

Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex

Peter Kok,¹ Janneke F.M. Jehee,¹ and Floris P. de Lange^{1,*}

¹Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, 6500 HB Nijmegen, The Netherlands

*Correspondence: floris.delange@donders.ru.nl

<http://dx.doi.org/10.1016/j.neuron.2012.04.034>

SUMMARY

Prior expectations about the visual world facilitate perception by allowing us to quickly deduce plausible interpretations from noisy and ambiguous data. The neural mechanisms of this facilitation remain largely unclear. Here, we used functional magnetic resonance imaging (fMRI) and multivariate pattern analysis (MVPA) techniques to measure both the amplitude and representational content of neural activity in the early visual cortex of human volunteers. We find that while perceptual expectation reduces the neural response amplitude in the primary visual cortex (V1), it improves the stimulus representation in this area, as revealed by MVPA. This informational improvement was independent of attentional modulations by task relevance. Finally, the informational improvement in V1 correlated with subjects' behavioral improvement when the expected stimulus feature was relevant. These data suggest that expectation facilitates perception by sharpening sensory representations.

INTRODUCTION

Top-down expectations about the visual world can facilitate perception by allowing us to quickly deduce plausible interpretations from noisy and ambiguous data (Bar, 2004). However, the neural mechanisms of this facilitation are largely unknown.

A theory that has gained growing popularity in the last decade surmises that vision can be cast as a process of hierarchical Bayesian inference, in which higher order cortical regions provide guidance to lower levels, thereby facilitating sensory processing (Friston, 2005; Lee and Mumford, 2003; Summerfield and Kochlin, 2008; Yuille and Kersten, 2006). Within this framework, it has been put forward that higher order regions may suppress the predictable, and hence redundant, neural responses in early sensory regions that are consistent with current high level expectations (Mumford, 1992; Murray et al., 2002; Rao and Ballard, 1999), resulting in a sparse and efficient coding scheme (Jehee et al., 2006; Olshausen and Field, 1996). An alternative possibility is that higher order regions may rather “sharpen” sensory representations in early cortical areas, by suppressing lower order neural responses that are *inconsistent* with current expectations (Lee and Mumford, 2003). This could

be done either directly, through inhibitory feedback, or indirectly, by excitatory feedback to neurons representing the expected feature, which in turn engage in competitive interactions with alternative representations at the lower level (Spratling, 2008). Such a coding scheme would result in a “sharpening” of the population response in early sensory regions for expected percepts. It should be noted that both these mechanisms are incorporated in a more recent model of predictive coding (Friston, 2005), which posits two functionally distinct subpopulations of neurons, encoding the conditional expectations of perceptual causes and the prediction error, respectively (Jehee and Ballard, 2009; Rao and Ballard, 1999). In this scheme, high-level predictions explain away prediction error, thus silencing error neurons, while neurons encoding sensory causes rapidly converge on the (correctly) predicted causes, yielding a relatively sharp population response.

While empirical studies have provided some empirical support for both the above scenarios by showing a reduction of neural activity in early sensory regions as a result of top-down expectation (Alink et al., 2010; den Ouden et al., 2009; Kok et al., 2011; Meyer and Olson, 2011; Murray et al., 2002; Summerfield et al., 2008; Todorovic et al., 2011), these studies could not adjudicate between these two models and answer the question of *how* top-down expectation alters sensory processing. Here, we capitalize on the fact that these hypotheses make opposite predictions about how expectation changes the amount of information present in these regions. If expectation operates by suppressing neural responses that are consistent with the current expectation, the activity reduction in early sensory cortex should be accompanied by a reduction of the sensory representation in this region. If, on the other hand, expectation sharpens the population response, the activity reduction in early sensory cortex should be accompanied by an improved sensory representation in this region. We adjudicated between these hypotheses by noninvasively measuring neural activity and representational content in the early visual cortex of human volunteers, using functional magnetic resonance imaging (fMRI) and multivariate pattern analysis (MVPA) techniques (Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2005). Our results provide evidence for a sharpening account of expectation, in which overall neural activity is reduced, yet the stimulus representation is enhanced by expectation.

