fMRI Repetition Suppression: Neuronal Adaptation or Stimulus Expectation?

Jonas Larsson and Andrew T. Smith

Department of Psychology, Royal Holloway, University of London, Egham, TW20 0EX, UK

Address correspondence to Jonas Larsson. Email: jonas.larsson@rhul.ac.uk.

Measurements of repetition suppression with functional magnetic resonance imaging (fMRI adaptation) have been used widely to probe neuronal population response properties in human cerebral cortex. fMRI adaptation techniques assume that fMRI repetition suppression reflects neuronal adaptation, an assumption that has been challenged on the basis of evidence that repetition-related response changes may reflect unrelated factors, such as attention and stimulus expectation. Specifically, Summerfield et al. (Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Egner T. 2008. Neural repetition suppression reflects fulfilled perceptual expectations. Nat Neurosci. 11:1004–1006) reported that the relative frequency of stimulus repetitions and non-repetitions influenced the magnitude of repetition suppression in the fusiform face area, suggesting that stimulus expectation accounted for most of the effect of repetition. We confirm that stimulus expectation can significantly influence fMRI repetition suppression throughout visual cortex and show that it occurs with long as well as short adaptation durations. However, the effect was attention dependent: When attention was diverted away from the stimuli, the effects of stimulus expectation completely disappeared. Nonetheless, robust and significant repetition suppression was still evident. These results suggest that fMRI repetition suppression reflects a combination of neuronal adaptation and attention-dependent expectation effects that can be experimentally dissociated. This implies that with an appropriate experimental design, fMRI adaptation can provide valid measures of neuronal adaptation and hence response specificity.

Keywords: attention, cerebral cortex, human, neuroimaging, vision

Introduction

In the primate visual cortex, neural responses to repeated stimuli are often attenuated relative to the response to a single stimulus. This phenomenon, known as repetition suppression, has been observed both in single-unit recordings in non-human primates (Desimone 1996) and in functional magnetic resonance imaging (fMRI) studies in humans (Henson and Rugg 2003). Repetition suppression is believed to be closely related to neuronal adaptation—a reduction in neuronal firing rates over time to a constant stimulus. Because neurons adapt most strongly to their preferred stimuli and adapt more weakly or not at all to stimuli that do not drive the neurons (Vautin and Berkley 1977; Hammond et al. 1985; Marlin et al. 1988), the strength of adaptation varies with stimulus parameters in a manner that reflects neuronal stimulus selectivity. This property underlies the widespread use of adaptation and repetition suppression as tools to infer neuronal stimulus selectivity with fMRI (often referred to as fMRI adaptation) (Grill-Spector and Malach 2001). A common application of this method involves measuring the fMRI responses to sequentially presented pairs of stimuli that are identical along some feature dimension (e.g., orientation) and comparing the magnitude of

those responses against responses to stimulus pairs that differ in that feature dimension. If responses to the identical pairs (stimulus repetitions) are reduced relative to the nonidentical pairs (stimulus non-repetitions), then this is interpreted as evidence of neuronal selectivity for the stimulus feature, whereas similar responses to both repetitions and non-repetitions is taken as evidence of a lack of selectivity for the feature in question (Larsson et al. 2006; Smith and Wall 2008; Rokers et al. 2009).

The conventional view of fMRI repetition suppression as being due to neuronal adaptation was challenged by Summerfield et al. (2008). They proposed that rather than reflecting a reduced response to stimulus repetitions as a result of adaptation, the difference between responses to repeated and non-repeated stimuli measured by fMRI could be explained as a stronger response to the non-repeated stimuli that reflected the mismatch between observed and expected stimuli. Summerfield et al. (2008) based this proposal on the observation that the strength of repetition suppression in the fusiform face area (FFA) measured by fMRI depended on the relative frequency of stimulus repetitions versus non-repetitions. When stimulus repetitions were more frequent (and therefore more expected), the magnitude of response suppression was greater than when stimulus repetitions were infrequent (or unexpected). They suggested that the stronger response to the unexpected stimuli was not evidence of neuronal adaptation but reflected a prediction error signal. This interpretation was motivated by predictive coding theories (Rao and Ballard 1999) which posit that perception relies on matching feedback prediction signals from higher order areas with the sensory feedforward signals; when the prediction matches the input, the prediction error signal is low, whereas when there is a mismatch between input and prediction signals, the error signal is large. Because the prediction error reflected how well the stimulus input matched the expected percept, Summerfield et al. termed this phenomenon "perceptual expectation." Consistent with this interpretation, other studies have shown that stimulus expectation can significantly modulate neuronal responses measured by fMRI (Summerfield et al. 2006; Summerfield and Koechlin 2008).

However, stimulus expectation does not appear to modulate the magnitude of neuronal repetition suppression measured by single-unit or local field potential recordings in monkey inferotemporal cortex (IT) (Kaliukhovich and Vogels 2010), supporting the conventional interpretation of repetition suppression as reflecting neuronal adaptation. The discrepancy between this result and that of Summerfield et al. (2008) could imply that fMRI repetition suppression is particularly susceptible to stimulus expectation effects. Thus, it is possible that fMRI repetition suppression may not only be driven largely by perceptual expectation but may not even reflect changes in underlying neural activity reliably. If fMRI repetition suppression does not primarily reflect neuronal adaptation, it would call into

question the validity of using fMRI adaptation designs to infer neuronal selectivity. Such a conclusion would, however, only be warranted if the perceptual expectation effects observed by Summerfield et al. (2008) were to generalize to other types of fMRI adaptation designs than the specific one used in that study. In particular, the Summerfield study used what is sometimes referred to as a "rapid" or "short-term" adaptation design, which relies on measuring responses to pairs of brief (<500 ms) stimuli shown in rapid succession, separated by a very brief (0–250 ms) interval. Such designs are sometimes contrasted with "long-term" adaptation designs, which measure responses to single stimuli after medium to long-term (e.g., 4–100 s) exposure to an adapter stimulus. Although there are many parallels between long- and short-term adaptation (or repetition suppression), there is also evidence suggesting that these 2 phenomena may tap into different neuronal mechanisms. Fang et al. (2007) showed that face adaptation (measured with fMRI) exhibits different sensitivity to viewpoint changes depending on the duration of adaptation and orientation-selective adaptation in V1 is only evident with long adaptation durations (Boynton and Finney 2003; Fang et al. 2005; Larsson et al. 2006), and in the macaque middle temporal area (MT), short- and long-term adaptation differ in their spatial specificity (Priebe et al. 2002; Kohn and Movshon 2003). Given these differences, it is possible that perceptual expectation may influence these 2 types of repetition suppression in different ways.

Moreover, there is a close but poorly understood relationship between perceptual expectation and attention, which may have influenced the results of Summerfield et al. (2008). It is well known that both spatial and feature-based visual attention can strongly modulate visually evoked fMRI responses, and there is growing behavioral and physiological evidence that attention may also modulate neuronal adaptation itself (Eger et al. 2004; Yi and Chun 2005; Vuilleumier et al. 2005; Henson and Mouchlianitis 2007). Attentional modulation of stimulus-evoked responses can potentially confound the interpretation of fMRI studies that measure adaptation processes. For example, early fMRI studies of the motion aftereffect (Tootell et al. 1995) were later shown to have incorrectly interpreted the effect of attention as evidence of rebound from adaptation (Huk et al. 2001). To minimize such confounds, it is common practice to use an attentionally demanding task in order to equate attention across stimulus conditions, so that any attentional effects are matched across conditions. A widely employed strategy uses tasks that require subjects to attend to the stimuli and discriminate or detect some stimulus feature—for example, in the Summerfield et al. (2008) study, subjects were required to respond to infrequently occurring upside-down faces. While this task may have been effective at ensuring constant attention to the stimuli, by directing subjects' attention toward the stimuli it may also have made them more aware of differences in the relative frequency of each stimulus condition (repetitions vs. non-repetitions). In other words, by directing attention to the stimuli, the task may have amplified the effects of perceptual expectation, raising the question of whether the observed perceptual expectation effects were specific to the particular task. It is even possible that the effects of perceptual expectation observed by Summerfield et al. (2008) were not due to a prediction error signal at all but merely reflected attentional modulation: If participants attended more to the novel or infrequent trials than the frequent ones, it could explain the relatively greater response to the infrequent trials, regardless of whether these trials were stimulus repetitions or non-repetitions.

