

Published in final edited form as:

Neuroimage. 2011 May 15; 56(2): 627–642. doi:10.1016/j.neuroimage.2010.09.037.

Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns

Denis Chaimow^a, Essa Yacoub^b, Kamil Ugurbil^b, and Amir Shmuel^{b,c,*}

^aMax-Planck Institute for Biological Cybernetics, Spemannstr. 41, Tuebingen, Germany

^bCenter for MR Research, University of Minnesota Medical School, 2021 6th St. SE, Minneapolis, MN, USA

^cMontreal Neurological Institute, Departments of Neurology and Neurosurgery and Biomedical Engineering, McGill University, Montreal, QC, Canada

Abstract

Multivariate machine learning algorithms applied to human functional MRI (fMRI) data can decode information conveyed by cortical columns, despite the voxel-size being large relative to the width of columns. Several mechanisms have been proposed to underlie decoding of stimulus orientation or the stimulated eye. These include: (I) aliasing of high spatial-frequency components, including the main frequency component of the columnar organization, (II) contributions from local irregularities in the columnar organization, (III) contributions from large-scale non-columnar organizations, (IV) functionally selective veins with biased draining regions, and (V) complex spatio-temporal filtering of neuronal activity by fMRI voxels. Here we sought to assess the plausibility of two of the suggested mechanisms: (I) aliasing and (II) local irregularities, using a naive model of BOLD as blurring and MRI voxel sampling.

To this end, we formulated a mathematical model that encompasses both the processes of imaging ocular dominance (OD) columns and the subsequent linear classification analysis. Through numerical simulations of the model, we evaluated the distribution of functional differential contrasts that can be expected when considering the pattern of cortical columns, the hemodynamic point spread function, the voxel size, and the noise. We found that with data acquisition parameters used at 3 Tesla, sub-voxel supra-Nyquist frequencies, including frequencies near the main frequency of the OD organization (0.5 cycles per mm), cannot contribute to the differential contrast. The differential functional contrast of local origin is dominated by low-amplitude contributions from low frequencies, associated with irregularities of the cortical pattern. Realizations of the model with parameters that reflected best-case scenario and the reported BOLD point-spread at 3 Tesla (3.5 mm) predicted decoding performances lower than those that have been previously obtained at this magnetic field strength. We conclude that low frequency components that underlie local irregularities in the columnar organization are likely to play a role in decoding. We further expect that fMRI-based decoding relies, in part, on signal contributions from large-scale, non-columnar functional organizations, and from complex spatio-temporal filtering of neuronal activity by fMRI voxels, involving biased venous responses. Our model can potentially be used for evaluating and optimizing data-acquisition parameters for decoding information conveyed by cortical columns.

Introduction

Recent studies have demonstrated that multivariate machine learning algorithms can decode visual stimuli from functional MRI (fMRI) data (Haxby et al., 2001; Kamitani and Tong, 2005; Haynes and Rees, 2005a). Using gradient-echo (GE) blood oxygenation level dependent (BOLD) fMRI data obtained at 3T, these algorithms decoded information thought to be mediated by cortical columns. This result seems to be surprising given the large size of the voxels ($3 \times 3 \times 3 \text{ mm}^3$) relative to the mean cycle length of columns (2 mm or less for ocular dominance columns (ODCs) and orientation columns in humans). This result is even more surprising considering the relatively wide point-spread function of GE BOLD fMRI signals at 3T (~3.5 mm; Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007).

The mechanism by which low-resolution imaging decodes information represented at a fine scale relative to the voxel size is not clear. In the following we mention five alternative mechanisms that have been hypothesized (we believe the terms we use are appropriate for describing these mechanisms, although the original publications may have used different terms). (I) Aliasing of high spatial-frequency components of the columnar organization by the large voxels has been suggested (Boynton, 2005). The “aliasing” mechanism, also termed the “hyperacuity” mechanism (Op de Beeck, 2010), involves components of the columnar organization with frequencies higher than the Nyquist frequency of the MRI sampling process, that were thought to contribute to the sampled voxels. (II) It was hypothesized that random, local variations and irregularities in the functional organization contribute to decoding (Kamitani and Tong, 2005, 2006; Haynes and Rees, 2006; Kriegeskorte and Bandettini, 2007). The argument is that due to the irregular underlying columnar pattern, each voxel overlaps columns with different preferences unequally, resulting in biases towards specific preferences. If irregularities exist, the columnar organization cannot consist of one single spatial-cortical frequency: it is likely to involve a distribution of frequencies, including frequencies lower than the main frequency of the organization (Rojer and Schwartz, 1990). Note that these components, with frequencies lower than the main frequency of the columnar organization, may be present even if the overall preferences represented by the columns are distributed equally across the investigated cortical region. Indeed, Swisher et al. (2010) and Shmuel et al. (2010) demonstrated contributions from low frequency components of the functional columnar organization to decoding. (III) Very low spatial frequencies, reflecting large-scale components of the organization were proposed to play a role too (Op de Beeck, 2010). These include the oblique and radial effects (Sasaki et al., 2006; Furmanski and Engel, 2000) associated with the representation of orientation, and the higher amplitude response to stimulation of the contra-lateral eye associated with the representation of ODCs (Tychsen and Burkhalter, 1997).

Alternatively (IV), draining regions that cover cortical maps and columns non-homogeneously may cause selective responses of their corresponding blood vessels (Kamitani and Tong, 2005, 2006; Gardner et al., 2006; Kriegeskorte and Bandettini, 2007; Shmuel et al., 2010). In this scenario, henceforth termed “biased draining regions,” selective signals from macroscopic blood vessels can be captured by large voxels; therefore, they can contribute to the decoding of stimuli encoded at the resolution of cortical columns. Evidence in support of this phenomenon was provided by Gardner et al. (2006) and Shmuel et al. (2010). Lastly (V), Kriegeskorte et al. (2010) introduced a model in which fMRI voxels sample neuronal activity as complex spatio-temporal filters. These authors described how such a model can account for representation of high-frequency components of the cortical maps by the sampled voxels, and for decoding of information conveyed by cortical columns. Note that the functionally selective responses of veins demonstrated by Gardner et al. (2006) and Shmuel et al. (2010) constitute a specific scenario of the more general concept of

interpreting fMRI sampling as spatio-temporal filtering of neuronal activity. Irrespective of the exact mechanisms, all five proposed mechanisms mentioned above reflect neuronal selectivity. Even though the exact spatial information is lost, the signals are expected to originate at the neuronal level.

In order to assess the plausibility of the aliasing (hyperacuity) mechanism and the contributions of low-frequency components of the columnar organization, it is necessary to quantify their respective expected biases and the corresponding classification performances. In this current study, we aimed to create a model that can be used for studying the mechanisms underlying fMRI-based decoding of features represented in cortical columns. In addition, we sought to evaluate the distribution of responses, differential contrasts, and classification performance that can be expected when using large voxels under realistic conditions. The realization of these objectives can support the planning of studies involving decoding.

To address these objectives, we developed a model to image a region with a fine-scale organization of cortical columns, followed by decoding. The model first creates a realistic pattern of ODCs organization. Next, the model addresses the responses of neuronal assemblies within this organization to specific stimulus conditions. The spatial features of the BOLD response are then considered, followed by modeling the process of voxel sampling. In the subsequent decoding portion of the model, we show that classification performance can be predicted from quantities obtained within the model. Specifically, decoding performance is fully characterized by the distribution of differential contrasts, the noise level, and the number of analyzed voxels.

Using our model, we demonstrate the dependence of differential contrast and classification performance on parameters of the studied functional organization including the sharpness and irregularity of the cortical map. We further evaluate the dependence of differential contrast and classification performance on parameters of the data acquisition process, including the BOLD point spread function (PSF) and voxel size, and the number of voxels. Lastly, we compare results obtained by the model to those obtained in decoding studies.

Methods

Overview

We developed a model that enables the prediction of classification performance as a function of several parameters of interest. The model is based on linear classification. Linear classification has been used in previous fMRI-based decoding studies in the form of linear discriminant analysis (LDA) (e.g. Haynes and Rees, 2005a,b) or linear support vector machines (Kamitani and Tong, 2005). Here we briefly describe the structure of the model, and the different stages it involves. Variables and parameters of the model are presented in Table 1. All mathematical derivations and details of the model can be found in the Appendix.

Imaging model

The goal of the imaging model was to model the distribution of voxel-wise differential responses (Fig. 1). We use the term “contrast range” to describe how large the expected differential responses are on average. To quantify the contrast range, we used the standard deviation of the distribution of single voxel differential responses. This is a measure of how much contrast between stimulation conditions can be expected. It will be used later on in calculating the expected classification performance.

Realistic patterns of ocular dominance columns

The spatial pattern of cortical columns was modeled by spatial filtering of 2D Gaussian white noise (Rojer and Schwartz, 1990). The structure of the resulting pattern depends on the shape of the filter. An anisotropic band-pass filter was used, which yields realistic patterns of elongated ODCs. The ODC filter was parameterized by the main pattern frequency ρ , which determines the width of the columns. ρ was set to 0.5 cycles/mm corresponding to a column width of 1 mm (Yacoub et al., 2007). Two additional parameters, irregularity (δ) and branchiness (ϵ), were employed in order to control the level of pattern irregularities, orthogonal and parallel respectively, to the ODC columns. When not otherwise noted, parameters δ and ϵ were set to 0.3 cycles/mm and 0.4 cycles/mm, respectively. These numbers were based on the analysis of macaque ODC maps (Rojer and Schwartz, 1990) which we scaled to fit the spatial frequency of ODCs in humans. Here we assumed that human ODC maps have a very similar structure, only scaled in space according to Horton et al. (1990), Adams et al. (2007) and Yacoub et al. (2007).

The filter was normalized so that the output had a standard deviation of 1. The filtered noise was passed through a sigmoidal non-linearity with parameter α that controlled the sharpness of the transitions from one column to the adjacent columns (Rojer and Schwartz, 1990). When not otherwise noted, we used $\alpha = 4$, resulting in a moderate level of sharpness.

Neuronal response

The neuronal response was defined on an arbitrary scale from 0 to 1, where 0 stands for no response and 1 represents a maximal response. The two stimulation conditions were assumed to produce opposing patterns of neuronal responses proportional to their respective preferences as defined by the ODC map.

BOLD response

The spatial characteristics of the BOLD response were modeled as a convolution of the neuronal response with a two-dimensional BOLD point spread function (Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007). The width of the convolution kernel was parameterized using the full width at half maximum (FWHM) measure. A second parameter, β , stood for the absolute scaling of the kernel. Its role was to generate realistic response amplitude values such that the maximal neuronal response results in a steady state BOLD response of amplitude β . Following a realistic best-case scenario approach β was chosen to be 5% (Krüger et al., 2001; Boynton et al., 1996; our own experience). We assumed that residual head motion was comparable between studies that reported PSF at 3 Tesla and decoding studies. Therefore, rather than directly accounting for residual motion, our model considers residual head motion implicitly through the convolution with the BOLD PSF.

