Contribution of large scale biases in decoding of direction-of-motion from high-resolution fMRI data in human early visual cortex

A. Beckett [a], J.W. Peirce [a], R.-M. Sanchez-Panchuelo [b], S. Francis [b], D. Schluppeck [a,*]

[a] Visual Neuroscience Group, School of Psychology, University of Nottingham, Nottingham, UK
[b] SPMMRC, School of Physics and Astronomy, University of Nottingham, Nottingham, UK

A B S T R A C T

Previous studies have demonstrated that the perceived direction of motion of a visual stimulus can be decoded from the pattern of functional magnetic resonance imaging (fMRI) responses in occipital cortex using multivariate analysis methods (Kamitani and Tong, 2006). One possible mechanism for this is a difference in the sampling of direction selective cortical columns between voxels, implying that information at a level smaller than the voxel size might be accessible with fMRI. Alternatively, multivariate analysis methods might be driven by the organization of neurons into clusters or even orderly maps at a much larger scale. To assess the possible sources of the direction selectivity observed in fMRI data, we tested how classification accuracy varied across different visual areas and subsets of voxels for classification of motion-direction. To enable high spatial resolution functional MRI measurements (1.5 mm isotropic voxels), data were collected at 7 T. To test whether information about the direction of motion is represented at the scale of retinotopic maps, we looked at classification performance after combining data across different voxels within visual areas (V1–3 and MT+/V5) before training the multivariate classifier. A recent study has shown that orientation biases in V1 are both necessary and sufficient to explain classification of stimulus orientation (Freeman et al., 2011). Here, we combined voxels with similar visual field preference as determined in separate retinotopy measurements and observed that classification accuracy was preserved when averaging in this 'retinotopically restricted' way, compared to random averaging of voxels. This insensitivity to averaging of voxels (with similar visual angle preference) across substantial distances in cortical space suggests that there are large-scale biases at the level of retinotopic maps underlying our ability to classify direction of motion.

© 2012 Elsevier Inc. All rights reserved.

Introduction

For many cortical areas, cells with similar selective properties are arranged into clusters on the cortical surface (Mountcastle et al., 1957). In mammalian primary visual cortex, for example, cells with similar preference for stimulus orientation (Hubel and Wiesel, 1963) or ocular dominance (Hubel and Wiesel, 1969) are organized into *columns*, groups of cells with similar properties arranged perpendicular to the cortical surface.

In the human brain, the size of *ocular dominance columns* (ODC) has been estimated from post-mortem samples of patients that lost sight in one eye prior to death (Adams et al., 2007). These columns were found to be around twice the size of ODCs in monkeys (approximately 1 mm in humans compared to ≈400 μm in the macaque). Accordingly, if orientation columns in the human brain are scaled proportionally with the ocular dominance columns, estimates from the macaque indicate that they should be around half a millimeter wide (Gardner, 2010), far smaller

than the resolution generally available with conventional functional MRI. Recently, the development of ultra high-field, high resolution fMRI techniques has allowed measurements aimed at direct visualization of columnar organization in the human brain in vivo, showing both ODC (Cheng et al., 2001; Yacoub et al., 2007) and orientation-selective columns (Yacoub et al., 2008). In addition, a recent report has suggested the existence of cortical columns selective for *axis of motion* in the human equivalent of macaque area MT (Zimmermann et al., 2011).

In typical fMRI experiments performed at 3 T, the resolution required to measure activity from columns directly is not achievable. Most protocols at 3 T use approximately 2–3 mm isotropic voxels (8–27 mm^3), and therefore the activity of each voxel will reflect the summed activity of a large number of cells with different selectivities. Selective responses of individual neurons are therefore likely to cancel out, rendering the responses in voxels non-selective. Simply increasing the spatial resolution doesn't necessarily lead to improvements, as it is accompanied by a loss in signal-to-noise ratio (SNR); so maintaining the balance between signal strength and sampling resolution generally precludes the direct study of cortical columns at lower field strengths. Furthermore, blood oxygenation level dependent (BOLD) responses measured with fMRI (at any field strength)

* Corresponding author. Fax: +44 115 9515324.
  E-mail address: denis.schluppeck@nottingham.ac.uk (D. Schluppeck).

are only an indirect reflection of the underlying neural responses (Logothetis and Wandell, 2004; Ogawa et al., 1990). The inherent blurring of the measured responses by the hemodynamics (in time and space) will further reduce the very fine detail information, such as the pattern of cortical columns on the cortical surface.

Recent results using multivariate classification techniques have suggested that it may be possible to extract fine detail information relating to columnar organization from the pattern of voxel responses (Haynes and Rees, 2005; Kamitani and Tong, 2005). One of the mechanisms that have been proposed for this is that each voxel samples a mixture of different columns. The distribution of these sampled columns is unlikely to be uniform (Boynton, 2005), which may lead to slight biases for a given visual property in different voxels. While these biases may not lead to significantly different responses between stimulus classes at the univariate (voxel-wise) level, when the activity of a large number of voxels is considered together, these small, correlated, biases could lead to separable responses for different stimulus classes, and underlying neural selectivity can therefore be inferred. Multivariate classifiers have successfully been used to associate patterns of activity in visual cortex with different stimulus orientations (Haynes and Rees, 2005; Kamitani and Tong, 2005), directions of motion (Kamitani and Tong, 2006) and object class (Op de Beeck et al., 2006). In addition, the technique has now been applied in a wide variety of cognitive neuroscience studies to make inferences about the involvement of cortical areas in higher-level cognition (Haynes and Rees, 2006; Kahnt et al., 2010). The possibility that classification-techniques allow fMRI to provide sensitivity to fine neural patterns at a resolution beyond the resolution imposed by sampling (Kriegeskorte et al., 2010), has been dubbed *fMRI-hyperacuity* (Op de Beeck, 2010). If such hyperacuity in fMRI is achievable using multivariate techniques, it could provide an exciting opportunity to overcome the issues of sampling resolution and hemodynamic blurring, and allow more direct studies of neural responses.

Although fine-spatial information such as orientation columns on the cortical surface has been hypothesized as driving these classification results, the relationship has not been directly demonstrated and has proved controversial. Any bias from an uneven distribution of columns would be expected to be too small to measure given (a) the amount of variability in the BOLD signal and (b) the fact that in the case of orientation columns only a 20%–40% difference in response between preferred and orthogonal orientation has been shown for a given column using optical imaging (Fukuda et al., 2006). When combined with the very slight biases in the distribution of selective columns expected in voxel of typical sizes, this would suggest a specificity signal with a Contrast to Noise Ratio (CNR) of around 0.01 or less (Gardner, 2010), further implying much poorer performance than that generally found in classification studies.

It has also been shown that Gaussian blurring of the voxel patterns prior to classification analysis (Op de Beeck, 2010) or down-sampling data to lower resolution after acquisition (Gardner et al., 2006) does not harm performance, suggesting that fine detail information is not the source of the classification (although see Kamitani and Sawahata, 2010 and Kriegeskorte et al., 2010 for a discussion of why fine detail information may still be available in blurred images). Another suggested source of classification performance is draining veins with some form of selectivity (Gardner, 2010; Shmuel et al., 2007; Thompson et al., 2010). A study performed at 3 T using Spin Echo fMRI, which is thought to have reduced contributions from larger draining veins (Thompson et al., 2010), found reduced classification performance compared to data acquired with a more standard Gradient Echo (GE) sequence. The authors used this as evidence that classification at 3 T was driven by draining veins. It should be noted, however, that an alternative explanation for the reduction in classification accuracies could be that the spin-echo signal has intrinsically lower contrast-to-noise (Harmer et al., 2012). In our hands, for example, simple retinotopic measurements using 1.5 mm isotropic voxels at 7 T generally

require 3–4 times the amount of data when using spin-echo based imaging compared to gradient echo imaging (data not shown).

Multivariate classification could also be driven by coarse-scale information, such as a global bias for cardinal or radial orientations or directions (Clifford et al., 2009; Raemaekers et al., 2009; Sasaki, 2007; Sasaki et al., 2006). For example, in the orientation domain it has been demonstrated that a bias for radial orientations covaries with retinotopy, that classification remains above-chance even after averaging voxels together based on their retinotopy, and that removing the map-based component of the orientation signal by projecting it out of the data reduces classification accuracy (Freeman et al., 2011). These results were suggested to demonstrate both the sufficiency and necessity of a retinotopically-organized preference for radial orientation for successful classification of orientation. However, it should be noted that classification of stimuli with radially balanced orientation components, such as glass patterns, is also possible (Mannion et al., 2009; Seymour et al., 2009), so the exact nature of the contribution of coarse scale components is a subject of debate (Clifford et al., 2011).