A different strategy for controlling attention that circumvents some of the problems associated with subjects directing attention toward the stimuli is to use a task that diverts attention from the stimuli of interest—for example, a demanding discrimination task at the center of gaze while presenting (unattended) stimuli in the periphery (Larsson et al. 2006; Ashida et al. 2007). While this method has the potential drawback that responses tend to be weaker to unattended stimuli because the task is independent of the stimuli under study, it is easier to equate performance (and attentional demands) across stimulus conditions. Moreover, if a sufficiently demanding task is used, subjects frequently report being unaware of subtle differences in stimulus conditions, even to the extent of failing to detect differences between stimulus repetitions and non-repetitions (Larsson et al. 2006). Since subjects would be more unlikely to detect differences in the relative frequency of different trial types under these conditions, no trial type would be more expected (at least consciously). This raises the question whether perceptual expectation effects would still be observed if attention were diverted away from the stimuli (making subjects less likely to notice differences in the frequency of repetitions vs. non-repetitions). If as suggested by Summerfield et al. (2008), perceptual expectation effects were the sole explanation for fMRI repetition suppression, and if these effects depend on attention being directed toward the stimuli, then diverting attention away from the stimuli would predict that repetition effects would disappear altogether. Indeed, consistent with this prediction, some studies have found that face-evoked fMRI responses only adapt (reduce in magnitude) when subjects attend to the stimuli (Eger et al. 2004; Henson and Mouchlianitis 2007). However, it should be noted that other studies have found that although the magnitude of adaptation is reduced in the absence of attention, it does not disappear altogether (Vuilleumier et al. 2005).

In this study, we used the experimental paradigm used by Summerfield et al. (2008) to investigate the influence on perceptual expectation effects of 1) long adaptation durations and 2) attention. To test whether perceptual expectation effects would generalize to long-term adaptation, we replicated the experiments but modified the design to use long (4 s) adaptation durations and measured the responses to stimulus probes separately from the responses to the adapter (as opposed to measuring responses to pairs of stimuli). To test whether perceptual expectation effects depended on attention, the experiments were repeated with 2 different attentionally demanding tasks—one that required subjects to attend to the stimuli and another that required them to attend away from the stimuli. Moreover, we measured fMRI responses across a range of visual areas from V1 to FFA to test whether the effects of perceptual expectation and attention would be specific to areas selective for the stimuli used (faces).

The results showed that perceptual expectation effects are not specific to short-term adaptation designs but can be observed also when using long adaptation durations. Moreover, these effects were found in all visually responsive areas, not just in those selective for faces. However, the perceptual expectation effects were dependent on attention being directed to the stimuli: Once attention was diverted away from the stimuli, the effects disappeared but, crucially, significant repetition suppression remained. These results support the interpretation of fMRI repetition suppression as reflecting neuronal adaptation,

rather than being due to perceptual expectation or prediction error. However, the results also highlight the potential risk of attention and/or perceptual expectation confounding the interpretation of fMRI adaptation data and emphasize the importance of carefully designed attentional control tasks to minimize such effects.

Materials and Methods

Subjects
Eight subjects (all females) aged between 20 and 38 took part in the experiment. All subjects were naive to the purpose of the experiments. Ethical approval was obtained from the research ethics committee of the Department of Psychology at Royal Holloway, University of London. Subjects gave informed written consent to participate, and the experiments were undertaken in compliance with safety guidelines for magnetic resonance imaging (Kanal et al. 2002).

Stimuli
Stimuli were grayscale face photographs of 414 males and 95 females, each taken from 4 slightly different frontal views and viewing distances against a neutral (gray) background. Images were obtained from the following face image databases in the public domain: the Psychological Image Collection at Stirling (PICS) (http://pics.psych.stir.ac.uk); Georgia Tech Face Database (http://www.anefian.com/research/face_reco.htm); Database of Faces, AT&T Laboratories Cambridge (http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html); Yale Face Database B, Yale University (http://cvc.yale.edu/projects/yalefacesB/yalefacesB.html); Faces94 collection of facial images, Department of Computer Science, Essex University (http://cswww.essex.ac.uk/mv/allfaces/faces94.html); Faces 1999 database, Computational Vision, Caltech (http://www.vision.caltech.edu/html-files/archive.html); Indian Face Database, (http://vis-www.cs.umass.edu/~vidit/IndianFaceDatabase); and BioID face database, Friedrich-Alexander University of Erlangen-Nuremberg (http://ftp.uni-erlangen.de/pub/facedb). Automatic face detection software (Graham and Allinson 1998) was used to extract a rectangular region from each image encompassing the individual's face. The extracted region was normalized to 25% root mean square luminance contrast. For the main (adaptation) experiments, face stimuli were scaled to subtend 14 × 14 degree visual angle and presented against a uniform gray background.

For the FFA localizer and retinotopic mapping experiments (see Identification of visual area ROIs), the stimuli consisted of 16 collages of 35 different face images (not identical to those used in the main experiment) created by the above-described procedure and masked by a circular (for the localizer experiment) or wedge-shaped (for retinotopic mapping) aperture. Apertures extended to an eccentricity of 15°. For the localizer experiment, scrambled versions of the same images were created by previously published procedures (Larsson and Heeger 2006).

MRI Acquisition
Visually evoked cortical blood oxygenation level-development fMRI responses were measured by T_2^*-weighted gradient-recalled echoplanar imaging on a 3 T whole-body MR scanner (Magnetom Trio; Siemens, Erlangen, Germany) equipped with a custom 8-channel posterior-head array coil (Stark Contrast, Erlangen, Germany). Functional MRI data were acquired from 19 oblique slices roughly parallel to the calcarine sulcus and covering the occipital and temporal cortex (voxel size 3 × 3 × 3 mm, time repetition [TR] = 1500 ms, time echo [TE] = 34 ms, flip angle = 85°). On each session, a whole-brain anatomical MR volume was acquired and used for spatial coregistration of data across sessions (voxel size 1 × 1 × 1 mm, MPRAGE sequence, TR = 1830 ms, time to inversion [TI] = 1100 ms, TE = 5.6 ms, flip angle = 11°). In a separate session, a high-resolution high-contrast T_1-weighted anatomical MR volume of each subject was acquired (voxel size 1 × 1 × 1 mm, MDEFT sequence [Deichmann 2006], TR = 7.9 ms, TI = 910 ms, TE = 2.5 ms, flip angle = 16°) and used for cortical surface reconstruction (Larsson 2001).