MR imaging process and voxel sampling

The MR imaging process was modeled as sampling the k -space representation of the BOLD response patterns at discrete steps determined by the field of view and the voxel width, and a subsequent discrete Fourier transform (Haacke et al., 1999). The responses to the two conditions were subtracted to result in a voxelized differential response pattern.

Prediction of classification performance

We analyzed classification performance of a linear discriminant classifier. Hypothetical fMRI responses (percent change relative to baseline) of n voxels were considered as n -dimensional vectors associated with one of two stimulation conditions. We assumed that the voxels responded with amplitudes sampled from two multivariate normal distributions, each of which was associated with one stimulus condition. Each distribution was characterized by

its multivariate mean, reflecting the expected (in the sense of statistical expectation) voxel-wise relative responses, and by its covariance matrix representing all sources of noise. The distributions of noise associated with different stimulation conditions and in different voxels were all assumed to be equal, and independent of each other.

Expected classification performance was estimated by calculating the expected fraction of vectors classified correctly as being associated with the stimulus condition of their origin. A linear classifier partitions the feature space into two regions separated by a decision boundary. The fraction of correctly classified vectors from one stimulation condition equals the integral of the corresponding probability density function over the feature space region associated with that condition.

Differential responses and contrast range

The expected multivariate difference of voxel-wise responses \vec{d} determines the position of the decision boundary relative to the two distributions. \vec{d} was approximated using the standard deviation of the expected distribution of single voxel differential responses (referred to as “contrast range”), and the number of voxels. Eq. 1 in the Appendix shows classification performance as a function of contrast range, the number of voxels, and the noise level.

Overall contrast-to-noise ratio

Contrast range, the number of voxels, and noise level were combined into one measure of overall contrast-to-noise ratio (OCNR). OCNR is proportional to contrast range and the square root of number of voxels. It is inversely proportional to the noise level. Overall contrast-to-noise ratio completely determines the classification performance (Eq. 2 in the Appendix) and is directly related to the Fisher criterion in linear discriminant analysis.

Noise

The relative noise level σ is the standard deviation of all signal changes not related to an external stimulus, relative to baseline. It is the inverse of time-course signal to noise ratio (tSNR). Noise dependence on voxel size was modeled using the following formula from Triantafyllou et al. (2005).

$$\frac{1}{\sigma} = tSNR(V) = \frac{\kappa \times V}{\sqrt{1 + \lambda^2 \times \kappa^2 \times V^2}},$$

where V is the voxel volume, λ is a field and scanner independent constant governing the relation between temporal SNR and image SNR, and κ is the proportionality constant between volume and image SNR that is field strength and hardware dependent. Both constants were estimated by fitting this equation to the data given in Table 3 of Triantafyllou et al. (2005) using a Trust-Region non-linear least squares algorithm in MATLAB (The Mathworks, Inc., 2007). Based on the fitting, we set $\lambda = 0.01297$ and $\kappa = 6.641$. Note that the tSNR values in Table 3 of Triantafyllou et al. (2005) were obtained using TR = 5.4 s. In section C of the Appendix we show how the modeled tSNR values from Triantafyllou et al. (2005) were modified to tSNR values expected with different TRs.

Model implementation

We implemented the model using numerical simulations in MATLAB (The MathWorks Inc., Natick, MA, USA). We simulated a square area with a field of view between 48 mm \times 48 mm and 192 mm \times 192 mm, depending on the specific simulation. The latter relatively

large field of view was necessary for obtaining a high enough k -space resolution when studying the contributions of different spatial frequencies. The area was divided into 1024×1024 evenly spaced points.

We ran numerical simulations of the model components described above (Fig. 2) while varying different parameters. Contrast range was computed by calculating the standard deviation over a simulated differential voxel pattern response. Contrast range values obtained in multiple runs were averaged in order to increase the robustness of the results. Single frequency contributions to contrast range were computed by restricting the spatial frequency representation of the ODC pattern to a small range of absolute frequencies Δk around the frequency under investigation. The obtained contrast range was divided by Δk resulting in an estimate of contrast range per frequency unit.

Results

We aimed to analyze the mechanisms underlying decoding of information represented in a fine-scale functional organization using large voxels and a relatively wide point spread function. Classification performance depends on the differential contrast between stimulation conditions, the number of voxels, and the relative noise level (see Eq. 1 in the Appendix). In this section we briefly introduce the model, and demonstrate its function by means of a numerical realization. We then study how the differential contrast depends on BOLD point spread and voxel size. Next, we evaluate the frequency components of the neuronal ODC organization that are reflected in fMRI voxels, and therefore potentially contribute to decoding. We demonstrate the effects of the BOLD PSF and the MR imaging process on these frequency components. We demonstrate how voxel-size specific noise, functional contrast, and number of voxels combine to a measure of overall CNR that determines classification rate. In the last section, we evaluate the dependence of classification performance on parameters of the functional columnar organization.

Contrast range

The model—In order to quantify the functional contrast at the single voxel level, we developed a model of imaging cortical columns, specifically for ODCs. The end result of the model is a distribution of single voxel differential responses. The modeled voxel differential responses follow a distribution with zero mean. The standard deviation of the distribution of differential responses reflects the dispersion of condition specific contrasts (here, contrast between responses to left and right eye stimulation) present in a set of imaged voxels. The larger this standard deviation, the larger the contrast values that exist in the specific distribution. Through the rest of the manuscript, the standard deviation of the distribution of differential functional contrast will be referred to as the “contrast range.”

Numerical realization—Fig. 2 presents a numerical realization of the model using a BOLD point spread with FWHM of 3.5 mm and a voxel size of 3 mm. It is evident that the results of both the BOLD response stage and the subsequent voxel sampling show condition specific patterns. Nonetheless, these patterns do not directly reflect the structure of the ODC pattern, which is dominated by higher spatial frequencies. In addition, the functional contrasts following the BOLD response and voxel sampling stages are very small.

Dependence of contrast range on BOLD point spread and voxel size—We simulated differential response patterns while varying voxel width and BOLD point spread width. We computed the contrast range from these patterns and plotted the contrast range as a function of BOLD PSF width and voxel width (Fig. 3). The contrast range decreased with increasing width of the BOLD PSF (Fig. 3A and B) and with increasing voxel-width (Fig.

3C and D). Qualitatively, the effects of BOLD point-spread width and of voxel width are similar, as reflected in the approximately symmetric pattern in Fig. 3E. Assuming infinitesimally small voxels, with a BOLD point spread FWHM of 3.5 mm the contrast range drops to 0.09%, which is ~2% of its expected value (4%) if there was no spread (Fig. 3A). The effect of voxel sampling is similar. Assuming no effect of BOLD point-spread, at a voxel width of 3 mm the contrast range drops to 0.16%, ~4% of its value (4%) using infinitesimally small voxels (Fig. 3C). With narrow BOLD PSF or with small voxels, changes in the other parameter (voxel size or BOLD PSF, respectively) have substantial effects on contrast range (Fig. 3A, C, E). In contrast, for wide BOLD point spreads or large voxels, the effect of varying the other parameter is not as pronounced (Fig. 3B, D, E). At a point spread of 3.5 mm, the contrast range is almost independent of voxel size (Fig. 3D). Taken together, BOLD point spread with FWHM of 3.5 mm and voxel width of 3 mm, which are typical to BOLD imaging at 3T, reduced the contrast range to 0.08%, ~2% of its original value (Fig. 3B, D, and E).

Frequency contributions to contrast range and aliasing—We have shown that the contrast range is considerably reduced by the BOLD point spread and sampling with large voxels. We next sought to estimate the relative contributions of different frequency components of the ODC organization to the contrast available for decoding (Figs. 4 and 5).

To this end, we first considered the effect of the MRI data-acquisition and reconstruction processes. MRI voxels are often thought of as taking the shape of a rect-function in the image space (Fig. 4A, in cyan). However, MRI is not equivalent to integrating the signal over the area of a rect-function-like voxel. Instead, MRI samples the k -space at discrete steps up to the Nyquist frequency, which is the inverse of twice the voxel width. This is equivalent to integrating the signal in the image space as weighted by a sinc-function (Fig. 4A, in blue). In other words, a more precise model of a voxel in image space follows a sinc-function (Haacke et al., 1999; See also here, Section B.4 of the Appendix). Fig. 4B presents the frequency-space representation of a 3 mm wide rect-voxel, a 3 mm wide sinc-voxel, and the frequency content of a realistic neuronal ODC organization. To obtain the latter, we calculated the contributions of different spatial frequency components to the contrast range by decomposing the ODC map into its spatial frequency components.

Fig. 4C presents the frequency components of the ODC organization that remain following the voxel sampling process for 3 mm rect-voxels (cyan) and 3 mm sinc-voxels (blue), assuming infinitesimally small BOLD PSF. Rect-voxel sampling reduces the contrast range across all pattern frequencies (Fig. 4C, cyan curve). It reduces the contributions to contrast range of multiples of the sampling frequency (0.33 cycles/mm, 0.66 cycles/mm) more than it does for other frequencies. However, its effect on the relative contributions of frequency components lower or higher than the main frequency of the organization (0.5 cycles/mm) is small (compare to Fig. 4B, gray curve). The contrast depends almost entirely on frequencies that are higher than the Nyquist frequency ($f_{\text{Nyquist}} = 0.167$ cycles/mm). It includes significant contributions from frequencies near the main frequency of the ODC organization (0.5 cycles/mm).

In contrast, when using sinc-voxels the contrast beyond the Nyquist frequency drops sharply (Fig. 4C, blue curve), and it is completely eliminated beyond the Nyquist frequency that corresponds to the diagonal of the k -space ($0.167 \cdot 2$). Contributions from most of the frequency components of the ODC pattern are eliminated. All information from frequencies around the main frequency of the organization (0.5 cycle/mm) is lost. Only contributions from frequencies lower than the main frequency of the organization, that are present due to the irregularity of the ODC pattern, prevail.

Panels D and E in Fig. 4 present frequency contributions to the functional contrast as a function of varying voxel width for rect-voxels and sinc-voxels, respectively. True for both types of voxels, the contrast with origin in the main frequency of the organization decreases with increasing voxel width. However, while 3–4 mm wide rect-voxels still carry functional contrast with origin in that frequency, sampling with sinc-voxels wider than ~1.4 mm eliminates it completely.

In Fig. 4 we considered contrast contributions from various frequency components while assuming infinitesimally small BOLD PSF. Next we studied the effect of the BOLD PSF on the frequency contributions to contrast range. Fig. 5A presents the frequency representation of a Gaussian PSF with FWHM of 3.5 mm (red curve), along with the frequency representation of a sinc-voxel and the ODC organization. The BOLD point spread, even when assuming infinitesimally small voxels, acts as a strong low pass filter (Fig. 5A and B). High frequencies of the columnar pattern are filtered out almost completely. A convolution with a realistic BOLD PSF therefore shifts the distribution of the frequency components that contribute to the functional contrast towards lower frequencies (Fig. 5B).

Fig. 5C shows that convolving the neuronal response with a 3.5 mm BOLD PSF prior to MRI sampling diminishes the differences between the frequency components captured by sinc-voxels and rect-voxels. In both cases, only very low frequencies prevail.

