

Concurrent Repetition Enhancement and Suppression Responses in Extrastriate Visual Cortex

Vincent de Gardelle^{1,2}, Monika Waszczuk¹, Tobias Eegner³ and Christopher Summerfield¹

¹Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, UK ²Laboratoire de Psychologie de la Perception, Université Paris Descartes, Paris 75006, France ³Center for Cognitive Neuroscience, Duke University, Durham, NC 27708, USA

Address correspondence to V. de Gardelle. Email: vincent.gardelle@gmail.com

Visual cortical responses are usually attenuated by repetition, a phenomenon known as repetition suppression (RS). Here, we use multivoxel pattern analyses of functional magnetic resonance imaging (fMRI) data to show that RS co-occurs with the converse phenomenon (repetition enhancement, RE) in a single cortical region. We presented human volunteers with short sequences of repeated faces and measured brain activity using fMRI. In an independently defined face-responsive extrastriate region, the response of each voxel to repetition (RS vs. RE) was consistent across scanner runs, and multivoxel patterns for both RS and RE voxels were stable. Moreover, RS and RE voxels responded to repetition with dissociable latencies and exhibited different patterns of connectivity with lower and higher visual regions. Computational simulations demonstrated that these effects must be due to differences in repetition sensitivity, and not feature selectivity. These findings establish that 2 classes of repetition responses coexist within 1 visual region and support models acknowledging this distinction, such as predictive coding models where perception requires the computation of both predictions (which are enhanced by repetition) and prediction errors (which are suppressed by repetition).

Keywords: face perception, fMRI, predictive coding, repetition enhancement, repetition suppression

Introduction

When a stimulus is repeated, the neural response elicited is reduced, a phenomenon termed “repetition suppression” (RS). RS is thought to be a fundamental property of brain responses (Grill-Spector et al. 2006). Brain imaging studies have used this phenomenon as a methodological tool in many domains, as for instance it allows to dissect the distinct neural levels of stimulus processing (Vuilleumier et al. 2002; Rotshtein et al. 2005) or to assess stimulus processing under implicit or subliminal condition (Henson et al. 2000, 2002; Dehaene et al. 2001; Naccache and Dehaene 2001; Turk-Browne et al. 2006; Kouider et al. 2010).

Surprisingly, although RS is typically observed in extrastriate regions when a visual stimulus repeats, the converse phenomenon, “repetition enhancement” (RE) has nevertheless been reported under some circumstances. For example, RE may occur under conditions of low visibility (Turk-Browne et al. 2007) or when unfamiliar stimuli are repeated (Henson et al. 2000), and a recent study has shown that within the same cortical region, repetition effects for unfamiliar stimuli can turn from enhancement to suppression when the number of stimulus repetition increases (Muller et al. 2012). These findings prompt the question of whether RS and RE may reflect the response properties of 2 functionally distinct neural populations that co-exist within the same cortical

region, with some neural elements’ responses enhanced and some suppressed by repetition. This architecture would be consistent with the proposal that perception relies both on prediction signals (that would be enhanced by repetitions) and on prediction error signals (that would be reduced with repetitions; e.g. Mumford 1992; Friston 2005).

Demonstrating the co-existence of “positive-going” and “negative-going” responses to repetitions seems challenging at first sight. First of all, previous studies involving single-neuron electrophysiology or whole-brain imaging have reported mostly RS in stimulus-selective sensory regions, and reports of RE seem to depend on the use of unfamiliar stimuli or degraded perceptual conditions. Secondly, standard brain imaging techniques of fMRI data are not suited to establish the existence of intermingled populations of suppression and enhancement voxels. Indeed, face repetition might elicit consistent enhancement in a minority of face-responsive voxels, but spatial smoothing will erase this response if these are surrounded by a majority of voxels showing RS. Consistent with this possibility, when we re-analyzed the unsmoothed data of a previously collected fMRI data set, we found that only around 65% of voxels in an independently defined extrastriate face-responsive region were suppressed by face repetition. This prompts the question of whether measurements taken from the remaining 35% of voxels are simply noise, or whether they index a functionally significant RE signal.

To arbitrate among these possibilities, we conducted a new study using fMRI in conjunction with multivoxel pattern analysis to assess whether concurrent RE and RS signals could be found simultaneously in the ventral visual stream. We found that approximately one-third of all fusiform-face area (FFA) voxels displayed RE responses while two-thirds exhibited RS responses. Crucially, the separation between RS and RE voxels was consistent across scanner runs, and multivoxel patterns associated with both RS and RE were independently stable over time: In other words, “repetition sensitivity” was a significant characteristic of visual responses in our experiment. Moreover, when we compared the timing of the blood oxygenation level dependent (BOLD) response for RS and RE voxels, and their pattern of connectivity with other regions, we found that they differed reliably. Together, these findings suggest that voxels exhibiting predominantly suppressed and enhanced responses to repetition make distinct contributions to the computations underlying visual perception.

Materials and Methods

Participants

Eighteen right-handed volunteers (6 women, age range 18–35), reporting normal- or corrected-to-normal vision and

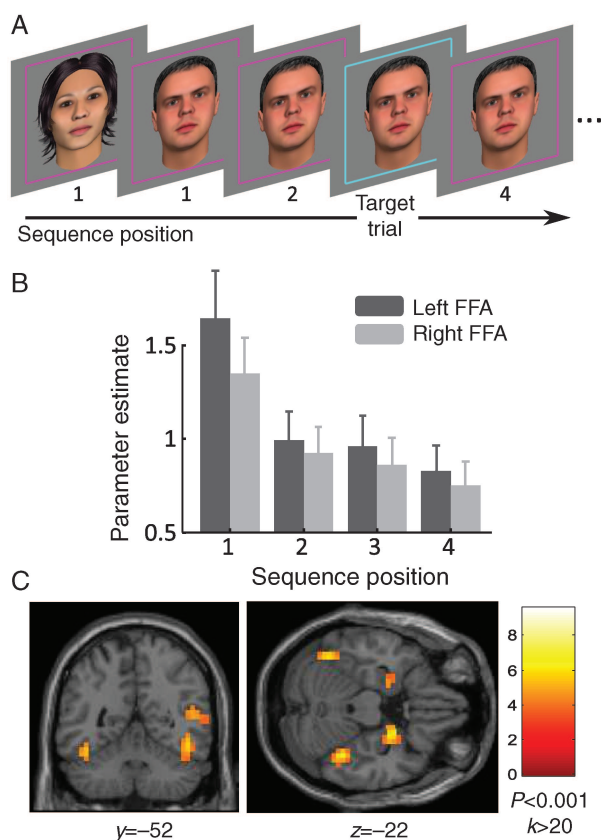


Figure 1. (A) The trial sequence. Each unique face exemplar occurred on 14 consecutive trials. All trials were separated by a blank screen (2–6 s, jittered). (B) Sequence position effect. Univariate BOLD responses to faces occurring in sequence positions 1–4, in the left FFA (dark bars) and right FFA (light bars). Error bars represent standard error of the mean (SEM) across subjects. (C) Face localizer results. Group-level comparison (faces > houses) rendered on coronal ($y = -52$) and axial ($z = -22$) slices of a standard brain. Results are thresholded at $P < 0.001$ (uncorrected) and $k > 20$. The red-white scale indicates t -values.

no history of psychiatric or neurologic illness, were recruited from Oxford University. They provided informed consent and received £25 compensation for their time. The experiment was approved by the local ethics committee.