Experimental Design
Each subject took part in 2 main (face adaptation) experiments, run in separate sessions. The 2 experiments differed in the task used to control spatial attention (see Spatial Attention Conditions) but were otherwise identical. The structure of the experiments was modeled on the experiments of Summerfield et al. (2008) but modified to use a long-term adaptation design. An event-related experimental design was used (Fig. 1A). On each trial, we measured the visually evoked fMRI response to brief (1 s) presentations of a face stimulus probe following a 4 s adapter face stimulus and a 0.5 s interstimulus interval. To minimize low-level contrast and luminance adaptation, adapter and probe stimuli consisted of a series of 4 different images of the same individual but taken from slightly different viewing angles and distances (see Stimuli), alternating in random order every 200 ms. There were 3 trial types: SAME, DIFF, and BLANK. On SAME trials, the same 4 images (of the same individual) were used for both adapter and probe stimuli, on DIFF trials, the adapter and probe stimuli were of different individuals, and on BLANK trials, only the adapter stimulus was shown (permitting estimation of the average response to the adapter stimulus alone). Intertrial intervals were randomly varied between 0.5 and 3.5 s in steps of 1.5 s. To minimize long-term priming effects, each trial used images of a unique face such that images of individual faces were never repeated across trials within a scanning session. The same sequence of face stimuli was used in both attention conditions (see Spatial attention conditions).
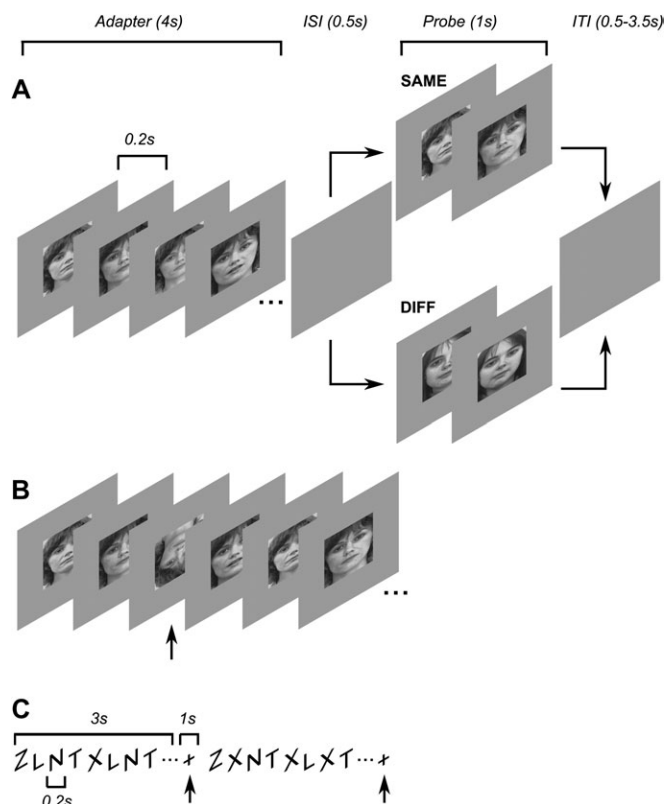

Figure 1. Experimental design and stimuli. (A) Event-related adaptation design. Each trial began with a 4 s adapter stimulus (consisting of multiple different images of the same person randomly alternating every 0.2 s) followed by a 1 s probe stimulus. On SAME trials, the probe stimulus consisted of images of the same person as the adapter stimulus; on DIFF trials, the probe stimulus was of a different person. Each trial ended with a blank screen that lasted between 0.5 and 3.5 s. (B) Focused attention task (FOCUS condition). Subjects fixated a central cross while monitoring the face stimuli and were instructed to respond to infrequently occurring inverted faces by pressing a response button. The inverted face was shown for 100 ms and could appear at any time during either of the adapter or probe epochs. (C) Diverted attention task (DIVERT condition). Independent of and asynchronous with the face stimuli, subjects monitored a stream of letters at the center of the display and counted the number or target letters (X) among the distractors (Z, L, N, T). After 3 s, the letters were replaced by a cross, cueing subjects to indicate how many targets (between 0 and 3) they had seen.

We employed the method used by Summerfield et al. (2008) to manipulate perceptual expectation for stimulus repetitions relative to stimulus non-repetitions, by varying the relative frequency of stimulus repeat (SAME) and non-repeat (DIFF) trials between (but not within) scans (runs). On half of the scans (FREQ scans), there were 3 times as many SAME trials (27) as DIFF trials (9); on the remaining scans (INFREQ scans), the relative numbers of SAME and DIFF trials were reversed. Four scans of each type were run within each scanning session. Scans alternated between the 2 types, counterbalancing the order of FREQ and INFREQ scans across subjects and sessions; for half the sessions, the FREQ scan was run first, and for the remaining sessions, the INFREQ scan was run first. The number of BLANK trials (9) was kept constant across scans. The order of trials was randomized such that on each trial, the likelihood of the preceding trial being a particular type was proportional to the relative frequency of that trial type. Each scan consisted of 45 trials and lasted 337.5 s.

Spatial Attention Conditions

Two different behavioral tasks, analyzed as separate conditions, were used to control and equate spatial attention load across trials and stimulus repetition frequencies. The tasks were run in separate scanning sessions on different days. The order of attention conditions was counterbalanced across subjects to ensure any differences between conditions could not be due to order effects (e.g., due to long-term priming or learning). The first task (FOCUS condition), which required subjects to focus attention on the face stimuli, was very similar to the task used by Summerfield et al. (2008). Subjects fixated a cross presented at the center of gaze throughout the scan, while simultaneously covertly attending to and monitoring the face stimuli for brief (200 ms) and rare (on average 9 per scan) inverted faces, to which they responded by pressing a response key as quickly as possible (Fig. 1B). Visual feedback was provided by briefly changing the color of the fixation cross to green if subjects responded within 1 s (no feedback was given for missing responses). The inverted face could appear at any time that a face was shown during a trial (i.e., either during adapter or probe stimulus presentations).

In the second task (DIVERT condition), attention was diverted away from the face stimuli and toward the center of gaze by a demanding rapid serial visual presentation (RSVP) task at fixation that was identical across scans and stimulus conditions (Larsson et al. 2006). Subjects were required to count the number (0–3) of target letters "X" in a rapid (200 ms/letter) stream of distractor letters shown at the center of gaze and respond by pressing 1 of 4 response keys at the end of each RSVP trial, indicated by the display of a cross at fixation (Fig. 1C). Feedback was provided by changing the color of the fixation cross to green following correct responses and to red following incorrect responses. Each RSVP trial lasted 4 s (3 s letter stream followed by 1 s response time window), and trials were run back-to-back. The timing of the RSVP trials was independent of and asynchronous with the face stimuli. We have previously shown that this task is effective at controlling and equating spatial attention across different trials and stimulus conditions in experiments using similar event-related designs (Larsson et al. 2006; Montaser-Kouhsari et al. 2007). Importantly, unlike the FOCUS task, the DIVERT task required constant attention at the center of gaze regardless of whether a face stimulus was present or not and was perceived as highly attentionally demanding by subjects.

MRI Preprocessing

Functional image volumes acquired at different time points were spatially aligned using motion-correction software (FSL). Data for each scanning session were aligned across sessions by coregistering them with high-resolution anatomical MR images of each subject's brain using custom software (Nestares and Heeger 2000). Cortical surface models of each individual subject's brain (used for visualization and visual area identification) were reconstructed from the high-resolution anatomical MR images using the public domain software SurfRelax (Larsson 2001).

fMRI Data Analysis

Data from the adaptation scans were analyzed separately for individual subjects and visual area regions of interest (ROIs) (see Identification of

visual area ROIs) using custom software written in Matlab. First, the average response time courses to the adapter and probe stimuli for each of the 3 trial types (SAME, DIFF, and BLANK) were estimated by linear deconvolution. Second, the response amplitudes to individual probe stimuli were estimated by a general linear model, using estimates of the average responses to adapter and probe stimuli, to model the fMRI response to each probe stimulus separately. Third, a perceptual expectation index—a measure of the effect of perceptual expectation on the amount of fMRI repetition suppression—was computed from the probe response amplitudes for each combination of stimulus repetition frequency (FREQ vs. INFREQ) and attentional control task (FOCUS vs. DIVERT).