Classification performance—Contrast range, the number of voxels and the level of noise can be combined into a single measure of overall-contrast-to-noise ratio (OCNR). Overall contrast-to-noise ratio is proportional to contrast range, the square root of the number of voxels, and the square root of the number of averaged volumes (assuming time-independent noise; see section A of the Appendix). It is inversely proportional to the noise level.

In order to calculate classification performance, we modeled noise according to Triantafyllou et al. (2005). We predicted time-course SNR (tSNR; see section C of the Appendix) for a TR of 2 s. tSNR increases with increasing voxel width (Fig. 6A). Fig. 6B presents the dependence of classification performance on the voxel size for a BOLD point spread width of 3.5 mm. We varied the in-plane width of the voxel, while holding the slice thickness constant at 3 mm and keeping a constant number of voxels. We considered voxel size dependent noise, as demonstrated in Fig. 6A. The expected classification performance using 100 voxels, each of which covering a volume of $(3 \text{ mm})^3$, resulting in a voxel volume dependent noise level of 1.5% (tSNR = 68, TR = 2 s) relative to the temporal mean of the baseline, and a BOLD PSF of 3.5 mm, was 61% (with chance level being 50%). Fig. 6C shows how classification performance depends on overall contrast-to-noise (OCNR) ratio. The logarithmic scaling of the OCNR-axis illustrates that increases in OCNR result in only moderate increases in decoding performance. When using 100 $(3 \text{ mm})^3$ voxels with a BOLD PSF of 3.5 mm, OCNR is 0.55. In this range, a factor of two improvement in OCNR results in 10% increase in decoding performance. In order to obtain 75% correct classifications, an overall CNR of 1.3 is needed. To obtain 95% correct classifications, the overall CNR needs to reach 3.3.

Sharpness of the ODC organization—The results reported thus far were based on ODC maps with a moderate, realistic sharpness ($\alpha = 4$; Fig. 7B). In order to assess the effect of smoother and sharper transitions between neighboring columns on the contrast range and classification performance, we simulated smooth ODC patterns that were not passed through a sigmoidal non-linearity (Fig. 7A) and binary ODC patterns (Fig. 7C), representing the two extreme alternatives along the pattern sharpness domain. We then computed contrast range as a function of voxel width and BOLD point spread. Qualitatively,

the resulting patterns of contrast range were similar across all three versions of ODC organizations (Fig. 7, middle row). All three patterns demonstrated approximately symmetric roles of BOLD PSF and voxel width, similar to those demonstrated in Fig. 3. Quantitatively, for large BOLD point spreads and/or large voxel sizes, ODC maps with sharper transitions produced larger contrast ranges compared to their counterparts with smoother transitions.

For 3 mm wide voxels and a 3.5-mm BOLD PSF, a binary ODC map produced a contrast range of 0.15% (70% classification performance with 100 voxels and a TR of 2 s), and an ODC map with intermediate sharpness level ($\alpha = 4$) produced a contrast range of 0.08% (61% correct classification with 100 voxels and a TR of 2 s). This is compared to 0.015% (52% classification performance with 100 voxels and a TR of 2 s) for the smooth ODC map model.

Irregularities in the ODC organization—Local variations and irregularities in cortical maps were proposed as a possible source of selective signals available for decoding. We therefore sought to study the effect of irregularities in the ODC pattern on classification performance. To this end, we varied the parameters δ (irregularity) and ϵ (branchiness) that control the level of pattern irregularities orthogonal and parallel to the axis of anisotropy of the ODC organization.

Fig. 8A demonstrates the dependence of the ODC pattern on the irregularity (δ) and branchiness (ϵ) parameters. High values of δ make the pattern of the ODC more irregular along the axis orthogonal to their major anisotropy axis, introducing wide regions in space that are biased towards one of the two eyes. In contrast, higher values of ϵ decrease local biases by interfering with the regular structure orthogonal to the columns. Panels B and C in Fig. 8 support this intuitive description. They show that classification performance increases with increases in irregularity (δ ; Fig. 8B), and decreases with increases in branchiness (ϵ ; Fig. 8C). Fig. 8D demonstrates that the effect of varying irregularity on classification performance is more pronounced than the corresponding effect of branchiness.

Discussion

Summary of the results

We developed a model of imaging cortical columns and subsequent decoding of information conveyed by them. When considered separately, the width of the BOLD point spread function and the width of the sampled voxels were found to be important factors in determining the functional contrast and classification performance (Fig. 3). BOLD PSF and the voxel width act as low-pass filters in a comparable manner. We analyzed the contributions of single spatial frequency components to the functional contrast and classification with parameters routinely used at 3 Tesla. The results ruled out contributions of aliasing of information represented at high spatial frequency corresponding to the main frequency of the columnar organization or higher frequencies (Figs. 4 and 5). Not only these high-frequency components are filtered out by the BOLD PSF, also all frequencies higher than the Nyquist frequency are discarded by the MR imaging process. Modeling MRI voxels as sinc-functions removes aliased sub-voxel signals, since they are not part of the k -space sampling, whereas the BOLD PSF further attenuates contributions from high-frequencies that are still within the range of frequencies sampled in the k -space. Therefore, all locally generated contrast useable by a classifier, although very low in amplitude, is caused by random variations and irregularities of the columnar organization, which contribute to low frequency components of this organization. Increasing these irregularities improves classification performance (Fig. 8).

Assumptions, simplifications, and upper bound of classification performance

Exclusive consideration of basic mechanisms—We aimed to develop a model that would show the levels of contrast and classification performance that can be expected considering basic mechanisms. By “basic mechanisms” we refer to the integration of signals that an MRI voxel overlaps, while considering the BOLD point spread (i.e., voxel as a compact kernel, and BOLD-as-blurring model, Kriegeskorte et al., 2010), the process of voxel sampling, and noise. Therefore, of the mechanisms proposed to account for decoding, our model evaluates (I) “aliasing of the main frequency components of the organization” and (II) “contributions of irregularities in the columnar organization,” but not “very low-frequency large-scale components of the organization,” “selectivity of draining veins” and “complex spatio-temporal filters.” Because we aimed to consider basic mechanisms exclusively, we refer to our model as a “naive” model. The results of this naive model are intended to serve as baseline when evaluating more complex mechanisms that potentially contribute to successful decoding of information conveyed by cortical columns.

Simplifications leading to best-case scenario of classification performance—

We made several simplifying assumptions, which cause overestimation of classification performance. These simplifications and assumptions, described in more detail below, include: (1) binary (and separately, smooth) representation of ocular dominance columns and maps, (2) uncorrelated noise model, (3) a perfectly learned model, and (4) the employment of an optimal decision boundary. Therefore, our model offers a best-case estimate (or an upper bound) of classification performance when considering basic mechanisms.

(1) Binary ODC representation: Our model included a non-linearity introduced in the process of simulating the ODC maps, which produced spatial transitions of varying degrees of sharpness. Quantitatively, for large point spreads and/or large voxel sizes, ODC maps with sharper transitions produced larger contrast ranges compared to their counterparts with smooth transitions (Fig. 7). For 3-mm-wide voxels and a 3.5-mm point spread of BOLD response, a binary ODC map produced a contrast range of 0.15% (70% classification performance with 100 voxels and a TR of 2 s).

A binary ODC map, consisting of columns with neurons responding exclusively to either the left or the right eye, is not realistic. Therefore, assuming a binary map contributes to our approach of estimating an upper bound for classification performance.

(2) Uncorrelated noise model: Our model does not consider spatial correlation of noise between voxels. In reality, the noise in a subset of the voxels would be correlated, in part depending on their spatial distance. This will decrease the effective degrees of freedom, and result in decoding performance comparable to that obtained with a reduced number of voxels. Similarly, when considering averaging of volumes before classification, we assume the noise to be uncorrelated in time, which maximizes the SNR gains achieved by averaging. Therefore, in considering independent noise, our model overestimates classification performance.

(3) Perfectly learned model: In our analysis of classification performance, we assumed that the estimated means are equal to the real means of the response distributions. This situation corresponds to a perfectly learned model. In reality, there will be differences between the estimated and the real means of the classified patterns, which will decrease classification performance. The choice of classification algorithm will have an effect on how well the model is learned. The more data used for learning, the closer the estimated means will be to

the real means. Our model reflects this asymptotic limit, conforming to our approach of modeling the best case scenario.

(4) Choice of classification framework and optimal decision boundary: We assumed that evoked responses to stimulation of the left or the right eye follow two respective normal distributions in each voxel. Linear classification is the simplest and optimal choice for classifying this type of data. We applied a decision boundary perpendicular to the line separating the means of the two distributions, that results in a minimum-error-rate classification (Duda et al., 2006), in line with our best-case scenario approach. While this boundary is optimal when considering a perfectly learned model, it is also the decision boundary that linear classifiers such as linear discriminant analysis or linear support vector machines (SVM) would converge to, given a large enough data-set available for learning. Thus, our choice of optimal decision boundary follows our approach of modeling the best-case scenario.

Ocular dominance columns vs. orientation columns—Here we analyzed decoding of information conveyed by ODCs. These were the basis for a study that decoded the visual percept during binocular rivalry (Haynes and Rees, 2005b). Other decoding studies were based on orientation columns (Kamitani and Tong, 2005; Haynes and Rees, 2005a). Orientation is not a binary stimulus dimension: it varies continuously. Furthermore, orientation columns in monkeys have slightly higher spatial frequencies than ODCs (Obermayer and Blasdel, 1993). These differences are expected to decrease differential contrast obtained after considering BOLD point spread and voxel sampling. However, we have shown that classification performance at 3 Tesla depends solely on information represented at spatial frequencies lower than the main frequency of the organization. It may well be the case that differences between ODCs and orientation columns, such as the main spatial frequency and the arrangement of columns (anisotropic and isotropic, respectively) have negligible effects on decoding. In contrast, the exact nature of the small and seemingly irrelevant low frequency signals associated with the two organizations may play a key role in decoding. Indeed, ongoing preliminary simulations show that similar columnar patterns with only subtle differences in their low frequency content can result in very different decoding performances.

Voxel selection and number of voxels—In multivariate classification it is often beneficial to reduce the number of features (voxels) (Pereira et al., 2009). Voxels can be either selected based on condition-unspecific criteria such as their location or general response strength. Alternatively, condition-specific criteria, such as differential contrast between conditions, may be employed in order to optimize decoding performance while reducing the number of voxels.

In our current model, we did not include condition-specific voxel selection. However, our model can be extended to include forms of voxel selection. This can be done, for example, by taking into account changes in the distribution of voxel differential contrasts due to voxel selection (e.g., removing the voxels with the lowest contrast).

The decoding studies to which we compare our model selected voxels according to cortical position relevant to the paradigm and response to a localizer (Kamitani and Tong, 2005) or a measure of response magnitude to the group of stimuli (Haynes and Rees, 2005a) (in a second step, the latter study employed a condition-specific criterion in order to further reduce the number of voxels below 100). The purpose of this kind of voxel selection is to obtain functionally responsive voxels in the gray matter of V1, in accordance with the assumption we employed in our model (voxels localized in gray matter) following a best-case scenario approach.