RESULTS

During each trial, subjects were presented with two consecutively presented grating stimuli. Before each trial, we induced




Figure 1. Experimental Paradigm and Results

On each trial, subjects were presented with two consecutive grating stimuli, differing slightly in terms of orientation, contrast, and spatial frequency. In separate blocks, subjects performed either an orientation task (“Was the second grating rotated clockwise or anticlockwise with respect to the first?”) or a contrast task (“Was the second grating of higher or lower contrast than the first?”). The grating stimuli were preceded by an auditory cue, which predicted (with 75% validity) the overall orientation of the gratings ($\sim 45^\circ$ or $\sim 135^\circ$).

an expectation about the overall orientation ($\sim 45^\circ$ or $\sim 135^\circ$) of these gratings by means of an auditory cue (Figure 1 and Experimental Procedures). Subjects had to perform either an orientation task on the stimuli (indicate whether the second grating was slightly tilted clockwise or anticlockwise with respect to the first) or a contrast task (indicate whether the second grating had higher or lower contrast than the first), thereby manipulating the task relevance of the expectation.

Behavioral Results

Behavioral data confirmed that subjects were able to discriminate small differences in orientation (3.5° with 81.8% accuracy) and contrast (4.5% with 75.1% accuracy). Angular and contrast differences between the two gratings were manipulated throughout the experiment by an adaptive staircase procedure, for trials containing expected and unexpected orientations separately (see Supplemental Experimental Procedures available online). This was done to rule out a potential confound of task difficulty with the effects of expectation on neural activity. For the orientation task, the staircase procedure adjusted the angle difference to a smaller value for expected than unexpected trials (mean angle difference of 3.4° versus 3.8° : $t_{17} = 2.8$, $p = 0.013$), while keeping accuracy roughly equated (81% versus 84%: $t_{17} = -1.9$, $p = 0.070$), suggesting that expectation had a facilitatory effect on perceptual performance. For the contrast task, there was a nonsignificant trend toward slightly smaller contrast differences for trials containing expected than unexpected orientations (mean contrast difference of 4.3% versus 5.0%: $t_{17} = 1.9$, $p = 0.075$), while accuracy was again roughly equated (74% versus 78%: $t_{17} = -1.9$, $p = 0.077$). Reaction times did not differ between conditions (orientation task: mean RT = 761 ms, for both expected and unexpected trials; contrast task: mean RT = 765 ms for expected, and 767 ms for unexpected trials).

fMRI Results

Neuroimaging data showed that gratings with an expected orientation evoked a reduced response in primary visual cortex, compared to gratings with an unexpected orientation (Figure 2A, bars), in line with previous results (Alink et al., 2010; den Ouden et al., 2009). This neural suppression by expectation was robustly present during both tasks ($F_{1,17} = 14.3$, $p = 0.002$) and did not differ between tasks ($F_{1,17} = 1.4$, $p > 0.1$). This expectation-induced suppression was also observed in V2 and V3 (Figure S1A). There were no overall activity differences in these regions between tasks (all $F_{1,17} < 1$, $p > 0.1$), which is expected

given that these regions are involved in processing both contrast and orientation of stimuli.

Next, we asked whether the reduction of activity in V1 was paired with a decrease or increase in representational content (or stimulus information) in this area. In order to investigate this issue, we used MVPA methods (see Experimental Procedures) to classify the overall orientation of the two gratings presented in each trial ($\sim 45^\circ$ or $\sim 135^\circ$). If orientation classification performance is selectively enhanced/reduced for expected gratings (compared with unexpected gratings), then this would imply that expectation increases/decreases the orientation-selectivity of responses in V1. First, in line with earlier reports (Jehee et al., 2011; Kamitani and Tong, 2005), we found that task relevance enhanced orientation classification accuracy: accuracy was overall higher during the orientation task than during the contrast task ($F_{1,17} = 8.2$, $p = 0.011$; Figure 2A). Critically, despite the reduction in neuronal response, MVPA orientation classification accuracy was further improved for gratings with an expected orientation, compared to an unexpected orientation ($F_{1,17} = 8.3$, $p = 0.010$, Figure 2A). The effects of task relevance and prior expectation were additive and did not interact ($F < 1$, $p > 0.1$). These results were obtained using the 150 most stimulus-responsive voxels (as determined through an independent functional localizer; see Supplemental Experimental Procedures), but the effects were largely independent of the amount of voxels selected (Figures 2B and 2C). Unlike in V1, expectation did not significantly affect orientation classification accuracy in V2 and V3 (Figure S1). This difference between V1 and higher-order visual areas might be due to stimulus characteristics (e.g., the high spatial frequencies in the grating stimuli may have preferentially activated V1), or they might represent a real difference in the extent to which top-down expectation affects representations in V1 versus V2 and V3, as has been previously suggested (Smith and Muckli, 2010).