Using random dot motion stimuli where the direction of global motion rotated by 360° over a given period, Raemaekers et al. (2009) demonstrated a radial bias for motion in areas V1–V3, but not in hMT (the human homologue of macaque area MT). Clifford et al. (2009) also showed a radial bias for motion in areas V1–V3 using motion defined contours of varying orientation, with the bias for radial motion being independent of the orientation of the contour, although that study did not consider area hMT. To test whether the coarse scale signal suggested by these experiments may be responsible for the classification seen in previous studies, in the current study we assess the effect of averaging based on retinotopy across voxels to test the sufficiency of such a coarse scale preference for radial direction for successful classification. Previous classification studies found some evidence of a retinotopic bias (Swisher et al., 2010), but no effect of removing radial biases for orientation and motion on classification accuracy (Kamitani and Tong, 2005, 2006). There is therefore also still a debate about this question in the literature.

Here, we performed a motion-classification experiment at ultra-high field strength (7 T), which allowed us to use 1.5 mm isotropic voxels (volume, 3.375 mm^3), much smaller than a previous study of direction classification (Kamitani and Tong, 2006), performed at 3 T, which used 3 mm isotropic voxels (volume, 27 mm^3). The voxel size in our study is still too large to sample the direction of motion columns directly, but the reduced voxel volume should lead to a smaller subset of direction selective columns being sampled, which in turn should increase the biases in the voxel sampling of direction columns, and therefore the selectivities of the voxels themselves (cf Fig. 1 in Boynton, 2005). One would expect the reduction in voxel size to increase classifier accuracy if columnar information is critical and voxels are straightforwardly sampling the columns (Kriegeskorte et al., 2010). However, increasing the size of a voxel also increases noise cancelation due to signal averaging (Kamitani and Sawahata, 2010). Due to this trade-off between sampling specificity and signal to noise ratio, the benefit expected by reducing voxel size is not completely clear a priori.

The aim of the current study was to identify the possible information sources used by the classification process. The higher resolution afforded by 7 T fMRI should allow a stronger direction-selective bias for each voxel, which is expected to improve performance if classification is driven by columnar-level information. Following Freeman et al. (2011), we therefore tested whether a retinotopic bias for radial directions of motion exists that is *sufficient* to allow classification.

Methods

Subjects

Five subjects experienced in fMRI experiments and with normal vision participated in this study with written consent. Procedures

were conducted with approval from the University of Nottingham ethics committee. Each subject participated in several scan sessions: one session at 7 T to characterize responses to drifting dot motion stimuli and to obtain retinotopic (visual field angle) maps at the same spatial resolution. For each subject a T1-weighted (MPRAGE, 1 mm isotropic) anatomical image of the whole brain was acquired in a separate scanning session at 3 T. In addition ROIs for the key visual areas were identified via retinotopic mapping at 3 T.

Stimuli and task

Stimuli were generated using MATLAB (Mathworks, Natick, MA) and MGL (available at http://justingardner.net/mgl) on a Macintosh Intel computer. Stimuli were projected from an Epson EMP-8300NL LCD Projector onto a back-projection screen close to the bore of the magnet. Subjects viewed the projected stimuli through a set of mirror-glasses (maximum visible eccentricity of 16.2° of visual angle).

Two experimental paradigms were used in each scan session: a traveling wave paradigm in which visual stimulation created waves of activity across cortical regions with a retinotopic map of visual space, and a block paradigm with motion in one of eight possible directions. In the traveling wave paradigm, we used a motion retinotopy stimulus (Huk et al., 2002), consisting of a motion-defined wedge that rotated slowly through the visual field. The wedge was defined in a 24.6° diameter circular aperture filled with white dots on a black background (dot density 5 dots per degree^2 of visual angle, dotsize = .1°). At a given point in time, dots within a 90° wedge moved toward and away from fixation, while the rest of the dots within the aperture remained stationary. The wedge advanced 15° every second, completing a full revolution each 24 s, and completed 5 revolutions each scan (total, 2 min per scan). During each scan, the subjects performed a demanding two-interval forced-choice contrast discrimination task at fixation to control for attention. Subjects completed 2 retinotopy scans per session. Retinotopy stimuli evoked a traveling wave of activity in retinotopically organized visual areas. For each voxel, we calculated the temporal phase (as well as coherence and amplitude) of the best-fitting sinusoid at the stimulus-alternation frequency to the BOLD signal. The phase value corresponds to the preferred angular position in the visual field represented by that voxel (Fig. 1).

The block-classification stimuli also consisted of white dots in a circular aperture (diameter 24.6° of visual angle). For these stimuli, all the dots drifted in the same direction at the same time for blocks of 16 s (order randomized), following the paradigm of Kamitani and Tong (2006), with a block of fixation at the beginning and end of each scan (total, 2 min 40 s per scan). The dots had a lifetime of 200 ms. Subjects completed 16–20 of the classification runs per scan-session. As with the retinotopy scans, subjects performed a contrast discrimination task at fixation to control for attention.

We also collected an additional data session for two of the subjects using split-hemifield stimuli as a control. The stimuli were identical to the block-classification stimuli, except that the visual field was split into two apertures along the vertical meridian, with the dots in each aperture drifting in different directions, independently chosen, for each 16 s block. This was done to exclude the possibility that eye movements (which could not be measured in the 7 T scanner) could explain the classification results; in these stimuli no consistent pattern of eye movements could lead to reliable decoding of direction (Swisher et al., 2010).

Magnetic resonance imaging (7 T)

High-resolution magnetic resonance imaging data were collected on a 7 T system (Achieva, Philips) using a volume transmit coil and a 16-channel receive coil (Nova Medical, Wilmington, MA). To minimize head motion, we stabilized participants with a customized MR-compatible vacuum pillow (B.u.W. Schmidt, Garbsen, Germany) and foam padding.

Functional data for 2 subjects were obtained using three-dimensional gradient echo (T2* weighted) echo-planar imaging with the following parameters: TE = 25 ms, TR = 85 ms, FA = 22°, dynamic scan time (time to acquire one MRI volume) = 2 s, parallel imaging with a SENSE reduction factor r = 2.35 in the foot–head direction, and r = 2 in the anterior–posterior direction (Pruessmann et al., 1999). The remaining subjects data were obtained using two-dimensional gradient-echo (T2* weighted) echo-planar imaging with the following parameters: TE = 25 ms, TR = 2000 ms, FA = 78°, parallel imaging (SENSE reduction factor 2.35 in the foot–head direction, 2 in the anterior–posterior direction). The spatial resolution of the acquired data was 1.5 mm isotropic and the field of view (FOV) were 192 mm, 192 mm, and 51 mm in the anterior–posterior (ap), right–left (rl) and foot–head (fh) directions, respectively. Magnetic field inhomogeneity was minimized using a local, image-based shimming approach (Poole and Bowtell, 2008; Wilson et al., 2002), described elsewhere (Sanchez-Panchuelo et al., 2010).

In each scan session, we acquired functional data in 16–20 scans: 2 scans comprising 60 volume acquisitions (2 min) for the traveling wave design, and 14–18 scans comprising 80 volume acquisitions (2 min 40 s) for the block design. Subjects were given several seconds of rest between consecutive scans.

To simplify alignment of the functional MRI data to the whole head anatomy (MPRAGE), we also obtained a high-resolution
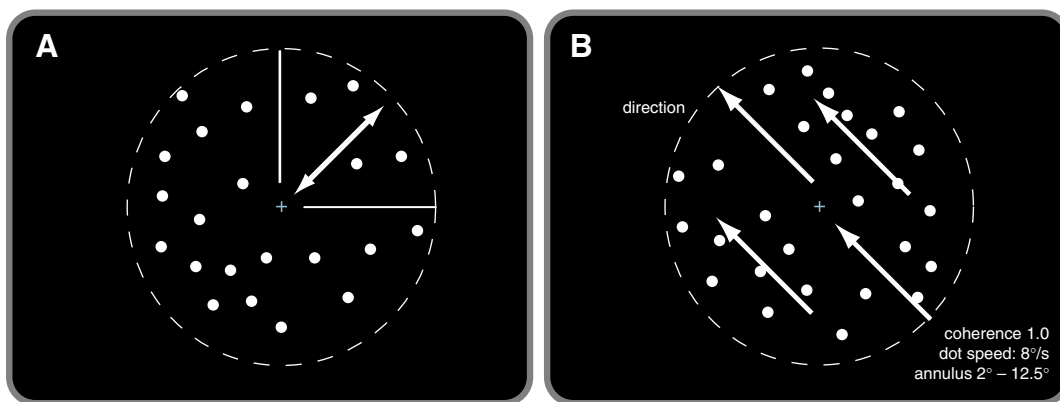

Fig. 1. Examples of stimuli used in (A) retinotopy and (B) classification scans. (A) For the retinotopy (visual field angle) scans, dots with a 90° wedge drifted away and toward fixation, with the wedge rotating through the visual field over a 24 s period. (B) In the classification scans, all the dots in the field drifted in a common direction, with the direction (white arrow) changing every 16 s. Note that the dot density and size are not to scale in this illustration. In the actual stimulus display, dot density was 5 dots/deg^2 and dot size was .1°.

