NeuroImage 180 (2018) 280–290

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Cortical feedback signals generalise across different spatial frequencies of feedforward inputs

Yulia Revina [1], Lucy S. Petro, Lars Muckli *

Centre for Cognitive Neuroimaging, Institute of Neuroscience & Psychology, College of Medical, Veterinary and Life Sciences, University of Glasgow, G12 8QB, UK

ARTICLE INFO

Keywords:
fMRI
Spatial frequency
Cortical feedback
V1

ABSTRACT

Visual processing in cortex relies on feedback projections contextualising feedforward information flow. Primary visual cortex (V1) has small receptive fields and processes feedforward information at a fine-grained spatial scale, whereas higher visual areas have larger, spatially invariant receptive fields. Therefore, feedback could provide coarse information about the global scene structure or alternatively recover fine-grained structure by targeting small receptive fields in V1. We tested if feedback signals generalise across different spatial frequencies of feedforward inputs, or if they are tuned to the spatial scale of the visual scene. Using a partial occlusion paradigm, functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA) we investigated whether feedback to V1 contains coarse or fine-grained information by manipulating the spatial frequency of the scene surround outside an occluded image portion. We show that feedback transmits both coarse and fine-grained information as it carries information about both low (LSF) and high spatial frequencies (HSF). Further, feedback signals containing LSF information are similar to feedback signals containing HSF information, even without a large overlap in spatial frequency bands of the HSF and LSF scenes. Lastly, we found that feedback carries similar information about the spatial frequency band across different scenes. We conclude that cortical feedback signals contain information which generalises across different spatial frequencies of feedforward inputs.

1. Introduction

Visual processing in cortex relies on ascending feedforward projections, descending feedback projections and intra-areal signals. Top-down feedback signals from higher cortical areas contextualise input from the bottom-up feedforward stream. Despite the importance of feedback for visual perception, we have much to discover about the contextual surrounding scene information that is transmitted to neurons in early visual areas such as primary visual cortex (V1). Primary visual cortex has small receptive fields and therefore has the capacity to process images at a fine-grained spatial scale. Higher visual areas have larger receptive fields which are spatially invariant. Therefore, feedback from higher visual areas could provide coarse, spatially invariant information about the global scene structure. Alternatively, feedback to V1 could contain fine-grained structure by targeting the region's small receptive fields.

Feedforward visual input is decomposed into different spatial frequency (SF) bands, with low spatial frequencies (LSFs) conveying coarse information, and high spatial frequencies (HSFs) providing the fine-grained details (e.g. De Valois et al., 1982; Wilson and Bergen, 1979). However, the spatial scale of feedback signals is less well understood. Do feedback signals generalise across different spatial frequencies of feedforward inputs, or are they tightly tuned to the spatial scale of the specific visual scene?

An important source of contextual information is global scene representation (Bar, 2004; Oliva and Torralba, 2007). Therefore, feedback may use coarse information carried via LSF. Coarse signals can carry information about the overall structure of the scene, and this rough draft can help to segment the scene and boost the recognition process of scene categories or the identification of objects within the scene. Thus it is plausible that feedback carries information about global scene structure using LSF information (Bar et al., 2006; Oliva and Torralba, 2006). Consistent with the hypothesis that LSF information contributes to top-down signals, studies have shown that LSF information is processed faster than HSF, and can influence perception before fine-grained information is processed (e.g. Bar, 2003; Bar et al., 2006; Breitmeyer, 2014;

Abbreviations: FB, Feedback; FF, Feedforward; HSF, High spatial frequency; LSF, Low spatial frequency.

* Corresponding author. 58 Hillhead Street, Glasgow, G12 8QB, UK.
E-mail addresses: yrevina@ntu.edu.sg (Y. Revina), lucy.petro@glasgow.ac.uk (L.S. Petro), Lars.Muckli@glasgow.ac.uk (L. Muckli).
1 Present address: Division of Psychology, 14 Nanyang Drive, Nanyang Technological University, 637332, Singapore.

https://doi.org/10.1016/j.neuroimage.2017.09.047
Received 15 March 2017; Received in revised form 1 September 2017; Accepted 21 September 2017
Available online 22 September 2017

Kveraga et al., 2007). Furthermore, top-down feedback originates from neurons with larger receptive fields and diverges to many finely tuned receptive fields of neurons closer to the entrance level of visual processing. Such a coarse-to-fine projection might carry only information captured in the LSF content.

An alternative hypothesis is that feedback to early visual areas recovers HSF precision. The coarse, abstract contextual message of higher areas is translated to fine-grained feedback signals within early visual areas. The Reverse Hierarchy Theory (Ahissar and Hochstein, 2004; Hochstein and Ahissar, 2002) suggests that feedback to lower areas is necessary to gain precision in perception. Initially a coarse gist of the scene is processed, but during difficult perceptual tasks top-down processes enable the access of fine-grained low-level features. These "vision with scrutiny" processes focus attention to specific low-level units. In this way, feedback connections add details to our initial perception of a scene. This suggests that feedback has access to the most fine-grained information in the lowest areas, as these have the smallest receptive fields coding for basic features. Therefore, despite originating from large receptive fields, feedback may contain fine-grained information at the level of V1. In addition, global scene structure may also be coded in HSFs (Oliva and Torralba, 2006; Walther et al., 2011). Feedback contains information about individual scenes as well as category, suggesting some fine-grained information may remain in feedback (Morgan et al., 2016).

Due to the retinotopic nature of V1, it is possible to investigate voxel information patterns specific to regions of a visual image. We can study feedback signals by isolating them from the feedforward inputs by partially occluding the image. The V1 representation of occluded regions of the visual field contains contextual information (Ban et al., 2013; Shushruth, 2011; Smith and Muckli, 2010; Sugita, 1999) fed back to the superficial layers (Muckli et al., 2015). Using a partial occlusion paradigm (Smith and Muckli, 2010), functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA, see Mur et al., 2009; Norman et al., 2006) we investigated whether feedback to V1 contains coarse or fine-grained information by manipulating the spatial frequency of the scene surround outside the occluded region.

2. Materials and methods

2.1. Subjects

Thirty five subjects from the University of Glasgow participated in the experiment (17 males; mean age: 24.63 years, range: 17–42 years). We paid subjects for participation. Subjects provided informed written consent and the local ethics committee at the University of Glasgow approved the experiment (#CSE01063). We excluded one subject due to uncorrectable motion artefacts (>1.5 mm) and another subject due to missing visual cortex activation in one run (below threshold). Therefore, we report results from 33 subjects (16 males; mean age: 24.48 years, range: 17–42 years).

2.2. Stimuli

2.2.1. Feedback vs feedforward condition

To test feedback signals in the absence of feedforward stimulation, we used an occlusion paradigm previously employed by Smith and Muckli (2010), Muckli et al. (2015) and Morgan et al. (2016). For the feedback conditions, we occluded the lower right image quadrant with a white rectangle. The white rectangle spanned 11.6° x 9.2° and we placed it 0.5° of visual angle diagonally from the centre of the scene. For ten subjects who saw a different set of images, the occluded region spanned 15.6° × 11.6° visual angle. In the feedforward conditions, we presented only the corresponding quadrant that was occluded in the feedback conditions (Fig. 1).

Voxels in the occluded region do not receive informative, scene-specific feedforward input (Fig. 1E). However, large surround receptive fields of these voxels capture feedforward information from the stimulated surround (Smith and Muckli, 2010; Muckli et al., 2015). Thus, voxels in the occluded region receive information fed back from higher areas; we refer to this condition as 'feedback'. In the stimulated quadrant, voxels receive direct feedforward input and feedback triggered by information in the quadrant; for simplicity we refer to this condition as 'feedforward'.

2.2.2. Scenes

We aimed to induce contextual associations, which are particularly strong in natural scenes (Bar, 2004). We used two natural scene images for each participant and varied these across the different experiments to generalise our results across different images (Fig. 1A). Scenes were 600 × 480 px, which corresponded to 24° × 19.2° visual angle. For five subjects in the Small SF Overlap group (two in 0.65/1.30 cpd; three in 0.81/1.62 cpd) and for five subjects in the Large SF Overlap group (two in 0.81/2.03 cpd; three in 0.97/2.43 cpd), different scenes were used (a classical concert scene and a New York street scene, Fig. 1A, bottom section). These scenes spanned 800 × 600 px, which corresponded to 32° × 24° visual angle. Each scene was filtered to create a high spatial frequency (HSF) and a low spatial frequency (LSF) version (Fig. 1A). In creating these HSF and LSF scenes, we explored a variety of HSF and LSF cut-offs. Therefore, some participants viewed HSF and LSF scenes which shared, to various extents, a subset of SFs, whilst others viewed scenes not sharing any SFs (see Table 1 for SF cut-offs and number of subjects presented with each combination). We did this to investigate how specific the feedback signals are to the SF band of the surround. If a large amount of SF information needs to be shared between HSF and LSF version of the scene for HSF and LSF feedback to be similar, this would suggest that feedback is tightly tuned to the SF band of the surrounding scene.