Stimuli

Face stimuli were created using FaceGen (Singular Inversions, Ontario, Canada): 200 images (400×400 pixels) represented faces (half males) of variable age, with hair, with a frontal view. Pictures were in color and not degraded. All stimuli were presented centrally on a gray background. The faces were surrounded by a blue or pink square frame (Fig. 1A), which appeared 500 ms before the face onset. Stimulus presentation was done using the PsychToolBox for Matlab.

Procedure

Participants viewed a sequence of faces surrounded by colored frames and were asked to make a gender judgment (using a magnetic resonance imaging [MRI]-compatible response device) only when the frame had a target color (9% of all trials), while passively viewing the other faces. The target color and the response mapping were indicated at the beginning of each block, and counterbalanced across subjects. On target trials, participants received auditory feedback

following their response: A high-pitched tone (800 Hz) for correct responses within 1 s and a low-pitched tone (400 Hz) for incorrect or slow (>1 s) responses. Inter-trial intervals were uniformly jittered between 2 and 6 s. The experiment involved 6 blocks of 9 min each, for a total of 546 trials per participant. Each face was used only once, in a sequence of 1–4 successive presentations (denoted as sequence positions 1–4 in what follows). In the whole experiment, there were 15, 24, 35, and 96 sequences of length 1–4, respectively. Target trials were balanced across positions in the mini-sequences and were equally likely to involve a gender change or not when occurring at the beginning of the mini-sequence.

Localizer Task

After the main task, participants performed a localizer task consisting of 12 alternating “face” and “house” blocks, in which 15 (either houses or faces) images were presented in a sequence, and participants had to press a button at immediate stimulus repetitions (“1-back task”), of which, there were 0, 1, or 2 per block. Stimuli were presented for 750 ms followed by 250 ms of fixation, and 10-s periods were inserted between blocks. Images were 300×300 pixels, black-and-white photographs on a black background.

fMRI Data Acquisition and Preprocessing

MRIs were acquired with a 3T Siemens VERIO scanner with a 32-channel head coil using a standard echo-planar imaging sequence. Whole-head T_2^* -weighted echo-planar images were continuously acquired with a repetition time of 2 s, echo time of 30 ms. We acquired 270 volumes per block, plus 3 dummy scans discarded before the analyses. Each volume included $64 \times 64 \times 36$ voxels of $3 \times 3 \times 3$ mm. A high-resolution T_1 -weighted structural image was also obtained (voxel size = $1 \times 1 \times 1$ mm). For standard preprocessing and univariate statistical analyses, we used SPM8 (Wellcome Department of Cognitive Neurology, London, United Kingdom). All other analyses were done with custom scripts for Matlab (Mathworks, Natick, MA, United States of America). We also used xjview (<http://www.alivelearn.net/xjview>) to visualize the data, construct mask and conjunction images. For each participant, we first realigned all functional images, then we co-registered (rigid body transformation) the subject’s anatomical scan to the mean functional image, and then co-registered the participant’s data to the Montreal Neurological Institute (MNI) template brain. We then normalized each subject’s data to the template brain space, using segmented probabilistic maps for gray matter, white matter, and cerebro-spinal fluid. Functional images were resampled ($4 \times 4 \times 4$ mm voxels) and spatially smoothed (8-mm full-width half-maximum (FWHM) Gaussian kernel). Note, however, that in our univariate processing stream, these normalization and smoothing stages preceded both first- and second-level statistical analyses, as in standard approaches. However, for multivariate analyses, normalization and smoothing followed the first-level statistical analysis, in order to maintain the details of the local patterns in individual subjects.

Univariate Analysis

Our univariate analyses used a generalized linear model (GLM) approach. A 128-s temporal high-pass filter was applied to remove low-frequency scanner artifacts. Temporal

autocorrelation in the time series data was estimated using restricted maximum-likelihood estimates of variance components using a first-order autoregressive model (AR-1), and the resulting non-sphericity was used to form maximum-likelihood estimates of the activations. Our GLM included regressors coding for onsets and durations of stimuli or events, which were then convolved with the canonical hemodynamic response function (HRF) and regressed against the observed fMRI data. Experimental blocks were modeled using separate regressors, and constant terms for each block were included. Additionally, motion parameters were included as nuisance variables. In our localizer task, we used regressors to code for face and house blocks. Face > house contrast images were computed in each subject and subjected to a second-level *t*-test across subjects, to determine a set of face-sensitive regions of interest (ROIs) to provide a priori constraints to our subsequent analyses of the repetition effects in the main task. In our main task, we used 4 regressors coding for the occurrence of faces in position 1–4, and a fifth regressor indicating the target trials (if there were any) in the block. After estimation of the model, for each subject the beta parameter for the 1–4 positions was aggregated as a linear “repetition effect” using a $[-3, -1, +1, +3]$ contrast over sequence positions 1–4. A simple *t*-test across subjects was then used at the second level. A negative effect would indicate RS, and a positive effect would indicate RE.

ROI Definitions

We used the localizer data to define our (left and right) FFA and amygdala ROIs as 10-mm-radius spheres centered on local maxima (around the anatomically defined fusiform gyrus and amygdala) of the *t*-statistic map for the face > house comparison across subjects; Fig. 1). The MNI coordinates for these ROIs were $[-42, -56, -18]$ (left FFA), $[42, -44, -22]$ (right FFA), $[-18, -8, -18]$ (left amygdala), and $[22, -4, -14]$ (right amygdala). For the connectivity analyses (see below), we defined 2 other ROIs, in the left posterior middle occipital gyrus (MOG, centre of mass $[-27.3, -93.8, 1.5]$) and right anterior middle temporal gyrus (MTG, $[64, -9.4, -15.3]$). We selected them as clusters showing RS (MOG) or RE (MTG) in the main univariate analyses.

Multivariate Analyses

To preserve the details of the multivoxels patterns, we applied subject-to-template normalization and spatial smoothing to our functional images only after the first-level statistic stage for our multivariate analyses. Within each block, the raw data timeseries for each voxel were temporally high-pass filtered (128 s), and normalized to a mean of zero and standard deviation of one. Beta parameters for each voxel were assessed by multiplying these timeseries with the pseudo-inverse of a temporally filtered design matrix, which was formed by our 5 regressors (indicating sequence positions 1–4 and target trials) convolved with the canonical HRF and by the nuisance parameters. No prewhitening or correction for serial auto-correlation was used. We computed the repetition effect as a $[-3, -1, +1, +3]$ contrast over the beta parameters sequence positions 1–4 and defined RS (respectively RE) voxels as showing negative (respectively positive) responses to this contrast.

We now describe 2 analyses that we have carried out on these data. Importantly, in these analyses, the selection process and the statistical tests applied are independent. At the group level, we used *t*-tests across participants on the resulting Z-scores.

In the “sign-consistency” analysis (Fig. 2), we assessed whether voxels showing RS in $N-1$ blocks would be relatively more likely to also show RS in the remaining block. To do so, we computed the quantity $(P(\text{RS}|\text{RS}) - P(\text{RS}))/P(\text{RS})$, where $P(\text{RS})$ is the probability that voxels in 1 block show RS, and $P(\text{RS}|\text{RS})$ the probability that they show RS given that they already do so in the remaining $N-1$ blocks. We obtained a Z-score by comparing this quantity to the distribution of this statistic for 1000 datasets obtained by permuting the voxels in the $N-1$ blocks. Note that this corrected statistic was then identical for RS and RE voxels: If RS tended to remain the same, the RE voxels necessarily had to remain the same as well.