$(0.5 \times 0.5 \times 1.5 \text{ mm}^3)$ T2*-weighted anatomical volume at the end of each scan session. These images had the same slice prescription and coverage as the functional data. Parameters: TE = 11 ms, TR = 608 ms, FA = 32°. Example images are shown in Fig. 2. These scans took approximately 2 min. The T2* anatomy was used for nonlinear registration between the distorted EPI images and (intrinsically undistorted) anatomy for each subject, and to calculate the warp fields necessary to correct for distortions (Klein et al., 2009). We used the nonlinear registration as a final step for visualization to project statistical maps calculated in the classification analysis (in the space of the originally acquired data) to be displayed on the subjects' flattened cortical surfaces (in the un-distorted anatomical space), as well as allowing regions of interest (ROIs) defined in undistorted anatomy-space to be used in analyzing the 7 T EPI (BOLD) data.

fMRI preprocessing

The T2*-weighted anatomical volumes acquired for each scanning session were aligned to the whole-head anatomical volume of each subject's brain using a robust image registration algorithm (Nestares and Heeger, 2000). FMRI data were motion-corrected within and between scans for each session using standard motion correction algorithms. We also applied linear trend removal to the time course at each voxel. Note that no temporal or spatial smoothing was applied to the data. Data for the traveling wave scans for each subject from each session were averaged together, voxelwise. Patterns for classification were created by shifting the time series by 4 s (2 TRs) to account for the hemodynamic lag (for details see Kamitani and Tong, 2005, 2006), and temporally averaging the MRI signal in each voxel across each of the 16 s blocks, leading to one value per voxel per block. The responses for each voxel were then z-score normalized for each individual scan to minimize baseline differences between scans before passing data into classification analysis.

Traveling wave analysis

The data from the retinotopy scans was analyzed using standard Fourier-based (traveling wave) methods (Huk et al., 2002; Larsson and Heeger, 2006). Briefly, the averaged voxel-wise time-series from the motion-defined retinotopy scans were fitted with a sinusoid matching the period of the stimulus rotation (24 s). The phase of the sinusoid that best fitted the time series corresponded to the angular-position of the stimulus in the visual field that best drove responses for that voxel; the coherence between the sinusoid and the time series indicated the extent to which each voxel's activity was modulated by the retinotopy stimulus.

Multivariate classification analysis

Regions of interest for early visual areas (V1–V3) for the classification analysis were drawn from a separate session of retinotopy at 3 T to allow identification on the undistorted cortical surface. The ROIs were transformed into the (distorted) EPI space for each subject and scanning session using the nonlinear alignment between the EPI images and the T2* weighted anatomy. Voxels from each ROI were chosen on the strength of their response to the within-session (but independent) retinotopic localizer, by selecting the 500 voxels with the highest coherence values.

A linear Support Vector Machine (SVM) was used to classify the patterns of activity in each ROI according to motion direction. The input patterns were created by averaging the activity in each voxel across each of the 16 s blocks of stimulus motion, leading to one value per direction per run for each voxel.

To allow comparison with previously published results, we used the same method for building a direction decoder as the Kamitani and Tong (2005, 2006). A 'linear ensemble detector' for direction of motion was used to calculate a weighted sum of the voxel intensities in the input patterns: this was done by optimizing the weights such that the detector for each direction maximized its output for its preferred direction, and reduced it for other directions. The calculation for a linear ensemble detector for direction θ_k is a linear function of voxel inputs $\mathbf{x} = (x_1, x_2, \ldots, x_d)$

$$g_{\theta_k}(\mathbf{x}) = \sum_{i=1}^{d} w_i x_i + w_0 \tag{1}$$

where w_i is the weighting for each voxel i and w_0 is the bias. To construct the required weighting function for each direction, linear discriminant functions were calculated for each pairwise combination of directions using linear support vector machines (SVM). These pairwise discriminant functions were a weighted sum of the voxel inputs in the patterns, satisfying

$$g_{\theta_k \theta_l}(\mathbf{x}) > 0 \tag{2}$$

$$g_{\theta_k \theta_l}(\mathbf{x}) < 0 \tag{3}$$

g being greater than 0 if $\mathbf{x}$ is a pattern of fMRI response induced by direction θ_k, and less than 0 for direction θ_l. A linear SVM found the optimal weights and bias for each discriminant function based on the
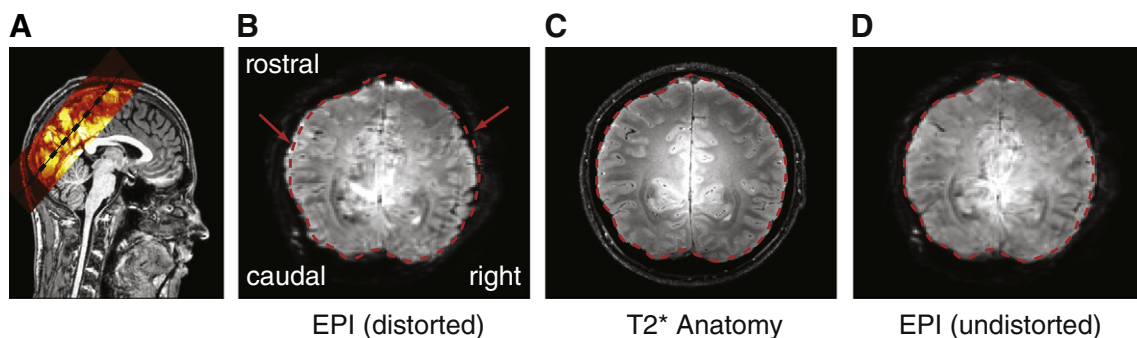

Fig. 2. Examples of MRI images. (A) Sagittal slice of T1-weighted MPRAGE anatomy (gray scale) showing location of the functional imaging stack (red color map). Coverage was optimized occipital and occipito-temporal areas. Dashed black line, location of slices shown in B–D. (B) Gray scale, mid-stack image showing a typical example of EPI data obtained 7 T. Dashed red line, outline of brain for corresponding, undistorted, T2* anatomy image. Arrows indicate areas of residual geometric distortions that cause voxels in the EPI data to be displaced from their true anatomical location. (C) Corresponding slice of high-resolution T2* weighted anatomy; these images are not affected by distortions in the way EPI data are; dashed red line as in B. (D) EPI image after un-distortion procedure; note the close correspondence between corrected EPI data and outline derived from T2* anatomical image. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

training dataset. The pairwise functions for a given direction versus all other directions were added together to yield the linear detector function for that direction.

$$g_{\theta_k}(\mathbf{x}) = \sum_{m \neq k} g_{\theta_k \theta_m}(\mathbf{x}) \tag{4}$$

This function becomes greater than zero when the input vector $\mathbf{x}$ (voxel pattern of fMRI response) is one induced by its preferred direction θ_k. The *test* patterns are assigned to a given direction label based on which detector function yields the largest output.

To test the accuracy of the classifier, we used 'leave one run out' cross validation. The classifier was trained using the data on all but one of the runs, and then its accuracy at classifying the patterns from the remaining run was assessed. This process was repeated for each run in turn. This process assures the statistical independence of the test and training data, as they are drawn from separate runs. To test whether the classification accuracy in the ROI analysis was significantly above chance, we performed a non-parametric permutation test. 1000 classification analyses were performed with the labels indicating which patterns of activity related to which direction of motion shuffled. This produced a distribution of classification accuracies expected under the null hypothesis that these patterns did not relate to the perceived direction of motion. Accuracies from the correctly labeled dataset were considered significant if they were higher than the 95th percentile of this null distribution (p<0.05, one tailed permutation test).