Each group of subjects viewed HSF and LSF scenes with either a *Small Overlap* in their SF bands, a *Large Overlap* or no shared SF information (*Gap* and *No Overlap* groups). We chose these particular cut-offs as previous studies have indicated V1 preference for SFs of around 0.68–2 cycles per degree of visual angle (cpd, Haynes and Rees, 2005; Henriksson et al., 2008; Tong et al., 2012). In the *Small Overlap* groups, there was an overlap of 1 octave in the spatial frequency bands. For two *Large Overlap* groups there was an overlap of 1.3 octaves. We tried several cut-offs so as not to restrict our results to one particular SF cut-off, but rather a particular overlap ratio. The *Gap* group had a "gap" of 1 octave not shared by either stimulus. The *No Overlap* group did not have any shared SFs nor a gap. Each subject only saw one HSF-LSF cut-off pair. Fig. 1A shows stimuli for the 0.81/1.62 cpd *Small Overlap* and the 0.97/2.43 *Large Overlap* groups. Inline Supplementary Fig. S1 shows the amplitude spectra of each stimulus.

2.2.3. Occluded region mapping

Twice per run, subjects viewed three contrast-reversing checkerboards (5 Hz). The checkerboards either covered the inner rectangular part of the occluded region (*Target* mapping – approximately 1.5° diagonally from fixation, 10.1° × 7.7° visual angle or 14.9° × 10.9° for the ten subjects who saw the larger scenes) or the border between the lower right quadrant and the rest of the stimulus (*Surround* mapping). There were two types of surround stimuli – *Near Surround* (approximately 0.5° diagonally from fixation, 11.6° × 9.2° visual angle or 15.6° × 11.6°) and *Inside Border* (approximately 1.5° diagonally from fixation, 11.6° × 9.2° visual angle, or 2.5° from fixation and 14.2° × 10.2° visual angle for ten subjects) (Fig. 1B and C).

2.3. Task and procedure

Subjects were familiarised with the unfiltered non-occluded images with a short practice run (non-filtered non-occluded stimuli shown in a random order 10 times each with a duration of 1 s) prior to entering the scanner. For visual stimulation, we used MRI compatible goggles

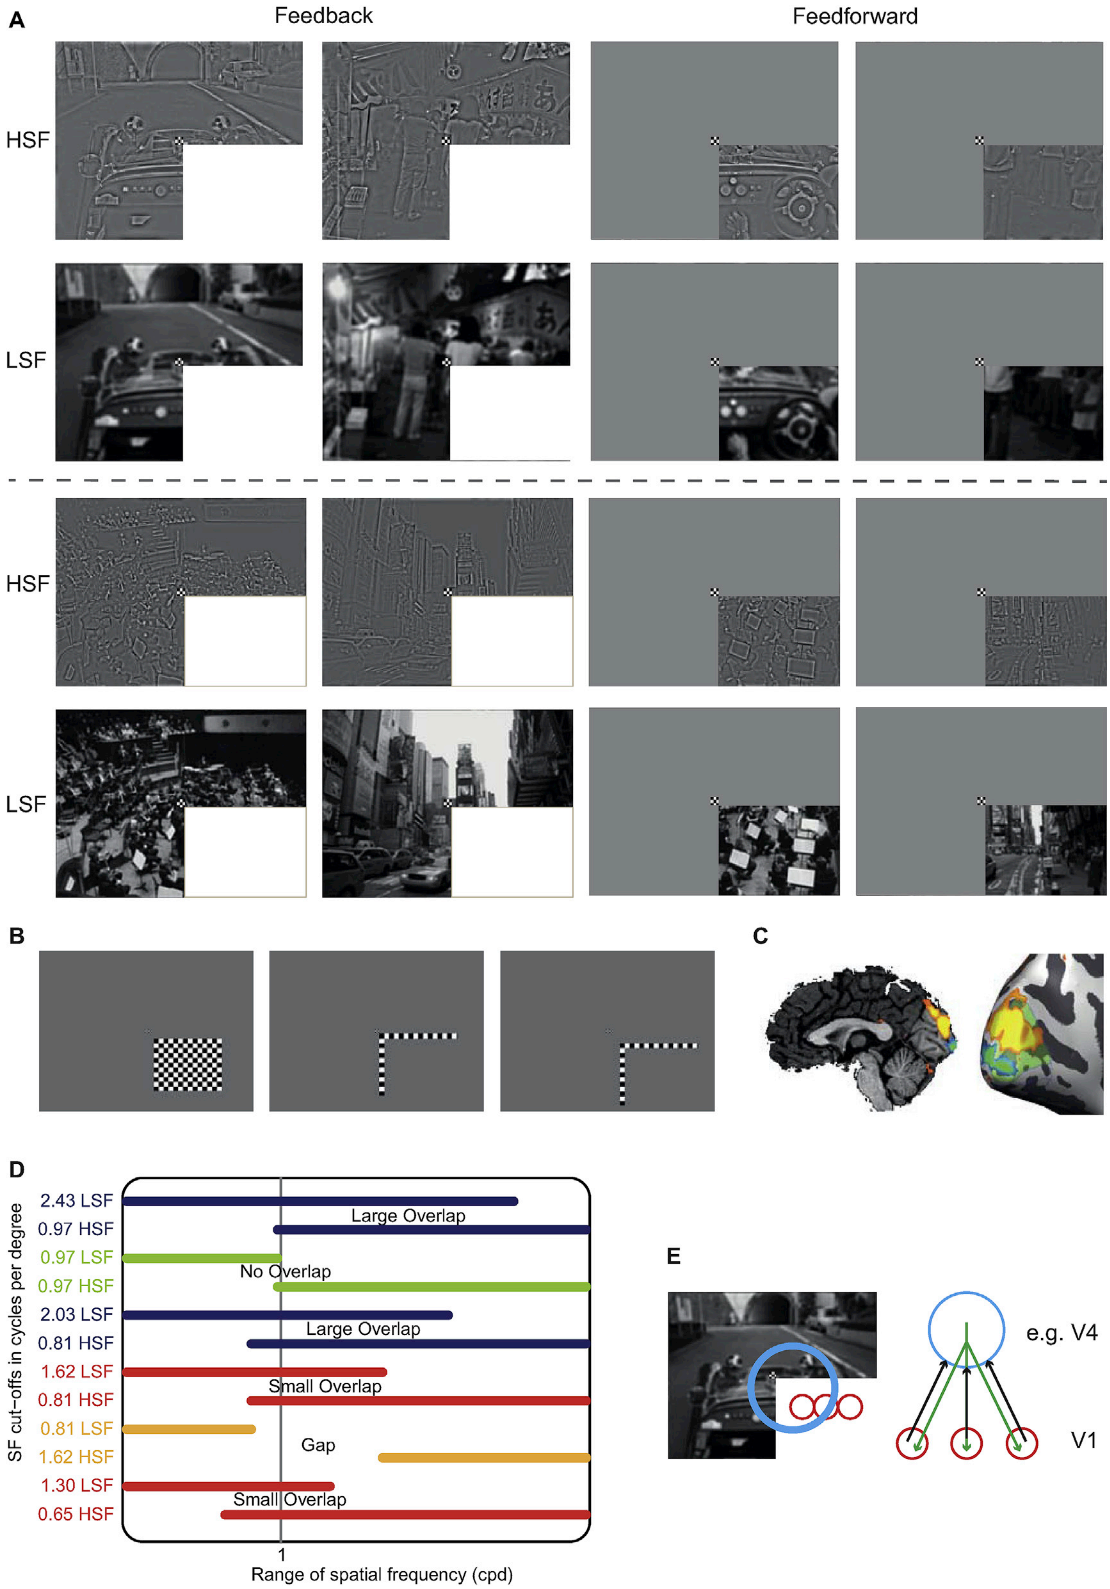

Fig. 1. Example stimuli, shown here for the *Small Overlap* and *Large Overlap* pairs. In feedback conditions the lower right image quadrant was occluded with a white rectangle, while in feedforward conditions only the corresponding quadrant was visible. A) Each scene was filtered to produce a high spatial frequency (HSF) or a low spatial frequency (LSF) version. The example images represent the *Small Overlap* condition (HSF scene filtered at 0.81 cpd and LSF scene filtered at 1.62 cpd, top section, car and people scenes) and *Large Overlap* condition (HSF scene filtered at 0.97 cpd and LSF scene filtered at 2.43 cpd, bottom section, New York and concert scenes). B) Checkerboards were used to retinotopically map the occluded region in V1: *Target* (left), *Near Surround* (middle) and *Inside Border* (right). C) The contrast map shows the retinotopic representation of the occluded quadrant (*Target* – *Near Surround*) in the occipital cortex, with V1 shaded in green on the inflated visualization. D) Graphic representation of the ranges of spatial frequencies contained in each Overlap group. See also Table 1. E) Diagrammatic representation of the receptive fields in the occluded region and the origin of feedback and feedforward signals. Feedforward signals (black arrows) from V1 neurons (red) reach higher visual areas with larger receptive fields (blue). Feedback signals (green arrows) send information back to V1, thus being able to inform V1 neurons in the occluded region about information in the stimulated surround. The example V1 cells have uniform white space as feedforward information; however cells in higher areas (e.g. V4), which also receive information from the non-occluded part of the image, can feed back contextual information to these V1 cells.