Response time courses and amplitudes

Average ROI response time courses to each of the 3 trial types (SAME, DIFF, and BLANK) were computed for each subject and stimulus condition separately (FREQ vs. INFREQ × FOCUS vs. DIVERT) using linear deconvolution (Burock and Dale 2000). For each ROI, the mean fMRI response time course vector Y (averaged across voxels and concatenated across scans) was converted to percent signal change, detrended by high-pass filtering (cutoff 0.03 Hz), and fit with a linear model:

$$Y = X_1 b_1 + \varepsilon \tag{1}$$

The first column of the design matrix X_1 (all 1's) modeled the mean response. The next 16 columns (corresponding to 24 s) modeled the average response to the SAME trials as follows. The first of these columns had a 1 at the onset of every SAME trial and 0 elsewhere; each of the next 15 columns was a copy of the previous column, shifted one time point down (with the restriction that responses were not modeled across scans). The responses to the DIFF and BLANK trials were modeled in the same way in the next 32 columns. Multiple regression was used to compute a least-squares estimate of the beta weights vector $\hat{b}_1$:

$$\hat{b}_1 = (X_1^T X_1)^{-1} X_1^T Y \tag{2}$$

.

This yielded an estimate of the mean and standard error of the fMRI response amplitude at each of the 16 time points (0–22.5 s) following trial onset for each trial type. Time courses were averaged across subjects for visualization (Figs 2 and 3).

For each subject and ROI, the estimated response time course for the BLANK trials corresponded to the average response to the adapter stimulus alone. By subtracting this time course from the mean time course of the SAME and DIFF trials, we obtained an estimate of the average response time course to the probe stimuli alone (Figs 2 and 3). The average time courses for the adapter and probe stimuli were used to estimate response amplitudes to individual probe stimulus presentations as follows. The estimated adapter and probe response time courses were fit with a model of the hemodynamic response function (HRF) (a difference of 2 gamma functions) separately. A design matrix X_2 was then constructed with the first column (all 1's) modeling the mean response, the second column modeling the average adapter response across all trial types (created by convolving the estimated HRF for the adapter response with a vector with 1 at the onset time of each trial and 0 elsewhere), and the remaining columns modeling the responses to individual probe stimuli (one column per probe, created by convolving the average HRF for the probe stimuli with a vector having a value of 1 at the onset of the probe stimulus for a single trial and 0 elsewhere). In this model, the response to the adapter stimuli was treated as a covariate of no interest that was constant across trials, whereas responses to individual probe stimuli were modeled separately, permitting estimation of standard errors and confidence intervals of the probe responses and perceptual expectation indexes by bootstrapping (see Perceptual expectation index). The response time course Y was then fit with this model by multiple regression as described above to yield a second vector of beta weights b_2, in which the weights corresponding to individual probe stimuli represented the response amplitudes for each probe stimulus presentation. This model explained between 29% and 81% of the variability in the measured fMRI responses (mean $R^2 = 0.56$), suggesting the model described the data adequately. Beta weights representing response amplitudes were averaged across subjects to yield a mean and standard error of the response amplitude for each ROI and stimulus

condition. For each stimulus condition, the amount of repetition suppression (reduction in fMRI response for stimulus repetitions) was given by the difference between mean response amplitudes for DIFF and SAME trials. For the focused attention condition, we included all trials in the analysis, including those trials in which an inverted face was shown, as these were rare and distributed randomly across trial types. To confirm that this did not bias our analysis, we also analyzed the data with those trials excluded, with virtually identical results.

Perceptual Expectation Index

We quantified the effect of stimulus repetition frequency, or perceptual expectation, by computing for each ROI, subject, and attention condition a "perceptual expectation index" (PI). This index measured the difference in the magnitude of repetition suppression (the mean difference in response amplitudes to DIFF and SAME trials) between FREQ and INFREQ scans expressed as a proportion of the average repetition suppression magnitude across all scans:

$$PI = \frac{D_{\text{freq}} - D_{\text{infreq}}}{\left| D_{\text{freq}} \right| + \left| D_{\text{infreq}} \right|} \quad (3)$$

where D_{freq} and D_{infreq} are the differences in fMRI response amplitudes between DIFF and SAME trials for the FREQ and INFREQ scans, respectively. The perceptual expectation index could vary between −1 and 1. A PI close to zero would indicate that the amount of repetition suppression was independent of stimulus repetition frequency (i.e., no difference between FREQ and INFREQ scans), as would be expected if repetition suppression reflected neural adaptation but not perceptual expectation. Conversely, a PI close to 1 would mean that all the repetition suppression in FREQ scans could be attributed to perceptual expectation (i.e., there would be no repetition suppression in INFREQ scans). Values in between 0 and 1 would indicate that the observed repetition suppression reflected a combination of neural adaptation and perceptual expectation. A PI significantly less than zero would indicate an effect of perceptual expectation opposite to that reported by Summerfield et al. (2008), such that repetition suppression would be greater when stimulus repetitions were rare or unexpected. Bootstrapping (Efron and Tibshirani 1993) was used to compute confidence intervals of the sample mean PIs for each visual area ROI for every combination of stimulus repetition frequency and attention condition (FREQ vs. INFREQ × FOCUS vs. DIVERT), and for the difference in mean PIs between the 2 attention conditions (FOCUS vs. DIVERT). On each bootstrapping iteration, samples were taken with replacement from the individual fMRI probe response amplitudes corresponding to each combination of repetition frequency and attention condition for each subject (the number of samples being equal to the actual number of measurements for each condition). From these samples, the mean PI (or mean difference between FOCUS and DIVERT PIs) was computed for each subject and averaged across subjects and the process repeated 1000 times to yield an estimate of the sampling distribution of the PI for each stimulus condition and visual area ROI. Standard errors and 95% confidence intervals for the sample mean PI were computed by identifying the limits of the corresponding percentiles from these distributions. The statistical significance of PIs being greater than 0 was estimated as the proportion of the resampled distribution that was less than or equal to this value (one-tailed test). Significant differences between PIs in the FOCUS and DIVERT conditions were assessed with a permutation test. The mean PIs for each subject and condition were randomly labeled as coming from either of the 2 conditions and the mean difference between the resampled groups computed. This process was repeated for each unique combination of labels to generate a resampled distribution of the mean difference between the 2 conditions under the null hypothesis that PIs for the 2 attentional conditions were the same. For each ROI, differences between PIs were considered significant if less than 5% of all resampled differences were greater than or equal to the observed difference (one-tailed test).

Identification of Visual Area ROIs

Standard phase-encoded retinotopic mapping methods were used to identify borders between retinotopic visual areas corresponding to reversals in visual field maps (Sereno et al. 1995). In a separate scan at the end of each scanning session, subjects viewed collages of grayscale face stimuli (see Stimuli) masked by a wedge-shaped aperture extending to 15° eccentricity from a central fixation cross. The aperture slowly rotated across the visual field with a period of 24 s. Data were collected for 10.5 cycles (255 s) of the rotating wedge stimulus; data for the first half-cycle were discarded before analysis. The coherence and phase of the response to the stimulus at the stimulus alternation frequency was computed for each voxel and visualized on flattened surface representations (flat maps) of the posterior cortical surface. Boundaries between 10 retinotopic visual areas (V1, V2, V3, hV4, LO1, LO2, V3A/B, V7, MT, and VO1) were identified along reversals in the response phase on these flat maps. We used the nomenclature and scheme for defining visual areas of Wandell et al. (2007) and Larsson and Heeger (2006); however, the conclusions of the study do not depend on the particular parcellation scheme used. The consistency of visual area definitions across sessions was confirmed by comparing boundaries drawn from the first session's data against data from the second scanning session; these agreed for all subjects.