To test whether a coarse-scale retinotopic preference for direction of motion, for example a preference for radial motion, was sufficient for direction classification, the classification analysis was repeated after averaging the input voxels based on the phase of their response to the retinotopic stimulus. Voxels were assigned to bins, each bin corresponding to a separate range of polar angles. The time-series of the voxels within each bin were then averaged to yield a smaller number of 'super-voxels' (Freeman et al., 2011), which were then used to generate the patterns for classification analysis detailed above. This process was repeated for bins of varying width. To test whether this retinotopic scale signal allowed successful classification, we repeated the averaging process, but shuffled the phase values for each voxel prior to binning and

averaging, leading to equivalent levels of averaging that did not preserve the retinotopic signal.

Results

Retinotopic maps

The phase values from the retinotopy scans were nonlinearly aligned to the whole head anatomy and displayed on the flattened cortical surfaces for each subject to determine how well they matched with the visual areas identified in a separate retinotopy session. Phase values progressed cleanly across the surface, with reversals at the borders between visual areas (Fig. 3). Note that all the analyses were performed in native (distorted) EPI space to avoid issues of data interpolation. The nonlinear alignment was done as a final stage for visualization of the data on flat maps and surface representations of the cortical hemispheres.

The average of two retinotopy (visual field angle) scans was sufficient to produce reliable visual field maps for subsequent analyses. An example of the timecourses from two small ROIs (spanning voxels in the upper and lower portions of the calcarine sulcus) is shown in Fig. 3D. FMRI responses from these small (1.5 mm isotropic) voxels were robust and on the order of several % signal change.

Classification

In our multivariate analysis, we found above-chance accuracies for direction of motion classification in areas V1–V3 and MT+/V5 for all 5 subjects. Permutation tests confirmed this for each individual subject. As our experiment was based on 8 possible directions of motion, the chance level for *proportion correct* was 0.125. In line with previous results, but perhaps surprisingly given previous single cell studies, classification accuracy was higher in earlier visual areas (V1–V3) than in area MT+/V5 (Fig. 4), which is known to show strong selectivity for motion at the single-cell level (Snowden et al., 1992).

To test whether a retinotopic bias (or signal) was *sufficient* to drive motion classification, we followed the analysis of Freeman et al. (2011) and performed an analysis in which we binned the voxels in each ROI
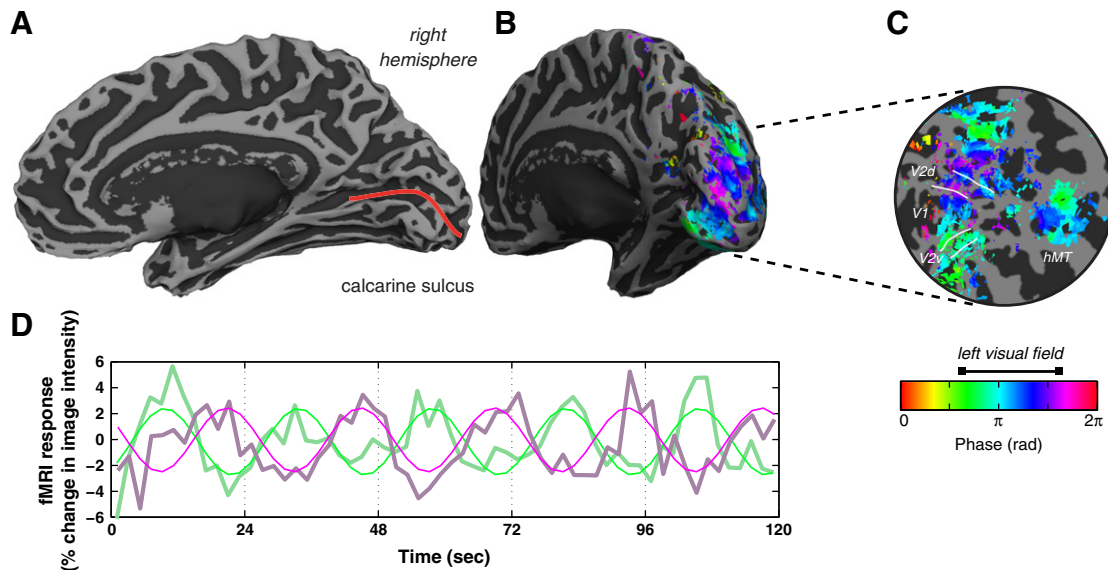

Fig. 3. Example data from in-session retinotopy scans for determining visual field angle. (A) Partially inflated right hemisphere of subject 1 showing gyral/sulcal anatomy. Red line, location of calcarine sulcus. (B) Same hemisphere as in A. Colors correspond to phase of best-fitting sinusoid to fMRI response elicited by the wedge of moving dots stimulus. Phase values are only displayed for voxels that exceeded a coherence threshold at p<.01 after correcting for multiple comparisons. Note that EPI data have been distortion-corrected to allow display on partially inflated and flattened cortical surfaces, which are derived from undistorted anatomy images. White lines and labels, borders and locations of areas V1, V2d, V2v, and human MT+. Legend, phase values from Fourier-based analysis and approximate range of values corresponding to stimulus locations in the left visual field. (D) Thick lines, sample fMRI time series from two voxels in different parts of the retinotopic map of V1 (green = upper visual meridian, magenta = lower visual meridian). Thin lines, best fitting sinusoid at stimulus frequency. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
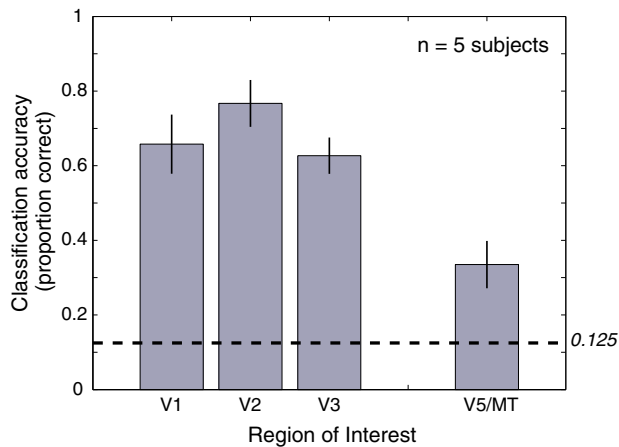
Fig. 4. Mean classification accuracies of 8-way direction classification in retinotopically defined ROIs. Error bars, standard error across 5 subjects. Dashed line indicates chance performance, p(correct) = 0.125.

based on the phase of their response to a retinotopic stimulus (a rotating wedge of moving dots). This essentially groups the voxels by their preferred visual field location. Next, we averaged the signal of the voxels in each bin. The number of bins we used ranged from 3 to 1400 (corresponding to bin widths of 120° to 0.26° of polar angle). For each number of bins we separately performed a classification analysis with 'leave-one-run-out' cross-validation. As a control analysis to this averaging scheme, we repeated the same analysis but randomly assigned the voxels to each bin (by shuffling the phase values associated with each voxel), and re-tested classification performance.

If there is a bias in the voxel preferences for direction of motions, such that preferred directions of motion are correlated with particular visual field angle preferences, then we would expect classification following averaging that respects the retinotopic maps to be less disrupted. For both averaging methods, decreasing the number of bins used (i.e. increasing the range of polar angles combined) led to a decrease in classification performance (see Fig. 5). However, for the averaging scheme in which we respected the preferred visual field location for voxels—in effect preserving any information at the level of a coarse-scale map—performance dropped off less quickly. Classification performance remained above chance until the width of the bins was considerably larger than in the random averaging scheme (Fig. 5). For example, in V1 the retinotopic binning method was at 0.5 accuracy at a bin width of ≈30°, while the random method had equivalent accuracy with an equivalent width of ≈3°. Therefore there was a roughly tenfold increase in the number of voxels (as judged by the amount of polar angle represented) that could be averaged together under the retinotopic binning method to yield an equivalent level of performance. This preserved performance in the retinotopic condition is represented by a rightward shift along the logarithmic bin-width axis. This pattern was repeated across V1–V3, and was not apparent in hMT+/V5.