Table 1
Spatial frequency (SF) cut-offs used for each pair of stimuli and the SF overlap groups these corresponded to. Values are expressed in cycles per degree (cpd). See also Fig. 1D.

	High pass filter cut-off	Low pass filter cut-off	SF Overlap
Pair 1 (n = 4)	1.62	0.81	Gap
Pair 2 (n = 6)	0.97	0.97	No overlap
Pair 3 (n = 2)	0.65	1.30	Small
Pair 4 (n = 10)	0.81	1.62	Small
Pair 5 (n = 2)	0.81	2.03	Large
Pair 6 (n = 9)	0.97	2.43	Large

(NordicNeuroLab) with 800 × 600 px screen resolution, which corresponded to 32° × 24° visual angle. The stimulus was a natural visual scene presented on a grey background. For each subject there were 8 types of trial (2 scenes, high or low SF, occluded [feedback, FB] or non-occluded [feedforward, FF]). In each 12 s trial the stimulus was flashed on and off (200 ms on/200 ms off) 30 times. For ten subjects who saw the larger images in each 12 s trial the stimulus flashed on and off (200 ms on/200 ms off) 28 times (11.6 s + variable fixation to make each trial a total of 12 s to account in uncertainty in timing). This flashing procedure increases the signal to noise ratio relative to continuous presentation due to more repetitions of stimulation (Kay et al., 2008) and leads to a larger BOLD response (Boynton et al., 1996). Each different type of trial was presented sequentially, with the trial order randomized in each sequence of 8 trials. Each sequence lasted 96 s (8 × 12 s). A 12 s fixation period was included before and after each sequence of trials. Each experimental run lasted 10 min 48 s, consisting of four trial sequences and two mapping sequences (each sequence consisted of *Target* and the two *Surrounds*). There were four experimental runs in total. The subjects' task was to fixate on a central checkerboard and report a fixation colour change with a button press. Subjects pressed a different button, depending on whether the colour change occurred during scene 1 or scene 2. The purposes of the task were to ensure that the subject attended the stimuli and to minimize eye movements.

After the experimental runs, we performed a retinotopic mapping procedure to allow us to estimate the borders of the early visual areas V1-3. The mapping procedure was a standard polar-angle protocol consisting of either wedge shaped checkerboards arranged in a "bow-tie" or a single wedge which started in the right horizontal meridian and rotated clockwise (12 rotations per scan, wedge angle: 22.5°, scan time: 13 min 28 s for the single wedge and 7 min 4 s for the bow-tie).

2.4. MRI acquisition

We collected MRI data using a 3 T Siemens Tim Trio System with a 12-channel head coil. We measured blood oxygen level dependent (BOLD) signals with an echo-planar imaging sequence (echo time: 30 ms, repetition time: 1000 ms, field of view 210 mm, flip angle: 62°, 18 axial slices). The spatial resolution for functional data was 3 × 3 × 3 mm. Each experimental run had 648 volumes. Retinotopic mapping consisted of 424 (bow-tie) or 808 volumes (single wedge). We positioned the 18 slices to maximize coverage of occipital cortex. We also recorded a high resolution 3D anatomical scan (3D Magnetization Prepared Rapid Gradient Echo, 1 × 1 × 1 mm resolution, 192 volumes).

2.5. MRI data processing

We corrected functional data for the experimental runs and retinotopic mapping runs for slice time (cubic spline interpolation) and 3D motion (Trilinear/Sinc interpolation), temporally filtered (high-pass filtered at 6 cycles with GLM-Fourier and linearly detrended), and spatially normalized data into Talairach space with BrainVoyager QX 2.8 (Brain Innovation). Subsequently, we used the anatomical data to create an inflated cortical surface and functional data were overlaid.

2.6. Voxel selection and analysis

Excessive subject movement between runs is likely to affect correspondence between voxels from one run to another. This movement could affect our analysis because we selected our region of interest (ROI) based on the averaged functional data of all 4 runs. To determine whether there was good alignment between functional data covering the visual areas, we calculated an alignment value for each subject by measuring Pearson's correlation in a ROI in the visual cortex between the four functional runs. This was a measure of anatomical alignment to ensure we did not lose power during statistical analyses due to poor alignment between the four functional runs that subjects performed. Correlations were performed in a ROI covering the early visual cortex using intensity values from an anatomical representation of the first volume of the functional data of every run. High correlations would suggest a close anatomical alignment between the 4 runs. The median alignment value across the subjects was 98% and single subject values ranged from 92% to 99%.

For our analysis, we selected voxels in V1 that corresponded to the occluded lower right image quadrant. We identified this non-feedforward stimulated region of V1 using our checkerboard mapping conditions; we used a general linear model (GLM) contrast of the *Target* against the *Near Surround*, as described previously in Smith and Muckli (2010). The ROI was selected from activation in V1 only. To further minimize spillover activity from neighbouring stimulated areas, voxels from the ROI were then selected for analysis based on the *Target* t-values being greater than 2, difference between *Target* and *Near Surround* t-value being greater than 2 and the mean t-value for feedback conditions being lower than the mean t-value for the feedforward conditions. This aimed to exclude voxels responding to the stimulation in the surround. Any voxels responding to the image surround would be more active in the feedback condition compared to the feedforward condition, as the surround in the feedforward condition was a grey background. In addition, we removed any subjects from the subsequent group analyses, who did not have above chance classifier performance for both of the feedforward conditions. Inability to decode the scenes in the stimulated control condition may be indicative of subjects not fixating properly, falling asleep, and so on. It would not be meaningful to assess feedback classifier performance (or lack of) in such cases. This further criterion removed 4 subjects (2 in Small and 2 in Large Overlap), and we therefore present results for 29 subjects in the subsequent *Results* section.

2.7. Controls for MVPA analysis

2.7.1. Analyses with extended boundary around the occluded region

We ran a separate analysis with a more stringent method of voxel selection to further ensure our findings of scene information in the quadrant were not due to spillover activity from the feedforward surround. We selected our region of interest in BrainVoyager using a GLM contrast of *Target* mapping being higher than both the *Near Surround* and the *Inside Border*. In addition, prior to classification analyses we selected voxels with t-values fitting the same criteria as described above. Two further subjects were removed from this analysis due to not having above threshold activation in V1 using the more conservative ROI definition, giving a total of 27 participants for this particular analysis.

2.8. Multivariate pattern classification analysis

We entered voxels matching the above-mentioned criteria into the linear classifier (Support Vector Machine [SVM], using the LIBSVM toolbox in MATLAB, Chang and Lin, 2001). We trained the classifier to decode the two scenes in each condition. For cross-classification analyses we trained the classifier on one experimental condition and tested on the other. The classifier used single-trial activity patterns (beta values) for training, and was then tested on either "single trial" (8 trials x 4 sequences = 32 separate trials) or "average block" activity patterns for each of the 8 trial types (average of the 4 repetitions). In other words, for the average block analysis, the training was the same (single trials of three runs, 32 trials in each run) but the testing was done on the average per stimulus condition (e.g. HSF Feedback) of the fourth run. For both types of analyses, we trained the classifier on 3 of the runs and tested on the remaining run (i.e. one-run-out cross-validation).

We used bootstrapping and permutation analysis to get a robust average value out of a small set of individual values, and to test how well the classifier would perform when the labels are randomly assigned. There were four classifier performance values for each condition for each subject, as we were able to train the classifier on four different folds — three runs of training and one run to test the classifier. We bootstrapped (1000 samples) the classifier performances for individual subjects, in order to estimate the single subject mean. We did not calculate confidence intervals (CIs) for this step. We then bootstrapped (1000 samples) these individual subject mean values to estimate the group means and confidence intervals (CIs) for each condition. The CIs covered the middle 95% of the distribution. Classifier performances were deemed to be significantly above chance (50%) if the 95% CIs did not contain chance-level performance. We computed differences between group classifier performances via a permutation test (1000 samples) of the differences between the group means (p values not corrected for multiple comparisons). We shuffled the observed values across conditions 1000 times, and calculated the absolute differences between the conditions. If the actual observed difference was in the top 5% of the differences distribution, then we deemed our conditions to be significantly different from each other.

3. Results

3.1. Both HSF and LSF scene surrounds induce meaningful feedback

First, we tested whether the spatial frequency filtered surround induced meaningful information in non-stimulated V1. We trained the Support Vector Machine (SVM) classifier to decode between the two scenes using voxel patterns responding to the lower right quadrant. We used SVM classification performance as an estimate of whether the quadrant contained informative feedback signals about the scene. We present classifier performance for the single trial analysis in the main manuscript text. We show corresponding classifier performances for the average block analysis in the Inline Supplementary Fig. S2–5.