Classification performance in decoding studies is higher than the modeled upper bound—For 3 mm wide voxels and a 3.5 mm point spread of BOLD response, a binary ODC map produced a contrast range of 0.15%, and correct classification rate of 70% (Fig. 7) with 100 voxels and a TR of 2 s. An intermediate sharpness ($\alpha = 4$) introduced to the ODC map produced a contrast range of 0.08% (61% classification performance with 100 voxels and a TR of 2 s), compared to 0.015% (52% classification performance with 100 voxels and a TR of 2 s) for the smooth ODC map model. As discussed above, a binary ODC map is not realistic. Nonetheless, it gives an upper bound for classification performance (70%). Considering that $\alpha = 4$ is much more likely to reflect realistic ODC patterns, and taking into account all other best-case scenario approximations, we expect that realistic classification performance based solely on basic mechanisms and 100 voxels is in the range of 55–65%.

A previous study that considered ODCs (Haynes and Rees, 2005b) obtained ~75% correct classification. Haynes and Rees classified binocular rivalry percepts projected onto a model based on training with monocular stimulation and stable perception (Haynes and Rees, 2005b, Fig. 4C). Our estimated realistic classification (55–65%) is significantly lower than that obtained in this study, although the modeled best-case scenario performance with a binary ODC map (70%) is comparable to the one obtained by Haynes and Rees (2005b). Note however, that this study used a TR of 1.3 s and only 50 voxels for classification. With this TR and number of voxels, our model predicts a best-case scenario classification performance of 64% for a binary ODC pattern and 57% for a realistic ODC pattern (ODC model with intermediate sharpness).

Our estimated realistic- and best-case scenario classification performances are lower than those obtained for two orthogonal orientation stimuli using LDA (~80%; Haynes and Rees, 2005a). This study used a TR of 1.3 s and 100 voxels, for which our model predicts decoding performance of 69% for the binary ODC map and 60% for the realistic ODC map with intermediate sharpness ($\alpha = 4$).

Our estimated realistic- and best-case scenario classification performances are lower than those obtained for two orthogonal orientation stimuli using a linear support vector machine (~96%; Suppl. Fig. 4, Kamitani and Tong, 2005). However, Kamitani and Tong (2005) averaged 8 volumes together before classification, which is expected to increase decoding performance. Taking this effect into account, our model predicts a decoding performance of 98% for the binary ODC map and 86% for the ODC map with intermediate sharpness ($\alpha = 4$). Although this latter classification rate (86%) that considers a realistic ODC sharpness depends on several best-case scenario assumptions, it is still lower than the actual results (96%) obtained by Kamitani and Tong (2005).

Overall, the modeled classification performance is lower than what has been obtained in decoding studies at 3 Tesla, although it considers several best-case scenario assumptions. In the rest of this section we discuss mechanisms of fMRI-based decoding and possible reasons for these differences.

Mechanism of fMRI-based decoding of information conveyed in cortical columns

Aliasing is not possible in MRI: sampling with sinc-function voxels—Whereas typically, MRI voxels are considered to be squares, they are in fact more accurately described as sinc-functions in the space domain (Haacke et al., 1999). This more accurate description rules out spatial aliasing of subvoxel-scale signals in MRI.

If the imaging PSF (not to be confused with the BOLD PSF) is considered to be a rect-function, then in the Fourier domain, the MR signal is described by a sinc-function

multiplied with the Fourier representations of the columnar organization and associated BOLD PSF (Figs. 4 and 5). Since the “ripples” in the tails of the sinc-function extend infinitely, this means that spatial frequency components higher than the MRI Nyquist frequency (sub-voxel) can contribute to the measured signal in k-space; in other words, subvoxel-scale signals are spatially aliased into lower spatial frequencies in the reconstructed image. These sub-voxel signals are further attenuated by the low-pass BOLD PSF, which acts as an anti-aliasing filter (Fig. 5).

However, a better characterization of the imaging PSF is as a sinc-function in the space domain, not the Fourier domain (Haacke et al., 1999). This means that the Fourier domain representation of the signal is a rect-function multiplied with the Fourier representations of the columnar organization and the BOLD PSF. A rect-function has compact support, meaning that high spatial frequency components are zeroed out, and cannot be spatially aliased into lower frequencies in the reconstructed image. In this more accurate model of the imaging process, it is impossible for MRI to be sensitive to sub-voxel, supra-Nyquist scale signals, regardless of the BOLD PSF.

As described by Greenspan (2009), MRI super-resolution is impossible in the phase and frequency encode directions, as MRI is inherently band limited in these directions. Mayer and Vrscaj (2007) suggested that while super-resolution may be technically possible in the Phase-Encoding direction, it can at best contribute only a very limited amount of additional information. Therefore, band limitations of the imaging and reconstruction processes prevent or limit detection of sub-voxel supra-Nyquist signals. These band limitations hold for fMRI and fMRI-based decoding (Swisher et al., 2010) and rule out, under the assumption that an MRI voxel acts as a compact kernel (rather than a spatio-temporal filter), the possibility of sub-voxel scale contributions via aliasing as a contributing mechanism to decoding.

The effects of voxel size and the PSF of the imaging signal—We found strong dependence of classification performance on the point spread of the imaging signal, especially when small voxels are used. This result can be explained by the substantial decrease in functional contrast with increasing point-spread (Fig. 3A and B).

The BOLD point spread and the voxel sampling have very similar effects on the functional contrast: both act as low-pass filters, reducing information conveyed by higher frequencies. Nonetheless we found that for large point spreads, the voxel width has almost no effect on functional contrast (Fig. 3D). In contrast, for large voxel widths, increasing BOLD PSF still decreases the functional contrast (Fig. 3B). The reason for this is that MR voxel sampling simply discards frequencies higher than the Nyquist frequency but leaves lower frequencies untouched. Therefore it has a very small effect when high frequencies are already filtered out by the BOLD point spread. In contrast, the BOLD PSF reduces contributions at every frequency, including lower frequencies.

The classification performance obtained when considering a 3.5-mm wide point-spread was lower than previously reported (Kamitani and Tong, 2005; Haynes and Rees, 2005a,b). This phenomenon suggests that rather than considering the reported mean point-spread, one needs to consider possible variability of the point-spread in space (Kriegeskorte et al., 2010). Along these lines, it is possible that previous decoding studies relied in part on data from cortical sites in which the PSF was significantly lower than 3.5 mm, while excluding data associated with wider PSF. This can be done implicitly by the learning algorithm, by assigning high-weights to data from voxels with selective responses that are presumably associated with narrow PSF.

Yet another possibility is that the reported BOLD point spread at 3 Tesla was overestimated. One reason for such overestimation could be the relatively large voxels ($2 \times 2 \times 2 \text{ mm}^3$) used in these studies (Parkes et al., 2005). Using large voxels for sampling introduces low-pass filter properties that can contribute to an overestimated point spread width. Based on our previous analysis of this effect (Fig. 8 in Shmuel et al., 2007), we expect that the mean GE-BOLD PSF at 3 Tesla is approximately 3 mm. Our imaging model assumes the BOLD point spread width to be only BOLD response related, with low-pass contributions from the voxel sampling process considered independently.

The convolution with the BOLD PSF in our model cannot be compared to spatial smoothing of already obtained fMRI data (Op de Beeck, 2010; Swisher et al., 2010; Kamitani and Sawahata, 2010). In our imaging model, the convolution with the PSF precedes both the voxel sampling and the consideration of noise. Therefore, the reduced classification performance obtained here following the convolution with the BOLD PSF is not in disagreement with the findings on the effect of spatial smoothing on classification rate (Swisher et al., 2010; Op de Beeck, 2010; Kamitani and Sawahata, 2010).

Irregularities/low spatial frequency components of columnar organizations in V1

—It was hypothesized that random, local variations and irregularities in the functional organization contribute to decoding (Kamitani and Tong, 2005, 2006; Haynes and Rees, 2006; Kriegeskorte and Bandettini, 2007). The argument is that due to the irregular underlying columnar pattern, each voxel overlaps columns with different preferences unequally, resulting in biases towards specific preferences. Irregularities are thought to be manifested through components of the columnar organization with frequencies higher and lower than the main frequency of the organization (Rojer and Schwartz, 1990). These components may be present even if the overall preferences represented by the columns are distributed equally across the investigated cortical region.

Here, we have shown that signal from the main frequency of the columnar organization (0.5 cycles/mm) cannot contribute to decoding (Fig. 4). The only local contributions to contrast range arise from frequency components that are considerably lower than the main frequency of the columnar organization and are lower than the Nyquist frequency which corresponds to the diagonal in k-space (Fig. 5). These low frequencies, in conjunction with frequencies higher than the main frequency of the organization, underlie random variations and irregularities in the columnar pattern. Indeed, varying the content of these irregularities had a strong effect on decoding performance (Fig. 8).

Higher classification performance in decoding studies could be explained if we considerably underestimated low frequency components in the ODC pattern. Following Rojer and Schwartz (1990) we used a filter composed of two Gaussians to model ODC columns. There are indications (Rojer and Schwartz, 1990; Blasdel et al., 1995) that the spatial frequency spectra of real ODC columns correspond to a more heavy-tailed filter functions than the Gaussian filter. In other words, ODC organizations are expected to include higher contributions of low spatial frequencies than those we modeled. Note that we refer here to low-spatial frequency components caused by local random variations of ODCs, even when considering equal representations of the two eyes at the more global level. This could potentially be a source for larger contrast contributions by low frequencies, and would imply improved classification performance over those obtained with the Gaussian-filter based maps we analyzed here.

Experimental evidence for significant, local rather than global, contributions of low-spatial-frequency components to the pattern of ODCs was demonstrated by Shmuel et al. (2010). Figs. 2–4 in this paper show OD patterns following low-pass filtering (cycles shorter than 4

mm were filtered out). Note significant contributions of low-frequency components to the differential maps (panel A in Figs. 2–4, Shmuel et al., 2010); these low-frequency components carry discriminative power (panel B). Whereas some of these eye-selective broad structures correspond to macroscopic blood vessels, others correspond to regions in which gray matter contributions dominate (panel C in Figures 2, 4, 5 S1 and 5 S3, Shmuel et al., 2010). We expect that the latter are caused by local variations in the ODC pattern. Similarly, Swisher et al. (2010) reported that, in cat visual cortex, reliable orientation bias could still be found at spatial scales of several millimeters. In the human visual cortex, the majority of orientation information imaged at a resolution of $1 \times 1 \times 1 \text{ mm}^3$ was found on scales of millimeters (Swisher et al., 2010).

Large-scale organizations in V1—Additional contributions from very low-frequency components to decoding of the stimulated eye could be of a more global origin, e.g. higher response amplitude to the contra-lateral eye in V1. This mechanism was not evaluated by our model. Such higher response amplitude could result from unequal representations of the two eyes, termed ‘nasotemporal asymmetry’ (Tychsen and Burkhalter, 1997).