To quantify the effect of decreasing the number of bins used in the 'retinotopic' (polar angle) compared to 'random' averaging, we fitted an exponential growth function of the form

$$y(x) = p_c + (A - p_c)[1 - \exp(-x/B)] \qquad (5)$$

where y is the classification accuracy as a function of x (the number of bins), the parameters A and B are the asymptote and the exponential constant in units of number of bins, respectively, and $p_c = 0.125$ is the chance level. A function of this form was chosen as it has been shown to fit well the pattern of results obtained by incrementing the number of voxels used in classification analysis (Mannion et al., 2009). We used a nonparametric bootstrapping technique to estimate 95% (BC$_a$,

bias corrected and accelerated) confidence intervals for the fitting parameters A and B. Note that differences in the exponential constant between the two conditions appear as shifts on the log-scale plots used here. Also note that an increase in the number of bins (super-voxels) used corresponds to a decrease in the number of polar angles averaged together.

We computed 5000 bootstrapped replications for each fitted function by resampling the residuals; the resulting distributions were used to generate a confidence interval for each parameter, without assumptions about the distribution of the raw data or the residuals (Efron and Tibshirani, 1993). The parameter A, representing the asymptote of each curve, did not significantly differ across the two averaging schemes, whereas the exponential growth parameter B showed a significant difference in the two averaging schemes in early visual areas (Fig. 6). Note that the error bars in those plots represent 95% confidence intervals. We considered fitting parameters with non-overlapping confidence intervals as statistically significantly different.

In a control experiment, we tested (a) whether the ability to classify direction of motion of stimuli could be due to a specific pattern of eye movements and (b) whether the pattern of results was particular to our choice of full-field stimuli. The stimuli had the same properties as in the main experiment, with the modification that we split the display into left and right hemifields and displayed moving stimuli whose directions were independently drawn for each block. We added this control in part because technical limitations prevented us from measuring eye movements in the 7 T scanner. ROIs corresponding to the left and right visual fields were analyzed and binned separately over the polar angle ranges represented in each ROI. We found the same pattern of results for these stimuli. Firstly, the direction of motion could be decoded from the patterns of activity in the visual cortex as is evident by the asymptotic classification accuracy in Fig. 7, which was significantly above chance level. Secondly, binning and averaging voxels according to their visual field preference showed preserved performance compared to averaging after random binning, as is evident from the rightward shift of the blue curve in the log-plot in Fig. 7.

Discussion

We used high-resolution fMRI (1.5 mm isotropic, 3.375 μl voxels) to investigate the mechanisms underlying multivariate classification (or *decoding*) of direction of motion in human visual cortex. If successful classification is in fact based on information arising at the scale of the retinotopic map, then averaging voxels together in a way that preserves this information should be sufficient to preserve classification. Conversely this is not the signal driving classification, then this form of averaging will show no benefits compared to others.

Using stimuli and methods originally utilized by Kamitani and Tong (2006), we found above chance classification of the observed motion of drifting dot patterns in areas V1–V3 and MT+/V5, with higher levels of accuracy in the earlier visual areas. The finding that area MT+/V5 shows lower classification accuracy than earlier visual areas, despite macaque MT+/V5 containing a greater proportion of direction selective cells, was suggested by Kamitani and Tong (2006) possibly to arise from the smaller number of voxels available from MT+/V5 compared to V1, as when they reduced the number of voxels used in early visual areas, classification accuracies were equivalent. However, given the higher resolution available at 7 T, we were able to increase the number of voxels available for classification from MT+/V5, without improving the classification accuracy. This would seem to indicate that a simple lack of features is not the cause of this difference in performance. Other factors that could lead to reduced performance even if equivalent (or greater) selectivity exists in MT+/V5 are the arrangement or distribution of direction selective columns leading to very small voxel-wise biases (Kamitani and Tong, 2006; Malonek et al., 1994), or differences in the amplitude of BOLD response in different parts of the cortex (Smith et al., 2010). Due to the wide range of possible factors involved in determining
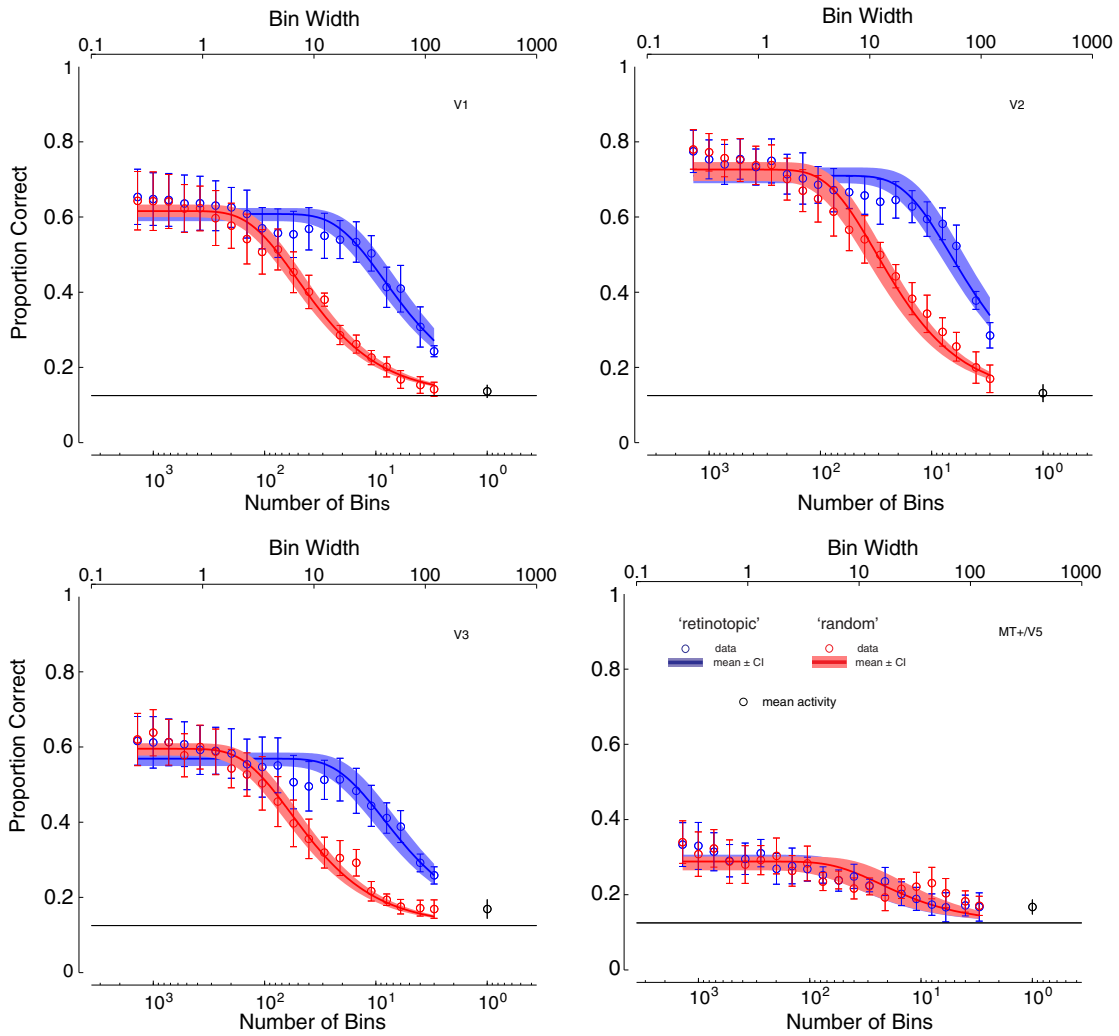
Fig. 5. Effect of binning and averaging 500 voxels from V1–V3 and MT+/V5 based on their visual angle preference. Each panel shows data from a different ROI (V1, V2, V3, and hMT+/V5). Blue symbols, binning and averaging by preferred visual field angle with increasing bin size (note log scale). Blue lines and shaded regions, exponential fit and 95% confidence intervals (see Methods for details), Red symbols, equivalent averaging but with shuffled preferred visual field angle labels, therefore providing a randomized control. Red lines and shaded regions, exponential fit and 95% confidence intervals for shuffled control. Dashed black line, chance level proportion correct (0.125). Black symbols in each panel, control indicating that classification accuracy drops to chance level when using mean fMRI response across each ROI. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

classification accuracy in a given area, there are many difficulties in comparing classification performance across cortical areas and inferring differences in selectivity from classification results.