Collapsing across the different SF cut-off groups, classifier performance for decoding between the two scenes was above chance for both HSF and LSF scenes, in both feedback and feedforward conditions for the single trial analysis (Fig. 2A; for average block analysis, see Inline Supplementary Fig. S2). Decoding during feedforward (FF) conditions was significantly higher than during feedback (FB) conditions for high spatial frequency (*FF HSF*: 78.77%, confidence interval [CI] [0.0463 0.0442] vs *FB HSF*: 57.97%, CI [0.0431 0.0442], $p < 0.001$; and low spatial frequency (*FF LSF*: 82.76%, CI [0.0442 0.0388] vs *FB LSF*: 59.49%, CI [0.0366 0.0399], $p < 0.001$) conditions. There was no difference in classifier performance between HSF and LSF conditions, for either the feedback ($p = 0.552$) or feedforward stimuli ($p = 0.225$). For these classification results presented separately for Concert/New York and Car/People stimuli, see Inline Supplementary Fig. S6.

Fig. 2B shows the single trial classification performance for decoding between the two scenes in each of the different cut-offs for HSF and LSF stimuli in the feedback conditions. Fig. 2C shows the same for the feedforward conditions. CIs are not shown for conditions with only one data point.

3.2. Similarity of feedback across HSF and LSF stimuli

Secondly, we tested whether the classifier can generalise over spatial frequencies, in other words, decode between the scenes even when it was trained on HSF and tested on LSF (and vice versa). We trained the classifier on either the HSF version of the two scenes and tested on LSF, or trained on LSF and tested on HSF. We performed this analysis for the different Overlap groups, to see how the amount of shared spatial frequency information between the HSF and LSF version of the scene would affect this generalisation. We predicted that if feedback is specific to the SF range of the scene surround then the larger the overlap the better the classifier would perform since there would be more shared information between HSF and LSF. Alternatively, if feedback is similar across different SF surrounds, then we predicted we would see similar levels of generalisation across all Overlap groups.

3.2.1. Training on HSF and testing on LSF

In feedforward conditions, classifier performance was above chance for all Overlaps (Fig. 3A, right, for average block analysis, see Inline Supplementary Fig. S3). For feedback conditions, classifier performance was above chance only for the *Large Overlap* group (Fig. 3A, left; *Gap*: 52.34%, CI [0.0781, 0.0547], *No Overlap*: 54.17%, CI [0.0625 0.0677], *Small Overlap*: 54.06%, CI [0.0406 0.0406] and *Large Overlap*: 58.33%, CI [0.0764 0.0799]).

We had fewer subjects in the *Gap* and *No Overlap* conditions than in the *Small* and *Large Overlap* conditions. To increase statistical power in the *Gap* and *No Overlap* conditions, we combined these two groups of subjects and re-ran our classification analyses. This grouping is conceptually motivated by the hypothesis that common spatial frequency information is required for successful generalisation across spatial frequency (i.e. significant cross-classification), and neither of these stimuli had shared spatial frequency information; therefore we can consider them together. In this new group of "*Gap/No Overlap*", we found that cross-classification from HSF to LSF was still at chance level (53.44%, CI [0.0500 0.0500]) and therefore we do not think our result was due to a lack of power in those groups. We keep the *Gap* and *No Overlap* conditions separate for the remainder of the analyses.

3.2.2. Training on LSF and testing on HSF

For feedforward conditions, classifier performance was above chance for *Gap*, *Small Overlap* and *Large Overlap* (Fig. 3B, right). For feedback conditions (Fig. 3B, left), classifier performance was above chance for all the Overlap groups (*Gap*: 58.59%, CI [0.0391 0.0313], *No Overlap*: 56.77%, CI [0.0521 0.0521], *Small Overlap*: 53.44%, CI [0.0313 0.0313], *Large Overlap*: 57.29%, CI [0.0486 0.0590]).

The results show that the classifier can generalise over spatial frequencies. We can train on one spatial frequency and decode the scenes presented in another spatial frequency. We see that a lot of shared information in terms of spatial frequency is not necessary for cross-classification, as we could cross-classify in the *Gap*, *No Overlap* and *Small* conditions. However, *Large Overlap* is the only condition that was above chance for both directions of cross-classification (HSF to LSF, and LSF to HSF), which may suggest that the classifier generalises better across SF in the presence of more shared information. However, the fact that cross-classification works in the other conditions (*Gap*, *No Overlap* and *Small Overlap* for LSF to HSF cross-classification), suggests generalisation can still occur without a large amount of shared information.

3.3. Does feedback carry spatial frequency information not related to a specific scene?

Next, we were interested in what information spatial frequency related feedback is transmitting. Is the information specific to the scene in question or does feedback carry spatial frequency information that is similar across different scenes with the same spatial frequency content? In other words, does feedback carry any information about the spatial frequency of the surround that is not specific to a particular scene? We trained the classifier to decode between HSF vs LSF on Scene 1 and tested whether it could decode between HSF vs LSF on Scene 2 (and vice versa). Classifier performance was above chance for both feedback and feedforward conditions, and for both directions of cross-classification (Fig. 4,