In the “pattern-consistency” analyses (Fig. 3), we assessed whether the response pattern of RS voxels was stable across runs. We selected RS voxels from $N-1$ blocks and computed for these voxels the Pearson’s correlation between their responses in the $N-1$ blocks and their responses in the

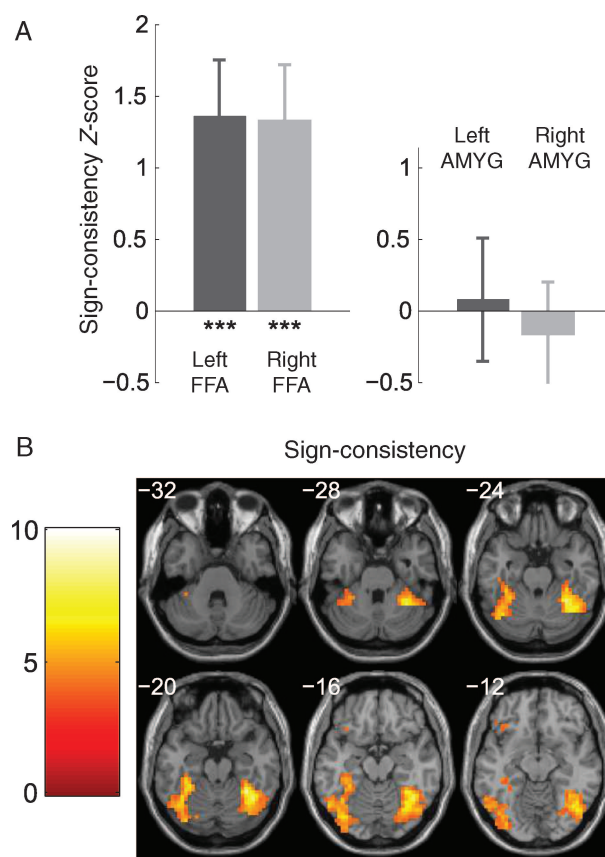


Figure 2. (A) Segregation of RS and enhancement voxels. Mean Fisher’s Z-score in the FFA and amygdala (AMYG), for the measure of sign consistency (see Materials and Methods). Error bars represent SEM across subjects; stars indicate significance against a null effect ($***P < 0.001$). (B) Searchlight analysis. Regions showing significant sign consistency ($P < 0.001$, $k > 20$) rendered on 6 transverse slices ($z = -32$ to $z = -12$) of a template brain. The color bar indicates the value of the *T*-statistic at the group level.

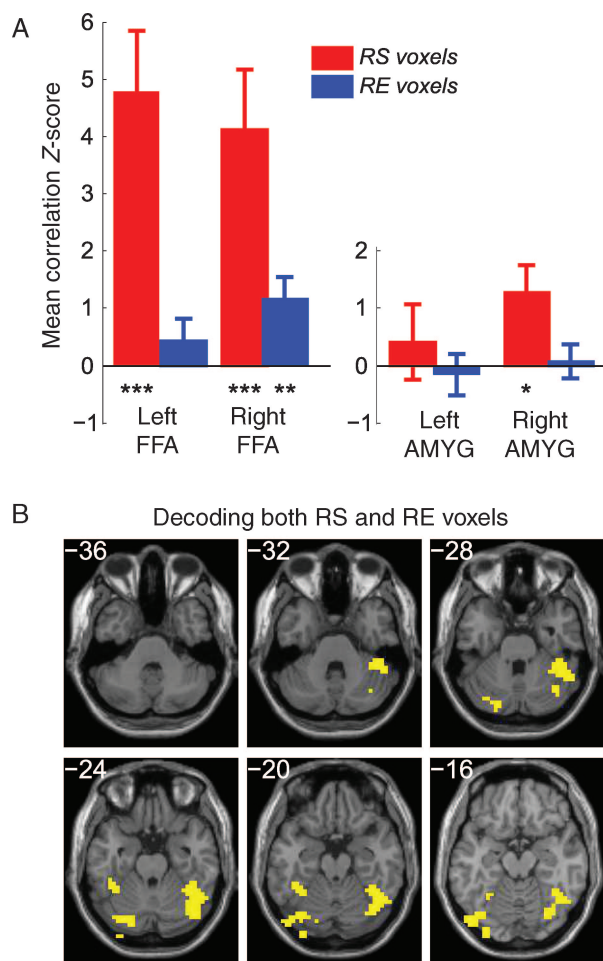


Figure 3. (A) Multivariate pattern consistency for RS and enhancement voxels. Mean Fisher's Z-score for the voxel-wise correlation (see Materials and Methods) assessing pattern consistency in the FFA and amygdala (AMYG), separately for RS voxels (in red) and RE voxels (in blue). Error bars represent SEM across subjects; stars indicate significance against a null correlation (* $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$). (B) Conjunctive searchlight analysis. Regions showing significant pattern consistency of both RS voxels ($P < 0.005, k > 20$) and RE voxels ($P < 0.005, k > 20$), rendered on 6 transverse slices ($z = -36$ to $z = -16$) of a template brain.

remaining block. Repeating this leave-one-out procedure on the N different blocks gave us N estimates of the correlation, which were converted to Fisher's Z-scores and averaged for each subject. The procedure was also applied (separately) to RE voxels.

Searchlight

Although initially these analyses were applied to the FFA only, we could apply it to other face-sensitive regions identified in the localizer (e.g. the amygdala), and eventually to the whole brain. To do so, we used a "searchlight" approach in which a 10-mm-radius sphere is moved across the whole brain defining an ROI within which the same statistics are assessed and saved at the central voxel of the sphere. We generated full statistic maps for each participant, which were then normalized to the template MNI brain, re-sampled at $4 \times 4 \times 4$ mm, spatially smoothed (8-mm FWHM Gaussian kernel) and finally subjected to t -tests across participants at the group level.

FIR Analyses

To assess the response latency (Fig. 4) or the connectivity (Fig. 5) of the RS and RE voxels in the FFA, we had to examine in further details their hemodynamic responses. To this end, we ran a separate analysis in which we used a finite impulse response (FIR) model applied to the mean BOLD activity, separately for the RS and RE voxels. In each subject and each block, we first identified RS voxels (respectively RE voxels) as showing negative (respectively positive) responses to the repetition effect (from the main univariate analysis run on unsmoothed data). We then extracted the raw timeseries for each of these voxels and applied high-pass temporal filtering and normalization (mean of 0, standard deviation of 1). We then fit a GLM (using Matlab `glmfit` routine) in which the timeseries of the average activity of these voxels was predicted by 100 regressors, corresponding to 20 FIRs for each of the 5 conditions (4 stimulus position + target trials). The FIRs' onsets were at $-2, 0, 2, \dots, 36$ s poststimulus onset, and their associated parameter estimates thus indicate the timecourse of the hemodynamic response, with respect to each condition. No prewhitening or correction for serial auto-correlation was used in this analysis. For the control analyses presented in Figure 6 (see simulation data and analyses), we multiplied these FIR timecourses with the canonical HRF to assess the parameter estimate for each sequence position.