To assess whether a coarse-scale signal would be sufficient for successful classification, rather than a high spatial frequency signal such as columnar distribution, we repeated the classification after
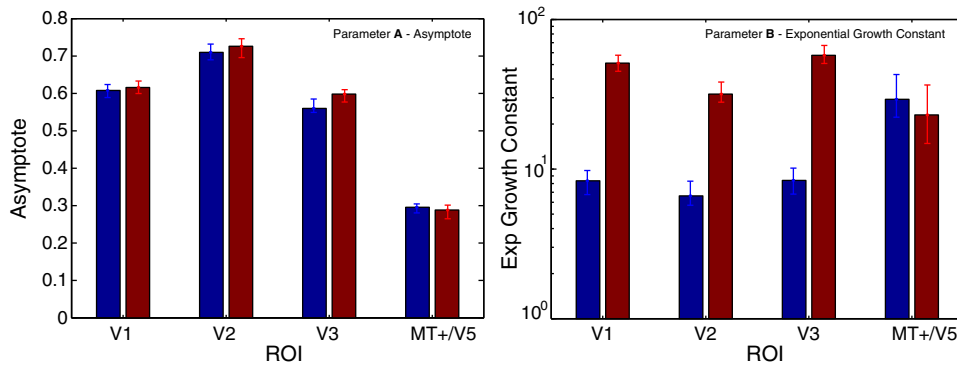

Fig. 6. Mean parameters *A* (left panel) and *B* (right panel) from exponential fit to the data in Fig. 5. Error bars represent the 95% Confidence Intervals from the bootstrapping procedure. Asymptotic classification accuracy (top panel) in the 'retinotopic' and 'random' binning and averaging schemes is the same across all ROIs, while the scale parameter *B* (representing the number of bins required for successful classification) significantly differs (as judged by the 95% CIs from the bootstrapping procedure) between 'retinotopic' and 'random' binning for V1-V3 −, but not MT +.
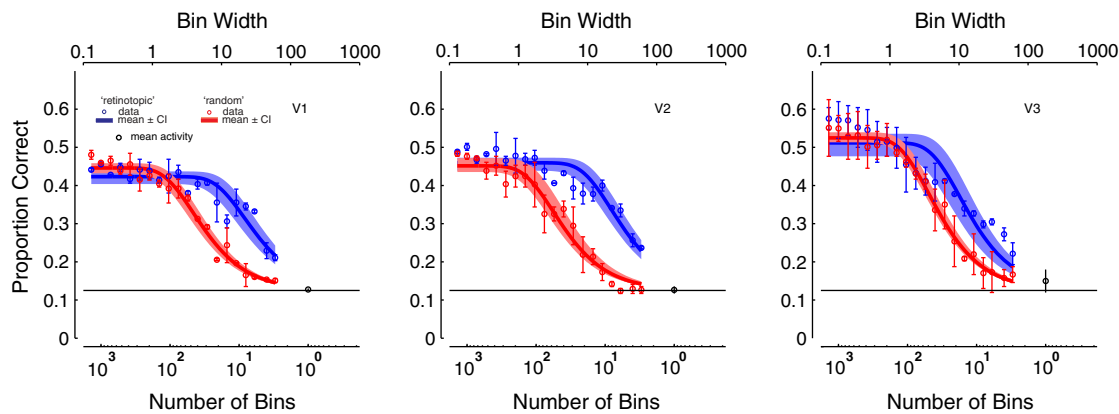
Fig. 7. Control experiment using split-hemifield stimuli. Same conventions as in Fig. 5. Graph shows the effect of binning and averaging 500 voxels from V1–V3 based on their visual field preference. Blue symbols, binning and averaging by preferred visual field angle with increasing bin size (note log scale). Red symbols, equivalent averaging but with shuffled preferred visual field angle labels, therefore providing a randomized control. Lines, fit to data using Eq. (5). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

averaging voxels within an ROI based on their preferred visual field angle. This process removes any of the high-frequency information, such as the slight biases for different directions of motion that each voxel would obtain from the distribution of motion selective columns. Because our sampling resolution (1.5 mm isotropic) was larger than that needed to resolve the expected pattern of columnar organization, the spatial distribution of any small biases would appear random (Kriegeskorte et al., 2010). While averaging together responses from different voxels, small biases should cancel out while any information that is correlated with retinotopy (visual field angle preference) should be preserved. There are several possible mechanisms that could explain the presence of this coarse-scale information, including biases for cardinal and radial directions of motion (Schluppeck and Engel, 2003; Raemaekers et al., 2009). A bias for cardinal motion directions is unlikely: when the mean amplitude of the ROI was used for classification, classification accuracy was at chance. To check for a radial bias, we examined the largest weight that each voxel or supervoxel contributed to each direction detector, and observed how these weights varied depending on the polar angle of visual field the voxel or super-voxel represented, based on it's retinotopic phase or the range of phases that super-voxels bin encompassed (Fig. 8). When the voxel weights were displayed on the cortical surface, no pattern could be observed, but as the voxels were binned over increasing ranges of polar angle to form the super-voxels, biases for directions of motion radial to the supervoxel bin's mean polar angle began to emerge. This is consistent with the previous findings that any radial biases are quite small, and only directly observable using a combination of methods and sufficient amounts of data (Freeman et al., 2011; Raemaekers et al., 2009), or when the data have been blurred considerably (Swisher et al., 2010).

We found that averaging voxels together based on their visual field preference preserved classification performance compared to random averaging of the equivalent number of voxels. By fitting curves to the data, we were able to calculate 2 parameters for each ROI/binning method combination, with the A parameter representing the asymptote and the B parameter representing the exponential growth constant, a change which represents a shift of the curve along the log-scale x axis in Fig. 5. The lower the B value, the lower the number of bins required for a given level of classification; a smaller number of bins correspond to a greater number of voxels averaged. The B parameter was significantly lower (as measured by bootstrapping) for the retinotopic-binning method compared to the random-binning method in V1–V3, indicating that a smaller number of bins were required for successful classification in this condition. Averaging voxels together in a way that preserves the signal at the scale of the retinotopic map allows significantly improved performance compared to averaging together the equivalent number of

voxels at random, indicating that there is information at this spatial scale sufficient for successful classification. It should be noted that the fitted curves do not completely capture the fine details in the results for the retinotopic averaging condition. For example, the curve for retinotopically averaged data shows a small dip that is not exactly matched by the fitted curve and may hint that retinotopic averaging has effects at more than one scale. This dip in the data could perhaps indicate that some information at the fine spatial scale exploited by the classifier is removed by the retinotopic averaging process. However, the dominant pattern in the data which shows the advantage of the retinotopic averaging condition over the random averaging condition does seem to indicate that a signal at a scale closer to the retinotopy than the direction-selective columns is sufficient for successful classification.

Our results are comparable to the findings of Freeman et al. (2011), who found almost unaffected performance even when bin width was increased to $\approx 60°$ of visual field angle (see their Fig. 5a). However, for both our averaging schemes, 'retinotopic' and 'random', classification performance dropped sharply for the largest bin sizes. There are several possible reasons for this difference. Firstly, we used blocks of motion-defined stimuli rather than slowly rotating gratings and secondly, the phase values we used for binning voxels for preferred visual field angle were based on 2 within-session scans, rather than a separate full-session of retinotopy, which may have added variability to our voxelwise visual field measurements. It may be the case that a radial bias exists, but is not strong enough to allow classification with a small number of bins, or it may be the case that the classifier is not driven by a 'radial' bias, but some other coarse scale signal. Nevertheless, it is still the case that a signal exists that is preserved when the averaging process takes into account the underlying retinotopy (see also Fig. 8) and not preserved when voxels are averaged at random.

A slight bias for cardinal directions of motion has been reported Schluppeck and Engel (2003), and recent studies suggest biases for radial directions of motion in human visual cortex (Clifford et al., 2009; Raemaekers et al., 2009; Sun et al., 2012). Given that Freeman et al. (2011) demonstrated that a radial orientation bias was both *necessary* and *sufficient* for classification of orientation in V1, an analogous radial bias for direction of motion is a likely candidate signal underlying the motion classification. Another intriguing possibility is that orientation signals in the form of motion 'streaks' have been suggested as the signal driving classification of non-opponent motion (Apthorp et al., 2010). If this were the case, a radial bias for orientations could be the mechanism underlying both orientation and (non-opposing) motion classification.