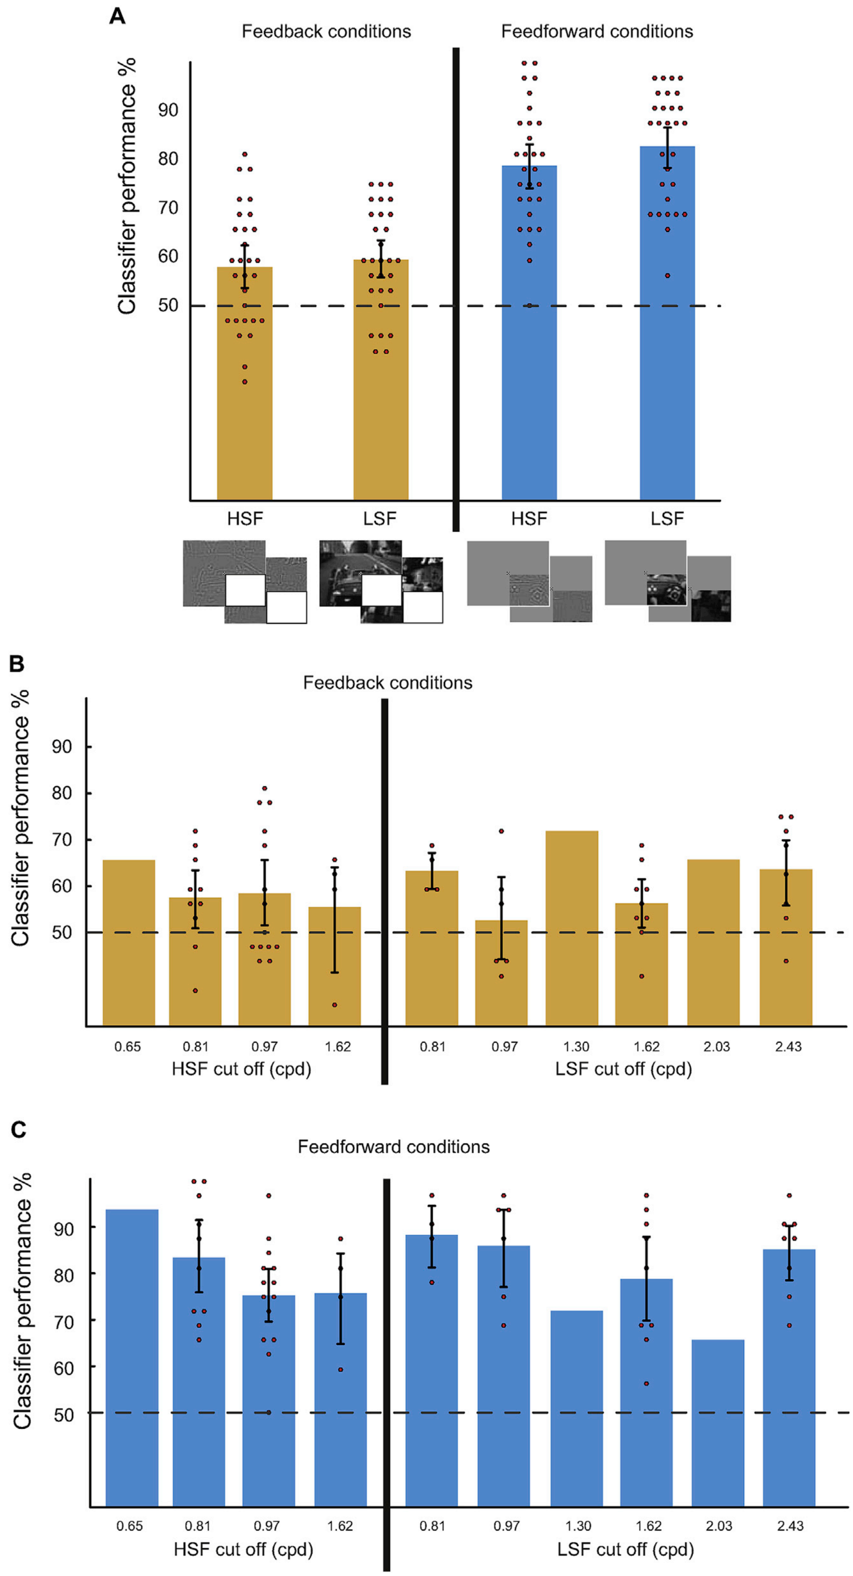

Fig. 2. Classification performance for decoding the two scenes in HSF and LSF conditions, for feedback and feedforward stimuli. Chance level is 50%. Lines represent 95% confidence intervals around the bootstrapped mean (1000 bootstrap samples of individual subjects' performances). Classifier performance is significantly above chance at $\alpha = 0.05$ (not corrected for multiple comparisons) if the confidence intervals do not intersect with the chance line. The small red circles represent individual subjects' results. A) Classifier performance for HSF and LSF conditions, collapsed over different SF cut-offs. N = 29. Images are example stimuli used for a subset of subjects. B) Classifier performance split by different SF cut-offs for feedback conditions. C) Classifier performance split by different SF cut-offs for feedforward conditions.

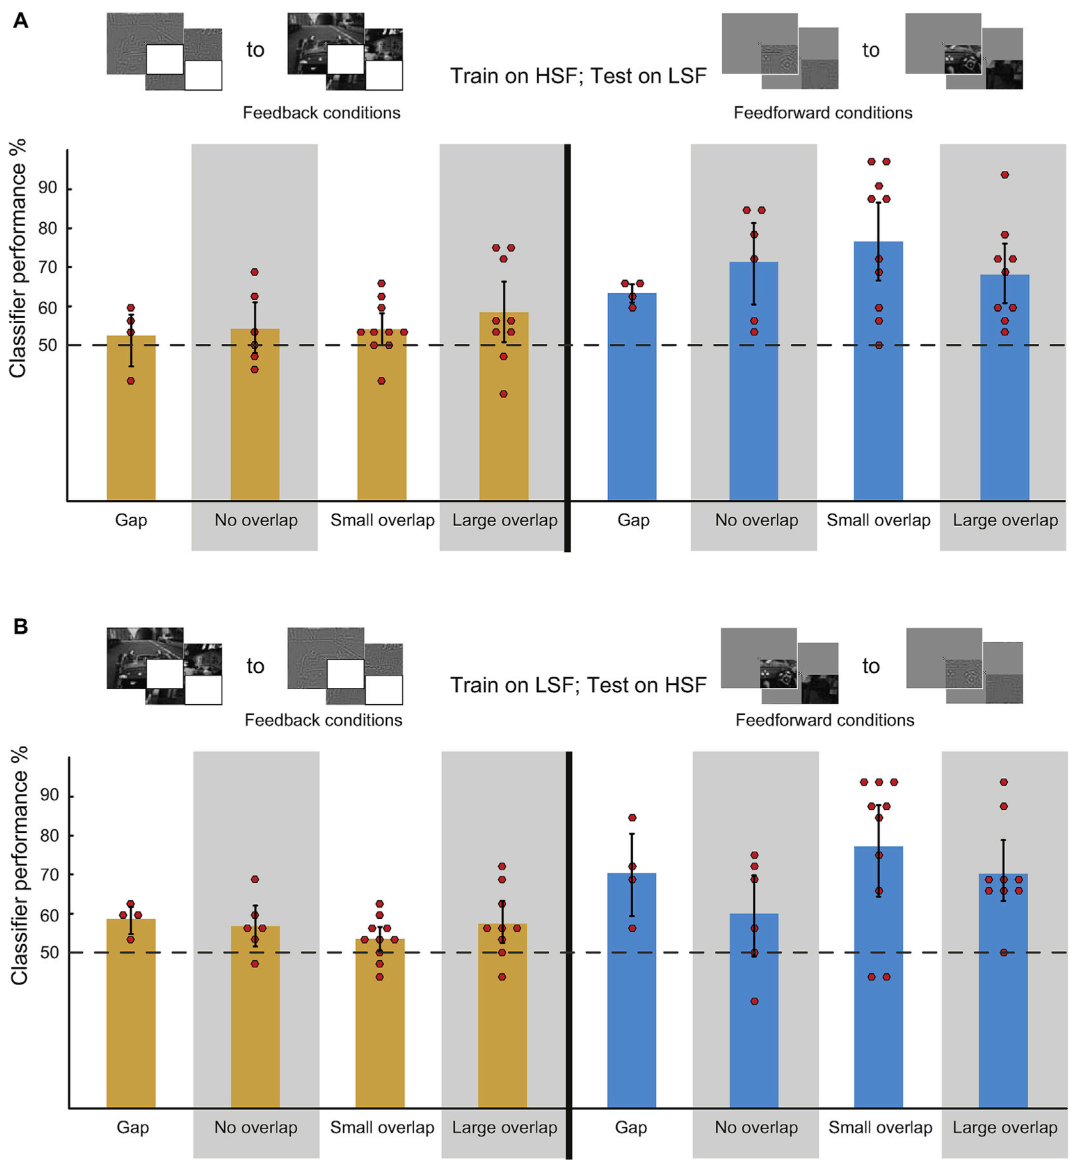

Fig. 3. Cross-classification performance for training to decode the two scenes in one SF and testing in the other, for different Overlap groups. Chance level is 50%. Lines represent 95% confidence intervals around the bootstrapped mean. Classifier performance is significantly above chance at $\alpha = 0.05$ (not corrected for multiple comparisons) if the confidence intervals do not intersect with the chance line. The small red circles represent individual subjects' results. *Gap*, n = 4; *No Overlap*, n = 6; *Small Overlap*, n = 10; *Large Overlap*, n = 9. Images are example stimuli used for a subset of subjects. A) Classifier performance for training on HSF and testing on LSF. B) Classifier performance for training on LSF and testing on HSF.

Scene 1 to Scene 2, Feedback: 60.78%, CI [0.0323 0.0323], Feedforward: 74.25%, CI [0.0377 0.0409]; Scene 2 to Scene 1, Feedback: 58.41%, CI [0.0302 0.0269], Feedforward: 71.55%, CI [0.0431 0.0453]; for average block analaysis, see Inline Supplementary Fig. S4). Feedback appears to carry information about high or low spatial frequency that is similar across different scenes. In other words, there is some degree of similarity between the information in, for instance, the occluded region of the HSF version of Scene 1 and the occluded region of the HSF version of Scene 2.

3.4. Lack of similarity between feedback and corresponding feedforward information

Are feedback signals similar to the corresponding feedforward information? For example, are feedback signals in the HSF condition similar to the activity pattern relating to the HSF feedforward stimulation, or do they carry different information? To test feedback and feedforward similarity, we first trained the classifier to decode between the two scenes on feedback conditions and tested on feedforward conditions (and vice versa), for both HSF and LSF (Fig. 5A; for average block analysis, see Inline Supplementary Fig. S5). Classifier performance was at chance level for both HSF and LSF scenes, and for both directions of cross-classification. This suggests that the information feedback provides in the occluded region is different to the corresponding feedforward information. Secondly, we trained the classifier to decode HSF vs LSF on feedback conditions and tested its ability to decode in the feedforward conditions (and vice versa, Fig. 5B). Classifier was above chance only for Scene 2 when training on feedback and testing on feedforward (54.53%, CI [0.0302 0.0345]).

3.5. Analyses with more stringent criteria

Analyses using a more conservative method of voxel selection in the occluded region, using both the *Near Surround* and the *Inner Border* mapping (see Methods), led to a similar pattern of results (see Inline Supplementary Fig. S7), suggesting our results are unlikely to be due to "spillover" activation from the surround feedforward stimulation.