Low-frequency large-scale organizations of a more global nature that may contribute to decoding of orientation are the radial bias (Sasaki et al., 2006) and the oblique effect (Furmanski and Engel, 2000). The radial bias is an overrepresentation of orientations in cortical positions in which these orientations are retinotopically radial relative to the center of the visual field. It introduces very low frequency components on top of the low frequency components caused by local random variations as described above. The oblique effect is an overrepresentation of cardinal orientations (horizontal and vertical) compared to oblique orientations. This effect is expected to introduce very low-frequency, large-scale differences between the response to cardinal and oblique orientations; it may contribute to distinguishing between these two groups of orientations. Consistent with these expectations, Swisher et al. (2010) reported contributions to decoding of orientation in the human visual cortex from larger-scale spatial biases exceeding 1 cm.

Functional selectivity of macroscopic blood-vessels and complex spatio-temporal filters—As mentioned above, we developed a model of basic mechanisms that estimates contributions to functional contrast and classification from aliasing and low-frequency components caused by random variations in the columnar organization. Our model does not consider contributions from functionally selective macroscopic blood vessels (Gardner et al., 2006; Shmuel et al., 2010). Therefore, the differences between our modeled classification performance and those obtained in previous decoding studies could be accounted for, in part, by contributions of macroscopic blood vessels to decoding. Lastly, Kriegeskorte et al. (2010) introduced the hypothesis that a voxel’s BOLD response can be modeled as a complex spatio-temporal filter of neuronal activity. Assuming that this hypothesis proves true, it may account for part of the differences between previously measured- and our modeled classification performance.

Conclusions

Under the assumptions of MRI voxels acting as compact kernels, BOLD-blurring of neuronal activity, and imaging parameters used at 3 Tesla, spatial frequencies as high as the main frequency of ODCs (0.5 cycles per mm) cannot contribute to decoding of stimulus features represented in cortical ODCs. Variations in the ocular dominance maps captured by lower frequencies constitute the only local component that conveys significant information on the stimulated eye. The contrasts contributed by these low frequencies are very small though, insufficient for accounting for classification performance reported at 3 Tesla. We expect that lower frequency, larger scale pattern variations (e.g., due to higher-amplitude

responses to the contra lateral eye; and oblique and radial effects in the orientation domain) contribute significantly to fMRI based classification. We expect, in addition, that mechanisms not considered in the current model, e.g. functionally biased venous responses, spatially-variable point spread, and complex spatio-temporal filtering of neuronal activity play significant roles in decoding.

Acknowledgments

We thank Bruce Pike, Peter O'Connor, Ze-Shan Yao, Javeed Shaikh, Debra Dawson, Lars Omlor and Sebastian Schmitter for their helpful comments. Supported by a Max-Planck Society fellowship awarded to DC, NIH grants P41 RR08079, P30 NS057091, R01-MH070800 and R01-EB000331, Natural Sciences and Engineering Research Council of Canada grant 375457-09, Human Frontier Science Program grant RGY0080/2008, and by the Canada Research Chairs program.

Appendix A: Performance of a linear classifier

Let n be the number of voxels. Consider a voxel response map as an n -dimensional data vector \vec{x}_t . Each vector \vec{x}_t is sampled from one of two normal distributions stimulation $\mathcal{N}(\vec{\mu}_A, \Sigma_A)$ and $\mathcal{N}(\vec{\mu}_B, \Sigma_B)$ corresponding to the two conditions.

The means $\vec{\mu}_A = (\mu_A^1, \mu_A^2, \dots)$ and $\vec{\mu}_B = (\mu_B^1, \mu_B^2, \dots)$ characterize the expected voxel-wise activation in all voxels under each condition.

We assume that the data is zero centered in the sense that $\vec{\mu}_A = -\vec{\mu}_B$. Defining the expected differential activation $\vec{d} = (\vec{\mu}_B - \vec{\mu}_A)$, we can write $\vec{\mu}_A = -\frac{\vec{d}}{2}$ and $\vec{\mu}_B = +\frac{\vec{d}}{2}$.

The covariance matrices of the two distributions $\Sigma_A = \Sigma_B = \sigma^2 I$ characterize the relative noise, which is assumed to be independent and identically distributed between voxels.

For classification, we project each data vector \vec{x}_t onto the normalized vector $\vec{w} = \frac{\vec{d}}{\|\vec{d}\|}$ pointing in the direction of the line connecting the means of the two distributions. This results in one-dimensional variables y_t given by:

$$y_t = \vec{w}^T \vec{x}_t = \frac{\vec{d}^T}{\|\vec{d}\|} \vec{x}_t.$$

These resulting variables y_t will also be normally distributed according to $\mathcal{N}(m_A, s^2)$ or $\mathcal{N}(m_B, s^2)$, depending on which condition their corresponding activation vectors \vec{x}_t were associated with. The distribution means are given by $m_A = \vec{w}^T \vec{\mu}_A = -\frac{\|\vec{d}\|}{2}$ and $m_B = \vec{w}^T \vec{\mu}_B = \frac{\|\vec{d}\|}{2}$. The variance is given by $s^2 = \vec{w}^T \sigma^2 \vec{w} = \|\vec{w}\|^2 \sigma^2 = \sigma^2$.

If $y_t < 0$ then x_t is classified as belonging to A , otherwise x_t is classified as belonging to B .

The expected percentage p of correct classifications does not change if we restrict our analysis to responses coming from one condition only, due to the symmetry of conditions. Without loss of generality we can choose condition A and compute p as the expected fraction of y_t associated with condition A ($y_t \sim \mathcal{N}(m_A, s^2)$) that is also classified as coming from condition A ($y_t < 0$).

$$p = \int_{-\infty}^0 f_{\mathcal{N}\left(-\frac{\|\vec{d}\|}{2}, \sigma^2\right)}(y) dy = \int_{-\infty}^{\frac{\|\vec{d}\|}{2\sigma}} f_{\mathcal{N}}(y) dy,$$

where $f_{\mathcal{N}}$ is the probability density function of \mathcal{N} .

We define the contrast range c to be the standard deviation of the distribution of differential contrasts that can be obtained in a single voxel:

$$c = \sqrt{\frac{1}{n} \sum_{i=1}^n d_i^2} = \frac{1}{\sqrt{n}} \|\vec{d}\|.$$

The equation for p reduces to

$$p = \int_{-\infty}^{\frac{\sqrt{nc}}{2\sigma}} f_{\mathcal{N}}(y) dy. \quad (1)$$

Defining the overall contrast-to-noise ratio to be $OCNR = \frac{\sqrt{nc}}{\sigma}$ we get:

$$p = \int_{-\infty}^{\frac{OCNR}{2}} f_{\mathcal{N}}(y) dy. \quad (2)$$

$OCNR$ is related to Fisher's criterion $\frac{\|\vec{d}\|^2}{2\sigma^2} = \frac{OCNR^2}{2}$. Fisher's criterion measures the ratio between within-class variance and between-class variance and is maximized in linear discriminant analysis.

When averaging multiple volumes before classification it is possible to reduce the noise level σ to σ_{avr} . Assuming that temporal noise is uncorrelated, the reduced noise level is

$\sigma_{avr} = \frac{\sigma}{\sqrt{t}}$, where t is the number of averaged volumes. The overall contrast-to-noise-ratio is then $OCNR = \frac{\sqrt{nt}c}{\sigma}$.

Appendix B: Definition of the model components

We define each single step of the model as a transformation with input variables denoted as x and output variables denoted as y . In general these are two dimensional fields (real-valued functions on \mathbb{R}^2), representing quantities in two-dimensional image space. $\vec{r} \in \mathbb{R}^2$ denotes spatial position and $\vec{k} \in \mathbb{R}^2$ denotes coordinates in k-space.

B.1. ODC model

The ODC pattern is modeled by filtering Gaussian spatial white noise according to Rojer and Schwartz (1990). The shape of the filter is defined in k-space as the sum of two two-dimensional Gaussian functions (reflecting the symmetry of k-space):

$$\tilde{F}_{ODC}(\vec{k}) := e^{-\left(\frac{2\log 2}{\epsilon^2}k_1^2 + \frac{2\log 2}{\delta^2}(k_2 - \rho)^2\right)} + e^{-\left(\frac{2\log 2}{\epsilon^2}k_1^2 + \frac{2\log 2}{\delta^2}(k_2 + \rho)^2\right)},$$

where ρ is the principal frequency determining the column width. δ is the width (full width at half maximum) of each Gaussian parallel to the filter orientation. δ determines the variation in column width. ϵ is the width (full width at half maximum) of each Gaussian orthogonal to the filter orientation. ϵ determines the *branchiness* of the columns.

In order for the ODC maps to have the same variance as the Gaussian white noise, we normalize the filter:

$$F_{ODC}(\vec{k}) := \frac{\tilde{F}_{ODC}(\vec{k})}{\sqrt{\int_{\mathbb{R}^2} \tilde{F}_{ODC}^2(\vec{k}') d\vec{k}'}}$$

Using the spatial representation of the filter $f_{ODC} := \mathcal{F}^{-1}(F_{ODC})$, the transformation that creates the ODC pattern $y(\vec{r})$ from the white noise input $x(\vec{r})$ is defined as:

$$y(\vec{r}) = x(\vec{r}) * f_{ODC}(\vec{r}).$$

We then pass the output of the filter, now denoted by x , through a sigmoidal non-linearity that controls the sharpness of transitions in neighboring ODCs (Rojer and Schwartz, 1990):

$$y = \frac{1}{1 + e^{-ax}}.$$

B.2. Neuronal response

The response of the neuronal population depends on the stimulus condition. We assume a maximal response of 1 for monocular neurons when stimulated through their preferred eye.

Using the ODC map $x(\vec{r})$ as the input, we obtain the two condition specific responses $y_A(\vec{r})$ and $y_B(\vec{r})$.

$$\begin{aligned} y_A(\vec{r}) &= \frac{1}{2} + \frac{x(\vec{r})}{2} \\ y_B(\vec{r}) &= \frac{1}{2} - \frac{x(\vec{r})}{2}. \end{aligned}$$

B.3. BOLD response

The BOLD response is modeled by convolving the neuronal response with a Gaussian spatial impulse response function given by:

$$f_{BOLD}(\vec{r}) = \frac{\beta}{2\pi\sigma_{BOLD}^2} \cdot e^{-\frac{r_1^2+r_2^2}{2\sigma_{BOLD}^2}},$$

where σ_{BOLD} defines the spatial width of the response. It is related to the full width at half maximum of the response by $FWHM=2\sqrt{2\ln 2} \cdot \sigma_{BOLD} \approx 2.35 \cdot \sigma_{BOLD}$. The response magnitude β is the maximal response corresponding to a neuronal response of 1.

The BOLD response $y(\vec{r})$ elicited by the neuronal response $x(\vec{r})$ then given by:

$$y(\vec{r}) = x(\vec{r}) * f_{BOLD}(\vec{r}).$$

B.4. Voxel sampling

Sampling of a voxel is modeled according to the MRI measurement process (Haacke et al.