The FFA ROI was identified using a standard FFA localizer stimulus (faces vs. scrambled faces) (Kanwisher et al. 1997), run before the adaptation scans in both sessions. Stimuli were the same as those used for retinotopic mapping but shown within a 15° wide annular aperture around fixation (see Stimuli). The face stimuli alternated with scrambled versions of the same stimuli every 24 s. Data were collected for 255 s (5.25 cycles); data from the first quarter cycle were discarded prior to analysis. The FFA was defined operationally as the region in fusiform gyrus and adjacent cortical tissue that responded at the stimulus alternation frequency with a coherence of 0.2 or greater with a phase corresponding to a preference for faces over scrambled faces. This was based on both scanning sessions to avoid any potential session bias in the definition of the ROIs. The FFA ROIs were masked to exclude voxels that showed significant response modulation to the retinotopic stimulus as several retinotopic visual areas (e.g., VO1) also showed a preference for faces over scrambled faces.

Results

Focused Attention Condition

In the frequent repetition scans (FREQ), visually evoked responses in FFA exhibited robust repetition suppression: responses to probe stimuli that were identical to the adapter stimuli (stimulus repetitions; SAME trials) were significantly smaller than responses to probe stimuli that differed from the adapter (stimulus non-repetitions; DIFF trials) (Fig. 2A) (paired t-test, t_7 = 4.65, P < 0.01). Although FFA exhibited the greatest response amplitudes and repetition suppression, as expected if the repetition suppression reflected face-selective neuronal adaptation in this area, all visual areas examined showed significant repetition suppression in the frequent repetition scans (Fig. 2B) (paired t-test, P < 0.01 for all areas). This may reflect adaptation to low-level image components and/or top-down feedback to these areas from face-selective regions such as the FFA.

When stimulus repetitions were infrequent (INFREQ scans), the amplitude of repetition suppression was significantly reduced in all visual areas (Fig. 2) except VO1 (paired t-test, P < 0.05 in all areas), thus showing that the results of Summerfield et al. (2008) for short-term adaptation pertain also for long adaptation durations. The reduction in adaptation amplitudes was attributable both to a weaker response to the stimulus non-repetitions (DIFF trials) and a stronger response to the SAME trials (cf. left and right panels in Fig. 2) in the INFREQ scans than in the FREQ scans. In 6 of the 11 areas examined (V1, V2, V3, V3A, LO1, and MT), there remained no significant difference between responses to SAME and DIFF trials in the INFREQ scans (paired t-test, all P > 0.1), but in the remaining areas, such as the FFA, repetition suppression was weaker but not absent. However, the mean

response to the SAME trials was never significantly greater than mean response to DIFF trials in either of the FREQ or INFREQ scans in any visual area (Fig. 2B) (paired *t*-test, *P* > 0.2), as would have been expected if the difference in response amplitudes between DIFF and SAME trials only reflected the relative frequency of each trial type. This pattern of responses suggests that the observed differences between the responses to SAME and DIFF trials were not solely due to neural adaptation (which depended only on trial type, i.e., SAME vs. DIFF and would have predicted no differences in response due to stimulus repetition frequency) or solely due to perceptual expectation (which depended only on the frequency of each trial type, i.e., FREQ vs. INFREQ and would have predicted no differences in response due to trial type) but rather reflected a combination of both, in all the visual areas examined (see also Perceptual expectation index).

Diverted Attention Condition

When attention was diverted away from the face stimuli, visually evoked fMRI responses were attenuated compared with the focused attention condition, but significant repetition suppression was still observed in most areas (Fig. 3). However,
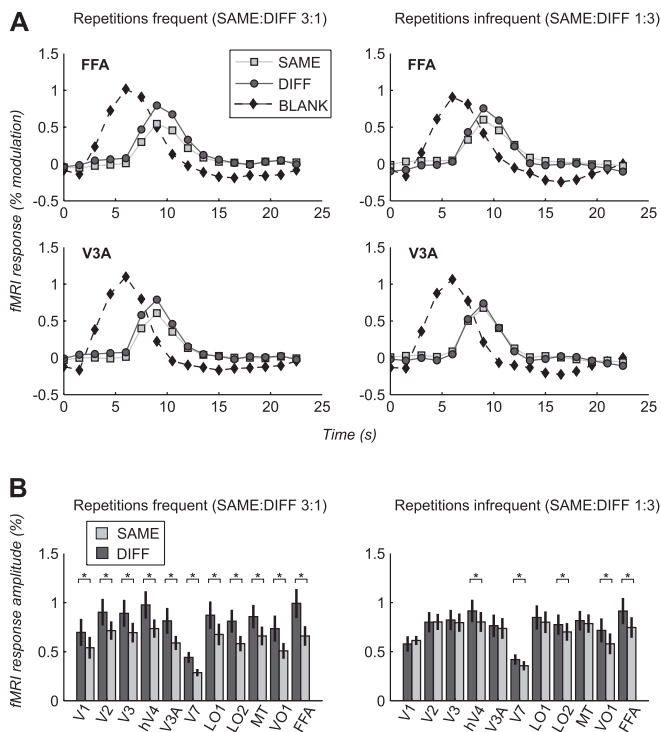
in this condition, we observed no significant differences in the magnitude of repetition suppression between scans with frequent stimulus repetitions (FREQ scans) and scans in which such repetitions where infrequent (INFREQ scans) (paired *t*-test, *P* > 0.05 in all areas), suggesting that perceptual expectation did not influence the degree of repetition suppression in this condition. Indeed, there was not even a nonsignificant trend, the adaptation effect being greater, if anything, in INFREQ trials (Fig. 3B). Average performance on the diverted attention task was well above chance level (0.25) for all subjects and conditions (average performance 0.60, range 0.41–0.87). Neither performance nor response times differed significantly between the frequent and infrequent repetition conditions (repeated measures analyses of variance across scans, both *P* > 0.2), suggesting that the task was effective at equating spatial attention across conditions.

Perceptual Expectation Index

In the focused attention condition, perceptual expectation indexes (PIs) were significantly greater than zero (*P* < 0.01, one-tailed, bootstrapped estimate of confidence limits) in all areas except VO1, meaning that the repetition suppression observed in these areas reflected in part perceptual expectation due to stimulus repetition frequency (Fig. 4). Mean PIs were around 0.5 in most areas, consistent with repetition suppression reflecting a combination of perceptual expectation and neural adaptation. In contrast, perceptual expectation indexes were not significantly greater than zero in any visual area (*P* > 0.2, one-tailed, bootstrapped estimate of confidence
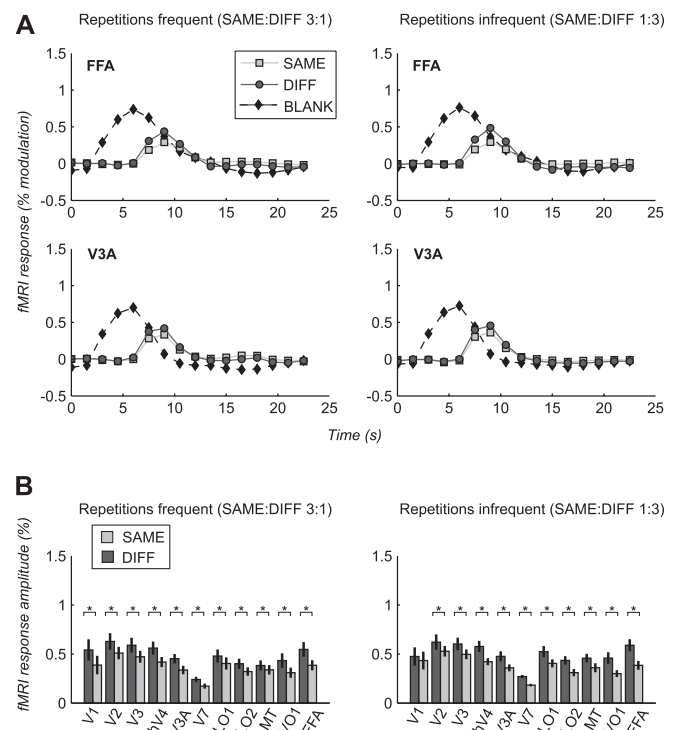
Figure 2. Perceptual expectation modulates repetition suppression when stimuli are attended. (*A*) Average stimulus-evoked response time courses to adapter (BLANK) and probe (SAME, DIFF) stimuli in FFA (upper panels) and V3A (lower panels) for frequent repetition runs (left hand panels) and infrequent repetition runs (right hand panels) in the focused attention condition (FOCUS). For SAME and DIFF trials, only the response components evoked by the probe stimuli are shown (obtained by subtracting the response to the BLANK trials from the overall trial-triggered response; see Materials and Methods). Error bars (standard error of the mean across subjects) are smaller than plot symbols. Repetition suppression (larger responses to DIFF than to SAME trials) is greater when stimulus repetitions are frequent than when they are infrequent. (*B*) Average response amplitudes to probe stimuli across visual areas for the SAME and DIFF trials for frequent repetition runs (left hand panels) and infrequent repetition runs (right hand panels) in the FOCUS condition. The magnitude of repetition suppression (larger responses to DIFF than to SAME trials) is greater in the frequent repetition runs than in the infrequent repetition runs. Error bars, standard error of the mean across subjects. Asterisks (*) indicate significant repetition suppression (greater response to DIFF than to SAME trials, paired *t*-test, *P* < 0.05).