Peak Latency Analyses in the FFA

For each subject and condition, the FIR timecourses were extracted from the right FFA and the linear contrast was then computed over stimulus positions. For each subject, the peak

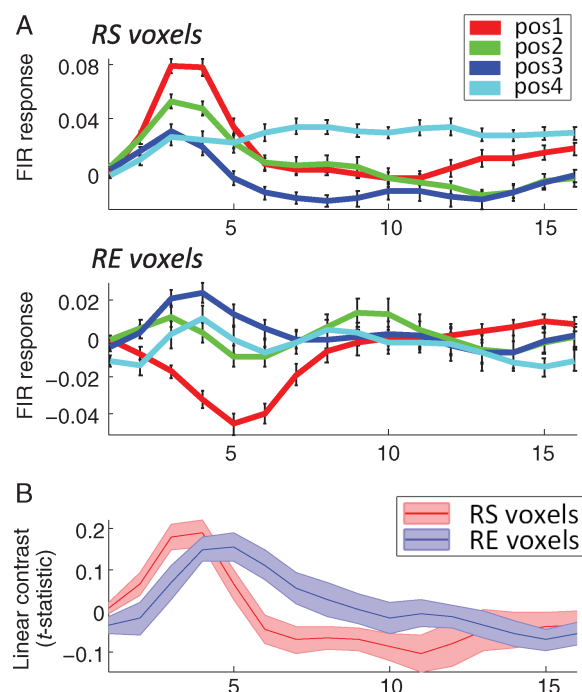


Figure 4. (A) Temporal response of RS and RE voxels. Average BOLD response of RS and RE voxels in the right FFA to the different stimulus position (positions 1–4), assessed at different time points (scans of 2 s) after stimulus onset using a FIR model. The timecourses were baselined and smoothed in time (see methods). (B) Timing of the repetition effect. Time course of the t -statistic of the linear effect across stimulus position.

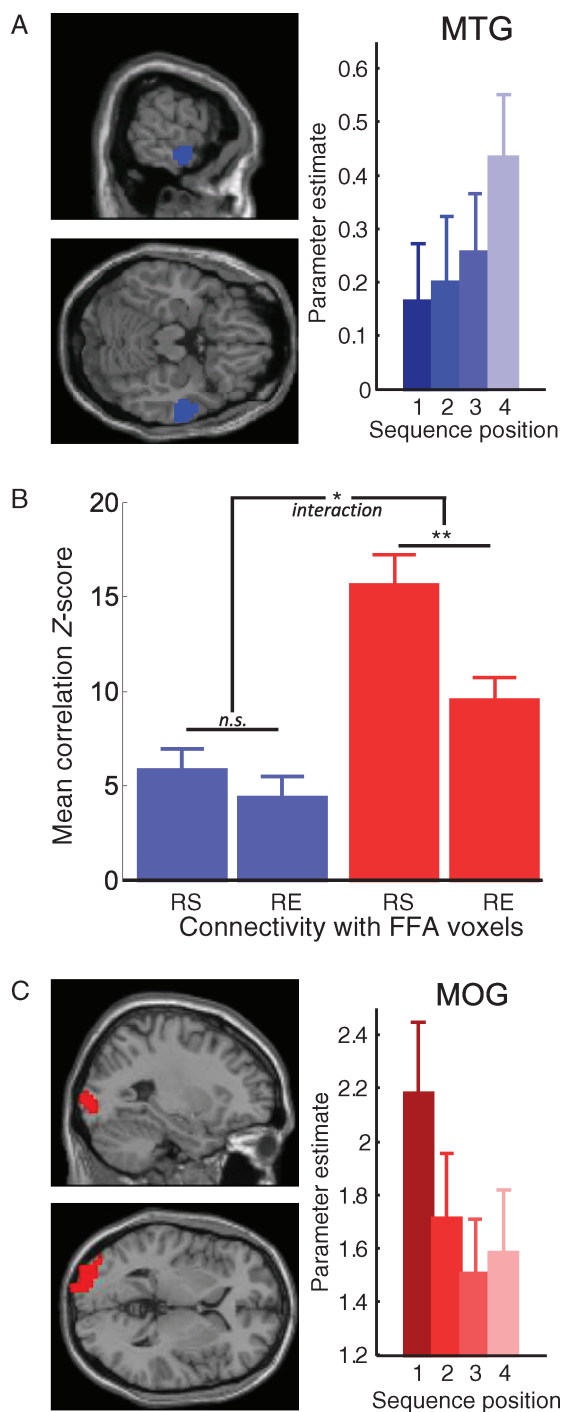


Figure 5. Connectivity analyses between FFA voxels and MTG and MOG. (A–C) Rendering of the MTG (A, in blue) and MOG (C, in red) ROI on a template brain, with their respective univariate BOLD responses to faces occurring in sequence positions 1–4. (B) Mean correlation Z-scores between the residual activities (from a FIR model) in the MTG (blue bars) and MOG (red bars) and the residual activities in the RS and RE populations of the FFA.

latency was defined as the point at which the statistic reached a maximum in the [2 20] seconds period poststimulus onset (we checked that no latencies fell at the limit of this interval). A *t*-test was used to compare the peak latency for RS and RE voxels, across subjects. The same analysis was run in the amygdale as a control ROI. In Figure 4, the presented

timecourses are zero-baselined (using the average of the first 3 datapoints) for plotting purposes only.

Connectivity Analyses

The connectivity between 2 regions can be seen as the extent to which the unexplained activities (e.g. the residuals in the FIR model described above) co-vary between these 2 regions. Here, we assessed connectivity between the FFA RS and RE voxels and a “lower” ROI (MOG) and a “higher” ROI (MTG) within the visual ventral stream. We extracted the timeseries of the residuals for each ROI in each subject and assessed 4 correlations between these timeseries (RS/MOG, RE/MOG, RS/MTG, and RE/MTG). We then applied Fisher’s *r*-to-*Z* transformation and used the resulting *Z*-scores in statistical tests (*t*-tests, analysis of variance) across subjects (Fig. 5). The same analysis was run in the amygdale as a control ROI (Supplementary Fig. S1).

Simulated Data and Analyses

We simulated 2 datasets, each constituted by 18 virtual subjects, with 100 voxels for each subject, in 2 fMRI runs. The first dataset (repetition sensitivity) was generated under the hypothesis that all voxels had the same preference for the stimulus, but some were responding positively to repetitions (RE) and some negatively (RS). In the second dataset (“face selectivity”), all voxels contributed with reduced responses to repeated stimuli but had different preferences for the stimulus. At each voxel k (1–100), the response to face i (1–4) in run r (1 or 2) was noted $y_{k,i,r}$ and generated as the sum of 3 components, the main face response (noted c_k), the additional repetition effect (noted a_k), and some independent and identically distributed random noise (noted $n_{k,i,r}$). If we note $N(\mu, \sigma)$ a random number drawn from a normal distribution of mean μ and standard deviation σ , our datasets can be summarized in the following equation for all k , i , and r : $y_{k,i,r} = c_k - i \cdot a_k + n_{k,i,r}$, where the parameters are:

1. For the repetition sensitivity model: $a_k = N(0.5, 1)$, $c_k = C = 2$, $n_{k,i} = N(0, 6)$.
2. For the face selectivity model: $a_k = A = 0.5$, $c_k = N(2, 1)$, $n_{k,i} = N(0, 6)$.