Interestingly, Raemaekers et al. (2009) reported a radial bias in early visual cortex, but not for MT+/V5. In our data, there is no strong evidence of an advantage for retinotopic averaging over random averaging in MT+/V5 suggesting a weakened or absent contribution
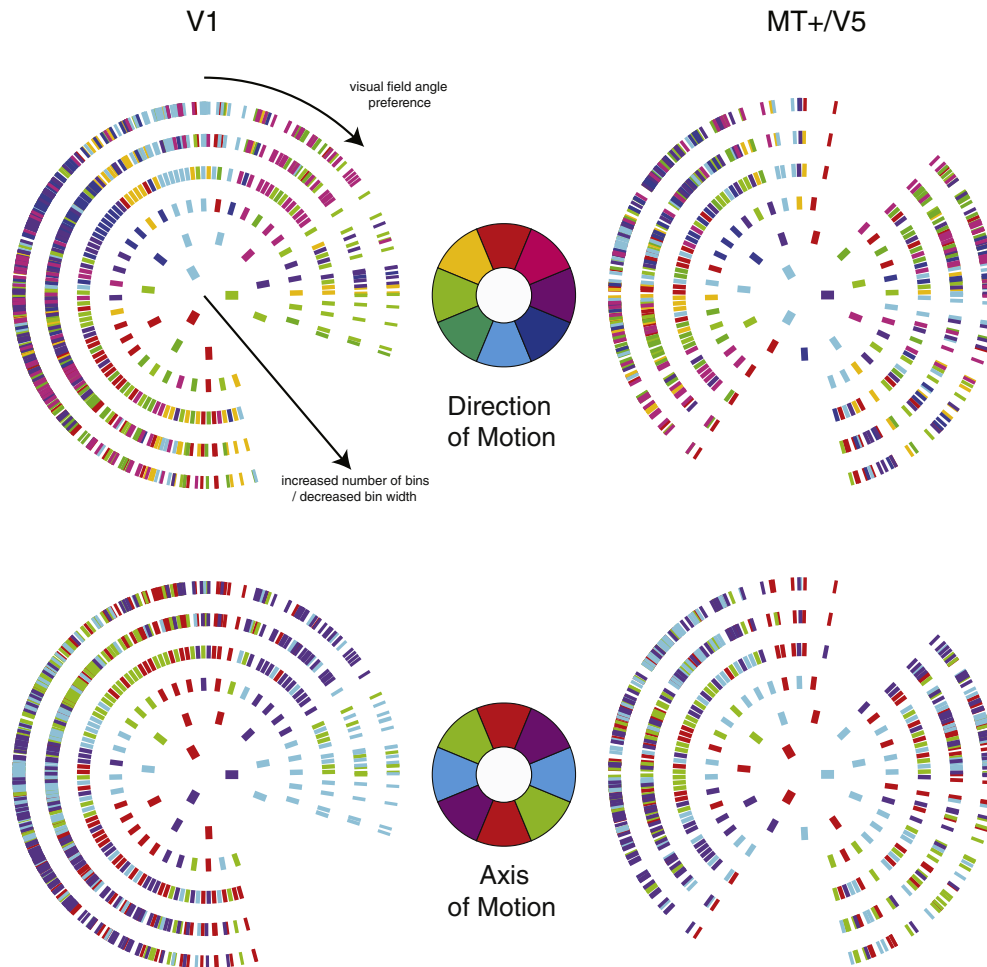
Fig. 8. Largest weights (colors) for 'direction of motion' or 'axis of motion' detectors as a function of visual field preference for (super-)voxels, based on using 500 voxels for each ROI. Left column shows data for V1; right column for hMT+. Top, analysis considering 'direction of motion' (0–360°); bottom, 'axis of motion' (0–180°). Polar plots show weights for (super-)voxels according to their phase value from the retinotopic localizer for a example subject. Angle (theta) for each symbol, indicates the phase value from the retinotopic localizer (mean across component voxels making up super-voxels). Eccentricity of each ring in the plot (R) corresponds to different levels of binning: from outer ring (no binning) toward the center of the plot (binning with large phase bins). Color coding shows which direction detector for that (super-)voxel has the largest weighting. As the bin widths increase (toward center), a bias for radial directions of motion becomes apparent. The analysis for axis of motion reveals a stronger radial bias (cf bottom panels versus top panels). The gaps in the plot which appear as un-represented visual field angles, are due to voxels with small coherence values which were therefore not included in the 500 voxels per ROI.

from radial bias in driving classification performance in this region. However, there are some limitations to this interpretation. Comparing results across different visual cortical areas is problematic. While neurons in MT +/V5 are generally more selective for direction of motion than those in V1 (as in the macaque), differences in the columnar organization or the amplitude of the BOLD response may lead to lower classification accuracies. The ability to decode stimulus properties from signals in a given area may (but not need) indicate selectivity to those properties in a given area. Conversely, however, a failure of decoding cannot be taken as evidence for a lack of selectivity (Bartels et al., 2008).

One possible mechanism for inducing a retinotopic signal that could lead to artifactual classification is a systematic relationship between different stimulus directions and eye movements (Kamitani and Tong, 2006). Although we had subjects perform a demanding contrast discrimination task at fixation during our experiment scanning, we did not monitor their eye movements due to technical limitations of our high-field MRI setup. We could therefore not be sure that subjects had not moved their eyes during stimulus presentations. To control for this possible artifact, we collected additional data and performed a follow-up analysis on 2 subjects. In these experiments, we used split hemifield stimuli, i.e. separate fields of random dots in each visual hemifield, moving in different randomly drawn directions. This approach has been used previously to minimize possible effects of eye movements (Swisher

et al., 2010). In that control experiment, classification performance was preserved in V1. Also, binning and averaging voxels by visual field preference (as in the main experiment) still showed a benefit compared to random averaging (see data for V1–V3 in Fig. 7). Therefore, the effects we describe here are unlikely to be due to any information afforded by eye movements correlated to the stimuli.

Although we have shown that a coarse-scale signal is sufficient for above-chance classification of motion direction, we cannot show that such a retinotopically-organized radial bias is necessary for successful classification, that it is the only information used, or that classification results obtained at lower field strengths work in the same way. A high-field fMRI study in which fMRI signals were used to decode ocular dominance in V1 (Shmuel et al., 2010) suggested that at coarser scales of analysis, vasculature and coarse-scale information not only dominates, but also found contribution from high-frequency gray-matter signals. Because there are known differences in the characteristics of the BOLD signal at different field strengths (Gati et al., 1997; Ogawa et al., 1998; Yacoub et al., 2001), it is difficult to draw strong conclusions about the signals that drive classification at 3 T from our data. A recent hypothesis has also emerged about how the voxel sampling of neural activity may be described as a complex spatiotemporal filtering via the underlying vasculature (Kriegeskorte et al., 2010). This model describes a more general framework of how voxels average the activity of the underlying

cortex. One of the particulars of this description is that hyper-band information (i.e. information beyond the sampling resolution) can alias into the voxel-sampled image of the cortex which may explain the seemingly paradoxical hyperacuity even in blurred or downsampled images. While our study indicates that a coarse-scale map is available for classification, it does not preclude the possibility that hyper-band information may be available via other mechanisms.

It should also be noted that, as in the case of radially balanced orientation stimuli (Mannion et al., 2009), classification of radially balanced, rotational stimuli is also possible (Kamitani and Tong, 2006). A bias for radial directions of motion would not be expected to contribute to classification of this kind of stimuli, so a radial bias may only contribute to radially-unbalanced stimuli.

In conclusion, we found that a coarse-scale, retinotopically organized signal is sufficient to allow the 'decoding' of the direction motion of a cloud of drifting dots at 7 T. This adds to a range of evidence that multivariate analysis techniques may not always be based on high-resolution information (such as columnar distribution) but may in fact based on coarser scale signals such as global biases or vasculature (Freeman et al., 2011; Op de Beeck, 2010; Shmuel et al., 2010), and that care should be taken in drawing conclusions from successful multivariate classification in a given area.

Acknowledgments

References

Adams, D., Sincich, L., Horton, J., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. J. Neurosci. 27, 10391.

Apthorp, D., Bahrami, B., Schwarzkopf, D., Kaul, C., Alais, D., Rees, G., 2010. Motion streaks in the brain: an fMRI study. Perception 39, 137.

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

Boynton, G., 2005. Imaging orientation selectivity: decoding conscious perception in V1. Nat. Neurosci. 8, 541–542.

Cheng, K., Waggoner, R., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. Neuron 32, 359–374.

Clifford, C., Mannion, D., McDonald, J., 2009. Radial biases in the processing of motion and motion-defined contours by human visual cortex. J. Neurophysiol. 102, 2974.

Clifford, C., Mannion, D., Seymour, K., McDonald, J., Bartels, A., 2011. Are coarse-scale orientation maps really necessary for orientation decoding? [eletter]. J. Neurosci.

Efron, B., Tibshirani, R., 1993. An Introduction to the Bootstrap. Chapman & Hall. http://www.jneurosci.org/content/31/13/4792/reply.