If the expectation-induced reduction of neural activity reflects a sharpening of neural activity, it might be expected that the effect is strongest in neurons preferring orientations different from the currently presented orientation, while neurons preferring the presented orientation are relatively unsuppressed. To examine this, we calculated the expectation suppression separately for voxels preferring the presented and the non-presented orientation (see Supplemental Experimental Procedures for details). Indeed, expectation suppression was significantly greater in the latter set of voxels, in line with a sharpening account of expectation ($t_{17} = 2.2$, $p = 0.039$; Figure S2A). This

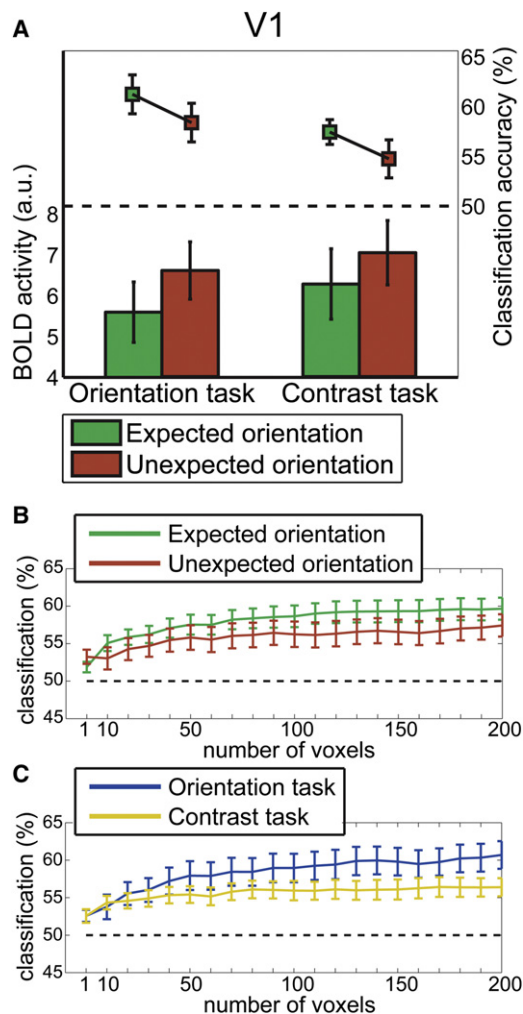


Figure 2. Effects of Expectation and Task Relevance in V1

(A) Gratings with an expected orientation (i.e., the orientation predicted by the auditory cue) evoked less activity in primary visual cortex than gratings with an unexpected orientation, irrespective of the task-relevance of orientation. In contrast, MVPA orientation classification accuracy of the grating orientation ($\sim 45^\circ$ or $\sim 135^\circ$) in V1 was higher for expected than unexpected orientations. Orientation classification performance was also overall higher during the orientation task (where orientation was task-relevant) than during the contrast task (where orientation was task irrelevant). Error bars indicate SEM (B and C) Effects of expectation (B) and task-relevance (C) on MVPA orientation classification accuracy were consistent across a wide range of selected voxels. Error bars indicate SEM.

account further predicts a quantitative relationship between orientation preference and expectation suppression: when the preference of a voxel for the presented orientation is stronger, the expectation suppression should be smaller. This prediction was confirmed by a significant negative correlation across voxels between their preference for the presented orientation and the corresponding expectation suppression effect ($r = -0.292$, $p < 0.001$; Figure S2B).

Is the expectation-induced reduction of neural activity and increase in representational content relevant for perception?