We ran univariate analyses by computing the linear contrast across the repetition sequence at each voxel in both runs, and we sorted voxels as RS versus RE in a post hoc manner, depending on whether their response was negative versus positive. We computed our multivariate analyses by using the same procedure as for our real data. We report (Fig. 6) the average correlation coefficient for each “virtual subject” for the pattern-consistency analysis, for both RS and RE voxels, and *t*-tests across subjects.

Results

Participants viewed a continuous stream of faces in which each trial-unique exemplar was repeated 1–4 times consecutively (Fig. 1A), while performing an incidental gender-judgment task on a subset of target faces indicated by a colored frame (these target trials were modeled separately and excluded from our analyses). As expected, most participants exhibited a high level of accuracy in the simple gender-judgment task but some ($N=4$) did seem to have forgotten the stimulus-button contingencies over the time of the

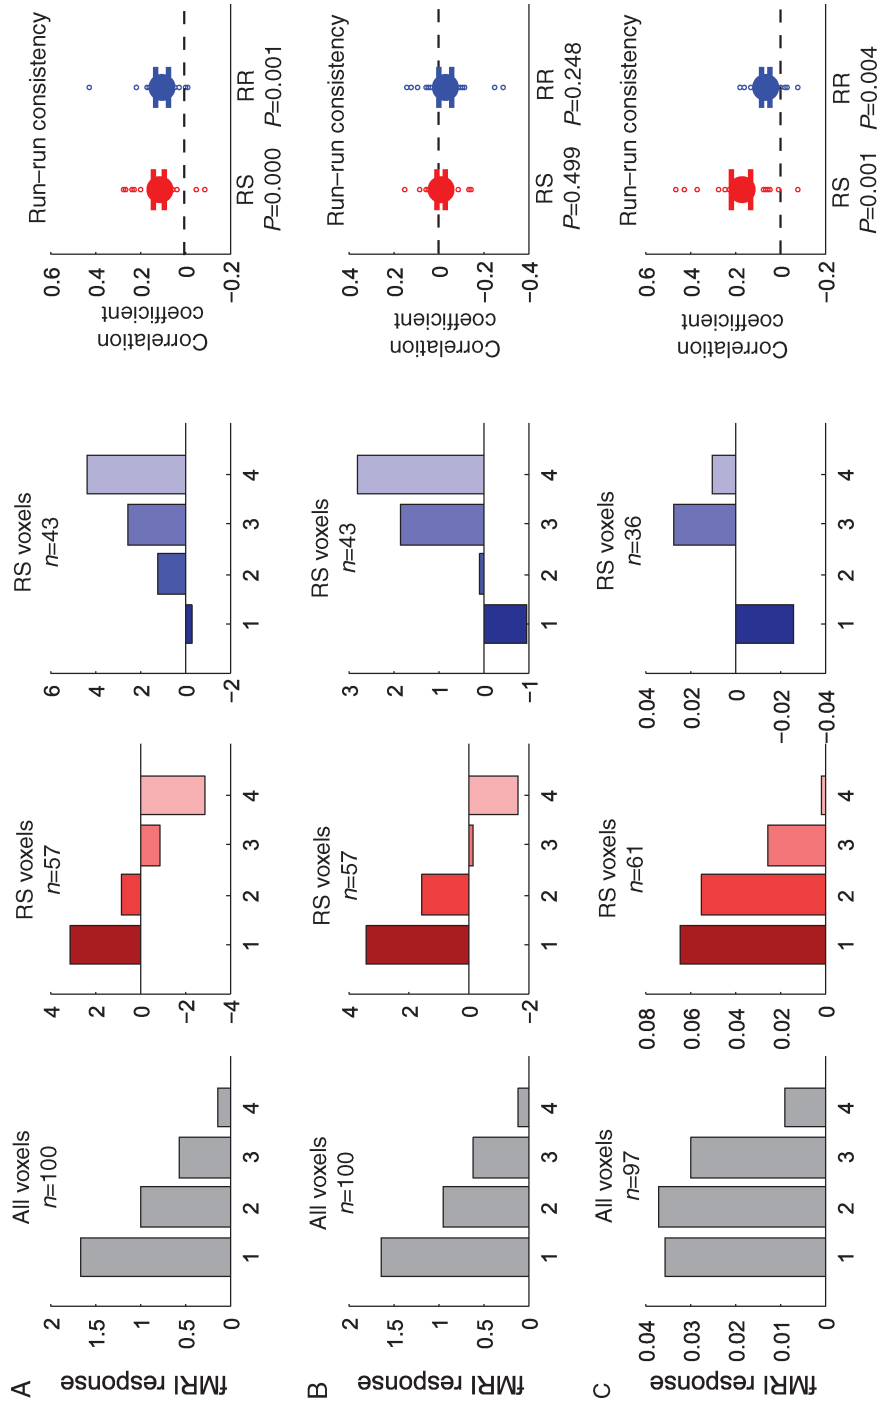


Figure 6. Univariate and multivariate analyses applied to simulated and real data. (A) Simulated data for the repetition sensitivity model, (B) simulated data for the face selectivity, and (C) real data. The 3 bar graphs show univariate responses (y-axis: Arbitrary units) as a function of sequence positions (1–4), separately for all voxels (gray), RS voxels (red), and RE voxels (blue). The rightmost graph shows the correlation coefficient r for each virtual subject, along with the group average and standard error, separately for RS voxels (red) and RE voxels (blue); P -values refer to the t -test against zero correlation across subjects. The data in C are calculated from the fHR responses in Figure 3 multiplied with a canonical HRF profile.

experiment, possibly because of the rare occurrence of target trials. Overall, mean accuracy on target trials was 76%. Accuracy on target trials did not depend on stimulus position ($F < 1$), which suggest that even if there might have been occasional lapses of attention during the experiment, they would have been equally distributed between the different conditions of interest.

Univariate Analyses

We focused our analyses on an extrastriate face-sensitive region defined in an independent localizer task (hereafter, FFA). Standard univariate analyses revealed an anticipated effect of RS: FFA activity on average decreased over consecutive repetitions of each exemplar (Fig. 1). In all the following analyses, we quantify this repetition effect as the linear contrast across sequence positions 1–4. Strong RS was observed (i.e. a negative-going response to the linear contrast) in the standard univariate analysis on spatially smoothed data, in both the left ($t_{(17)} = 4.08$, $P < 0.001$) and right FFA ($t_{(17)} = 5.25$, $P < 0.001$). However, when we omitted the spatial smoothing step from the standard univariate processing stream, approximately one-third (on average 37%) of FFA voxels exhibited a numerically positive-going response to the linear contrast, that is, a gradual enhancement across repetitions.

Sign-Consistency Analysis

We reasoned that if the 37% of observed RE responses simply reflect measurement noise, we would not expect the RS versus RE status of each voxel to be consistent across runs. By contrast, if RS and RE voxel populations reflect distinct functional contributions to visual perception, the sign of each voxel's response should be consistent across scanner runs. We found strong evidence supporting the latter hypothesis. The sign-consistency analyses (see Materials and Methods and Fig. 2A) revealed that the segregation between RS and RE voxels was reliable: Voxels classified as RS (or as RE) from their response in $N - 1$ blocks were more likely to be of the same type in the remaining block, both in the left FFA ($t_{(17)} = 3.42$, $P < 0.005$) and right FFA ($t_{(17)} = 3.50$, $P < 0.005$). Thus, voxels showing RS and enhancement were segregated in the FFA in a consistent manner. This was not the case in the amygdala for instance, where no consistent segregation was found between RS and RE voxels. Note however that from the results of this analysis alone, one cannot distinguish the contribution of RS and RE voxels.