Freeman, J., Brouwer, G., Heeger, D., Merriam, E., 2011. Orientation decoding depends on maps, not columns. J. Neurosci. 31, 4792–4804.

Fukuda, M., Wang, P., Moon, C., Tanifuji, M., Kim, S., 2006. Spatial specificity of the enhanced dip inherently induced by prolonged oxygen consumption in cat visual cortex: implication for columnar resolution functional MRI. NeuroImage 30, 70–87.

Gardner, J., 2010. Is cortical vasculature functionally organized? NeuroImage 49, 1953–1956.

Gardner, J., Sun, P., Tanaka, K., Heeger, D., Cheng, K., 2006. Classification analysis with high spatial resolution fMRI reveals large draining veins with orientation specific responses. Society for Neuroscience Abstracts, pp. 614–640.

Gati, J., Menon, R., Ugurbil, K., Rutt, B., 1997. Experimental determination of the bold field strength dependence in vessels and tissue. Magn. Reson. Med. 38, 296–302.

Harmer, J., Sanchez-Panchuelo, R., Bowtell, R., Francis, S., 2012. Spatial location and strength of bold activation in high-spatial-resolution fMRI of the motor cortex: a comparison of spin echo and gradient echo fMRI at 7 t. NMR Biomed. 25, 717–725.

Haynes, J., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nat. Neurosci. 8, 686–691.

Haynes, J., Rees, G., 2006. Decoding mental states from brain activity in humans. Nat. Rev. Neurosci. 7, 523–534.

Hubel, D., Wiesel, T., 1963. Shape and arrangement of columns in cat's striate cortex. J. Physiol. 165, 559.

Hubel, D., Wiesel, T., 1969. Anatomical demonstration of columns in the monkey striate cortex. Nature 221, 747–750.

Huk, A., Dougherty, R., Heeger, D., 2002. Retinotopy and functional subdivision of human areas MT and MST. J. Neurosci. 22, 7195.

Kahnt, T., Heinzle, J., Park, S., Haynes, J., 2010. The neural code of reward anticipation in human orbitofrontal cortex. Proc. Natl. Acad. Sci. 107, 6010.

Kamitani, Y., Sawahata, Y., 2010. Spatial smoothing hurts localization but not information: pitfalls for brain mappers. NeuroImage 49, 1949–1952.

Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nat. Neurosci. 8, 679–685.

Kamitani, Y., Tong, F., 2006. Decoding seen and attended motion directions from activity in the human visual cortex. Curr. Biol. 16, 1096–1102.

Klein, A., Andersson, J., Ardekani, B., Ashburner, J., Avants, B., Chiang, M., Christensen, G., Collins, D., Gee, J., Hellier, P., 2009. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. NeuroImage 46, 786–802.

Kriegeskorte, N., Cusack, R., Bandettini, P., 2010. How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatiotemporal filter? NeuroImage 49, 1965–1976.

Larsson, J., Heeger, D., 2006. Two retinotopic visual areas in human lateral occipital cortex. J. Neurosci. 26, 13128.

Logothetis, N., Wandell, B., 2004. Interpreting the BOLD signal. Annu. Rev. Physiol. 66, 735–769.

Malonek, D., Tootell, R., Grinvald, A., 1994. Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT. Proc. R. Soc. London, Ser. B 258, 109.

Mannion, D., McDonald, J., Clifford, C., 2009. Discrimination of the local orientation structure of spiral glass patterns early in human visual cortex. NeuroImage 46, 511–515.

Mountcastle, V., Davies, P., Berman, A., 1957. Response properties of neurons of cat's somatic sensory cortex to peripheral stimuli. J. Neurophysiol. 20, 374–407.

Nestares, O., Heeger, D., 2000. Robust multiresolution alignment of MRI brain volumes. Magn. Reson. Med. 43, 705–715.

Ogawa, S., Lee, T., Kay, A., Tank, D., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. 87, 9868–9872.

Ogawa, S., Menon, R.S., Kim, S.G., Ugurbil, K., 1998. On the characteristics of functional magnetic resonance imaging of the brain. Annu. Rev. Biophys. Biomol. Struct. 27, 447–474 (http://www.annualreviews.org/doi/pdf/10.1146/annurev.biophys.27.1.447).

Op de Beeck, H., 2010. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? NeuroImage 49, 1943–1948.

Op de Beeck, H., Baker, C., DiCarlo, J., Kanwisher, N., 2006. Discrimination training alters object representations in human extrastriate cortex. J. Neurosci. 26, 13025.

Poole, M., Bowtell, R., 2008. Volume parcellation for improved dynamic shimming. Magn. Reson. Mater. Phys., Biol. Med. 21, 31–40 http://dx.doi.org/10.1007/s10334-007-0102-2.

Pruessmann, K., Weiger, M., Scheidegger, M., Boesiger, P., et al., 1999. Sense: sensitivity encoding for fast MRI. Magn. Reson. Med. 42, 952–962.

Raemaekers, M., Lankheet, M., Moorman, S., Kourtzi, Z., van Wezel, R., 2009. Directional anisotropy of motion responses in retinotopic cortex. Hum. Brain Mapp. 30, 3970–3980.

Sanchez-Panchuelo, R., Francis, S., Bowtell, R., Schluppeck, D., 2010. Mapping human somatosensory cortex in individual subjects with 7T functional MRI. J. Neurophysiol. 103, 2544.

Sasaki, Y., 2007. Processing local signals into global patterns. Curr. Opin. Neurobiol. 17, 132–139.

Sasaki, Y., Rajimehr, R., Kim, B., Ekstrom, L., Vanduffel, W., Tootell, R., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. Neuron 51, 661–670.

Schluppeck, D., Engel, S.A., 2003. Oblique effect in human mt + follows pattern rather than component motion. J. Vis. 3, 282.

Seymour, K., Clifford, C., Logothetis, N., Bartels, A., 2009. The coding of color, motion, and their conjunction in the human visual cortex. Curr. Biol. 19, 177–183.

Shmuel, A., Raddatz, G., Chaimow, D., Logothetis, N., Ugurbil, K., Yacoub, E., 2007. Multi-resolution classification analysis of ocular dominance columns obtained at 7 tesla from human V1: mechanisms underlying decoding signals. Society for Neuroscience Annual Meeting.

Shmuel, A., Chaimow, D., Raddatz, G., Ugurbil, K., Yacoub, E., 2010. Mechanisms underlying decoding at 7 t: ocular dominance columns, broad structures, and macroscopic blood vessels in V1 convey information on the stimulated eye. NeuroImage 49, 1957–1964.

Smith, A., Kosillo, P., Williams, A., 2010. The confounding effect of response amplitude on MVPA performance measures. NeuroImage 56, 525–530.

Snowden, R., Treue, S., Andersen, R., 1992. The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. Exp. Brain Res. 88, 389–400.

Sun, P., Gardner, J.L., Costagli, M., Ueno, K., Waggoner, R.A., Tanaka, K., Cheng, K., 2012. Demonstration of tuning to stimulus orientation in the human visual cortex: a high-resolution fMRI study with a novel continuous and periodic stimulation paradigm. Cereb. Cortex (http://cercor.oxfordjournals.org/content/early/2012/05/31/cercor.bhs149.full.pdf+html).

Swisher, J., Gatenby, J., Gore, J., Wolfe, B., Moon, C., Kim, S., Tong, F., 2010. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. J. Neurosci. 30, 325.

Thompson, R., Correia, M., Cusack, R., 2010. Vascular contributions to pattern analysis: comparing gradient and spin echo fMRI at 3T. NeuroImage 56, 643–650.

Wilson, J.L., Jenkinson, M., Jezzard, P., 2002. Optimization of static field homogeneity in human brain using diamagnetic passive shims. Magn. Reson. Med. 48, 906–914.

Yacoub, E., Shmuel, A., Pfeuffer, J., Van De Moortele, P., Adriany, G., Andersen, P., Vaughan, J., Merkle, H., Ugurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 tesla. Magn. Reson. Med. 45, 588–594.

Yacoub, E., Shmuel, A., Logothetis, N., Ugurbil, K., 2007. Robust detection of ocular dominance columns in humans using Hahn spin echo BOLD functional MRI at 7 tesla. NeuroImage 37, 1161–1177.

Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. Proc. Natl. Acad. Sci. 105, 10607.

Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P.F., Feinberg, D., Adriany, G., Chaimow, D., Shmuel, A., Ugurbil, K., Yacoub, E., 2011. Mapping the organization of axis of motion selective features in human area mt using high-field fMRI. PLoS One 6, e28716.