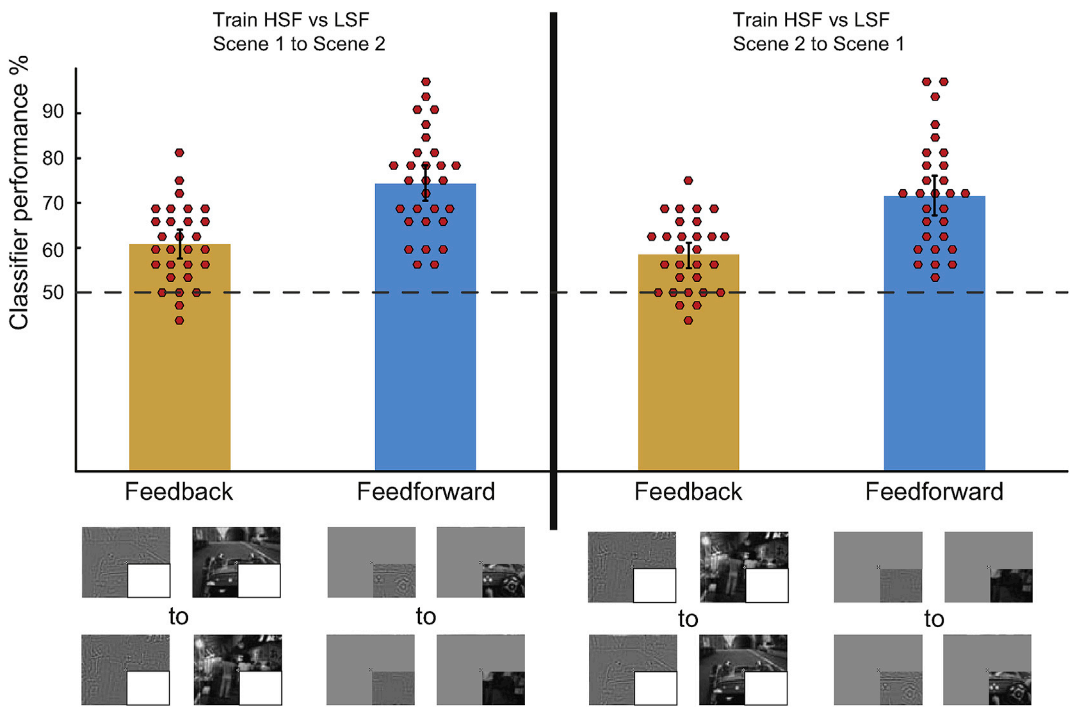

Fig. 4. Cross-classification performance for training to decode HSF vs LSF in Scene 1 and testing in Scene 2 (and vice versa), for feedback and feedforward conditions. Chance level is 50%. Lines represent 95% confidence intervals around the bootstrapped mean. Classifier performance is significantly above chance at α = 0.05 (not corrected for multiple comparisons) if the confidence intervals do not intersect with the chance line. The small red circles represent individual subjects' results. N = 29. Images are example stimuli used for a subset of subjects.

4. Discussion

The present study aimed to investigate the spatial frequency information carried by feedback signals to an occluded region of the visual field in V1. First, we have replicated the findings of Smith and Muckli (2010) and Muckli et al. (2015), by showing that non-stimulated V1 receives contextual feedback from the surrounding regions of the scene, and this can occur even with reduced information, in other words, when the surrounding scene only contains information in certain spatial frequencies. Second, we show that feedback carries information about both low and high spatial frequencies, suggesting that it transmits both coarse and fine-grained information, respectively (Fig. 2). Third, we found feedback signals containing LSF information are similar to feedback signals containing HSF information, even without a large overlap in spatial frequency bands of the LSF and HSF scenes (Fig. 3). Fourth, feedback carries similar information about the spatial frequency band across different scenes (Fig. 4). Finally, we demonstrate that the information in the occluded region does not represent a direct filling-in of the missing feedforward input (Fig. 5).

4.1. Feedback contains both coarse and fine-grained information

We found scene-specific information patterns in the occluded region, with both LSF and HSF surrounds. This suggests that both LSF and HSF surround information gave rise to meaningful feedback signals. Our results are in line with the flexible usage hypothesis (Schyns and Oliva, 1997), which proposes that both HSF and LSF information can be processed first and demands of the task can bias the visual system to attend to the most informative scale (e.g. Oliva and Schyns, 1997; Schyns and Oliva, 1999, 1994; Sowden et al., 2003). In addition to theories suggesting LSFs are important for providing contextual information and contribute to top-down expectations (e.g. Bar, 2003; Bar et al., 2006; Breitmeyer, 2014; Kveraga et al., 2007), we show that HSF scene information is also sufficient to trigger contextual feedback. It is plausible that LSF is an important source of contextual information in natural viewing when all spatial scales are available. However, in our filtered scenes, informative context was presented only in one type of spatial scale. Thus, in the HSF stimulus, high spatial frequencies were the only informative spatial scale and therefore this spatial scale was used for top-down context. Walther et al. (2011) found a similarity between brain activity in relation to line drawings (HSF) and coloured photographs of the same image (full SF spectrum) in several brain regions including V1, suggesting that impoverished HSF information is sufficient for scene identification. Rajimehr et al. (2011) showed that the parahippocampal place area (PPA), which processes scenes and spatial context (Bar and Aminoff, 2003), responds preferentially to HSF information. Kravitz et al. (2011) have shown that PPA represents spatial aspects of scenes. Various extrastriate regions contribute to the feedback of scene information to V1. PPA is one candidate region amongst others that may feed back high spatial frequency information of visual scenes to V1. However, exploring the extent of cortical areas contributing to the feedback signals of visual scene information as well as defining their functional role is beyond the scope of this manuscript.

Oliva and Torralba (2006) argue that scene gist might be provided by global scene structure, but which might not necessarily use LSF. Walther et al. (2011) found that they could cross-classify from line drawings to coloured photographs, and vice versa. However, the coloured photographs had a full SF spectrum and therefore there was shared information between the two types of image. To address how broad feedforward and feedback signals are in terms of spatial frequency tuning, we tested cross-classification from one SF type to another when only some or none of the SF information was shared between the two versions. We could cross-classify from one SF to the other, in both feedforward and feedback conditions even when there was no overlap in terms of the spatial frequency spectrum between the LSF and HSF version (for example, in the *Gap* condition), suggesting brain activity patterns were similar for the LSF and HSF version, commensurate with a broad tuning. Training on LSF and testing on HSF worked better than vice versa as cross-classification was successful for all Overlap groups. Training on HSF and testing on LSF, on the other hand, was only successful for the