1999), by sampling the k-space representation of the signal $\tilde{X}(k_1, k_2) = \mathcal{F}(\tilde{x}(\vec{r}))$ at discrete steps and calculating the inverse discrete Fourier transform. The resulting sampled signal is

$$\tilde{y}_{(l,m)} = w^2 \Delta k^2 \sum_{p=-N}^{N-1} \sum_{q=-N}^{N-1} X(p\Delta k, q\Delta k) e^{i\pi pqm/N^2},$$

where $\tilde{y}_{(l,m)}$ is the signal associated with the voxel with indices (l, m) , w is the voxel width and $2N$ is the number of sampled points along one dimension in k-space.

In order to obtain the signal change $y_{(l,m)}$ relative to the baseline of an MRI-sampled signal, we consider a spatially constant baseline pattern of amplitude b and a pattern $x(\vec{r})$ of change relative to baseline.

The signal to be sampled is $\tilde{x}_{stim}(\vec{r}) = b + b \cdot x(\vec{r})$ during stimulation and $\tilde{x}_{bl}(\vec{r}) = b$ during baseline.

It follows then, that the sampled change relative to baseline is

$$\begin{aligned} y_{(l,m)} &= \frac{\tilde{y}_{(l,m)}(\tilde{x}_{stim}(\vec{r}))}{\tilde{y}_{(l,m)}(\tilde{x}_{bl}(\vec{r}))} - 1 \\ &= \frac{\tilde{y}_{(l,m)}(b + b \cdot x(\vec{r}))}{\tilde{y}_{(l,m)}(b)} - 1 \\ &= \frac{b \cdot w^2 + b \cdot \tilde{y}_{(l,m)}(x(\vec{r}))}{b \cdot w^2} - 1 \\ &= \frac{1}{w^2} \tilde{y}_{(l,m)}(x(\vec{r})) \\ &= \Delta k^2 \sum_{p=-N}^{N-1} \sum_{q=-N}^{N-1} X(p\Delta k, q\Delta k) e^{i\pi pqm/N^2}. \end{aligned}$$

To obtain the signal value of one voxel, we pick without loss of generality the center voxel at $l=0, m=0$:

$$y_{(0,0)} = \Delta k^2 \sum_{p=-N}^{N-1} \sum_{q=-N}^{N-1} X(p\Delta k, q\Delta k).$$

For $N \gg 1$, we can apply integration instead of summation, taking into account that

$N\Delta k = \frac{1}{2w}$. Our voxel sampling process is then modeled by

$$y = \lim_{N \rightarrow \infty} y_{(0,0)} = \int_{-\frac{1}{2w}}^{+\frac{1}{2w}} \int_{-\frac{1}{2w}}^{+\frac{1}{2w}} X(k_1, k_2) dk_1 dk_2.$$

Using a rect-function in k-space, we can drop the integration boundaries:

$$y = \int_{\mathbb{R}^2} \text{rect}(wk_1) \text{rect}(wk_2) X(k_1, k_2) dk_1 dk_2.$$

The integral of a function over the entire k-space equals the value of its Fourier transform at 0. Furthermore, we replace the product in k-space by a convolution in image space, and calculate its value at 0 taking the symmetry of the sinc-function into account:

$$\begin{aligned} y &= \mathcal{F} [\text{rect}(wk_1) \text{rect}(wk_2) X(k_1, k_2)] (0) \\ &= \left(x(\vec{r}) * \left(\frac{1}{w^2} \text{sinc}\left(\frac{\pi r_1}{w}\right) \text{sinc}\left(\frac{\pi r_2}{w}\right) \right) \right) (0) \\ &= \frac{1}{w^2} \int_{\mathbb{R}^2} x(\vec{r}') \left(\text{sinc}\left(\frac{\pi(r'_1 - r_1)}{w}\right) \text{sinc}\left(\frac{\pi(r'_2 - r_2)}{w}\right) \right) dr_1 dr_2 \Big|_{\vec{r}'=0} \\ &= \frac{1}{w^2} \int_{\mathbb{R}^2} x(\vec{r}) \cdot \left(\text{sinc}\left(\frac{\pi r_1}{w}\right) \text{sinc}\left(\frac{\pi r_2}{w}\right) \right) dr_1 dr_2. \end{aligned}$$

The last line shows that the signal sampled by a voxel can be regarded as an integral over image space weighted with a sinc-function centered on the voxel.

B.5. Differential activation

The differential activation y is obtained by subtracting the activations of the two conditions:

$$y = x_1 - x_2.$$

Appendix C: Time-course signal to noise ratio and its dependence on repetition time

Time-course signal to noise ratio tSNR is modeled using the following formula (Triantafyllou, et al. 2005):

$$\begin{aligned} \text{tSNR} &= \frac{\text{SNR}_0}{\sqrt{q + \lambda^2 \cdot \text{SNR}_0^2}} \\ &= \frac{\kappa \cdot V}{\sqrt{1 + \lambda^2 \cdot \kappa^2 \cdot V^2}}, \end{aligned} \quad (3)$$

where SNR_0 is the image SNR, V is the voxel volume, λ is a field and scanner independent constant governing the relation between temporal SNR and image SNR, and κ is a field strength and hardware dependent proportionality constant between volume and image SNR.

When the repetition time TR is short, the longitudinal magnetization does not fully recover resulting in a lower signal and therefore lower $\text{SNR}_0(\text{TR})$ relative to the maximally obtainable $\text{SNR}_0^{\text{max}}$ for infinite TR. Using the Ernst angle as the excitation angle, $\text{SNR}_0(\text{TR})$ is related to $\text{SNR}_0^{\text{max}}$ according to Haacke et al. (1999):

$$\begin{aligned}\text{SNR}_0(\text{TR}) &= \text{SNR}_0^{\text{max}} \cdot \sqrt{\frac{1-e^{-\frac{\text{TR}}{T_1}}}{1+e^{-\frac{\text{TR}}{T_1}}}} \\ &= \text{SNR}_0^{\text{max}} \cdot \sqrt{\tanh\frac{\text{TR}}{2T_1}}.\end{aligned}$$

If SNR_0 is given for a specific $\text{TR}=\widetilde{\text{TR}}$, it follows that $\text{SNR}_0(\text{TR})$ for any TR is

$$\text{SNR}_0(\text{TR}) = \text{SNR}_0(\widetilde{\text{TR}}) \cdot \sqrt{\frac{\tanh\frac{\text{TR}}{2T_1}}{\tanh\frac{\widetilde{\text{TR}}}{2T_1}}}.$$

Inserting this result into equation 3 we get:

$$\begin{aligned}\text{tSNR}(\text{TR}) &= \frac{\text{SNR}_0(\widetilde{\text{TR}}) \cdot \sqrt{\frac{\tanh\frac{\text{TR}}{2T_1}}{\tanh\frac{\widetilde{\text{TR}}}{2T_1}}}}{\sqrt{1+\lambda^2 \cdot \text{SNR}_0(\widetilde{\text{TR}})^2 \frac{\tanh\frac{\text{TR}}{2T_1}}{\tanh\frac{\widetilde{\text{TR}}}{2T_1}}}} \\ &= \frac{\kappa_{\widetilde{\text{TR}}} \cdot V \cdot \sqrt{\frac{\tanh\frac{\text{TR}}{2T_1}}{\tanh\frac{\widetilde{\text{TR}}}{2T_1}}}}{\sqrt{1+\lambda^2 \cdot \kappa_{\widetilde{\text{TR}}}^2 \cdot V^2 \frac{\tanh\frac{\text{TR}}{2T_1}}{\tanh\frac{\widetilde{\text{TR}}}{2T_1}}}}\end{aligned}\quad (4)$$

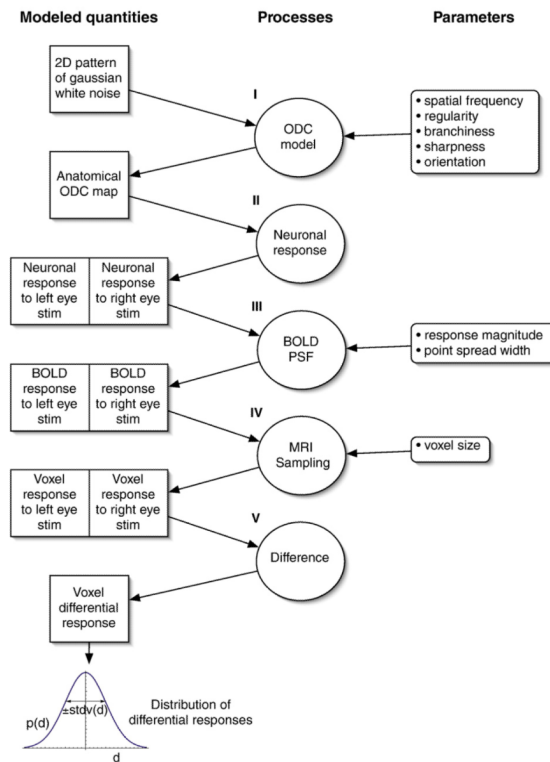
where $\kappa_{\widetilde{\text{TR}}}$ is the constant κ estimated for data acquired using the repetition time $\widetilde{\text{TR}}$.

References

- Adams DL, Sincich LC, Horton JC. Complete pattern of ocular dominance columns in human primary visual cortex. *J. Neurosci.* 2007; 27:10391–10403. [PubMed: 17898211]
- Blasdel G, Obermayer K, Kiorpes L. Organization of ocular dominance and orientation columns in the striate cortex of neonatal macaque monkeys. *Vis. Neurosci.* 1995; 12:589–603. [PubMed: 7654611]
- Boynton GM. Imaging orientation selectivity: decoding conscious perception in V1. *Nat. Neurosci.* 2005; 8:541–542. [PubMed: 15856054]
- Boynton GM, Engel S, Glover G, Heeger DJ. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* 1996; 16:4207–4221. [PubMed: 8753882]
- Duda, RO.; Hart, PE.; Strok, DG. *Pattern classification*. 2nd ed. Wiley-Interscience; 2006.
- Engel S, Glover G, Wandell BA. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex.* 1997; 7:181–192. [PubMed: 9087826]
- Furmanski CS, Engel SA. An oblique effect in human primary visual cortex. *Nat. Neurosci.* 2000; 3:535–536. [PubMed: 10816307]
- Gardner, JL.; Sun, P.; Tanaka, K.; Heeger, DJ.; Cheng, K. Classification analysis with high spatial resolution fMRI reveals large draining veins with orientation specific responses. Society for Neuroscience Meeting; Atlanta, GA, USA. 2006.
- Greenspan H. Super-Resolution in Medical Imaging. *Comput. J.* 2009; 52:43–63.
- Haacke, ME.; Brown, RW.; Thompson, MR. *Magnetic resonance imaging: physical principles and sequence design*. Wiley-Liss; 1999.