Figure 3. Perceptual expectation effects disappear when attention is diverted away from stimuli. Response time courses (*A*) and response amplitudes (*B*) as in Figure 2, but for the diverted attention condition (DIVERT). Repetition suppression magnitudes do not differ between frequent and infrequent runs (see text for details). Asterisks (*) indicate significant repetition suppression (greater response to DIFF than SAME trials, paired *t*-test, *P* < 0.05).
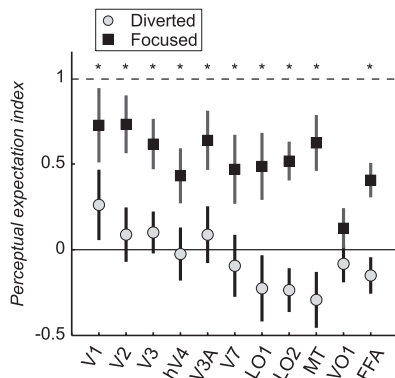
Figure 4. Perceptual expectation modulates repetition suppression magnitude in most visual areas. Perceptual expectation indexes across areas for the focused attention condition (FOCUS) and the diverted attention condition (DIVERT). Perceptual expectation indexes were significantly greater than 0 ($P < 0.05$, one-tailed permutation test) in all areas except VO1 in the focused attention condition but were not significantly different from 0 in any area in the diverted attention condition. Dashed line indicates an index of 1 (repetition suppression entirely due to perceptual expectation); solid line an index of 0 (repetition suppression entirely due to response adaptation). Error bars, bootstrapped 68% confidence intervals across subjects and trials (see Materials and Methods). Asterisks (*) indicate significant differences between perceptual expectation indexes for the focused and diverted attention conditions (one-tailed permutation test, $P < 0.05$).

limits) in the diverted attention condition, implying that diverting attention from the stimuli removed the effects of perceptual expectation. A direct comparison of perceptual expectation indexes showed a significant difference ($P < 0.05$, one-tailed permutation test) between the 2 attention conditions in all areas except VO1, where the difference between PIs in the 2 conditions was nonsignificant ($P > 0.2$) (Fig. 4).

In summary, the results suggest that perceptual expectation can strongly modulate fMRI response amplitudes but that such effects are dependent on spatial attention being directed toward the stimuli. When attention was diverted away from the stimuli, perceptual expectation effects disappeared, but significant repetition suppression was still observed. These results suggest that when stimuli are unattended, fMRI repetition suppression primarily reflects actual neural response adaptation, but when attention is focused on the stimuli, the observed repetition suppression may indeed be dominated by the effects of perceptual expectation.

Discussion

There is no doubt that direct physiological measurements of neuronal adaptation can provide useful information about neuronal response properties (Vautin and Berkley 1977; Hammond et al. 1985; Kohn and Movshon 2003; Priebe et al. 2010). Potentially, measuring adaptation with fMRI offers a powerful means to derive similar information in the context of population responses in the human brain, across multiple brain regions and in a range of experimental contexts. A fundamental issue, therefore, is to what extent repetition suppression measured with fMRI is a veridical measure of neuronal adaptation. Although there is some evidence that the 2 measures generally agree for low-level stimulus features (e.g., contrast and orientation; Gardner et al. 2005; Fang et al. 2005; Larsson et al. 2006) in early visual areas, much less is known about the degree of correspondence between neuronal adaptation and fMRI repetition suppression in higher visual areas

and/or for higher order features, such as objects and faces. Moreover, the assumptions underlying fMRI repetition suppression techniques has been criticized as several studies have found significant discrepancies between the selectivity of adaptation and neuronal tuning properties measured in single neurons (Tolias et al. 2005; Sawamura et al. 2006; Krekelberg et al. 2006; Bartels et al. 2008; Verhoef et al. 2008). Some of these discrepancies are likely not specific to adaptation measures but may reflect differences between single-unit and population-based (e.g., fMRI) measures of neuronal activity (Logothetis and Wandell 2004). Nonetheless, it is evident that adaptation-based measures of selectivity need not directly map onto neuronal response selectivities measured in single neurons and need to be interpreted with caution. Indeed, the finding that perceptual expectation can strongly modulate fMRI repetition suppression to face stimuli both for long-term adaptation (present study) and short-term adaptation (Summerfield et al. 2008) is evidence that repetition suppression measured by fMRI need not reflect neuronal adaptation but may largely be a confound of stimulus expectation. As the effects of perceptual expectation are present with long as well as short durations of adaptation and have also previously been observed with standard block designs that do not rely on adaptation (Summerfield and Koechlin 2008), it appears that these effects are unrelated to neuronal adaptation. Moreover, in this study, perceptual expectation effects were observed in all visually responsive areas regardless of their selectivity for face stimuli, implying that these effects were not stimulus selective. Hence, not only are the effects of perceptual expectation unrelated to adaptation, they also reveal little about the underlying neuronal selectivity. Because the perceptual expectation effects were as large as or larger than those attributable to adaptation, without a means to independently quantify and account for these effects, any interpretation of the observed repetition suppression as being evidence of face-selective neuronal adaptation (implying face-selective neuronal responses) would therefore be in error.

Although these results might, like those of Summerfield et al. (2008), seem to invalidate the use of fMRI adaptation methods, our results also show that the effects of perceptual expectation were only apparent when stimuli were attended and disappeared when attention was diverted away from the stimuli by an attentionally demanding task. Under these conditions, significant response suppression (reduced responses to repeated stimuli) was still observed (Fig. 3), as would be expected if fMRI adaptation reflected neuronal adaptation. This result is consistent with the finding that perceptual expectation does not modulate repetition suppression measured in single units in macaque IT when monkeys engage in a fixation task (Kaliukhovich and Vogels 2010) assuming such a task focuses attention on the fixation spot. (Given that monkeys were only rewarded for trials in which they maintained stable fixation, this does not seem an unreasonable assumption.) Overall, the results suggest that repetition suppression measured by fMRI reflects a combination of perceptual expectation effects (reflecting the relative frequency of stimulus repetitions vs. non-repetitions), the magnitude of which depends on whether subjects attend to the stimuli or not, and neuronal adaptation, which is not influenced by perceptual expectation. Moreover, with the appropriate experimental design (controlling and diverting attention away from the stimuli), the component due to neuronal adaptation can be measured in isolation. Importantly, our results also suggest that perceptual expectation

alone cannot explain fMRI adaptation (a possibility suggested by Summerfield et al. (2008)), as predicted by the extensive body of physiological evidence for neuronal adaptation at all levels of the visual system.