To explore this issue, we assessed the relationship between behavioral and neural effects of expectation. We quantified orientation discrimination thresholds separately for expected and unexpected gratings during the orientation task. If expectation-induced behavioral benefits are linked to increased representational content in V1, we would expect a correlation between intersubject variation in the expectation-induced reduction in orientation discrimination threshold (behavioral improvement) and the expectation-induced improvement in MVPA orientation classifier performance (neural improvement). Indeed, we observed such a correlation ($r = 0.53$, $p = 0.023$; Figure 3A). Since the orientation discrimination threshold was directly related to the angle difference between gratings, due to the staircase procedure, we applied the same analysis to the data from the contrast task, and found no such relationship there ($r < 0.01$, $p = 0.990$; Figure 3B). This precludes an explanation of our results in terms of physical stimulus differences, since these were roughly equal between tasks (see Supplemental Experimental Procedures). Further analyses confirmed that differences in MVPA orientation classification accuracy were not related to physical stimulus differences (see Supplemental Experimental Procedures for a full description). First, no across-subject correlations were found between stimulus differences and MVPA orientation classification accuracy, neither within nor between expectation conditions (Figure S3). Second, there were no within-subject correlations between trial-by-trial orientation angle differences and MVPA accuracy, for either expected ($r = -0.02$, $t_{17} = -1.3$, $p = 0.220$) or unexpected ($r = -0.04$, $t_{17} = -1.1$, $p = 0.287$) gratings. Third, binning trials on the basis of angle differences revealed no differences in MVPA orientation classification in any of the early visual cortex regions for “large orientation difference versus small orientation difference” (Figure S2C). Fourth, an analysis in which stimulus differences between conditions were removed by selecting a subset of the trials still revealed significantly higher orientation classification accuracy for expected than unexpected grating orientations (62.5% versus 57.5%, $t_{13} = 2.1$, $p = 0.028$; Figure S2D). Finally, we ran a control experiment in which stimulus attributes were exactly equalized between tasks and expectation conditions. Again, a valid expectation of the orientation angle of the gratings led to a reduction in BOLD response amplitude ($F_{1,7} = 7.2$, $p = 0.016$), but an increase in MVPA orientation classification accuracy, in V1 ($F_{1,7} = 3.6$, $p = 0.050$; Figure S2E). Together, these results preclude an explanation of our results in terms of within-trial stimulus differences between conditions.

In order to investigate the relationship between effects of top-down expectation and stimulus repetition, we separately calculated MVPA orientation classification accuracy for trials containing the same (“repetitions”) or different (“alternations”) grating orientations (45° or 135°) as the previous trial (see Supplemental Experimental Procedures). This analysis revealed a significant main effect of expectation ($F_{1,17} < 9.0$, $p = 0.008$), but not of repetition ($F_{1,17} < 1$, $p > 0.1$). This precludes an explanation of our results in terms of repetition effects. Interestingly, there was also a marginally significant interaction between expectation and repetition ($F_{1,17} = 4.2$, $p = 0.056$), indicating that the effect of expectation on MVPA classification accuracy was larger for alternation than for repetition trials (Figure S2F).