Pattern-Consistency Analysis

Thus, we then asked whether or not each of these subpopulations makes stable functional contribution to visual perception, by assessing whether or not it responds with consistent response patterns over scanner runs. Our pattern-consistency analyses (see Materials and Methods and Fig. 3A) revealed significant run–run correlations in the spatial pattern of FFA BOLD activity, independently for RS and RE voxels. Indeed, RS voxels considered in isolation were significantly correlated across runs, in both the left FFA ($t_{(17)} = 4.56$, $P < 0.001$) and right FFA ($t_{(17)} = 4.01$, $P < 0.001$). Critically, RE voxels also exhibited a consistent pattern in the right FFA ($t_{(17)} = 3.08$, $P < 0.01$). Again, we used the amygdala as a control ROI and here we found no consistent pattern for the RE voxels. Thus, in the right FFA, both RS and RE voxels

showed independently a multivariate profile that was stable across time. These findings are in agreement with the view that these effects reflect true functional differences in the computations performed by the underlying neural populations and argue against the possibility that RE voxels are simply reflecting measurement noise from a general RS population. Note that all these analyses do not give rise to circularity, as the ROI is defined in a separate localizer session (and from an orthogonal, face > house contrast), and as the leave-one-out procedure we used here ensured independence between the selection criterion and the statistical test employed.

Searchlight

We then used a whole-brain searchlight approach (methods) to ascertain how specific these effects were with respect to regions involved in face individuation. Sign consistency was found bilaterally at the posterior occipital sites (BA 18 and 19) and extended along the visual stream and fusiform gyri (Fig. 2B), but also along the dorsal stream with peaks in the precuneus (left $[-26, -68, 34]$, right $[22, -68, 30]$), and finally at a bilateral cluster in the middle frontal gyrus (left peak $[-42, 20, 42]$, and right peak $[42, 16, 50]$). When we looked for regions in which both RS and RE voxels exhibited significant pattern consistency across runs ($P < 0.001$ uncorrected, $k > 20$ for both maps), we identified 2 clusters ($k > 20$) which were the right fusiform region and the precuneus ($[-26, -72, 34]$). The left fusiform region also emerged at a more liberal threshold in this conjunction analysis ($P < 0.005$ uncorrected, $k > 20$ for both maps, Fig. 3B). No other brain region, including those identified by the localizer (such as the amygdala, Fig. 3A and Supplementary Fig. S1), showed consistency for both RS and RE. For completeness, we also report as Supplementary Material the whole-brain results for both the univariate effects and the multivariate pattern-consistency effects (Supplementary Fig. S2).

Peak Latency Analyses

Two further observations support the claim that RS and RE voxels make dissociable functional contributions to information processing in the FFA. First, we used a FIR model to assess the HRFs associated with RS and RE voxels (see Materials and Methods and Fig. 4) to look for potential differences in the shape of the HRF between these 2 voxel populations. Crucially, examination of the peak latency for the linear contrast across position revealed that RS voxels responded faster than RE voxels (mean peak latency: 7.1 vs. 10 s after stimulus onset, $t_{(17)} = 3.01$, $P < 0.01$). This was not the case for other face-responsive regions, like the amygdala, where the latency of the HRF responses of the RS and RE populations was indistinguishable ($P > 0.1$, Supplementary Fig. S1). Further examination of the temporal profiles suggests that the latency effect observed could reflect a delayed peak for the negative response of RE voxels in the sequence position 1. We therefore compared directly the peak latency for the response of RS (positive peak) and RE (negative peak) voxels to the first face (sequence position 1). A t -test across participants indeed confirmed that the peak response to the first face was delayed for RE compared with RS voxels ($t_{(17)} = 2.95$, $P < 0.01$), which constitutes a significant difference in time between the 2 populations of voxels.

Connectivity Analyses

Secondly, we tested whether RS and RE voxels in the FFA differed in their pattern of functional connectivity, in particular with other regions showing suppression and enhancement responses. Indeed, according to predictive theories of perception, neural units sensitive to predictions (i.e. RE voxels) and prediction error (RS voxels) differ in their connectivity, with error signals flowing forward in a bottom-up fashion, and predictions fed back from higher cortical stages in a top-down manner. We defined a lower visual ROI (MOG) and a higher visual ROI (MTG), which exhibited main univariate effects of RS and enhancement, respectively, and measured their connectivity with the RS and RE voxels in the FFA (Fig. 5). Connectivity between 2 regions was assessed by correlating the timeseries of their residuals after removing the main response to the 4 stimulus positions using a FIR model (see Materials and Methods) as described previously (Summerfield et al. 2006; Norman-Haignere et al. 2012). Critically, the expected interaction between the FFA populations (RS vs. RE voxels) and the other visual ROIs (MOG vs. MTG) was significant in the ANOVA ($F_{1,17}=7.79$, $P=0.02$). More specifically, the lower MOG region was connected more with RS than with RE voxels ($t_{(17)}=3.50$, $P<0.005$), while no difference was found for the connectivity with the higher MTG region ($t_{(17)}=1.67$, $P>0.10$). This interaction occurred on top of 2 main effects, indicating that residuals of both FFA sub-populations were correlated more with the residual of the lower visual region (MOG vs. MTG: $F_{1,17}=43.22$, $P<0.001$), and that the RS voxels were more connected to other regions than the RE voxels (RS vs. RE: $F_{1,17}=11.82$, $P<0.005$). Further post hoc *t*-tests indicated that all correlations were significantly different from zero (all $P<0.005$). For completeness, we carried out similar connectivity analyses substituting the amygdala for the FFA, but found no differences in connectivity (all $P>0.1$, Supplementary Fig. S1).

Simulation Analyses

Could our results be an artifact of differing degrees of face selectivity among FFA voxels? In other words, could our stability results be observed if all voxels in the FFA were equally suppressed by repetitions (RE being a consequence of the noise) but differed in their main responses to faces? To formally compare these face selectivity and repetition sensitivity accounts of our multivariate findings, we generated simulated datasets under both hypotheses and analyzed them in the same way as our human fMRI data (see Materials and Methods and Fig. 6). In the repetition sensitivity dataset (Fig. 6A), all voxels had the same face selectivity but variable repetition sensitivity. In the face selectivity dataset (Fig. 6B), all simulated voxels were suppressed to an identical degree by repetitions, but they exhibited different sensitivity to the main effect of face stimulation. Random noise was added to all voxels' responses in 2 runs, and voxels were classified post hoc as RS or RE from their response to the linear contrast across the sequence, akin to the treatment of the empirical data. Each dataset comprised 18 virtual subjects, each contributing with 100 voxels, for comparison with our real fMRI data (Fig. 6C).

Both datasets recreated the effects shown in our univariate analyses, including a positive response to faces across all responses, and an increasing attenuation of this response with repetition. Both simulations also predicted the

negative response to the initial face in voxels classified as RE, which is a natural consequence of those voxels with a negative response to the first face being more likely to have an overall positive-going slope (i.e. a selection bias). Crucially, however, is that only in the repetition sensitivity dataset did multivariate analyses reveal significant sign consistency ($P<0.001$) and pattern consistency for both RS voxels ($P<0.001$) and RE voxels ($P=0.001$), while in the face selectivity dataset these statistics were all centered on zero. Additionally, as for the observed data, the multivariate pattern consistency of RE voxels was weaker, probably because RE voxels were less numerous overall.