Implications for Studies of fMRI Adaptation

Our results have important implications for the application of fMRI adaptation techniques and interpretation of results obtained by these techniques. In recent years, fMRI adaptation has been widely used as a tool to measure neuronal selectivity for a range of visual stimulus parameters, such as orientation (Fang et al. 2005; Larsson et al. 2006), motion (Huk et al. 2001; Lingnau, Ashida, et al. 2009), stereoscopic depth (Smith and Wall 2008), or object shape (Grill-Spector et al. 1999; Kourtzi and Kanwisher 2000), and also to identify populations of neurons having more abstract response properties (e.g., number representations [Cohen Kadosh et al. 2007; Piazza et al. 2007; Notebaert et al. 2010] or mirror neurons [Chong et al. 2008; Lingnau, Gesierich, et al. 2009]). Many of these previous fMRI adaptation studies have either not attempted to control spatial attention or used tasks that directed attention toward the stimuli. One conclusion from the present study is that studies in which the tasks require subjects to attend to the stimulus features being studied may be subject to confounding effects of perceptual expectation, even though the tasks themselves are attentionally demanding (such as the one used in the FOCUS condition in this study). To the extent that the relative frequencies of stimulus repetitions and non-repetitions were unequal, it is thus likely that the results of some of these previous studies may have reflected in part perceptual expectation effects, rather than (as assumed) pure neuronal adaptation. If, as in the present study, the effects of perceptual expectation were as large as or larger than those attributable to neuronal adaptation, some of these studies may therefore erroneously have interpreted such effects as evidence of selectivity for a particular property in a region of cortex (or, conversely, may have failed to identify true adaptation). It may be that our expectation effects are greater than those in most studies because of our unusual use of a 3:1 ratio of repetitions and non-repetitions, but the possibility cannot be ruled out that even small differences in relative stimulus frequency could give rise to such effects. Moreover, it is possible that the magnitude of expectation effects depends on the stimulus set used and may be greater for complex stimuli (such as the faces used in this study) than for low-level patterns (e.g., gratings or random dots). While the problem of expectancy effects might limit the interpretability of some previous fMRI adaptation studies, our finding that robust adaptation can be observed even in the absence of directed attention suggests a simple strategy for avoiding confounding effects of perceptual expectation, by using a task that diverts attention away from the stimuli. Hence, studies that have used demanding tasks diverting attention away from the stimuli are less likely to have been susceptible to this confound. An open question is whether the effect of attention on expectation is mediated primarily by spatial or feature-based attention (or a combination). Given that the task used in this study would likely have manipulated both types of attention, we cannot from the present results determine which one is more critical.

A possible reason why many adaptation studies have used tasks that require subjects to attend to stimuli may be that early studies suggested that higher order visual areas, such as the FFA, did not show repetition suppression in the absence of attention (Eger et al. 2004; Henson and Mouchlianitis 2007). In contrast, our data show that it is in fact possible to evoke and measure robust stimulus-specific repetition suppression in these areas also when attention is diverted away from the stimuli, confirming similar results by others (Vuilleumier et al. 2005). This result is also consistent with single-unit measurements of neuronal adaptation in macaque IT showing that the magnitude of repetition suppression is unaffected by whether the monkey is performing an attentionally demanding task or passively fixating (De Baene and Vogels 2010). However, as one effect of attention is to boost fMRI stimulus-evoked responses (Gandhi et al. 1999; Somers et al. 1999), it may simply be more difficult to detect repetition suppression when attention is diverted from the stimuli, which may explain some of the differences between different studies of the effect of attention on adaptation. Considering that the present study was carried out on standard research-grade MRI equipment, the loss in sensitivity from using a diverted attention task is not insurmountable but can in principle be compensated for by increasing the number of trials and/or scan time.

Predictive Coding or Attention?

Although the conclusions of this study have the greatest implications for the interpretation, validity, and optimal design of fMRI adaptation experiments, the results are also relevant for predictive coding theories that originally motivated the study of Summerfield et al. (2008). Our results show that attention was necessary for perceptual expectation effects to be measurable. Previous studies of perceptual expectation have acknowledged the close relationship between attention and perceptual expectation, and the difficulty of dissociating them experimentally, but have not explicitly examined whether perceptual expectation requires attention (Summerfield and Koechlin 2008; Summerfield et al. 2008). Our observation of a tight relationship between the 2 phenomena suggests an alternative explanation that does not depend on predictive coding. Although the perceptual expectation effects observed when subjects attended to the stimuli are consistent with a prediction error signal riding on top of neuronal adaptation, an alternative explanation of this effect is simply that it reflected spatial attention. It is likely that subjects attended more strongly to the stimuli on novel or infrequent trials, resulting in a stronger fMRI response for those trials regardless of whether they were repetitions or non-repetitions. While this is in some sense also an effect of stimulus expectation, it differs from the expectation effect described by Summerfield et al. (2008), in that it ascribes the increase in activity on infrequent trials to an attentionally driven increase in the gain of stimulus-evoked neuronal responses, rather than reflecting an error signal from a process matching input to top-down predictions about the stimulus. An alternative possibility is that the expectation effect reflects a global (nonstimulus specific) attentional signal (Donner et al. 2008). Given the large body of evidence for attentional modulation of the activity of single neurons and the lack of corresponding data on prediction-error signals at the individual neuron level, the former explanation seems more parsimonious. Hence, predictive coding mechanisms, while theoretically appealing and consistent with our results, are not necessary to explain the modulatory effects of perceptual expectation on fMRI adaptation as these effects can readily be explained in terms of top-down attentional

modulation. However, as the present experiment was not designed to distinguish between these 2 hypotheses, further studies will be required to resolve this issue.

Notes

Conflict of Interest: None declared.

References

Ashida H, Lingnau A, Wall MB, Smith AT. 2007. fMRI adaptation reveals separate mechanisms for first-order and second-order motion. J Neurophysiol. 97:1319-1325.

Bartels A, Logothetis NK, Moutoussis K. 2008. fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. Trends Neurosci. 31:444-453.

Boynton GM, Finney EM. 2003. Orientation-specific adaptation in human visual cortex. J Neurosci. 23:8781-8787.

Burock MA, Dale AM. 2000. Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach. Hum Brain Mapp. 11:249-260.

Chong TT-J, Cunnington R, Williams MA, Kanwisher N, Mattingley JB. 2008. fMRI adaptation reveals mirror neurons in human inferior parietal cortex. Curr Biol. 18:1576-1580.

Cohen Kadosh R, Cohen Kadosh K, Kaas A, Henik A, Goebel R. 2007. Notation-dependent and -independent representations of numbers in the parietal lobes. Neuron. 53:307-314.

De Baene W, Vogels R. 2010. Effects of adaptation on the stimulus selectivity of macaque inferior temporal spiking activity and local field potentials. Cereb Cortex. 20:2145-2165.

Deichmann R. 2006. Fast structural brain imaging using an MDEFT sequence with a FLASH-EPI hybrid readout. Neuroimage. 33:1066-1071.

Desimone R. 1996. Neural mechanisms for visual memory and their role in attention. Proc Natl Acad Sci U S A. 93:13494-13499.

Donner TH, Sagi D, Bonneh YS, Heeger DJ. 2008. Opposite neural signatures of motion-induced blindness in human dorsal and ventral visual cortex. J Neurosci. 28:10298-10310.

Efron B, Tibshirani R. 1993. An introduction to the bootstrap. Boca Raton (FL): Chapman & Hall/CRC.

Eger E, Henson RNA, Driver J, Dolan RJ. 2004. BOLD repetition decreases in object-responsive ventral visual areas depend on spatial attention. J Neurophysiol. 92:1241-1247.