- Haxby J, Gobbini M, Furey M, Ishai A, Schouten J, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*. 2001; 293:2425–2430. [PubMed: 11577229]
- Haynes J, Rees G. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci*. 2005a; 8:686–691. [PubMed: 15852013]
- Haynes J, Rees G. Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol*. 2005b; 15:1301–1307. [PubMed: 16051174]
- Haynes J, Rees G. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci*. 2006; 7:523–534. [PubMed: 16791142]
- Horton JC, Dagi L, McCrane E, de Monasterio F. Arrangement of ocular dominance columns in human visual cortex. *Arch. Ophthalmol*. 1990; 108:1025–1031. [PubMed: 2164380]
- Kamitani Y, Sawahata Y. Spatial smoothing hurts localization but not information: pitfalls for brain mappers. *Neuroimage*. 2010; 49:1949–1952. [PubMed: 19559797]
- Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat. Neurosci*. 2005; 8:679–685. [PubMed: 15852014]
- Kamitani Y, Tong F. Decoding seen and attended motion directions from activity in the human visual cortex. *Curr. Biol*. 2006; 16:1096–1102. [PubMed: 16753563]
- Kriegeskorte N, Bandettini PA. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage*. 2007; 38:649–662. [PubMed: 17804260]
- Kriegeskorte N, Cusack R, Bandettini P. How does an fMRI voxel sample the neuronal activity pattern: compact-kernel or complex spatiotemporal filter? *Neuroimage*. 2010; 49:1965–1976. [PubMed: 19800408]
- Krüger G, Kastrup A, Glover GH. Neuroimaging at 1.5 T and 3.0 T: comparison of oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med*. 2001; 45:595–604. [PubMed: 11283987]
- Mayer GS, Vrscaj ER. Measuring information gain for frequency-encoded super-resolution MRI. *Magn. Reson. Imaging*. 2007; 25:1058–1069. [PubMed: 17707168]
- Obermayer K, Blasdel GG. Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci*. 1993; 13:4114–4129. [PubMed: 8410181]
- Op de Beeck HP. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? *Neuroimage*. 2010; 49:1943–1948. [PubMed: 19285144]
- Parkes LM, Schwarzbach JV, Bouts AA, Deckers RHR, Pullens P, Kerskens CM, Norris DG. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn. Reson. Med*. 2005; 54:1465–1472. [PubMed: 16276507]
- Pereira F, Mitchell T, Botvinick M. Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage*. 2009; 45:S199–S209. [PubMed: 19070668]
- Roger A, Schwartz E. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol. Cybern*. 1990; 62:381–391. [PubMed: 2110005]
- Sasaki Y, Rajimehr R, Kim BW, Ekstrom LB, Vanduffel W, Tootell RBH. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*. 2006; 51:661–670. [PubMed: 16950163]
- Shmuel A, Chaimow D, Raddatz G, Ugurbil K, Yacoub E. Mechanisms underlying decoding at 7 T: Ocular dominance columns, broad structures, and macroscopic blood vessels in V1 convey information on the stimulated eye. *Neuroimage*. 2010; 49:1957–1964. [PubMed: 19715765]
- Shmuel A, Yacoub E, Chaimow D, Logothetis NK, Ugurbil K. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. *Neuroimage*. 2007; 35:539–552. [PubMed: 17306989]
- Swisher JD, Gatenby JC, Gore JC, Wolfe BA, Moon C-H, Kim S-G, Tong F. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J. Neurosci*. 2010; 30:325–330. [PubMed: 20053913]
- Triantafyllou C, Hoge R, Krueger G, Wiggins C, Potthast A, Wiggins G, Wald L. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage*. 2005; 26:243–250. [PubMed: 15862224]

- Tychsen L, Burkhalter A. Nasotemporal asymmetries in V1: ocular dominance columns of infant, adult, and strabismic macaque monkeys. *J. Comp. Neurol.* 1997; 388:32–46. [PubMed: 9364237]
- Yacoub E, Shmuel A, Logothetis NK, Ugurbil K. Robust detection of ocular dominance columns in humans using Hahn Spin Echo BOLD functional MRI at 7 Tesla. *Neuroimage.* 2007; 37:1161–1177. [PubMed: 17702606]

**Fig. 1.**

Model overview. In stage 1, ODC maps are modeled by spatial filtering of white noise. In stage 2 we simulate the neuronal response to right-or left eye stimulation. In stage 3 the neuronal response is convolved with a BOLD point-spread function. In stage 4 the BOLD response is transformed into a voxel pattern. In stage 5 the difference between the responses to the two stimulation conditions is computed in order to obtain a voxel pattern of differential response. All voxels in this pattern create the distribution of differential response contrast values. This distribution is characterized by its standard deviation, which reflects the range of contrasts in the set of imaged voxels (= “contrast range”).

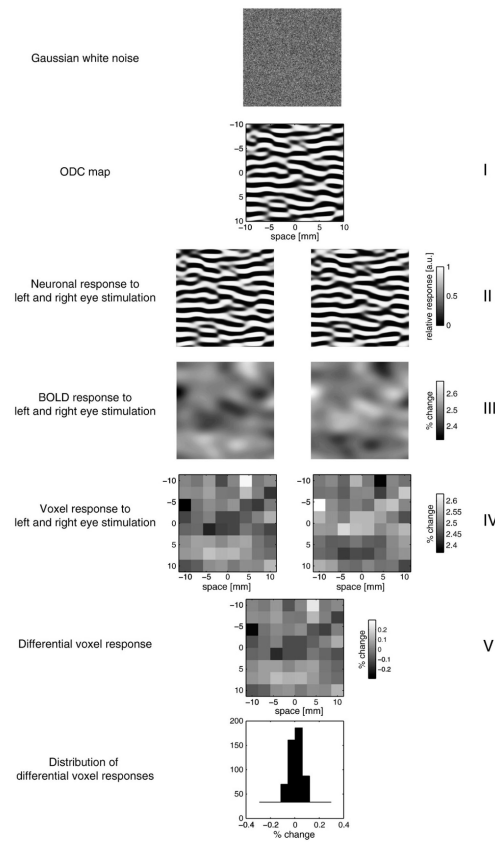
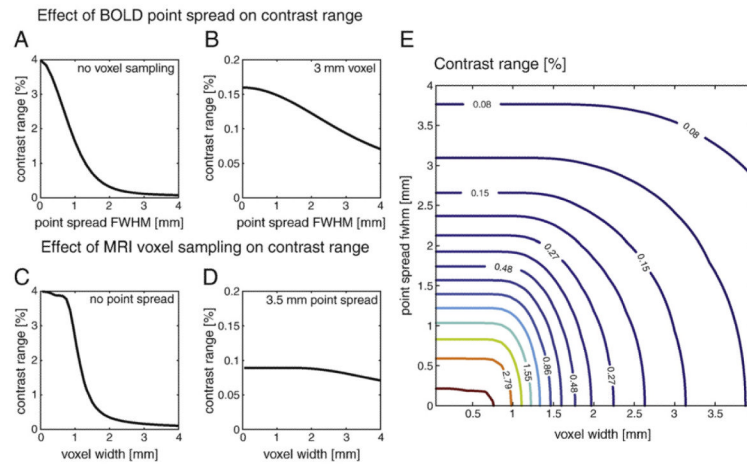
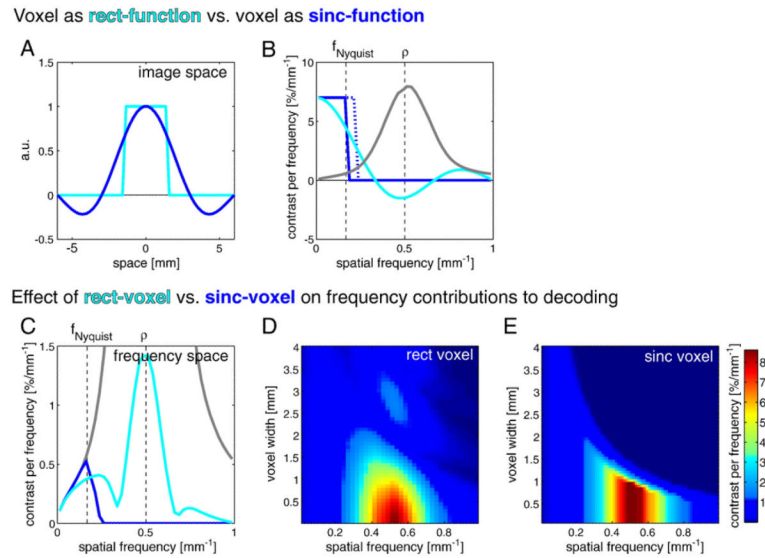


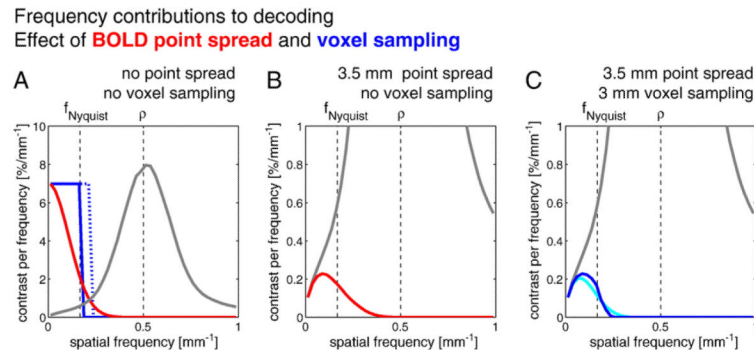
Fig. 2. Numerical realization of the model. The figure presents a numerical simulation of the model. Each row shows the results of one single stage of the model: the Gaussian white noise input, the ODC map, the neuronal response, the BOLD response, the voxel response, and the differential voxel response. The BOLD response and the voxel response show patterns that differentiate the stimulation conditions, although they do not seem to reflect the spatial organization of the ODC pattern.

**Fig. 3.**

Dependence of contrast range on voxel width and BOLD point spread. Contrast range is defined as the standard deviation of the distribution of differential responses (percent change relative to baseline). Contrast range in a set of imaged voxels is presented as a function of FWHM of BOLD point spread (A and B), voxel width (C and D) or both (E). In A, the voxel width is infinitesimally small, while in B it is held constant at $w = 3$ mm. In C, the BOLD point spread is assumed to be infinitesimally small, while in D it is held constant at FWHM = 3.5 mm. Contrast range decreases fast with increasing voxel size and increasing point spread width.

**Fig. 4.**

Comparing the MR imaging process that relies on sinc-shaped voxels to integrating over rect-shaped voxels. The figure presents the contributions of spatial frequency components in the ODC pattern to the range of contrasts in the set of imaged voxels which are sampled as integral over the voxel area (rect-function in image space) or as a sinc-function weighted integral in the image space. The contrast range per frequency (standard deviation of the distribution of differential responses) was computed by restricting the k -space representation of the ODC pattern to different spatial frequencies and calculating the resulting contrast range. (A) The image space representation of a 3-mm wide MRI sinc-shaped voxel (in blue) and the corresponding 3 mm rect-voxel (in cyan). (B) The spatial frequency representations of the MRI imaging process (in blue), voxel as a rect-function (in cyan), and the frequency components of the ODC organization (in gray). “ f_{Nyquist} ” refers to the Nyquist frequency (0.167 cycles/mm for 3 mm voxel). The dotted blue line represents the higher frequencies sampled in k -space along the diagonal rather than along the shorter main coordinate axes. (C) The effect of voxel sampling on the imaged frequency components of the ODC organization. MRI 3 mm wide sinc voxel sampling (in blue) is compared to sampling by integrating over a 3-mm rect-shaped voxel in image space (in cyan). The original ODC frequency components presented in B are shown in gray for comparison. The sinc-shaped voxel is frequency-band limited, while the rect-function is not. With sinc-shaped voxels, contributions from frequencies higher than the Nyquist frequency are sampled along the diagonal in k -space up to a frequency equal to 1.4 times the Nyquist frequency. With rect-shaped voxels, frequencies higher than the Nyquist frequency contribute to contrast range by means of aliasing. (D) Frequency contributions for varying rect-shaped voxel size (BOLD PSF effects were not applied). Aliasing contributions can be observed here at frequencies with cycle lengths larger than twice the voxel width. (E) Frequency contributions for varying sinc-shaped voxel size (BOLD PSF effects were not applied). In contrast to the frequency contributions seen with rect-shaped voxels (D), no contributions can be observed at frequencies with cycle lengths larger than twice the voxel width.