Discussion

We isolated 2 sub-populations of voxels in face-sensitive extrastriate visual cortex, one whose responses were suppressed by repetition of face stimuli (RS voxels), and one whose responses were enhanced by repetition (RE voxels). We observed that RS and RE voxels can be dissociated in different ways: Their segregation in sign were consistent and their response profiles were separately correlated across measurements (stability), they responded with a different temporal profile (latency) and showed different patterns of correlations with other brain regions in the ventral pathway (connectivity). These findings are consistent with any theory of perception in which information processing in the sensory neocortex relies upon 2 distinct types of signals, some increasing by repetitions and some decreasing with repetitions. After discussing our different results, we will describe one such framework called predictive coding (Mumford 1992; Friston 2005) which does rely upon 2 types of signals carrying predictions or representations (that would be enhanced by repetition) and prediction errors (suppressed by repetition), and which we would argue can account for all of the present findings.

In the FFA, sign-consistency analyses revealed that RS and RE voxels tended to remain the same across measurements. Over and above this consistent segregation in sign, the pattern-consistency analyses showed that the response profiles within each of these RS and RE populations were both consistent across measurements. Simulation analyses ruled out the possibility that these stability results depended on variability in face selectivity, rather than variability in repetition sensitivity. All face stimuli were trial unique in our experiment, so this result cannot be driven by bottom-up responses to specific stimulus exemplars. Moreover, we avoided using any normalization or proportional scaling across runs in our imaging analyses, to ensure run-run correlations in RS and RE voxels were assessed in a completely independent fashion. This finding is consistent with a growing body of reports of RE responses in the recent literature—for example, that unfamiliar faces prompt primarily RE (Henson et al. 2000; Muller et al. 2012), or that RE is observed for low visibility stimuli (Turk-Browne et al. 2006). Moreover, they show that in the FFA 2 types of response are elicited at the same time by the same face stimulus.

As a control region, we focused on the responses in the amygdala, which also exhibited a preference for face stimuli (in the localizer data) and an overall RS response (in the univariate analysis on smoothed data). However, no evidence for a functional segregation between RS and RE voxels in the amygdala was observed in any of the analyses conducted (stability, connectivity, and latency). This demonstrates that

the segregation of RS and RE voxels was not a trivial consequence of the global properties of regions showing face preference and repetition effects. Rather, within these regions, it was specific to the fusiform gyrus.

Outside of these face-sensitive regions, our sign-consistency analysis revealed that the RS and RE voxels could be segregated also at more posterior occipital sites (e.g. BA 18 and 19), suggesting that early visual regions may also implement this dissociation between 2 types of computational units. In the primary visual cortex (BA 17), however, we found no evidence for this dissociation, but neither did we find RS for faces. Whilst adaptation effects for low-level stimuli in V1 are well described (e.g. Weigelt et al. 2012), it is possible that the face stimuli used in our study may not have been the most efficient for eliciting adaptation in the primary visual cortex.

In the FFA, voxels showing RS and enhancement responses were also dissociated in terms of their relative connectivity with lower and higher visual regions, with RS voxels but not RE voxels correlating preferentially with lower visual regions. This result suggests that RS voxels receive more bottom-up information from early occipital regions extracting low-level visual properties from the stimulus than top-down information from higher visual regions that encode face-related information (Kriegeskorte et al. 2007). The RE voxels, by contrast, may be receiving as much information from higher visual regions as from lower ones, which, might help to achieve a better individuation of the current face as repetitions allow this information to accumulate. Our connectivity analyses used a FIR model to remove evoked activity, a rigorous approach that guards against misfitting of canonical basis functions to our data, or generally a systematic influence of the main response of the ROIs tested (Summerfield et al. 2006; Norman-Haignere et al. 2012). Besides, our control analysis in the amygdala rules out the possibility that 2 regions showing the same response profile (e.g. RS voxels and the lower visual region) shall automatically remain correlated even after removal of their response profile.

Our final result is that RS and RE voxels also differ in terms of their response latency, which offers further evidence that they make distinct contributions to visual computation. In particular, RS voxels respond with reduced latency and show an earlier effect of repetition in comparison with RE voxels. However, we interpret this result with some caution, for 2 reasons. First, the temporal resolution of fMRI is poor and the link to the underlying latency of neuronal processes is questionable. Secondly, it is unclear whether the latency effect is about the linear contrast per se or simply about the response to the first face only. Further work is needed to clarify this issue.

Within a face-sensitive fusiform region, we could dissociate 2 sub-populations of voxels that showed different response profiles to face repetitions. Here, we show that, when looking closely at the responses of individual voxels without spatial smoothing, one could detect sparse but consistent RE responses that could have been previously “masked” by the dominant RS response in their local neighborhood. How are these sub-populations organized? In further analyses of the current dataset (Supplementary Analyses and Fig. S3), we found that the 2 sub-populations of RS and RE voxels seemed to show some consistent spatial organization rather than being homogeneously intermingled. However, the spatial resolution of fMRI is limited in the current study that used 3×3 mm voxels, and a better characterization of these local

spatial relations might be best achieved using high-field fMRI that allows for a better spatial resolution, or using neurophysiological tools.

Ongoing debate surrounds the nature of the neural computations underlying visual perception. The classic bottom-up theories whereby neurons act as sensors tuned to features, shapes, and objects (Hubel and Wiesel 1959; Tanaka 1996) have been challenged by evidence that visual neurons are sensitive to contextual influences lying beyond their receptive field (Allman et al. 1985; Angelucci and Bullier 2003; Smith and Muckli 2010), and their activity is strongly modulated by past stimulation history, even over long lags (van Turennout et al. 2000; Vuilleumier et al. 2002; Kouider et al. 2009). These findings have prompted the theory that visual computation depends both on bottom-up sensory input and on contextual signals that bias perception toward a particular interpretation of the visual world (Mumford 1992; Rao and Ballard 1999; Bar 2004; Friston 2005; Gilbert and Sigman 2007). For example, “predictive coding” proposes that visual processing depends on the interplay between top-down expectation (or representation) signals and bottom-up surprise signals, which would be processed by distinct units at each stage of the cortical hierarchy (Friston 2005). How does this framework account for our findings?

The predictive coding framework naturally accounts for the dissociation between 2 types of visual units, prediction units and surprise units, which can be mapped respectively on the RE and RS voxels that we have dissociated here. Because stimuli that occur repeatedly come to be expected, in the predictive coding framework repetition should elicit greater expectations, and enhanced activity in expectation units, that is, a RE response. On the other hand, a stimulus that conforms to expectations elicits less surprise, so repetitions should also elicit a reduced response (i.e. RS) in units encoding the surprise signals. Recent computational modeling estimated that prediction error responses outweigh prediction signals by a factor of 2:1 (Egner et al. 2010). Although the reasons for the weaker contribution of RE signals to visual responses are still unclear, it is interesting to note that this proportion is remarkably similar to our present observation that only one-third of voxels showed RE. Interestingly, predictive coding can also account for our connectivity and latency results. Indeed, we found RS voxels to be preferentially connected with lower visual regions and to respond faster, while RE voxels responded later and were relatively more connected with higher visual regions, consistent with the idea that neural computations in the FFA gradually reconcile surprise (RS) signals flowing forward with prediction (RE) signals coming from higher regions. However, predictive coding also argues that the reconciliation of prediction and prediction error signals is an iterative process that occurs on a finer time-scale than is measurable with fMRI, so this interpretation remains speculative.