Fang F, Murray SO, He S. 2007. Duration-dependent fMRI adaptation and distributed viewer-centered face representation in human visual cortex. Cereb Cortex. 17:1402-1411.

Fang F, Murray SO, Kersten D, He S. 2005. Orientation-tuned fMRI adaptation in human visual cortex. J Neurophysiol. 94:4188-4195.

Gandhi SP, Heeger DJ, Boynton GM. 1999. Spatial attention affects brain activity in human primary visual cortex. Proc Natl Acad Sci U S A. 96:3314-3319.

Gardner JL, Sun P, Waggoner RA, Ueno K, Tanaka K, Cheng K. 2005. Contrast adaptation and representation in human early visual cortex. Neuron. 47:607-620.

Graham DB, Allinson NM. 1998. Characterizing virtual eigensignatures for general purpose face recognition. In: Wechsler H, Phillips PJ, Bruce V, Fogelman-Soulie F, Huang TS, editors. Face recognition: from theory to applications. NATO ASI Series F, Computer and Systems Sciences. Berlin (Germany): Springer-Verlag. p. 445-456.

Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzchak Y, Malach R. 1999. Differential processing of objects under various viewing conditions in the human lateral occipital complex. Neuron. 24:187-203.

Grill-Spector K, Malach R. 2001. fMR-adaptation: a tool for studying the functional properties of human cortical neurons. Acta Psychologica. 107:293-321.

Hammond P, Mouat GS, Smith AT. 1985. Motion after-effects in cat striate cortex elicited by moving gratings. Exp Brain Res. 60:411-416.

Henson RNA, Rugg MD. 2003. Neural response suppression, haemodynamic repetition effects, and behavioural priming. Neuropsychologia. 41:263-270.

Henson RN, Mouchlianitis E. 2007. Effect of spatial attention on stimulus-specific haemodynamic repetition effects. Neuroimage. 35:1317-1329.

Huk AC, Ress D, Heeger DJ. 2001. Neuronal basis of the motion aftereffect reconsidered. Neuron. 32:161-172.

Kaliukhovich DA, Vogels R. Forthcoming 2010. Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. Cereb Cortex. 21:1547-1558.

Kanal E, Borgstede JP, Barkovich AJ, Bell C, Bradley WG, Felmlee JP, Froelich JW, Kaminski EM, Keeler EK, Lester JW, et al. 2002. American College of Radiology White Paper on MR Safety. Am J Roentgenol. 178:1335-1347.

Kanwisher N, McDermott J, Chun MM. 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci. 7:4302-4311.

Kohn A, Movshon JA. 2003. Neuronal adaptation to visual motion in area MT of the macaque. Neuron. 39:681-691.

Kourtzi Z, Kanwisher N. 2000. Cortical regions involved in perceiving object shape. J Neurosci. 20:3310-3318.

Krekelberg B, Boynton GM, van Wezel RJA. 2006. Adaptation: from single cells to BOLD signals. Trends Neurosci. 29:250-256.

Larsson J. 2001. Imaging vision: functional mapping of intermediate visual processes in man [PhD thesis]. Stockholm (Sweden): Karolinska Institutet.

Larsson J, Heeger DJ. 2006. Two retinotopic visual areas in human lateral occipital cortex. J Neurosci. 26:13128-13142.

Larsson J, Landy MS, Heeger DJ. 2006. Orientation-selective adaptation to first- and second-order patterns in human visual cortex. J Neurophysiol. 95:862-881.

Lingnau A, Ashida H, Wall MB, Smith AT. 2009. Speed encoding in human visual cortex revealed by fMRI adaptation. J Vis. 9:3.1-14.

Lingnau A, Gesierich B, Caramazza A. 2009. Asymmetric fMRI adaptation reveals no evidence for mirror neurons in humans. Proc Natl Acad Sci U S A. 106:9925-9930.

Logothetis NK, Wandell BA. 2004. Interpreting the BOLD signal. Annu Rev Physiol. 66:735-769.

Marlin SG, Hasan SJ, Cynader MS. 1988. Direction-selective adaptation in simple and complex cells in cat striate cortex. J Neurophysiol. 59:1314-1330.

Montaser-Kouhsari L, Landy MS, Heeger DJ, Larsson J. 2007. Orientation-selective adaptation to illusory contours in human visual cortex. J Neurosci. 27:2186-2195.

Nestares O, Heeger DJ. 2000. Robust multiresolution alignment of MRI brain volumes. Magn Reson Med. 43:705-715.

Notebaert K, Nelis S, Reynvoet B. 2010. The magnitude representation of small and large symbolic numbers in the left and right hemisphere: an event-related fMRI study. J Cogn Neurosci. 23:622-630.

Piazza M, Pinel P, Le Bihan D, Dehaene S. 2007. A magnitude code common to numerosities and number symbols in human intraparietal cortex. Neuron. 53:293-305.

Priebe NJ, Lampl I, Ferster D. 2010. Mechanisms of direction selectivity in cat primary visual cortex as revealed by visual adaptation. J Neurophysiol. 104:2615-2623.

Priebe NJ, Churchland MM, Lisberger SG. 2002. Constraints on the source of short-term motion adaptation in macaque area MT. I. The role of input and intrinsic mechanisms. J Neurophysiol. 88:354-369.

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.

Rokers B, Cormack LK, Huk AC. 2009. Disparity- and velocity-based signals for three-dimensional motion perception in human MT+. Nat Neurosci. 12:1050-1055.

Sawamura H, Orban GA, Vogels R. 2006. Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the fMRI adaptation paradigm. Neuron. 49:307-318.

Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB. 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science. 268:889-893.

Smith AT, Wall MB. 2008. Sensitivity of human visual cortical areas to the stereoscopic depth of a moving stimulus. J Vis. 8.10(1): 1-12.

Somers DC, Dale AM, Seiffert AE, Tootell RB. 1999. Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. Proc Natl Acad Sci U S A. 96:1663-1668.

Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J. 2006. Predictive codes for forthcoming perception in the frontal cortex. Science. 314:1311-1314.

Summerfield C, Koechlin E. 2008. A neural representation of prior information during perceptual inference. Neuron. 59:336-347.

Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Egner T. 2008. Neural repetition suppression reflects fulfilled perceptual expectations. Nat Neurosci. 11:1004-1006.

Tolias AS, Keliris GA, Smirnakis SM, Logothetis NK. 2005. Neurons in macaque area V4 acquire directional tuning after adaptation to motion stimuli. Nat Neurosci. 8:591-593.

Tootell RB, Reppas JB, Dale AM, Look RB, Sereno MI, Malach R, Brady TJ, Rosen BR. 1995. Visual motion aftereffect in human cortical area MT revealed by functional magnetic resonance imaging. Nature. 375:139-141.

Vautin RG, Berkley MA. 1977. Responses of single cells in cat visual cortex to prolonged stimulus movement: neural correlates of visual aftereffects. J Neurophysiol. 40:1051-1065.

Verhoef B-E, Kayaert G, Franko E, Vangeneugden J, Vogels R. 2008. Stimulus similarity-contingent neural adaptation can be time and cortical area dependent. J Neurosci. 28:10631-10640.

Vuilleumier P, Schwartz S, Duhoux S, Dolan RJ, Driver J. 2005. Selective attention modulates neural substrates of repetition priming and "implicit" visual memory: suppressions and enhancements revealed by FMRI. J Cogn Neurosci. 17:1245-1260.

Wandell BA, Dumoulin SO, Brewer AA. 2007. Visual field maps in human cortex. Neuron. 56:366-383.

Yi D-J, Chun MM. 2005. Attentional modulation of learning-related repetition attenuation effects in human parahippocampal cortex. J Neurosc. 25:3593-3600.