**Fig. 5.**

Contributions of pattern spatial frequency components to the contrast range. The figure presents contributions of spatial frequency components in the ODC pattern to the range of contrasts in the set of imaged voxels, in a format similar to that used in Fig. 4. In A–C, different subsets of the model were used to illustrate their respective effects. (A) The frequency contributions reflect the spectrum of the ODC pattern (gray curve, with no voxel sampling, and no BOLD point spread). This spectrum is dominated by the main pattern frequency (0.5 cycles/mm). Due to the irregularity of the pattern, significantly higher and lower frequencies contribute to the pattern as well. The spatial frequency representation of the MR voxel sampling process (blue) and of the BOLD point spread (red) are shown for comparison. The dotted blue line represents the higher frequencies sampled in k-space along the diagonal rather than along the shorter main coordinate axes. (B) The effect of a BOLD point spread with FWHM=3.5 mm on the imaged frequency components (in red; voxel sampling effects were not applied) is that the BOLD response acts as a low pass filter. The contrast range is dominated by low frequency pattern components. The original ODC frequency contributions presented in A are shown in gray for comparison. (C) Frequency specific contributions for rect-function voxel sampling versus MRI sinc-shaped voxel sampling with 3 mm wide voxels following the convolution in image space with a BOLD point spread with a FWHM (in image space) of 3.5 mm. The BOLD point spread acts as low pass filter, removing aliased high frequency contributions in the rect-shaped voxel sampling, making the result of both sampling models more comparable.

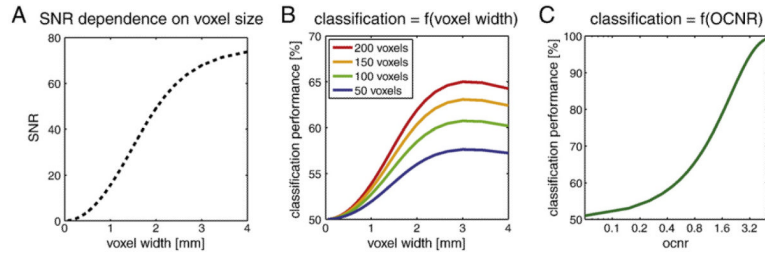


Fig. 6. Classification performance. (A) Time-course SNR (tSNR) as a function of voxel width at 3 Tesla. Noise levels were computed following Triantafyllou et al. (2005), using TR = 2 s (see section C in the Appendix). B presents classification performance as a function of in-plane voxel width. The slice thickness was held constant at 3 mm. A BOLD point spread FWHM of 3.5 mm was applied; Voxel volume dependent noise levels at 3 Tesla were computed following Triantafyllou et al. (2005), modified for a TR of 2 s. Classification performance is presented in units of percent correct classification and is plotted for 50, 100, 150 and 200 voxels. In C, classification performance is shown as a function of overall contrast-to-noise ratio. Classification performance depends on the contrast range, the number of voxels and the relative noise level. All three factors can be combined into one quantity: the overall contrast-to-noise ratio. The overall contrast to noise ratio is proportional to the contrast range and to the square root of the number of voxels. It is inversely proportional to the noise level.

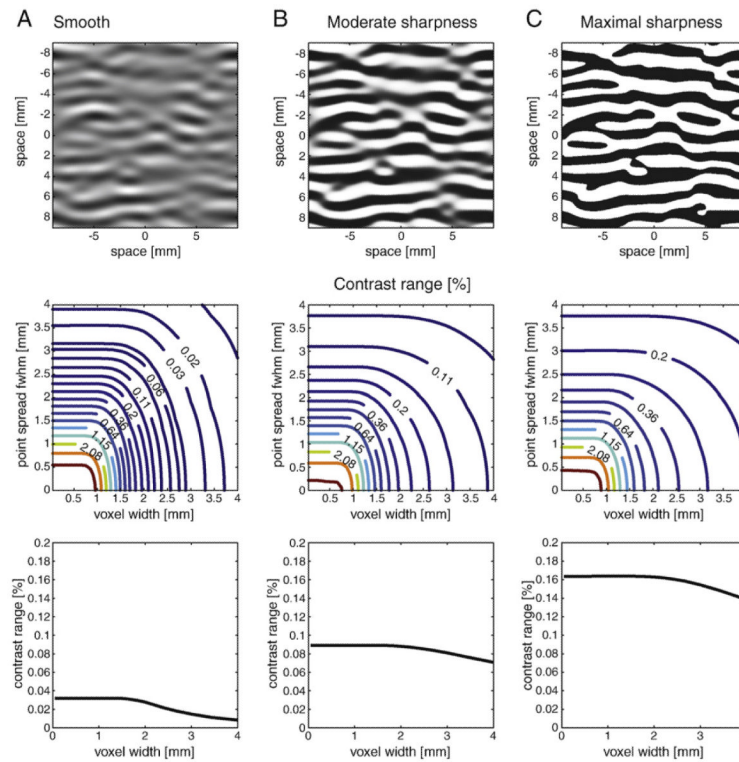
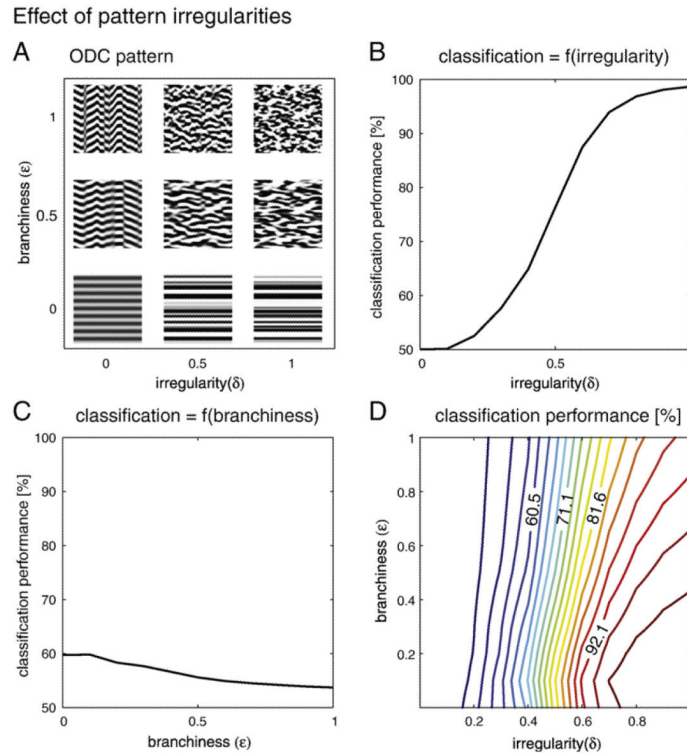


Fig. 7.

The effect of varying ODC pattern sharpness on the sampled contrast range. The sharpness of transitions between ODCs can be modeled by a sigmoidal non-linearity with a degree controlled by parameter α . The figure shows the effect of this non-linearity on the model results by using different alpha values. The top row shows simulated ODC patterns. The middle row shows the corresponding contrast ranges as a function of voxel size and point spread, in a format similar to that used in Fig. 3. The bottom row shows contrast range as a function of voxel width while the BOLD point spread width is held constant at 3.5 mm. (A) shows a smooth ODC pattern. This pattern was obtained directly from the filtered white noise. (B) shows an intermediate level of sharpness ($\alpha = 4$). This pattern is the most realistic of the three ODC patterns presented here. Therefore, it was used in the analysis throughout the rest of the paper. The dependence of contrast range here on voxel width and BOLD point spread is qualitatively similar to that obtained with the smooth ODC pattern. However, the contrast ranges are significantly higher than those obtained using the smooth pattern. (C) shows a binary pattern with sharp edge transitions between neighboring columns (α approaches infinity). The qualitative results are similar to those presented in A and B but the contrast ranges are even larger than those obtained with the intermediate-level sharpness.

**Fig. 8.**

The effect of pattern irregularities on classification rate. The irregularity of the ODC pattern is varied to study its effect on classification performance. Classification performance is predicted for 100 voxels and a TR of 2 s. (A) demonstrates the effects of the irregularity parameter (δ) and the branchiness parameter (ϵ) on the ODC pattern. The panel presents different patterns resulting from combinations of δ and ϵ values of 0, 0.5 and 1. (B) shows classification performance as a function of irregularity (δ) with branchiness (ϵ) held constant at 0.4. Increasing irregularity leads to increasing classification performance, since larger contributions of low frequency components are introduced into the ODC pattern. (C) shows classification performance as a function of branchiness (ϵ) with irregularity (δ) held constant at 0.3. With increasing branchiness, classification performance decreases, because branchiness counteracts the effect of low frequency biases introduced by the irregularity. (D) shows classification performance as a function of irregularity (δ) and branchiness (ϵ). Classification performance depends on irregularity more than on branchiness.

Table 1

Variables and parameters of the model.

Variable	Description	Formula
ρ	Main frequency of ODC pattern	
δ	Pattern irregularity, variations orthogonal to ODC bands	
e	Pattern branchiness, variations parallel to ODC bands	
a	Sharpness parameter of the sigmoidal non-linearity in ODC	$f(x) = \frac{1}{1 + e^{-ax}}$
σ_{BOLD}	Bold point spread width	$FWHM = 2.35 \cdot \sigma_{\text{BOLD}}$
β	Maximal BOLD response	
w	Voxel width	
$\vec{\mu}_z$	Mean multivariate voxel-wise response to condition z	$\vec{\mu}_z = (\mu_{1,z}, \mu_{2,z}, \dots)$
tSNR	Time-course SNR	
σ	Time-course noise in a single voxel (standard deviation of signal change during baseline)	$\sigma = \frac{1}{tSNR}$
\vec{d}	Multi-voxel mean difference between conditions	$\vec{d} = \vec{\mu}_1 - \vec{\mu}_2$
d_i	Single voxel difference between conditions	$\vec{d} = (d_1, d_2, \dots)$
c	Contrast range	$c = \sqrt{\langle d_i^2 \rangle}$
OCNR	Overall contrast-to-noise ratio	$\frac{\sqrt{nt}c}{\sigma}$
n	Number of voxels	
t	Number of averaged volumes	