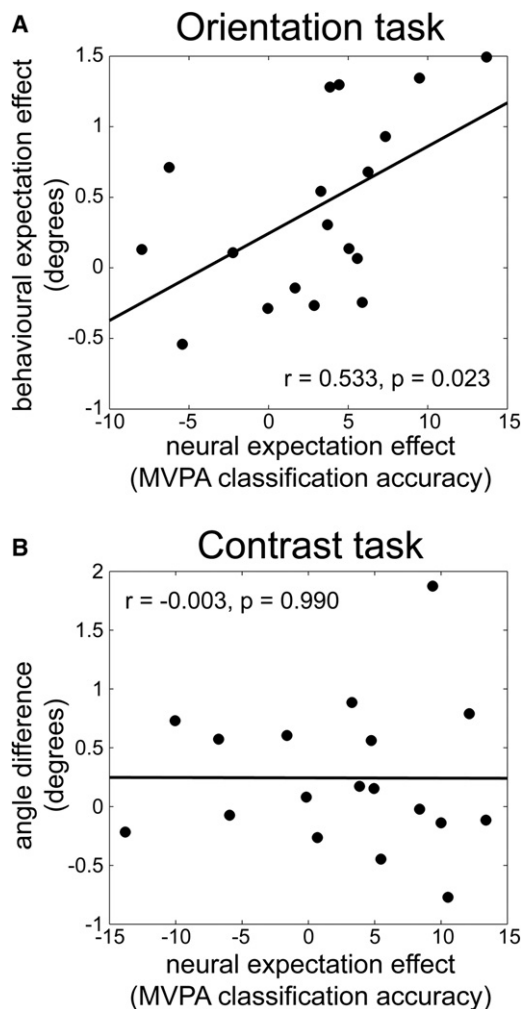


Figure 3. Correlation between Neural and Behavioral Improvement by Expectation

(A) Expectation-induced improvement in MVPA orientation classification accuracy (on the x axis) correlated with behavioral improvement induced by expectation (on the y axis) during the orientation task. Behavioral performance was indexed by the mean difference in orientation angle between the two gratings (see Behavioral Results). (B) Expectation-induced improvements in MVPA orientation classification accuracy did not correlate with angle difference between grating 1 and grating 2 during the contrast task. This precludes an explanation of the correlation found between neural and behavioral improvement on the orientation task in terms of stimulus differences, since similar stimulus differences were present during the contrast task (see Supplemental Experimental Procedures).

This could potentially be explained by the fact that repetition of a stimulus in itself already sharpens stimulus representations (Desimone, 1996; Moldakarimov et al., 2010), reducing the effect of any additional top-down sharpening, while the opposite is true for alternation trials.

DISCUSSION

We observed a striking dissociation between the effects of expectation on the amount of neural activity and the information

that can be gleaned from the neural activity pattern. Whereas expectation leads to suppressed responses in V1, it concurrently enhances the amount of information about the orientation of the stimulus. Crucially, this pattern of results is exactly what is predicted by the “sharpening” hypothesis of expectation, in which bottom-up sensory signals that are incongruent with prior expectations are relatively suppressed (Lee and Mumford, 2003). The sharpening hypothesis was further corroborated by the fact that the expectation-induced reduction of neural activity was less strong in voxels preferring the currently presented orientation than in voxels preferring the nonpresented orientation.

This sharpening has behavioral ramifications for fine-grained orientation discrimination performance, reflected by the fact that subjects were better able to judge small differences in orientation for expected gratings. This behavioral improvement correlated with the improvement of the neural signal as measured by MVPA, in line with theories of representational sharpening due to perceptual priming (Desimone, 1996; Moldakarimov et al., 2010) and adaptation (Wang et al., 2010).

These results are also in line with recent predictive coding models (Friston, 2005; Rao and Ballard, 1999; Spratling, 2008), in which separate populations of neurons within a cortical region code the current estimate of sensory causes (predictions) and the mismatch between this estimate and incoming sensory signals (prediction error). Here, we did not manipulate the prior expectation of the occurrence or omission of stimuli (grating stimuli were present in all trials), but the likelihood of the stimulus having a certain feature (i.e., orientation). This calls for a slightly more sophisticated model of hierarchical Bayesian inference that allows for a representation of uncertainty in terms of the precision of future events, an issue which has been addressed recently within the framework of predictive coding (Feldman and Friston, 2010). Bayes-optimal inference in this setting relies upon top-down predictions about the certainty or precision of events that will occur and suggests that prediction error neurons are selectively biased in a top-down manner following a cue. Simulations within this framework suggest that anticipation enhances early prediction error responses to valid stimuli compared to invalid stimuli. Crucially, this prediction error can be cancelled out more quickly, reducing the overall amount of activity, consistent with the reduction in the amplitude of V1 responses under the predictive coding model. However, it also suggests that the signal-to-noise ratio of prediction error responses is enhanced when valid or anticipated targets are processed. In other words, there should be representational sharpening. In this scheme, top-down expectations about future events increase the gain of prediction error neurons encoding the expected stimulus feature, leading to a quick resolution of prediction error if the input matches the expectation (Feldman and Friston, 2010; Summerfield and Koehlin, 2008). If, on the other hand, the expectation is violated, a large prediction error will ensue, leading to an increase in neural activity (Alink et al., 2010; den Ouden et al., 2009; Kok et al., 2011; Meyer and Olson, 2011; Todorovic et al., 2011). Also, the activity pattern in prediction neurons will contain a mixture of neurons coding the expected (due to top-down biasing) and the actually presented (due to bottom-up input) orientations, yielding a noisy population response.