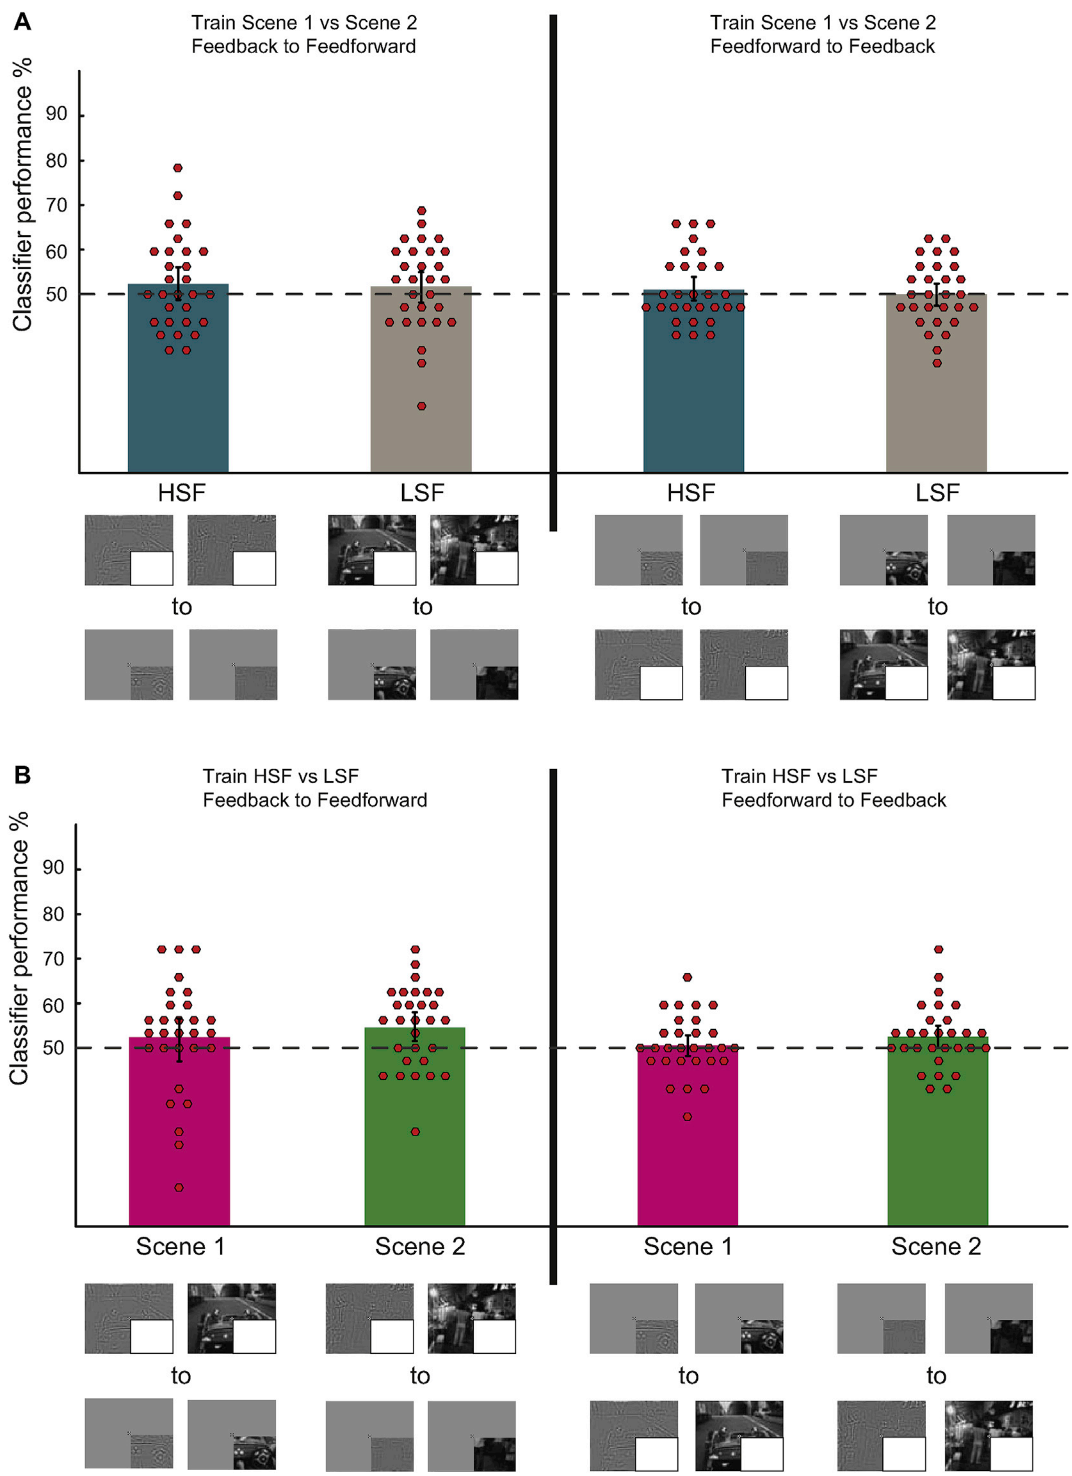

Fig. 5. Cross-classification performance for training on feedback and testing on feedforward conditions (and vice versa). Chance level is 50%. Lines represent 95% confidence intervals around the bootstrapped mean. Classifier performance is significantly above chance at $\alpha = 0.05$ (not corrected for multiple comparisons) if the confidence intervals do not intersect with the chance line. The small red circles represent individual subjects' results. N = 29. Images are example stimuli used for a subset of subjects. A) Classifier performance for decoding Scene 1 vs Scene 2 in HSF and LSF conditions. B) Classifier performance for decoding HSF vs LSF in Scene 1 and Scene 2.

Large Overlap group. This means that although a large overlap is not necessary for generalization, the direction of cross-classification is important. Training the classifier on a coarse signal and classifying a more fine-grained signal is easier than vice versa.

Both HSF and LSF surrounds elicited meaningful feedback and since it was similar for the two versions of the scene, one hypothesis is that the surround elicits feedback containing the same general template (in some SF band) for the particular image, regardless of the SF it is presented in. Since we find that both HSF and LSF image surrounds could induce meaningful information in the occluded region, we can at least say that feedback transmits information about fine-grained and coarse image features, even if we cannot say how fine-grained or coarse this transmitted information is. However, we would argue that feedback for the HSF and LSF scene is not identical as we could also decode between HSF and LSF in the occluded region and generalise this decoding to another scene. This suggests that the tuning of feedback does depend at least somewhat on the surround of the scene; presumably mostly but maybe not exclusively HSF for HSF surround and similarly LSF for LSF surround.

We have demonstrated that on the one hand feedback signals share a similarity between the LSF and HSF scenes even when there is little shared information in terms of the spatial frequency of the scene surround. On the other hand, the HSF and LSF surrounds elicited informative feedback signals equally well, and we could also decode between HSF and LSF versions of the same scene. This suggests that feedback may contain coarse information about the scene, but nevertheless retain some fine-grained properties. This is in line with previous work (Morgan et al., 2016) that showed that feedback contains information about both the category of a scene (such as a forest, corresponding to coarse structure) as well as about individual scenes within a category (fine-grained structure of a particular example of a forest scene). However, we should clarify that we do not infer a double dissociation about LSF and HSF in relation to category and exemplar. HSFs also likely contain information diagnostic of category. Rather than this double dissociation, we suggest that a visual feature diagnostic of (either perceptual or semantic) category could be contained in either LSF or HSF, and this might depend on the task, the length of exposure to the image and the individual category, among other things (Schyns and Oliva, 1994, 1997). It is possible that LSF is used mainly for category (e.g. cityscapes typically contain vertical lines of buildings which might be captured by LSF), but that HSF is better suited for discriminating individual exemplars (e.g. to distinguish New York from Paris).

One possible limitation is that we also saw SF generalisation for the feedforward conditions. Even when there was no shared SF information in the HSF and LSF version, we could cross-classify from HSF to LSF, and vice versa. This suggests that our SF cut-offs were not wide enough to lead to completely separate brain activity patterns in the feedforward regions. Since feedback signals are likely to be even coarser than feedforward signals (Muckli et al., 2015), it is unsurprising we could generalise across SF in feedback conditions as well. As noted before, feedforward conditions had bottom-up stimulation which still includes some feedback. Therefore, a generalization in the feedforward conditions could be attributed in part to the feedback. With a wider Gap condition, we might be able to better probe how broad feedback signals are in comparison to feedforward. It may be that a gap of around two octaves or more is needed to avoid an overlap in neural signals (Sowden and Schyns, 2006). Alternatively, this generalisation could be possible because there is a relationship between where the object boundaries are in the different SF bands. A blurry or a sharp edge is still the same edge, and hence the brain activity pattern is similar, if the scene representation in the visual system is related to figure-ground segregation and object identification.

4.2. Feedback signals do not correspond to a direct filling-in of the missing feedforward information

We saw that feedback is meaningful for both HSF and LSF scenes. However, how do the activity patterns in the occluded region compare to those in the corresponding feedforward region of the scene? Are feedback signals in the HSF and LSF scenes similar to the corresponding feedforward signals? To answer this question we trained the classifier on feedback conditions and tested whether it can use the same information to decode the stimuli in the feedforward conditions. We did not find a similarity between feedback and feedforward signals, suggesting that feedback signals do not represent a direct filling-in of the feedforward information. This is in contrast to the findings of Smith and Muckli (2010). However, they used a full scene as the feedforward condition, whereas in the present study we used a feedforward quadrant with no surround. We have previously demonstrated (unpublished observations, Revina et al., 2016) that this feedback and feedforward similarity depends on the amount of surrounding scene information. This is because removing the scene surround outside the feedforward quadrant removes the contextual surrounding feedback, which drives this similarity effect. This finding might be surprising if we consider that feedback has been implicated in transmitting predictions and expectations about the scene (e.g. Bastos et al., 2012; Clark, 2013; Friston, 2010; Kok et al., 2012; Rao and Ballard, 1999), and we might therefore expect feedback to represent the missing feedforward information. It is possible that the missing scene information is still represented, but in a different format. For example, it may be that the information is coarser in terms of its content because of the larger visual fields in higher visual areas or less precise retinotopically (e.g. de-Wit et al., 2012) or because feedback and feedforward signals project to different cortical layers (Muckli et al., 2015; Rockland and Pandya, 1979). Muckli et al. showed using high resolution fMRI that during normal visual stimulation, feedforward information peaks in the mid-layers, while contextual feedback information peaks in the superficial layers.

4.3. Low level properties

One question arises about the level of information that feedback transmits. Does it code for high level information, for example the scene category, or does it have some low level information, such as information about the spatial frequency band that is not specific to a particular scene category? We were able to decode HSF vs LSF on Scene 1 and generalise this to decoding HSF vs LSF on Scene 2. This similarity of the SF information between the different scenes suggests that feedback carries general information about the spatial frequency band of the surround that is unrelated to the specific structure of the scene in question. However, this finding may be explained by differences in contrast since we did not equalise the images for contrast. Kauffmann et al. (2015) showed that HSF and LSF images which were equalised for contrast activated higher-level scene-selective regions differently than images where the contrast of the HSF and LSF was not equalised. Therefore, it would be useful to further test our finding with SF filtered scenes that are better matched for contrast.

4.4. Perceptual filling-in and amodal completion

None of the presented stimuli triggered a perceptual filling-in, in other words, the subjects perceived the missing quadrant as a white rectangle. The occlusion paradigm is therefore different to paradigms using Kanizsa figures (e.g. Kok et al., 2016) or neon colour spreading where the non-stimulated region is perceptually filled-in. Our paradigm relates to amodal completion where the subjects do not directly perceive, but rather infer a presence of the continued stimulus behind the occluder, from knowledge about natural objects and scenes. Filling-in (modal completion) and inference (amodal completion) may have different neuronal processes, discussed, for example, in Lawrence et al. (2017).