To sum up, we have shown that distinct populations of voxels in the FFA exhibit consistent suppressed and enhanced responses to repeated faces. While incommensurate with bottom-up accounts of neural repetition effects (e.g. presented in Grill-Spector et al. (2006)), our 3 results (stability, connectivity, and latency) are all consistent with the predictive coding view of visual perception (Mumford 1992; Friston 2005). Our findings thus contribute to a growing literature that supports predictive coding as a model of perception (Murray

et al. 2002; Summerfield et al. 2008, 2011; den Ouden et al. 2009; Alink et al. 2010; Egner et al. 2010; Meyer and Olson 2011; Kovacs et al. 2012) and supports the view that visual processing, like its counterpart in the dopaminergic reward system (Schultz et al. 1997; Schultz and Dickinson 2000), depends on the interplay between of prediction and prediction error signals (Ullman 1995; Deco and Rolls 2005; Friston 2005; Spratling 2008). This suggests that predictions and prediction errors may form part of a general computational mechanism that is employed across neocortical and subcortical regions alike (Rushworth et al. 2009).

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>.

Funding

This study was supported by a Wellcome Trust Grant to C.S. and T.E. Funding to pay the Open Access publication charges for this article was provided by the Wellcome Trust.

Notes

Conflict of Interest: None declared.

References

Alink A, Schwiedrzik CM, Kohler A, Singer W, Muckli L. 2010. Stimulus predictability reduces responses in primary visual cortex. *J Neurosci.* 30:2960–2966.

Allman J, Miezin F, McGuinness E. 1985. Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annu Rev Neurosci.* 8:407–430.

Angelucci A, Bullier J. 2003. Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? *J Physiol Paris.* 97:141–154.

Bar M. 2004. Visual objects in context. *Nat Rev Neurosci.* 5:617–629.

Deco G, Rolls ET. 2005. Attention, short-term memory, and action selection: a unifying theory. *Prog Neurobiol.* 76:236–256.

Dehaene S, Naccache L, Cohen L, Bihan DL, Mangin JF, Poline JB, Riviere D. 2001. Cerebral mechanisms of word masking and unconscious repetition priming. *Nature Neurosci.* 4:752–758.

den Ouden HEM, Friston KJ, Daw ND, McIntosh AR, Stephan KE. 2009. A dual role for prediction error in associative learning. *Cereb Cortex.* 19:1175–1185.

Egner T, Monti JM, Summerfield C. 2010. Expectation and surprise determine neural population responses in the ventral visual stream. *J Neurosci.* 30:16601–16608.

Friston K. 2005. A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci.* 360:815–836.

Gilbert CD, Sigman M. 2007. Brain states: top-down influences in sensory processing. *Neuron.* 54:677–696.

Grill-Spector K, Henson R, Martin A. 2006. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci.* 10:14–23.

Henson R, Shallice T, Dolan R. 2000. Neuroimaging evidence for dissociable forms of repetition priming. *Science.* 287:1269–1272.

Henson RN, Shallice T, Gorno-Tempini ML, Dolan RJ. 2002. Face repetition effects in implicit and explicit memory tests as measured by fMRI. *Cereb Cortex.* 12:178–186.

Hubel DH, Wiesel TN. 1959. Receptive fields of single neurones in the cat's striate cortex. *J Physiol.* 148:574–591.

Kouider S, de Gardelle V, Dehaene S, Dupoux E, Pallier C. 2010. Cerebral bases of subliminal speech priming. *Neuroimage.* 49:922–929.

Kouider S, Eger E, Dolan R, Henson RN. 2009. Activity in face-responsive brain regions is modulated by invisible, attended faces: evidence from masked priming. *Cereb Cortex.* 19:13–23.

Kovacs G, Iffland L, Vidnyanszky Z, Greenlee MW. 2012. Stimulus repetition probability effects on repetition suppression are position invariant for faces. *Neuroimage.* 60:2128–2135.

Kriegeskorte N, Formisano E, Sorger B, Goebel R. 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc Natl Acad Sci USA.* 104:20600–20605.

Meyer T, Olson CR. 2011. Statistical learning of visual transitions in monkey inferotemporal cortex. *Proc Natl Acad Sci USA.* 108:19401–19406.

Muller NG, Strumpf H, Scholz M, Baier B, Melloni L. 2012. Repetition suppression versus enhancement—it's quantity that matters. *Cereb Cortex.* In press.

Mumford D. 1992. On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol Cybern.* 66:241–251.

Murray SO, Kersten D, Olshausen BA, Schrater P, Woods DL. 2002. Shape perception reduces activity in human primary visual cortex. *Proc Natl Acad Sci USA.* 99:15164–15169.

Naccache L, Dehaene S. 2001. The priming method: imaging unconscious repetition reveals an abstract representation of number in the parietal lobes. *Cereb Cortex.* 11:966–974.

Norman-Haignere SV, McCarthy G, Chun MM, Turk-Browne NB. 2012. Category-selective background connectivity in ventral visual cortex. *Cereb Cortex.* 22:391–402.

Rao RP, Ballard DH. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neurosci.* 2:79–87.

Rotshtein P, Henson RN, Treves A, Driver J, Dolan RJ. 2005. Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nature Neurosci.* 8:107–113.

Rushworth MF, Mars RB, Summerfield C. 2009. General mechanisms for making decisions? *Curr Opin Neurobiol.* 19:75–83.

Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science.* 275:1593–1599.

Schultz W, Dickinson A. 2000. Neuronal coding of prediction errors. *Annu Rev Neurosci.* 23:473–500.

Smith FW, Muckli L. 2010. Nonstimulated early visual areas carry information about surrounding context. *Proc Natl Acad Sci USA.* 107:20099–20103.

Spratling MW. 2008. Reconciling predictive coding and biased competition models of cortical function. *Front Comput Neurosci.* 2:4.

Summerfield C, Greene M, Wager T, Egner T, Hirsch J, Mangels J. 2006. Neocortical connectivity during episodic memory formation. *PLoS Biol.* 4:e128. doi:10.1371/journal.pbio.0040128.

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

Summerfield C, Wyart V, Johnen VM, de Gardelle V. 2011. Human scalp electroencephalography reveals that repetition suppression varies with expectation. *Front Hum Neurosci.* 5:67.

Tanaka K. 1996. Inferotemporal cortex and object vision. *Annu Rev Neurosci.* 19:109–139.

Turk-Browne NB, Yi DJ, Chun MM. 2006. Linking implicit and explicit memory: common encoding factors and shared representations. *Neuron.* 49:917–927.

Turk-Browne NB, Yi DJ, Leber AB, Chun MM. 2007. Visual quality determines the direction of neural repetition effects. *Cereb Cortex.* 17:425–433.

Ullman S. 1995. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb Cortex.* 5:1–11.

van Turennout M, Ellmore T, Martin A. 2000. Long-lasting cortical plasticity in the object naming system. *Nature Neurosci.* 3:1329–1334.

Vuilleumier P, Henson RN, Driver J, Dolan RJ. 2002. Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neurosci.* 5:491–499.

Weigelt S, Limbach K, Singer W, Kohler A. 2012. Orientation-selective functional magnetic resonance imaging adaptation in primary visual cortex revisited. *Hum Brain Mapp.* 33:707–714.