The effects of top-down expectation were observed alongside the previously observed improvements in neuronal representation as a function of task relevance (Jehee et al., 2011; Kamitani and Tong, 2005), and indeed, the effects of task-relevance and expectation were additive. However, while there were no overall differences in BOLD activity between the different tasks, expectation induced a strong reduction of BOLD activity. This suggests that expectation and attentional task-set may be partly distinct processes, as has been previously argued (Summerfield and Egner, 2009).

Although the relationship between neuronal excitation and inhibition and the hemodynamic (or metabolic) response is equivocal and multifaceted (Logothetis, 2008), the activity reductions observed here for expected stimuli likely reflect a reduction of neural activity. This is in line with recent neurophysiological studies in monkeys and humans, highlighting that valid expectations lead to a reduction in spiking activity (Meyer and Olson, 2011) as well as gamma-band oscillatory activity (Todorovic et al., 2011). Additionally, a recent combined hemodynamic/neurophysiological study reported hemodynamic and metabolic downregulation neuronal inhibition in the visual cortex of monkeys (Shmuel et al., 2006).

In sum, our data provide evidence for how expectations facilitate perceptual inference in a noisy and ambiguous visual world by sharpening early sensory representations.

EXPERIMENTAL PROCEDURES

Subjects

Twenty healthy right-handed individuals (sixteen female, age 22 ± 4 , mean \pm SD) with normal or corrected-to-normal vision gave written informed consent to participate in this study, in accordance with the institutional guidelines of the local ethics committee (CMO region Arnhem-Nijmegen, The Netherlands). Data from one subject were excluded due to excessive head movement, and one subject was excluded due to failure to comply with task instructions.

Stimuli

Grayscale luminance-defined sinusoidal grating stimuli were generated using MATLAB (MathWorks, Natick, MA) in conjunction with the Psychophysics Toolbox (Brainard, 1997), and displayed on a rear-projection screen using a luminance-calibrated EIKI projector (1,024 \times 768 resolution, 60 Hz refresh rate). Gratings were displayed in an annulus (outer diameter: 15° of visual angle, inner diameter: 3°), surrounding a fixation point. The auditory cue consisted of a pure tone (450 or 1,000 Hz), presented over MR-compatible earphones.

Experimental Design

Each trial consisted of an auditory cue, followed by two consecutive grating stimuli (Figure 1). The two grating stimuli were presented for 500 ms each, separated by a blank screen (100 ms). The auditory cue consisted of either a low- (450 Hz) or high-frequency (1000 Hz) tone, which predicted the orientation of the subsequent grating stimuli ($\sim 45^\circ$ or $\sim 135^\circ$) with 75% validity. The contingencies between cues and gratings were flipped halfway through the experiment, and the order was counterbalanced over subjects.

In separate runs (128 trials, ~ 14 min), subjects performed either an orientation or a contrast discrimination task on the two gratings. The first grating had an orientation of either 45° or 135° (\pm a Gaussian jitter, drawn from a normal distribution with mean = 0 and standard deviation = 1) and a luminance contrast of 80%. The second grating differed slightly from the first in terms of both orientation and contrast. When performing the orientation task, subjects had to judge whether the second grating was rotated clockwise or anticlockwise with respect to the first grating. In the contrast task subjects had to judge whether the second grating had lower or higher contrast than

the first one. Subjects indicated their response using an MR-compatible button box.

The orientation and contrast differences between the two gratings were determined by an adaptive staircase procedure, separately for trials containing expected and unexpected orientations. This was done to yield comparable task difficulty and performance ($\sim 75\%$ correct) for the different conditions (Supplemental Experimental Procedures). All subjects completed four runs (two of each task, order was counterbalanced over subjects) of the experiment, yielding a total of 512 trials.