4.5. Conclusion

Using pattern analysis techniques we probed the information content of cortical feedback signals, and show that they contain both high and low spatial frequency information about the surrounding scene. Further to behavioural studies demonstrating that both HSF and LSF information can be used for scene processing depending on which spatial scale is informative, we show on a neural level that both HSF and LSF scene surrounds can contribute to top-down contextual feedback. In addition, we demonstrate that although we can decode between HSF and LSF feedback, there are also similarities in feedback for the two versions of the scene, suggesting its tuning is quite broad. Finally, we find that feedback information is not a direct filling-in of the missing feedforward input.

Funding

This work was supported by the European Research Council grant (ERC StG 2012_311751-'Brain reading of contextual feedback and predictions' awarded to LM) and BBSRC DTP Studentship (BBSRC BB/

J013854/1 awarded to YR). This project has also received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No. 720270 (HBP SGA1).

Acknowledgements

We thank Hans-Christian Rath and Frances Crabbe for assistance with data collection. We thank Philippe G. Schyns, Andrew T. Morgan and Matthew A. Bennett for useful discussion.

Appendix A. Supplementary data

Supplementary data related to this article can be found at https://doi.org/10.1016/j.neuroimage.2017.09.047.

References

Ahissar, M., Hochstein, S., 2004. The reverse hierarchy theory of visual perceptual learning. Trends Cognitive Sci. 8, 457–464.
Ban, H., Yamamoto, H., Hanakawa, T., Urayama, S.-I., Aso, T., Fukuyama, H., Ejima, Y., 2013. Topographic representation of an occluded object and the effects of spatiotemporal context in human early visual areas. J. Neurosci. 33, 16992–17007.
Bar, M., 2003. A cortical mechanism for triggering top-down facilitation in visual object recognition. J. Cognitive Neurosci. 15, 600–609.
Bar, M., 2004. Visual objects in context. Nat. Rev. Neurosci. 5, 617–629.
Bar, M., Aminoff, E., 2003. Cortical analysis of visual context. Neuron 38, 347–358.
Bar, M., Kassam, K.S., Ghuman, A.S., Boshyan, J., Schmid, A.M., Dale, A.M., Hämäläinen, M.S., Marinkovic, K., Schacter, D.L., Rosen, B.R., Halgren, E., 2006. Top-down facilitation of visual recognition. Proc. Natl. Acad. Sci. 103, 449–454.
Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. Neuron 76, 695–711.
Boynton, G.M., Engel, S.A., Glover, G.H., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci. 16, 4207–4221.
Breitmeyer, B.G., 2014. Contributions of magno- and parvocellular channels to conscious and non-conscious vision. Philos. Trans. R. Soc. Lond., B, Biol. Sci. 369, 20130213.
Chang, C.C., Lin, C.J., 2001. LIBSVM: a Library for Support Vector Machines. Available at: https://www.csie.ntu.edu.tw/∼cjlin/libsvm. (Accessed 14 March 2017).
Clark, A., 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav. Brain Sci. 36, 181–204.
De Valois, R.L., Albrecht, D.G., Thorell, L.G., 1982. Spatial frequency selectivity of cells in macaque visual cortex. Vis. Res. 22, 545–559.
de-Wit, L.H., Kubilius, J., Wagemans, J., Op de Beeck, H.P., 2012. Bistable Gestalts reduce activity in the whole of V1, not just the retinotopically predicted parts. J. Vis. 12 (11), 1–14.
Friston, K., 2010. The free-energy principle: a unified brain theory? Nat. Rev. Neurosci. 11, 127–138.
Haynes, J.-D., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat. Neurosci. 8, 686–691.
Henriksson, L., Nurminen, L., Hyvärinen, A., Vanni, S., 2008. Spatial frequency tuning in human retinotopic visual areas. J. Vis. 8 (5), 1–13.
Hochstein, S., Ahissar, M., 2002. View from the top: hierarchies and reverse hierarchies in the visual system. Neuron 36, 791–804.
Kauffmann, L., Ramanoël, S., Guyader, N., Chauvin, A., Peyrin, C., 2015. Spatial frequency processing in scene-selective cortical regions. NeuroImage 112, 86–95.
Kay, K.N., Naselaris, T., Prenger, R.J., Gallant, J.L., 2008. Identifying natural images from human brain activity. Nature 452, 352–355.
Kok, P., Bains, L.J., Mourik, T., van Norris, D.G., F.P., de Lange, 2016. Selective activation of the deep layers of the human primary visual cortex by top-down feedback. Curr. Biol. 26, 371–376.
Kok, P., Jehee, J., de Lange, F.P., 2012. Less is more: expectation sharpens representations in the primary visual cortex. Neuron 75, 265–270.
Kravitz, D.J., Peng, C.S., Baker, C.I., 2011. Real-world scene representations in high-level visual cortex: it's the spaces more than the places. J. Neurosci. 31, 7322–7333.
Kveraga, K., Boshyan, J., Bar, M., 2007. Magnocellular projections as the trigger of top-down facilitation in recognition. J. Neurosci. 27, 13232–13240.
Lawrence, S.J., Formisano, E., Muckli, L., de Lange, F.P., 2017. Laminar fMRI: Applications for Cognitive Neuroscience. NeuroImage, in press. https://doi.org/10.1016/j.neuroimage.2017.07.004.
Morgan, A.T., Petro, L.S., Muckli, L., 2016. Cortical feedback to V1 and V2 contains unique information about high-level scene structure. BioRxiv. https://doi.org/10.1101/041186.
Muckli, L., De Martino, F., Vizioli, L., Petro, L.S., Smith, F.W., Ugurbil, K., Goebel, R., Yacoub, E., 2015. Contextual feedback to superficial layers of V1. Curr. Biol. 25, 2690–2695.
Mur, M., Bandettini, P.A., Kriegeskorte, N., 2009. Revealing representational content with pattern-information fMRI—an introductory guide. Soc. Cognitive Affect. Neurosci. 1–9.
Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cognitive Sci. 10, 424–430.
Oliva, A., Schyns, P.G., 1997. Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. Cogn. Psychol. 34, 72–107.
Oliva, A., Torralba, A., 2006. Building the gist of a scene: the role of global image features in recognition. Prog. Brain Res. 155, 23–36.
Oliva, A., Torralba, A., 2007. The role of context in object recognition. Trends Cognitive Sci. 11, 520–527.
Rajimehr, R., Devaney, K.J., Bilenko, N.Y., Young, J.C., Tootell, R.B., 2011. The "parahippocampal place area" responds preferentially to high spatial frequencies in humans and monkeys. PLoS Biol. 9, e1000608.
Rao, R.P., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat. Neurosci. 2, 79–87.
Revina, Y., Petro, L.S., Denk-Florea, C., Blum, S., Kriegeskorte, N., Muckli, L., 2016. Meaningful feedback to occluded V1 is improved by increasing local information in the surround. J. Vis. 16 (12), 572. https://doi.org/10.1167/16.12.572. Visual Sciences Society Meeting, St Pete's Beach, Florida, USA.
Rockland, K., Pandya, D., 1979. Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. Brain Res. 179, 3–20.
Schyns, P.G., Oliva, A., 1994. From blobs to boundary edges: evidence for time- and spatial-scale-dependent scene recognition. Psychol. Sci. 5, 195–200.
Schyns, P.G., Oliva, A., 1997. Flexible, diagnosticity-driven, rather than fixed, perceptually determined scale selection in scene and face recognition. Perception 26, 1027–1038.
Schyns, P.G., Oliva, A., 1999. Dr. Angry and Mr. Smile: when categorization flexibly modifies the perception of faces in rapid visual presentations. Cognition 69, 243–265.
Shushruth, S., 2011. The Role of Extrastriate Feedback in Contextual Computations in the Primate Primary Visual Cortex. Doctoral Thesis. University of Utah.
Smith, F.W., Muckli, L., 2010. Nonstimulated early visual areas carry information about surrounding context. Proc. Natl. Acad. Sci. 107, 20099–20103.
Sowden, P.T., Schyns, P.G., 2006. Channel surfing in the visual brain. Trends Cognitive Sci. 10, 538–545.
Sowden, P.T., Özgen, E., Schyns, P.G., Daoutis, C., 2003. Expectancy effects on spatial frequency processing. Vis. Res. 43, 2759–2772.
Sugita, Y., 1999. Grouping of image fragments in primary visual cortex. Nature 401, 269–272.
Tong, F., Harrison, S.A., Dewey, J.A., Kamitani, Y., 2012. Relationship between BOLD amplitude and pattern classification of orientation-selective activity in the human visual cortex. Neuroimage 63, 1212–1222.
Walther, D.B., Chai, B., Caddigan, E., Beck, D.M., Fei-Fei, L., 2011. Simple line drawings suffice for functional MRI decoding of natural scene categories. Proc. Natl. Acad. Sci. 108, 9661–9666.
Wilson, H.R., Bergen, J.R., 1979. A four mechanism model for threshold spatial vision. Vis. Res. 19, 19–32.