 1, seven in experiment 2, four in experiment 3, and two in experiment 4 (*SI Methods*).

Paradigm. Participants were presented with partially occluded natural visual scenes (Fig. 1*A*). There were three such scenes, one instance of each (Fig. S1*A*). In experiments 1 and 2, the same scenes also were presented nonoccluded as a control. Participants had to maintain fixation and detect a change in frame color (experiment 1, block design), perform one-back repetition (experiment 2, rapid-event–related design), or detect a change in the color of the fixation marker (experiments 3 and 4, block design). Independently, we mapped the cortical representation of the target region in V1/V2 (Fig. 1*B* and *SI Methods*) by presenting contrast-reversing checkerboards (4 Hz) in either the target or the surround mapping position (Fig. 1*B*) in a standard block-design experiment. Vertices were selected from the defined regions of V1 and V2 that met the criteria of significant effect for target alone and a nonsignificant effect for the surround alone (*SI Methods*).

Structural and Functional MRI. MRI was performed at 3 Tesla using standard MRI parameters. Anatomical data were transformed to Talairach space, and the cortical surface was reconstructed. fMRI time series were preprocessed using standard parameters (no smoothing) and coaligned to an anatomical dataset (*SI Methods*). A general linear model was used to estimate the activity patterns on each single block (or trial).

Pattern Classification. We trained two linear pattern classifiers, independently for occluded and control trials (*SI Methods*), to learn the mapping between a set of brain-activity patterns and the presented scene. We then tested the classifiers on an independent set of test data (leave one run out cross-validation). We chose input features (vertices) randomly from the set that met our mapping criteria, initially pooled across V1/V2 but later also split by V1/V2 (*SI Methods*). We also performed trial-type generalization analyses in which the classifiers (in experiments 1 and 2) were trained on one type of trial (e.g., occluded) and were tested on the other type of trial (e.g., control).

1. Bar M (2004) Visual objects in context. *Nat Rev Neurosci* 5:617–629.
2. Murray SO, Boyaci H, Kersten D (2006) The representation of perceived angular size in human primary visual cortex. *Nat Neurosci* 9:429–434.
3. 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.
4. Williams MA, et al. (2008) Feedback of visual object information to foveal retinotopic cortex. *Nat Neurosci* 11:1439–1445.
5. Serences JT, Boynton GM (2007) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55:301–312.
6. Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L (2010) Stimulus predictability reduces responses in primary visual cortex. *J Neurosci* 30:2960–2966.
7. Muckli L (2010) What are we missing here? Brain imaging evidence for higher cognitive functions in primary visual cortex V1. *Int J Imaging Syst Technol* 20:131–139.
8. 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.
9. Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
10. Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
11. 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.
12. 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.
13. Duda RO, Hart PE, Stork DG (2001) *Pattern Classification* (John Wiley & Sons, New York), 2nd Ed.
14. Hastie T, Tibshirani R, Friedman J (2001) *The Elements of Statistical Learning* (Springer, New York).
15. Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
16. Hawkins J, Blakeless S (2004) *On Intelligence* (Times Books, New York).
17. McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102:419–457.
18. Nadel L, Samsonovich A, Ryan L, Moscovitch M (2000) Multiple trace theory of human memory: Computational, neuroimaging, and neuropsychological results. *Hippocampus* 10:352–368.
19. 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.
20. Lee TS, Mumford D (2003) Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 20:1434–1448.
21. Yuille A, Kersten D (2006) Vision as Bayesian inference: Analysis by synthesis. *Trends in Cognitive Sciences* 10:301–308.
22. Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian inference. *Annu Rev Psychol* 55:271–304.
23. Peelen MV, Fei-Fei L, Kastner S (2009) Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature* 460:94–97.
24. den Ouden HE, Friston KJ, Daw ND, McIntosh AR, Stephan KE (2009) A dual role for prediction error in associative learning. *Cereb Cortex* 19:1175–1185.
25. den Ouden HE, Daunizeau J, Roiser J, Friston KJ, Stephan KE (2010) Striatal prediction error modulates cortical coupling. *J Neurosci* 30:3210–3219.
26. Douglas RJ, Martin KAC (2004) Neuronal circuits of the neocortex. *Annu Rev Neurosci* 27:419–451.
27. Douglas RJ, Martin KAC (2007) Mapping the matrix: The ways of neocortex. *Neuron* 56:226–238.
28. Angelucci A, Bressloff PC (2006) Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. *Prog Brain Res* 154:93–120.
29. Angelucci A, et al. (2002) Circuits for local and global signal integration in primary visual cortex. *J Neurosci* 22:8633–8646.
30. Maunsell JH, van Essen DC (1983) The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J Neurosci* 3:2563–2586.
31. Vincent JL, et al. (2007) Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447:83–86.
32. Carandini M, et al. (2005) Do we know what the early visual system does? *J Neurosci* 25:10577–10597.
33. Masland RH, Martin PR (2007) The unsolved mystery of vision. *Curr Biol* 17:R577–R582.
34. Olshausen BA, Field DJ (2005) How close are we to understanding v1? *Neural Comput* 17:1665–1699.
35. Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
36. Miyawaki Y, et al. (2008) Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron* 60:915–929.
37. Walther DB, Caddigan E, Fei-Fei L, Beck DM (2009) Natural scene categories revealed in distributed patterns of activity in the human brain. *J Neurosci* 29:10573–10581.
38. Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
39. Rao RP, Ballard DH (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2:79–87.
40. Mumford D (1992) On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol Cybern* 66:241–251.

NEUROSCIENCE