Subsequent to the main experiment, subjects performed a functional localizer task, consisting of flickering gratings, and a retinotopic mapping session (Supplemental Experimental Procedures).

fMRI Acquisition Parameters

Functional images were acquired using a 3T Trio MRI system (Siemens, Erlangen, Germany), with a T2*-weighted gradient-echo EPI sequence (TR/TE = 1,500/30 ms, 26 transversal slices, voxel size $2 \times 2 \times 2$ mm, interslice gap 20%, GRAPPA acceleration factor of 3). Anatomical images were acquired with a T1-weighted MP-RAGE sequence, using a GRAPPA acceleration factor of 2 (TR/TE = 2,300/3.03 ms, voxel size $1 \times 1 \times 1$ mm).

fMRI Data Preprocessing

We used SPM5 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London, UK) for image preprocessing and analysis. The first six volumes of each subject's data set were discarded to allow for T1 equilibration. All functional images were spatially realigned to the mean image, yielding head movement parameters which were used as nuisance regressors in the general linear model (GLM), and temporally aligned to the first slice of each volume. The structural image was coregistered with the functional volumes.

BOLD Amplitude Analyses

For univariate analyses, functional images were spatially smoothed with an isotropic Gaussian kernel with a full-width at half-maximum (FWHM) of 4 mm. Data of each subject were modeled using an event-related approach, within the framework of the GLM. Regressors representing the different conditions (the two tasks and the two expectation conditions) were constructed by convolving the onsets of the first grating in each trial with a canonical hemodynamic response function (HRF) and its temporal and dispersion derivatives (Friston et al., 1998). Instruction screens were included as regressors of no interest, as were head motion parameters and their first-order derivatives (Lund et al., 2005). Finally, the data were high-pass filtered (cutoff 128 s) to remove low-frequency signal drifts.

Freesurfer (<http://surfer.nmr.mgh.harvard.edu/>) was used to identify the boundaries of retinotopic areas in early visual cortex, using well-established methods (DeYoe et al., 1996; Engel et al., 1997; Sereno et al., 1995). Within each retinotopic ROI (V1, V2, and V3), we identified the 150 most stimulus-responsive voxels according to their response to the grating stimulus in the independent functional localizer session. Parameter estimates for each condition were averaged over these voxels. The resulting averaged parameter estimates for the (canonical HRF) regressors comprised the data for the second level analysis (i.e., at the between-subject level).

Multivoxel Pattern Analyses

For multivoxel pattern analyses (MVPA), functional images were not spatially smoothed. Again, the data of each subject were modeled using an event-related approach, but here each trial was modeled by a separate regressor, convolved with a canonical HRF. The exact same voxels were used as for the BOLD amplitude analysis, but now parameter estimates were not averaged over voxels. This procedure yielded a pattern of voxel activations for each single trial. T values (i.e., parameter estimates divided by unexplained variance) obtained for each voxel comprised the data for further analysis (Misaki et al., 2010). These patterns were analyzed using MVPA classification methods (Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2005). Specifically, we classified the orientation of the observed gratings based on the pattern of BOLD activation in early visual areas (V1, V2, and V3). Classification performance can be seen as an indication of the amount of orientation information available in the BOLD signal, such that relative

changes therein can be informative about the effects of expectation and task relevance (Jehee et al., 2011). Linear support vector machines were applied to a subset of the trials, designated as the training set, in order to find a linear discriminant function. Subsequently, the remaining trials (the test set) were classified as containing one of the two orientations, dependent on the outcome of applying the discriminant function to the accompanying voxel activation pattern (see Supplemental Experimental Procedures for details).

SUPPLEMENTAL INFORMATION

Supplemental Information includes three figures and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2012.04.034>.

ACKNOWLEDGMENTS

We thank Hakwan Lau and Tobias Donner for helpful discussions and comments on the manuscript and Paul Gaalman for MRI support. This study was supported by the Netherlands Organisation for Scientific Research (NWO VENI 451-09-001 awarded to F.P.d.L.).

Accepted: April 27, 2012

Published: July 25, 2012

REFERENCES

- Alink, A., Schwiedrzik, C.M., Kohler, A., Singer, W., and Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *J. Neurosci.* *30*, 2960–2966.
- Bar, M. (2004). Visual objects in context. *Nat. Rev. Neurosci.* *5*, 617–629.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spat. Vis.* *10*, 433–436.
- den Ouden, H.E.M., Friston, K.J., Daw, N.D., McIntosh, A.R., and Stephan, K.E. (2009). A dual role for prediction error in associative learning. *Cereb. Cortex* *19*, 1175–1185.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proc. Natl. Acad. Sci. USA* *93*, 13494–13499.
- DeYoe, E.A., Carman, G.J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., and Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc. Natl. Acad. Sci. USA* *93*, 2382–2386.
- Engel, S.A., Glover, G.H., and Wandell, B.A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* *7*, 181–192.
- Feldman, H., and Friston, K.J. (2010). Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* *4*, 215.
- Friston, K.J. (2005). A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *360*, 815–836.
- Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., and Turner, R. (1998). Event-related fMRI: characterizing differential responses. *Neuroimage* *7*, 30–40.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* *293*, 2425–2430.
- Haynes, J.D., and Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* *8*, 686–691.
- Jehee, J.F., and Ballard, D.H. (2009). Predictive feedback can account for biphasic responses in the lateral geniculate nucleus. *PLoS Comput. Biol.* *5*, e1000373.
- Jehee, J.F.M., Rothkopf, C., Beck, J.M., and Ballard, D.H. (2006). Learning receptive fields using predictive feedback. *J. Physiol. Paris* *100*, 125–132.
- Jehee, J.F.M., Brady, D.K., and Tong, F. (2011). Attention improves encoding of task-relevant features in the human visual cortex. *J. Neurosci.* *31*, 8210–8219.
- Kamitani, Y., and Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* *8*, 679–685.
- Kok, P., Rahnev, D., Jehee, J.F.M., Lau, H.C., and de Lange, F.P. (2011). Attention reverses the effect of prediction silencing sensory signals. *Cereb. Cortex*.
- Lee, T.S., and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* *20*, 1434–1448.
- Logothetis, N.K. (2008). What we can do and what we cannot do with fMRI. *Nature* *453*, 869–878.
- Lund, T.E., Nørgaard, M.D., Rostrup, E., Rowe, J.B., and Paulson, O.B. (2005). Motion or activity: their role in intra- and inter-subject variation in fMRI. *Neuroimage* *26*, 960–964.
- Meyer, T., and Olson, C.R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. *Proc. Natl. Acad. Sci. USA* *108*, 19401–19406.
- Misaki, M., Kim, Y., Bandettini, P.A., and Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage* *53*, 103–118.
- Moldakarimov, S., Bazhenov, M., and Sejnowski, T.J. (2010). Representation sharpening can explain perceptual priming. *Neural Comput.* *22*, 1312–1332.
- Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol. Cybern.* *66*, 241–251.
- Murray, S.O., Kersten, D., Olshausen, B.A., Schrater, P., and Woods, D.L. (2002). Shape perception reduces activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA* *99*, 15164–15169.
- Olshausen, B.A., and Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* *381*, 607–609.
- Rao, R.P.N., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* *2*, 79–87.
- Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* *268*, 889–893.
- Shmuel, A., Augath, M., Oeltermann, A., and Logothetis, N.K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nat. Neurosci.* *9*, 569–577.
- Smith, F.W., and Muckli, L. (2010). Nonstimulated early visual areas carry information about surrounding context. *Proc. Natl. Acad. Sci. USA* *107*, 20099–20103.
- Sprattling, M.W. (2008). Predictive coding as a model of biased competition in visual attention. *Vision Res.* *48*, 1391–1408.
- Summerfield, C., and Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends Cogn. Sci.* *13*, 403–409.
- Summerfield, C., and Koechlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron* *59*, 336–347.
- Summerfield, C., Trittschuh, E.H., Monti, J.M., Mesulam, M.M., and Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nat. Neurosci.* *11*, 1004–1006.
- Todorovic, A., van Ede, F., Maris, E., and de Lange, F.P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *J. Neurosci.* *31*, 9118–9123.
- Wang, Q., Webber, R.M., and Stanley, G.B. (2010). Thalamic synchrony and the adaptive gating of information flow to cortex. *Nat. Neurosci.* *13*, 1534–1541.
- Yuille, A., and Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* *10*, 301–308.